

The triple unity of life: closure, code and the topology of living form

Juan Duan¹, Songbo Zhang^{2*}

¹ Schistosomiasis National Key Specialty, Hunan Provincial Institute of Schistosomiasis Control (Hunan Provincial Third People's Hospital), Yueyang, China

² Yueyang Vocational Technical College, Yueyang, Hunan, China

* Corresponding author: songbozhang@yvtc.edu.cn

Manuscript prepared for the Journal of Theoretical Biology, October 2026

Abstract

What is common to every living thing? This paper develops a framework — three universalities of unequal strength, closed into a regulatory loop — as a Lakatosian research programme with a declared hard core, a graded protective belt, and a ranked positive heuristic. Four new results carry the weight. (1) Monte-Carlo margin analysis: measured rotor/gradient pairs sit at a mean 12.5th percentile of a neutral null under four registered null models, so c-rings run close to the thermodynamic floor as a statistical fact, with the mammalian c8 rotor the tightest-constrained case. (2) An exhaustive design table of 9,936 three-node signed loop topologies: self-maintenance is generic, robustness is graded by sign class, the all-positive class is the most robust, and the canonical loop sits mid-class. (3) A basal-leak extension reconciles the exactly absorbing dead state with the design table: spontaneous ignition is rare (6.8% of parameter sets) and gated by a smooth critical-leak line, so genesis requires founder events; spatially, pinning is structural and vesicle-like dispersal unlocks invasion above a critical coupling. (4) The same loop, applied to a second ignition with independent historical constraints, accounts for eukaryogenesis: the prokaryotic genome ceiling, the non-transferability of membrane-buried genes, and the singularity and irreversibility of the event. Honest negatives are reported throughout. Eleven predictions with frozen protocols close the paper.

Keywords: chemiosmosis; ATP synthase; genetic code; error threshold; persistent homology; design table; autopoiesis; self-organised criticality; origin of life; protocell; research programme

Classification: Systems Biology; Biophysics

1. Introduction


An earlier formulation of this framework closed with a claim we can now afford to examine: that the physics is “specific enough to fail — and that, so far, it has not”. The obvious rejoinder was immediate and correct: *has not failed yet* is weak comfort when the framework has not been pushed to where it would fail. This manuscript is our answer. It does three things the earlier formulation could not.

First, it states its own logical structure. Following Lakatos, the theory is now presented as a **hard core** (the three unities and the loop, protected from direct falsification and answerable only as a whole), a **protective belt** (parameterisations, thresholds, metric choices — adjustable by design), and a **positive heuristic** (a ranked programme of constructions and experiments). Section 2 declares which is which. This is not defensive machinery; it is the difference between a theory and a mood.

Second, it converts three assertions into statistical statements with null models — the single most effective repair available to a framework of this kind. Does anything actually run close to the thermodynamic floor, or merely above it (Sect. 3.2)? Is the closed loop just any positive-feedback cycle, or does the C1C2C3 coupling occupy a privileged region of design space (Sect. 6)? Are metabolic graphs topologically distinguishable from signalling graphs when the analyst is not allowed to see the labels (Sect. 5.4)? Each question now has a number attached, and one of the numbers is negative — the blind clustering fails, and we report that failure as the honest provenance of the surviving claim.

Third, it borrows a method wholesale from the literature on exhaustive network design. Ma, Tang and collaborators answered “which network topologies can achieve biochemical adaptation?” by enumerating every three-node signed topology, simulating each over sampled parameters, and reading off a finite design table [1]. We ask the corresponding question for the present framework — *which three-node signed loops can sustain themselves?* — and let the enumeration decide whether the C1C2C3 coupling is privileged, common, or trivial (Sect. 6). The same paradigm has been productive for cell-cycle dynamics [2], for the yeast G1/S switch [3], and, with recurrent neural networks replacing hand-written kinetics, for function-to-topology maps generally [4].

Everything else from the earlier formulation survives here: the three unities of unequal strength; the loop C1 (closure enables a gradient), C2 (the gradient pays for fidelity), C3 (fidelity makes structure heritable); the evidence grades A/B/C/D attached to every claim; and the open archives — PDB, UniProt, KEGG, CrossRef — from which every number is derived. Eighty-five references were CrossRef-verified in the initial analysis; eighteen more are verified here, including the autopoiesis and chemoton lineage whose absence was correctly flagged as a genealogical gap.



Closed-loop attractor requires all three: remove one layer and the loop collapses

evidence discipline: A empirical · B inference · C model · D interpretation

The triple unity as a closed loop. Three layers of unequal strength — language (structure), mechanism (energy) and symbol (information) — are joined by the couplings C1, C2 and C3. C1: the topology of the compartment enables a sustained proton-motive force. C2: ATP flux pays the thermodynamic cost of proofreading and code fidelity. C3: the error-minimising code rebuilds structure from a template, closing the loop. Evidence grades: A empirical, B inferred, C model, D interpretive.

2. The structure of the theory: hard core and protective belt

A foundational theory earns that status progressively, and the process has a known shape [5–6]. It is worth being explicit about where this framework stands in that shape.

Hard core. Four propositions, not individually falsifiable but answerable only as a unit:

HC-1 (mechanism). Autonomous energy metabolism is transmembrane and rotary: ion-motive gradients drive ATP synthesis through the F/A/V synthase superfamily.

HC-2 (symbol). Cellular life reads one genetic code with a homologous decoder, and the code's error-minimising architecture is a retained property of all natural variants.

HC-3 (language). The structures life depends on — closed membranes, cyclic reaction networks, conserved linkages — are classified by topological invariants, and these invariants carry functional information.

HC-4 (the loop). The three are coupled: closure enables the gradient (C1), the gradient pays for fidelity (C2), fidelity makes structure heritable and evolvable (C3). A living system is a self-maintaining system realising this loop with positive entropy production.

Protective belt. Everything quantitative that the hard core needs in order to touch data, each with an explicit adjustability grade:

Table B. The protective belt, with adjustability grades. A failed prediction first motivates adjustment here; repeated failure across the belt demotes the adjacent hard-core proposition.

Belt element	Grade	Adjustability
G_ATP physiological range (50–65 kJ mo)	medium	scenario-dependent; mitochondrial matrix values are lower
(S) feature set and clustering metric	high	normalisation choices change cluster counts
Loop-model rate laws and parameters	high	any reasonable kinetics; structure is the claim
Margin null ($c \sim U(4,20)$, $p \sim U(150,250)$ mV)	medium	a neutral default; alternatives in the package
Threshold constants in the founder model	high	illustrative; only the threshold structure is claimed

Negative heuristic. Author and critical readers agree not to attack HC-1...HC-4 with a single anomaly in the belt. The cost of this protection is responsibility: the hard core must pay rent in the form of novel, corroborated predictions (Sect. 9), and if the belt adjustments become ad hoc without limit, the programme degenerates — Lakatos’s own criterion for abandonment.

Positive heuristic. The ranked programme of Sect. 10: complete the design table at four nodes and with RNN-searched kinetics [4]; measure the in-situ tunnelling slope P2; run the founder-ignition protocol on protocell platforms P6; and the derivative predictions P7–P10.

We flag one boundary honestly. “Hard core” is a methodological status, not a truth certificate: it organises retreat and advance; it does not confer immunity from history. If the design table had shown that self-maintenance ignores sign structure, or that the loop’s predictions are reproducible from any single layer alone, HC-4 would be the first proposition we would surrender.

3. The mechanism unity (energy layer)

3.1 Closure is the precondition for a gradient

Proton flux near a membrane obeys the Nernst–Planck equation

$$J_p = -D \left(\nabla c + \frac{cF}{RT} \nabla \Psi \right),$$

with the steady-state condition that the flux be divergence-free. Two geometries behave in qualitatively different ways.

An **open sheet** has a boundary through which normal flux need not vanish. The edge of a lipid bilayer exposes hydrocarbon to water, an expensive defect; the resulting line tension drives pore nucleation, and the boundary behaves like a Robin condition with a positive leakage coefficient. For the corresponding elliptic problem, the maximum principle bounds the interior by the source divided by the leak: the more the edge leaks, the smaller the steady gradient can be, and in the leak-dominated limit the gradient vanishes [B].

A **closed surface** has no boundary at all. The Fredholm compatibility condition for the steady problem then reduces to a single statement — total source balances total sink — and the gradient becomes a storable resource rather than a lost cause [B]. This is Proposition E1:

E1. Maintenance of a macroscopic proton-motive force requires topological closure of the membrane.

An honest caveat, raised by a critical reader of an earlier draft: bilayer sheets in water close spontaneously, so “open membrane” is hardly a long-lived state of nature, and the empirical content of E1 might seem thin. We accept the observation and sharpen the consequence. Closure is not merely *allowed*; it is *energetically inevitable*, and the numbers are large. A flat disc of radius 100 nm carries $2\pi R\Gamma$ of edge energy; at a measured line tension of $\Gamma = 25$ pN [7], that is about $3.8 \times 10^3 k_B T$ released on closing [A for the tension, B for the conclusion]. Pore nucleation, the reverse process, costs $\pi\Gamma^2/\Sigma$ — 150 to 6,100 $k_B T$ for physiological tensions (Fig. 8a) — which is why intact membranes hold their contents. E1 is thus best read not as a constraint that nature must satisfy but as a free lunch nature has already eaten: the topology comes first because it is cheap.

3.2 The coupling window, as a margin distribution

The thermodynamic floor is unchanged: to make one ATP, n_H protons must return through the rotor,

$$n_H F \Delta p \geq \Delta G_{ATP},$$

and $n_H = c/3$ for an F-type synthase with a c-ring of c subunits. An initial analysis presented eight measured proton-motive forces and eight c-ring stoichiometries and asserted that “every system satisfies its own floor”. Two readers of an earlier draft put a finger on the same wound: at $G_{ATP} = 55$ kJ mol⁻¹ the mammalian rotor requires $p = 214$ mV, while measured mitochondrial p is 180–200 mV. The assertion and the numbers were in tension, and no amount of prose could dissolve it.

The repair is to stop treating the margin as a binary. Define the thermodynamic margin

$$M = n_H F \Delta p - \Delta G_{ATP},$$

propagate uncertainty through it ($G_{ATP} \sim U(50,65)$ kJ mo for bacterial and chloroplast systems, $U(46,58)$ for the mitochondrial matrix scenario; $p \sim N(\cdot)$ with the compiled measurement uncertainties), and report the probability that the margin is positive. Table 1 gives the result. Three readings matter.


First, the bacterial and chloroplast rotors are comfortable: *B. subtilis* ($p = 220$ mV) has $P(M>0) = 0.96$ with a median margin of +13 kJ mo; thylakoid runs at 0.97. Second, the mammalian case is exactly what those readers suspected and is now stated as such: at $p = 200$ mV, the rotor has $P(M>0) = 0.46$ under matrix G scenarios and median margin 0.5 kJ mo — the floor passes *through* the measured operating point. Liver mitochondria ($p = 190$ mV) sit at $P(M>0) = 0.28$. The mammalian rotor is not a counterexample to the floor; it is the **tightest-constrained case in the data set**, and the framework treats it as such, with an explanation tree that is now part of the theory rather than a footnote: (a) matrix G_{ATP} sits at the low end of the physiological band; (b) the ring co-evolved with adenine nucleotide and P_i carriers, so the relevant free-energy accounting includes transport terms beyond the rotor [8]; (c) partial uncoupling (slip) and proton leak shave the effective p at the enzyme below the bulk value [9]. Each branch is testable; none is assumed.

