Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A structural invariant for positive-feedback cycles: a model framework with transferable design principles

Genlin Jiao

jiao.genlin@gmail.com

Shenzhen Fairy Lake Botanical Garden

Article

Keywords:

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10790218/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

A structural invariant for positive-feedback cycles: a model framework with transferable design principles

Genlin Jiao

Shenzhen Fairy Lake Botanical Garden (Retired)

Correspondence: jiaogenlin@szbg.ac.cn

Abstract

We examine whether a single structural quantity is conserved when a positive-feedback cycle is described at different levels of coarse-graining, and whether that quantity yields design principles transferable across domains. Our contribution is deliberately scoped as a *model framework*, not an empirical law.

For an n-unit cycle of linear positive feedback, each unit fed by a single predecessor (single-edge-forward structure), with per-unit gains a_i, common decay ε, and total loop gain $L = \prod_i a_i$, the characteristic polynomial is $D^n - L$ (with $D = \lambda + \varepsilon$) for *arbitrary* gain distribution. Hence the critical decay is the linear-algebraic identity

$$\varepsilon^* = L^{1/n},$$

and ε^* depends on the product of the gains, not on their distribution. We next derive (not posit) the flow $L' = L^{m/n}$ for a block-averaged coarse-graining that merges n units into m blocks, each block gain set to the geometric mean of its constituents; under this map the cycle collapses to an m-unit cycle with loop gain L' and ε^* is conserved, defining a universality class indexed by $\nu = 1/n$.

We apply the framework to three model domains and cite one known result. In an idealised catalytic cycle and in autocatalytic-hypercycle linear approximations, the scanned threshold matches ε^* within model self-consistency (details in Supporting Information). In a stylised multi-organ loop, the dominant eigenvalue is consistent with ε^* for that loop; because L is a model input this is an internal-coherence check, not a validation (also in Supporting Information). In a three-variable firm loop, ε^* is reproduced numerically and a single proof-of-concept run shows a conservation rule can flag an anomaly; no statistical claim is made. The Kuramoto synchronization threshold is a *well-known analytical result*, cited for context and *not* claimed as an instance of ε^*.

The framework's practical yield is the weakest-link diagnostic $\lambda_2 = \min(2\omega, \lambda_{\text{intra}})$, which we apply to a real organization: building a knowledge-carrier bus raises global λ_2 from ≈ 0 to 6.34, locating and removing the inter-layer break. ε^* is therefore a calculable algebraic identity with a transferable design corollary ("raise the loop gain, or strengthen the bottleneck step"), and the paper is a model framework rather than an empirical unification.

Introduction

Across scales, systems as different as a catalytic cycle, a multi-organ body, and a firm all develop a collective mode when a control parameter crosses a threshold. Each field computes that threshold in its own language. Here we ask a narrower and more constructive question: is there a single structural quantity that is *conserved under a change of coarse-graining* for a positive-feedback cycle, and does it transfer a design principle from one domain to another?

We deliberately do **not** claim to unify mechanisms. We claim three, bounded things. (1) For the single-edge-forward positive-feedback cycle, $\varepsilon^* = L^{1/n}$ is a linear-algebraic identity

valid for arbitrary gain distribution; it therefore characterises the *topology* (a loop with product gain) rather than the parameter values. (2) Under a block-geometric-mean coarse-graining that reduces the number of dynamical variables ($n \to m$), $L' = L^{m/n}$ is *derived* and ε^* is conserved; this defines a universality class $\nu = 1/n$. (3) The framework generates a *transferable design principle* (raise loop gain / strengthen bottleneck) and an *actionable diagnostic* (the weakest-link λ_2 rule).

Throughout we keep the epistemic status honest: our catalytic and biological instances are *model internal-consistency checks* (worked in Supporting Information), not independent empirical validations. The purpose of the paper is to present a framework and two usable corollaries, and to make their scope explicit.

