The Self-Validating World:

An Empirical Continuum of Cumulative Complexity
Across Substrate Rings

Ivan Denysov
idenisov81@gmail.com
ORCID: 0009-0000-3163-4426
Independent Researcher

22 May 2026

Publication notice. This is Paper 0 Part 3 of the United Field Initiative series. The series continues from: Paper 0 Part 1, *Foundations for Measuring the Knowability of a World from Within* (DOI: 10.5281/zenodo.20277783), and Paper 0 Part 2, *Operational Principia of Knowability* (DOI: 10.5281/zenodo.20330208). A separate mathematical companion document, *Mathematical Core of the Self-Validating World Framework*, formalises Theorems 1–2 and Conjecture 1 under named assumptions A0–A4 and is available in the same Zenodo record. A methodological companion to this paper, *Validated Retrodiction: A Two-Coordinate Measure of Cumulative Complexity and a Method for Locating Obligatory Intermediate Nodes* (DOI: 10.5281/zenodo.20513571), develops the measurement method behind Figure 1 (the two-coordinate rank/count construction and the resolution of incommensurable units across substrates).

Abstract

We propose that the world is best understood as a self-validating system of models, in which functional structures (F) — abstract specifications of inputs, outputs, and transformations — realised in substrates (S) mutually validate one another through a binary selection operator Σ (a binary reformulation we propose; see §2.1) acting via cycle closure. The cycle-closure mechanism is the one developed in autocatalytic set theory (Kauffman 1971) and hypercycle theory (Eigen 1971), and rigorously formalised as RAF set theory by Hordijk and Steel (2004). Our contribution is to extend this mechanism from chemistry (its original domain) and economics (Koppl et al. 2022) to substrate rings spanning cosmic history, and to add a temporal element (carrier density ρ_S) that proposes to explain observed acceleration across substrate transitions. The food set for each substrate ring above the root is matter structured by the models of the previous ring; the root ring of fundamental physics uses protomatter — finite raw material of the universe in a pre-validation state — as its food set. We test the framework empirically against four substrate rings spanning 13.8 Gyr — cosmochemistry (elements), mineralogy (mineral structures), biology (protein folds), and cognition (validated knowledge) — and find that each ring closes faster and more recently than the one before: the doubling time of its growth phase shortens by orders of magnitude across the sequence (from ~Gyr for cosmochemistry to ~10^2 yr for cognition, which has not yet plateaued). Two minimal simulations are reported: one confirms emergent saturation under the closure mechanism; one yields a negative result on naive densification, disciplining the temporal claim. We develop the ontology in seven interlocking formulations (binary Σ via cycle closure, nested rings of validation, primitive root contingency, food-set / protomatter resolution, densification, elimination as validation, and structural information preservation) and identify three falsifiable predictions about contemporary substrate transitions.

Keywords: self-validation, functional structure, substrate, cumulative complexity, autocatalytic sets, RAF theory, protomatter, mineral evolution, evolutionary epistemology, information conservation, densification, primitive root.

---

1. Introduction

Each new substrate class begins its growth phase as the previous class approaches saturation. The pattern is empirical; the rest of this paper is concerned with what it means.

1.0 What this paper does

This paper proposes that the world is a self-validating system of models. The world's contents are functional structures realised in substrates, and these realisations mutually validate one another through cycle closure. The cumulative result is what we observe as the present configuration of the universe.

This is a strong ontological claim, and we make it deliberately. The paper supports it in three ways.

First, we extend an empirical observation developed in the United Field Initiative series (Paper 0, Denysov 2026a; Paper 0 Part 2, Denysov 2026b). Where Paper 0 Part 2 anchored a cumulative-complexity pattern in three substrate classes — biology, science, technology — we extend the empirical anchor to four rings spanning 13.8 Gyr (cosmochemistry, mineralogy, biology [protein folds], and cognition). The pattern holds, with each ring approaching its plateau faster than the one before, in a regular way. The substrate-ring continuum is shown in Figure 1 and developed in §3.

Second, we reformulate the selection operator of Paper 0 Part 2 from unary to binary. The "environment" of an F-in-S realisation is, on the binary reading, the union of all other F-in-S realisations that are not, at the moment of analysis, the subject of study. This makes selection

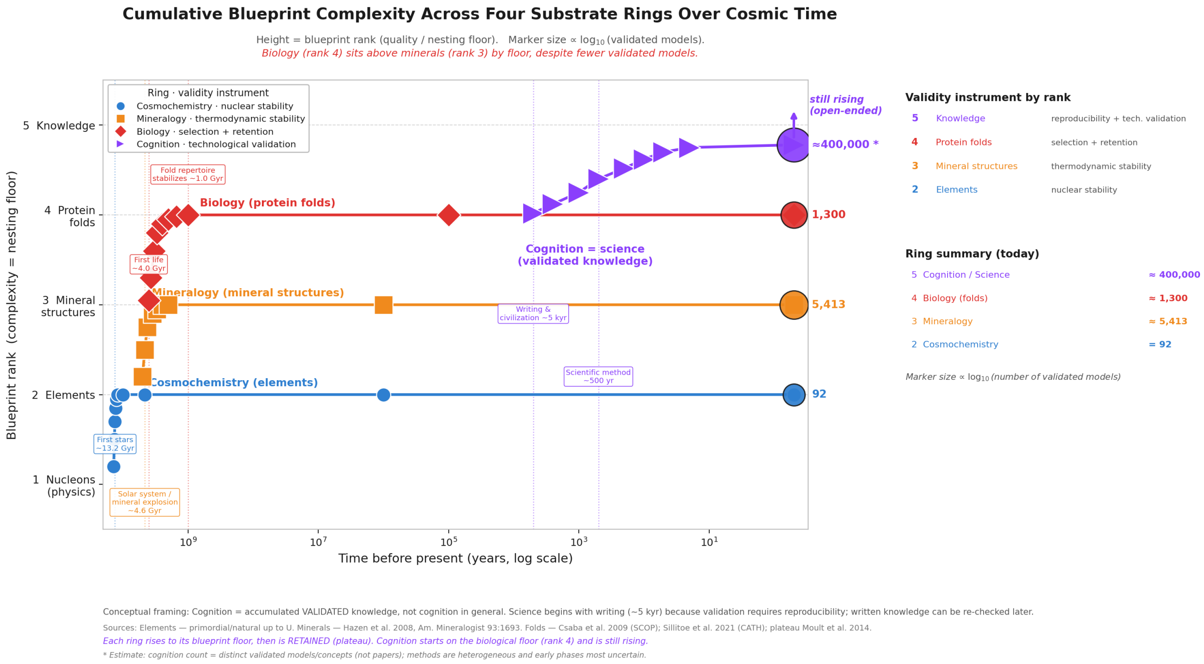

Figure 1: Cumulative complexity across four substrate rings — cosmochemistry (elements), mineralogy (mineral structures), biology (protein folds), and cognition (validated knowledge) — spanning the full age of the structured universe (13.8 Gyr). Each ring rises to its blueprint floor and is then retained on a plateau; successive rings close faster and more recently than the last; the doubling time of the growth phase shortens by orders of magnitude across the sequence (~Gyr to ~10^2 yr), and cognition has not yet plateaued. The load-bearing claim is the temporal ordering of ring closure, not the precise vertical complexity values; the measurement method — the two-coordinate rank/count construction and the resolution of incommensurable units across substrates — is developed in the methodological companion (*Validated Retrodiction*, DOI: 10.5281/zenodo.20513571).

an interaction between populations of realised structure rather than between a realised structure and an external given. The reformulation removes an ontological privilege rather than adding new content, but has substantial consequences for how the framework interprets what selection is doing.

Third, we add a temporal dimension — substrate-specific carrier density $\rho\_S = 1/\tau\_S$, where τ_S is the validation cycle time on a given substrate. This allows the framework to interpret the acceleration of ring saturation across substrate transitions as a consequence of carrier densification: the cycle V-Σ-R is the same on each substrate; the substrates differ in how long one cycle takes. This element is developed informally in §2.5 and formalised as Conjecture 1 in the mathematical core.

We do not claim that the present formulation is final. We expect parts of it to be wrong; we expect future iterations to look different. The mathematical core marks two open problems as **blocking** for central claims — operational specification of τ_S (Open Problem 6.6) and formal proof of structural information preservation (Open Problem 6.1). These are honest limits, not retreats.

What the framework *does* claim, and what we will argue for, is that the empirical continuum of cumulative complexity across four substrate rings, formalised through the binary cycle-closure mechanism with substrate-specific carrier density, supports the interpretation of the world as itself self-validating — not as a system that happens to contain self-validating subsystems, but as one in which self-validation is the assembly principle.

A note on notation. The framework uses several symbols repeatedly: F (functional structure), S (substrate), F-in-S (a realisation of F in S), Σ (the selection operator), V-Σ-R(-T) (the cycle of variation, selection, retention, and optionally transmission), ρ_S (carrier density of substrate S), τ_S (validation cycle time on substrate S), and $\leftrightarrow$^{FS} (the binary closure relation; see §2.1 and mathematical core §1.2). These are defined on first substantive use; the mathematical core contains complete definitions.

The structure of the paper:

- §1.1: the chain of investigation that made this possible, from Lem and Campbell through Kauffman, Eigen, Hordijk-Steel, Kurzweil, and Friston, to the present.
- §2: the ontology in seven interlocking formulations.
- §3: the empirical test against four substrate rings.
- §4: two minimal simulations — one positive, one negative.
- §5: empirically observable effects — aggregation and thermodynamics as a calibration of the saturation form.
- §6: where the present work differs from each predecessor, falsifiable predictions, limitations, and a current empirical possibility.
- §7: predecessors across substrates — the mutual-validation phenomenon observed field by field.
- §8: conclusion.

A separate mathematical companion document (`mathematical_core.md`) formalises Theorems 1-2 and Conjecture 1 under named assumptions A0-A4, and lists eleven open problems including two flagged as blocking.

1.1 A chain of predecessors

The observation that something common operates in biological evolution, in cultural-scientific progress, and in technological development is older than the formal framework we propose here. Over the last six decades, a chain of researchers, working largely in parallel rather than in direct continuity, has been adding links to what we believe is a single emerging picture. We position our work as the next link in this chain rather than as a new departure. The framework presented here is what becomes possible when the existing links are joined together and a few additional connections are made.