Third — and this is the genuinely new, falsifiable content — the margins are *systematically smaller* than a neutral null would predict. Under the zero model ($c \sim U(4,20)$, $p \sim U(150,250)$ mV, same G band), the median absolute margin is 25.3 kJ mo; the six real rotor/gradient pairs with both quantities known have a median absolute margin of 4.5 kJ mo, and four of the six lie below the null's 12th percentile (mean percentile 12.5). Rotor stoichiometry and operating gradient are not independent draws from their physiological ranges: they are paired so as to run close to the thermodynamic floor. This is the energy layer's first non-trivial statistical prediction, it is graded **A/B**, and it now carries a frozen robustness protocol (Appendix A), registered before recomputation. Four null models were fixed in advance, with a decision rule: retain the claim if every null puts the mean percentile below 25; demote to "suggestive" if any exceeds 35. The nulls: N0, the uniform model above; N1, c drawn from the empirical set of resolved c-rings {8, 10, 11, 13, 14, 15} rather than uniformly; N2, p drawn log-uniformly over [150, 250] mV; N3, a phylogenetically stratified null in which each system's rotor is drawn from its own clade's c-ring set (mitochondrial 8–10, bacterial 10–15, chloroplast 13–15) and the percentile is computed within clade. Results: N0 12.48, N1 18.15, N2 12.66, N3 19.03 — the claim is retained under all four. A bootstrap over the six systems gives a 95% CI for the mean percentile of [3.6, 21.4]: the interval excludes the null median (50), and it is also honestly wide — six systems are six systems, and the text keeps saying so. The selection effect (only systems with both quantities measured can enter) is bounded rather than waved at: if the two c-ring systems lacking paired p


measurements sat at the null median, the mean percentile would rise to 21.9, still under the retention line; five missing systems would push it to 29.5, past retention but short of demotion — and five is more such systems than exist. A Bayesian contrast agrees: a floor-seeking generative model outscores the independent null by a Bayes factor of 19–78 across prior bandwidths — corroborating rather than primary, since at $n = 6$ the factor is sensitive to the likelihood form; the primary evidence is the frozen-protocol percentile (Appendix A). The explanation tree (a)–(c) above is, finally, counted as what it is: three belt adjustments spent on one anomaly. Under the programme’s own surrender criteria (Sect. 10.3) that is short of the ad-hoc limit, and each branch names its own kill; the count is kept so that the next adjustment cannot be added silently. The registered extension stands: as more c-ring structures and in-vivo p measurements accumulate, the percentile should stay low — and if future pairs cluster at the null’s median under every null model, the “tuned coupling” reading dies and the belt adjusts.

Thermodynamic margins as probabilities, not verdicts. $M = n_H F p_{GATP}$ with $G_{ATP} \sim U(50,65)$ kJ mo (cytosolic scenario) or $U(46,58)$ (mitochondrial matrix scenario) and $p \sim N(\cdot)$ truncated to $[120, 280]$ mV; 2 draws per system. $P(M>0)$ is the fraction of draws with a positive margin. The mammalian rotor is the tightest-constrained case — a finding, not a contradiction — with the explanation tree given in the text. Systems without both quantities measured are listed for completeness. Inclusion rule: a system enters the percentile statistic only when a resolved c-ring stoichiometry and a measured p both exist; $c = 11$ and 12 rings are known in other organisms but lack paired p and are excluded — the rule is frozen in the Appendix A protocol, not chosen per analysis.

System	p (mV)	c	$n_H = c/3$	G scenario	P(M>0)	median M (kJ mo)
<i>H. salinarum</i>	230 25	—	—	50–65	—	—
<i>B. subtilis</i>	220 20	10	3.33	50–65	0.96	+13.3
Cardiac mitochondrion	200 15	8	2.67	46–58	0.46	0.5
Liver mitochondrion	190 15	8	2.67	46–58	0.28	3.1
<i>E. coli</i>	180 15	10	3.33	50–65	0.52	+0.4
Yeast mitochondrion	180 15	10	3.33	46–58	0.84	+5.9
Thylakoid (<i>spinach</i>)	160 15	14	4.67	50–65	0.97	+14.6
<i>Spirulina platensis</i>	—	15	5.00	50–65	—	—
Thermoalkaliphile	—	13	4.33	50–65	—	—



The coupling-ratio window. (a) The thermodynamic floor $n_H = \Delta G_{ATP}/(F\Delta p)$ for four values of the ATP free energy; the shaded band is the measured biological range of p . Markers show real c -ring stoichiometries ($c/3$). (b) Reported proton-motive forces across the three domains (Table 1).




Thermodynamic margins. (a) Monte-Carlo margins M for six systems with both quantities measured: bars span the 5th–95th percentile of M , dots the median; squares give $P(M > 0)$ (upper axis). Green: comfortable margin; red: margin straddling zero. (b) Distribution of $|M|$ under the neutral zero model ($c \sim U(4,20)$, $p \sim U(150,250)$ mV, same G band); coloured lines mark the real systems' median $|M|$. Real rotors sit at a mean 12.5th percentile of the null — the coupling ratio and the operating gradient are paired, not independent.

3.3 Geometric sensitivity of electron transfer

Electron transfer along the respiratory chain is governed by non-adiabatic Marcus theory,

$$k_{ET} = \frac{2\pi}{\hbar} |H_{AB}|^2 \frac{1}{\sqrt{4\pi\lambda k_B T}} \exp \left[-\frac{(\Delta G^0 + \lambda)^2}{4\lambda k_B T} \right],$$

with the electronic coupling decaying exponentially with donor-acceptor distance, $|H_{AB}| \propto e^{-\beta d/2}$, hence $k_{ET} \propto e^{-\beta d}$. In proteins the decay constant is $1.0\text{--}1.4 \text{ \AA}^{-1}$; the canonical measurement is the ruthenium-modified azurin series, $\beta = 1.1 \text{ \AA}^{-1}$ [10], and the same scaling organises rates across dozens of redox proteins [11–12]. A rate falls an order of magnitude every $\ln 10/\beta \approx 2.1 \text{ \AA}$ (Fig. 4).



The Moser–Dutton ruler. Electron-transfer rate decays exponentially with distance; the solid line is the ruler with $\beta = 1.1 \text{ \AA}^{-1}$ measured in ruthenium-modified azurin [10], and the shaded band spans the protein literature range $\beta = 0.9\text{--}1.4 \text{ \AA}^{-1}$. The vertical line marks the empirical 14 Å ceiling between redox cofactors found in a structural survey of all redox proteins of known fold [12]; longer transfers are always bridged by cofactor chains. The green band marks the turnover range of respiratory enzymes: tunnelling at or below 14 Å is fast enough to stay out of the way of catalysis.

The structural survey behind the 14 Å ceiling is the key empirical fact [12–13]: proteins do not evolve optimal tunnelling pathways, they evolve *short ones*, and

where the distance is longer, they insert cofactors. The consequence for this framework is that the Ångström-scale architecture of a protein — and of the supercomplexes that assemble it — is a load-bearing parameter of energy production [A]. It also yields the sharpest experimental target of the mechanism layer (P2, Sect. 6): if cofactor spacing is tuned to sit near the exponential penalty, then perturbing complex spacing should degrade respiration exponentially, with a decay constant directly comparable to the tunnelling .

3.4 Scale: the energy–size problem and its two solutions

Power that depends on membrane flux is bounded by the surface-to-volume ratio,


$$P_{ATP} \leq \eta \cdot \frac{S}{V} \cdot J_{proton}$$

and for a sphere $S/V = 3/R$: the bigger the cell, the less power per unit volume. Endosymbiosis is the internal solution; the fractal transport networks of macroscopic organisms are the external one, delivering the quarter-power scaling $B \propto M^{3/4}$ [14]. Both are, in the language of this paper, the mechanism layer borrowing topology from the structure layer — and Section 7.6 makes the internal solution carry historical weight: the scaling constraint is what routes eukaryogenesis through endosymbiosis and nothing else.

3.5 One motor, one graded ancestry

Proposition E. *All life with autonomous energy metabolism — bacteria, archaea, eukaryotes — makes ATP by rotary catalysis driven across a membrane; F-, A- and V-type ATP synthases form one homologous superfamily; and the coupling ratio n_H of every lineage is confined to the thermodynamic window of Sect. 3.2.*

Homology is a graded claim, so we tested it with sequences. We retrieved sixteen - and -subunit sequences of F-, A- and V-type synthases (bacteria, mitochondrial, chloroplast, archaeal and vacuolar; accessions in Table S1) and computed pairwise local-alignment identities (BLOSUM62; Fig. 5a). Three readings follow. First, and paralogs *within* a lineage are only 26–30% identical (*E. coli* 26.2%, human mitochondrial 26.9%, archaeal 30.1%, human V-type 25.9%): the duplication that created the two catalytic subunits is ancient, predating the radiation of the family. Second, orthologs across vast phylogenetic distances are much closer (73.0% between *E. coli* and human mitochondrial), which is what a shared, slowly-drifting engine looks like. Third, the archaeal A-type subunit is closer to the eukaryotic V-type (53.0%) than to the bacterial F-type (30.1%) — the signature of an A/V clade sharing a more recent ancestor with each other than either shares with F. A neighbour-joining tree on the same identities recovers exactly this topology (Fig. 5b). Proposition E is supported at the sequence level, but as a *graded* homology claim, not an identity claim [A].



*Homology of the rotary motor across the three domains. (a) Pairwise local-alignment identities of 16 - and -subunit sequences of F-, A- and V-type ATP synthases (UniProt accessions in Table S1). Within-lineage / paralogy runs 26–30% (boxed values), while cross-domain orthology runs far higher (*E. coli* vs human mitochondrial : 73%). (b) Neighbour-joining tree of the same identity matrix. The α and β subunits each form their own clade — the duplication predates the F/A/V radiation — and the archaeal A-type falls adjacent to the eukaryotic V-type (53% identity), consistent with the known affinity of A- and V-type pumps.*

3.6 The boundary of the hard core: parasites, viruses, and one repaired argument

Two boundary cases are carried over; one is repaired.

Degenerate parasites (mycoplasmas, *Rickettsia*, *Buchnera*) remain losses of the mechanism, not counterexamples; Proposition E speaks of autonomous energy metabolism.

Viruses. An earlier formulation argued that the virus, carrying only the symbol layer, is “an argument for the necessity of the triple”. One objection, which we accept, calls this circular: that viruses need a host loop is a restatement of what they are, not independent evidence. The repaired statement is narrower and weaker, and therefore better: viruses are the framework’s *boundary condition* — systems excluded from HC-4 by construction — and their existence poses the question the framework must answer elsewhere: what, mechanically, makes a symbol system unable to go it alone? The answer is in HC-1 and HC-3 (a symbol system with no membrane and no decoder cannot close the loop), but the observation of viruses no longer counts as evidence for it. The hard core earns its keep from Sect. 6 and Sect. 8, or not at all.

Falsification of E. A free-living organism whose ATP synthesis does not depend even indirectly on a transmembrane ion gradient; or an ATP synthase, in an organism with autonomous energy metabolism, not homologous with F/A/V. Neither is known.

4. The symbol unity (information layer)

4.1 The channel

The coding channel maps 64 codons onto 21 symbols (20 amino acids plus termination), so a codon carries at most 6 bits. The redundancy is not uniform: the third position wobbles, the first two carry the load. This layering is what lets one table serve genomes whose GC content varies by tens of percent while the amino-acid repertoire stays fixed — a property of channel design, not an accident of the particular table [A].

4.2 The energetic price of fidelity

Uncorrected molecular recognition caps single-step discrimination at an error of order 10^{-1} ; measured translation error is about 10^{-2} [15]. Hopfield and Ninio showed how a proofreading step — delay, discriminate, discard, fuelled by one high-energy bond — squares the discrimination, and an n -step cascade raises the discrimination to the $(n+1)$ -st power [16–17]. The thermodynamic price has a floor: each order-of-magnitude reduction in error costs at least $RT \ln 10 \approx 5.9$ kJ mol⁻¹ at 310 K. A translation event spends about four high-energy phosphate bonds, of which proofreading takes roughly a quarter [18]. With the Landauer bound $k_B T \ln 2$ per bit underneath it all [19–20], the message of this subsection is that fidelity is purchased: the symbol layer is embedded in the mechanism layer, and the exchange rate is known [B].


4.3 The code is optimal — with the alternative hypothesis named

The recomputation stands: standard code $M_s = 7.465$, best 0.007% of block-respecting random codes; seven natural variant codes all in the best 0.01%, yeast and invertebrate-mitochondrial codes marginally better than the standard code (Fig. 7). Two boundary conditions are now stated with the force the objection demanded.

First, the metric is protocol-sensitive. Excluding terminator-involving mutations (the Haig–Hurst protocol, and the right choice because the vertebrate mitochondrial code has five stops) changes variant-code scores dramatically; the percentile we report is a property of *metric plus code*, not of the code alone.

Second, error minimisation is one of (at least) two live explanations of code structure. The co-evolution/history hypothesis — codon assignments frozen as the amino-acid inventory expanded — also produces apparent error minimisation, because early-appearing amino acids occupy chemically central positions. The

earlier formulation did not mention this; a critical reader did. The honest statement: the recomputation shows the standard architecture is robust to reassignment, which any origin story must accommodate; it does not adjudicate between selection on error minimisation and structural inheritance from co-evolution. Both remain live [21–23].



Error minimisation of natural codes (protocol of Haig & Hurst 1991; terminator mutations excluded per code; block-respecting random codes). (a) Null distribution of M_s with the standard code (solid) and seven variant codes (dotted). (b) Percentiles: every natural code tested lies in the best 0.01%.