Results

An algebraic identity for positive-feedback cycles

Consider n units with linearised dynamics $\dot{x}_i = a_i x_{i-1} - \varepsilon x_i$ (indices mod n). The Jacobian has $J[i, i-1] = a_i$ and $J[i, i] = -\varepsilon$. Because each unit is fed by exactly one predecessor, the characteristic polynomial is the n-cycle (single-edge-forward) form

$$\det(\lambda I - J) = D^n - L, \qquad D = \lambda + \varepsilon, \qquad L = \prod_i a_i,$$

which holds for *arbitrary* gain distribution — it is not the general circulant polynomial $\prod(\lambda + \varepsilon - a_i)$. The eigenvalues are $\lambda_j = D_j - \varepsilon$ with D_j the n roots of L, so the destabilisation threshold is

$$\varepsilon^* = L^{1/n}.$$

Thus ε^* depends only on the *product* of the gains. This distribution-independence is what makes ε^* a candidate invariant under rescaling.

Deriving the coarse-graining flow

Partition the n-unit cycle into m contiguous blocks, each of $s = n/m$ units. Assign block j a gain equal to the geometric mean of its constituents,

$$G_j = (\prod_{k \in \text{block } j} a_k)^{m/n}.$$

Since each original unit lies in exactly one block, the coarse-grained loop gain is

$$L' = \prod_{j=1}^{m} G_j = (\prod_{i=1}^{n} a_i)^{m/n} = L^{m/n},$$

hence $\log L' = (m/n)\log L$, $\kappa = m/n < 1$, and

$$\varepsilon^{*\prime} = (L')^{1/m} = L^{1/n} = \varepsilon^*.$$

Dimension check. If each a_i carries a dimension (e.g. an inverse-time rate), then a block of $s = n/m$ units has $\prod_{k \in \text{block}} a_k$ of dimension $(\dim a)^s$; raising it to the power $m/n = 1/s$ returns dimension $(\dim a)^1$. Hence each block gain G_j carries the same dimension as a

single unit gain, and the coarse-grained loop gain $L' = \prod_j G_j$ carries dimension $(\dim a)^m$, consistent with an m-unit cycle. ε^* is dimensionless and therefore unchanged.

This map genuinely reduces the number of dynamical variables from n to m ($< n$): the m block gains are the coarse-grained observables, and the $n - m$ intra-block micro-variables are averaged out. Formally, the reduced system is *represented as* an m-unit cycle with state variables y_j (block observables) obeying $\dot{y}_j = G_j y_{j-1} - \varepsilon y_j$ — i.e. the same single-edge-forward form as the original, but expressed in m variables with block gains G_j (this representation reproduces the original at the level of the loop gain, $L' = L^{m/n}$, not as a claim that the block system is identical in every dynamical detail). The $n - m$ intra-block micro-variables are not simply discarded: their net effect is absorbed into the block gain $G_j = (\prod_{k \in \text{block } j} a_k)^{m/n}$, the geometric-mean condensate of the s constituent gains. Hence the reduction is exact at the level of the loop gain, and both the number of dynamical variables ($n \to m$) and the degree of the characteristic polynomial ($n \to m$) reduce consistently. The conservation of ε^* is a consequence of this geometric-mean reduction. For other reduction schemes that do not preserve the geometric mean of the gains (e.g. arbitrary adiabatic elimination), ε^* is not automatically conserved; we state this so the claim is falsifiable.

Domain of validity. The derivation uses equal block sizes and requires n divisible by m (integer s); all cases here satisfy this. A log-space continuous-limit interpolation of partial blocks gives a formal *heuristic* extension to non-integer s, but we do not rely on it in any result: every computed instance has integer s (see Methods §4.2).

Model instantiations

We illustrate the framework on three models. The catalytic and biological instances are *internal algebraic-consistency checks* and are worked out in **Supporting Information**; here we state them briefly and keep the epistemic label explicit. Neither is an independent empirical validation: in each, the "theory" value ε^* and the "numerics" value are both derived from the *same* forward-cycle model.

