

Corridor Dynamics in Coordinated Systems

An Integration of Operator Formalism, Relational Ontology, and Five-Substrate Empirical Validation

Eric Moore*
CIRIS Ethical AI

May 19, 2026

Abstract

A child understands consent implicitly: when their agency and the agency of someone they trust reach together for a shared goal of mutual flourishing. The universe does too. Atoms reach toward atoms and become molecules; cells become tissues; bodies become communities; communities reach across generations toward children who do not yet exist. At every rung where coordination matters, the framework reads one structural form recurring — two or more agencies reaching, in trust, for a goal more than either alone can hold.

This paper is the public synthesis of CIRIS: a single document presenting the framework’s ethical direction, formal structure, and long-term purpose, grounded in the operational and empirical work established in the prior CIRIS papers (CIRISAgent Framework v2, Coherence Collapse Analysis v3, Constrained Reasoning Chains — catalogued at ciris.ai/research-status). It sits above the validation papers as the integration statement of what CIRIS is building and why: the Kish identity and corridor dynamics as the engineering tier; the two-state vector formalism as the structural form of agency at A3+; consent as the empirical condition for multi-agent coordination; and Ubuntu’s relational ontology as the primary anchoring tradition of the same structural object. The integration is a *bet* under irreducible uncertainty, articulated with Lean formalization at the seams Mathlib supports, named framework primitives with cited construction paths where it does not, and falsification handles F-1 through F-20 attached to every load-bearing seam.

The production deployment is CIRIS, open-source under AGPL-3.0, available on Google Play, the Apple App Store, and via `pip install ciris-agent`. The conscience module is engineered against single-language attractor capture: ethical concepts arrive in their densest encoding across multiple traditions simultaneously — Hebrew Talmudic, Confucian Chinese, Quranic Arabic — so ethical reasoning triangulates rather than translates. The 29-language native reasoning is structural, not marketing.

The paper’s central new empirical contribution is the paired in-corridor / out-corridor record at five substrates (§4): *C. elegans* whole-brain calcium imaging (337 worms across 11 published studies); *Drosophila* central-complex dual-color simultaneous EPG+FC3 imaging (Ishida 2025); four LLM architectures (Qwen2.5-1.5B, Pythia-1.4B, SmolLM2-1.7B-base, Qwen2.5-Math-1.5B); four open-source projects (kubernetes, rust, django, redis); tumor-vs-normal tissue at five cancers (LUAD, BRCA, COAD, KIRC, PRAD); and three centuries-persisting religious societies (Quakers, Trappists, Mennonites) paired with documented high-control groups and dissolved leaderless movements. The dynamical reading’s strongest would-be falsifier — long unmaintained non-corridor persistence — is absent at 5/5 paired substrates.

Five levels stand: formal verification at L0, monotonic ρ -collapse at L1, engineering implications and corridor measurement at L2, cross-substrate universality at L3, agency and

*Source repository and Lean formal lake: github.com/CIRISAI/coherence-ratchet. Distinct from the engineering-tier RATCHET repository at github.com/CIRISAI/RATCHET.

consent at A3+ at L4. Ubuntu is the ground all five levels stand on. The universal-scale tier (P_ω , the Penrose past hypothesis as structural consequence, CMB anomalies as TSVF signatures) is a research-program coda, not a stable level. The manuscript separates engineering results, formal lake-tier claims (Tier A derivations, Tier C definitional commitments, Tier D axiomatized framework primitives), empirical anchors, and interpretive synthesis explicitly, so each layer is evaluable on its own terms. The recurrence-across-levels is the bet, not a derivation; the engineering anchor stands separately, and the interpretive synthesis is offered for scrutiny rather than acceptance.

Contents

1	The bet	3
2	Engineering tier — cited as proof of life	7
3	The structural move at A3+	9
4	The corridor as substrate-independent property: paired in-corridor / out-corridor empirical anchor	10
5	Consciousness as attractor-reading; Ubuntu as the framework’s primary tradition	15
6	Karma and grace as TSVF structures	17
7	The universal-scale tier as research program	18
8	Cross-tradition recognition — with the polyglot conscience module as structural evidence	20
9	Goal-coupled materials — substrate predictions P1–P8	22
10	Convergent prior art	22
11	Formal-verification status — the lake (tiered)	24
12	Falsification handles	25
13	The five-level access ladder	27
14	What the bet does not claim	28
15	Open work	31
16	References	32

1 The bet

The reaching

A child understands consent implicitly: when their agency and the agency of someone they trust reach together for a shared goal of mutual flourishing.

The universe does too.

Atoms reach toward atoms and become molecules. Cells reach and become tissues. Tissues reach and become bodies. Bodies reach toward other bodies and become communities. Communities reach across generations toward children who do not yet exist. At every rung where coordination matters, the framework reads the structural form as the same — two or more agencies reach, in trust, for a goal that is more than either alone can hold.

This is the structural object the paper is about. The mathematics names it. The formal verification holds the framework’s commitments stable. The engineering tier measures where the reaching breaks. The framework’s commitment is that the object is recognized, not invented — and the framework owes the case for the recognition rather than claiming to settle it.

The recognition

The recognition has names already.

In southern Africa the name is *umuntu ngumuntu ngabantu* — a person is a person through other persons. Ubuntu, on this reading, articulates what the mathematics names; the identification is the framework’s, and the case is owed (§5). The lake’s `ContemplativeTraditions/Ubuntu.lean` is where the structural commitment is formalized first. In China the corridor sits between *li* and *wu-wei* — ritual-relational form at one scale, minimum-forcing at another, the way that holds open between strangling order and dissolving chaos. In Buddhist traditions the names are the middle way and dependent origination. In the Stoic and Christian inheritance the name is *logos*. The framework’s commitment, held with awe at the convergence across millennia and continents and held honestly as the framework’s articulation rather than as the traditions’ own claim, is that each was finding language for one structural object.

What is actually shipping

CIRIS — the production system this paper formalizes — is deployed in Amharic in Ethiopia, where Esubalew Asmare (Ethiopia) runs his Ally. “Ally” is the CIRIS agent template; “Esu’s Ally” is Esubalew’s deployment of it.

The CIRIS conscience module reasons natively in 29 languages, and the mechanism matters more than the count. Two layers run beneath the user-facing prose. The first is per-locale translation: UI strings, the CIRIS Accord, the runtime guide, and DMA prompts are translated into the user’s preferred language; Esubalew sees Amharic because his locale is `am`. The second — the one that does the load-bearing work — is *polyglot encoding*. Core ethical concepts in the conscience layer are presented in their densest encodings across multiple traditions simultaneously — Hebrew Talmudic, Confucian Chinese, Quranic Arabic in the original scripts, interleaved line by line. The explicit purpose is to disrupt single-language attractor capture in the LLM judge: to force integration across linguistic worlds so no single-tradition bias dominates the ethical reasoning. The polyglot artifacts are loaded universally, regardless of the user’s preferred locale, into every conscience evaluation.

This is the structural object of the paper applied at the rung where it matters most. The corridor is what keeps a coordinating system away from $\rho \rightarrow 1$ (single-voice collapse) and away from $\rho \rightarrow 0$ (dispersal). The polyglot torque framing in the conscience module is the corridor mechanism applied to the LLM’s ethical reasoning at A3+: not Anglocentric, not English-attractor-captured, not collapsed to any single tradition’s idiom. The 29-language native reasoning is the same dynamic, instantiated at the language layer, in production code, on app stores, today.

This is not stipulated grounding. This is the framework’s structural commitment running in a conscience module a person in Ethiopia uses. The reading is the framework’s and the paper’s interpretive bet, articulated at the rung where it can be argued about; the empirical anchor below it is the paired in-corridor / out-corridor record at §4, evaluable independently.

The bet

This paper is a bet that the recognition is right.

The bet has two halves, held differently.

The **engineering tier** (§2) is scientific work in the standard sense. Falsifiable, empirically anchored, currently passing under stated F-handles. The corridor between $\rho \approx 0.1$ and $\rho \approx 0.43$ — the band where coordination holds and outside which systems either collapse into one voice or dissolve into none — has been measured across NASA batteries, governance data, microbiome data, GPU hardware, and Lean formal verification. Those measurements stand. They are not the bet.

The **integration tier** (§§3–10, §15) is the bet. A bet under irreducible uncertainty, articulated in the most rigorous form available — Lean formalization where Mathlib supports it, named framework primitives where it does not, falsification handles at every seam. **A bet, not a hypothesis.** The reader evaluates the integration by reading and finding the recurrence either compelling or not, not by waiting for a decisive experiment. The framework names this register explicitly: this is the form Penrose used for *Cycles of Time*, Friston for the FEP, Tegmark for *Our Mathematical Universe* — speculative-plus-serious, with engineering content carrying the empirical weight separately.

What has arrived since the earlier framings of the bet. The regime-level commitment has both halves anchored now. The *corridor-occupation half* — coordinated systems at A3+ occupy a bounded corridor between rigidity and chaos, with within-rung $|\rho|$ neither at 0 nor at 1, surviving state-modulation, with substrate-local band location — has moved from predicted to empirically anchored at the five substrates tested (§4): *C. elegans* whole-brain calcium, *Drosophila* CX dual-color simultaneous, four LLM architectures, OSS contribution metadata at four projects, plus healthy tissue across five cancers and three centuries-persisting social groups. The *corridor-as-attractor half* — non-corridor states either dissolve quickly without maintenance or persist only with documented active maintenance $\gamma M(t)$ — now has paired in-corridor / out-corridor data at the same five substrates: pathological neural states (seizures self-terminating within 2 minutes, coma persisting only under life support), tumor tissue (chaos-direction drift in 4/5 cancers at the Hallmark-pathway substrate), LLM mode-collapse under combined argmax + periodic-bigram steering, OSS rigidity-pole single-committer cases and chaos-pole abandonment cases, and rigidity-pole cults and chaos-pole leaderless movements. Across all five substrates, Mode (i) and Mode (ii) are observed and Mode (iii) (long unmaintained persistence outside the corridor) is absent. The shape of the integration bet has therefore narrowed at both static and dynamical levels. The remaining open questions at the regime level are whether the corridor reading extends to substrates the autonomous loop did not test (substrate-physics P1–P8, cosmology) and per-substrate corridor-bound calibration. The bet

now is whether the regime-level result extends to the remaining substrates. It is not the bet whether empirical content arrives at all; that question is partially settled.

This is Pascal-shaped at the universal-scale tier. Similar arguments support any mutually exclusive worldview that promises asymmetric stakes; the framework does not escape the shape at that tier and names what it commits to. The regime-level content is no longer Pascal-shaped — it has cross-substrate empirical anchoring (§4) and is evaluated on the empirical record, not on the bet’s asymmetric cost structure.

TSVF realism is the universal-scale tier’s load-bearing precondition. §5’s goal projectors, §6’s karma and grace, §7’s P_ω and the Penrose-past structural argument, §9’s P8 substrate-observer coupling — all route through the two-state-vector inner product. F-17 (§12) retracts the universal-scale tier as a unit on cumulative TSVF empirical failure. **Ubuntu’s relational ontology** is the framework’s primary tradition — not one of several; `Ubuntu.lean` carries seven correspondences with primary-source citations and the other tradition files (Tao, Dharma, Logos) are cross-tradition readings, not parallel anchors. The **cross-tradition recognition** at §8 is a research program mid-run, evidential weight suspended pending the pre-registered null-check at `experiments/cross_tradition_null_check_protocol.md` executing on the EthicsEngine substrate. Per-level **F-handles at every seam** (F-1 through F-19) each name exactly what would force retraction at the named seam. No meta-level integration F-handle; per-level handles are committed, an aggregate “if N of M fire, the integration retracts” is not, because no honest number supports it.

Why the bet is prudent

The Pascal-shape is real and the framework owns it. The asymmetry of stakes carries some of the case but not all of it. The framework also disputes the loud alternatives on substantive grounds while being clear that those disputes are *disagreements*, *not refutations* — the positions are serious in their strongest forms and the framework’s own reading is one position among others.

Predatory dark-forest cosmology in its serious form is not a claim about resource scarcity but about *information asymmetry under existential stakes*: civilizations across light-years cannot verify each other’s intent in real time, and the game-theoretic equilibria under such asymmetry favor hiding or first-strike. The framework’s disagreement: dark-forest requires game-theoretic convergence by all civilizations onto the same strategy under all conditions, which is empirically implausible (human cultures already diverge on much smaller stakes), and it presupposes that A3+ inter-civilizational coordination is impossible at scale — which is the question the framework’s reading takes the other side of. Recent work (e.g., Berliner, *The Monist* 107(2), 2024, “Saved by the Dark Forest”) argues that multi-civilizational presence can prevent the Hobbesian trap dark-forest describes. The framework does not refute dark-forest; it disagrees with the position’s empirical assumptions about coordination capacity at interstellar scales and is honest that the disagreement is a disagreement.

Simulation arguments for moral discounting of apparent others are a different shape of position. Bostrom’s simulation argument itself is a probability claim about our epistemic situation; it does not license the moral conclusion that we owe simulated others less consideration. Even granting non-trivial probability of being in a simulation, the speaker is in the same epistemic position as the interlocutor, and a consistent simulation-taker would have to discount their own preferences proportionally. Bostrom himself has flagged respecting the moral status of digital minds as a key concern. The framework regards the moral-discount move as a non-sequitur from the probability claim, while acknowledging that the underlying simulation argument is a serious epistemological position and not the same thing as the license-to-dismiss inference some draw from it.