4.4 The cost of freezing, and the error threshold

Changing the meaning of one codon requires rewriting every genomic site that carries it — of order sites per genome — a fitness barrier so large that the code is, for practical purposes, frozen [B]. “Frozen” can be made exact with the quasispecies model [24]. For a single-peak landscape with master fitness σ , a genome of length L and per-base mutation rate μ has per-genome error rate $U = 1 - (1 - \mu)^L$, and the master sequence survives only if


$$U < U_c = 1 - \frac{1}{\sigma}.$$

We verified the collapse with an exact one-way-degradation Wright–Fisher simulation ($L = 50$, $\mu = 5$, population 5,000; Fig. 7a): the master frequency is 0.36 at $U = 0.72$, 0.10 at $U = 0.78$, and zero by $U = 0.83$, bracketing the analytic threshold $U_c = 0.80$ [A]. The threshold tracks $U_c = 1 - 1/\sigma$ across $\mu = 2$ –8 (the simulated collapse point stays within one -grid step of the analytic value) and is stable from $L = 50$ to $L = 100$ at fixed μ , so the headline numbers are not grid or parameter artefacts.

The value of the exercise is that symbol collapse becomes a number, and the number can be compared to real genomes (Fig. 7b). RNA viruses, with per-base rates of $\mu \sim 0.1$ per replication [25–26] and small genomes, sit near $U = 0.1$ –1: close to the threshold, and the largest riboviral genomes (~ 26 –32 kb) sit right at it. DNA microbes hold $U \sim 0.003$ [27], orders of magnitude below. The extreme case is the

viroid — a ~400-nucleotide non-coding RNA whose measured substitution rate, 1/400 per site, is the highest reported for any replicon and is survivable only because its information content is trivial [28]. And the one lineage that breaks the viral ceiling, the nidoviruses at ~30 kb, does so by encoding a proofreading exonuclease [29–30] — the mechanism layer buying back symbol-layer fidelity, which is coupling C2 operating in plain sight.

Read. The error threshold is the boundary at which the symbol layer can hold an alphabet at all. Cross it, and there are no messages — genetic information dissolves into a mutation cloud. The threshold is why the code is frozen, and the frozen code is why the threshold is rarely discussed: cellular life lives comfortably below it because it bought proofreading early and permanently.



The error threshold, in simulation and in nature. (a) Master-sequence frequency against per-genome error rate for the single-peak quasispecies ($\sigma = 5$, $L = 50$): the exact deterministic curve, the analytic collapse point $U_c = 1 - 1/\sigma = 0.80$, and the Wright-Fisher simulation (population 5,000; three replicates per point; one-way degradation). The simulated collapse brackets the analytic threshold. (b) Per-genome error rate against genome size for real systems (log-log). The dashed horizontals are the thresholds for $\sigma = 2, 5$ and 20. RNA viruses (shaded band, per Sanjuán et al. 2010) sit at $U \sim 0.04-1$ with genomes below ~32 kb; the viroid point [28] is the extreme; DNA microbes [27] sit orders of magnitude below every threshold. The nidoviruses extend past the riboviral size ceiling by encoding proofreading (nsp14 ExoN), the C2 coupling made enzymatic.

4.5 Proposition I

Proposition I. Every cellular organism shares one codon table and a homologous decoder, and the table is a near-optimal solution to the problem of minimising the phenotypic cost of mutation.

Graded **A** for universality within cellular life, **A** for error minimisation (as recomputed above), and **B** for the extension to variant codes as a class. **Falsification:** a natural code worse than the 99.99th percentile of random codes under a stated protocol; or a cellular decoder not homologous with the ribosome. Neither is known.

What remains genuinely open — and is *not* claimed here — is why the particular table was fixed historically; the framework carries its robustness, not its biography.

5. The language unity (structure layer)

5.1 Topological typing, and how to compute it

To any living structure S we assign an invariant vector

$$\tau(S) = (\chi(S), \beta_1(S), \beta_2(S), \text{Lk}(S), \{b_i, d_i\}),$$

with the Euler characteristic, β_k the Betti numbers, Lk the linking number, and $\{b_i, d_i\}$ a persistence barcode (births and deaths of homology classes along a filtration) [31–32]. The power of τ is homotopy invariance: it is blind to continuous deformation, so it records *how* a structure is connected rather than *where* its atoms sit — which is exactly why one family of descriptors can compare a lipid vesicle, a mitochondrion and a metabolite network. Two computational points matter for biology. First, the descriptors are computable from ordinary experimental data: a PDB file supplies a C point cloud, and a Vietoris–Rips filtration on it yields the barcode in seconds on a laptop [33–34]. Second, they are robust — a claim we test directly in Sect. 5.5, because the standard objection to topological descriptors is noise.

5.2 Membranes: the Gauss–Bonnet ledger

The Helfrich bending energy of a membrane is

$$F = \int_M \left[\frac{\kappa}{2} (2H - c_0)^2 + \bar{\kappa} K \right] dA,$$

and for a closed surface the Gauss–Bonnet theorem gives $\int_M K dA = 2\pi\chi = 4\pi(1 - g)$, so the Gaussian term is


$$\bar{\kappa} \int_M K dA = 4\pi\bar{\kappa}(1 - g),$$

a *topological constant*. No smooth deformation changes it; only a change of genus — fission, fusion, pore formation — rewrites it. Membrane scission is therefore a topological event with a price tag of $\Delta F \sim 4\pi/\bar{\kappa}$ per unit change in genus, and the price is paid in mechanochemical currency [35].

The accounting is quantitative (Fig. 8b). Measured Gaussian-curvature moduli are $\bar{\kappa} \approx -(0.4\text{--}1.5)\kappa$ with bending rigidities $\kappa = 11\text{--}24 k_B T$ for common lipids [36–38], giving barriers of $55\text{--}450 k_B T$. A dynamin- or Drp1-mediated fission hydrolyses on the order of $10\text{--}50$ GTP [39–40], i.e. roughly $210\text{--}1,050 k_B T$: the enzyme budget and the topological toll are the same order of magnitude. The coincidence is not trivial. It

means that neither the lipid composition nor the enzyme alone sets the price of division; both sit within a factor of a few of the Gauss–Bonnet toll, which is what a system tuned to the edge of feasibility should look like. Where low- lipid compositions are used, the barrier falls to tens of $k_B T$ and protein-free fission becomes thermally accessible — as observed in model and bacterial systems [A for the moduli; B for the inference].

Read. The number of handles is the unit of account. Cell division rewrites one entry in the genus ledger, and physics sets the price of the entry; the cell pays it in GTP.



Membrane topology has a price, and cells can pay it. (a) Pore energy $E(r) = 2\pi\Gamma r$ for three combinations of line tension and membrane tension from the experimental literature; the barrier $E(r)$ ranges from ~ 150 to $\sim 6,000$ $k_B T$, which is why closed compartments persist, and why closure (which removes edge energy of $\sim 3,800$ $k_B T$ for a 100 nm disc) is spontaneous. (b) The topological scission barrier $4\pi|\bar{\kappa}|$ computed from measured bending rigidities ($\kappa = 11$ – 24 $k_B T$, $\bar{\kappa}$ from 0.4 to 1.5) against the mechanochemical budget of dynamin (10–50 GTP ~ 10 – $1,050$ $k_B T$). The two are the same order of magnitude: the price of changing the genus ledger and the enzyme’s wallet match.

5.3 Chains and their wrapping

For a closed DNA molecule, $Lk = Tw + Wr$ [41–42], where only Lk is topological and only topoisomerases may change it — type I in steps of 1, type II in steps of 2 [41–43]. Physiological superhelical density is $\sigma = (Lk - L_0)/L \approx -0.06$ [A]. The discretisation is the same one we met in the rotor: gene regulation inherits a quantised topological switch, and each notch costs ATP.

5.4 Networks, blind: a claim relocated by its own null

A first pass, from 12 pathway graphs, claimed that metabolic graphs are loop-rich and signalling graphs tree-like. A critical reader identified the flaw precisely: with one graph carrying the “tree” side of the dichotomy and a functional judgement (“those MAPK loops are cross-talk, not stoichiometric cycles”) doing the classifying,


the claim risked circularity. The repair is a blind analysis at scale, and it produced a result we did not expect.

We parsed 329 KGML pathway graphs from three organisms (human, *E. coli*, yeast) spanning three KEGG top-level classes — metabolism (MET), genetic information processing (GEN) and environmental information processing/signalling (ENV) — computed six topology features without reference to labels (V, E, loop density $/V$, algebraic connectivity, directed-cycle edge fraction, mean degree), and clustered blindly.

The class-level statistic survives, and survives strongly: directed-cycle edge fraction averages 0.23 in metabolism and 0.014 in signalling, with genetic-information graphs at essentially zero; loop density $/V$ is 0.86 for MET against 0.19–0.20 for GEN and ENV (Fig. 9a). Metabolic graphs are full of cycles; the other classes are not. The mechanism, moreover, is now a topological rather than functional criterion: directed cycles (feedback loops and reversible stoichiometric steps) are what metabolic graphs have and the others lack, which replaces the old post-hoc appeal to “cross-talk”.

The naive unsupervised clustering, however, **fails**. Average-linkage on standardised features against the three-class labels gives ARI 0 (indeed slightly negative at every k we tried), and the best three-cluster solution is degenerate: one singleton plus one cluster containing 259 metabolites, 4 genetic and 62 signalling graphs (Fig. 9b). In other words, “metabolic graphs are cycle-rich” is a real and large class-level effect, but it does not follow that topology alone partitions pathways into their functional classes — the feature space is multidimensional and the classes overlap in it. Against degree-preserving rewired nulls, real graphs carry *fewer* cycles than their nulls (median directed-cycle $z = 1.1$; $z = 2.6$), which is itself informative: pathway graphs are sparser and less connected than their degree-preserving randomisations, so their cycles are not an artefact of degree distribution.

We report this as the manuscript’s most instructive negative. The claim that survives is smaller than the first pass’s: topological features encode *how much cyclic self-maintenance a network contains*, not which functional class it belongs to. That is a claim about HC-3’s weakest clause, and it is exactly the kind of claim the hard core can carry.




KEGG graphs analysed blind. (a) Directed-cycle edge fraction against loop density, coloured by class after the fact: metabolic graphs (blue) separate strongly from signalling (red) and genetic-information (orange) graphs on both axes. (b) Blind clustering confusion matrix: the naive unsupervised solution does not recover the classes (ARI 0). (c) z-scores against degree-preserving and directed rewiring nulls: real graphs carry fewer cycles and lower connectivity than their nulls. The failure has a reading worth recording. The class statistic separates the categories on single coordinates (directed-cycle fraction, β_1/V), but blind clustering must use all six features jointly, and the classes overlap on the others: within-class heterogeneity in size and density exceeds the between-class separation in the remaining dimensions — a manifold-overlap problem, not a measurement failure. The corollary is methodological rather than embarrassing: recovering the dichotomy blind needs supervised or topology-aware feature engineering (projecting on the cycle statistic itself is the obvious first move), which turns the negative result into the design of the fair benchmark that P4 needs.

5.5 Typing real structures with persistent homology


The strongest form of Proposition T is a demonstration: compute for real structures and show that it *types* them. We analysed 19 experimental structures spanning three orders of magnitude in chain length (46–5,936 residues): small compact monomers (crambin, BPTI, ubiquitin), all- and / folds (myoglobin, haemoglobin, lysozyme, CheY, adenylate kinase, kinase domain), α -barrels (GFP), a structured RNA (tRNA^{Phe}), the 24-mer hollow sphere of ferritin (biological assembly), large cage assemblies (GroEL–GroES, the 20S proteasome) and the rotary motor (PDB identifiers in Table S1). Vietoris–Rips persistent homology was computed on C (or phosphate) point clouds, subsampled to 350–500 points, with filtration radii of 24 Å (40 Å for ferritin) [33].

Three results (Fig. 10, 11). **First, robustness.** Adding Gaussian noise with up to 3 Å to the C coordinates leaves the correlation of the Betti curves with the noise-free curves above 0.94 for every structure tested (Fig. 11c). The descriptor is among

the most noise-tolerant statistics computable from a point cloud — the objection is answerable in one plot. **Second, typing.** Unsupervised average-linkage clustering of the normalised Betti curves partitions the 19 structures into three classes with clean functional meaning: compact globular folds; hollow or anisotropic bodies (the ferritin sphere, the L-shaped tRNA); and large multi-subunit cages (proteasome, GroEL–GroES,). The separation is driven by loop content () and its characteristic scale, and by the presence of long-lived classes of tens of Ångströms in the cages. Ferritin is the clearest single case: one 172-residue chain shows almost no signal, while the assembled 24-mer shows a strong, long-lived second homology class — the invariant detects the *assembly*, not the chain, which is the cross-scale property Proposition T claims [A for the computation; C for the cluster count]. **Third, honest limits.** Feature magnitudes correlate with size (Spearman up to 0.9 for raw extrema), so the typing must be read as a shape claim after normalisation, not a size claim; the subsampling and filtration radius are modelling choices; and 19 structures demonstrate the principle rather than benchmark it. A pre-registered benchmark on a fixed PDB subset is the follow-up experiment this result earns (P4).