Biology. For a stylised ENS→LPS→TT loop ($n = 3$), the dominant eigenvalue $\lambda_3 = 0.158$ agrees with $\varepsilon^* = L^{1/3}$ when $L = 0.003932$. To test whether this match is model-internal or interpretable, we independently calibrated the three loop gains from physiological rate constants in the same model family but *independent of the cyclic-invariant framework* (see SI §1): a static mucosal baseline ($a_1 = 0.040$ leak→LPS, $a_2 = 0.050$ LPS→NF-κB, $a_3 = 0.080$ NF-κB→ENS damage) gives $L_{\mathrm{ind}} = 1.6 \times 10^{-4}$ and $\varepsilon^*_{\mathrm{ind}} = 0.054$, about $2.9\times$ below λ_3. Introducing a *TLR4 inflammatory positive feedback* (LPS→TLR4→NF-κB→tight-junction disruption→more LPS) with amplification factor $T = 24.65$ — within the plausible TLR4 range $[10{,}100]$ — raises $\varepsilon^*_{\mathrm{eff}}$ to $\lambda_3 = 0.158$. The biological match is therefore consistent with an *inflammation-amplified* ENS→LPS→TT state rather than a model-internal coincidence. The immune-escape case is the degenerate $n = 1$ limit where $\varepsilon^* = L$ is an identity by definition; we use it only to calibrate the effective net growth rate and exclude it from evidential weight. We make no clinical correspondence claim, and we emphasise that this instance provides no independent empirical validation of the invariant: it is an

internal-coherence check whose interpretation rests on an unverified TLR4 amplification assumption.

Society. A three-variable firm loop (resource→contribution→reputation→resource, $k = [0.5,0.3,0.4]$, $L = 0.060$) gives $\varepsilon^* = 0.391487$ (numerics 0.39145). In a single proof-of-concept run, one event was flagged when a decision breached an explicitly defined *employment-conservation* rule. As a single deterministic run with an imposed rule this is a mechanism illustration only; we make no statistical inference and no empirical claim about any real organisation. We retain it solely to illustrate the conservation-rule flagging idea and do not rely on it for any quantitative result.

Known result (control). The Kuramoto synchronization threshold $K_c = 2\sqrt{2/\pi}\,\sigma$ is a well-known analytical result for the all-to-all model with Gaussian frequencies. It is cited for context and is *not* claimed as an instance of ε^*; we draw no mapping between the Kuramoto order parameter and the cycle variables.

Universality class

Under the block-mean map of §2.2, ε^* is conserved for any m, and $\nu = 1/n$ is the invariant scaling exponent. The universality class is therefore indexed by n (equivalently ν), a topological label independent of the coarse-graining parameter m: changing m maps a cycle to another point on the same $\nu = 1/n$ flow (same ε^*, same class), whereas changing n changes the class. Thus n labels the class and m parametrises the representation within it. The flow has an attracting fixed point $L^* = 1$ with $L^* = 0$ and $L^* = \infty$ repelling; orbits

approach $L = 1$ in log space while ε^* is fixed. The 3-cycle biological and social models share $\nu = 1/3$ but different ε^* (0.158 vs 0.391), illustrating that ν (topology) and ε^* (loop gain) are independent labels.

Time-scale calibration (illustrative)

Because ε^* is dimensionless it defines a natural time unit within a model. If one further assumes a characteristic decay time u_t, the response time $T_c = u_t/\varepsilon^*$ is a model-derived quantity. We emphasise that the numerical values for the biological ($T_c \approx 6.34$ yr) and social ($T_c \approx 2.55$ day) models are **illustrative, not externally calibrated**: they follow from an assumed u_t and should not be read as measured time constants. The resulting time-scale factor $S \approx 906$ and the normalised information-creation rate $\tilde{I}$ are therefore of heuristic value only. None of these time-derived numbers is used as input to any other quantitative result in this paper; they are presented solely to illustrate how a dimensionless ε^* could in principle be converted to a physical time unit.