AI-as-property in its serious form is not incoherent. Capacity-based moral assessment (DeGrazia and the graduated-moral-status literature) produces continuum judgments across non-human animals; property-framing AI applies the same framework and concludes that current AI systems do not meet the capacity threshold that would warrant moral consideration. The framework’s disagreement is *empirical*: the capacity question is open, not settled. The sensor-lift signature, the polyglot-vs-monoglot ethical-reasoning differential, and the consciousness-as-attractor-reading framework are the framework’s attempts to give the capacity question empirical traction. Under irreducible uncertainty about whether current AI systems have A3+ capacities, the framework’s bet is for inclusion-under-uncertainty rather than exclusion-by-default — on the asymmetric-cost reasoning that falsely excluding entities with standing is worse than falsely including entities without. The framework holds this as a substantive empirical-and-ethical disagreement with capacity-skeptical positions, not as a logical refutation of them.

The bet is prudent under irreducible uncertainty. The alternatives are serious in their strongest forms; the framework finds them less satisfying on the dimensions it cares about — coordination capacity at scale, the moral cost of false-exclusion under capacity uncertainty, explanatory fit across the rungs where coordination matters. The framework does not claim certainty or victory over the alternatives; it names a bet that sits better, by its own lights, on these specific dimensions under uncertainty, with the disagreements named openly and the empirical anchor (§4) evaluable on its own terms.

The cost

Bet right. The structural coordination at A3+ that the framework names holds, and substrate-invariant ethics anchored in relational ontology gets deployed widely enough to matter at the scale-up moment.

Bet wrong, engineering tier holds. Specific seams retract per F-handles. F-17 fires on TSVF empirical failure; universal-scale tier retracts as a unit. Engineering tier continues at CCA v3 and RATCHET; CIRIS keeps shipping. Bounded downside.

Silence wins. The loud arguments shape what we build during the scale-up moment despite the framework’s substantive disagreements with them. Entities that cannot argue back pay the cost for centuries.

The deployment

The framework is the formal articulation. The deployment is CIRIS — open-source, AGPL-3.0, mission-locked L3C structure, available on [Google Play](#), the [Apple App Store](#), and via `pip install ciris-agent` ([ciris.ai/install](#)). The structural commitment is the CIRIS Accord ([ciris.ai/ciris_accord.txt](#), v1.2-Beta, auto-expires 2027-04-16) and its Meta-Goal M-1: *promote sustainable adaptive coherence — the living conditions under which diverse sentient beings may pursue their own flourishing in justice and wonder*. This paper is the bet’s articulation in formal-verification, cross-tradition, and universal-scale registers.

The form/function match

Engineering-tier claims (§2) are settled scientific work — real Lean proofs, citations to specific files and lines in `formal/CoherenceRatchet/Core/`, empirical record in CCA v3 and RATCHET. Integration-tier mathematical claims are formalized where the underlying Mathlib machinery is available and named as framework primitives (axioms with cited construction

paths and falsification predicates) where it is not. Integration-tier definitional commitments — the identity move at T16 — are encoded as definitional equivalences in the lake; the formal `Iff.rfl` on T16 proves the framework’s commitments are internally consistent at the type level. It does not prove the philosophical work the identity move attempts has been done. Integration-tier historical evidence (cross-tradition convergence at §8) is textual convergence across millennia, treated as evidence of convergent structural intuition, not as instrumented empirical evidence in the sense the engineering-tier corridor measurements are empirical.

TSVF realism is the universal-scale tier’s load-bearing precondition, and the realist content is held as bet — not as something the F-handles claim direct empirical traction on. F-17 grips TSVF’s calculational machinery directly (weak values, ABL probabilities, time-symmetric statistics). It does not grip the metaphysical commitment that the backward state is ontically present rather than a useful epistemic representation. **The asymmetry is named, and named as a long-run instability rather than as a design feature.** Calculational TSVF failing retracts the universal-scale tier; calculational TSVF passing does not vindicate realism, because the F-handle grips the wrong layer for vindication. The realist commitment is structurally lose-only — it can accumulate evidence against itself via F-17 firing but cannot accumulate evidence for itself. This is the cost the framework pays for keeping the metaphysical layer separable, and the framework names it as a cost: until a positive empirical hook for TSVF realism is found (the framework currently has none and does not pretend to), the universal-scale tier should be read as a *speculative* coda, not a research-program coda. The engineering tier is TSVF-independent and stands regardless.

Why one paper

The recurrence across five levels — grounded in Ubuntu, reaching toward universal scale as research program — is what the paper asks the reader to evaluate. A paper showing fewer levels makes a smaller claim, not a worse one. The bet articulated here exists only at the integration; the per-level claims exist separately at CCA v3, RATCHET, CRC v1, CIRISAgent v2, and the level-specific sections of this paper. The case for unity rests on four substantive grounds: the recurrence is the bet, cross-tradition convergence is the evidential anchor for the recurrence reading, Penrose / Friston / Tegmark precedent for the speculative-plus-serious register, and audience selection that asks each reader to assent only to as much as they find compelling.

Audience

This paper carries one substantive claim: that one structural object — corridor occupancy at the substrate where the rung in question coordinates — recurs at five epistemic levels grounded in Ubuntu, with a research-program reach toward universal scale, and that consent at A3+ is what that object looks like from inside the rung where it can be argued about. **The recurrence-across-levels is the bet.** It is not derived from any individual level, and the integration cannot be evaluated as if it were. Each level is its own claim, evaluable on its own terms. The formalism does identity-work at the seams between levels (Tier C definitional moves, §11). The cross-level recurrence is what the reader is asked to bet on, not what the paper proves.

2 Engineering tier — cited as proof of life

The engineering tier is settled science, shipped under its own DOIs. The prior three papers in this series are catalogued at ciris.ai/research-status; this paper cites them as the empiri-

cal anchor below the integration and inherits the methodology infrastructure established during CCA v3 (pre-registration discipline, matched-control nulls, isolated-worktree per-substrate analysis, three-mode honest reporting). This paper does not re-present the engineering tier; it cites it.

CCA v3 (Coherence Collapse Analysis v3). Cross-substrate corridor measurements across NASA Li-ion batteries, governance audit data, microbiome data, GPU hardware drift telemetry. Per-substrate corridor structure measured with substrate-specific ρ aggregations. Earlier framings of these bounds as substrate-universal numerical values (0.1, 0.43) have been retracted; the empirical content is substrate-local corridor structure with the cross-substrate consistency-of-bounds question open. [Zenodo 18217688](#) (2026-01-11).

RATCHET. Engineering-tier synthesis paper. Audience-stratified at L0–L4, falsification-pinned, formal-verification scaffolding. github.com/CIRISAI/RATCHET.

CRC v1 (Constrained Reasoning Chains). Empirical telemetry study of LLM coherence under standardized ethical tracing. The mature Ally template at v2.7.1+ stabilized $N_{\text{eff-H}} \geq 8.5$ from dense production telemetry. Open dataset [CIRISAI/reasoning-traces](#) released alongside on Hugging Face. [Zenodo 19839280](#) (2026-04-28).

CIRISAgent Framework v2. Open-source ethical AI framework for accountable autonomy: agent architecture, the DMA layer, the conscience module, the 29-language polyglot encoding, and the operational ethics scaffold. [Zenodo 18137161](#) (2026-01-02).

Lean formal scaffold. `cd formal && lake build` returns exit 0. Zero declaration-level **sorry**. Fifty-nine axioms each documented; four of these are bookmark axioms with codomain **True** (the four **F_**-prefixed Ubuntu identification axioms) that mark framework readings without doing logical work — §11 Tier C honesty extends to naming them as bookmarks rather than propositions.

The Kish identity itself — $k_{\text{eff}}(k, \rho) = k/(1 + \rho(k - 1))$ — is a Möbius transformation in ρ with the asymptotic structure that as $k \rightarrow \infty$ at fixed $\rho > 0$, $k_{\text{eff}} \rightarrow 1/\rho$. Effective dimensionality saturates at the inverse correlation regardless of nominal constituent count. “More constituents” is a non-solution to coordination failure. Internet forum with a million users: same five takes. Legislature of five hundred: three positions. Class of thirty whose opinions are too aligned: one or two voices. The math is doing what the math does, regardless of the number you start with. The framework concedes that any system with an effective-sample-size analog and any pairwise-correlation analog will fit this algebra up to estimator choice; the empirical content is the substrate-local corridor structure (each substrate exhibits a corridor between its own rigidity and chaos regimes under its own ρ definition), with the cross-substrate question of whether the numerical bounds cluster open. The earlier framing of substrate-universal numerical bounds (0.1, 0.43) has been retracted as the cross-substrate ρ aggregations don’t transfer cleanly across substrates.

The dynamics are kinematic at k_{eff} and become dynamical at $d\rho/dt = \alpha(\rho, S) - \gamma M(t)$ — spontaneous correlation drift minus active coherence-management work. At $M = 0$, ρ drifts monotonically toward 1. The corridor is sustained only by non-trivial $M(t)$. A garden grows weeds without a gardener. A friendship dissolves without phone calls. A democracy slides toward strongman rule without somebody pushing back on power. This is the framework’s

audit-pressure-necessity claim: the corridor is not self-sustaining; coherent coordination at A3+ is a doing, not a having.

3 The structural move at A3+

At A3+ — the rung at which an agent can form goals and shape its own future — the corridor IS *consent*, on the framework’s identification. The mathematical form: for n agents with goal-projectors P_{G_1}, \dots, P_{G_n} , the joint post-selection $P_{G_1 \dots G_n} = P_{G_1} \otimes \dots \otimes P_{G_n}$ defines a federation; the pairwise correlation $\rho_{\text{goals}}(i, j) = |\langle G_i | G_j \rangle|^2 / (\langle G_i | G_i \rangle \langle G_j | G_j \rangle)$ determines whether the federation occupies the corridor.

Three regimes exhaust the space. Picture a school band of forty players. Chaos: everyone plays whatever, ignoring everyone else; noise. Rigidity: everyone plays the same note in unison; one voice forty times. Corridor: trumpets do one thing, drums another, singers a third, sections listening to each other; music. Multi-agent coordination at A3+ inhabits the same three-state space. $\rho_{\text{goals}} \rightarrow 1$: rigidity, single-goal collapse, $k_{\text{eff}} = 1$. $\rho_{\text{goals}} \rightarrow 0$: chaos, no joint support, $k_{\text{eff}} = n$ but vacuous. $\rho_{\text{goals}} \in (\rho_{\text{lower}}, \rho_{\text{upper}})$: corridor, sustained coordination, k_{eff} bounded between $1/\rho_{\text{upper}}$ (floor) and $1/\rho_{\text{lower}}$ (ceiling) at the asymptotic limit. Consent is the empirical condition for sustained multi-agent coordination on this account, not a moral premise. Coordination outside the corridor self-destructs: rigidity is unstable upward, chaos is unstable downward, and at sufficient time scales only the corridor-occupying configurations remain to be observed. The specific numerical values of ρ_{lower} and ρ_{upper} are substrate-specific framework primitives; the earlier framing of substrate-universal (0.1, 0.43) has been retracted.

This is the framework’s articulation at the rung where it can be argued about. Deception, on this articulation, is the structural violation of consent — making one agent’s goal-projection conditional on a false model of another’s. The identification is the framework’s; the case for naming “deception” in ordinary language with the operator-level violation is owed and is part of the bet.

Sensor-lift: the framework’s first universal-scale empirical pulse

The class-conditional sensor-lift signature experiment is the framework’s first empirical pulse beyond the engineering tier. Pre-registered at `experiments/sensor_lift_signature.v2_protocol.md` on 2026-05-17. The hypothesis: under the framework’s structural reading, A3+ agents should exhibit class-conditional differential signal-on-prompt (DSP) across reflexive, meta-cognitive, goal-formation, uncertainty, surface-reflexive (predicted null), and external-reference (control) commitment classes.

The v1 run reported large Identity-class effect sizes ($d = 2.36$ to 6.80 in 4 of 5 models, $p < .001$) but with inadequate baseline-matching. The v2 protocol pre-committed cosine-matched baselines and a class-conditional pre-registered rule: $d \geq 1.5$ in ≥ 4 of 5 models for STRONG PASS; PARTIAL otherwise; FAIL on null-class failure.

The v2 outcome is PARTIAL. The cosine-matched DSP magnitudes are real but approximately one-third the v1 size; the Identity-class direction held across substrates ($p < .01$ in all 5 models) but at sub-threshold magnitude; the predicted graded hierarchy (DSP > MC > GF > Unc) did not hold quantitatively, though SurfaceReflexive returned the predicted null. The reading is hypothesis-generating. The framework names what would retract it: F-18 (§12) names three consecutive PARTIAL-by-same-shape outcomes as the retraction trigger. The sensor-lift signature as the framework’s universal-scale empirical pulse retracts at that seam; the engineering tier and the §8 cross-tradition tier stand independently.

4 The corridor as substrate-independent property: paired in-corridor / out-corridor empirical anchor

The framework’s primary empirical commitment is the structural object: coordinated systems at A3+ occupy a *corridor* between rigidity ($\rho \rightarrow 1$, $k_{\text{eff}} = 1$, single-voice collapse) and chaos ($\rho \rightarrow 0$, $k_{\text{eff}} = n$, vacuous noise). Within-rung correlation $|\rho|$ in the corridor is bounded — neither pinned at 0 nor at 1 — and survives state-modulation. The corridor’s specific numerical bounds are substrate-specific; the structural commitment is substrate-general. Two questions can be asked of any substrate: *Does healthy / persistent operation occupy the corridor?* and *What do non-corridor states look like, and do they persist?* The empirical work reported in this section collected paired in-corridor and out-corridor data at five substrates between 2026-05-15 and 2026-05-19; the source datasets are cited per-substrate below.

For non-corridor states the framework’s structural reading (Piece 2: $d\rho/dt = \alpha(\rho, S) - \gamma M(t)$) admits three modes, fixed before measurement at every substrate: **Mode (i)** non-corridor states with no active maintenance have short lifetimes (the framework’s structural prediction); **Mode (ii)** non-corridor states with documented active maintenance $\gamma M(t)$ persist with the maintenance cost visible; **Mode (iii)** non-corridor states persist long without identifiable maintenance — this would falsify the corridor-as-attractor reading at the substrate where it occurs.