Persistent homology of four experimental structures (PDB; point clouds of C or phosphate atoms). Top row: Betti numbers (blue), (orange) and (red) against filtration radius . Myoglobin is compact and loop-poor; the GFP -barrel carries a strong class (the enclosed volume); the ferritin 24-mer shows the long-lived cavity of the hollow sphere; the 20S proteasome, a large barrel chamber, shows elevated loop content across scales. Bottom row: (orange) and (red) persistence barcodes for GFP and ferritin; long bars are the robust topological features that survive coordinate noise (Fig. 8c).



Topological typing is unsupervised, meaningful and robust. (a) Average-linkage clustering of the 19 structures by correlation-normalised Betti curves separates compact monomers from hollow bodies and from large cage assemblies (colour-coded clusters). (b) Loop content (area under the curve) against size (log residues), coloured by cluster: large cage assemblies (1PMA proteasome, 1OEL GroEL–GroES, 1BMF) separate from compact monomers at comparable and even larger sizes — the separation is shape, not size. (c) Robustness: correlation of the Betti curves with the noise-free curves after adding Gaussian noise (= 1–3 Å) to the coordinates; five noise replicates per point; all structures stay above 0.94 at 3 Å.

5.6 Proposition T

Proposition T. The structures that life depends on are classified by one shared set of topological invariants, with three distinguishable roles: the closed membrane (finite and positive) is the precondition of energy compartmentalisation; cyclic networks ($\neq 1$) are the seat of material and logical self-maintenance; and conserved linkages (Lk) are the topological packaging of information.

Falsification. A critical biological function that depends irreducibly on absolute geometric coordinates rather than adjacency and cannot be captured by β ; or a structural class whose invariant distribution is unrelated to function. Section 5.5 shows that the second clause is testable and has so far survived its first test.

6. The closed loop and its design space

6.1 Three couplings

The couplings C1–C3 are each graded **B** (inferences from established results): closure bounds the gradient (Sect. 3.1); ATP flux bounds fidelity (Sect. 4.2); the code's error-minimising architecture bounds the mutational cost of rebuilding

structure (Sect. 4.3). Together they say that a living system is a self-maintaining system satisfying C1–C3 with strictly positive entropy production, $\dot{S} > 0$, in Prigogine’s sense [44]. The loop is not merely verbal: it can be written as dynamics and pushed to its bifurcations.

6.2 A minimal loop with a threshold and a founder requirement

Consider the three-variable model

$$\begin{aligned}\dot{m} &= b i (1 - m) - w m, \\ \dot{e} &= \rho m (1 - e) - \ell e - \gamma e i, \\ \dot{i} &= a e i (1 - i) - \delta i,\end{aligned}$$

where m is membrane closure, e the normalised proton-motive force and i the fraction of maintenance machinery under coded control; b is the structure-energy coupling, w the leak, and the remaining parameters the rates of building, wearing, replicating and degrading. The couplings reproduce the loop: closure enables the gradient ($m \rightarrow e$), the gradient enables coding ($e \rightarrow i$), coding rebuilds closure ($i \rightarrow m$).

Steady-state continuation and stability analysis give the phase portrait (Fig. 12). The living state — finite closure, gradient and coding — exists only above a critical coupling $_SN \approx 0.31$ and is absent below it: a saddle-node bifurcation. The dead state is locally stable *everywhere*, including above the threshold. Two consequences follow, and both are substantive.


The first is that the origin of life is a **threshold, not a drift**: below $_SN$ there is no living attractor to approach, however long one waits. The second is the **founder requirement**: even above the threshold, the loop cannot ignite from infinitesimal perturbations, because the dead state is stable. Ignition requires a transient founder event large enough to cross the basin boundary. We quantified this with an ignition curve (Fig. 12d): a founder pulse of size ≈ 0.2 ignites the loop for any $_SN$, while a small founder (≈ 0.05) needs ≈ 0.9 — a scaling relation between founder size and environmental coupling. The parameterisation is illustrative [C]; the *structure* of the result — threshold plus rare event — is generic for any loop in which each variable’s production requires the others.

The founder requirement and the design table of the next subsection can now be stated in one currency. Give every variable a small basal synthesis leak:


$$\begin{aligned}\dot{m} &= b i (1 - m) - w m + \varepsilon, & \dot{e} &= \rho m (1 - e) - \ell e - \gamma e i + \varepsilon, & \dot{i} \\ &= a e i (1 - i) - \delta i + \varepsilon.\end{aligned}$$

At $\varepsilon = 0$ this is the model above, with its exactly absorbing dead state; at $\varepsilon > 0$ the dead state is replaced by a small positive steady state whose stability is not guaranteed — and the Michaelis–Menten engine of Sect. 6.3, whose basal enzymes produce each node at a small rate even without regulators, is exactly the $\varepsilon > 0$ regime sampled parameter-wise rather than fixed. Continuation gives a smooth, decreasing

critical-leak line (Fig. 13a): the dead state loses stability above $\rho = 0.35$, falling to $\rho = 2.00$. The reconciliation is itself a result: **the founder requirement is robust to basal chemical noise**. A prebiotic environment whose basal synthesis rates sit below the line still needs a discrete founder event; only above the line can genesis be spontaneous. The line quantifies how clean prebiotic chemistry must be before seeding becomes unnecessary, and it is P6's newest sub-clause: measure the basal reaction rates in a protocell platform and predict whether seeding is required.



The basal-leak phase boundary. (a) Critical basal-leak level against coupling for the minimal loop of Sect. 6.2: below the line the dead state is the only local attractor near the origin and genesis needs a founder event; above it, ignition from the near-dead state is spontaneous. The design-table engine (basal enzymes) operates in the $\rho > 0$ regime; the founder analysis operated at $\rho = 0$. (b) Median robustness Q of the three sign classes across the sweep $\{0, , , \}$: the ranking all-positive $>$ one-negative $>$ two-negative holds at every level, so the design-table conclusion does not depend on the basal-enzyme convention.



The closed loop as a dynamical system. (a) Ignition above the threshold ($\rho = 0.9$): from a small seed, closure m , gradient e and coding i converge to the living steady state. (b) Steady-state branches: the living branch (red, stable; grey, unstable) appears at a saddle-node bifurcation $\rho_{SN} \approx 0.31$ and vanishes below it; the dead state is stable for all. (c) Phase diagram in the coupling–leak plane: the living attractor exists only in the region of strong coupling and weak leak. (d) The founder requirement: ignition threshold ρ_{ign} against founder-pulse size. A founder of $s_0 = 0.2$ ignites at the bifurcation itself; small founders need coupling up to three times stronger. The origin of life is a threshold crossed by a rare event, not a gradual drift.

6.3 The design table: which loops can sustain themselves?

6.3.1 The question, made askable

Ma, Tang and colleagues turned “what network can adapt?” into an enumeration: all 16,038 three-node signed topologies, parameter sets each, functional region defined by two measurable quantities (sensitivity > 1 , precision > 10), and the result that every robustly adapting topology contains one of two motifs [1]. The parallel question here is *which three-node signed loops can sustain themselves?* — and the value of asking it is that the answer, whatever it is, tells us whether the C1C2C3 coupling is privileged, common, or trivial. An earlier formulation’s minimal ODE could not answer this; it could only show that *its own* loop bifurcates.

Function, operationalised — and split. An earlier formulation of the criterion folded two requirements into (i): “the all-low state is unstable *and* there is a stable all-high state”. A critical reader was right to call this a conflation: if all-low is unstable, ignition is spontaneous and no founder event is needed — which contradicts the absorbing-state analysis of Sect. 6.2, and the contradiction is between two of our own statements, not between us and a critic. The fix is to split the criterion and report the two functions separately. Terms, defined once: **self-maintenance** (persistence) = a stable all-high state ($m, e, i^* 0.1$) reachable from random initial conditions; **load-bearing** = at least one directed cycle whose breaking collapses the all-high state; **spontaneous genesis** = ignition from the near-dead state (0.01, 0.01, 0.01) — the design-table analogue of all-low losing stability; **self-repair** = recovery under knockdown (each variable lowered to 30% of its steady value recovers to within 80%). Robustness Q counts parameter sets (of 200 per topology) satisfying persistence and load-bearing; the spontaneous-genesis count is reported alongside rather than baked into the definition. **Knockdown protocol (registered correction).** Knockdowns are applied from the sampled true steady state — the mean of the random-IC endpoints. The first implementation applied them from a fixed out-of-range state, which is the artefact source behind its erroneous self-repair claims (Sect. 6.3.3); the correction is registered here before the results are read.

Method. All = 19,683 three-node signed topologies; Michaelis–Menten conversion dynamics as in Ma et al. (2009), with basal enzymes added exactly where a node lacks a regulator of one sign; edges with $k_{cat} [, 10]$ and $K_M [, 1]$ sampled log-uniformly; 200 parameter sets per topology; 5 initial conditions per parameter set; RK4 integration to $t = 50$ with convergence excision. The basal-leak parameter of Sect. 6.2 is swept through $\{0, 10^{-3}, 10^{-2}, 10^{-1}\}$, implemented as a basal production term $\mu_j(1-x_j)$ on every node: $\mu_j = 0$ reproduces the V3 engine exactly, and the sweep asks whether the ranking below is an artefact of the basal-enzyme convention. A convergence check (200 $\rightarrow$ 2,000 parameter sets on a random 50-topology subset) confirms the Q medians are stable at 200 (repository; Appendix A of the supplement). Because a topology with no directed cycle cannot satisfy (ii) by construction, the search space is the 19,574 cycle-containing topologies; we report the full set and, separately, the 9,936 containing a Hamiltonian three-cycle (the topologies commensurate with a three-layer loop).

6.3.2 What the enumeration found

Ninety-nine point eight per cent of Hamiltonian-cycle topologies (9,915 of 9,936) support a load-bearing self-maintaining state under at least one parameter set. **Self-maintenance is generic.** Any signed three-cycle, with almost any additional wiring, can be tuned to sustain an all-high state — which is exactly what a critical reader who called the loop model “a consistency check” suspected, and the enumeration settles the point against us. The hard core’s loop is *not* privileged by the mere fact of self-maintenance.

The split settles the founder question, and settles it in the theory's favour. Persistence is unchanged by the split — the full enumeration's 9,915 of 9,936 functional topologies qualify under either wording, and all 1,983 topologies of the stratified subsample remain functional — the subsample (286 all-positive, 664 one-negative, 1,033 two-negative, in proportion to the full table) is where the sweep and the criterion split are computed, while the full-table counts and medians quoted alongside come from the complete 200-parameter enumeration — while **spontaneous genesis is rare and leak-gated**: $= 0$: 26,802 of 396,600 parameter sets (6.8%) ignite from the near-dead state; $= 10^{-3}$: 7.3%; $= 10^{-2}$: 10.2%; $= 10^{-1}$: 22.4%. By sign class at $= 0$: 11.2% (all-positive), 10.2% (one-negative), 3.3% (two-negative); the all-positive class — the one the theory cares about — rises to 43.7% at $= 10^{-1}$, so even the most ignition-prone class remains founder-dependent at zero leak. The critical-leak line of Sect. 6.2 is exactly this boundary seen from the enumeration side: basal flux below $_c()$ leaves genesis a rare discrete event, and even a leak ten times the $_c$ scale leaves most of design space founder-dependent. The founder requirement is therefore *strengthened* by the enumeration, not contradicted by it: the same table that says self-maintenance is generic says that getting there from the dead state is a separate, rare event.


Robustness, however, is graded, and here the structure does bite. Ranking by the median Q of each sign class:

Robustness of self-maintenance by sign structure (9,936 Hamiltonian-cycle topologies, 200 parameter sets each). The claim is about the class, not an individual: the canonical C1C2C3 loop (+++) sits at $Q = 30$, below the median of its own class. The ranking is preserved at every basal-leak level $\{0, 10^{-3}, 10^{-2}, 10^{-1}\}$ on the stratified subsample (medians 65/40/19, 66/41/19, 69/46/24, 53/39/30 across the four levels — ordered throughout, compressing as the leak floods the design space), and preserved again at 2,000 parameter sets on the 50-topology convergence subset (Fig. 13b).