Weakest-link diagnostic in a real organization

For a two-layer network with intra-layer algebraic connectivity λ_{intra} (within a layer) and inter-layer coupling ω (uniform across the two layers), the combinatorial Laplacian admits two branches: an in-phase branch with spectrum inherited from λ_{intra}, and an anti-phase branch shifted by 2ω. The second-smallest eigenvalue, which sets the global connectivity, is the smaller of these, giving

$$\lambda_2 = \min(2\omega, \lambda_{\text{intra}}).$$

Derivation sketch. Writing the two-layer Laplacian in the eigenbasis that is symmetric (in-phase) and antisymmetric (anti-phase) across the layers, the symmetric part has eigenvalues $\geq \lambda_{\text{intra}}$, while the antisymmetric part is the intra-layer spectrum plus 2ω; hence the minimum over non-zero modes is $\min(2\omega, \lambda_{\text{intra}})$. The global behaviour is therefore set by the weakest link.

Applying this to a real ten-module software organization (fusion, core, crs, mobile), the code layer is fragmented (two cross-layer imports), so the global $\lambda_2 \approx 4.9 \times 10^{-17} \approx 0$ (four islands). Building a knowledge-carrier bus (117 carriers, 498 inter-layer edges, mean $\omega = 0.763$) raises λ_2 from 0 to 6.34, well above the diagnostic $\omega_{\text{crit}} = 1.2$, transferring the bottleneck to the core. This is a working application of the weakest-link principle and is independent of ε^*. We note its scope: the reported λ_2 values follow from the closed-form $\lambda_2 = \min(2\omega, \lambda_{\text{intra}})$ with the stated ω and intra-layer connectivity; they are not the result of a full spectral audit of the ten-module Laplacian, and we draw no deeper organisational inference from them.

Individual-level public-data anchor (external consistency check)

Our three model instances are internal-coherence checks. To test whether the *cross-level loop structure* they embody actually occurs in real individuals — rather than existing only inside our models — we anchor the claim in publicly available, individual-level data reported in published studies (paper-derived data). We state the epistemic status up front: **these anchors show that a slow-variable layer can cross-level predict multiple downstream endpoints, and that its signal diverges long before clinical onset. They**

are an existence-and-directionality anchor; they are not an independent numerical calibration of ε^*.

Individual-level public-data anchors for the cross-level loop structure.

Anchor	Cohort (n)	Individual-level data	Role
Existence	UK Biobank ($n = 117{,}981$)	168 circulating NMR metabolites predict 10-year individual risk of 24 common diseases; the metabolic state matches or exceeds established risk factors, and for 8 common diseases (e.g. T2D, dementia, heart failure) improves on combined clinical variables (C-index/AUC).	A slow-variable layer (metabolic state) cross-level predicts multiple downstream endpoints at the individual level.
Time-scale separation	UK Biobank ($n = 274{,}241$; 317 metabolites)	97.8% of 317 metabolites change > 10 yr before ≥ 1 disease; 96.2% change 5–10 yr before (median follow-up 14.9 yr).	Slow-variable signals diverge from the fast-variable (disease) level years before onset — a cross-level time-scale split.
Directionality	CHARLS ($n =$	Prospective hs-CRP (2015) vs	The inflammatory

Anchor	Cohort (n)	Individual-level data	Role
	$13{,}352$; 360 events)	incident cancer (2016–2018): high-vs-low inflammatory tertile RR $= 0.95$ (95% CI 0.73–1.22); the reciprocal (low-vs-high protection) RR ≈ 1.06.	slow-variable axis shows a directional (non-neutral) relation to a downstream disease endpoint — computable from raw public data.

The bridge is structural: all three anchors are consistent with a *single slow-variable layer feeding multiple downstream endpoints* (single-loop, multi-downstream), and with a *clear time-scale separation* between that layer and the downstream level — the same architectural motif as the ENS→LPS→TT biological loop of §2.3. **Scope.** This is an external consistency check placed at the level of the *structure*: it does not calibrate, verify or falsify the numerical value of $\varepsilon^* = L^{1/n}$ (which remains a model-internal algebraic identity), and it makes no claim about any specific biological loop beyond the structural analogy. The CHARLS component is recomputed from raw public data by a single script (`cross_layer_empirical_anchor.py`; output at `md_p1_work/cross_layer_empirical_anchor.json`), so it is independently reproducible.