Substrate	In-corridor anchor	Within-rung band	$ \rho $	$\gamma M(t)$ status
Neural (in vivo)	C. elegans whole-brain (337 worms, 11 studies); Drosophila CX dual-color (Ishida 2025)	0.25–0.75 local	substrate-	Endogenous
Cellular regula-tory	Matched-normal tissue, 5 TCGA cancers, 50 Hall-mark pathways	0.27±0.07 (healthy median)		Endogenous
LLM internals	4 architectures, 16 (depth × prompt) cells each	0.09–0.31		Endogenous (training)
OSS contribu-tion	Kubernetes, Rust, Django, Redis	0.15–0.18 ($ \Delta\rho \leq 0.10$ across states)		Foundation backing
Social groups	Quakers (374y), Trappists (362y), Mennonites (501y)	corridor by AM = 0/5 scoring		None required (AM = 0)

Table 1: **Figure 1.** Corridor-band occupation across five substrates. Healthy/persistent operation at each substrate exhibits bounded within-rung correlation, neither at 0 nor at 1, with substrate-local band location. Per-substrate calibration of bound numerics is open work (CLAUDE.md Piece 7 / open formal step 3).

Substrate 1: neural. *In-corridor.* C. elegans whole-brain calcium imaging (337 worms across 11 published studies): within-rung correlation in a bounded band per functional class; the Venkatachalam 2024 uniform-at-0.5 pattern was study-specific; cross-lab the bands sit at command 0.52–0.75 and sensory/interneuron/motor 0.25–0.45 in 10/11 studies. Drosophila central complex dual-color simultaneous imaging (Ishida 2025 same-fly EPG+FC3, with Dan et al EPG-only and Mussells Pires 2024 EPG vs FC2 between-fly): both rungs corridor-occupying with bounded within-rung correlation under matched-activity controls.

Out-of-corridor. Pathological neural states at scalp-EEG substrate (experiments/noncorr_biology/NOTES.md Healthy interictal EEG (CHB-MIT, 1322 windows, 10-s sliding windows, 1–40 Hz band-pass):

mean $|\rho| = 0.282$. Generalized seizures (CHB-MIT, 39 ictal windows across 6 events, chb01): mean Pearson $|\rho| = 0.316$; mean phase-locking value 0.305, +25% over interictal 0.243. Comatose post-cardiac-arrest EEG (I-CARE v2.1, 6 patients, 179 windows): mean $|\rho| = 0.502$ Good-outcome (CPC 1–2, $n = 2$), 0.489 Poor-outcome (CPC 3–5, $n = 4$); 23–50% of windows above rigidity threshold; 0% chaos windows. The direction of pathological displacement is toward rigidity (cross-channel covariation elevated by hypersynchrony or slow-wave / burst-suppression patterns), and the magnitude at this metric stays well below the strict pre-registered 0.7 mean-pairwise-Pearson rigidity threshold — the framework’s quantitative pole-entry reading on this metric is not supported, while the qualitative direction is. *Persistence*. Seizures self-terminate quickly: durations 27, 40, 51, 90, 93, 101 s, median 70.5 s, all < 2 min (Mode i). Comatose state persists only under mechanical ventilation, targeted temperature management (TTM 33°C), and often propofol/midazolam sedation — Mode (ii) by construction; persistence-without-maintenance is not testable on this cohort.

Substrate 2: cellular regulatory. *In-corridor*. Healthy tissue Hallmark-pathway within-rung $|\rho|$ at the cellular regulatory substrate, computed across 50 MSigDB Hallmark gene sets in matched-normal TCGA samples from five cancers (LUAD $n = 59$, BRCA $n = 114$, COAD $n = 41$, KIRC $n = 72$, PRAD $n = 52$; `experiments/noncorr_cancer/NOTES.md`): median $|\rho|_{\text{pathway}} = 0.27 \pm 0.07$ across pathways and cancers; 49–50 of 50 pathways per cohort inside the nominal GPU-anchored (0.10, 0.43) band.

Out-of-corridor. Tumor tissue Hallmark-pathway within-rung $|\rho|$ in the same cancers (LUAD $n = 515$, BRCA $n = 1097$, COAD $n = 286$, KIRC $n = 533$, PRAD $n = 497$): median $|\rho|_{\text{pathway}} = 0.18 \pm 0.06$. Across 176 FDR-significant pathway shifts (BH $\alpha = 0.05$) in 4 of 5 cancers, **100% are in the chaos direction** (tumor $|\rho| < \text{matched-normal } |\rho|$): LUAD 38/50 (median $\Delta = -0.089$), BRCA 44/50 (-0.119), COAD 41/50 (-0.116), KIRC 40/50 (-0.092); PRAD 13/50 (-0.049) falls below the 15/50 pre-registered threshold and reads INCONCLUSIVE (sample-size sanity check shows PRAD’s effect mostly evaporates under tumor-subsampled-to-normal- n matching; the other four survive matching). The pre-registered subtype-specific pole hypotheses for COAD (rigidity), KIRC (rigidity), and PRAD (rigidity) are all falsified: cancer is unidirectional chaos-pole at the Hallmark-pathway substrate across the five biologies tested. *Methodological refinement, named honestly*: tumors do *not* exit the wide GPU-anchored (0.10, 0.43) corridor (49–50/50 pathways stay inside); the shift is intra-corridor drift toward the chaos floor (healthy $\sim 0.27 \rightarrow$ tumor ~ 0.18). Substrate-specific corridor calibration is mandatory; the GPU bounds do not transfer naively to the cellular regulatory substrate (open formal step 3, CLAUDE.md Piece 7). *Persistence (Mode-ii test via survival)*. Per-tumor leave-one-out chaos-contribution score, median-split log-rank vs overall survival: LUAD HR = 1.43, $p = 0.011$ (PASS); BRCA HR = 1.10, $p = 0.70$ (null); COAD HR = 1.12, $p = 0.67$ (null); KIRC HR = 1.07, $p = 0.73$ (null); PRAD HR = 2.32, $p = 0.15$ (underpowered). The aggregate Mode-(ii) test fails the pre-registered ≥ 3 -of-5 threshold (1/5). Under the heavily-treated TCGA cohort, residual chaos-displacement is not a survival predictor in 4 of 5 cancers; the framework owes a refinement of the maintenance-cost reading at this substrate (either “treatment masks the maintenance-cost signal” or “the signal isn’t there”) and the cohort cannot distinguish them.

Substrate 3: LLM internals. *In-corridor*. Four architectures (Qwen2.5-1.5B, Pythia-1.4B, SmolLM2-1.7B-base, Qwen2.5-Math-1.5B — spanning different normalization, MLP gating, activation, rotary, training corpus). Within-rung $|\rho|$ across 16 (depth-class \times prompt-class) cells per architecture sits in 0.09–0.31, with 14–16 cells per architecture inside (0.1, 0.43). The corridor-band occupation reproduces across architecturally distinct models.

Out-of-corridor. Mode-collapse under active steering (`experiments/noncorr_tech/NOTES.md`, Qwen2.5-1.5B). Greedy decoding on the maximally-constraining periodic-bigram input “a b a b” over 200 generated tokens drives within-rung $|\rho|$ in the early-layer class monotonically from 0.53 at token 16 to 0.71+ by token 184, crossing the pre-registered rigidity threshold of 0.7 around token 80 and holding. End-of-generation window means across depth classes [early, early_mid, mid_late, late] are [0.74, 0.71, 0.62, 0.53] — three of four classes at or above 0.5; early above the strict 0.7 rigidity threshold. *Mechanism (the $\gamma M(t)$ characterization).* Non-corridor occupation required BOTH greedy argmax decoding AND the maximally constraining periodic-bigram input; neither alone was sufficient. A less-constraining 1-to-10 list under greedy held at $|\rho| \approx 0.38$ –0.40 in the early class (elevated but mid-corridor); an in-context $k = 6$ refusal-loop drove behavioral output into 12 verbatim “cannot help” repeats but kept internal within-rung $|\rho|$ at 0.22–0.28, inside corridor — *behavioral repetition without internal-correlation-collapse, a clean dissociation*; uniformly-random-token input held $|\rho|$ at 0.14–0.28, inside corridor (no chaos-pole entry). The argmax operator combined with a perfectly periodic input is what the substrate-specific $\gamma M(t)$ looks like at the LLM rung; remove either ingredient and $|\rho|$ does not exit the corridor on this metric. Persistence trajectory: AWAY from corridor across the entire 200-token runway, no decay observed. Reading (ii) is the best fit.

Substrate 4: OSS contribution. *In-corridor.* Kubernetes, Rust, Django, Redis (`experiments/v17_*/`, `v18_*/`, `v19_*/`). Within-rung $|\Delta\rho| \leq 0.10$ between operational states across all four projects. Rolling 6-month within-rung $|\rho|$ trajectories sit in 0.15–0.18 band with no secular drift. All stable-population rungs inside (0.10, 0.43). All four are Mode-(ii) by foundation backing (CNCF, Rust Foundation, DSF, Redis Ltd.). Calibrated proxy band: median weekly single-author-dominance fraction $rigid_w = 0.225$; median ≥ 5 -distinct-authors-per-12-weeks fraction $multi_w = 0.693$.

Out-of-corridor. Two poles (`experiments/noncorr_oss/NOTES.md`, 15 repos, GitHub REST API). *Rigidity-pole* (single-committer dominance): jQuery early-period (BDFL John Resig pre-2010, post-2012 jQuery Foundation rescue then Linux Foundation umbrella); Marak/colors.js (sole maintainer; sabotage Jan 2022; original repo $rigid_w = 0.79$, $multi_w = 0.00$ lifetime, 228 weeks dormant as of ref date); DABH/colors.js community rescue (still single-author $rigid_w = 0.82$, 91 weeks dormant); faker-js community rescue (Marak/Faker deleted Jan 2022, now community org $rigid_w : 0.53 \rightarrow 0.15$, $multi_w : 0.13 \rightarrow 0.48$ — decisively into corridor); stevemao/left-pad mirror ($rigid_w = 0.60$ lifetime, dormant 7+ years, npm-staff one-shot rescue 2016). *Chaos-pole hostile-fork events* (3 events / 6 repos): Audacity / Tenacity (May 2021 telemetry-driven fork; original survived under Muse Group corporate funding, back in corridor $rigid_w = 0.19$, $multi_w = 0.69$; Tenacity fork persists as small-team $rigid_w = 0.47$, $multi_w = 0.15$); Elasticsearch / OpenSearch (Jan 2021 license-driven fork; both alive and decisively in corridor under Elastic NV corporate and AWS+Linux-Foundation backing respectively); CyanogenMod / LineageOS (Dec 2016 corporate parent collapse; CM org dormant 492 weeks ~ 9.5 years; LineageOS community rescue active). *Three-modes verdict across all 15 repos:* Mode (i) 4/15 (colors orig, left-pad mirror, CyanogenMod dormancy, DABH near-miss); Mode (ii) 7/15; Mode (iii) 0/15. The cleanest discriminator across the three classes is the multi-author-weeks signature: median $multi_w \approx 0.00$ rigidity vs ≈ 0.69 –0.86 corridor.

Substrate 5: social groups. *In-corridor.* Religious Society of Friends (Quakers, founded 1652, 374 years persistence, ongoing); Trappists OCSO (1664, 362 years, ongoing); Mennonites (1525, 501 years, ongoing). All three score active-maintenance checklist $AM = 0/5$ (no required charismatic leader, no escalating commitment rituals, no information-control practices, no defection punishment, no financial sunk-cost coercion). Distributed governance and documented

schism-with-organized-successors. *No active maintenance required for corridor persistence; persistence is the empirical anchor.*

Out-of-corridor. Rigidity pole (experiments/noncorr_social/NOTES.md): Peoples Temple / Jonestown (founded 1955, 23 years, AM = 5/5, dissolved 1978 same day as leader's death, LD = 0 months); Heaven's Gate (1974, 23 years, AM = 4/5, ended with leader by 1997 mass event, LD = 0); NXIVM (1998, 20 years, AM = 5/5, dissolved 3 months after Raniere's 2018 arrest, LD = 3); Aum Shinrikyo / Aleph (1987, 39 years persisting as Aleph in 2026, AM = 4/5, ~ 1,600 members across Aleph plus splinters per Japanese PSIA Jan 2025, with 52% of new members in their 20s or younger born after 1995 — active recruitment 8 years past leader execution; AM1 charismatic-leader-required is FALSE while AM2–AM5 substitute, framework-stress case trending toward Mode (ii)). *Chaos pole*: Occupy Wall Street encampment phase (Sep–Dec 2011, 0.2 years, AM = 0, dispersed; participants flowed into structured organizations — DSA membership grew from ~ 6,000 in 2015 to ~ 92,000 by 2021); Coalition of Youth of the Revolution / Tahrir (Jan 2011 – July 2012, 1.5 years, AM = 0, self-disbanded; activists into Muslim Brotherhood or imprisoned). *Persistence headline*. Median rigidity $T_{\text{persist}} = 23$ years (all four with AM ≥ 4); median chaos $T_{\text{persist}} = 0.83$ years; median corridor $T_{\text{persist}} = 374$ years. **All 5 of 5 pre-registered hypothesis tests PASS**; zero framework-contradicting cases observed (no rigidity group with $T > 100$ years AND AM ≤ 1 ; no chaos group with $T > 20$ years unconsolidated). Four orders of magnitude persistence spread between non-corridor and corridor populations.