Sign class of the three-cycle	Functional topologies	Median Q	Reading
all-positive (+++)	1,431	63	most robust design
one negative (++)	3,321	40	intermediate
two negatives (+)	5,163	18	least robust of the three-cycles
two-cycle-only topologies	0	—	never load-bearing alone
self-loop-only topologies	0	—	never load-bearing alone

Two consequences. First, the *sign class* containing the framework's loop is the most robust class — and the claim must be stated at the class level, because the canonical loop itself sits at $Q = 30$, at the class median (63) rather than above it. The likely reading is that the canonical loop is the minimal three-edge member of its class, with no auxiliary edges whose parameter draws can add redundancy [C]; a design table in the strict sense is one in which the theory's own loop is an entry with

a rank, and the rank is *mid-class*. (The corrected-engine re-run moves the canonical loop to $Q = 31$ — one parameter set apart; the full-table value 30 is used throughout.) Second, the rank ordering is the *opposite* of the one adaptation produced, where incoherent feedforward loops beat negative feedback [1], and this difference is not an embarrassment but a definitional consequence: persistence (an on-state that survives perturbation and dilution) rewards positive reinforcement, whereas adaptation (returning to a fixed output under a persistent input) requires negative control. Frameworks that unify persistence and regulation must therefore expect at least two design classes, which is what the next subsection finds.



Design table for self-maintenance. (a) Number of functional topologies per sign class (log scale, bars) and median robustness Q (dashed line): all-positive cycles are both frequent and the most robust as a class; two-cycle and self-loop-only topologies are never load-bearing. (b) Distribution of $\log_{10} Q$ across functional topologies, with the canonical C1C2C3 loop marked: the loop sits at its class median ($Q = 30$ vs class median 63), which is the claim's honest position.

6.3.3 What the corrected engine taught us: self-repair is not separate — genesis is


An earlier run of the enumeration claimed, from the strictest criterion — knockdown self-repair — that persistence and homeostasis occupy different regions of design space, and that the canonical loop had *zero* recovering parameter sets. **The corrected engine overturns that claim, and we report the overturning rather than bury it.** With knockdowns taken from the true sampled steady states, the all-positive class repairs best of all: median self-repair counts Q_p of 56 (against persistence $Q = 65$), 25 ($Q = 40$) and 10 ($Q = 19$) for the all-positive, one-negative and two-negative classes — 86%, 63% and 53% of their own load-bearing parameter sets. A 30% knockdown of one variable of a mutually reinforcing loop stays inside the all-high basin, and the loop re-amplifies it: *positive feedback is what restores positive feedback*. What the first implementation measured was an artefact of testing recovery from an out-of-range state.

What survives — and it is the part that matters — is a separation with a firmer boundary. Genesis is separated from persistence by the $_c()$ line: spontaneous ignition is rare (6.8% of parameter sets at $_c = 0$) even where persistence and self-

repair are generic. And adaptation — returning a *regulated output* to a set-point against a persistent input, Ma et al.’s own criterion — is a genuinely different design problem with a genuinely different sign requirement [1]. Persistence, self-repair and adaptation are three functions, not two, and only the first two share a sign preference. This is an old idea with a modern quantitative form — the interlinked fast-and-slow positive and negative feedback loops that make cell-fate decisions reliable [2,45–46] — and the corrected enumeration places the present framework inside it rather than beside it. HC-4 is a claim about persistence, and it should be stated as such. Two predictions follow (P7 and P8 in Sect. 9): four-node tables should show that adaptation co-localises with negative feedback while persistence and repair co-localise with positive reinforcement, and the mixed topologies should win only on the adaptation criterion. One boundary condition on everything in this subsection: it is a three-node table. Design-space volume grows super-exponentially with node count, and nothing here rules out a topological phase change at four nodes — which is precisely why P7 sits at the front of the positive heuristic rather than in the discussion.

6.4 The basin surface and the bifurcation

For the canonical loop we computed the global basin of the live state on a 1 grid of initial conditions (basin_surface.json), and the picture is a clean slice of the design table: the live basin fills the large-, small-leak corner, its volume falls smoothly to zero as $_SN \rightarrow 0.31$, and the dead state is stable everywhere — there is no parameter region where an infinitesimal perturbation revives the loop. A saddle-node normal-form check confirms the bifurcation is of saddle-node type (the Jacobian has a simple zero eigenvalue with non-degeneracy conditions satisfied), so the threshold has the standard square-root scaling in the distance to threshold. The ignition curve alone gives a one-parameter slice; the basin surface gives its geometry, and the geometry is what the founder experiment (Sect. 7) works on.



Basin volume of the live attractor on a grid of initial conditions. Contours mark $B = 0.1 \dots 0.9$; the live basin exists only for strong coupling and weak leak, and vanishes at the saddle-node.

7. Genesis: absorbing dead states, pinned founders, and the company we keep

7.1 The dead state is exactly absorbing

The minimal loop has a property stated but not exploited in earlier formulations: the dead state is not merely attracting, it is **absorbing**. Both the coding variable and the closure variable are produced in proportion to each other ($d_i = e_i, d_m = i$), so from exactly $i = 0$ no trajectory ever leaves. A pulse in the gradient field alone cannot ignite the loop; genesis requires a founder event that seeds all three variables *simultaneously*. This is the most literal possible reading of “founder event”, and it has a clean experimental face: protocells must start with membrane, gradient and coded catalysis already co-located, not with any one of the three. The statement is exact at zero basal flux, and the critical-leak line $\ell_c()$ of Sect. 6.2 bounds it: below the line the dead state remains the only attractor near the origin and the founder requirement stands; above it ignition is spontaneous. Clean prebiotic chemistry keeps the founder requirement; dirty chemistry removes it. Both regimes

are measurable (P6), which is what makes the boundary a prediction rather than a patch.


7.2 Founder patches: pinning without dispersal, invasion with it

We put the founder on a 4040 lattice: a disc of radius r is initialised with $(m, e, i) = 1$, the gradient diffuses ($D = 0.05$) to its neighbours, and the system relaxes for 400 time units. Results (Fig. 16):

1. **Viability is generic.** Every founder disc with $r \geq 2$ and $\gamma \geq 0.2$, at every coupling tested ($\gamma = 0.9, 1.2, 1.5$), settles into a persistent live island. There is no minimum radius for survival.
2. **The islands are pinned, not invasive.** The live area at $t = 400$ equals the founder area to within a few per cent across all conditions. No founder ever expands — the live region neither grows nor shrinks. The boundary is stabilised by diffusive dilution: at the rim the gradient leaks away faster than closure can sustain it, so e falls below the ignition value e_a and the front stalls. Invasion of the dead environment would require the interface condition $e_{\text{rim}} > e_a$, which the present parameters never meet.
3. **The founder size sets the surviving structure's size, linearly — without dispersal.** Because islands are pinned, founder radius maps one-to-one onto the scale of the living compartment — which is the spatial content of P6: the protocell radius R_c is not a critical nucleus below which life fails, but the *size of the compartment that will persist*. This relation is a property of the no-dispersal regime and is broken by dispersal (point 5).
4. **Pinning is structural, not parametric.** A 1-D reaction–diffusion version of the same model (only the gradient e diffuses, $D = 0.05$, as on the lattice; m and i immobile) confirms the diagnosis and sharpens it: the front stays fixed at every coupling tested, including $\gamma = 4.0$ — five times the bifurcation value. The mechanism is compositional, not a matter of thresholds almost met: ignition requires m, e and i co-located, but only e leaks into the dead neighbourhood, and the decisive variable i is immobile *and* exactly absorbing at zero — so the leaked gradient finds nothing to catalyse. “Live = founder” is a theorem of this model class, not a coincidence of parameters, and no parameter scan within the class can change it.
5. **Dispersal unlocks invasion.** Give the coding variable a non-zero mobility ($\gamma = 0.02$, vesicle-like: division deposits diluted offspring containing all three variables). The front unlocks: $v = 0$ at $\gamma = 0.6$ (where the founder itself cannot persist — consistent with the extinction phase of the lattice diagram), $v = +0.059$ at $\gamma = 0.9$, $+0.074$ at $\gamma = 1.2$, $+0.086$ at $\gamma = 1.5$, rising monotonically with coupling (Fig. 17b). On the lattice, the same effect appears as a three-phase diagram in (γ, p_{div}) (Fig. 17a): below the persistence threshold the patch goes extinct regardless of dispersal ($\gamma = 0.6$); above it, $p_{\text{div}} = 0$ pins the island exactly as in the no-dispersal case, while division probabilities above 0.02 – 0.05 produce genuine invasion, with live area growing monotonically in both γ and p_{div} . The invasion boundary is the spatial face of the ignition curve: coupling

sets whether the loop can run; dispersal sets whether it can spread. What was a liability without dispersal is now a controlled, falsifiable contrast — protocell platforms with and without a coupled growth–division module sit on opposite sides of it (P6).

One more reading deserves its sentence. In the hydrothermal-vent picture, compartment walls partition m and i while the pore network transports what leaks [48–50] — the pinning result is, in that light, an argument *for* compartmentalised origins: in free solution the same diffusive dilution that pins the front would simply carry the parts away. The lattice says something a beaker cannot: without a dispersal mechanism, chemistry that is alive locally stays where it was born.




Founder patches and the invasion boundary. (a) Three-phase diagram in the coupling–division plane (4040 lattice; budding deposits a diluted offspring seed when local maturity exceeds threshold): extinct ($\rho = 0.6$), pinned (low ρ), invading (high ρ); the no-dispersal run’s all-pinned result is the $\rho = 0$ row. (b) 1-D front velocity against coupling with the coding variable immobile (all pinned, including $\rho = 4.0$) and with $\rho = 0.02$ (invasion above $\rho = 0.7$; speed rises monotonically with coupling). Pinning is structural; dispersal unlocks it.

7.3 Self-organised criticality: the wrong shape of explanation, and we say so

Self-organised criticality [47] offers the most serious competing picture of genesis: a near-critical system needs no founder event, because perturbations propagate as scale-free avalanches with no characteristic size. The framework’s picture predicts the opposite — a characteristic threshold, and sub-threshold perturbations that die.

The spatial model decides for the framework, and for a blunt reason. Because the dead state is absorbing, sub-threshold gradient pulses produce *no* avalanches at all: we simulated 400 pulses per condition at three couplings and four pulse amplitudes and recorded zero ignition events from pulses alone at every setting. The avalanche-size distribution is not a power law with an exponent to be fitted; it is

empty below the founder threshold and finite above it. Self-organised criticality is not merely unsupported here — it cannot describe *this model class*: with an absorbing dead state, zero ignitions from sub-threshold pulses is close to a definitional result, not an independent verdict. The honest statement is narrower and stronger at the same time: the framework predicts a characteristic threshold where SOC predicts scale-free avalanches, and the discriminating observable is the sub-threshold fluctuation size distribution (exponential truncation versus power law), which is what P6b registers. Absorbing states are, of course, the heart of the modern SOC literature — absorbing-state phase transitions and directed percolation [66] — and a driven version of this model (basal leak above the critical line of Sect. 6.2, or stochastic re-seeding) sits precisely in that universality question; the framework’s claim is about genesis from the clean dead state, not about driven absorbing-state systems in general. The test that would distinguish the two pictures in a real protocell (the size distribution of sub-threshold fluctuations: exponential truncation versus power law) is stated as P6b.



Founder patches on a 4040 lattice. (a) Outcome class versus coupling and founder radius: viable (pinned) islands for every founder disc $r \geq 2$ at every coupling, no expansion anywhere. (b) Live area at $t = 400$ against founder area across couplings: islands are pinned (live = founder), and founder size sets the size of the persistent compartment.

7.4 Origins: the loop in diachrony

The synchronic loop says how life persists. The diachronic question — how the loop came to be closed — is where an earlier formulation of this framework was silent, and we address it here as conjecture with testable components (graded B/C/D, not A).

(i) Closure first, and cheaply. Amphiphiles in water close spontaneously; closure removes edge energy (Sect. 2.1). The structure layer comes first, not by selection but by free-energy minimisation [A/B] — Lane and Martin’s point [48–50]: geochemically maintained gradients in alkaline hydrothermal compartments make the first step nearly free.

(ii) A gradient, and its leak. A closed compartment holding an abiotic or photochemical proton source can maintain a transient gradient; leak limits it (Sect. 2.1). This stage is chemistry requiring only closure plus a source [A].