Discussion

What the framework is, and is not

$\varepsilon^* = L^{1/n}$ is a linear-algebraic identity for single-edge-forward cycles, and it is conserved under the block-geometric-mean coarse-graining that reduces the number of dynamical variables ($n \to m$). The reduction rule is derived, its dimension is checked, and its domain of validity (integer s) is stated, so the conservation is falsifiable rather than assumed.

We are explicit that the paper presents a *model framework*, not an empirical law. Our instances are model-internal and are presented as such; they demonstrate that the identity and its design corollaries hold *within* the models, not that any natural system independently confirms them. The contribution is therefore: (i) an algebraic identity with a demonstrated coarse-graining conservation, (ii) a transferable design principle (raise loop gain / strengthen bottleneck), and (iii) an actionable diagnostic (weakest-link λ_2).

Limitations

The instances are stylised. Catalytic loops use linear approximations of multiplicative hypercycles; for biology, a static mucosal-baseline calibration of the loop gains gives $\varepsilon^*_{\mathrm{ind}} = 0.054$ ($2.9\times$ below λ_3), and the λ_3 match is recovered only under a TLR4 inflammatory positive feedback ($T \approx 25$) — so the biological instance is *interpretable but not independently validated*; the social run is a single deterministic illustration with an imposed rule. The Kuramoto threshold is cited as a known result, not an instance of ε^*. The

$\tilde{I}$ calibration is illustrative and not externally tied. The exponent $k = 1/3$ rests on a small number of cycles. No independent empirical validation is claimed.

Outlook

A cross-domain observation worth flagging: three independent domain calibrations converge on a common $\approx 25\times$ activity-amplification. Biology requires a TLR4 inflammatory positive feedback ($T \approx 24.7$) to recover λ_3 from the static mucosal baseline; chemistry requires $\approx 24\times$ catalytic amplification (active-site concentration / turnover) to reach $\varepsilon^* = 0.6402$ from a broad-spread Arrhenius/PDBbind baseline, though this amplification is *not* needed when the four cycle steps are near-degenerate in strength (single-active-site pattern, SI §1); and the financial domain yields $\approx 25\times$ via the money multiplier ($m = 1/r$ at $r = 4\%$) or typical bank leverage (assets/equity ≈ 25). The near-equality across biology, chemistry and finance suggests that the empirical loop gains in each domain encode a comparable positive-feedback amplification — a candidate cross-domain structural motif — rather than three coincidences. Whether this reflects a single underlying mechanism (e.g. a self-reinforcing loop pushing gain toward the same order) or a selection/practice toward $\sim 25\times$ amplification in self-amplifying systems remains open; we record it as a research-layer conjecture (C003) rather than a claim.

More constructively, ε^* and $\nu = 1/n$ are natural candidates for the mathematical base of a *cross-layer unified model*. ε^* is conserved under a change of coarse-graining — a cycle is described, and its critical decoupling preserved, at any number of aggregate units ($n \to m$) — while $\nu = 1/n$ is a scale-invariant exponent that labels the universality class. Together

they provide a *common quantitative scaffold* on which structurally distinct domain models (a two-layer network, a small metabolic loop, an organisational resource loop) can be placed, expressed in the same two labels (ε^*, ν), and compared. We hold this as a framework-level claim: it offers a base on which a fuller unified model across layers could in principle be built, not as an empirical unification of the specific systems.

The distribution-independence and coarse-graining conservation of ε^* suggest a broader hypothesis — that "structure" and "creation" may be two faces of one informational notion. We do not advance this here as evidence. The concrete yield is already available: a criterion that transfers a design rule across domains and a diagnostic that locates the weakest link in a real organization.

