# Mathematical Core of the Self-Validating World Framework

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

22 May 2026

**Publication notice.** This is the mathematical companion to Paper 0 Part 3 of the United Field Initiative series. The main manuscript is *The Self-Validating World: An Empirical Continuum of Cumulative Complexity Across Substrate Rings* (Denysov 2026), available in the same Zenodo record. The United Field Initiative 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). The framework also relates to Paper 1, *Quantitative Limits of Internal Observer Knowledge* (DOI: 10.5281/zenodo.20123642), discussed in §5.

---

## Purpose

This document records the mathematical core of the self-validating world framework. It contains definitions, a list of named assumptions (A0-A4), two theorems with proof sketches, one conjecture with explicit empirical content, an examination of which RAF results actually carry over to our pairwise closure relation $\leftrightarrow\hat{}\{\mathrm{FS}\}$, and a statement of relationship to Paper 1 Theorem 1.

The framework's pairwise closure relation $\leftrightarrow\hat{}\{\mathrm{FS}\}$ is **inspired by** but **not equivalent to** canonical RAF closure (see §1.2 and §1.8). Examination of inheritance shows that **no RAF results are formally inherited** in the strict sense of proving statements directly about $\leftrightarrow\hat{}\{\mathrm{FS}\}$-defined structures; RAF results either await operational specification per substrate (R1), are downgraded to conceptual analogy (R2-R4), or are available but unused (R5).

Theorems 1 and 2 are mean-field phenomenological derivations under explicitly named assumptions (A0: short-cycle dominance; A1: mean-field independence/uniformity/constant generation rate, with three sub-assumptions A1a/A1b/A1c; A2: independence of input coverage and novelty exhaustion). Conjecture 1 (formerly stated as Theorem 3) proposes a mechanism for cross-substrate acceleration via Assumption A3 (substrate-invariant per-cycle rate). **Conjecture 1 is currently descriptive, not predictive**: its empirical content is uncalibrated and may not survive principled operational specification of $\tau\_S$ per substrate (Open Problem 6.6, blocking).

Two open problems are flagged as **blocking** for central framework claims: 6.1 (formal proof of structural information preservation, blocks §5 complementarity with Paper 1) and 6.6 (operational specification of $\tau$_S, blocks Conjecture 1 empirical content).

The named-assumption discipline (§1.9) follows the pattern of Paper 0 Part 2 (P0-P5). Each derivation invokes specific assumptions explicitly. The status of each result (theorem, conjecture, inheritance pending, conceptual analogy) is explicitly marked.

The mathematical core is intended to support but not constrain the empirical and ontological content of the main manuscript.

---

# Section 1. Definitions

## 1.1 Primitives

We work with the following primitives:

- **Functional structure** F: an abstract specification of inputs, outputs, and an internal transformation. Formally, F = (I_F, O_F, t_F) where I_F is a set of input indices, O_F is a set of output indices, and t_F is a transformation (not used in the present derivations but included for completeness).
- **Substrate** S: a domain of possible realisations. Each F can be realised in one or more S; a realisation is denoted (F, S) or F-in-S.
- **Time** t: a discrete or continuous parameter ordering events. In the present derivations time is treated as continuous on the substrate's own scale. Time itself is a derived concept: it is the order of $\Sigma$ events, not an independent parameter (see §1.7).
- **Food set** food_S: the structured carriers from which F-in-S realisations are built. For substrate S = root, food_root = protomatter (see §1.6).
- **Realised configuration** R(t): a finite multiset of F-in-S realisations at time t. Cardinality N(t) := |R(t)|.

The primitives are not all independent in the deep sense. F and S are inseparable in the empirical world: every realised configuration is F-in-S, never F alone or S alone. The framework treats them as logically distinguishable for purposes of formal derivation while acknowledging their empirical inseparability.

## 1.2 Mutual validation: the pairwise F-in-S closure relation $\leftrightarrow\hat{}\{FS\}$

Two realisations $F_1$-in-$S_1$ and $F_2$-in-$S_2$ are said to be **mutually validating** if there exists a closed functional cycle in which each realisation's outputs supply (directly or through intermediaries within the realised configuration) the inputs of the other. We denote this binary relation by $\leftrightarrow\hat{}\{FS\}$ ("pairwise F-in-S closure") to distinguish it from canonical RAF closure.

**Why the distinction matters.** This binary closure relation is **inspired by** autocatalytic set theory (Kauffman 1971, 1986) and hypercycle theory (Eigen 1971; Eigen and Schuster 1977-1978), and informed by their rigorous formalisation as RAF (Reflexively Autocatalytic and Food-generated) set theory (Hordijk and Steel 2004). However, $\leftrightarrow\hat{}\{FS\}$ is **not** the same as RAF closure. Three differences should be noted:

1. **Pairwise vs. global.** RAF closure is a property of a *set* as a whole (every reaction catalysed by some element of the set; every reactant producible within the set from food). Our $\leftrightarrow\hat{}\{FS\}$ is a *binary relation* between pairs of F-in-S. From pairwise $\leftrightarrow\hat{}\{FS\}$ one does not automatically get RAF-closure of the underlying set — isolated mutually-validating pairs may exist that do not form a closed network in the strict RAF sense.

2. **Food set treatment.** RAF closure makes the food set explicit in its definition (F-generated condition). Our $\leftrightarrow\hat{}\{FS\}$ does not mention food directly; the food set is treated separately in §1.6.

3. **F-generated condition.** RAF requires that all reactants in the set be producible from food via reactions within the set. Our $\leftrightarrow\hat{}\{FS\}$ implies this for the pair (each F's outputs supply the other's inputs through cycles within R) but does not impose it globally on R.

Graph-theoretically, $\leftrightarrow\hat{}\{FS\}$ is closer to "$F_1$ and $F_2$ belong to the same strongly connected component in the directed graph of input/output dependencies in R" than to canonical RAF closure. The relation is in the same conceptual family but technically distinct.