(iii) Coding, and the founder barrier. For coding to begin, replication must beat the error threshold (Sect. 4.4) and the loop must be ignited (Sect. 6.2). Viroids show that an error rate of 2.5 per site is survivable only by informationally trivial genomes [28]; ribozyme estimates give a similar message [51]. The earliest replicons were therefore small and error-prone, and the transition to large genomes required fidelity — required the loop. The founder effect then implies that fidelity could not have been assembled gradually from the dead state. On this account the origin of life is a **threshold crossed by a rare founder event**, and the two conditions are coupled by an ignition curve (Fig. 12d) that protocell experiments could, in principle, measure [B/C]. Protocell systems with coupled growth and division already exist as experimental platforms [52–53], and the ignition curve is, to our knowledge, a novel quantitative target for them (P6).

(iv) The triple unity as an order parameter. In the model, “aliveness” is not a variable but a *phase*: the living state is an attractor of the coupled dynamics, distinguished from the dead state by finite values of all three order parameters. This reframing — life as a phase of the loop, not a substance or a list — is the interpretive payload of the framework, and we grade it **D**.

7.5 The company we keep: a repaired genealogy

An earlier comparative table omitted the tradition that anticipated the loop most directly, and a critical reader noted the gap by name. The genealogy, stated properly:

- **Autopoiesis** [54–55]: a living system is a network that continuously produces its own components, with the boundary as a product of the network. The present framework’s closure term m is an autopoietic reading made dynamical; the difference is quantifiable — the framework’s version of “self-producing” is a bifurcation condition, not a definition.
- **The chemoton** [56]: three coupled subsystems — metabolism, template replication, membrane — each a cycle. This is, almost line for line, the triple unity; the present framework adds the currencies (kJ mo, bit, k_BT) and the couplings C1–C3, and now a design table.
- **(M,R)-systems** [57]: metabolism and repair, each requiring the other, with the category-theoretic machinery of realisation. The present loop’s C2/C3 pair is the quantitative shadow of Rosen’s (M,R) interdependence.
- **Dissipative structures** [44]: the positive-entropy-production criterion is adopted wholesale, and now attached to a bifurcation structure.
- **Free-energy principle** [58], **integrated information theory** [59], **metabolic theory of ecology** [60], **network motifs** [61]: the comparative accounting in Sect. 10, Table 4.

The repair is genealogical honesty, not appropriation: autopoiesis and the chemoton said the qualitative thing first; the claim here is the quantitative one, and it is falsifiable, which is the only kind of claim a hard core is entitled to.

7.6 The loop applied to a second transition: eukaryogenesis

One question pressed harder than any other: what does the framework do besides the origin? There is a second ignition with far better historical constraints — eukaryogenesis — and the framework's three layers bite on it simultaneously. The case is cheap to run because it carries no new theory: every load-bearing element is published (the Lane–Martin energy budget [48,49,62], CoRR [63], the Moser–Dutton tunnelling ruler [11,12]), and the discipline is the one Section 4.3 established for the genetic code: name the alternative hypothesis and keep it in view [65].

Energy: the scaling flip [B]. Section 3.4's ceiling locks prokaryotic genomes: with the plasma membrane as the only coupling membrane, ATP supply scales as r^2 while maintenance scales as r^3 , so genome size has a hard ceiling of order 10^6 genes — exactly where the α -proteobacterial ancestor sits [62]. Internalised, replicating coupling membranes break the scaling: mitochondrial number, and with it inner-membrane area, grows with volume — area with r^2 [2, 3] for cristae folding, a directly testable Cryo-ET number — so supply and maintenance scale together and the ceiling lifts; eukaryotic genome sizes exceed the prokaryotic range by one to two orders of magnitude, matching the 15–30-fold ATP gain. The uniqueness claim must be stated carefully: **among the solutions that respect the scaling constraint, internalised replicating membranes are the only class observed or proposed.** Autogenous alternatives do not change the surface-to-volume geometry; the giant bacteria that evade the ceiling do it by membrane folding and polyploidy — paying exactly the cost the constraint demands, without the replicating organellar membrane.

Information: the non-transferable set [B]. Endosymbiotic gene transfer cleared 99.6% of the partner genome (4,000–13 genes). What stays is not historical accident: a gene whose product works away from the membrane transfers readily; a gene for a subunit buried in the inner membrane does not, because electron-tunnelling rates decay by roughly an order of magnitude per Ångström [12], so remote synthesis fails; CoRR adds the temporal half — redox-coupled in-place expression answers in seconds, gene-transfer regulation in minutes [63]. The two-genome chimaera is a partition drawn by the physics of the membrane, not a residue.

History: one event, necessarily fixed [C for the model, B for the verdict]. Place the partnership on a coupling–symmetry continuum (,); the merged state is an absorbing corner [C]. Eukaryogenesis requires a conjunction — internalisation, metabolic complementation, an EGT mechanism, an energy feedback — each individually rare. The historical record constrains the compound probability from both sides: endosymbiosis itself is frequent (plastids arose several times) while eukaryogenesis happened once, which rules out both a high and a negligible

compound probability — an interval statement, not a point estimate. We can measure neither the attempt rate nor the per-attempt success, and the anthropic caveat stands: the question is only asked in a universe where the event happened. Once achieved, the energy gain drives the fixation probability toward one [64]; after gene transfer, reversal requires rebuilding a lost genome while forfeiting the energy gain — an absorbing state, Dollo-irreversible. Later symbioses add modules to an existing platform and cannot restart the merge: priority closes the door.

Eukaryogenesis by step and verdict. The structure is the framework’s own — threshold, rare founder, absorbing lock — applied to a second ignition whose historical constraints are independent of the first’s.


Step	Verdict	Basis
Endosymbiosis as soil	necessary	prokaryotic symbiosis is frequent (plastids, several origins) [A]
The specific pairing	contingent	the conjunction is rare; the record brackets its probability between strict bounds [B]
Fixation	necessary	energy gain drives 1 [64] [B]
Irreversibility	necessary	CoRR lock + absorbing state + Dollo [B]
Uniqueness of the path	semi-necessary	scaling-decoupling solution unique in its class + priority effect [B]

The event’s occurrence is a rare realisation (contingent); its path is confined to one class (semi-necessary); its fixation and irreversibility are forced (necessary). That is the framework doing work on a problem it did not set, with every load-bearing element published and the alternatives named — which is what makes it a programme rather than an origin story.

8. The loss of the loop: fragmentation and

If the loop is what holds a living system together, its local failure should be visible in the structure layer’s connectivity. Mitochondrial networks in aged cells fragment — the reticulum breaks into discontinuous pieces — and percolation-like transitions have been measured in mitochondrial networks [67–68]. As a minimal model, we pruned edges at random from a 28 × 20 lattice and computed the algebraic connectivity [69] and the probability of connectivity (Fig. 18). Both decline smoothly at first and then collapse: falls by half with ~8% of edges pruned and reaches zero within a few percent of additional pruning, while connectivity is lost well below the infinite-lattice bond-percolation threshold because the finite network has no infinite cluster to fall back on. The practical reading is that *p* is a *fragile* order parameter at fragmentation — small changes in connectivity produce large changes in *p* — which is precisely the property a health monitor should have. The numbers here are model results [C], and they are not a small-grid artefact: on a 56x40 lattice the normalised curve $\langle p \rangle / \langle p \rangle_0$ tracks the 28x20 result, the half-collapse point moving from $p = 0.24$ to $p = 0.20$. On a 112x80 lattice full connectivity is lost

already by $p = 0.05\text{--}0.10$: with $\sim 9,000$ cells, any single cell whose four edges are pruned isolates the graph — the finite-size face of the same fragility that makes a good early-warning monitor. The empirical test (against age or mortality slope across species) remains open (P8).



Fragmentation collapses algebraic connectivity. (a) of a 28 20 lattice under random edge pruning (mean SD over 20 replicates). halves by $p = 0.08$ and collapses to zero by $p = 0.2$; the dashed line marks the bond-percolation threshold of the infinite square lattice. (b) Probability that the pruned network remains connected. The collapse of precedes and predicts the loss of connectivity, which is the property wanted of a network health monitor in the mitochondrial-fragmentation prediction P8.

9. Predictions, with protocols

Predictions with status and protocols. P2, P5 and P6 carry pre-registered designs in the supplement; P7–P10 are new in this version and are deliberately “surprising” in the sense the programme needs — each could embarrass the framework.

#	Prediction	Status	Protocol / test
P1	Autonomous-metabolism organisms use F/A/V synthases and gradients inside the c-ring window; real rotor/gradient pairs sit near the floor (mean percentile of null < 25 under every registered null).	AB ; verified here (12.5th percentile; robust across N0–N3, means 12.5–19.03); extension open.	Add c-rings and in-vivo p as available; recompute under the frozen protocol of Appendix A. Failure if the percentile drifts to 50 under every null.
P2	Perturbing supercomplex spacing d lowers OCR as $\exp(-d)$, $1.1 \text{ \AA}$.	Open.	Cryo-ET spacing + Seahorse OCR on assembly mutants; pre-registered slope test (supplement).
P3	Variant codes in the best $\sim 0.01\%$ of block-preserving	A , verified for 8 codes.	Extend to all NCBI tables under the

	random codes.		stated metric.
P4	Topological typing separates structures into functionally meaningful classes at scale; the three-class structure survives a fair benchmark.	A/C , demonstrated at $n = 19$; benchmark protocol pre-registered.	CATH-stratified 500-structure sample; cluster number fixed a priori; noise threshold 0.9.
P5	Cross-species metabolic-rate residuals correlate with proofreading/genome-maintenance cost.	Open; confounded.	Paired comparative dataset; controls for size, temperature, activity. Programme, not claim.
P6	Protocell founders: viability is generic above co-location of (m,e,i); without dispersal, compartments are pinned and founder size sets compartment size (R_c mapping); no sub-threshold avalanches (exponential truncation, not power law); with a growth-division module, invasion resumes above a critical coupling and its speed rises with coupling.	C , modelled here; now with a spatial form and a dispersal contrast.	GUV/protocell platforms [53]; measure founder radius vs surviving compartment, the sub-threshold size distribution, and spatial spread with vs without a coupled growth-division module.
P7	Four-node design tables: adaptation co-localises with negative feedback, while persistence and knockdown-repair co-localise with positive reinforcement; mixed topologies win only on the adaptation criterion.	New, model-level (C); the persistence/repair half is settled at three nodes (Sect. 6.3.3).	Extend the enumeration to four nodes; RNN-searched kinetics as an independent generator [4].
P8	Senescent mitochondrial networks lose before losing connectivity; collapse precedes function loss.	New, B/C .	Time-resolved imaging + TDA; and as leading indicators.
P9	Topologically defective membranes (high genus) sustain gradients for shorter times than topologically trivial ones of equal composition.	New, C .	Microfluidic protocells with tunable topology; pH/voltage probes.
P10	Under proliferation selection, cancer cells' codon-usage bias shifts away from the usage	New, B .	TCGA pan-cancer vs matched normal: cosine distance

pattern co-initial with the standard code's error-minimising structure. (M_s is a property of the code table, which cancers share; codon-*usage bias* is the variable that can move.)

between the usage-bias vector and the M_s -optimal direction, with expression-level and GC-content controls.

P11 Synthetic reconstruction of endosymbiosis yields genome-capacity expansion only when scaling decoupling, EGT and in-place redox control are jointly present (single conditions insufficient); any half-eukaryotic intermediate shows an energy budget with transition scaling between and ; across species, the retained mitochondrial gene set correlates with the fraction of respiratory-chain subunits buried in the inner membrane.

New, **B/C**.

Synthetic-biology reconstruction protocols; Cryo-ET morphometry of existing eukaryotic cristae () and of candidate intermediates; cross-species mitochondrial gene-content regression [12,63].

10. The programme: comparative accounting and next steps

10.1 Against the neighbours

Comparative accounting. The framework is not offered as a replacement for any of these; it is offered as the quantitative joint at which three of them meet.

Framework	Core claim	Where this framework is stronger	Where it is weaker
Free-energy principle [58]	Life minimises variational free energy	Concrete mechanisms (chemiosmosis, code, topology) and measurable currencies	FEP is more general and covers cognition
Integrated information theory [59]	Consciousness is integrated information	Grounds the information layer physically	IIT addresses experience; this framework does not
Self-organised criticality [47]	Life sits at a critical point without tuning	Gives thresholds and design tables	SOC is simpler and, where applicable, parameter-free
Metabolic theory of	Metabolic rate sets ecological	Places MTE inside the	MTE is empirically

ecology [60]	scale	mechanism layer	deeper on ecology
Autopoiesis / chemoton [54,56]	Self-production defines life	Quantified, falsifiable, computational	Those traditions are conceptually richer about organismic identity

10.2 The positive heuristic, ranked

Near term (computational). Complete the four-node design table (P7) with both MM kinetics and RNN-searched generators; run the CATH-stratified TDA benchmark (P4); extend the margin statistics (P1). These need no wet lab and each produces a citable result.