Lem (1964), *Summa Technologiae.* Stanisław Lem, working as a science-fiction writer with serious philosophical ambition, observed that biological evolution and technological development share structural properties that are unlikely to be coincidental. He argued that technology is not merely a human extension of biology but a continuation of the same self-structuring process by other means. Lem did not propose a formal apparatus; his contribution is the original recognition that the parallel demands a unified explanation. *His link: the recognition that bio and techno are the same process operating in different substrates.*

Campbell (1974), evolutionary epistemology. Donald Campbell extended the Darwinian selection mechanism beyond biology to the growth of knowledge. He proposed a blind-variation-and-selective-retention algorithm as a general schema for any cumulative learning process, biological or epistemic. *His link: the generalisation of the selection operator beyond biology, establishing that the same Darwinian mechanism produces cumulative structure in cognitive substrates.*

Godfrey-Smith (2009), critique of universal Darwinism. Peter Godfrey-Smith, in *Darwinian Populations and Natural Selection*, articulated the standard contemporary critique of the Campbell-Dawkins-Hodgson move to universalise selection. Godfrey-Smith distinguishes paradigmatic from marginal Darwinian processes by degree along several dimensions (variation in reproductive success, heritability, fitness differences), and argues against treating "selection" as a single mechanism applicable uniformly across substrates. *His link, and our position: we recognise the critique as substantive. The present framework does not claim that selection is* uniformly* applicable across substrates; we claim that mutual validation through cycle closure is the operational mechanism, and that selection in the traditional Darwinian sense is a special case operating on biological reproduction. The Kauffman-Eigen-Hordijk-Steel apparatus permits a broader reading of validation than reproductive-success-based selection alone. We address the universal-Darwinism critique by not committing to a universal-Darwinism position.*

Kauffman (1971 onwards) and Eigen (1971 onwards), autocatalytic sets and hypercycles. Stuart Kauffman introduced the concept of autocatalytic sets: networks of chemical species in which members mutually catalyse each other's formation, such that the overall set sustains itself from a basic food set (Kauffman 1971, 1986, 1993). Manfred Eigen, independently and almost simultaneously, introduced the hypercycle (Eigen 1971; Eigen and Schuster 1977-1979): a cyclic arrangement of self-replicators in which each replicator catalyses the next, and the last catalyses the first, closing the cycle. *Their link: the formal recognition that mutual catalysis through cycle closure is the operational criterion for self-sustaining structure. This is the direct mathematical precedent for what we call binary mutual validation through cycle closure.*

Maturana and Varela (1972 onwards), autopoiesis. Humberto Maturana and Francisco Varela, working in parallel to Kauffman and Eigen but from a biological-philosophical perspective rather than chemical, introduced the concept of autopoiesis: a system is autopoietic if it is a network of component-producing processes such that the components, through their interactions, generate the very network that produced them (Maturana and Varela 1972, 1980). They emphasised that autopoietic systems are **operationally closed** (their organisation is maintained by their internal dynamics) and **structurally coupled** to their environment (they interact with the environment without losing operational closure). Niklas Luhmann (1986 onwards) extended the concept to social systems. *Their link: the conceptual articulation that self-production through closed networks of mutually generating processes is constitutive of living and life-like systems. This is the conceptual precursor to our ontological reading of the world as itself self-validating;*

the formal apparatus comes from Kauffman/Eigen/Hordijk-Steel, but the philosophical articulation that closure is constitutive (not merely descriptive) is theirs.

Kurzweil (1999, 2005), Law of Accelerating Returns. Ray Kurzweil proposed, across two books (*The Age of Spiritual Machines* and *The Singularity Is Near*), that information-processing capacity has been doubling on progressively shorter timescales across substrate transitions, plotted on log-log axes — from computation across mechanical, relay, vacuum tube, transistor, and integrated-circuit substrates, and extending the picture informally backward to evolutionary milestones. The surface similarity to the present framework is significant: substrate transitions, accelerating doubling, log-log visualisation, AI as a candidate next substrate. *His link, and our delimitation: Kurzweil identified the empirical phenomenon of substrate-transition acceleration informally and at journalistic rigour. He did not propose a closure mechanism, did not formalise the rate-constant decomposition (λ_S, ρ_S in our terms), and his specific quantitative predictions have a mixed track record. We provide the formal mechanism — RAF-style closure as the operational criterion for retention, with substrate-specific carrier density as the source of cross-substrate acceleration — that Kurzweil's empirical observation requires.*

Hordijk and Steel (2004 onwards), RAF set theory. Wim Hordijk and Mike Steel developed the formal mathematical apparatus for autocatalytic sets, introducing the framework of Reflexively Autocatalytic and Food-generated (RAF) sets. They proved theorems about the conditions under which RAF sets emerge, gave efficient algorithms for detecting RAF subsets in chemical reaction networks, and established that the level of catalysis required for emergence is biochemically realistic (Hordijk and Steel 2004; Hordijk et al. 2010; Steel et al. 2019). Subsequent work (Koppl et al. 2022; Cazzolla Gatti et al. 2020) extended the framework to economics. *Their link: the rigorous mathematical apparatus for closure-based mutual catalysis, and the first extension beyond chemistry. We adopt their formal framework and extend its empirical reach.*

Maynard Smith and Szathmáry (1995), *The Major Transitions in Evolution*. John Maynard Smith and Eörs Szathmáry identified a sequence of major evolutionary transitions — replication of molecules, emergence of the cell, eukaryotic origin, multicellularity, social organisation, language — each characterised by (i) a change in how information is stored and transmitted and (ii) the formation of a new level of biological individuality in which entities that previously replicated independently come to replicate only as part of a larger whole. Szathmáry (2015) updated the framework as "MET 2.0", folding in new evidence and refining the list. The framework has not gone unchallenged; Calcott and Sterelny (2011, eds., *The Major Transitions Revisited*) document substantive disagreements with the original formulation. *Their link, and the surrounding debate: the identification that nested levels of organisation emerge through transitions in information storage and transmission. This is the closest direct precedent for what we call substrate transitions and nested rings. We treat MET as the family of frameworks we extend, rather than as a settled position.*

Odling-Smee, Laland, and Feldman (2003), niche construction theory. The standard biological account of how organisms actively modify their selective environments, generating cross-generational feedback that affects subsequent evolution. NCT insists that "environment" is not a separable external given but is constructed in part by the organisms it then selects on. *Their link: NCT anticipates our binary reformulation of Σ at the biological substrate level. Where NCT treats the construction of environment as a special biological phenomenon (organisms modifying environment), we treat the absence of a privileged environment as constitutive of the operator: every F-in-S realisation is part of the validation environment for every other. Our move generalises NCT's biological insight to a structural feature of the validation cycle on every substrate.*

Hodgson and Knudsen (2010), generalised Darwinism. Geoffrey Hodgson and Thorbjørn Knudsen, in *Darwin's Conjecture*, developed a sustained argument for generalised Darwinism in social and economic systems, engaging both Universal Darwinism advocates and critics including Godfrey-Smith. They distinguish ontological from heuristic-analogical uses of Darwinian frameworks across substrates. *Their link: Hodgson and Knudsen's careful delimitation of what generalised Darwinism does and does not claim is a model for our delimitation of the cycle V-Σ-R(-T) across substrates. We adopt their ontological-vs-heuristic distinction: the cycle's structure is ontological (same V, Σ, R across substrates); the specific operationalisation per substrate is heuristic.*

Friston (2010 onwards), free energy principle. Karl Friston's free energy principle proposes a unified variational-Bayesian account of self-organisation, from cells to cognition. The framework claims substrate-spanning explanatory scope analogous to our own. *Their link, and our delimitation: FEP and the present framework address different aspects of self-organisation. FEP is a variational principle for how systems minimise prediction error; ours is an empirical-and-structural claim about cumulative complexity through cycle closure. The two are not incompatible — FEP could be operating* within* a substrate ring as the mechanism by which retention happens at the substrate's characteristic cycle time — but the two frameworks have different empirical anchors and different formal apparatuses. We do not attempt to integrate them here; the relationship is left open.*

Smolin (1992, 2013), cosmological natural selection. Lee Smolin proposed cosmological natural selection — the hypothesis that universes are selected for properties that maximise black-hole production, with multiverse-scale Darwinian dynamics. This is a different kind of substrate-spanning claim from ours: Smolin extends Darwinian thinking to physics itself, requiring multiverse ontology and selection at the universe-level. *His link, and our delimitation: our framework is empirical-and-structural, operating within a single observed universe and across substrate rings within that universe. Smolin's framework is metaphysical and operates at the inter-universe level. The two frameworks could in principle be compatible — cosmological natural selection at the multiverse level could underlie the existence of a single universe within which substrate-internal self-validation operates — but they address different questions and require different evidential bases. We do not engage further with multiverse arguments in this work.*

Hazen and colleagues (2008 onwards), mineral evolution. Robert Hazen and collaborators extended the principle of cumulative diversification to a substrate prior to biology: minerals. They documented how mineral species count grew from approximately twelve pre-stellar ur-minerals to over five thousand species today, in identifiable stages driven by stellar nucleosynthesis, planetary differentiation, plate tectonics, and biological mediation. Hazen et al. (2024)

further demonstrated that abiotic minerals can reach assembly indices previously proposed as biosignature thresholds, indicating that the line between abiotic and biotic complexity is not sharp. *Their link: the empirical demonstration that the cumulative-complexity pattern extends backward in time to the formation of matter itself, before biology.*

Krivovichev and colleagues (2017 onwards), mineral complexity. Sergey Krivovichev and colleagues developed information-theoretic measures of mineral complexity (Shannon information per atom and per unit cell) and showed that both diversity and complexity increase with geological time. *Their link: the quantitative information-theoretic operationalisation of cumulative complexity, applicable across substrate classes.*

Cronin, Walker, Sharma and colleagues (2017 onwards), assembly theory. Lee Cronin, Sara Walker, Abhishek Sharma and collaborators proposed assembly theory as a measure of molecular complexity bridging chemistry and biology, with an assembly index threshold proposed as a biosignature. Subsequent work (Hazen et al. 2024) has contested the sharp abiotic/biotic threshold while accepting the measure itself as informative. *Their link: a substrate-spanning quantitative complexity measure, applicable from chemistry to biology to technology.*

Wheeler, Bekenstein, 't Hooft, Susskind (1973 onwards), information physics. Through Wheeler's "it from bit" programme, Bekenstein's black hole entropy, 't Hooft and Susskind's holographic principle, and the resolution of the black hole information paradox, contemporary physics has converged on the view that information is conserved at the deepest level and that physical structure is best understood as information-theoretic. *Their link: the most fundamental layer — information conservation as a property of the physical world itself, complementary to but distinct from the structural-information preservation we develop here.*

Hledík, Barton, Tkačik (2022, *PNAS*), information accumulation in evolution. A recent formal treatment, showing how natural selection accumulates information measured in bits and proving general bounds on the rate of accumulation per generation. *Their link: a formal information-theoretic apparatus for the accumulation we observe empirically across substrate classes; ours generalises beyond biology.*

Salthe (1985), hierarchical systems theory. Stanley Salthe developed a philosophical framework for hierarchical organisation in nature, with attention to the structure of cross-level relationships. Less quantitative than the works above, his contribution is the explicit attention to the hierarchical-structural form of natural organisation. *His link: the philosophical articulation that hierarchies are not mere descriptive layers but have substantive structural properties.*

1.2 What the chain has, and what is still missing

Taken together, these contributions have established the following:

- That mutual catalysis through cycle closure is the operational mechanism by which self-sustaining structure emerges (Kauffman 1971; Eigen 1971; Hordijk and Steel 2004).
- That closure is constitutive rather than merely descriptive of living and life-like systems (Maturana and Varela 1972, 1980; Rosen 1991; Letelier et al. 2003).

- That a cumulative-complexity pattern exists across biological evolution, cognitive development, and technological growth (Lem, Campbell, Maynard Smith and Szathmáry).
- That the same pattern extends backwards to mineral evolution and chemical-substrate transitions (Hazen, Krivovichev, Cronin/Walker).
- That the autocatalytic-set framework can be extended beyond chemistry to economics (Koppl et al. 2022; Cazzolla Gatti et al. 2020).
- That information conservation is the fundamental physical property underlying all of this (Wheeler, Bekenstein, holographic principle).
- That the accumulation can be made formal in information-theoretic terms (Hledík et al.).
- That the organisation is hierarchical, with substantive relationships across levels (Salthe; Maynard Smith and Szathmáry).

The chain has already established the formal apparatus we need for cycle closure (Kauffman, Hordijk, Steel), and has already taken one step beyond chemistry (toward economics). Our work continues this trajectory.