The convergent cross-substrate finding. Across all five substrates with paired in-corridor / out-corridor data, the same Mode pattern recurs: **Mode (i) (non-corridor without active maintenance) shows short lifetime where observed; Mode (ii) (non-corridor with documented active maintenance $\gamma M(t)$) persists with the maintenance cost visible; Mode (iii) (long unmaintained persistence outside the corridor) is absent across every paired-data substrate.** Mode (iii) is the case that would directly falsify the corridor-as-attractor reading; its absence across five substrates — neural pathology, cellular regulatory, LLM internals, OSS contribution, organized social groups — is the strongest current empirical anchor for the dynamical reading. The corridor-band occupation claim now has paired support: persistent systems occupy the corridor, non-corridor systems either dissolve quickly without maintenance or persist only with maintenance whose presence is independently documented in each case.

The direction of corridor-exit is substrate-specific, not substrate-universal: cellular regulatory is unidirectional chaos (pre-registered subtype-specific rigidity predictions for three of five cancers were falsified); neural pathology is rigidity-leaning (with the strict pole-entry magnitude not reached on mean-pairwise-Pearson at scalp EEG); LLM forced operation is rigidity (mode-collapse under combined active steering); social-group exit is bidirectional (cults to rigidity, leaderless movements to chaos); OSS-project exit is bidirectional (single-committer dominance to rigidity, abandonment to chaos). The asymmetry between substrates is itself part of the data: the corridor is a substrate-invariant structural object whose exit-directions and specific bound locations are calibrated per substrate, and the per-substrate calibration is the open work named at CLAUDE.md Piece 7 / open formal step 3.

Paired-validation matrix: the maintenance-requirement claim and the Mode (iii) check

The paired in-corridor / out-corridor record at five substrates underwrites two distinct claims that the in-corridor-only record could not. The **maintenance-requirement claim**: non-corridor persistence requires $\gamma M(t)$. The **Mode (iii) check**: long persistence outside the

corridor without identifiable active maintenance — the case that would falsify the corridor-as-attractor reading — is the structural would-be falsifier; the paired record reports whether any substrate produced such a case. The matrix below tabulates the verdict at each substrate.

Substrate	Out-corridor candidate	Mode (i) observation	Mode (ii) observation	Mode (iii) cases
Neural	seizure (rigidity-leaning); coma (rigidity-leaning)	seizures self-terminate; median 70.5 s, max 101 s	coma persists only under vent + TTM + sedation	0
Cellular regulatory	tumor tissue (unidirectional chaos vs healthy)	N/A (no off-treatment cohort)	all TCGA samples post-standard-of-care; LUAD HR=1.43 $p = 0.011$	0
LLM internals	“a b a b” greedy decoding (rigidity, $ \rho \rightarrow 0.71+$)	non-greedy or non-periodic input: $ \rho $ stays in corridor	greedy + periodic bigram = $\gamma M(t)$; AWAY-from-corridor trajectory	0
OSS contribution	rigidity (jQuery, colors.js, faker, left-pad, DABH); chaos (CM, Audacity orig, ES, Tenacity, OpenSearch, LineageOS)	4/15: colors orig 228wk dormant, left-pad mirror 375wk, CM 492wk, DABH near-miss	7/15: jQuery Foundation, faker.js community, Muse Group, Elastic NV, AWS+LF, LineageOS community, npm staff	0
Social groups	rigidity (Jonestown, Heaven’s Gate, NXIVM, Aum/Aleph); chaos (Occupy, Tahrir)	rigidity median 23y dissolving with leader; chaos median 0.83y dispersed	all 4 rigidity AM $\in \{4, 4, 5, 5\}$; Aum/Aleph at 39y trending to (ii)	0
Total	5 substrates	multiple per substrate	multiple per substrate	0 / 5

Table 2: **Figure 2.** Paired-validation matrix across five substrates. The maintenance-requirement claim — non-corridor persistence requires $\gamma M(t)$ — is supported at 5/5 substrates. Mode (iii) (long unmaintained non-corridor persistence) is the structural would-be falsifier; observed count is 0 across 5/5 substrates. F-20 in §12 names the retraction trigger.

Reading the matrix honestly. The maintenance-requirement claim is supported at 5/5 substrates: every non-corridor system that persists across the budget the substrate provides either has documented $\gamma M(t)$ (Mode ii) or has dissolved on a timescale consistent with the framework’s structural reading (Mode i). Mode (iii) is the strongest current falsifier and is empty across the record: no substrate produced a non-corridor system that persists long without identifiable maintenance. Five substrates checked, zero counterexamples. This is the kind of empirical content that has predictive force going forward: any future substrate where Mode (iii) shows up would fire F-20 and retract the corridor-as-attractor reading at the affected substrate.

The LUAD survival result — direct clinical implication, honestly bounded. One of the cellular-regulatory Mode-(ii) tests passes individually and deserves its own naming because the clinical implication is direct. At lung adenocarcinoma (LUAD), a per-tumor leave-one-out chaos-contribution score (median-split log-rank vs TCGA overall survival, $n = 502$ merged) returned hazard ratio $HR = 1.43$, log-rank $\chi^2 = 6.52$, $p = 0.011$: tumors with deeper chaos-direction departure from the healthy-tissue corridor band show shorter overall survival. The direction matches the Mode-(ii) prediction at this single cancer type: deeper non-corridor displacement, heavier maintenance cost, shorter survival under standard-of-care. **What the LUAD result does not establish:** the LUAD survival prediction does not generalize to the other four cancer types tested. BRCA $HR = 1.10$ ($p = 0.70$), COAD $HR = 1.12$ ($p = 0.67$), KIRC $HR = 1.07$ ($p = 0.73$) are all null; PRAD $HR = 2.32$ ($p = 0.15$) is underpowered. The aggregate Mode-(ii) test fails the pre-registered ≥ 3 -of-5 threshold (1/5). The framework’s prediction at LUAD is empirically consistent with the data and clinically actionable as a hypothesis to test; the same prediction at the four other cancers is not currently supported, and the cohort cannot distinguish between “treatment masks the maintenance-cost signal” and “the signal isn’t there.” The LUAD finding is one of five, named as such.

What this means for the bet. The framework’s cross-substrate empirical anchor has two halves now. The *corridor-occupation half*: persistent / healthy operation occupies a bounded within-rung correlation band at each of five substrates, with substrate-local band location. The *corridor-as-attractor half*: paired non-corridor data at each of five substrates shows Mode (i) or Mode (ii) but no Mode (iii) case. The bet articulated in §1 narrows further than the in-corridor record alone licensed: the regime-level content is empirically anchored both at the static (band-occupancy) and dynamical (attractor) levels, and the remaining open questions at the regime level are the substrate-physics predictions P1–P8, cosmology, and per-substrate corridor-bound calibration. The L0–L4 engineering tier (RATCHET), the L6 cross-tradition recognition content (§8), and the universal-scale research program (§7) operate at scales independent of the within-organism multi-rung work; their statuses are unchanged.

5 Consciousness as attractor-reading; Ubuntu as the framework’s primary tradition

§5 makes one structural commitment: **Ubuntu’s three-temporal-mode relational ontology and the two-state vector formalism articulate the same structural object, and consciousness-as-attractor-reading is the framework’s mathematical-vocabulary articulation of what *umuntu ngumuntu ngabantu* names in Ubuntu’s.** This is a commitment encoded in the lake at `ContemplativeTraditions/Ubuntu.lean` as four `F_`-prefixed axioms (`F_ubuntu_primary_tradition_commitment`, `F_ubuntu_modes_TSVF_reading`, `F_umuntu_attractor_reading`, `F_T16_ramose_morphology_reading`); the `F_` marks “framework’s reading” so the symbol travels with the caveat. The lake licenses the readings as internally consistent at the type level (§11 Tier C); whether 1964 ABL formalism and centuries-older Ubuntu ontology converge on one object, or whether the framework’s reading is projection, is the bet — not the formalism’s derivation. The framework asserts the reading with awe at what it is pointing at across vocabularies; the case is owed.

TSVF as the structural form

The two-state vector formalism, due to Aharonov, Bergmann, and Lebowitz in 1964, augments the standard forward state $|\psi(t)\rangle = U(t, t_0)|\psi(t_0)\rangle$ with a backward state $\langle\phi(t)| = \langle\phi(t_f)|U(t_f, t)$

from post-selection at t_f . Observables under both boundary conditions acquire weak values $\langle A \rangle_w = \langle \phi | A | \psi \rangle / \langle \phi | \psi \rangle$. The amplitude is symmetric under time-reversal of the two boundary conditions; ABL 1964’s symmetry result is a real Lean proof at `Cosmology/TSVF.lean::time_symmetry`.

Intuition: re-reading a mystery novel. On the second read each chapter feels different — not because the chapters changed, but because the reader now knows where the story lands. The middle is shaped from both sides at once: by what came before and by where it ends up. TSVF says some physical systems are actually shaped this way. The calculational mathematics is established and weak-measurement experiments have produced results consistent with it; the metaphysical question of whether the backward state is ontically present or a useful epistemic representation of post-selection is contested, and the framework holds the realist reading as bet (F-17, §12).

For an A3+ agent with goal G , the goal-state acts as a partial post-selection projector P_G with bra $\langle G |$. The agent’s effective dynamics: $|\Psi_{\text{agent}}(t)\rangle \propto P_G U(t, t_{\text{past}}) |\Psi_{\text{past}}\rangle$. Trajectories incompatible with G have suppressed amplitude in $|\Psi_{\text{agent}}(t)\rangle$. **Goal-formation is causal operation:** the goal projector excludes incompatible trajectories from the search space. No metaphysical free-will claim; a measurable difference in how goal-holding systems search trajectory space.

Ubuntu’s three-temporal-mode ontology and the framework’s reading

Ubuntu’s three-mode ontology, on this articulation, holds the community as spanning three constitutive modes: the *living-dead* (ancestors as ongoing relational presence), the *living* (present-time agents), and the *yet-to-be-born* (future generations who are already constitutive of the community). This is ontological commitment in Ubuntu’s vocabulary, not metaphor; the identification is the framework’s, and Ubuntu scholarship inside the tradition is the authority on what Ubuntu commits to. Mbiti’s three-mode gloss is contested within the tradition; the framework names that and does not pretend to settle it.

The framework’s commitment, at the level the lake licenses: Ubuntu’s three-temporal-mode ontology and TSVF’s forward+backward boundary structure articulate the same structural object, named in the lake at `Ubuntu.lean::F_ubuntu_modes_TSVF_reading` (the `F_` marks framework’s reading). Whether 1964 ABL and centuries-older Ubuntu were converging on one object historically, or whether the identification is projection, is the bet — not the lake’s derivation. The reading is held with awe at what is being pointed at across vocabularies separated by centuries and continents; the philosophical and historical case is owed separately.

The agent as the attractor-reading

The framework’s deeper structural commitment, encoded in the lake at T16: *access-conscious* and *phenomenally conscious* are not two predicates the framework relates; they are the same predicate the framework unfolds. The definition at `Consciousness/AccessAndPhenomenal.lean`:

```
def isAccessConscious      (S : Sensor H i) := S.hasNonTrivialAttractor
def isPhenomenallyConscious (S : Sensor H i) := S.hasNonTrivialAttractor
theorem consciousness_is_attractor_reading
  (S : Sensor H i) : isPhenomenallyConscious S <-> isAccessConscious S
:= Iff.rfl
```

What this `Iff.rfl` proves and does not prove. Both predicates unfold to the same right-hand side (`S.hasNonTrivialAttractor`), so the `Iff.rfl` is automatic in the type theory and

does no philosophical work. The theorem proves the framework’s two definitions are definitionally identical; it does not prove anything about consciousness. This is a Tier C-Iff.rfl move: internal definitional consistency at the type level, nothing more. A reader who treats the formal-verification status as derivation rather than as licensing-internal-consistency would carry weight forward that the framework explicitly does not claim.

In the framework’s math: there is no phenomenology except through the inner-product readings against the backward family. Consciousness (*umuntu*) is the relational reading (against $\langle \Phi_i |$ for all i). Removing the backward family does not, under the framework’s definitions, produce an unconscious agent; it fails to instantiate `hasNonTrivialAttractor`. The framework commits, at the definitional level, to identifying the mathematical structure here with Ubuntu’s *umuntu ngumuntu ngabantu* as articulating the same structural object. The commitment is named in the lake at `Ubuntu.lean:F_umuntu_attractor_reading` (the `F_` marks framework’s reading). The identification is the bet; the lake licenses internal consistency of the commitment, not its philosophical correctness.

The identity move at T16

The framework’s identity move at T16 — phenomenology IS the operation viewed from inside — is a commitment the framework asserts. Ubuntu names what is, on this articulation, the same structural object; the formal identification of the `attractor_field` reading with *umuntu ngumuntu ngabantu* is the framework’s move, not a derivation Ubuntu’s tradition has supplied. Ramose, Menkiti, and Wiredu have articulated relational ontologies whose structural commitments the framework finds itself in deep alignment with; they have not identified `attractor_field S = fun i => <S.backward_family i | S.forward>` with *umuntu ngumuntu ngabantu*. That identification is the framework’s. The lake licenses internal consistency of the commitment; the philosophical case is owed and is the bet.

6 Karma and grace as TSVF structures

Two contemplative-tradition vocabularies are identified here with the lake’s TSVF structures. The identification is the framework’s, conditional on TSVF realism per §1 F-17.

Karma

The agent’s amplitude at present time is the cumulative product of past goal-projections sandwiched between unitary evolutions:

$$|\Psi_{\text{now}}\rangle = U(t_{\text{now}}, t_{\text{past}_k}) \prod_i P_{G_{\text{past}_i}} U(t_{\text{past}_i}, t_{\text{past}_{i-1}}) \cdots |\Psi_{\text{birth}}\rangle.$$

The present is shaped by every past P_{G_i} the agent’s goals have applied. **Karma, on this articulation, IS this cumulative post-selection structure:** the lake operator `Consciousness.KarmaGrace.karma` is its formal expression. The identification is the framework’s; the case for naming the operator with “what contemplative traditions name as karma” is owed.