Mid term (existing data). Time-resolved mitochondrial TDA (P8) on published imaging; TCGA codon-usage analysis (P10); the nidovirus genome-window check from Sect. 4.4; cristae-fractal dimension morphometry from existing Cryo-ET volumes — the eukaryogenesis anchor of Sect. 7.6 is measurable today. These convert open archives into tests, which is the discipline this framework has adopted throughout.

Far term (wet lab, collaborations needed). P2 (Cryo-ET + respirometry), P6 (protocell founder protocol), P9 (topology-defective membranes). Each has a pre-registered design in the supplement and each is a genuine collaboration invitation.

10.3 What would falsify the hard core

Worth stating plainly, in the spirit of Lakatos’s own criterion: repeated, non-ad-hoc failure across the protective belt would force surrender of the adjacent proposition. Specifically — HC-2 dies if a natural code is found far outside the null’s extreme tail under every reasonable metric; HC-3 dies if the design table at four nodes shows that the loop’s functions are insensitive to topology, or if the TDA benchmark dissolves its classes; HC-4 dies if the loop’s predictions turn out to be reproducible from the mechanism layer alone at equal precision (the “one layer is enough” test); HC-1 requires the discovery of a genuinely non-chemiosmotic free-living metabolism.

11. Evidence discipline

Evidence grading of the main claims. The framework’s honesty is the claim: no A-level statement rests on a model, and no model is asked to carry more than its grade.

Layer	Claim	Grade	Basis
Mechanism	E1: closure precondition for a gradient	B	Elliptic regularity + edge-leak physics
Mechanism	Coupling-ratio window (Fig. 2, Table 1)	A	Thermodynamics + measured p and c-rings

Mechanism	14 Å cofactor ceiling; 1.1 Å	A	Langen et al. 1995; Page et al. 1999
Mechanism	F/A/V graded homology (Fig. 5)	A	UniProt sequences, this paper
Symbol	Code in best 0.007% of null (Fig. 6)	A	-code null, this paper
Symbol	Error threshold and real-genome margins (Fig. 6)	A/B	Eigen theory + Wright-Fisher + literature rates
Symbol	Freezing barrier of codon reassignment	B	Quasispecies arithmetic
Structure	Gauss-Bonnet ledger and scission budget (Fig. 8)	A/B	Measured moduli + Helfrich theory
Structure	Topological typing of 19 structures (Figs. 10–11)	A/C	This paper (cluster count is C)
Structure	Metabolic loop richness (Fig. 9)	A	KEGG graphs, this paper
Loop	C1–C3 couplings	B	Established results
Loop	Threshold-plus-founder origin dynamics (Fig. 12)	C/D	Minimal model; parameters illustrative
Loop	Basal-leak boundary $\varepsilon_c(\rho)$; persistence vs spontaneous genesis (Fig. 13)	C	Minimal model + corrected enumeration
Loop	Dispersal-unlocked invasion boundary (Fig. 17)	C	Lattice + 1-D front model
Loop	Life-as-phase reframing	D	Interpretation

Three boundaries deserve explicit statement. **What the framework does not claim:** it does not derive the genetic code from physics, nor explain why *this* table was fixed; it carries the code’s robustness, not its biography. **What the models are:** the loop model, the fragmentation model and the error-threshold simulation are C-grade — their *structure* (bifurcation, threshold-plus-founder, collapse) is the claim, their parameters placeholders. **What remains genuinely open:** the in-situ measurement of P2, the comparative test of P5, and the protocell test of P6.

12. Conclusion

The framework is now two things it was not. It is *structured*: a declared hard core, a graded protective belt, and a ranked heuristic whose first three entries need no new experiments and whose next three need only existing data. And it is *embarrassable*: the design table says self-maintenance is generic and locates the loop’s value in robustness ranking, not existence; the KEGG blind analysis says the loops-versus-trees claim survives as a statistic but not as a classifier; the margin analysis says the mammalian rotor runs so close to the thermodynamic floor that

the floor and the operating point are indistinguishable; the founder experiments say the dead state is exactly absorbing, that genesis needs a founder event below a measurable critical basal leak, and that spatial spread is blocked by the model's own composition — until dispersal is added, whereupon invasion resumes above a critical coupling. Each of these is a place where the framework could have been wrong and is not, and three of them (the closeness of the floor, the failure of blind classification, and the corrected engine's overturning of the earlier self-repair claim) are places where the framework's own earlier claims were too strong or its instruments too loose. And the loop does work on a second ignition — eukaryogenesis — where the historical constraints are independent of the first's. That is what a research programme looks like in its second decade of ambition and its first year of data.

The claim remains the one we began with: not that life is just physics, but that the physics is specific enough to fail.

Data and code availability

The complete package (analysis code, archived raw inputs, processed results, all 329 KGML graphs with a machine-readable inventory CSV, figures and build scripts) accompanies this manuscript and has been archived with a Zenodo DOI 10.5280/zenodo.23154575; the repository URL is <https://github.com/zhangsongbo2026/The-triple-unity-of-life>. Supplementary Table S2 lists a representative subset of the KEGG graphs; the complete 329-graph inventory (organism, class, size and topological features per graph) is the file `data/kgml_inventory.csv` in the data package.

All inputs are public archives: PDB accessions (Table S1), UniProt accessions (Table S1) and KEGG pathway identifiers (Table S2). The complete pipeline — retrieval, analysis, figures and manuscript build — is provided as deterministic Python scripts and archived inputs, released on GitHub (repository URL in the cover letter). All 105 primary references were verified against CrossRef before citation; the verification record (DOI, year, journal for every entry) ships with the data package.

Table S1. Experimental structures and sequences analysed. *Structures (PDB)*: 1MBN (myoglobin), 4HHB (haemoglobin), 1UBQ (ubiquitin), 2GB1 (protein G B1), 1AKI (lysozyme), 1TIM (triosephosphate isomerase), 1GFL (GFP), 1FHA (ferritin, biological assembly), 1PMA (20S proteasome), 1OEL (GroEL–GroES), 1BMF (-ATPase), 1AZU (azurin), 1EHZ (tRNA^{Phe}), 1R69 (repressor), 3CHY (CheY), 1AKE (adenylate kinase), 1BPI (BPTI), 1FMK (Src kinase), 1CRN (crambin). *Sequences (UniProt)*: P0ABB0 / P0ABB4 (*E. coli* /), P37808 / P37809 (*B. subtilis* /), P25705 / P06576 (human mitochondrial /), P07251 / P00830 (yeast mitochondrial /), Q57670 / Q57669 (*M. jannaschii* A-type /), Q60186 / Q60187 (*M. mazei* A-type /), P38606 / P21281 (human V-type A/B), P17255 / P16165 (yeast V-type Vma1/Vma2).

Table S2. KEGG pathway graphs analysed (KGML): eco00010, eco00020, eco02020; hsa00010, hsa00020, hsa00190, hsa04010; sce00010, sce00020, sce00190, sce04011. Graphs with fewer than five parsed nodes after reduction (the oxidative-phosphorylation KGMLs, which are reference maps rather than stoichiometric charts) were excluded and are listed here for completeness.

Acknowledgements

The maintainers of the PDB, UniProt, KEGG and CrossRef open archives made a framework of this kind testable; the readers of earlier drafts made it honest.

Author contributions (CRediT). S. Zhang: conceptualisation, methodology, software, formal analysis, investigation, data curation, and J.Duan: writing — original draft, writing — review and editing.

Declaration of interests. The author declares no competing interests.

Appendix A. Frozen protocol for the margin statistic (P1)

Registered before the present recomputation, in the sense of the pre-registration discipline used throughout this paper: the null models, thresholds and inclusion rules below were fixed before the recomputation and are not adjusted by its outcome.

- **Nulls.** N0: $c \sim U(4, 20)$, $p \sim U(150, 250)$ mV, $G \sim U(50, 65)$ (cytosolic) / $U(46, 58)$ (matrix) kJ mo. N1: c from the empirical set of resolved c-rings {8, 10, 11, 13, 14, 15} (uniform). N2: p log-uniform over [150, 250] mV. N3: clade-stratified c (mitochondrial {8, 10}, bacterial {10, 11, 15}, chloroplast {13, 14, 15}) with clade G bands; the percentile is computed within clade.
- **Inclusion.** A system enters the statistic only when both a resolved c-ring stoichiometry and a measured p exist; one-quantity systems are listed in Table 1 and used only for the selection-effect bound.
- **Decision rule.** Retain “paired near the floor” at grade A/B if the mean percentile is < 25 under every null; report as “suggestive, grade B” if any null exceeds 35; between 25 and 35, report both numbers and demote to B. Every report carries a bootstrap 95% CI; if the CI contains 50, the text must state that the achieved sample size cannot exclude neutrality.
- **Update rule.** Future (c , p) pairs enter under this protocol unchanged; any protocol change requires a new version note.
- **Outcome of the present analysis.** Retained: N0 12.48, N1 18.15, N2 12.66, N3 19.03; CI [3.6, 21.4]; selection bound 21.9 (two missing systems at the null median); Bayes factor 19–78.

Funding. This work received no external funding.

References

1. Ma W., Trusina A., El-Samad H., Lim W. A., Tang C. (2009). Defining Network Topologies that Can Achieve Biochemical Adaptation. *Cell* 138, 760-773. <https://doi.org/10.1016/j.cell.2009.06.013>
2. Li F., Long T., Lu Y., Ouyang Q., Tang C. (2004). The yeast cell-cycle network is robustly designed. *Proceedings of the National Academy of Sciences* 101, 4781-4786. <https://doi.org/10.1073/pnas.0305937101>
3. Yang X., Lau K. Y., Sevim V., Tang C. (2013). Design Principles of the Yeast G1/S Switch. *PLoS Biology* 11, e1001673. <https://doi.org/10.1371/journal.pbio.1001673>
4. Shen J., Liu F., Tu Y., Tang C. (2021). Finding gene network topologies for given biological function with recurrent neural network. *Nature Communications* 12, 3125. <https://doi.org/10.1038/s41467-021-23420-5>
5. Kuhn T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press
6. Lakatos I. (1970). Falsification and the methodology of scientific research programmes. *Criticism and the Growth of Knowledge* (Lakatos I, Musgrave A, eds), Cambridge University Press, 91-196
7. Dimova R., Stano P., Marques C. M., Walde P. (2019). Preparation methods for giant unilamellar vesicles. *The Giant Vesicle Book*, 3-20. <https://doi.org/10.1201/9781315152516-1>
8. Watt I. N., Montgomery M. G., Runswick M. J., Leslie A. G. W., Walker J. E. (2010). Bioenergetic cost of making an adenosine triphosphate molecule in animal mitochondria. *Proceedings of the National Academy of Sciences* 107, 16823-16827. <https://doi.org/10.1073/pnas.1011099107>
9. Rolfe D. F. S., Brand M. D. (1997). The Physiological Significance of Mitochondrial Proton Leak in Animal Cells and Tissues. *Bioscience Reports* 17, 9-16. <https://doi.org/10.1023/a:1027327015957>
10. Langen R., Chang I. J., Germanas J. P., Richards J. H., Winkler J. R., Gray H. B. (1995). Electron Tunneling in Proteins: Coupling Through a Strand. *Science* 268, 1733-1735. <https://doi.org/10.1126/science.7792598>
11. Moser C. C., Keske J. M., Warncke K., Farid R. S., Dutton P. L. (1992). Nature of biological electron transfer. *Nature* 355, 796-802. <https://doi.org/10.1038/355796a0>
12. Page C. C., Moser C. C., Chen X., Dutton P. L. (1999). Natural engineering principles of electron tunnelling in biological oxidation-reduction. *Nature* 402, 47-52. <https://doi.org/10.1038/46972>