**Consequence for inheritance.** Because $\leftrightarrow\hat{}\{FS\} \neq$ RAF closure, theorems proved for RAF sets do not automatically apply to $\leftrightarrow\hat{}\{FS\}$-defined subsets. Each invocation of an inherited RAF result downstream must check applicability. Where applicability is not established, the result is downgraded to "conceptual analogy" rather than "formally inherited" (see §1.8).

Formally, let R denote a realised configuration. Define the relation

$$F_1 \leftrightarrow_R^{\mathrm{FS}} F_2 \quad \Longleftrightarrow \quad \exists\, \text{path } P \text{ in } R \text{ such that } O_{F_1} \supseteq I_{P,1},\ O_{P,k} \supseteq I_{F_2},\ O_{F_2} \supseteq I_{P,k+1},\ \ldots,\ O_{P,m} \supseteq I_{F_1}$$

where P = (P_1, . . . , P_m) is a sequence of realisations in R forming a closed cycle that begins and ends at F_1 through F_2.

The case of direct mutual validation (m = 0, i.e., no intermediaries) is a special case in which O_{F_1} ⊇ I_{F_2} and O_{F_2} ⊇ I_{F_1}. This is the same condition that holds for two species in a direct autocatalytic pair.

**Cycle length restriction (Assumption A0).** For the derivations in §2-4, we assume that closure cycles have bounded length: |P| ≤ L_max where L_max is a substrate-specific constant. Without this restriction, the number of potential closure cycles through a candidate F scales as N^(L-1) where L is the maximum allowed cycle length, producing super-linear growth rather than the linear-in-N scaling assumed in Theorem 1. In practice, real validation cycles tend to be

short (direct mutual catalysis predominates in chemistry; small loops in ecology; brief peer-review chains in science), but the assumption should be made explicit.

**Remark on biomimetic justification.** The operational meaning of "validation" at every empirical substrate level is the closure of a cycle: chemical bond formation closes an energetic cycle (Kauffman, Eigen); predator-prey relations close an ecological cycle; theory-and-experiment close an epistemic cycle; producer-consumer relations close a technological cycle. The autopoiesis tradition (Maturana and Varela 1972, 1980; Rosen 1991 (M,R)-systems; Letelier et al. 2003 on self-referential closure) articulates this closure property as constitutive of living and life-like systems. The present framework treats cycle closure as a universal mechanism across all substrate classes.

### 1.3 Selection operator $\Sigma$

The selection operator is

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

where $\mathcal{P}(\mathcal{S})$ denotes the power set of possible substrate configurations.

$\Sigma$ takes a candidate population and an existing retained configuration, and returns those candidates that are mutually validated by elements of the existing configuration:

$$\Sigma(C, R) = \{F \in C \ : \ \exists\, F' \in R \text{ such that } F \leftrightarrow_R^{\text{FS}} F'\}$$

The retention operation updates R:

$$R(t + \Delta t) = R(t) \cup \Sigma(C(t), R(t))$$

where C(t) is the set of candidate F's generated at time t (the V-component of the cycle).

### 1.4 Carrier density

Each substrate has a characteristic **carrier density** $\rho$_S, defined operationally as the inverse of the time required to complete one $\Sigma$-evaluation on that substrate:

$$\rho_S := 1/\tau_S$$

where $\tau$_S is the validation cycle time on substrate S, measured in absolute time. Densification across substrate transitions is the empirical observation that $\rho$_{S_{N+1}} > $\rho$_{S_N} when S_{N+1} is the substrate of a higher ring built on S_N.

### 1.5 Feature space

Each substrate has a finite feature space F_max(S), defined as the set of distinct input/output indices accessible to F's realised in S. The cardinality |F_max(S)| bounds the number of distinct F's that can be realised in S without redundancy.

For combinatorial estimation, the number of F's realisable in S with bounded input/output sizes is approximately polynomial in |F_max(S)| with degree determined by the maximum cardinalities of I_F and O_F.

### 1.6 Food set

In the RAF tradition (Hordijk and Steel 2004), an RAF set requires a food set: building blocks not produced within the set itself but available from outside. We adopt this construct with a specific interpretation appropriate to the self-validating world ontology.

**Definition (food set for a substrate ring).** The food set for substrate ring N+1 is the set of structured carriers (matter, information embedded in matter) produced by the validated F-in-S realisations of substrate ring N. Formally, food_{N+1} = the realised F-in-S of ring N that supply inputs to candidate F's of ring N+1.

The food set is therefore **internal to the world but external to the specific substrate ring**. Each ring above the root has its predecessor as food set.

**Definition (protomatter and root food set).** Protomatter denotes matter in its pre-validation state — material that has the minimal properties (interaction capability, distinguishability) required to participate in self-validation but that has not yet been structured by any model. The root ring of fundamental physics has protomatter as its food set. Protomatter is part of the world (not external to it) but exists prior to any closed model.

**Remark.** The distinction between food set as quantity of structured carrier and the models that structured it (quality of structure) is essential. A substrate ring inherits a finite quantity of matter from its predecessor; the structural form of that matter is what makes it suitable as the substrate's food. Quantity and quality are independent in principle and may evolve on different timescales.

### 1.7 Time as ordering of Σ events

Time is treated in the formalism as a parameter, but the framework's ontological commitment is stronger: time is the ordering of Σ events, not an independent dimension along which events unfold. Before any Σ event, no temporal ordering exists; the root ring's closure is the first Σ event and is logically prior to time in the formal sense.

This commitment is consistent with relational interpretations of time in physics (Leibniz, Mach, Barbour) and with the absence of background time in canonical quantum gravity. The framework does not require this commitment for the derivations in Sections 2-4; it is recorded here for ontological completeness.

### 1.8 Results from RAF theory: what carries over to $\leftrightarrow\hat{}\{FS\}$

The mathematical apparatus of RAF set theory (Hordijk and Steel 2004; Hordijk et al. 2010, 2014, 2015, 2022; Steel et al. 2019, 2020) provides several results developed for canonical RAF closure. Because the present framework uses $\leftrightarrow\hat{}\{FS\}$ (a pairwise relation, not RAF closure proper — see §1.2), each RAF result must be examined for whether it actually applies to $\leftrightarrow\hat{}\{FS\}$-defined structures, or only motivates analogous reasoning. We classify each result accordingly.