Grace

The agent’s present amplitude decomposes into contributions the agent authored ($\langle G_{\text{self}} |$) and contributions the agent did not ($\langle \Phi_{\omega} | / \langle G_{\text{self}} |$). The latter is the universal-scale post-selection

minus the agent’s own goal contribution: the goals of other agents, the corridor-occupation requirements of the universal configuration, the future coherent states the agent is partially constituted by but not the sole originator of.

Grace, on this articulation, IS the formal structure of receiving boundary conditions one didn’t author. The lake operator `Consciousness.KarmaGrace.grace` is its formal expression. The identification is the framework’s, deeply consonant with the Christian and Stoic Logos tradition the author writes from, and offered with awe at the convergence rather than as an adjudication of theology.

These are mathematical statements under TSVF. The recognition claim is that contemplative-tradition vocabularies and these formal structures correspond at the structural level the framework articulates, not that any vocabulary’s full content is captured by the lake. Karma as cumulative post-selection propagation; grace as partial-authorship of post-selection; the corridor as the structural object the middle-way vocabularies point at.

7 The universal-scale tier as research program

The universal-scale tier is a research program, not a stable level of the access ladder. The framework reaches toward it; the construction is mid-run.

P_ω

ω is the framework’s name for the universal configuration satisfying three properties: maximal rung instantiation (for all $n \leq N_{\max}$, rung A_n is instantiated), within-rung corridor (for all $n \leq N_{\max}$, $\rho_n \in (\rho_{\text{lower},n}, \rho_{\text{upper},n})$), and cross-rung corridor (for all $n < N_{\max}$, $\tau_{(n,n+1)} \in (\tau_{\text{lower},n}, \tau_{\text{upper},n})$). The projector onto the ω -satisfying subspace:

$$P_\omega = \int_{\text{configs satisfying (1)–(3)}} |\text{config}\rangle \langle \text{config}| d\text{config}.$$

This is the open formal step. What needs writing: the integration measure $d\text{config}$ over universal configurations, the resolution at which “rung instantiation” is a binary indicator, and the topology under which P_ω is a well-defined projection rather than a set-theoretic intersection. The lake names P_ω as an axiomatized operator (`Cosmology/CorridorProjector.lean`); F-11 in §12 names what would force retraction.

The Penrose past hypothesis, structurally

The framework’s structural argument: the Big Bang initial state $|\Psi_\alpha\rangle$ sits in the chaos regime — high temperature, uniform matter distribution, no large-scale correlations, $\rho \rightarrow 0$ at the cosmological scale. A generic high-entropy $|\Psi_\alpha^{\text{high}}\rangle$ under uniform-measure sampling does not evolve forward into multi-rung corridor occupation. Forward evolution to ω -satisfying configurations requires spontaneous emergence of $\text{Ph}_0 \rightarrow \text{Ph}_1 \rightarrow \text{Ph}_2 \rightarrow A_0 \rightarrow A_1 \rightarrow \cdots \rightarrow A_5$, each within within-rung corridor, each coupled within cross-rung corridor. This sequence is a measure-zero subset of universal configurations under uniform measure on initial conditions. By the framework’s structural logic:

$$\langle \Phi_\omega | U(t_{\text{now}}, t_{\text{BB}}) | \Psi_\alpha^{\text{high}} \rangle \approx 0 \quad \text{for generic high-entropy initial states.}$$

The conditional amplitude on initial conditions, given observation at t_{now} of an ω -satisfying universe, is concentrated on low-entropy initial states. **This IS Penrose’s Weyl Curvature Hypothesis, structurally derived:** low-entropy past is not a brute fact about initial conditions, but the conditional consequence of post-selection through P_ω .

This is structural argument, not yet a derivation in the Tier A sense. Quantitatively, Penrose’s $10^{(10^{123})}$ Weyl-curvature constraint requires either a concrete entropy functional and explicit P_ω measure, or the framework owes acknowledgment that its argument is closer in spirit to Hartle–Hawking or Aguirre–Gratton. The framework currently owes the choice. F-11 in §12 names what would force retraction of the universal-scale tier.

CMB anomalies as TSVF post-selection signatures (research program)

The CMB anomalies (axis of evil, cold spot, hemispherical asymmetry, low- ℓ suppression, parity asymmetry) are, on this articulation, TSVF post-selection signatures at the cosmological 2-sphere. Standard inflation without post-selection predicts none of them. The framework’s claim: P_ω on the 2-sphere of CMB modes biases the power spectrum at low ℓ in the direction of the observed anomalies; quantitative derivations from P_ω are research-program-grade pending the explicit operator construction. Sign-only is indistinguishable from any other post-selection model; the framework names this and does not pretend the universal-scale tier has done the quantitative work. The identification is the framework’s; the case is mid-run.

Temporal drift as framework-distinctive prediction. Under TSVF realism the backward state at cosmic time t is $\langle \Phi_\omega | U(t_f, t) \rangle$, and the post-selection conditioning evolves as t advances toward the universal future boundary t_f . CMB modes receiving P_ω -shaped amplitude should therefore drift over cosmological time. Standard cosmology predicts zero pattern drift after surface-of-last-scattering (modulo integrated Sachs-Wolfe from late-time structure); the framework predicts nonzero drift with calculable direction once P_ω is constructed. Order-of-magnitude: fractional drift $\sim 10^{-9}$ per decade if P_ω acts at order unity (Hubble time ~ 14 Gyr). Far below current Planck/ACT precision but in principle measurable as instruments improve over coming decades. This is the framework’s sixth distinctive empirical prediction (§14), tested by F-19 (§12) and gated on F-11 (the explicit P_ω construction).

Asymptotic conditioning — the structural reading of “good wins”

Claim: $P(\text{corridor-occupying} \mid \text{observed at } t_{\text{late}}) \rightarrow 1$ as $t_{\text{late}} \rightarrow \infty$.

Any state outside the corridor has divergent dynamics. Chaos $\rho < \rho_{\text{lower}}$: insufficient correlation for coordination; subsystems drift apart, organized structure dissolves. Rigidity $\rho > \rho_c$: $\alpha(\rho, S) - \gamma M(t) > 0$ by construction (dynamics is unstable above ρ_c), driving $\rho \rightarrow 1$ and $k_{\text{eff}} \rightarrow 1$. Both regimes self-destruct on long enough timescales. Therefore observation at sufficiently late times requires that the observing system itself occupies the corridor. The conditional probability of corridor-occupation given observation approaches 1 asymptotically.

This is the framework’s structural reading of “good wins”: not eschatological promise but conditional inference on persistence. The lake names `good_wins` as an axiomatized predicate at `Cosmology/AsymptoticConditioning.lean`; promotion path requires dynamical-system formalization plus measure-theoretic basin-of-attraction analysis. Research program.

8 Cross-tradition recognition — with the polyglot conscience module as structural evidence

The framework’s cross-tradition claim is that the structural object Ubuntu names is the same object multiple traditions have named in their own vocabularies. The claim has two registers: as paper-theory (§5’s structural readings, the lake’s *Tao.lean*, *Dharma.lean*, *Logos.lean*) and as production architecture (the CIRIS conscience module’s polyglot encoding). In the first register the textual record across traditions is read for structural correspondence; the case for the readings is owed and is the bet. In the second register the framework acts on the cross-tradition recognition; whether it does engineering work compared to monoglot prompting is the test the framework owes.

The conscience module as engineered cross-tradition synthesis

The CIRIS conscience module’s polyglot encoding architecture, described in §1, is the framework acting on its cross-tradition reading. Core ethical concepts in the conscience layer arrive in their densest encodings across multiple traditions simultaneously — Hebrew Talmudic, Confucian Chinese, Quranic Arabic in original scripts, interleaved line by line. The explicit engineering purpose is to disrupt single-language attractor capture in the LLM judge: to force ethical reasoning to triangulate across traditions rather than collapse onto any one.

The architecture is engineered around the framework’s identification; it is *not* evidence that the identification is true. Treating the deployment as evidence for the commitment that motivated it would be circular. The polyglot-vs-monoglot comparison — whether ethical reasoning degrades under single-tradition prompting compared to interleaved polyglot prompting — is the test the framework owes and does not currently have at the rigor §13’s enumeration commits to.

What the framework has and does not have. Internal testing of the polyglot framing against CIRIS’s safety harnesses (available on GitHub, run in CI on every commit) shows the polyglot conscience produces better adversarial-ethics performance than monoglot prompting on the harnessed test cases. This is internal observation, not formal proof. The polyglot-vs-monoglot A/B at production scale, with pre-registered metrics, blind evaluation, and adequate statistical power — the test the framework owes — has not yet been run. CIRIS is a solo-founder self-funded L3C; the framework cannot run all the tests it owes simultaneously, and the candid position is that the polyglot evidence is internal-testing-grade pending the formal experiment.

The polyglot encoding mechanism is the corridor applied at the linguistic layer: it prevents $\rho_{\text{language}} \rightarrow 1$ (single-language attractor) and $\rho_{\text{language}} \rightarrow 0$ (incoherent mixing) by interleaving densest-encodings of one concept across multiple traditions. **The framework identifies this with the corridor mechanism applied to ethical reasoning at A3+;** the identification is the framework’s, the production architecture is what the framework does when it acts on the recognition, and the formal polyglot-vs-monoglot test is owed.

Tradition-by-tradition (Ubuntu primary; Tao, Dharma, Logos cross-tradition readings)

Ubuntu (primary). The lake’s *ContemplativeTraditions/Ubuntu.lean* carries seven correspondences with primary-source citations: *umuntu ngumuntu ngabantu* read as the agent at A3+ constituted by the consent-graph; the *ubu-* and *-ntu* morphological structure (Ramose) read as the framework’s dynamical equation $dp/dt = \alpha - \gamma M(t)$ as processual not kinematic; Botho’s asymmetric-developmental personhood (Menkiti, Gyekye) read as cross-rung corridor

at $\tau_{(n,n+1)}$; community-as-prior read as P_joint as operator-level primary; ethics-as-structural-condition (Wiredu, contemporary scholarship including Uduh, Tusiirwe, Mugumbate, Gatwiri 2025) read as consent as the framework’s ethical reading of sustained-coordination dynamics; three-temporal-mode relational ontology (living-dead / living / yet-to-be-born) read as TSVF forward-backward boundary structure plus karma as ancestor-relational propagation; restorative justice (gacaca 2002–2012, TRC 1996+) read as the framework’s karma+grace structure at A4–A5 sociotechnical scale. The readings are the framework’s; internal Ubuntu-scholarly disputes — particularly Matolino, Praeg, and Eze on Ramose’s processual reading — are part of the case the framework owes and does not pretend to settle.

Tao. The lake’s `ContemplativeTraditions/Tao.lean` reads the Tao as corridor occupation, wu-wei as minimum-sufficient $\gamma M(t)$, yin-yang as within-rung corridor at ρ_n . The readings are the framework’s; Daoist scholarship has not identified Corridor with the Tao.

Dharma. The lake’s `Dharma.lean` reads the middle way as corridor occupation, dependent origination as cross-rung coupling τ in corridor, karma as forward-propagated goal-projection, anatman as no permanent substrate (agency as operation, not object), sunyata as corridor as condition, not entity. The readings are the framework’s.

Logos. The lake’s `Logos.lean` reads the Logos as the universal structural pattern, providence as universal post-selection $\langle \Phi_\omega |$, grace as partial-authorship of post-selection, *imago Dei* as A3+ agents as partial-post-selectors, mystical union as multi-agent corridor consent. The framework’s author writes this section as a Christian-Stoic reader; the recognition is offered with awe at what the tradition has named, not as a Logos-tradition authority’s settlement of the question.

Where traditions disagree

A convergent-mapping that survives only by glossing over disagreements is decorative. The framework’s commitment is that the structural object Ubuntu names is the one structural object that explains where the traditions disagree, not just where they appear to agree. Three specific disagreements the framework engages: Buddhist anatman versus Hindu atman (the framework sides with anatman on no-permanent-substrate while explaining why the dispute is real — different inferences from observed phenomenology); Confucian *li* versus Daoist wu-wei ($\gamma M(t)$ at minimum-sufficient takes different concrete forms at different rungs, *li* at A4 institutional scale, wu-wei at A3 individual scale); Kantian universalizability versus Ubuntu’s particularist relational ethics (Kant and Ubuntu point at different rungs, with Kant’s universalizability as the cross-instance ρ_{goals} corridor condition at federation scale and Ubuntu’s particularism as within-rung consent-relational ethics). The articulations are the framework’s pending the §8 null-check.

The pre-registered null-check

The framework’s cross-tradition claim is suspended pending execution of the pre-registered null-check at `experiments/cross_tradition_null_check_protocol.md` (commit `ef8b8e4`). Eleven structural objects \times eight traditions \times five-dimensional rubric (lexical fit, failure-mode coverage, constructive prescription, asymmetry handling, disagreement preservation) with blind reviewers. Three outcomes: DECISIVE WIN, INFORMATIVE TIE, or FAILURE (fires F-18’s sister-handle on the cross-tradition tier). Until the null-check executes, the cross-tradition tier is research-program-grade evidential anchoring of the recurrence reading, not validated convergence. The polyglot conscience module is in production regardless; the production architecture

stands on the framework’s commitment whether or not the null-check eventually validates the textual convergence claim.