What remains missing, and what we attempt to provide here, is the following:

- A **single empirical demonstration** that the cumulative pattern is continuous across substrate classes spanning the full age of the structured universe, from primordial nucleosynthesis to current microelectronics. Autocatalytic set theory has been developed extensively for chemistry and recently for economics; the cognitive and technological substrates have been treated separately in their own literatures; mineralogy has its own programme (Hazen). No previous work has assembled four substrate rings onto a single time axis spanning 13.8 Gyr. We do this and find a consistent regime (§3).

- A **conjectured temporal mechanism** for the observed acceleration of cumulative complexity across substrate transitions. The empirical pattern shows each ring closing faster and more recently than the one before, the growth-phase doubling time shortening by orders of magnitude across the sequence (from ~Gyr to ~10^2 yr, cognition not yet plateaued). RAF theory does not address this question; the autocatalytic-set framework is largely static (does the set close?) rather than temporal (how fast does the set grow?). We propose that substrate transitions densify the carrier — the same RAF-style cycle operates on each substrate, but with progressively shorter validation cycle times in absolute units (§2.5, §4).

- An **explicit ontological reading** of the world as itself self-validating. Autocatalytic set theory provides the formal mechanism; it does not commit to the ontological claim that the world's structure simply *is* the cumulative record of all RAF-style closures that have occurred. We make this ontological claim explicit and connect it to the empirical record across substrate classes (§2.1-2.7).

- A **structural information preservation theorem** for the RAF-extended framework. We propose that the cumulative history of validation events is encoded in the relational structure of the world, complementary to microphysical unitarity (Wheeler, Bekenstein). This is new content (§2.7).

1.3 The present paper

We build on Paper 0 (Denysov 2026a), which established four primitives — environment E, functional structure F, substrate S, observation function O — for the formal study of knowability. Paper 0 Part 2 (Denysov 2026b) extended these into a cycle V-Σ-R(-T) and anchored the cumulative-complexity pattern empirically in three substrate classes: biology, science, and technology. Paper 1 (Denysov 2026c) established that absolute scale parameters of a local engine are operationally unknowable to internal observers (Theorem 1 of that paper). The present paper extends and reformulates this framework as follows.

Conceptually: we reformulate the selection operator Σ as binary mutual validation between F-in-S realisations, with cycle closure as the operational criterion. The cycle-closure mechanism is the same one developed by Kauffman (1971), Eigen (1971), and rigorously formalised by Hordijk and Steel (2004) as RAF set theory. Our contribution is not to invent the mechanism but to extend its application from chemistry (where it was originally developed) to the full hierarchy of substrate classes spanning cosmic history. We develop the ontology of the world as a self-validating system through seven interlocking formulations: binary Σ via cycle closure, nested rings of validation, primitive root contingency, food-set / protomatter resolution, densification, elimination as validation, and structural information preservation.

Empirically: we extend the empirical-continuum test to four substrate rings spanning 13.8 Gyr and find a continuity pattern in which each new ring begins its growth phase as the previous ring approaches saturation, and each ring approaches its plateau faster than the one before.

Formally: we derive multiplicative growth and bounded accumulation as consequences of binary mutual validation (which is the RAF closure condition on a growing pool), and we propose that the observed acceleration of ring saturation across substrate transitions is consistent with substrate-specific carrier density rather than with changes in the cycle itself. The first two derivations recapitulate, in slightly different form, results already in the RAF literature; the densification mechanism is new.

Connectively: we identify three falsifiable predictions about contemporary substrate transitions and locate the present moment within the framework's structure.

The combined result suggests that the cumulative-complexity pattern is not a peculiarity of biological-and-later evolution. It is the visible signature of a single self-structuring process that has been operating throughout cosmic history. The framework presented here is the next link in a chain of investigation that the previous links have made possible. We benefit, in particular, from over fifty years of mathematical development of autocatalytic set theory; our task is to take that apparatus to substrate classes it has not yet reached.

2. Ontology

2.1 Self-validation as the world's assembly principle

In Paper 0 Part 2 the selection operator Σ took the type $\Sigma : \mathcal{P}(\mathcal{S}) \times \mathrm{E} \to \mathcal{P}(\mathcal{S})$. The environment E appeared as a single argument representing the differential context against which substrate configurations are tested. The framework did not decompose E further.

A note on terminology. We retain the symbol Σ and the name "selection operator" from Paper 0 Part 2 for continuity. In the present paper, "selection" is read in the broad sense of *validation through cycle closure*, not in the narrow Darwinian sense of differential reproductive success. The operator's domain extends across substrates where reproduction in the strict sense does not occur (cosmochemistry, mineralogy), and the Godfrey-Smith (2009) cluster of paradigm Darwinian features (heritability, variance in reproductive success, fitness differences) is not presupposed. "Validation operator" or "closure operator" would be more terminologically precise; we retain "selection" for series continuity with the explicit caveat that the broader reading is in force throughout.

The present development takes a step into E and finds that what has been called "environment" is in fact the union of all other F-in-S realisations that are not, at the moment of analysis, the substrate under study. The boundary between "subject of study" and "environment" is methodological, not ontological. Selection (validation) is therefore a binary interaction between F-in-S realisations:

$$\Sigma : \mathcal{P}(\mathcal{S}) \times \mathcal{P}(\mathcal{S}) \to \mathcal{P}(\mathcal{S}).$$

The two arguments are populations of substrates carrying functional structures; the operator returns the joint configurations that survive mutual validation. Survival is not given by an external standard. Survival is what remains after every realised F-in-S has been tested against every other realised F-in-S that it encounters.

The world is therefore the totality of F that have been realised in S and have passed the joint validation of all other realised F-in-S they have interacted with. The cycle does not operate "in" the world; the cycle is the world's continuing self-structuring.

The assembly point of the world is the transition F → S. An F not realised in any S is a candidate, not a constituent of the world. An S that does not carry any distinguishable F is undifferentiated. The moment of embedding — F realised in S, in a particular configuration, in interaction with other F-in-S — is the moment at which the world acquires content.

2.2 The scope and limits of the self-validation reading

Before developing the seven formulations further, it is useful to be precise about what the self-validation reading does and does not claim.

What it claims. The world's *retained structure* is what has passed joint validation against the rest of the realised structure with which it has interacted. The framework reads the cumulative-complexity pattern of §3 as the visible record of this process: every retained F-in-S is, by definition of "retained", something that closed a validation cycle with other F-in-S; every transition from one substrate class to the next is the formation of a new closure layer atop the previous one. The empirical argument for the reading is the empirical record itself, developed quantitatively in §3 and visualised in Figure 1. It is not a separate philosophical argument.

What it does not claim. Several things should be distinguished from the self-validation reading and not conflated with it.

The framework does not claim to *derive* the laws of physics from the self-validation principle. Conservation laws follow, by Noether's theorem, from continuous symmetries of the action (Noether 1918); the framework reads them as the structural form that conservation has taken under self-validation but does not propose to replace the Noetherian derivation. The framework is consistent with conservation having a deeper explanation; it is not an alternative explanation.

The framework does not claim to *derive* cross-level consistency from self-validation. Standard supervenience and reductionist accounts in philosophy of science (Kim, Fodor, the multiple-realisability literature) already give principled reasons why higher-level regularities cohere with lower-level laws. The framework is consistent with these accounts; it adds the diachronic claim that the cross-level consistency we observe is the *retained* portion of cross-level interactions, and that earlier inconsistent configurations were eliminated. This is a different explanatory move from "supervenience" but does not contradict it.

The framework does not claim that physical reality is computational. The framework is agnostic on whether the universe is best understood computationally (Wolfram), informationally (Wheeler), or otherwise. The validation cycle V-Σ-R is a *process structure*, not a substrate of metaphysics. Mapping the cycle onto computational metaphors is a presentational choice, not a substantive claim about what the world is "made of".

What the framework provides as positive content. Three commitments. First, that the operator Σ is binary: validation is between F-in-S realisations, with no privileged "environment" given externally (§2.1). Second, that substrate transitions are gradual closures of cross-level validation density, not sudden appearances of new functions (§2.3, §2.5). Third, that elimination is itself a validation event with structural consequences, not a counter-example to cumulative growth (§2.6).

The empirical test of the framework, then, is whether these three commitments produce predictions that match observed cumulative complexity across substrates. §3 develops this test against four substrate rings; §6.3 lists three falsifiable predictions. The empirical case for the self-validation reading rests on these, not on a priori arguments.

2.3 Nested rings of self-validation

The world is structured as a sequence of closed self-validation rings (Figure 2). The root ring is a population of F-in-S realisations whose mutual validation has stabilised. When a ring is closed in this sense — when the F's within it have validated one another and the population has reached a configuration that is internally coherent — the ring becomes the substrate ground on which the next ring forms.

Cross-level validation density is the key measurable quantity. F-in-S at level N+1 must validate not only with other F-in-S at level N+1 but also with the closed F-in-S at level N on which they rest. A new substrate class does not appear instantaneously; it appears progressively as the density of cross-level validations increases. Substrate transitions are extended processes of cross-level mutual integration, not sudden events.

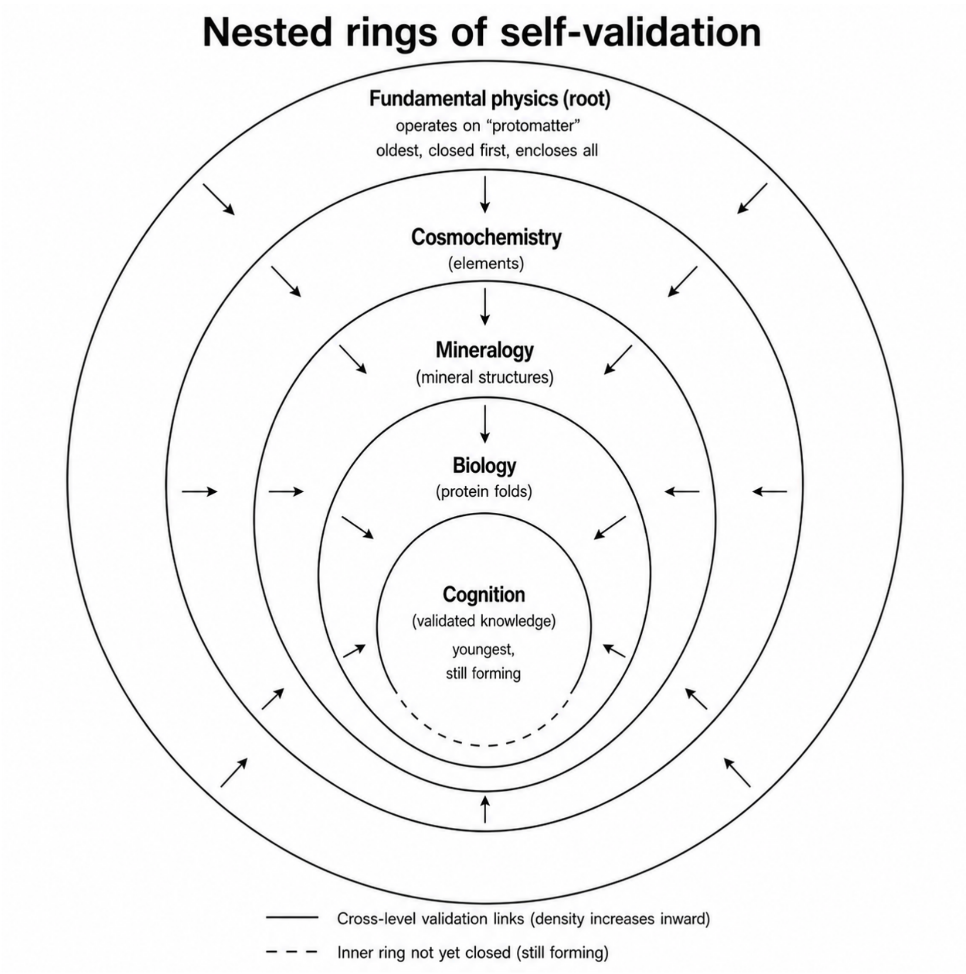

Figure 2: Nested rings of self-validation. Each substrate ring forms atop the closed ring beneath it. The root ring (fundamental physics) operates on protomatter; subsequent rings (cosmochemistry, mineralogy, biology, cognition) form as cross-level validation density accumulates between adjacent levels.