**R1 (closure detection — inheritance pending).** Hordijk and Steel (2004) gave a polynomial-time algorithm for detecting maximal RAF subsets within a catalytic reaction system. The algorithm requires a catalysis function explicitly mapping reactions to catalysts. For chemistry this is well-defined. For substrate domains such as science (theories validating theories) or technology (products validating products), what plays the role of the catalysis function is not formally specified. R1 is therefore **inheritance pending operational specification of catalysis per substrate**. For the present framework we adopt R1 as a procedure-template rather than as a directly applicable algorithm.

**R2 (emergence threshold — conceptual analogy).** Hordijk, Hein, and Steel (2010) showed that the level of catalysis required for RAF emergence in random CRS is approximately 1.5-2 catalysed reactions per molecule, far below the exponential threshold conjectured by Kauffman. This result was proved for random CRS under specific assumptions about catalysis probability distributions. The "each substrate has its own emergence threshold" generalisation we sometimes invoke is not derived from R2; it merely uses R2 as a conceptual analogy that thresholds exist and are tractable. R2 should be cited as **conceptual analogy**, not as an inherited theorem applied to non-chemical substrates.

**R3 (sub-RAF hierarchy — conceptual analogy).** Hordijk, Steel, and Kauffman (2012) showed that any RAF set can contain sub-RAFs as proper subsets, with definite hierarchy structure (irreducible RAF core; sub-RAFs ordered by containment). This is a theorem about RAF sets. $\leftrightarrow\hat{}\{FS\}$-defined subsets are not automatically RAF sets (§1.2). Therefore **R3 does not formally apply to $\leftrightarrow\hat{}\{FS\}$**; the nested-rings structure we invoke is motivated by R3 as a template, but a hierarchy theorem for $\leftrightarrow\hat{}\{FS\}$-defined subsets has not been proved here. R3 is downgraded to **conceptual analogy**, with the formal hierarchy of $\leftrightarrow\hat{}\{FS\}$-subsets listed as an open problem (see §6).

**R4 (boundaries and higher-level RAFs — conceptual analogy).** Hordijk and Steel (2018) showed that boundaries can be incorporated into RAF formalism, giving a mechanism for higher-level RAFs (e.g., protocells as RAFs-of-RAFs). We apply this to substrate-class transitions, treating each substrate ring as bounded by its food-set structure of the previous ring. The application is **conceptual analogy**, not directly inherited: substrate-class transitions are not literally chemical-mass boundaries, and the framework's "ring of rings" structure motivates but does not formally inherit from R4.

**R5 (algebraic semigroup — available, not currently used).** Loutchko (2022) gave an algebraic characterisation of self-generating CRS using semigroup models with partial order and idempotent addition. This provides additional formal tools that we do not currently require but

that are available for subsequent development.

**Concept-level precursors (acknowledged, not inherited).** Beyond the RAF formalism, the closure concept has been developed in several traditions that are conceptually relevant:

- **Maturana and Varela (1972, 1980), autopoiesis.** Closure as constitutive of living systems; operational closure with structural coupling to environment.
- **Rosen (1991), (M,R)-systems.** Metabolism-repair as self-referential closure. Category-theoretic in spirit but not formalised as inheritable theorems.
- **Letelier, Marin, Mpodozis (2003), self-referential systems.** Formal models of autopoiesis-related closures.

These traditions inform the ontological reading of the framework but their formal apparatuses are either insufficiently developed for direct inheritance, or developed in formalisms we do not currently use. We acknowledge them as part of the chain of investigation.

**Summary of inheritance status.** Of five RAF results, **none are formally inherited** in the sense that they prove statements directly about $\leftrightarrow\hat{}\{FS\}$-defined structures. R1 is procedure-template; R2-R4 are conceptual analogy; R5 is available-but-unused. This represents a **substantial methodological caveat** to the framework: the RAF apparatus motivates our approach but does not formally underwrite the present derivations. The novel theorems and conjecture below stand on the mean-field assumptions A0-A2 and the substrate-invariance hypothesis A3, not on inherited RAF theorems.

---

These five inherited results substantially shorten the formal work the present framework must do. The novel contributions of the framework — Theorems 1, 2 of Sections 2, 3 below, and Conjecture 1 of Section 4 — concern dynamics (growth rates, saturation, acceleration) that are not directly addressed by static RAF theory but that follow from combining the inherited closure apparatus with the substrate-specific carrier density $\rho\_S$.

### 1.9 Named assumptions (summary)

The framework's derivations rely on the following explicit assumptions. Each is named, numbered, and stated here for reference; each is invoked in the relevant theorem below.

- **A0 (short-cycle dominance).** Closure cycles $\leftrightarrow\hat{}\{FS\}$ have bounded length: $|P| \leq L\_max$ where $L\_max$ is a substrate-specific constant, *and* short cycles ($L = 0$ or 1) dominate the dynamics. The latter is the substantive part: A0 is not merely a technical bound but an assertion that pairwise and short-loop closures dominate over multi-step indirect ones. Substrate-specific operational specification of $L\_max$ is open (§6.6).