13. Moser C. C., Farid T. A., Chobot S. E., Dutton P. L. (2006). Electron tunneling chains of mitochondria. *Biochimica et Biophysica Acta (BBA) - Bioenergetics* 1757, 1096-1109. <https://doi.org/10.1016/j.bbabbio.2006.04.015>
14. West G. B., Brown J. H., Enquist B. J. (1997). A General Model for the Origin of Allometric Scaling Laws in Biology. *Science* 276, 122-126. <https://doi.org/10.1126/science.276.5309.122>
15. Allan Drummond D., Wilke C. O. (2009). The evolutionary consequences of erroneous protein synthesis. *Nature Reviews Genetics* 10, 715-724. <https://doi.org/10.1038/nrg2662>
16. Hopfield J. J. (1974). Kinetic Proofreading: A New Mechanism for Reducing Errors in Biosynthetic Processes Requiring High Specificity. *Proceedings of the National Academy of Sciences* 71, 4135-4139. <https://doi.org/10.1073/pnas.71.10.4135>
17. Ninio J. (1975). Kinetic amplification of enzyme discrimination. *Biochimie* 57, 587-595. [https://doi.org/10.1016/s0300-9084\(75\)80139-8](https://doi.org/10.1016/s0300-9084(75)80139-8)
18. Akashi H., Gojobori T. (2002). Metabolic efficiency and amino acid composition in the proteomes of *Escherichia coli* and *Bacillus subtilis*. *Proceedings of the National Academy of Sciences* 99, 3695-3700. <https://doi.org/10.1073/pnas.062526999>
19. Landauer R. (1961). Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development* 5, 183-191. <https://doi.org/10.1147/rd.53.0183>
20. Bennett C. H. (1982). The thermodynamics of computation—a review. *International Journal of Theoretical Physics* 21, 905-940. <https://doi.org/10.1007/bf02084158>
21. Crick F. (1968). The origin of the genetic code. *Journal of Molecular Biology* 38, 367-379. [https://doi.org/10.1016/0022-2836\(68\)90392-6](https://doi.org/10.1016/0022-2836(68)90392-6)
22. Woese C. R. (1965). On the evolution of the genetic code.. *Proceedings of the National Academy of Sciences* 54, 1546-1552. <https://doi.org/10.1073/pnas.54.6.1546>
23. Koonin E. V., Novozhilov A. S. (2009). Origin and evolution of the genetic code: The universal enigma. *IUBMB Life* 61. <https://doi.org/10.1002/iub.190>
24. Eigen M. (1971). Selforganization of matter and the evolution of biological macromolecules. *Die Naturwissenschaften* 58, 465-523. <https://doi.org/10.1007/bf00623322>

25. Sanjuán R., Nebot M. R., Chirico N., Mansky L. M., Belshaw R. (2010). Viral Mutation Rates. *Journal of Virology* 84, 9733-9748.
<https://doi.org/10.1128/JVI.00694-10>
26. Drake J. W., Holland J. J. (1999). Mutation rates among RNA viruses. *Proceedings of the National Academy of Sciences* 96, 13910-13913.
<https://doi.org/10.1073/pnas.96.24.13910>
27. Drake J. W. (1991). A constant rate of spontaneous mutation in DNA-based microbes.. *Proceedings of the National Academy of Sciences* 88, 7160-7164.
<https://doi.org/10.1073/pnas.88.16.7160>
28. Gago S., Elena S. F., Flores R., Sanjuán R. (2009). Extremely High Mutation Rate of a Hammerhead Viroid. *Science* 323, 1308-1308.
<https://doi.org/10.1126/science.1169202>
29. Eckerle L. D., Lu X., Sperry S. M., Choi L., Denison M. R. (2007). High Fidelity of Murine Hepatitis Virus Replication Is Decreased in nsp14 Exoribonuclease Mutants. *Journal of Virology* 81, 12135-12144.
<https://doi.org/10.1128/jvi.01296-07>
30. Smith E. C., Blanc H., Vignuzzi M., Denison M. R. (2013). Coronaviruses Lacking Exoribonuclease Activity Are Susceptible to Lethal Mutagenesis: Evidence for Proofreading and Potential Therapeutics. *PLoS Pathogens* 9, e1003565.
<https://doi.org/10.1371/journal.ppat.1003565>
31. Edelsbrunner, Letscher, Zomorodian (2002). Topological Persistence and Simplification. *Discrete & Computational Geometry* 28, 511-533.
<https://doi.org/10.1007/s00454-002-2885-2>
32. Zomorodian A., Carlsson G. (2004). Computing Persistent Homology. *Discrete & Computational Geometry* 33, 249-274. <https://doi.org/10.1007/s00454-004-1146-y>
33. Tralie C., Saul N., Bar-On R. (2018). Ripser.py: A Lean Persistent Homology Library for Python. *Journal of Open Source Software* 3, 925.
<https://doi.org/10.21105/joss.00925>
34. Xia K., Wei G. (2014). Persistent homology analysis of protein structure, flexibility, and folding. *International Journal for Numerical Methods in Biomedical Engineering* 30, 814-844. <https://doi.org/10.1002/cnm.2655>
35. Helfrich W. (1973). Elastic Properties of Lipid Bilayers: Theory and Possible Experiments. *Zeitschrift für Naturforschung C* 28, 693-703.
<https://doi.org/10.1515/znc-1973-11-1209>
36. Hu M., Briguglio J. J., Deserno M. (2012). Determining the Gaussian Curvature Modulus of Lipid Membranes in Simulations. *Biophysical Journal* 102, 1403-1410. <https://doi.org/10.1016/j.bpj.2012.02.013>

37. Rawicz W., Olbrich K., McIntosh T., Needham D., Evans E. (2000). Effect of Chain Length and Unsaturation on Elasticity of Lipid Bilayers. *Biophysical Journal* 79, 328-339. [https://doi.org/10.1016/S0006-3495\(00\)76295-3](https://doi.org/10.1016/S0006-3495(00)76295-3)
38. Nagle J. F., Tristram-Nagle S. (2000). Structure of lipid bilayers. *Biochimica et Biophysica Acta (BBA) - Reviews on Biomembranes* 1469, 159-195. [https://doi.org/10.1016/S0304-4157\(00\)00016-2](https://doi.org/10.1016/S0304-4157(00)00016-2)
39. Morlot S., Roux A. (2013). Mechanics of Dynamin-Mediated Membrane Fission. *Annual Review of Biophysics* 42, 629-649. <https://doi.org/10.1146/annurev-biophys-050511-102247>
40. Kraus F., Roy K., Pucadyil T. J., Ryan M. T. (2021). Function and regulation of the divisome for mitochondrial fission. *Nature* 590, 57-66. <https://doi.org/10.1038/s41586-021-03214-x>
41. White J. H. (1969). Self-Linking and the Gauss Integral in Higher Dimensions. *American Journal of Mathematics* 91, 693. <https://doi.org/10.2307/2373348>
42. Fuller F. B. (1971). The Writhing Number of a Space Curve. *Proceedings of the National Academy of Sciences* 68, 815-819. <https://doi.org/10.1073/pnas.68.4.815>
43. Champoux J. J. (2001). DNA Topoisomerases: Structure, Function, and Mechanism. *Annual Review of Biochemistry* 70, 369-413. <https://doi.org/10.1146/annurev.biochem.70.1.369>
44. Prigogine I., Nicolis G. (1971). Biological order, structure and instabilities. *Quarterly Reviews of Biophysics* 4, 107-148. <https://doi.org/10.1017/s0033583500000615>
45. Brandman O., Ferrell J. E., Li R., Meyer T. (2005). Interlinked Fast and Slow Positive Feedback Loops Drive Reliable Cell Decisions. *Science* 310, 496-498. <https://doi.org/10.1126/science.1113834>
46. Tsai T. Y. C., Choi Y. S., Ma W., Pomerening J. R., Tang C., Ferrell J. E. (2008). Robust, Tunable Biological Oscillations from Interlinked Positive and Negative Feedback Loops. *Science* 321, 126-129. <https://doi.org/10.1126/science.1156951>
47. Bak P., Tang C., Wiesenfeld K. (1987). Self-organized criticality: An explanation of the $1/f$ noise. *Physical Review Letters* 59, 381-384. <https://doi.org/10.1103/PhysRevLett.59.381>
48. Lane N., Martin W. F. (2012). The Origin of Membrane Bioenergetics. *Cell* 151, 1406-1416. <https://doi.org/10.1016/j.cell.2012.11.050>

49. Lane N., Allen J. F., Martin W. (2010). How did LUCA make a living? Chemiosmosis in the origin of life. *BioEssays* 32, 271-280. <https://doi.org/10.1002/bies.200900131>
50. Martin W., Russell M. J. (2003). On the origins of cells: a hypothesis for the evolutionary transitions from abiotic geochemistry to chemoautotrophic prokaryotes, and from prokaryotes to nucleated cells. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 358, 59-85. <https://doi.org/10.1098/rstb.2002.1183>
51. Kun Á., Santos M., Szathmáry E. (2005). Real ribozymes suggest a relaxed error threshold. *Nature Genetics* 37, 1008-1011. <https://doi.org/10.1038/ng1621>
52. Szostak J. W., Bartel D. P., Luisi P. L. (2001). Synthesizing life. *Nature* 409, 387-390. <https://doi.org/10.1038/35053176>
53. Zhu T. F., Szostak J. W. (2009). Coupled Growth and Division of Model protocell Membranes. *Journal of the American Chemical Society* 131, 5705-5713. <https://doi.org/10.1021/ja900919c>
54. Varela F., Maturana H., Uribe R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *Biosystems* 5, 187-196. [https://doi.org/10.1016/0303-2647\(74\)90031-8](https://doi.org/10.1016/0303-2647(74)90031-8)
55. Maturana H. R., Varela F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel Publishing Company
56. Gánti T. (2003). *The Principles of Life*. Oxford University Press
57. Rosen R. (1958). The representation of biological systems from the standpoint of the theory of categories. *The Bulletin of Mathematical Biophysics* 20, 317-341. <https://doi.org/10.1007/BF02477890>
58. Friston K. (2010). The free-energy principle: a unified brain theory?. *Nature Reviews Neuroscience* 11, 127-138. <https://doi.org/10.1038/nrn2787>
59. Tononi G. (2004). An information integration theory of consciousness. *BMC Neuroscience* 5, 42. <https://doi.org/10.1186/1471-2202-5-42>
60. Brown J. H., Gillooly J. F., Allen A. P., Savage V. M., West G. B. (2004). TOWARD A METABOLIC THEORY OF ECOLOGY. *Ecology* 85, 1771-1789. <https://doi.org/10.1890/03-9000>
61. Milo R., Shen-Orr S., Itzkovitz S., Kashtan N., Chklovskii D., Alon U. (2002). Network Motifs: Simple Building Blocks of Complex Networks. *Science* 298, 824-827. <https://doi.org/10.1126/science.298.5594.824>
62. Lane N., Martin W. F. (2010). The energetics of genome complexity. *Nature* 467, 929-934. <https://doi.org/10.1038/nature09486>

63. Allen J. F. (1993). Control of gene expression by redox potential and the requirement for chloroplast and mitochondrial genomes. *Journal of Theoretical Biology* 165, 609-631. <https://doi.org/10.1006/jtbi.1993.1210>
64. Kimura M. (1962). On the probability of fixation of mutant genes in a population. *Genetics* 47, 713-719.
65. Booth A., Doolittle W. F. (2015). Eukaryogenesis, how a fossil became a vital part of eukaryote cells. *Philosophical Transactions of the Royal Society B* 370, 20140333. <https://doi.org/10.1098/rstb.2014.0333>
66. Hinrichsen H. (2000). Non-equilibrium critical phenomena and phase transitions into absorbing states. *Advances in Physics* 49, 815-958. <https://doi.org/10.1080/00018730050198152>
67. Aon M. A., Cortassa S., O'Rourke B. (2004). Percolation and criticality in a mitochondrial network. *Proceedings of the National Academy of Sciences* 101, 4447-4452. <https://doi.org/10.1073/pnas.0307156101>
68. Figge M. T., Reichert A. S., Meyer-Hermann M., Osiewacz H. D. (2012). Deceleration of Fusion–Fission Cycles Improves Mitochondrial Quality Control during Aging. *PLoS Computational Biology* 8, e1002576. <https://doi.org/10.1371/journal.pcbi.1002576>
69. Fiedler M. (1973). Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal* 23, 298-305. <https://doi.org/10.21136/cmj.1973.101168>

Note on the manuscript's history. The version labels below refer to the private manuscript series; this is the first public version. The author list is S. Zhang and J.Duan in V2, V3 and V4 alike; no author change has occurred between versions. V3 added the design table, the blind KEGG re-analysis and the margin statistics. V4 adds: the basal-leak phase boundary $\epsilon()$ and the persistence/spontaneous-genesis criterion split, which reconcile the absorbing dead state with the design table; the dispersal-unlocked invasion analysis (pinning is structural; division unlocks it); the frozen four-null-model robustness protocol, selection bound and Bayesian contrast for the margin statistic; section renumbering throughout (Sects. 8–12); and the terminology and cross-reference audit.