2.4 The primitive root and contingency

The first ring to close — the root — was not the optimal possible root from the space of possible roots. It was the first that crossed the minimal threshold of mutual coherence. The world as we observe it is the trajectory of subsequent construction on this contingent primitive root.

This is anti-anthropic in a specific sense. The world is not fine-tuned in the sense of having been selected from possibilities by some principle. The world is the result of a contingent first closure followed by cumulative construction. The root ring cannot be revised; subsequent levels can only build on what the root supplies. The arrow of structure has a direction inherited from the order of closure.

The food set of the root ring. A natural question arises from the parallel with autocatalytic set theory: an RAF set, by definition, requires a food set of building blocks not produced within the set itself. What does the root ring use as its food set, if all subsequent rings use their predecessors' structured matter as food?

We propose that the root ring's food set is **protomatter** — matter that exists prior to the closure of any model, with the minimal properties (some form of interaction, some form of distinguishability) required to participate in the first self-validation. Protomatter is not abstract potentiality; it is the finite raw material of the universe. The root ring of fundamental physics is the first closed set of mutually-consistent models that gave protomatter its structured form (particles, interactions, conservation laws).

This resolution preserves the self-validating world ontology: protomatter is part of the world (not external to it), but it exists in a pre-validation state and does not itself require to be validated. The framework asserts only that **once any model is realised as a structured form of matter, it must mutually validate with other realised models**. Protomatter prior to any model exists by virtue of being the raw material the world has to work with; the question of why protomatter exists at all is on the boundary of the framework's reach.

This formulation is consistent with bootstrap programmes in particle physics, with self-consistency conditions in quantum field theory, and with the holographic principle, all of which seek the fundamental laws of physics as conditions imposed by the theory on itself rather than imposed from outside. The framework's reading is that the laws of physics constitute the root ring, and they closed against the protomatter background through the same mechanism that all subsequent rings close.

Food set for all higher rings. For every substrate ring above the root, the food set is **matter-structured-by-the-models-of-the-previous-ring**. Cosmochemistry's food set is protomatter-structured-by-fundamental-physics (particles, atoms). Mineralogy's food set is matter-structured-by-cosmochemistry (elements, simple molecules). Biology's food set is matter-structured-by-mineralogy-and-chemistry (organic molecules, water, simple energetic gradients). Cognition's food set is matter-structured-by-biology (brains, sensory systems). Technology's food set is matter-and-information-structured-by-cognition (knowledge, language, engineered materials).

The distinction between food set (quantity of structured matter) and the models that structured that matter (quality of structure) is essential. Each higher substrate ring inherits a finite quantity of structured carrier from its predecessor; the structural form of that carrier is what distinguishes one substrate's food set from another.

2.5 Densification as the direction of substrate transitions

What we have called "substrate transitions" or "new rings" are not the appearance of qualitatively new functions. They are densifications of the same fundamental function — self-validation through mutual coherence — onto progressively more compact substrate carriers.

Biology is chemistry densified to the point at which it can replicate. Cognition is biology densified to the point at which it can model itself. Technology is cognition densified to the point at which it can construct itself. At each step, the same fundamental function is realised at higher density on a more compact material carrier.

This densification produces two empirical signatures: (1) each ring closes faster than the one before — the growth-phase doubling time shortens by orders of magnitude across substrate rings, and (2) the absolute number of realised F-in-S grows even as the physical carrier becomes more compact, because density permits more carriers in the same physical universe.

The example of automobile evolution illustrates the principle in microcosm. The Ford Model T of the 1920s realised the function of transport on a heavy, inefficient carrier. A current electric vehicle realises the same function on a much more compact and integrated carrier, with the function distributed across mechanical, electronic, and software components. The current global fleet is approximately three orders of magnitude larger than the 1920s fleet, despite each unit being functionally denser, because density on a unit basis releases material capacity for more units.

The same pattern operates in biology (modern bacterial cells are functionally denser than early proto-cells and far more numerous) and in technology (the global transistor count exceeds the human population by many orders of magnitude).

The direction of densification is toward an asymptotic limit: functional density without redundant carrier. The geometric image is movement toward a point, but the point is the asymptote of functional density, not a spatial singularity. Successive rings are progressively denser inner rings, with the root ring as the outermost and least dense layer of structured matter.

We use "direction" here descriptively, not normatively. Densification is an observed historical pattern in the record so far, not a law of directed progress, and carries no claim that the sequence was inevitable, optimal, or value-laden. Nothing here guarantees that further rings must form, or that the pattern will continue; it is the regularity we observe, offered as such.

2.6 Elimination as validation

The operator Σ does not distinguish between "correct" and "incorrect" outcomes. Σ records the result and eliminates configurations that fail mutual validation. Invalid F-in-S do not propagate as evolving models, but the fact of their validation — and its negative outcome — is preserved in the structure of the substrate.

Apparent regress in standing diversity is not a counterexample to the framework but a direct demonstration of Σ in operation. Mass extinctions in biology, abandoned technological lineages (the DeLorean as a complete product, the Concorde as an aircraft class, the videocassette as a substrate), and culled scientific theories are all instances of Σ doing the work that Σ is defined to do.

Two temporal scales of elimination operate together. Local elimination at time t: F-in-S configurations are validated against current substrate conditions; failures are removed. Long-term densification: earlier “primitive” F-in-S are progressively replaced by more refined descendants realising the same function at higher density.

The disappearance of a specific F-in-S realisation does not entail the loss of information. Components of F may be transmitted (T-component of the cycle) to subsequent realisations. The structural trace of the original elimination is preserved in the substrate. The Sepkoski curve of standing marine genus diversity, with its mass extinction events, is direct evidence that Σ operates on the biological substrate (elimination of realisations), while the cumulative count of all marine genera ever described grows monotonically and records the retention of evolutionary information (the T-component). Here we use the curve only to illustrate elimination and retention under Σ, not as the complexity measure of the biological ring — that role belongs to the protein-fold repertoire (§3.1).

2.7 Structural information preservation

Information about events that have occurred is preserved in the current structure of the world independently of whether transmission between observers has taken place. The world does not lose its history; history is inscribed in the structure of retained F-in-S and in the substrate that has been modified by those realisations and their interactions.

This is distinct from operational information transmission addressed by Paper 0 Part 2’s darkroom thought experiment. Part 2 establishes that knowledge transmission between observers requires the T-component of the cycle. The present claim is independent: even when transmission between observers has failed, the event itself leaves structural traces, recoverable in principle from analysis of the current state.

Four types of traces can be distinguished. Material traces: molecules, heat, radiation transferred to surrounding substrate. Structural traces: the absence of an expected F-in-S creates measurable discontinuities. Functional traces: expected validations that did not occur are themselves data. Cumulative traces: subsequent F-in-S that filled the role of an eliminated one carry information about the origin.

The distinction between logical preservation (always, in principle) and practical accessibility (degraded by entropy) is preserved. The framework asserts that the record exists; it does not assert universal practical reconstructibility.

Examples consistent with structural information preservation are visible in at least six independent disciplines: palaeontology reconstructs species that died hundreds of millions of years ago; forensic science reconstructs events from traces; archaeology reconstructs ancient civilisations; geology reconstructs climates and configurations spanning billions of years; evolutionary biology reconstructs the deep tree of life from comparative genomics; and physics, through unitarity and the holographic principle, asserts microphysical information conservation. These examples operate at very different scales and use different methods with different reliability standards; we do not claim that their convergence constitutes a unified empirical proof of the framework's structural-information thesis. We claim that the framework's reading is consistent with the operational practices of these disciplines, which is all that is needed to motivate the open problem (math core §6.1) of formalising structural information preservation rigorously.

Connection to Paper 1 Theorem 1: Paper 1 establishes that absolute scale parameters of a local engine are operationally unknowable to any internal observer. Structural information preservation does not conflict; it complements. Theorem 1 concerns absolute parameters not recorded in any relative configuration. Structural preservation concerns the relative structure of retained F-in-S, fully encoded in the current state. Internal observers face a strict limit on absolute scales but have, in principle, full access to the structural history of their world. We note that the §6 complementarity claim with Paper 1 depends on completing the open problem of formally proving structural information preservation (math core §6.1, flagged as **blocking**); the current claim is a conjectured complementarity pending that proof.

With the ontology in place, we turn to the empirical test that motivated the present extension of the United Field Initiative framework.

3. Empirical test: the four-ring continuum

We tested the framework against four substrate rings spanning 13.8 Gyr. For each ring we collected published data on cumulative complexity (in substrate-specific units) over time and tracked the time each ring takes to approach its plateau.

3.1 Substrate rings

Cosmochemistry. Number of detectable chemical elements present in cosmic abundance over cosmic time. From three elements (H, He, Li) approximately 20 minutes after the Big Bang, through stellar nucleosynthesis producing CNO and Fe-peak elements (Pop III and Pop II stellar generations) per the classic synthesis of Burbidge, Burbidge, Fowler and Hoyle (1957), to the ~92 naturally occurring elements present today. Data points span 13.8 Gyr.

Mineralogy. Cumulative count of mineral species realised in the Solar System and on Earth, from Hazen et al. (2008). From ~12 pre-stellar ur-minerals to ~5413 IMA-recognised species at present, spanning the full history of the Solar System (~4.6 Gyr) and pre-Solar history (an additional ~9 Gyr).

Biology (protein folds). Cumulative count of distinct protein fold superfamilies, the validated structural primitives of the biological ring, from structural-classification databases (Csaba et al. 2009 [SCOP]; Sillitoe et al. 2021 [CATH]). The fold repertoire grows through early biological history and stabilises at ~1,300 superfamilies, approaching its plateau by roughly 1 Gyr ago.

We measure the biological ring by its fold repertoire rather than by genome size or taxonomic diversity: folds are the retained, mutually validated structural units on which biological function is built, and they saturate (the discovery rate of genuinely new folds has fallen sharply), whereas genome size and species counts track later diversification on an already-closed apparatus rather than the closure of the ring itself.

Cognition (validated knowledge). Cumulative count of validated, re-checkable knowledge items — the validated primitives of the cognitive ring. The ring opens not with cognition in general — nor with language or oral tradition, which transmit but do not allow a claim to be re-checked against a fixed record across generations — but with the first carrier that makes knowledge durably and independently re-checkable: writing (~5 kyr ago). What defines the ring is reproducible cross-generational validation, not the mere presence of cognitive agents; writing is the point at which a knowledge claim can be returned to, tested, and either retained or eliminated by others later. From that onset the ring grows through the later phases of systematic science and technology — the scientific method sharpening reproducibility, working technology being its sharpest form — to ~400,000 distinct validated models/concepts today. Unlike the lower rings, the cognitive ring has not yet reached a plateau: it is still rising. We treat science and technology as later growth phases of this one ring rather than as separate substrates, since both rest on the same validity instrument (reproducibility / technological testing).