- **A1 (mean-field independence, uniformity, constant generation).** Pairwise closures are statistically independent (A1a), uniformly distributed (A1b), and the candidate-generation rate c is approximately constant in N during the early phase (A1c). May fail qualitatively under scale-free topology.
- **A2 (independence of input coverage and novelty exhaustion).** The retention probability factorises as P(retention) ≈ P(i) × P(ii) where (i) is input coverage and (ii) is novelty exhaustion. Approximate (mildly anti-correlated in reality).
- **A3 (substrate-invariant per-cycle rate).** The product $\lambda\_S = c\_S \cdot p\_\{0,S\}$ is approximately substrate-invariant: $\lambda\_\{S'\} \approx \lambda\_S$ for any pair of substrates S, S'. Note: this is substrate-invariance of the *product* $\lambda$; the components c and p_0 may individually vary in compensating ways. Empirical content of A3 depends on operational specification of $\tau\_S$ per substrate (§6.6).
- **A4 (time-as-parameter for derivations).** The framework's ontology commits to time as the ordering of Σ events (§1.7), but derivations in §2-§4 use t as a continuous parameter via dN/dt. The reconciliation — that absolute time is reconstructed from substrate-specific cycle counts via $\rho\_S = 1/\tau\_S$ — is gestured at but not formalised. A4 assumes that this reconstruction is well-defined within each substrate. Cross-substrate comparison of doubling times (Conjecture 1) requires that the reconstruction is consistent across substrates, which is part of what operational specification of $\tau\_S$ must establish (§6.6).

A0-A2 underpin Theorem 1 and Theorem 2. A3 underpins Conjecture 1. A4 underpins the use of continuous t throughout the derivations. The assumptions are listed in increasing order of substantiveness and decreasing order of confidence: A0 is partly technical; A1 and A2 are standard mean-field approximations with qualitative caveats (scale-free topology, anti-correlation); A3 is the central interpretive hypothesis whose empirical content is currently uncalibrated; A4 is the philosophical reconciliation of time-as-ordering with time-as-parameter, awaiting formal treatment.

The named-assumption discipline follows the pattern of Paper 0 Part 2 (P0-P5). Naming assumptions keeps the framework's epistemic commitments visible and facilitates identification of which derivations break if which assumption fails.

---

## Section 2. Theorem 1 — Multiplicative regime under mean-field assumptions

**Background.** The exponential dynamics of mutually-catalysing populations was first formalised by Eigen (1971) in the quasispecies equation and by Eigen and Schuster (1977-1978) in hypercycle theory. Their derivation shows that under mass-action kinetics with mutual catalysis, populations grow exponentially in the early phase. Our derivation generalises this from chemical concentrations to counts of retained F-in-S realisations under cycle closure, with substrate-specific carrier

density. The mathematical structure is inherited from Eigen-Schuster; the substrate-spanning generalisation and the carrier-density factor are the present contribution.

RAF theory (Inherited Results R1, R2 in §1.8) addresses the **static** question of whether a closed network exists in a given CRS, and gives thresholds for its emergence. RAF theory does not directly address the **dynamical** question of how the network grows once closure has occurred. Theorem 1 addresses this dynamical question under mean-field assumptions.

**Statement.** Under Assumptions A0 (bounded cycle length, §1.2) and A1 (mean-field independence of pairwise closures, below), in the early phase of substrate evolution while |R(t)| ≪ N_max, the retention rate satisfies

$$\frac{dN}{dt} \approx \lambda \cdot N(t) \cdot \rho_S$$

for some constant $\lambda > 0$ depending on substrate-specific candidate generation rate and per-pair closure rate. Equivalently:

$$N(t) \approx N(0)\, e^{\lambda \rho_S t}$$

with doubling time $\tau$_d = ln(2) / ($\lambda\rho$_S).

**Assumption A1 (mean-field independence, uniformity, and constant generation rate).** The probability that a candidate F achieves cycle closure with the existing pool R(t) is approximately the sum of independent pairwise closure probabilities with each existing F' ∈ R(t):

$$P(F \leftrightarrow_R^{FS} \text{some } F' \in R) \approx \sum_{F' \in R} p_0(F, F') \approx p_0 \cdot N(t)$$

where p_0 is the average per-pair closure rate, treated as a substrate-specific empirical constant.

A1 bundles three conceptually distinct claims that the derivation requires: - **A1a (independence):** closures involving different F' ∈ R(t) are statistically independent (no correlated clustering of F's in feature space). - **A1b (uniformity):** the per-pair rate p_0 is approximately uniform across pairs (no preferential coupling structures with hub-dominated topology). - **A1c (constant generation rate):** the rate c at which the variation step V produces candidates does not itself scale super-linearly with N during the early phase and is approximately substrate-specific-constant.

A1a and A1b are logically separable. A1c is named because the derivation uses c implicitly (via $\lambda$ := c·p_0); naming it allows the substrate-invariant claim in A3 to be unpacked into substrate-invariance of c, p_0, or their product.

A1 may fail when: - F's accumulate clusters in feature space, making pairwise probabilities correlated (violates A1a). Clustering reduces the effective per-pair rate but in mild cases preserves linear-in-N scaling on average. - **Strongly heterogeneous (e.g., scale-free, hub-dominated) validation topology** dominates the network (violates A1b qualitatively, not just

quantitatively). Real catalytic and citation/production networks are often scale-free, not Poissonian. Power-law degree distributions appear empirically in biology (master regulators), science (citation networks), technology (platform products), and chemistry (hub catalysts). In such cases, the reduction to an average p_0 smooths over a structural property that may be central to the substrate's dynamics. Linear-in-N scaling may be replaced by topology-dependent scaling such as N · log N or N^{1+$\varepsilon$}. - The candidate distribution drifts as R grows (specialisation, niche construction), making p_0 effectively decrease with t and violating A1's static assumption. - Cycle length is unbounded (violates A0).

In practice, A1 is the standard mean-field approximation used throughout population dynamics (Eigen 1971; Lotka-Volterra; Verhulst), and we adopt it on the same epistemic footing. We **explicitly acknowledge** that for substrates with empirically known scale-free topology, A1b is qualitatively suspect, and the derivation's conclusions in such cases should be regarded as approximations whose quantitative form may differ from linear-in-N.

**Sketch of derivation under A0 and A1.** Consider a time interval [t, t + $\Delta$t]. During this interval, the variation step V produces candidates at rate c (substrate-specific candidate generation rate). Each candidate is evaluated for closure: under A1, the probability of closure with R(t) is approximately p_0 · N(t). Therefore the number of retained candidates per unit time is

$$\frac{dN}{dt} \approx c \cdot p_0 \cdot N(t) \cdot \rho_S$$

where $\rho$_S = 1/$\tau$_S converts internal cycle counts to absolute-time rates. Setting $\lambda$ := c · p_0 yields the stated result. Integration gives N(t) = N(0) exp($\lambda\rho$_S t) in the early phase. The doubling time follows.