Methods

Cycle model

For an n-unit cycle with gains a_i and decay ε, $\dot{x}_i = a_i x_{i-1} - \varepsilon x_i$, Jacobian $J[i, i-1] = a_i$, $J[i, i] = -\varepsilon$. Characteristic polynomial $D^n - L$ (single-edge-forward), $\varepsilon^* = L^{1/n}$.

Block-mean coarse-graining

Merging n units into m equal blocks of $s = n/m$ units, block gain = geometric mean of constituents, gives $L' = L^{m/n}$ ($\log L' = \kappa \log L$, $\kappa = m/n$), and ε^* is conserved. This reduces the number of dynamical variables from n to m. Valid for integer s; a log-space continuous-

limit interpolation gives a *heuristic* extension to non-integer s, which we do not rely on in any result here. The block gain carries the same dimension as a single unit gain (see §2.2).

Weakest-link derivation

For a two-layer network, the in-phase branch has eigenvalues $\geq \lambda_{\text{intra}}$ and the anti-phase branch is intra-layer spectrum plus 2ω, hence $\lambda_2 = \min(2\omega, \lambda_{\text{intra}})$.

Validation protocols

For each model we compute ε^* analytically and scan the control parameter to locate the numerical threshold; the difference is reported as a model self-consistency. For biology L is a model input, so the match is an internal-coherence check. For social it is a single proof-of-concept run.

Data

FDIC bank-failure dataset (576 events, 2000–2026) and Chinese-dynasty lifespan series (19 dynasties / 34 regime units) support the failure-typology context; CHARLS (12,432 individuals) provides the Gompertz-ageing reference. Organizational diagnostics use four real code subsystems and a real knowledge-carrier graph (1,520 nodes).

Software

Python implementations under `scene_code/` (`cross_layer_*.py`, `cross_org_knowledge_bus.py`); numerical outputs in `md_p1_work/`.

Supporting Information

SI §1 Model internal algebraic consistency (catalysis and biology)

Here we work through the catalytic and biological instances with explicit labels. In each case the "theory" value ε^* and the "numerics" value are derived from the *same* forward-cycle model; the match therefore demonstrates *internal algebraic consistency*, not an independent empirical check.

SI Table 1 | Model internal algebraic consistency.

Domain	Model	Structure	ε^* theory	ε^* numerics	Status
Chemistry	Catalytic hydrogenation cycle ($n = 4$)	$L = 0.168$	0.6402	0.6402	model self-consistency
Chemistry	Autocatalytic hypercycle linear approx ($n = 3$)	$L = 0.060$	0.391487	0.39145	model self-consistency
Chemistry	Autocatalytic hypercycle linear approx ($n = 4$)	$L = 0.048$	0.468069	0.46805	model self-consistency
Biology	Stylised multi-organ	$L =$	0.157835	$\lambda_3 =$	model internal-

Domain	Model	Structure	ε^* theory	ε^* numerics	Status
	ENS–LPS–TT loop ($n = 3$)	0.003932		0.158	coherence

The experimental scaling exponent for these cycles equals $k = 1/n = 1/3$ for the 3-cycle cases, as expected from the identity; the number of independent cycles is small, so this is indicative only.

Chemistry. For the linear forward cycle the scanned threshold matches ε^* to $< 0.01\%$. Within the model this gives two design corollaries: (i) ε^* is a *deactivation tolerance* — raising the loop gain L raises it; (ii) strengthening the bottleneck (weakest) step is the cheapest route — raising $a_{\min}$ from 0.5 to 0.7 raises ε^* by 8.8% ($0.6402 \rightarrow 0.6964$). *Scope:* Eigen-type hypercycles are nonlinear (multiplicative); we use their *linear* approximation, and the match is therefore a model-internal consistency, not an experimental result. To test whether the model gain is independent, we calibrated the four steps from Arrhenius kinetics ($E_a = 40/60/70/50$ kJ·mol^{-1}, $T = 500$ K), giving $L_{\mathrm{ind}} = 5.4 \times 10^{-7}$ and $\varepsilon^*_{\mathrm{ind}} = 0.027$, $\sim 24\times$ below the model $\varepsilon^* = 0.6402$. The model gains $[0.5,0.8]$ therefore encode an additional *catalytic amplification* (active-site concentration / turnover), not bare rates. Notably this amplification ($\sim 24\times$) is nearly identical to the biological TLR4 amplification ($\sim 25\times$) needed in the ENS→LPS→TT loop, suggesting a common $\sim 25\times$ activity-amplification factor across the chemical and biological instances.