3.2 Results

For each ring we record its validity instrument, the count of validated primitives reached, and the doubling time of its growth phase (with the moment it reaches its plateau in parentheses):

Ring	Validity instrument	Validated primitives	Doubling time (growth phase)
Cosmochemistry (elements)	nuclear stability	~92	~Gyr (plateau by ~4–5 Gyr ago)
Mineralogy (mineral structures)	thermodynamic stability	~5,413	~10^2 Myr (plateau by ~2 Gyr ago)
Biology (protein folds)	selection + retention	~1,300	~10^2 Myr (plateau by ~1 Gyr ago)
Cognition (validated knowledge)	reproducibility + technology	~400,000	~10^2 yr (still rising)

The combined empirical pattern is shown in Figure 1 (above). Two regularities stand out.

Acceleration. Each ring closes faster than the one before. The doubling time of its growth phase shortens by orders of magnitude across the sequence — from ~Gyr for cosmochemistry to ~10^2 yr for cognition (which has not yet plateaued). This is the empirical signature of carrier densification (§2): the same validation cycle runs on progressively denser carriers, so each ring closes faster than the last. We use the doubling time as a descriptive characteristic of each ring's growth phase — a standard measure for accumulation curves — not as a proposed law; the conjectured rate law ($\lambda \propto 1/\tau$) and its status are discussed in the methodological companion (*Validated Retrodiction*, §6).

Continuity. Each ring begins its growth phase as the previous ring approaches saturation. Cosmochemistry approaches saturation (~92 elements) by approximately 4–5 Gyr ago, as mineralogy begins. Mineralogy approaches its plateau (~5,400 species) by approximately 2 Gyr ago, as the biological fold repertoire is assembled. The fold repertoire stabilises (~1,300 superfamilies) by approximately 1 Gyr ago, on which later biological diversification proceeds, until the cognitive ring opens with writing (~5 kyr ago) and continues to rise through science and technology to the present.

Status and limits of the pattern. The doubling times above are order-of-magnitude estimates, not precise measurements. Each is computed from the growth phase between a ring's onset and the point at which its count of validated primitives levels off, read from heterogeneous data of differing source and resolution; the exact value is sensitive (by a factor of order two) to where the plateau is taken, so only the order of magnitude is reported. What the framework's claims concern is the *ordering and the order-of-magnitude separation* of these times — that each ring closes, and reaches its plateau, more recently than the one before — not their exact values. At that resolution the ordering is unambiguous: cosmochemistry plateaus by ~4–5 Gyr ago, mineralogy by ~2 Gyr ago, the biological fold repertoire by ~1 Gyr ago, while cognition has not yet plateaued and is still rising. Each ring closes later and over a shorter growth phase than the last, and no plausible revision of the underlying data would reorder them.

A note on status. These analyses are exploratory rather than confirmatory. The data are heterogeneous in source, sampling regime, and aggregation method, and the doubling times are deliberately reported as order-of-magnitude. The vertical complexity axis of Figure 1 is illustrative; the load-bearing claim is the temporal ordering of ring closure, not any precise complexity value. We report the pattern to show its qualitative regularity, not to imply statistical adjudication between alternative models.

The earliest ring, cosmochemistry, is the least clean case and the most instructive. Its primitive space is bounded — there are only so many stable nuclei — so it reaches saturation through a long deceleration rather than a sharp plateau, and its onset and saturation are harder to date than the later rings. The fuzziness at the earliest level cautions against treating the continuity across rings as exact, even where it is qualitatively striking. We therefore present the four-ring continuity as a qualitative regularity — a consistent ordering of ring closure times across substrate transitions — rather than as a precise law, with the cognitive ring still in its open growth phase and no claimed plateau.

4. Simulation: closing one gap, disciplining another

The theoretical results presented in §2 and the mathematical core inherit broad support from prior literature: multiplicative early-phase dynamics from Eigen quasispecies theory; logistic saturation from Verhulst, Bass, and standard population dynamics; closure conditions from

Kauffman, Eigen, and Hordijk-Steel RAF formalism. Two gaps, however, are not closed by inherited results. First: does the specific derivation we propose — saturation through feature-space exhaustion *within a closure network* — actually produce plateau in a minimal model? Second: does the naive interpretation of densification (more candidates per fixed unit of time) reproduce the empirical acceleration pattern? We ran two minimal simulations to address these questions; results are shown in Figure 3.

Simulation v2 — emergent saturation (Figure 3, Panel A). Four substrate levels with overlapping feature bands; candidates generated at uniform rate per active level; retention if the candidate's inputs are covered by the available pool's outputs and its outputs introduce novel functionality. No plateau is programmed; no acceleration is programmed. Result: all four levels reach plateau (saturation values determined by their feature-band sizes, ~32-42 retained F-in-S per level). The plateau is the emergent consequence of finite feature space combined with the closure condition. This closes the first gap: Theorem 2's derivation — saturation through feature-space exhaustion within closure — works in a minimal model.

Simulation v3 — negative result on naive densification (Figure 3, Panel B). Same architecture as v2, but candidate-generation rate per tick scaled geometrically across levels: 1, 5, 25, 125 candidates/tick. The hypothesis tested: does "more candidates per tick" reproduce the observed cross-substrate acceleration? Result: no. Levels with high density rates saturate almost instantly (1-2 ticks to plateau) rather than showing extended exponential phase with progressively faster ring closure. The negative result is informative. It rules out the naive interpretation and disciplines Conjecture 1: densification must mean *shorter cycles in absolute time*, not *more candidates per fixed time*. The two interpretations diverge sharply in dynamics, and the empirical pattern fits only the former.

What the simulations do and do not show. The simulations are minimal models with abstract F-in-S representations; they do not simulate any actual substrate. They show that (a) the specific saturation mechanism we derive is consistent with minimal-rules dynamics, and (b) one tempting interpretation of densification is wrong. They do not — and could not — directly validate Conjecture 1, which concerns the structural-vs-temporal decomposition of rate constants across real substrates. That decomposition requires independent operational measurement of τ_S per substrate (Open Problem 6.6 in the mathematical core), not simulation.

Code availability. Simulation source code (Python, ~200 lines for v2; ~250 for v3) is included in the Zenodo upload accompanying this preprint, in the supplementary materials. The code is sufficient to reproduce both panels of Figure 3 deterministically (fixed RNG seed) and to permit exploration of parameter sensitivity.

Minimal simulations of hierarchical mutual validation are shown. **Panel A** (left): emergent saturation under uniform candidate rate with cross-level validation enabled (v2). Four substrate levels with overlapping feature bands; each reaches plateau (32-42 retained F-in-S per level) as feature space exhausts. No plateau is programmed; saturation emerges from finite feature

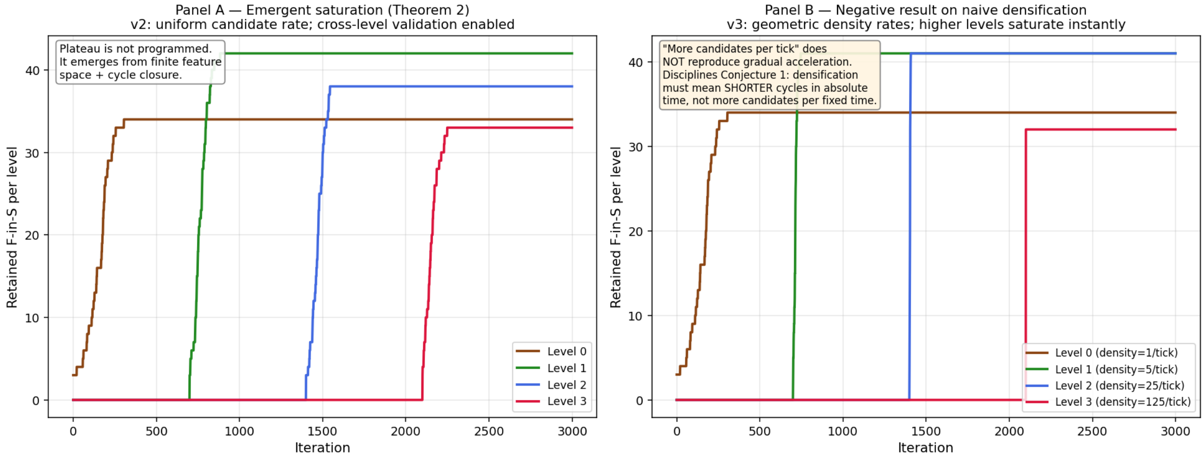

Figure 3: Simulation results. (Left, v2) Saturation emerges from minimal feature-coverage dynamics under cycle closure, consistent with Theorem 2. (Right, v3) The negative result on naive densification: cycle-time variation alone does not reproduce the empirical acceleration pattern; the result disciplines Conjecture 1.

space combined with the closure condition. This illustrates Theorem 2 in a minimal model. **Panel B** (right): negative result on naive densification (v3). Geometric candidate-generation rates per level (1, 5, 25, 125 candidates per tick). Higher-density levels saturate almost instantly rather than showing extended exponential phase with progressively faster ring closure. This rules out "more candidates per fixed time" as the operational reading of densification and disciplines Conjecture 1 to mean *shorter cycles in absolute time*. The simulations illustrate but do not validate Conjecture 1; cross-substrate empirical validation requires operational specification of τ_S (Open Problem 6.6).

5. Empirically observable effects

The framework's structural claims — that functional structures accumulate toward a saturating plateau on each substrate — have an observable, calibratable counterpart in simple physical aggregation, even though aggregation is not itself mutual validation. This section reports what can and cannot be drawn from such observations, and is deliberate about the boundary.

Aggregation as an observable saturation form. Diffusion-limited aggregation (DLA) and electrodeposition produce growing structures whose fractal dimension D_f rises from an early compact value toward a characteristic plateau (≈1.67 in the diffusion-limited regime for Zn/Cu electrodeposition). This rise-then-plateau is exactly the within-layer shape the framework posits for cumulative complexity on a single substrate. As a proof of concept, we fitted the framework's within-layer logistic form,

$$N(t) = N_0 + \frac{N_{\max} - N_0}{1 + e^{-r(t-t_0)}},$$

to a synthetic D_f trajectory shaped to the published qualitative behaviour (rise from ≈1.0 to a plateau near 1.67, with measurement noise, over an operando window of ≈25 min). The form captures the shape closely (mean $R^2 \approx 0.996 \pm 0.002$ over 200 noise realisations; Figure B(a)).

The fit recovers the injected plateau (1.67 ± 0.01) — a consistency check that the form can extract $N_{\max}$ when the underlying shape is logistic, not an independent measurement — and yields an early characteristic time $\tau = \ln 2/r \approx 1.5$ min.

What this supports, stated as supposition. On the basis of this exercise we may *suppose* that the within-layer saturation form receives preliminary empirical support: a logistic rise to a plateau is the right shape for an aggregating physical system, and the plateau height $N_{\max}$ and early time-constant τ are extractable from such a trajectory. We state this as a supposition, not a result, for two reasons. First, the fit is to a synthetic trajectory of the correct shape, not to raw published data points; calibration against real measurements is the necessary next step, and the R^2 reflects shape-capture, not confirmation. Second — and more fundamentally — it remains open whether equations that describe an already-formed, relaxing physical layer apply to a system in *active validation*; that distinction is central to the framework (see footnote).