**Caveats.** 1. The "early phase" approximation breaks down as N(t) approaches N_max; the full dynamics is logistic (Theorem 2). 2. The per-pair closure probability p_0 is treated as a substrate-specific empirical parameter, not derived from first principles. 3. The relation A1 is mean-field; deviations from A1 modify the scaling. We acknowledge that absorbing the question of A1's validity into "empirical estimation of p_0" effectively makes the theorem a postulate of mean-field constancy of p_0, not a derivation from RAF theory.

**Epistemic status.** Under A0 and A1, the multiplicative regime follows from mass-action-style kinetics with cycle closure as the validation criterion. This is the standard derivation in population dynamics and chemical kinetics (Eigen 1971; standard texts). The novelty in the present formulation is the **substrate-spanning generalisation** — the same form of equation applies across all substrate rings from cosmochemistry to cognition, with substrate-specific rate constants $\lambda$_S and carrier densities $\rho$_S — and the **explicit factoring of rate into structural ($\lambda$) and temporal ($\rho$) components**. The derivation itself is RAF-inspired but stands on the mean-field assumption A1, not on RAF theorems.

**Consequence.** Multiplicative growth — the P1 regime of Paper 0 Part 2 — is a derived consequence of binary mutual validation under mean-field assumptions, not an independent postulate. The principal mechanism is that each retained F becomes a potential closure partner

for subsequent candidates: the closure-partner pool grows with each retention, and the rate of new retention is proportional to the existing pool.

This is **consistent with** but not specifically confirmed by the empirical data of Track E, in which the four substrate rings show approximately log-linear growth (exponential in linear time) in their respective active phases. The exponential pattern in Track E is compatible with the mean-field derivation here, but it is equally compatible with any other dynamical model that produces exponential early-phase growth (standard population dynamics, replicator dynamics, Bass diffusion, etc.). The Track E data therefore corroborate but do not uniquely select the present derivation.

---

## Section 3. Theorem 2 — Bounded accumulation as derived

**Statement.** Under Assumption A2 (independence of input coverage and novelty exhaustion, below) and the assumptions A0, A1 of Theorem 1, as |R(t)| approaches the saturation value N_max determined by the substrate's feature space, the retention rate approaches zero:

$$\lim_{N(t)\to N_{\max}} \frac{dN}{dt} = 0$$

Equivalently, N(t) approaches an asymptote N_max from below. The transition from multiplicative regime to bounded regime occurs gradually as the feature space fills.

**Assumption A2 (independence of input coverage and novelty exhaustion).** The two retention conditions are treated as statistically independent: - Condition (i): a candidate's inputs are covered by the existing pool's outputs (cycle closure into R). - Condition (ii): a candidate's outputs introduce novel functionality not already in R.

A2 states that P(retention) ≈ P(i) × P(ii). In reality these conditions are mildly anti-correlated: an F that easily satisfies (i) (its inputs are common in R) may also have less novelty (its outputs are likely already represented). However, the qualitative form (logistic saturation) is robust to the anti-correlation; A2 affects the inflection point but not the existence of saturation.

**Sketch of derivation under A0, A1, A2.** Recall that a candidate F is retained iff (i) and (ii). Under A2:

- Condition (i) becomes increasingly easy as N(t) grows. In the limit N(t) → N_max where the substrate's outputs span the available feature space, almost any randomly generated candidate's inputs are covered. The retention probability under (i) approaches 1.
- Condition (ii) becomes increasingly hard as N(t) grows. As R(t) approaches feature space coverage, fewer combinations of outputs remain undiscovered. The probability that a random candidate's outputs introduce novel functionality decays approximately as (N_max - N(t)) / N_max (assuming uniform distribution of valid configurations over the feature space).

Combining under A2:

$$\frac{dN}{dt} \approx \lambda \cdot N(t) \cdot \rho_S \cdot \frac{N_{\max} - N(t)}{N_{\max}}$$

This is the logistic equation. It has the solution

$$N(t) = \frac{N_{\max}}{1 + ((N_{\max} - N(0))/N(0)) \cdot e^{-\lambda \rho_S t}}$$

which begins exponentially (Theorem 1) and approaches N_max asymptotically.

**Caveats.** 1. The probability decay (N_max - N(t)) / N_max assumes uniform distribution of valid configurations over the feature space; real substrates may have clustered distributions producing different functional forms. 2. The "feature space" is a working abstraction; the empirical analogue is substrate-specific and not always cleanly bounded (compare the finite Mendeleev table for mineralogy with the much larger but still bounded space of viable proteins for biology). 3. The model is mean-field; fluctuations near saturation are not captured. 4. A2 is approximate; anti-correlation between (i) and (ii) shifts the inflection point but preserves the logistic form qualitatively.

**Epistemic status.** Under A0, A1, A2, the bounded accumulation is a mean-field phenomenological derivation. The qualitative form (logistic S-curve) is standard in population dynamics (Verhulst 1838) and innovation diffusion (Bass 1969); the present novelty is the **derivation from feature-space exhaustion in a closure network** rather than from resource-carrying-capacity arguments.

**Consequence.** Plateau (P5 regime in Paper 0 Part 2) is derived as the long-time limit of the same dynamics that produces multiplicative growth in the early phase. The bounded accumulation regime is not postulated; it follows from the finiteness of the substrate's feature space combined with the closure dynamics.

This is consistent with the empirical observation that mineralogy approaches an asymptote near 5400 species, that the biological fold repertoire has been stationary near ~1,300 superfamilies for several hundred million years, and that the simulation of hierarchical mutual-validation produces saturation in each level without explicit programming.