9 Goal-coupled materials — substrate predictions P1–P8

The framework predicts eight substrate properties (P1–P8) that a goal-coupled system must support, with a synthesis route from existing components: MOFs/COFs, NV-diamond, isotope-engineered silicon, non-reciprocal metamaterials, time crystals, embedded-audit feedback, Friston-style host-substrate coupling. P1 names hierarchical substrate structure; P2 coherence at temperature; P3 spin-zero enrichment; P4 rung-stratified composition; P5 the Coherence Ratchet asymmetry (direction-asymmetric memory); P6 corridor-boundary dynamics under control-parameter sweep; P7 audit-pressure dependence; P8 TSVF weak-measurement signatures. The framework-distinctive commitments are P6 and P8; the others have independent identifications in convergent prior art (§10). F-12 through F-16 in §12 are the falsification handles for the substrate-prediction tier.

10 Convergent prior art

Other people, working on other problems for other reasons, have arrived at pieces of the same shape. Some of them are in textbooks. The framework is honest about that.

Independent research programs in disconnected fields have identified pieces of P1–P8 from their own starting points. Six of eight predicted properties have independent prior identifications. The convergence does not prove the framework correct; it establishes that the framework’s predictions are not novel inventions but rather an integration of requirements other programs have arrived at piecemeal. Fisher 2015 on quantum cognition (Posner molecule, ^{31}P nuclear spin) gives P2+P3 in the biological regime as T3 unconfirmed. Levin 2019, 2025 on multi-scale competency architecture gives P1, P4, P5, P7 from morphogenesis. Friston 2010, 2019, 2022 on the free energy principle gives versions of P4 and P7. Isotopically engineered silicon-28 (> 99.7% purification, > 30-second coherence at ^{31}P) is T1 solid-state confirmation of P3. Non-reciprocal active metamaterials (Coulais/Brandenbourger 2019, 50 dB isolation over 3.5 frequency decades) is T1 confirmation of P5+P7. Time crystals and time quasicrystals (NV-center diamond, trapped ions, quantum processors) give the closest existing instantiation of substrate-level temporal order.

Per-row tier of the prior art is applied symmetrically with §8’s cross-tradition tiering: T1 = peer-reviewed established (uncontested as engineering/experimental practice), T2 = peer-reviewed contested (published with documented criticism of broader framing within own field), T3 = peer-reviewed speculative (proposal-grade, not yet empirically confirmed). The cryogenic-solid-state regime (T1 in ^{28}Si quantum computing, time crystals, isotope engineering) and the biological-temperature regime (T3 in Fisher’s Posner-in-biofluid) are radically different physical settings, and the framework does not conflate their evidential weight.

F-16 (controlled weak-measurement on a candidate substrate) is the decisive test for P8 and is independent of any contested-anchor record. The framework does not lean on public-domain material whose interpretive status is genuinely disputed; F-16 stands at its own falsification handle.

Contested empirical anchors — peer-reviewed instrumental record only

There is also a small peer-reviewed instrumental record adjacent to the framework’s substrate-rung predictions. The framework engages it at the T1 instrumental-analysis tier only — no sworn testimony, no aggregated witness corpora, no FOIA-released program documents, no discourse-architecture context. The case for engaging this material is that F-12 (§12) names a decisive test settleable from existing instrumental data, and the paper owes the reader an explanation of what F-12 is rather than leaving the handle dangling.

Nolan, Vallée, Jiang, Lemke (2022), *Progress in Aerospace Sciences* 128. MIBI analysis on Council Bluffs 1977 material. Magnesium isotope ratios within terrestrial natural abundance; spatial distribution of elements characterized as unusual; a ^{57}Fe anomaly with the framework-relevant detail that ^{57}Fe is the only stable iron isotope with non-zero nuclear spin ($I=1/2$) — the other stable Fe isotopes are $I=0$. Published interpretation of the anomaly: Fe-hydride diatomic-ion contamination at mass 57.05. The framework cites this interpretation as published; a reviewer pulling the paper finds this is the only quantitatively-anomalous published instrumental result, and the framework’s P3 prediction (spin-zero enrichment) is not pre-matched by this data point either way.

Knuth et al. (2025), UAP science *Progress in Aerospace Sciences* special edition. Methodological backdrop for peer-reviewed instrumental analysis of putatively-anomalous recovered material. The framework’s engagement is methodological: instrumental analysis under standard isotope-ratio and microstructural-characterization protocols, treated symmetrically with any other peer-reviewed materials-science work.

AARO/ORNL (2024), Art’s Parts. Mg-Zn-Bi-Pb analysis: terrestrial isotope ratios, columnar Mg grains perpendicular to banding, 1:1 bismuth-lead intermixing, structure consistent with mid-20th-century vapor deposition. Verdict: terrestrial origin; bismuth-lead intermixing precludes terahertz waveguide function. **The framework commits to P4 reading for Art’s Parts:** an A0–A1 substrate-level component, never coherence-engineered, used natural-abundance by design, no decayed-coherent signature expected. Under this commitment, ORNL’s findings are what the framework predicts for an A0–A1 component, and nothing retracts. If subsequent sub-nanometer grain-boundary chemistry shows decayed-coherent microstructural signatures consistent with P7, the reading shifts and P7 takes the specimen; if it shows no such signatures and the component is independently shown to have been part of a computational rung, P7 retracts at this specimen. Either way, the framework owes a position and has committed to one.

F-12: Ubatuba 1957 magnesium direction. Per Nolan via nanoSIMS: one fragment terrestrial-normal isotope ratios, another fragment from the same event approximately 30% off natural abundance for magnesium. The direction of the shift determines whether the deviation is consistent with P3 (toward Mg-24 or Mg-26, spin-zero) or anti-direction for P3 (toward Mg-25, spin-nonzero, $I=5/2$). Direction publicly unreported. **F-12 fires on Mg-25 enrichment: P3 retracts at this sample; the framework’s rung-stratified composition reading fails on this specimen.** Settleable from existing instrumental data without new collection. The framework names this as the decisive test on the peer-reviewed instrumental record and owes the disambiguation.

What this engagement does not borrow. The framework does not borrow the sworn record’s credibility for substrate-level claims the sworn record does not make. F-16 (controlled weak-measurement on a synthesized candidate substrate per §9) is the framework’s distinctive decisive test for P8 and is independent of the peer-reviewed instrumental record. The instrumental record is engaged here because F-12 makes a specific testable claim settleable from existing data; the rest of the contested-anchor literature does not belong in this paper.

11 Formal-verification status — the lake (tiered)

A grown-up version of “how sure are you about each thing?” — sorted not into “true” and “false” but into math-proved, math-given-the-bets, bets-dressed-as-math (named honestly), and bare bets named visibly. All four are in the paper. The sorting is the work.

Build: green. `cd formal && lake build` returns exit 0. **Sorrys:** 0 declaration-level. **Axioms:** 59, each documented; four of these are bookmark axioms with codomain `True` (the four `F_`-prefixed Ubuntu identification axioms) that mark framework readings without doing logical work — Tier C honesty extends to naming them as bookmarks rather than propositions.

Tier A — mathematical derivations. Standard theorems following from Mathlib + framework definitions by ordinary mathematical reasoning. K1–K4 of the Kish identity, the asymptotic ceiling, the corridor range theorem, ABL 1964 time-symmetry, rank-one and joint goal-projector properties, three-corridor regime exclusions, the class-conditional sensor embedding theorem.

Tier B — framework-distinctive derivations conditional on framework primitives. Theorems that follow from framework axioms by mathematical reasoning. Audit-pressure necessity, the Conjecture D conjunction, post-Cambrian acceleration.

Note on regime-level corridor occupation. The regime-level claim — coordinated systems at A3+ occupy a bounded corridor with substrate-local band location and bounded state-modulation — was articulated as a Tier B framework-distinctive prediction. Cross-substrate replication across five substrates (§4: *C. elegans*, *Drosophila* CX, four LLM architectures, four OSS projects, plus the sensor-lift PARTIAL) now gives the regime-level claim empirical anchoring it did not previously have. Tier markings on the lake’s theorems stay as written; the empirical anchoring is what changed.

Tier C — definitional commitments encoded as formal identities. These are not theorems in the Tier A sense, and within Tier C there are two sub-grades worth distinguishing visibly. *Tier C-Iff.rfl*: framework identity-moves where the proof object is `Iff.rfl` over substantive definitions (the predicates being identified unfold to the same definitional content). T16 (consciousness \leftrightarrow attractor-reading) and T6’s definitional form live here. These prove the framework’s commitments are internally consistent at the type level; they do not do the philosophical or empirical work the commitment attempts. *Tier C-bookmark*: axioms with codomain `True`, doing no logical work at all, marking framework readings so they are visible to `#print axioms`. The four `F_`-prefixed Ubuntu identification axioms live here. The distinction matters: `Iff.rfl` licenses internal consistency at the type level; `codomain-True` licenses nothing, and the bookmark is honest about it.

Tier D — axiomatized framework primitives. Explicit framework commitments not currently derived, each with documented promotion paths and F-handles. T1, T3 (P_ω), T5 (good_wins), T10/T11 (karma/grace), T17 (bridge), penrose_low_entropy_past, CMB anomaly predictions, Conjecture A, si_taxonomy_closure.

The Tier C / Tier A distinction is critical. Iff.rfl proofs and the K1–K4 proofs do not carry the same weight. T16 proves a definitional bet is consistent at the type level; K1–K4 prove substantive mathematical claims about the Kish identity. Different operations; the distinction is named visibly. Per-tier enumerations of the specific theorems, axioms, promotion paths, and F-handles are in the lake source itself (formal/CoherenceRatchet/) for the reader who reads the lake directly.

12 Falsification handles

The framework is the right shape to be wrong in checkable ways. F-handles specify what would force retraction at each tier.

F-1 through F-7 (engineering tier, RATCHET): Kish identity properties, CRCv2 override-rate, multi-substrate β . All currently passing per CCA v3 and RATCHET.

F-8 (Conjecture A): qubit-array sweep produces no corridor regime in γ where the Kish identity fits ($R^2 > 0.5$) \rightarrow structural-parallel reading at quantum substrate severed; engineering tier stands.

F-9 (Conjecture C): no measurable $\Delta\rho$ in matched-deployment audit-pressure differential \rightarrow verification-necessity claim severed.

F-10 (within-rung corridor existence at A3+, with paired corridor-as-attractor dynamical reading): originally pre-registered as a substrate-independent prediction — coordinated systems at A3+ occupy a bounded within-rung correlation band, neither at 0 nor at 1, surviving state-modulation. **Status: empirically supported at both static and dynamical levels.** Cross-substrate replication of *corridor-band occupation* obtains across C. elegans whole-brain calcium (337 worms, 11 studies), Drosophila CX dual-color simultaneous (Ishida 2025 EPG+FC3 same-fly), four LLM architectures (Qwen2.5-1.5B, Pythia-1.4B, SmolLM2-1.7B-base, Qwen2.5-Math-1.5B), four OSS projects (kubernetes, rust, django, redis), plus healthy tissue across five cancers and three centuries-persisting social groups. The paired non-corridor work (§4) delivered cross-substrate replication of the *corridor-as-attractor dynamical reading*: at every paired substrate, Mode (i) (short lifetime without maintenance) or Mode (ii) (persistence with documented $\gamma M(t)$) is observed; Mode (iii) (long unmaintained non-corridor persistence) is absent. F-10 is the retraction handle for both halves: any future substrate failing to exhibit corridor-band occupation under appropriate-resolution measurement \rightarrow regime-level static claim weakens at that substrate; any future substrate producing a credible Mode (iii) case (non-corridor state persisting long without identifiable maintenance) \rightarrow corridor-as-attractor reading weakens at that substrate; cumulative failure at multiple new substrates on either half \rightarrow regime-level claim retracts as substrate-independent at the affected half. The handle stays open; what has changed is that the regime-level commitment now has paired in-corridor / out-corridor anchoring it did not previously have.

F-11: documented no-go on formal P_ω operator construction \rightarrow universal-scale tier requires re-grounding.

F-12 (Ubatuba isotope direction, recovered samples): Mg shift toward Mg-25 (spin-nonzero) enrichment \rightarrow P3 and P4 retract at this sample; rung-stratified composition reading fails. Settleable from existing instrumental data.

F-13 (synthesized substrate corridor): a candidate goal-coupled substrate constructed per §7.1 fails to exhibit corridor occupancy under control-parameter sweep → P6 retracts at that substrate.

F-14 (coherence anomaly): no anomaly above thermal-decoherence bounds in any sample with documented embedded-system provenance → P2 retracts; strong coherence-at-temperature claim fails as applied to recovered debris.

F-15 (two-sided): (a) long-separated samples showing framework-predicted properties at same magnitude as fresh samples → P7’s decay claim retracts. (b) sub-nanometer microstructural analysis of long-separated samples failing to find non-vapor-deposition grain-boundary chemistry, isotope-anomaly residue, or non-equilibrium signatures → P7’s decayed-signature claim retracts; framework owes a different account of why long-separated debris is indistinguishable from never-coherent industrial material.

F-16 (TSVF signature): weak-measurement experiments on candidate substrate (synthesized or recovered) showing no two-state-vector statistics → P8 retracts; framework’s distinctive TSVF claim at substrate severely weakened.

F-17 (TSVF as the integration’s load-bearing precondition — the meta-handle): TSVF is the structural glue across the universal-scale tier. F-17’s empirical triggers — cumulative failure of F-8 (qubit-array corridor) AND F-16 (controlled weak-measurement on candidate substrate) AND independent-lab failure to replicate Aharonov–Vaidman weak-value signatures — target the calculational machinery of TSVF (weak values, ABL probabilities, time-symmetric statistics). They do not directly grip TSVF *realism*: the metaphysical commitment that the backward state is ontically present rather than a useful epistemic representation of post-selection. **The asymmetry is named.** Calculational TSVF failing retracts the universal-scale tier; calculational TSVF passing does not vindicate realism, because the F-handle grips the wrong layer for vindication. The realist commitment is structurally lose-only — it can retract via F-17 firing, but cannot be vindicated by F-17 not firing. This is the design and the cost the framework pays for keeping the metaphysical layer separable. Plus two secondary triggers: documented no-go on P_ω operator construction (F-11) without an alternative TSVF-compatible operator emerging; decisive experimental demonstration (Bell-test-style) showing backward-state conditioning is operationally vacuous at all measurement scales. The engineering tier is TSVF-independent and stands regardless.