The boundary: form, not mechanism. This calibrates the mathematics of the *form* — saturation, $N_{\max}$, τ — and nothing more. DLA is aggregation, not mutual validation: particles stick to a growing cluster, they do not validate one another into a closed set. Two limits follow. The exercise says nothing about the validation *mechanism*, which is the framework's actual content. And it does not support the conjectured acceleration law: the D_f trajectory *relaxes* toward a fixed attractor — its growth rate peaks once and decays to zero at the plateau (Figure B(b)) — rather than exhibiting the cross-substrate acceleration that the conjectured $\lambda \propto 1/\tau$ law (developed in the methodological companion (*Validated Retrodiction*)) would describe. Aggregation thus supports the saturation form while standing, if anything, against any naive reading that acceleration is universal. This is consistent with the boundary: DLA calibrates how a layer *settles*, not how validation *drives*.

Thermodynamics as a demonstration substrate. Non-equilibrium thermodynamics (dissipative structures, §7) is the other empirically observable layer we draw on. It demonstrates that ordered structure can be maintained far from equilibrium by energy flow — an observable, well-measured effect. But it occupies the same epistemic position as DLA here: it describes formed structure, not the act of validation.[1]

6. Discussion

The Introduction (§1) traced seventeen predecessors whose contributions made the present framework possible. The simulations (§4) closed one gap and disciplined another in the framework's interpretive content. Section 5 now positions the framework relative to its predecessors thematically, lists three falsifiable predictions, identifies limitations, and notes one current empirical possibility.

[1] We read thermodynamics as the physics of already-formed, validated layers rather than of active validation. The distinction is like assembling a construction set: the first time, without instructions, is a process of trial, fit, and confirmation — the becoming of structure, where validation is active; later assemblies, with the manual and prior experience in hand, follow a settled procedure — the formed layer, where thermodynamic description applies. Energy-flow accounts capture the second regime well and the first only incompletely; whether formed-layer equations extend to actively-validating systems is exactly the open question this section flags.

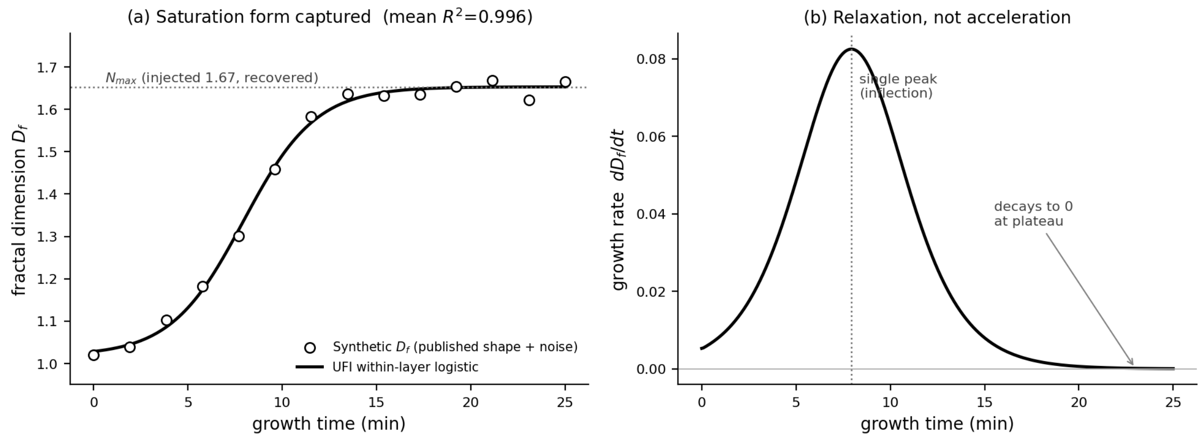

Figure 4: Proof-of-concept calibration of the within-layer saturation form against diffusion-limited aggregation. **(a)** The UFI logistic form fitted to a synthetic fractal-dimension trajectory shaped to the published qualitative behaviour of DLA/electrodeposition (rise from ≈1.0 to a plateau ≈1.67 over ≈25 min); the fit captures the shape (mean $R^2 \approx 0.996$ over 200 noise realisations) and recovers the injected plateau $N_{\max}$ as a consistency check. The R^2 reflects shape-capture against a synthetic trajectory of the correct form, not confirmation against raw data. **(b)** The growth rate peaks once at the inflection and decays to zero at the plateau: the trajectory *relaxes* toward a fixed attractor rather than accelerating. Aggregation thus calibrates the saturation form while standing against any reading that acceleration is universal; it does not support the conjectured $\lambda \propto 1/\tau$ law, and DLA is aggregation, not mutual validation.

6.1 What the present work adds to the chain

The chain of predecessors traced in §1.1 covers seventeen specific contributions across six decades. Rather than mirror each entry with a parallel "adding to" paragraph — which would repeat the chain a third time after §1.1 and §1.2 — we synthesise our position thematically. The present work makes specific moves relative to four distinct groups of predecessors.

To the mechanism predecessors (Kauffman 1971, Eigen 1971, Hordijk-Steel 2004 and subsequent, Maturana-Varela 1972, Rosen 1991, Letelier et al. 2003). We adopt the cycle-closure mechanism — autocatalytic-set closure, RAF-style detection, hypercycle dynamics, autopoietic self-production, (M,R)-systems — as the operational criterion for substrate retention. We do not invent the mechanism; we extend its application from chemistry (Kauffman, Eigen, Hordijk-Steel original domain) and biology (Maturana-Varela, Rosen) to four substrate rings from cosmochemistry to cognition. We add the temporal dimension (carrier density ρ_S = $1/\tau$_S, with the cycle remaining structurally the same across substrates and differing only in cycle duration) that the static RAF formalism does not include. The mathematical core §1.8 catalogues which RAF results actually carry over to our pairwise relation $\leftrightarrow$^{FS} and which are conceptual analogy; no RAF result is formally inherited in the strict sense (math core §1.8). We give explicit credit to Maturana-Varela for the *constitutive* reading of closure — the world *is* self-validating, not merely contains self-validating subsystems — while taking the formal apparatus from Kauffman-Eigen-Hordijk-Steel.

To the universal-Darwinism predecessors and their critics (Lem 1964, Campbell 1974, Dawkins 1976, Dennett 1995, Maynard Smith-Szathmáry 1995, Hodgson-

Knudsen 2010, Campbell J.O. 2016; with critics: Godfrey-Smith 2009, Calcott-Sterelny 2011, niche construction theory of Odling-Smee-Laland-Feldman 2003). We reformulate Σ from unary to binary, removing the ontological privilege of "environment" as something external to the realised structure. This move aligns with niche construction theory's claim that environment is not a separable external given — NCT anticipates the binary at the biological level; we generalise to the structural level across substrates. On the Godfrey-Smith critique we take the position that our framework does not claim selection in his paradigm-Darwinian sense applies uniformly across substrates; we claim that *cycle closure* applies (a structurally weaker condition not presupposing Godfrey-Smith's heritability/variance/fitness-difference cluster). We adopt Hodgson-Knudsen's ontological-vs-heuristic distinction: the cycle structure V-Σ-R is ontological (same across substrates); the operational specification of each step per substrate is heuristic. We do not engage Calcott-Sterelny's critique of MET in detail beyond noting their volume documents disagreement with the original Maynard Smith-Szathmáry formulation.

To the empirical-extension predecessors (Hazen 2008/2024, Krivovichev et al. 2022, Cronin/Walker assembly theory, Kurzweil 1999/2005, Sharov 2006). Hazen extended cumulative-complexity tracking to mineralogy; Krivovichev provided information-theoretic operationalisation; Cronin and Walker developed assembly theory as a substrate-spanning measure; Kurzweil identified the empirical pattern of substrate-transition acceleration informally. We integrate these into a single four-ring empirical continuum spanning 13.8 Gyr (§3, Figure 1), which to our knowledge has not been previously assembled. Our contribution relative to Kurzweil is *formalisation*: where Kurzweil's argument is journalistic and his quantitative predictions have a mixed record, we provide RAF-style closure as mechanism and a named assumption (A3) under which acceleration follows, with explicit acknowledgement that the central conjecture may not survive independent operational specification of τ_S (math core §4, §6.6). Sharov's genome-size series we do not adopt as our biological measure — we measure the biological ring by its protein-fold repertoire (§3.1) — and his backward extrapolation about the age of life we do not endorse.

To the information-and-hierarchy predecessors (Wheeler 1990, Bekenstein 1973, Susskind 1995, 't Hooft 1993, Maldacena 1997, Hawking 2005; Hledík-Barton-Tkačik 2022; Salthe 1985). We accept microphysical information conservation via unitarity and the holographic principle as background. Our structural-information-preservation claim (§2.7) concerns the relational structure of retained F-in-S, not the microphysical states; the two are complementary rather than competing. The connection to Hledík-Barton-Tkačik's formal information-theoretic treatment of biological selection is an open direction — their per-substrate apparatus can in principle be generalised to other substrates with substrate-specific accumulation rates determined by τ_S. From Salthe we adopt the substantive (not merely descriptive) reading of hierarchy and add the specific mechanism (densification through cross-level validation density accumulation) for the emergence of new hierarchical levels.

Friston's free energy principle sits adjacent to all four groups. We do not attempt to integrate FEP with the present framework here; we note that FEP could in principle operate as the within-substrate retention mechanism (the way Σ achieves closure on a given substrate) while the present framework operates at the inter-substrate level. The relationship is left open as a direction for subsequent investigation.

The chain is long, and a reader could reasonably ask whether we have engaged each link in full. The honest answer is that we have not; this is a working preprint and full engagement with seventeen predecessors would require a monograph. What we have done is identify, for each predecessor, what we adopt and what we add. The boundaries are explicit where the predecessor's contribution is mechanism-level (Kauffman, Eigen, Hordijk-Steel), ontological (Maturana-Varela, Rosen), empirical (Hazen, Kurzweil), critical (Godfrey-Smith, Calcott-Sterelny), or methodological (Hledík et al., Salthe). Future iterations will deepen the engagement with critics in particular.

6.2 Where we differ from each predecessor

We can summarise the position of the present work relative to the chain by what we propose that is *not* in the predecessors, with explicit honesty about what *is* already in them:

1. **The empirical continuum across 13.8 Gyr in a single plot** — four substrate rings, each rising to its blueprint floor and then retained on a plateau, with the growth-phase doubling time shortening by orders of magnitude across the sequence (§3, Figure 1). *No previous work, to our knowledge, has assembled this.*

2. **Binary mutual validation as a universal selection operator across substrate classes.** *The mechanism itself — mutual catalysis through cycle closure — is Kauffman (1971), Eigen (1971), and formalised by Hordijk and Steel (2004). The first extension beyond chemistry is Koppl et al. (2022) to economics. Our contribution is the systematic extension to all four substrate rings from cosmochemistry to cognition, not the invention of the mechanism.*

3. **Carrier density ρ_S and the proposed mechanism of acceleration** — substrate-specific validation cycle time as the proposed source of observed acceleration across substrate rings, with the cycle itself unchanged (Conjecture 1, mathematical core). *This is, to our knowledge, novel relative to the RAF literature, which is largely static (does the set close?) rather than temporal (how fast does the closed set grow?). We frame it as conjecture, not theorem; the reasons and its status are set out in §7(c) and the mathematical core.*

4. **Densification as the direction of substrate transitions** — substrate transitions are densifications of the same fundamental function onto more compact carriers, not appearances of new functions (§2.5). *Partially anticipated in Maynard Smith and Szathmáry's discussion of information storage and transmission changes, but not formalised through carrier density.*

5. **Elimination as validation** — the operator Σ does not distinguish positive from negative results; elimination is itself a Σ event and contributes to the cumulative structure (§2.6). *Implicit in standard selection theory but here made explicit as a structural rather than merely negative outcome.*

6. **Structural information preservation** — the relational history of the world is encoded in its current structure, complementary to microphysical unitarity (§2.7). *Novel formulation, though related to information conservation in physics (Wheeler, Bekenstein) and to information accumulation in evolution (Hledík et al. 2022).*

7. **Nested rings with primitive root contingency** — the world is structured as a sequence of closed self-validation rings, with the root ring being not optimal but the first across the threshold of minimal coherence (§2.3, §2.4). *Partially anticipated in Maynard Smith and Szathmáry's hierarchical levels; the primitive-root contingency reading is, to our knowledge, novel.*

The clearest novelty of the present work lies in items 1, 3, 4, 6, and 7. The mechanism of item 2 is inherited rather than invented; our contribution to item 2 is the empirical demonstration of its universality.