---

## Section 4. Conjecture 1 — Densification produces acceleration in absolute time

**Status note.** This section was originally formulated as Theorem 3. Independent review (22 May 2026) identified that without an operational definition of $\tau$_S that is independent of observed doubling times, the statement reduces from an explanatory theorem to a parameter-fitting identity. We therefore reframe it as a **conjecture** with explicit empirical content that admits

independent testing once operational definitions of $\tau$_S per substrate are developed. The mathematical structure is unchanged; only the epistemic status is downgraded to reflect the current state of the framework.

**Background.** RAF theory (Inherited Results R1, R2 in §1.8) addresses the static question of network closure. Theorems 1 and 2 of the present framework give the dynamical content within a single substrate (multiplicative early phase, logistic saturation). The empirical observation across the four substrate rings is that doubling times shorten across the ring sequence by orders of magnitude — from ˜Gyr (cosmochemistry) to ˜$10^2$ yr (cognition, still in its open phase). Conjecture 1 proposes a mechanism for this acceleration.

**Statement.** If substrate S' is a higher ring built on substrate S, and the carrier density satisfies $\rho$_{S'} > $\rho$_S (densification), then the doubling time of cumulative complexity on S' measured in absolute time is smaller than the doubling time on S in absolute time. The relationship is:

$$\frac{\tau_d^{(S')}}{\tau_d^{(S)}} = \frac{\lambda_S \rho_S}{\lambda_{S'} \rho_{S'}}$$

Under the additional **Assumption A3** (substrate-invariant per-cycle rate, see §1.9: $\lambda$_{S'} ≈ $\lambda$_S for any pair of substrates), this simplifies to:

$$\frac{\tau_d^{(S')}}{\tau_d^{(S)}} \approx \frac{\rho_S}{\rho_{S'}} = \frac{\tau_{S'}}{\tau_S}$$

**Sketch of derivation under A3.** Direct consequence of Theorem 1: $\tau$_d = ln(2) / ($\lambda\rho$) on each substrate. The ratio across substrates is the inverse ratio of the product ($\lambda\rho$). Under A3, the $\lambda$ terms cancel and the ratio of doubling times reduces to the ratio of cycle times.

**Assumption A3 (the central interpretive content).** The per-cycle rate constant $\lambda$_S is approximately substrate-invariant: $\lambda$_{S'} ≈ $\lambda$_S for any pair of substrates S, S'. Operationally, this says that the *structural* efficiency of one cycle V-Σ-R is the same on each substrate, and only the *temporal duration* of one cycle ($\tau$_S) varies. This is the densification hypothesis in mathematical form.

**Why this is a Conjecture and not a Theorem.**

1. Assumption A3 is **not derived** from the framework's primitives; it is the substantive proposal whose empirical consequences are to be tested.

2. The cycle time $\tau$_S is **not currently operationalised independently** of the observed doubling-time data. Without an operational specification of "what counts as one validation cycle" on each substrate that is independent of cumulative-complexity measurements, the relation $\tau_d \propto 1/\rho_S$ cannot be falsified: any observed doubling-time pattern can be matched by choosing $\tau$_S appropriately.

3. The conjecture's empirical content therefore depends on completing the operational specification programme (Open Problem 6.6). Once $\tau$_S is defined independently — for example

via independent measurement of geological cycle times, generational times, peer-review cycle times, production cycle times — then A3 becomes a falsifiable empirical claim. Until then, the conjecture is mathematically true under its own assumptions but empirically un-tested.

**Empirical content (provisional).** Using the empirical doubling times of Track E and **assuming A3**, the implied ratios of carrier densities across substrate transitions are:

| Transition | Doubling time ratio | Implied $\rho$ ratio (assuming A3) |
|---|---|---|
| Cosmochemistry → Mineralogy | ~3.5× | ~3× |
| Mineralogy → Biology (folds) | ~1.7× | ~2× |
| Biology (folds) → Cognition | ~$10^5$–$10^6$× | ~$10^5$–$10^6$× |

**Independent estimates of cycle times (back-of-envelope, with explicit operational choices).** The same row reads very differently depending on which operational choice of $\tau\_S$ is made:

| **Substrate** | **"Speciation/ transition cycle"** | **"Local production cycle"** | **Difference** |
|---|---|---|---|
| Cosmochemistry | ~$10^9$ yr (stellar generation) | ~$10^3$ yr (nucleosynthetic event) | $10^6$× |
| Mineralogy | ~$10^7$ yr (geological cycle) | days (crystallisation) | $10^{12}$× |
| Biology | ~$10^6$ yr (speciation) | ~1–10 yr (generation) | $10^5$–$10^6$× |
| Cognition | ~10–$10^2$ yr (paradigm shift / deployment) | ~1 yr (peer review / production cycle) | 10–$10^2$× |

If "speciation/transition cycle" is used as $\tau\_S$, the predicted ratios do not match the empirical doubling-time ratios. If "local production cycle" is used, they do not match either. The match depends on a specific (and currently uncalibrated) per-substrate choice of which cycle scale is operationally relevant.

This is precisely the problem the conjecture identifies: without independent calibration, the framework cannot distinguish between substrate-specific cycle definitions that fit the data and those that do not. The conjecture is **predictive once $\tau\_S$ is calibrated independently**; it is **descriptive otherwise**.

**What would constitute "independent" calibration — and a concern.** For Conjecture 1 to be supportable as an empirical claim, an operational principle for selecting $\tau\_S$ per substrate must be specified that does not itself reference the doubling-time data being explained. We do not currently have such a principle. The two natural unified choices presented above —

"speciation/transition cycle" and "local production cycle" — both fail to match the empirical pattern. If only substrate-specific eclectic choices work (a different cycle scale for each substrate, selected so as to fit the observed doubling time), then A3 is doing no explanatory work: it has been relabelled rather than supported. We are obliged to flag this. Conjecture 1 may already be in trouble even before further empirical work, unless a principled per-substrate operational definition of $\tau\_S$ can be supplied that is independent of doubling-time data. Identifying such a principle is the central task for any subsequent development of the framework's temporal claims.