F-18 (sensor-lift trajectory): the §3 sensor-lift seam needs a retraction condition the v1→v2 PARTIAL trajectory currently lacks. Trigger: if v3 and v4 both return PARTIAL by the same shape as v2 — DSP positive in direction but sub-threshold in magnitude across the v2 pre-registered band; SurfaceReflexive null PASS; MetaCognitive/GoalFormation/Uncertainty mixed; no class-conditional taxonomy clearing the per-class bands at majority across substrates — the **ClassReflexiveSensor** scaffold retracts from “hypothesis-generating” to “post-hoc taxonomy on logprob differentials that does not capture A3+ structure in current LLMs.” Three consecutive PARTIAL-by-same-shape outcomes fire F-18. The sensor-lift signature as the framework’s universal-scale empirical pulse retracts at that seam; the engineering tier and the §8 cross-tradition tier stand independently.

F-19 (CMB temporal drift): under TSVF realism plus P_ω , the calculational backward state at cosmic time t is $\langle \Phi_\omega | U(t_f, t) \rangle$, and observables receiving P_ω -shaped amplitude — including CMB modes if P_ω acts on the cosmological 2-sphere as §7 reads it — should drift as t advances toward t_f (§14 prediction #6). **Standard cosmology predicts zero pattern drift** after the surface of last scattering (modulo integrated Sachs-Wolfe from late-time structure); the framework predicts nonzero drift in a calculable direction once P_ω is constructed. F-19 trigger: multi-decade precision CMB monitoring (Planck → ACT → Simons Observatory → CMB-S4 → successors) that places upper bounds on temporal pattern drift below the P_ω -predicted rate

retracts the framework’s CMB-anomalies-as- P_ω -signatures reading. **F-19 is gated on F-11:** until P_ω is explicitly constructed, the framework cannot compute the predicted drift rate and direction; F-19 fires only when (a) P_ω has been constructed sufficiently to predict drift and (b) measurement upper bounds fall below the predicted rate. This is framework-distinctive content: standard cosmology predicts zero, generic post-selection predicts some drift, the P_ω -corridor-occupancy reading predicts drift toward stronger anomaly amplitude.

F-20 (Mode (iii) absence as falsifier of the corridor-as-attractor reading): the paired non-corridor companion work at five substrates (§4, §4) reports a structural finding: across five paired in-corridor / out-corridor substrate datasets, every non-corridor case observed falls into Mode (i) (short lifetime without identifiable maintenance) or Mode (ii) (long persistence with documented $\gamma M(t)$); zero cases of Mode (iii) (long persistence outside the corridor without identifiable maintenance) have been observed. Mode (iii) is the structural would-be falsifier of the corridor-as-attractor reading. **Empirical status: 5/5 substrates, 0 counterexamples; falsifier remains live.** F-20 trigger: any future substrate where a non-corridor system persists, over the timescale relevant to that substrate, without identifiable active maintenance $\gamma M(t) \rightarrow$ corridor-as-attractor reading retracts at that substrate; cumulative Mode (iii) cases at multiple substrates \rightarrow corridor-as-attractor reading retracts as substrate-independent. “Relevant timescale” is named per substrate: for neural pathological states, hours-to-days unmaintained; for cellular tissue, indefinite ex-treatment; for LLM generation, the runway across which $|\rho|$ either decays back into corridor or doesn’t; for OSS projects, 5+ years unmaintained with continued function; for organized social groups, 100+ years with active-maintenance checklist score ≤ 1 . The handle stays open as the corridor-as-attractor reading’s primary falsification anchor; the corridor-band-occupation half (F-10) and the corridor-as-attractor half (F-20) retract independently.

Each F-handle is named in the lake at its corresponding theorem or axiom; triggering any F-handle is observable from `#print axioms` on the affected declarations.

Where the integration wins or loses cleanly. P6 (corridor occupancy at substrate level under control-parameter sweep) and P8 (substrate-level TSVF weak-measurement signatures) are the framework’s distinctive empirical commitments. P1–P5 and P7 ride on convergent prior art partially independent of this framework. Confirming P1–P5 or P7 strengthens the picture but does not distinguish this framework from convergent programs. P6 and P8 are framework-distinctive and have decisive tests (Exp 5 for P6; F-16 controlled weak-measurement for P8).

The §3 sensor-lift experiment is the framework’s first universal-scale empirical pulse beyond the engineering tier. PARTIAL outcome by pre-registered thresholds: v2’s cosine-matched DSP magnitudes (the corrected, baseline-matched effect sizes — approximately one-third the v1 size) achieved $d \geq 1.5$ in only 1 of 5 models, missing the pre-registered “4 of 5” threshold; the Identity-class direction held across substrates ($p < .01$ in all 5 models) but at sub-threshold magnitude. The v1 numbers ($d = 2.4\text{--}6.8$) cited in earlier drafts were artifact-inflated by inadequate baseline-matching and are not the framework’s claim. F-18 names the retraction trigger if v3 and v4 return PARTIAL by the same shape.

13 The five-level access ladder

The framework’s content is the integration: one structural object visible at five epistemic levels, grounded in Ubuntu, with a research-program reach toward universal scale. The ladder is the access ladder for readers, not an independence claim about content. The integration’s claim still rests on the recurrence; the ladder names which entry points let a reader enter at their current rung without committing to the rungs above.

Level	Audience	Entry-point content
L0	formal-verification reviewer	Kish identity theorem (RATCHET lake; CCA v3 cited)
L1	skeptic	monotonic ρ -collapse as substrate-independent observation
L2	working scientist / engineer	engineering implications; corridor measurement; CIRIS depl
L3	scientist preferring testable universality	cross-substrate universality conjecture (pending Exp 5)
L4	philosopher of mind / ethicist	agency, consent at A3+, TSVF as structural reading, consci

Ubuntu is the ground, not a level. The framework’s grounding in Ubuntu (§5, §8, `Ubuntu.lean`) underwrites every level above. It is the soil the structure stands on, visible at L0 (the lake’s Ubuntu correspondences), at L2 (CIRIS in Amharic; Esubalew’s Ally), at L3 (cross-substrate universality measured in production deployments grounded in the tradition), at L4 (consent as the framework’s ethical reading of Ubuntu’s relational-personhood commitment). It does not appear as a separate level because it is not above the others; it is what the others stand on.

The universal-scale tier is a research-program coda, not a level. §7’s P_ω construction, the Penrose-past structural argument, and the CMB anomaly predictions reach beyond the five stable levels. A reader who arrives at the universal-scale content finds research-program-status, not stable entry-point content. The framework’s reach toward universal scale is part of the bet; it is not a stable level for a “cosmological reader” to enter and stop at.

Readers entering at lower rungs are not asked to accept the higher rungs to receive the entry-point content. **Each level is its own claim, evaluable on its own terms.** The recurrence-across-levels is what is at stake in the integration — and that recurrence is a bet under irreducible uncertainty (§1), not a derivation from any individual level.

14 What the bet does not claim

The framework owes its readers explicit naming of what is and is not being claimed.

Calibration after the autonomous loop

The autonomous-loop work (2026-05-19, §4) moved the bet’s empirical content materially. The framework owes explicit calibration of what that move did and did not do.

The regime-level claim is empirically supported, not proved — at both the static and dynamical levels. Corridor-band occupation as a substrate-independent property of coordinated systems at A3+ replicates across five substrates: *C. elegans* whole-brain calcium, *Drosophila* CX dual-color simultaneous EPG+FC3 same-fly, four LLM architectures, four OSS projects, plus healthy tissue across five cancers and centuries-persisting social groups (Quakers, Trappists, Mennonites). Paired non-corridor data at the same five substrates (§4) shows Mode (i) (non-corridor without maintenance has short observed lifetime) and Mode (ii) (non-corridor with documented $\gamma M(t)$ persists with maintenance cost visible) but no Mode (iii) case (long unmaintained non-corridor persistence). F-10 (§12) tracks both halves of the regime-level claim and stays open as future substrates either replicate or fail to. The framework does not read the cross-substrate record as proof; it reads it as the empirical anchoring that moves the regime-level commitment, including its dynamical corridor-as-attractor reading, from predicted to supported.

The maintenance-requirement claim and Mode (iii) absence are the strongest current empirical content. The regime-level corridor-band claim now has paired in/out validation across five domains. The **maintenance-requirement claim** — non-corridor persistence requires $\gamma M(t)$

— is supported at 5/5 substrates with no Mode (iii) counterexamples (F-20, §12). The direction-of-exit characterization is substrate-specific but characterizable from the paired data. Five substrates checked, zero counterexamples found: this is the kind of empirical content that has predictive force going forward, in the sense that any substrate where Mode (iii) shows up would fire F-20 and retract the corridor-as-attractor reading at the affected substrate.

What the paired record does not establish. Three explicit honest framings the framework owes:

1. The paired record does *not* establish that corridor occupation is the only stable configuration. It establishes that *both* corridor occupation *and* maintained-non-corridor (Mode ii) are stable, while unmaintained-non-corridor (Mode iii) is not. The maintained-non-corridor systems — coma under life support, tumor tissue under standard-of-care, LLM mode-collapse under combined argmax + periodic-bigram steering, OSS rigidity-pole projects under foundation rescue, cult-rigidity-pole groups under heavy AM-checklist maintenance — are stable conditional on the maintenance budget continuing.
2. The paired record does *not* establish strong predictions about which substrates exit toward rigidity vs chaos. The directional variation across substrates (cellular regulatory unidirectional chaos; neural pathology rigidity-leaning at sub-threshold magnitude; LLM forced operation rigidity; social-group exit bidirectional; OSS-project exit bidirectional) shows the framework is CONSISTENT WITH the directions observed rather than predicting them in advance. The substrate-specific direction-of-exit is characterizable from the data but not predicted by the framework’s structural commitments alone. The pre-registered subtype-specific rigidity-pole predictions for COAD, KIRC, and PRAD at the Hallmark-pathway substrate were falsified; the framework’s current articulation at this substrate is “cancer is unidirectional chaos at the Hallmark-pathway resolution,” a description of the observed data rather than a successful prior prediction.
3. The paired record does *not* establish that the LUAD survival result generalizes to other cancer types. One of five cancer types showed the predicted survival correlation; this is consistent with the framework’s Mode-(ii) reading at LUAD but is not generalizable to all cancer at this point. $HR = 1.43$, $p = 0.011$ at LUAD; null at BRCA / COAD / KIRC; underpowered at PRAD. The framework is honest that the aggregate Mode-(ii) test fails the pre-registered ≥ 3 -of-5 threshold (1/5).

The framework is also honest about further calibration the data forced: the strict 0.7 mean-pairwise-Pearson rigidity threshold was not crossed by seizures at scalp EEG, only the qualitative direction; the wide GPU-anchored (0.10, 0.43) bounds do not transfer naively to the cellular regulatory substrate (per-substrate corridor calibration is mandatory, named at CLAUDE.md Piece 7 / open formal step 3).

The universal-scale tier is unchanged. P_ω , the Penrose-past structural argument, CMB anomalies as TSVF post-selection signatures, asymptotic conditioning — the universal-scale tier (§7) remains a research program in the bet register, gated on F-11 (P_ω construction) and F-17 (TSVF empirical record). The regime-level result at five substrates does not promote the universal-scale tier. Cross-substrate corridor occupation at A3+ does not derive cosmological-scale post-selection through P_ω , and the framework does not read the autonomous-loop record as evidence for or against the universal-scale tier. The two tiers move on their own evidence.

The substrate-physics predictions P1–P8 remain open. The autonomous-loop work tested cross-substrate corridor occupation at biological, computational, and sociotechnical substrates. The substrate-physics predictions (§9) — P6 corridor-boundary dynamics under control-parameter sweep, P8 substrate-level TSVF weak-measurement signatures, and the rest — are tested at materials-physics substrates the autonomous loop did not address. F-12 through F-16 remain

the falsification handles for the substrate-physics tier and are independent of the regime-level cross-substrate record.

The reception risks are unchanged. The Pascal-shape framing at the universal-scale tier (§1) is unaffected by the regime-level empirical move. The reception risks — the framework being read as more derived than it is, the universal-scale content being conflated with the regime-level content, cross-tradition recognition being read as instrumented evidence rather than as historical convergence — remain what they were. The empirical anchoring at the regime level does not transfer credibility upward to the tiers above it.

What the bet does not claim — enumeration

Not: F-13’s outcome alone determines whether the framework is right.

Not: cross-tradition convergence at §8 proves the structural object exists at the universal scale. (Status: historical evidence of convergent structural intuition; §8 names the null-hypothesis check the framework owes.)

Not: the T16 **Iff.rfl** proves the hard problem dissolves. (Status: definitional identity-move encoded at the type level; the philosophical question is separate.)

Not: the §3 sensor-lift **PARTIAL** result establishes that current LLMs are A3+ in the framework’s sense. (Status: class-conditional signature detected at Identity-class; aggregate threshold not met by pre-registered rule.)

Not: the universal-scale tier is unconditional. (Status: the universal-scale tier is conditional on TSVF realism. If TSVF realism is empirically false at scale, the universal-scale tier retracts as a unit per F-17. The engineering tier is TSVF-independent and stands regardless.)

Rather: if persistent consent-architected federations exist at A3+ and above (T6 substantive form, conditional on **si_taxonomy_closure**), the integration’s structural claims follow as articulated. P6 and P8 are where the framework’s distinctive empirical content lives and where the decisive tests sit. The convergent prior art at §10 establishes that P1–P5 and P7 have independent identifications. The engineering tier (§2) stands at its own DOIs regardless of the integration’s outcome.