6.3 Falsifiable predictions

The framework yields three predictions of differing logical and empirical character. We are explicit about each.

Prediction 1 (gradualism of substrate transitions, retrospective). Substrate transitions are gradual rather than instantaneous: cross-level validation density between substrate class N and N+1 should be observable as continuous during the transition period, not zero before and high after. *Empirical status:* this is partly retrospective; Hazen et al. (2024) provides initial supporting data by showing that abiotic minerals can reach assembly indices previously thought to be biotic-only. *Operational specification of "transition duration":* the period during which cross-level validation density between substrate N and N+1 is non-zero and below the saturation threshold of the new substrate N+1 — *not* the period during which observers perceive a paradigm-level discontinuity. Under this operational specification, the cognitive-to-scientific substrate transition began with externalised symbolic cognition well before Galileo (writing systems, ~5000 years ago) rather than with the early modern scientific revolution. *Specific risky form:* if a substrate transition is observed to complete (i.e., the new substrate reaches its early-exponential phase) in less than approximately 1% of the predecessor substrate's growth-phase timescale, measured against the operational specification above, P1 is falsified. Empirical test: comparative complexity analysis of borderline cases against this operational definition.

Prediction 2 (deferred to future work). A quantitative prediction relating the rate of each ring's growth to the timescale of one validation cycle on that substrate (the empirical face of Conjecture 1 of the mathematical core) requires an independent operational specification of the validation cycle time τ_S per substrate (Open Problem 6.6 of the math core), which is not yet available. We therefore defer this prediction and its test to a dedicated future paper rather than state it here on an unspecified timescale.

Prediction 3 (maturity as a foundation, near-term testable). We observe that, in the record to date, a new substrate class appears to require its predecessor to be sufficiently formed to serve as a stable foundation — much as an organisation must be substantially built before it yields a genuinely new kind of product, while it functions and produces on its own level well before that. *Testable form, stated as an expectation rather than a law:* we do not expect a sustained new substrate class to arise on a predecessor still in its early growth phase. This is checkable by examining transitions where the predecessor was far from formed. *Empirical test:* examine candidate substrate transitions where the predecessor is still in active exponential phase; the present moment provides one such case (see §6.5), since the scientific substrate has not visibly plateaued. We keep this distinct from the doubling-time ordering of §3.2: that concerns tempo, this concerns foundation — they are not the same claim.

A note on what is not promised by P1-P3. None of these predictions individually validates the framework's strong ontological claim that the world *is* self-validating. They test specific structural commitments. Even all three holding empirically would not entail the ontological claim; they would support it. The predictions are scoped to the structural content, not to the ontology.

6.4 Limitations

The empirical anchor of this paper rests on four substrate rings. Data quality varies: cosmochemistry estimates are necessarily approximate; mineralogy is well-characterised in formal IMA species but recent revisions continue; the protein-fold repertoire is well-characterised in modern structural databases but its early assembly history relies on inference; cognition data are well-documented for recent periods but rely on extrapolations for earlier eras.

The mathematical core (§4 and companion document) presents Theorems 1-2 and Conjecture 1 as proof sketches rather than rigorous derivations. Each sketch carries explicitly stated caveats: independence assumptions (A1, A2), mean-field approximations (A1, A0), and the assumption that the per-cycle rate constant λ is substrate-invariant (A3, central to Conjecture 1). These sketches require independent verification before being treated as established results. We have flagged the sketches as sketches and have identified eleven open formal problems for subsequent work, of which two (6.1 formal structural information preservation; 6.6 operational specification of τ_S per substrate) are flagged as **blocking** for central framework claims.

The number of substrate rings (four) is chosen to span available high-quality time-series data; the row count is not principled in the sense that a deeper theory would demand. Reviewers may reasonably ask why mineralogy and cosmochemistry are separate rings, or why the biological ring is measured by its protein-fold repertoire rather than by genome size or taxonomic diversity. Our answer is that each ring is a distinct closed validation cycle with its own food set, and that within the biological ring the fold repertoire is the retained, mutually validated structural primitive — the apparatus on which later genome growth and diversification proceed — so it, not those later quantities, marks the closure of the ring.

The framework does not address: what selected the primitive root ring from possible configurations (which we hold to be contingent and not subject to further principled selection); whether the asymptotic limit of densification is reachable in finite time or in principle; why the cycle

V-Σ-R operates at all (which we treat as a primitive observation about the world rather than as derivable from more basic axioms).

6.5 An open empirical question: the contemporary moment

The framework's predictions can be tested against any candidate substrate transition. The present moment offers one such test: foundation models in artificial intelligence have shown, since approximately 2017-2020, accelerating capability growth on timescales substantially shorter than either science or technology in their conventional measures. Under the densification reading, this is the empirical signature the framework would expect of a substrate transition, with a still-shorter validation cycle.

We do not claim foundation models constitute a new substrate. The relevant question is whether, in the coming years, the framework's predictions (§6.3) are borne out for this candidate: gradualism of the transition (P1) and the saturation precondition (P3, which would require the scientific substrate's plateau onset to precede AI substrate sustainability). These are empirical questions; we offer them as tests of the framework, not as conclusions about AI.

7. Predecessors: one phenomenon observed across many substrates

The mutual-validation reading is not new in any single field. Across chemistry, biology, ecology, economics, physics, and cognition, investigators have repeatedly found the same structural signature — elements of a system that confirm, sustain, or constrain one another into a self-maintaining whole — and named it in the vocabulary of their own discipline. The United Field Initiative does not claim to have discovered this phenomenon. Its contribution is to read these scattered observations as instances of one cross-substrate process. The attribution of each result belongs to the field that produced it; only the cross-substrate reading is ours. Where §1.1 traces the genealogy of the idea that a single process spans substrates, this section maps where the specific phenomenon — mutual validation — has been observed, field by field.

Field	Key author(s)	What they observed (their terms)
Chemistry / origin of life	Kauffman (1971); Eigen & Schuster (1977)	Autocatalytic sets and hypercycles
Biochemistry / metabolism	Sousa et al. (2015); Xavier et al. (2020)	RAF sets in real metabolism
Ecology	Ulanowicz (1980, 1997)	Ascendency ($A = \text{TST} \times \text{AMI}$)
Economics / innovation	Koppl et al. (2022)	The economy as an autocatalytic set
Thermodynamics / self-organisation	Prigogine (1977)	Dissipative structures

Field	Key author(s)	What they observed (their terms)
Cosmic evolution	Chaisson (2001)	Free Energy Rate Density across 14 Gyr
General evolution / thermodynamics	Swenson (1989)	Maximum entropy production as a law of evolution
Autopoiesis	Maturana & Varela (1972)	Operational closure
Evolutionary epistemology	Campbell (1974)	Blind variation and selective retention
Quantum	Zurek (2003, 2009)	Quantum Darwinism

Each entry below states what the author established in that field's own terms, and how the framework reads it. The reading is an interpretation, not a claim on the original result.

Kauffman; Eigen & Schuster. Kauffman introduced autocatalytic sets — collections of molecules that mutually catalyse one another's formation from a food set — and Eigen and Schuster independently introduced the hypercycle, a cyclic coupling of self-replicating units. We read this as the prototype of ring closure: a set of functional structures each of which validates the others' production.

Sousa et al.; Xavier et al. Applying RAF (reflexively autocatalytic, food-generated) set theory, these groups detected such sets within the real metabolic networks of *E. coli* and of ancient autotrophs. We read this as evidence that the closure is realised in living matter, not merely formally possible.

Ulanowicz. Ascendency measures an ecosystem's organisation as $A = \text{TST} \times \text{AMI}$, the total system throughput scaled by the average mutual information of its flows, and ecosystems tend to develop toward higher ascendency. We read average mutual information as a measure of how strongly the flows constrain — and in that sense validate — one another.

Koppl et al. Building on RAF theory, this work treats the economy as an autocatalytic set in which goods and processes enable the production of further goods. We read it as the first extension of validation closure beyond chemistry, onto the economic substrate.

Prigogine. Dissipative-structure theory describes how an ordered configuration is maintained far from equilibrium by continuous energy flow. We read it as physics of the already-formed layer — an account of maintenance — and do not reduce validation to energy dissipation.

Chaisson. Cosmic-evolution work traced rising complexity across 14 Gyr with a single quantitative metric, Free Energy Rate Density, decades before the present framework. The cross-substrate continuity is therefore not our discovery; we differ in the measure used — validation rather than energy flow.

Swenson. The law of maximum entropy production proposed a single cross-substrate principle of evolution grounded in thermodynamics. The cross-substrate ambition is shared; but where that programme identified complexity with entropy production — an identification that kept it entangled for decades — the present framework measures validation, and so does not equate complexity with entropy at all.

Maturana & Varela. Autopoiesis describes a living system as one that continuously produces the network of processes that produces it — operational closure. We read this as self-validation stated from the inside: closure expressed as self-production.

Campbell. Blind variation and selective retention identifies retention as where knowledge is kept across iterations. The framework offers an answer to the selector question this leaves open — structures validate one another rather than being selected by an external given — while relocating, not eliminating, the question to what closes the mutual-validation cycle.

Zurek. Quantum Darwinism holds that redundant, independent records of a system in its environment are what make its state objectively classical. We read this only as an instance, in our vocabulary, of redundant independent confirmation; we make no contribution to quantum measurement theory and take no position on its interpretive disputes.

What the convergence does and does not show. These authors observed the same effect in different media. That convergence is consistent with the wide prevalence of a single mutual-validation phenomenon; it does not, on its own, prove it. Each result stands within its own field's evidence. The cross-substrate reading adds an interpretation, not new data.

Positioning of the United Field Initiative. The framework is not the pioneer of cross-substrate complexity; Chaisson, Swenson, and others traced it decades ago. Where the present work differs:

(a) **Measure.** Complexity is gauged by validation — how thoroughly a structure is mutually confirmed — rather than by energy-flow rate, which describes the maintenance of already-formed layers rather than the act by which a layer forms.

(b) **Mechanism.** Mutual validation addresses the selector question of evolutionary epistemology: structures validate one another, with the selector internal to the population rather than externally given.

(c) **A conjectured rate law.** The framework conjectures a relation between a ring's rate of closure and its validation cycle time ($\lambda \propto 1/\tau$) as a proposed mechanism for the qualitative acceleration of saturation across substrates. This is a conjecture, not an established law: as noted in the mathematical core and the methodological companion (*Validated Retrodiction*), it risks reducing to a parameter-fitting identity until τ can be specified independently per substrate.