**Testable prediction (conditional on operational $\tau\_S$ calibration).** Once $\tau\_S$ is operationally defined per substrate (Open Problem 6.6), the conjecture predicts that independently measured $\tau\_S$ ratios will match the doubling-time ratios above. If they match — A3 is supported and densification is confirmed as the mechanism of acceleration. If they do not match — A3 is falsified and the framework requires modification (perhaps by introducing substrate-dependent $\lambda\_S$).

**Consequence.** The empirical observation that doubling times shorten across substrate rings by orders of magnitude is, under the densification hypothesis, the visible record of the temporal step shortening across substrate transitions. The cycle remains the same; the carrier shortens it. This is consistent with simulation v2 (where uniform candidate-generation rate produces approximately constant doubling time in iteration counts across levels) and with the negative result of simulation v3 (which showed that "more candidates per fixed time" produces instant saturation rather than gradual acceleration, confirming that the correct interpretation of densification is shorter cycle in absolute time).

---

## Section 5. Relationship to Paper 1 Theorem 1

Paper 1 (Denysov 2026c) establishes:

**Theorem 1 (Paper 1).** For any local engine with absolute scale parameters dx (spatial) and dt (temporal), the knowability coefficient satisfies

$$\eta(dx) = \eta(dt) = 0$$

for any internal observer. Absolute scale parameters are operationally unknowable from inside the simulated world.

**Consistency with the present framework.** The present framework concerns the relative structure of retained F-in-S realisations and the cumulative information encoded in this structure. It does not address absolute scale parameters of the underlying substrate's dynamics. Therefore Paper 1 Theorem 1 and the structural information preservation claim of the present framework are concerned with different quantities:

- Paper 1 Theorem 1: absolute scales (dx, dt) of the underlying engine are unknowable.

- Present framework: the relative structure of retained F-in-S, and the cumulative history of $\Sigma$ events, is fully encoded in the current state of the world.

These two results are complementary. Internal observers face a strict limit on knowing absolute scales of their world but have, in principle, full access to the structural history of their world.

A combined statement, informally: **An internal observer cannot measure the absolute units of its world, but it can in principle reconstruct everything that has ever happened in its world from the current state of the world.**

---

## Section 6. Open problems

The following formal questions remain open and are deferred to subsequent work.

**6.1 [BLOCKING] Formal proof of structural information preservation.** A formal theorem stating that the current state of the world W(t) contains an injective encoding of all retained and eliminated F-in-S from times t' $<$ t. This requires a precise statement of what counts as "encoded" (Shannon information? Algorithmic information? Reconstructible by what class of decoder?) and a derivation under appropriate assumptions. We have working intuitions (six independent empirical disciplines operationally rely on it) but not a proof. **Blocking for §5 of the manuscript**, where the complementarity with Paper 1 Theorem 1 is claimed: without a proof of structural information preservation, the complementarity is asserted, not established. The §5 claim should currently be read as a conjectured complementarity awaiting formal proof.

**6.2 Rate of cross-level validation density buildup.** The framework asserts that substrate transitions occur as cross-level validation density accumulates. A formal model of this buildup — what determines the rate, what triggers the threshold for substrate self-sustainability, how it relates to internal saturation of the previous substrate — is not present in the current development. The threshold theorem of Hordijk-Steel (Inherited Result R2) and the boundary mechanism (R4) are relevant but have not been combined into a dynamical model.

**6.3 Connection to renormalisation group flow.** The structure of nested rings is suggestive of renormalisation group flow in statistical physics, in which different effective theories at different scales are connected. A formal derivation of our framework as an RG flow on the space of substrate descriptions would be valuable and is suggested as a direction.

**6.4 Carrier density as a derived quantity.** Currently $\rho$_S is introduced as a substrate-specific parameter. A formal account of *why* $\rho$_S takes the values it does — what physical, chemical, biological, or technological properties determine the validation cycle time on a given substrate — is missing. This is potentially deep: it may connect to information processing limits (Landauer's principle, Bremermann's bound, Bekenstein bound) at the most fundamental level.

**6.5 The primitive root closure dynamics.** The framework asserts that the first ring to close (the root) was contingent — first across the threshold rather than optimal. A formal model of

the threshold-crossing dynamics, and of what counts as "minimal coherence" sufficient to close the root ring against the protomatter background, is not provided. This connects to bootstrap programmes in particle physics and to the deepest layer of physics where the laws of physics themselves are taken as outputs of self-consistency conditions.

**6.6 [BLOCKING] Empirical validation of Assumption A3 via independent operational specification of $\tau\_S$.** Conjecture 1 rests on the assumption that the per-cycle rate constant $\lambda\_S$ is approximately substrate-invariant. The conjecture's empirical content depends on operationally specifying the validation cycle time $\tau\_S$ on each substrate **independently of the doubling-time data being explained**. Two natural unified choices ("speciation/transition cycle" vs "local production cycle") both fail to match the empirical doubling-time ratios when applied uniformly across substrates (see Conjecture 1 calibration table). Therefore: until an operational principle is identified that specifies $\tau\_S$ per substrate without circular reference to doubling-time data, A3 is **not falsifiable** and Conjecture 1 has descriptive rather than predictive content. This is a **blocking problem** for the empirical claim of Conjecture 1 — not a minor item of subsequent work. Identifying such a principle is the central task for any further development of the framework's temporal claims.

**6.7 Food-set depletion dynamics.** The framework treats each substrate ring's food set as a finite quantity of structured carrier inherited from the previous ring. Whether the depletion of this food set at any substrate level eventually limits the cumulative complexity reachable at that level, and on what timescale, is not derived. This question is potentially relevant to long-term trajectories of any substrate ring approaching its food-set boundary.

**6.8 Protomatter properties.** Protomatter is defined as matter in a pre-validation state with minimal properties required for self-validation participation. What these minimal properties are, and whether they admit further reduction, is not specified by the framework. This question is at the boundary of the framework's cognitive accessibility (compare Paper 1 Theorem 1 on absolute scale parameters) and may not admit a complete answer from within the framework itself.