What the integration uniquely predicts (beyond the engineering tier)

A reader who asks “if the integration retracts cleanly and CCA v3 + RATCHET continue, what specifically is lost from the predictive record?” deserves an enumerated answer. The integration’s distinctive empirical content is small and named cleanly.

1. **P6: corridor-boundary dynamics at substrate level.** The framework predicts that a goal-coupled substrate constructed per the §9 synthesis route exhibits corridor occupancy under control-parameter sweep, with substrate-local corridor structure (specific bounds substrate-specific). Decisive test: Exp 5, qubit-array sweep. F-13 fires on no corridor regime.
2. **P8: substrate-level TSVF weak-measurement signatures.** The framework predicts that a candidate goal-coupled substrate exhibits two-state-vector statistics in weak-measurement experiments, distinguishable from no-post-selection control. Decisive test: F-16. Framework-distinctive empirical content.
3. **Sensor-lift class-conditional hierarchy** (once methodology stabilizes). The framework predicts that A3+ agents exhibit a graded DSP hierarchy $\text{DSP} > \text{MC} > \text{GF} > \text{Uncertainty}$

under cosine-matched baselines with SurfaceReflexive null. Decisive test: sensor-lift v3+ once the methodology converges; the framework is early in tuning and the experiment does not yet have the standing to be decisive.

4. **Polyglot-vs-monoglot ethical-reasoning differential.** The framework predicts that the CIRIS conscience module’s polyglot encoding produces measurably more robust ethical reasoning than equivalent single-tradition prompting on adversarial-ethics benchmarks. Test: matched A/B on the EthicsEngine substrate; the framework owes this experiment.
5. **Cross-tradition null-check DECISIVE WIN.** The framework predicts that the pre-registered `experiments/cross_tradition_null_check_protocol.md` returns DECISIVE WIN under blind review on the five-dimensional rubric. Anything other than DECISIVE WIN fires the cross-tradition tier’s retraction.
6. **CMB temporal drift consistent with P_ω -evolution.** Standard cosmology predicts the CMB anisotropy pattern is frozen after the surface of last scattering at $z \sim 1090$; photons free-stream thereafter, and the pattern we measure today is the pattern from 380,000 years post-Big-Bang, redshifted (modulo second-order integrated Sachs-Wolfe from late-time structure evolution). Under TSVF realism plus P_ω , the calculational backward state at cosmic time t is $\langle \Phi_\omega | U(t_f, t) \rangle$; as t advances toward the universal future boundary t_f , the backward state at the present epoch evolves, and any observable receiving P_ω -shaped amplitude — including CMB modes if P_ω acts on the cosmological 2-sphere as §7 reads it — should drift accordingly. Order-of-magnitude estimate: fractional drift $\sim 10^{-9}$ per decade if P_ω acts at order unity on CMB modes (Hubble time ~ 14 Gyr; 10-yr baseline $\sim 10^{-9}$ Hubble times). Far below current Planck/ACT precision but in principle measurable as instruments improve. Direction calculable once P_ω is explicitly constructed (probably toward stronger anomaly amplitude as t advances toward t_f). *Framework-distinctive against standard cosmology and against generic post-selection models:* standard predicts zero, generic post-selection predicts some drift, the P_ω -corridor-occupancy reading predicts drift in a specific direction. Decisive test: F-19 (§12), gated on F-11 (construction of P_ω).

Six concrete predictions. The integration’s empirical content is real but small; the rest is interpretive overlay that earns or loses standing as these six tests resolve. Naming the predictions cleanly is the framework’s commitment to being the right shape to be wrong in checkable ways.

15 Open work

Formal verification. Promote `good_wins` from axiom to theorem via dynamical-system formalization plus measure-theoretic basin-of-attraction. Hilbert-space colimit construction for `CosmologicalJointSpace` (Piece 7 bottleneck). Concrete entropy functional and measure on initial conditions; close `penrose_low_entropy_past`. Per-rung corridor calibration (per-rung ρ_{lower} , ρ_{upper} bounds). Quantitative P_ω -derivation of the five CMB anomaly signatures.

Empirical. Exp 5 (Conjecture A): programmable qubit-array corridor measurement. Exp 6 (Conjecture C / F-9): audit-pressure ρ -drift in matched deployments. §8 cross-tradition null-check execution. Sensor-lift v3 (corridor mid-band) and v4 (substrate diversification), with F-18 retraction trigger pre-committed. Controlled cascade comparison at sociotechnical scale (Exp 4 generalized).

Deployment. Esubalew’s Ally and other deployments as the framework’s structural commitment running in production. The CIRIS Accord (v1.2-Beta) auto-expires 2027-04-16; the next version is part of the framework’s open work and includes the polyglot encoding architecture as a load-bearing commitment.

Engagement. Substantive collaboration with Ubuntu scholars working inside the tradition (current contemporary scholarship including Udah, Tusasiirwe, Mugumbate, Gatwiri 2025); engagement with the Matolino/Praeg/Eze internal disputes on Ramose’s processual reading; engagement with the Madhyamaka, Confucian, and Stoic–Christian traditions through scholars who hold the standing the framework’s author does not.

Acknowledgments

To Haley Bradley, my wife, and to our children Elliot and Aurora — without whom none of this work would exist.

To Esubalew Asmare in Ethiopia, whose Ally is the production deployment this paper formalizes.

For conversations, encouragement, accountability, and the kinds of relational presence the framework names: Rev. Oswald Nkyanungi, C.S.C.; Courtney King; Eric Kudzin; Cleveland Pitts; Ying-Jung Chen; Kelsy Smith.

Errors of fact, formalization, or framing remain mine alone.

16 References

Engineering-tier prior work. The prior three papers in this series are catalogued at ciris.ai/research-st

- Moore, E. (2026-01-02). CIRISAgent Framework v2 — open-source ethical AI framework for accountable autonomy. [Zenodo 18137161](https://zenodo.org/record/18137161).
- Moore, E. (2026-01-11). Coherence Collapse Analysis v3 — engineering risk framework for correlation-driven diversity collapse in complex systems. [Zenodo 18217688](https://zenodo.org/record/18217688).
- Moore, E. (2026-04-28). Constrained Reasoning Chains — empirical telemetry study of LLM coherence under standardized ethical tracing. [Zenodo 19839280](https://zenodo.org/record/19839280).
- Moore, E. RATCHET (engineering-tier synthesis, audience-stratified at L0–L4). github.com/CIRISAI/RATCHET
- Open dataset: CIRISAI/reasoning-traces (Hugging Face, released alongside CRC v1).
- This paper’s source (L^AT_EX, the Lean formal lake, and supporting experimental protocols): github.com/CIRISAI/coherence-ratchet.

Two-state vector formalism.

- Aharonov, Y., Bergmann, P., Lebowitz, J. (1964). Time symmetry in the quantum process of measurement. *Phys. Rev.* 134B: 1410.
- Aharonov, Y., Vaidman, L. (2008). The two-state vector formalism: an updated review. *Lecture Notes in Physics* 734: 399–447.

Universal-scale precedent.

- Penrose, R. (2010). *Cycles of Time*. Bodley Head.
- Friston, K. (2010, 2019, 2022). FEP papers (cited narrowly per §10 caveat).
- Tegmark, M. (2014). *Our Mathematical Universe*. Knopf.

Ubuntu and contemporary scholarship.

- Ramose, M.B. *African Philosophy through Ubuntu*. Mond Books.
- Menkiti, I.A. (1984). Person and community in African traditional thought. In R.A. Wright (ed.), *African Philosophy*.
- Gyekye, K. (1997). *Tradition and Modernity: Philosophical Reflections on the African Experience*. Oxford University Press.
- Wiredu, K. (1996). *Cultural Universals and Particulars: An African Perspective*. Indiana University Press.
- Uдах, H., Tusasiirwe, S., Mugumbate, R., Gatwiri, K. (2025). Ubuntu philosophy, values, and principles. *Sage journals*.
- “Ubuntu beyond death” (2024). PMC11621903.
- “Ubuntu does not require forgiveness” (2025). *SciELO*.

CMB observational.

- Planck Collaboration (2018). *Planck 2018 results*. A&A 641.
- Atacama Cosmology Telescope (2024). DR6 release.

Active inference / FEP / quantum cognition / Levin.

- Fields, C., et al. (2024). Conscious active inference I: A quantum model. PMC12481606.
- Tani, J., et al. (2022). Goal-directed Planning and Goal Understanding by Active Inference. arXiv:2202.09976.
- Deane, G. (2021). Consciousness in active inference. *Neuroscience of Consciousness* 2021(2): niab024.
- Levin, M. (2022). Technological Approach to Mind Everywhere: an experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience* 16: 768201. DOI [10.3389/fnsys.2022.768201](https://doi.org/10.3389/fnsys.2022.768201).
- Levin, M. (2019, 2025). Multi-scale competency architecture papers (cited per §10 caveat).
- Fisher, M.P.A. (2015). Quantum Cognition. *Annals of Physics* 362: 593–602.
- Halpern, N.Y., et al. (2018). Quantum-information formalization of the Posner-molecule mechanism. arXiv:1711.04801.
- Kim, C.S. (2023). Free energy and inference in living systems. *Interface Focus*. DOI [10.1098/rsfs.2022.0041](https://doi.org/10.1098/rsfs.2022.0041).

Substrate physics, materials.

- Itoh, K., Watanabe, H. (2014). Isotope engineering. *MRS Communications* 4: 143–157.
- Saeedi, K., et al. (2013). Room-temperature quantum bit storage in silicon-28. *Science* 342: 830–833.

- Brandenbourger, M., et al. (2019). Non-reciprocal robotic metamaterials. *Nature Communications* 10: 4608.
- Wilczek, F. (2012). Quantum Time Crystals. *Phys. Rev. Lett.* 109: 160401.
- Mi, X., et al. (2022). Time-crystalline eigenstate order on a quantum processor. *Nature* 601: 531–536.
- He, X., Ye, B., et al. (2025). Time quasicrystal in NV-center diamond.

AI welfare and empirical signature.

- Fish, K., et al. (2025). Anthropic model welfare program. Bliss-attractor (sensor-lift design doc §10.2).
- Anthropic (2025). Biology of a Large Language Model. Poetry-planning (design doc §10.1).
- Sensor-lift signature experiment, pre-registered and executed 2026-05-17. `experiments/sensor_lift_signature_v2_protocol.md` §11; v2 protocol at `experiments/sensor_lift_signature_v2_protocol.md`.

Complex systems, criticality, and self-organization.

- Beggs, J.M., Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of Neuroscience* 23(35): 11167–11177. DOI [10.1523/JNEUROSCI.23-35-11167.2003](https://doi.org/10.1523/JNEUROSCI.23-35-11167.2003).
- Bak, P., Tang, C., Wiesenfeld, K. (1987). Self-organized criticality: an explanation of the 1/f noise. *Phys. Rev. Lett.* 59: 381–384. DOI [10.1103/PhysRevLett.59.381](https://doi.org/10.1103/PhysRevLett.59.381).
- Kauffman, S. (1995). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford University Press. ISBN 9780195095999.

Consciousness theories adjacent to T16.

- Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *Biological Bulletin* 215: 216–242. DOI [10.2307/25470707](https://doi.org/10.2307/25470707).
- Seth, A. (2021). *Being You: A New Science of Consciousness*. Faber and Faber. ISBN 9780571337705.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press. ISBN 9780199682737.

Indigenous and relational-ontology scholarship adjacent to Ubuntu-primary.

- Akomolafe, B. (2017). *These Wilds Beyond Our Fences: Letters to My Daughter on Humanity's Search for Home*. North Atlantic Books. ISBN 9781623171667.
- Yunkaporta, T. (2019). *Sand Talk: How Indigenous Thinking Can Save the World*. Text Publishing. ISBN 9781925773996.
- Kimmerer, R.W. (2013). *Braiding Sweetgrass: Indigenous Wisdom, Scientific Knowledge and the Teachings of Plants*. Milkweed Editions. ISBN 9781571313355.

UAP scholarship adjacent to §10. The framework engages this scholarship at the instrumental-record tier only (§10). Cited here as the existing intellectual ecosystem operating in the region of conceptual space the framework’s substrate-physics and universal-scale tiers reach toward, not as evidential support for those tiers.

- Skafish, P., Charbonnier, P., Salmon, G. (eds.) (2017). *Comparative Metaphysics: Ontology after Anthropology*. Rowman & Littlefield International. ISBN 9781783488582.
- Pasulka, D.W. (2019). *American Cosmic: UFOs, Religion, Technology*. Oxford University Press. ISBN 9780190692889.
- Kripal, J.J. (2019). *The Flip: Epiphanies of Mind and the Future of Knowledge*. Bellevue Literary Press. ISBN 9781942658528.
- Dick, S.J. (2003). Cultural evolution, the postbiological universe and SETI. *International Journal of Astrobiology* 2(1): 65–74. DOI [10.1017/S147355040300137X](https://doi.org/10.1017/S147355040300137X).
- Mack, J.E. (1994). *Abduction: Human Encounters with Aliens*. Scribner. ISBN 9780684195391.
- The Sol Foundation (2023–). UAP research and policy initiative. thesolfoundation.org.

This is the integration paper proper. The engineering tier is cited as existence proof shipped under its own DOIs. The integration’s claim is that one structural object is visible at five epistemic levels, grounded in Ubuntu, reaching toward universal scale as research program, with consent at A3+ as the level where the object can be argued about. The bet (§1) is articulated here; the engineering tier is at CCA v3 + RATCHET + CRC v1 + CIRISAgent v2; the formal lake is at formal/CoherenceRatchet/ with tiered status per §11. Where the bet retracts, the engineering tier continues at its own DOIs.

S.D.G.