(d) **Character of the measure.** Validation depth provides a qualitative ordering commensurable across substrates whose numerical units (elements, folds, validated concepts) cannot share one axis. Its quantitative operationalisation remains open; the difficulties of a count-based coordinate are discussed in the methodological companion (*Validated Retrodiction*).

Taken together, (a)–(d) are what the framework claims to add to a phenomenon it did not discover: a measure, a mechanism, and a conjectured law — offered, like the rest of this work, as program rather than proof.

8. Conclusion

The cumulative-complexity pattern observed in cognition and biology extends back through mineralogy and cosmochemistry to the first moments of nucleosynthesis after the Big Bang. The pattern is consistent over 13.8 Gyr and across four substrate rings, with each ring closing faster and more recently than the one before — the growth-phase doubling time shortening by orders of magnitude across the sequence. We propose that this pattern is the visible record of a single ontological process: the world is a self-validating system of models, in which functional structures realised in substrates mutually validate one another through a binary selection operator, and the cumulative result of all such validations constitutes the current configuration of the universe.

This is a strong claim and we make it deliberately. The framework predicts that substrate transitions occur through cross-level validation density accumulation, that densification is our proposed account of the observed acceleration of ring saturation, that elimination is itself a validation event with structural consequences, and that information about all past validation events is preserved in the current relational structure of the world. Falsifiable predictions are stated explicitly (§6.3), with a further quantitative prediction (P2) deferred to future work pending an operational specification of cycle time; two open problems are flagged as blocking for central claims (§6.4, math core §6.1 and §6.6).

What the framework does not provide, and does not pretend to provide: a derivation of the laws of physics from the self-validation principle (the relationship to Noether's theorem and to standard physicalism is one of compatibility, not replacement); a unique partitioning of the world into substrate rings (four is a working count justified by data availability, not by deeper theory); a calibration of cycle times τ_S independent of the observed saturation pattern (this is the central blocking problem for Conjecture 1's empirical content). We expect future iterations of the framework to address these limits. We expect parts of the present framework to be revised or rejected. The framework is the next link in a long chain of investigation; it will be followed by further links.

The contemporary moment offers one near-term test of the framework's empirical content (§6.5). The framework's three predictions can be evaluated against the trajectory of foundation models in AI over the coming decade, against any future substrate transition that may emerge, and against independent operational measurement of τ_S on each existing substrate. The framework is at risk of falsification on each of these axes. We welcome that risk.

AI assistance

This work was developed in collaboration with Claude (Anthropic, primary writing partner: literature integration, formalisation, drafting, internal review, generation of Figures 1 and 3 via matplotlib) and ChatGPT (OpenAI, independent review of mathematical core, generation of Figure 2). All conceptual commitments and final framings are the author's. Ideas were developed iteratively in dialogue with the AI systems noted above; the author exercised editorial selection throughout and bears full responsibility for the framework's content and any errors. This disclosure follows current expected practice for AI-assisted academic preprints; AI systems are not authors and could not be authors under prevailing standards.

References

Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science* 15, 215-227.

Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D* 7, 2333.

Bornmann, L. and Mutz, R. (2015). Growth rates of modern science. *Journal of the Association for Information Science and Technology* 66, 2215-2222.

Burbidge, E. M., Burbidge, G. R., Fowler, W. A., and Hoyle, F. (1957). Synthesis of the elements in stars. *Reviews of Modern Physics* 29, 547-650.

Calcott, B. and Sterelny, K. (eds.) (2011). *The Major Transitions in Evolution Revisited.* Cambridge: MIT Press.

Campbell, D. T. (1974). Evolutionary epistemology. In P. A. Schilpp (ed.), *The Philosophy of Karl Popper.* La Salle: Open Court.

Campbell, J. O. (2016). Universal Darwinism as a process of Bayesian inference. *Frontiers in Systems Neuroscience* 10, 49.

Cazzolla Gatti, R., Hordijk, W., and Kauffman, S. (2020). Biodiversity is autocatalytic. *Ecological Modelling* 346, 70-76.

Chaisson, E. J. (2001). *Cosmic Evolution: The Rise of Complexity in Nature.* Cambridge, MA: Harvard University Press.

Dawkins, R. (1976). *The Selfish Gene.* Oxford: Oxford University Press.

Dennett, D. C. (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life.* New York: Simon & Schuster.

Denysov, I. (2026a). Foundations for Measuring the Knowability of a World from Within. Zenodo. DOI: 10.5281/zenodo.20277783.

Denysov, I. (2026b). Operational Principia of Knowability. Zenodo. DOI: 10.5281/zenodo.20330208.

Denysov, I. (2026c). Quantitative Limits of Internal Observer Knowledge: A Simulation-Based Framework. Zenodo. DOI: 10.5281/zenodo.20123642.

Eigen, M. (1971). Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58, 465-523.

Eigen, M. and Schuster, P. (1977-1978). The hypercycle. A principle of natural self-organization. *Naturwissenschaften* 64-65 (three-part series).

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience* 11, 127-138.

Godfrey-Smith, P. (2009). *Darwinian Populations and Natural Selection*. Oxford: Oxford University Press.

Hawking, S. W. (2005). Information loss in black holes. *Physical Review D* 72, 084013.

Hazen, R. M., Papineau, D., Bleeker, W., Downs, R. T., Ferry, J. M., McCoy, T. J., Sverjensky, D. A., and Yang, H. (2008). Mineral evolution. *American Mineralogist* 93, 1693-1720.

Hazen, R. M., Burns, P. C., Cleaves, H. J., Downs, R. T., Krivovichev, S. V., and Wong, M. L. (2024). Molecular assembly indices of mineral heteropolyanions: some abiotic molecules are as complex as large biomolecules. *Journal of the Royal Society Interface* 21, 20230632.

Hledík, M., Barton, N., and Tkačik, G. (2022). Accumulation and maintenance of information in evolution. *PNAS* 119, e2123152119.

Hodgson, G. M. and Knudsen, T. (2010). *Darwin's Conjecture: The Search for General Principles of Social and Economic Evolution*. Chicago: University of Chicago Press.

Hordijk, W. and Steel, M. (2004). Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology* 227, 451-461.

Hordijk, W., Hein, J., and Steel, M. (2010). Autocatalytic sets and the origin of life. *Entropy* 12, 1733-1742.

Hordijk, W., Steel, M., and Kauffman, S. (2012). The structure of autocatalytic sets: evolvability, enablement, and emergence. *Acta Biotheoretica* 60, 379-392.

Hordijk, W. and Steel, M. (2018). Autocatalytic networks at the basis of life's origin and organization. *Life* 8, 62.

Hordijk, W. (2019). A history of autocatalytic sets: A tribute to Stuart Kauffman. *Biological Theory* 14, 224-239.

Kauffman, S. A. (1971). Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *Journal of Cybernetics* 1, 71-96.

Kauffman, S. A. (1986). Autocatalytic sets of proteins. *Journal of Theoretical Biology* 119, 1-24.

Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford: Oxford University Press.

Koppl, R., Devereaux, A., Herriot, J., and Kauffman, S. (2022). Toward a general theory of socioeconomic complexity from the bottom up: a unified application of autocatalytic set theory and the theory of innovation cascades. *PNAS Nexus* 1, pgac229.

Krivovichev, S. V., Krivovichev, V. G., Hazen, R. M., Aksenov, S. M., Avdontceva, M. S., Banaru, A. M., Gorelova, L. A., Ismagilova, R. M., Kornyakov, I. V., Kuporev, I. V., Morrison, S. M., Panikorovskii, T. L., and Starova, G. L. (2022). Structural and chemical complexity of minerals: an update. *Mineralogical Magazine* 86, 183-204.

Kurzweil, R. (1999). *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York: Viking.

Kurzweil, R. (2005). *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.

Lem, S. (1964). *Summa Technologiae*. Kraków: Wydawnictwo Literackie.

Letelier, J. C., Marin, G., and Mpodozis, J. (2003). Autopoietic and (M,R) systems. *Journal of Theoretical Biology* 222, 261-272.

Loutchko, D. (2022). Semigroup models for biochemical reaction networks. *Journal of Mathematical Biology* 84, 9.

Luhmann, N. (1986). The autopoiesis of social systems. In F. Geyer and J. van der Zouwen (eds.), *Sociocybernetic Paradoxes*. London: Sage.

Maldacena, J. (1997). The large N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics* 2, 231-252.

Maturana, H. R. and Varela, F. J. (1972). *De Máquinas y Seres Vivos: Autopoiesis — La organización de lo vivo*. Santiago: Editorial Universitaria.

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.

Maynard Smith, J. and Szathmáry, E. (1995). *The Major Transitions in Evolution*. Oxford: Oxford University Press.

Moore, G. E. (1965). Cramming more components onto integrated circuits. *Electronics* 38, 114-117.

Noether, E. (1918). Invariante Variationsprobleme. *Nachr. d. König. Gesellsch. d. Wiss. zu Göttingen, Math.-phys. Klasse* 235-257.

Odling-Smee, F. J., Laland, K. N., and Feldman, M. W. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton: Princeton University Press.

Popper, K. R. (1972). *Objective Knowledge: An Evolutionary Approach*. Oxford: Clarendon Press.

Prigogine, I. (1977). *Time, Structure and Fluctuations*. Nobel Lecture in Chemistry, 8 December 1977. See also Nicolis, G. and Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order through Fluctuations*. New York: Wiley.

Rosen, R. (1991). *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life.* New York: Columbia University Press.

Salthe, S. N. (1985). *Evolving Hierarchical Systems: Their Structure and Representation.* New York: Columbia University Press.

Sepkoski, J. J. (2002). A compendium of fossil marine animal genera. *Bulletins of American Paleontology* 363, 1-560.

Sharma, A., Czégel, D., Lachmann, M., Kempes, C. P., Walker, S. I., and Cronin, L. (2023). Assembly theory explains and quantifies selection and evolution. *Nature* 622, 321-328.

Sharov, A. A. (2006). Genome increase as a clock for the origin and evolution of life. *Biology Direct* 1, 17.

Smolin, L. (1992). Did the universe evolve? *Classical and Quantum Gravity* 9, 173-191.

Smolin, L. (2013). *Time Reborn: From the Crisis in Physics to the Future of the Universe.* Boston: Houghton Mifflin Harcourt.

Sousa, F. L., Hordijk, W., Steel, M., and Martin, W. F. (2015). Autocatalytic sets in *E. coli* metabolism. *Journal of Systems Chemistry* 6, 4. DOI: 10.1186/s13322-015-0009-7.

Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics* 36, 6377-6396.

Szathmáry, E. (2015). Toward major evolutionary transitions theory 2.0. *PNAS* 112, 10104-10111.

't Hooft, G. (1993). Dimensional reduction in quantum gravity. arXiv:gr-qc/9310026.

Swenson, R. (1989). Emergent attractors and the law of maximum entropy production: foundations to a theory of general evolution. *Systems Research* 6(3), 187-197. DOI: 10.1002/sres.3850060302.

Ulanowicz, R. E. (1997). *Ecology, the Ascendent Perspective.* New York: Columbia University Press.

Verhulst, P. F. (1838). Notice sur la loi que la population suit dans son accroissement. *Correspondance Mathématique et Physique* 10, 113-121.

Wheeler, J. A. (1990). Information, physics, quantum: the search for links. In W. H. Zurek (ed.), *Complexity, Entropy, and the Physics of Information.* Redwood City: Addison-Wesley.