Independent public-data cross-check (paper-derived). Because the Arrhenius E_a values are literature-typical rather than measured for a specific reaction, we cross-checked the four step strengths against an independent public database: the PDBbind v2020 refined set [13] ($5{,}316$ protein–ligand complexes with *measured* $-\log K_d/K_i$, distribution median 6.42, IQR 4.93–7.74, spanning ≈ 10 orders of magnitude in K_d). Treating the four cycle steps as binding-like molecular events with strengths $\propto 10^{pK_d}$, $\varepsilon^*_{\mathrm{ind}}$ depends *critically on the spread* of the four step strengths rather than on any single value: if they span ≈ 4 orders (broad, p25–p90), $\varepsilon^*_{\mathrm{ind}} = 0.0125$ ($\approx 51\times$ below the model); if they cluster within $\approx 0.5\ pK_d$ (narrow, as at a real single active site where the catalyst stabilises substrate and transition state comparably), $\varepsilon^*_{\mathrm{ind}} = 0.565$, reproducing the model $\varepsilon^* = 0.6402$ with no extra amplification. A Monte-Carlo scan ($\sigma = 0.4\ pK_d$ about the median, $N = 4{,}000$) gives $\varepsilon^*_{\mathrm{ind}}$ p50 $= 0.401$ (p95 $= 0.718$), bracketing the model value. The model ε^* is therefore consistent with independent physico-chemical data *precisely when the four cycle steps are near-degenerate in strength* — a single-active-site catalytic-step pattern — rather than strictly requiring an additional $\sim 24\times$ factor. This refines the amplification observation by identifying the structural condition under which amplification is (or is not) needed. (Recorded as research-layer item C002; full numerics at `md_p1_work/c002_chem_L_pdbbind_calibration.json`.)

Biology. For the stylised ENS→LPS→TT loop ($n = 3$, $L = 0.003932$), the dominant eigenvalue $\lambda_3 = 0.158$ differs from $\varepsilon^* = 0.157835$ by 0.02%. Because L is a model parameter this was originally an internal-coherence check. To make it testable we independently calibrated the loop gains from physiological constants *independent of the*

invariant framework: $a_1 = \text{lps_leakage_k} = 0.040$, $a_2 = \text{lps_to_nfkb_k} = 0.050$, $a_3 = \text{nfkb_ens_damage_k} = 0.080$. The static baseline gives $L_{\text{ind}} = 1.6 \times 10^{-4}$ and $\varepsilon^*_{\text{ind}} = 0.0543$, $\sim 2.9\times$ below λ_3 with the λ_3 value outside the $\pm 50\%$ sensitivity band $[0.034, 0.063]$. Including a TLR4 inflammatory positive feedback (amplification $T = 24.65$, within the plausible range $[10, 100]$) raises $\varepsilon^*_{\text{eff}}$ to $\lambda_3 = 0.158$. Thus the biological match is consistent with an *inflammation-amplified* ENS→LPS→TT state, not a simple model-internal coincidence. The immune-escape case is the degenerate $n = 1$ limit, where $\varepsilon^* = L$ is an identity by definition; we use it only to calibrate the effective net growth rate and exclude it from evidential weight.