**6.9 Robustness of A1 under scale-free topology.** A1 (mean-field independence and uniformity) assumes Poissonian-style validation topology. Real substrates frequently exhibit scale-free (power-law) topology with hub-dominated structure: biological master regulators, scientific citation hubs, technological platform products. Whether Theorem 1's linear-in-N scaling survives qualitatively under such topology — or whether topology-dependent scaling ($N \cdot \log N$, $N\hat{}\{1+\varepsilon\}$, etc.) replaces it — has not been investigated. This may be approachable through extensions of network-science results for preferential attachment dynamics.

**6.10 Reconciliation of time-as-ordering with time-as-parameter.** §1.7 commits ontologically to time as the ordering of $\Sigma$ events, but the derivations in §2-§4 use t as a continuous parameter (dN/dt). The reconciliation — that absolute time is reconstructed from substrate-specific cycle counts via $\rho\_S = 1/\tau\_S$ — is gestured at but not formalised. This is connected to A4 and to 6.6: cross-substrate comparison of doubling times requires that the reconstruction is consistent across substrates, which itself requires the operational specification of $\tau\_S$ that 6.6 demands.

**6.11 Hierarchy theorem for $\leftrightarrow\hat{}${FS}-defined subsets.** The nested-rings structure of the

framework (§2.3 of the manuscript) is motivated by R3 (sub-RAF hierarchy of Hordijk-Steel-Kauffman 2012) but R3 is a theorem about RAF sets, not about $\leftrightarrow\hat{}\{FS\}$-defined subsets (§1.8). A proper hierarchy theorem for $\leftrightarrow\hat{}\{FS\}$-subsets — establishing that nested subsets of mutual-validation-closure structures have the same hierarchy properties as nested RAFs — has not been proved. Until such a theorem is established, the nested-rings structure of the framework rests on conceptual analogy with R3 rather than on inherited mathematics.

---

## Section 7. Summary

**Named assumptions (A0-A4, see §1.9):** - A0: short-cycle dominance - A1: mean-field independence (A1a), uniformity (A1b), and constant generation rate (A1c) - A2: independence of input coverage and novelty exhaustion - A3: substrate-invariant per-cycle rate constant (the product $\lambda = c \cdot p\_0$) - A4: time-as-parameter for derivations (philosophical reconciliation with §1.7 deferred)

**Status of RAF results (§1.8):** - R1 (closure detection): inheritance pending operational catalysis specification per substrate - R2 (emergence threshold): conceptual analogy, not formally inherited beyond chemistry - R3 (sub-RAF hierarchy): conceptual analogy; formal hierarchy theorem for $\leftrightarrow\hat{}\{FS\}$-subsets is Open Problem 6.11 - R4 (boundaries $\rightarrow$ higher-level RAFs): conceptual analogy for substrate transitions - R5 (algebraic semigroup): available, not currently used

**No RAF results are formally inherited** in the strict sense of proving statements directly about $\leftrightarrow\hat{}\{FS\}$-defined structures. The RAF apparatus motivates the framework but does not formally underwrite its derivations. The novel theorems and conjecture stand on the named assumptions A0-A4, not on inherited RAF theorems.

**Concept-level precursors (acknowledged):** - Eigen quasispecies (1971) and hypercycle (Eigen-Schuster 1977-1978) — direct mathematical antecedent for Theorem 1 - Maturana and Varela (1972, 1980), autopoiesis — concept-level precursor for self-validating ontology - Rosen (1991), (M,R)-systems — self-referential closure tradition - Letelier et al. (2003), formal self-referential models

**Novel contributions of the present framework:**

- **Theorem 1**: Under A0, A1, $\leftrightarrow\hat{}\{FS\}$-style mutual validation produces multiplicative early-phase dynamics with doubling time inversely dependent on substrate carrier density $\rho\_S$. The functional form is the standard Eigen-style result; the substrate-spanning interpretive factoring of rate into structural ($\lambda$) and temporal ($\rho$) components is the present novelty.
- **Theorem 2**: Under A0, A1, A2, bounded accumulation (logistic plateau) follows from finiteness of substrate feature space.

- **Conjecture 1**: Under A3, densification of carrier across substrate transitions yields the observed acceleration of doubling time in absolute time. **Currently descriptive, not predictive**: becomes empirically falsifiable only once $\tau\_S$ is operationally specified per substrate independently of doubling-time data (Open Problem 6.6, blocking). Two natural unified operational choices already fail; the conjecture may already be in trouble unless a principled substrate-specific operational definition is supplied.
- **Food set / protomatter resolution**: The bootstrap question raised by the RAF requirement of an external food set is resolved by identifying food sets of substrate rings above the root with structured-matter outputs of predecessor rings, and the root's food set with finite protomatter in a pre-validation state.

**Epistemic status, honestly stated.** Theorems 1 and 2 are mean-field phenomenological derivations under named assumptions (A0, A1, A2). They are standard mean-field results applied to a new closure criterion ($\leftrightarrow\hat{}\{FS\}$); their novelty lies in the criterion and the substrate-spanning interpretation, not in the derivations themselves. Conjecture 1 is a substantive proposal about the mechanism of cross-substrate acceleration; its empirical content is currently uncalibrated and may not survive principled operational specification of $\tau\_S$. The framework is a **mathematically informed theoretical programme with explicit assumptions and testable consequences** rather than a fully formalised mathematical theory.

**Blocking open problems** (must be addressed before central claims can be empirically supported): - **6.1** Formal proof of structural information preservation (blocks §5 complementarity claim with Paper 1) - **6.6** Operational specification of $\tau\_S$ per substrate (blocks Conjecture 1 empirical content)

**Eleven open formal problems** have been identified in total. Two are flagged as blocking.

The named-assumption discipline (A0-A4) follows the pattern of Paper 0 Part 2 (P0-P5). The relationship to Paper 1 Theorem 1 is **conjectured complementarity** (absolute scales unknowable per Paper 1; relative structure preserved per present framework, conjectured pending proof of 6.1).

## References


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

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

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.

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

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.

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