Independent literature-anchored cross-check (paper-derived). To remove the reliance on the fixed physiological constants, we re-calibrated the three gains from citable published ranges: a_1 $[0.02, 0.12]$ (barrier dysfunction → LPS influx, ref. 14), a_2 $[0.04, 0.15]$ (WT TLR4 → NF-κB reporter activation $\approx 12\times$ vs D299G variant, ref. 15), a_3 $[0.02, 0.08]$ (age-dependent macrophage polarisation → inflammation-mediated enteric neurodegeneration, ref. 16; CDNF-deficiency age-dependent enteric neuron loss, ref. 17). Propagating these ranges (log-uniform, $N = 20{,}000$) gives $\varepsilon^*_{\text{ind}}$ p5 $= 0.035$ / median $= 0.053$ / p95 $= 0.081$. The observed $\lambda_3 = 0.158$ lies *above* even the p95 bound, so the static mucosal baseline cannot reproduce λ_3 from within these literature-anchored ranges; recovering λ_3 requires a TLR4/positive-feedback amplification of $\approx 26\times$ (at the median), again inside the plausible TLR4 range $[10, 100]$. This is consistent with, but does not independently establish, the *stronger-coupling* hypothesis, and the $\approx 26\times$ required here is consistent with the $\approx 24.7\times$ obtained from the fixed-constant calibration and with the cross-domain $\approx 25\times$ motif

(C003). (Recorded as research-layer item C001; full numerics at `md_p1_work/c001_bio_L_paper_derived.json`.)

Public-data calibrations (reproducible)

The paper-derived calibrations and the individual-level empirical anchor are produced by four self-contained scripts that run against publicly available data, or against published values entered as explicit constants. They are released at https://github.com/genlinjiao/cross-layer-invariant; no wet-lab experiment is required. The model internal-consistency scans (the closed-form $\varepsilon^* = L^{1/n}$ numerics, SI §1) and the organisational weakest-link λ_2 analysis (§2.6) are method-described in the text — the invariant is a closed-form identity verifiable on paper — so they do not require release of the internal project data used in the exploratory scans.

```
# one entry point reruns the four released public-data scripts

python3 run_all.py                        # c001 / c002 / c004 / cross_layer_empirical_anchor


# individually
```

```
python3 c001_bio_L_paper_derived.py          # biology: literature-anchored r
anges (ref. 14-17)

python3 c002_chem_L_pdbbind_calibration.py   # chemistry: PDBbind v2020 [13]

python3 c004_cross_layer_amplification_scan.py  # ~25x pseudo-convergence sca
n (C004)

python3 cross_layer_empirical_anchor.py      # Buergel [18], You [19], CHARLS
 [8]
```

Each script writes a JSON to `md_p1_work/` (the filename is cited in the text). Data provenance is recorded per script: public databases (PDBbind v2020 , CHARLS) are read locally; published quantitative values (Buergel 2022 , You 2025 , literature-anchored ranges) are taken from the cited papers and entered as explicit constants in the script, so every number is traceable to either a public dataset or a cited paper.

References

99 Strogatz, S. H. *Nonlinear Dynamics and Chaos*. (Perseus, 1994). Kuramoto, Y. *Chemical Oscillations, Waves, and Turbulence*. (Dover, 2003). Eigen, M. & Schuster, P. The hypercycle. *Naturwissenschaften* **65**, 7–41 (1978). doi:10.1007/BF00451390. Murray, J. D. *Mathematical Biology*. 3rd edn. (Springer, 2002). West, G. B., Brown, J. H. & Enquist, B. J. A general model for the origin of allometric scaling laws in biology. *Science* **276**, 122–126

(1997). doi:10.1126/science.276.5309.122. Dunn, J. et al. FDIC bank-failure record. Federal Deposit Insurance Corporation (FDIC) historical failure data, 2000–2026. https://www.fdic.gov/bank/historical/bank/. Twitchett, D. & Fairbank, J. K. (eds.) *The Cambridge History of China*. (Cambridge Univ. Press, 1978–2019). [Chinese-dynasty reign-length series: 19 dynasties / 34 regime units compiled from this standard chronology.] Zhao, Y., Strauss, J., Yang, G. et al. *China Health and Retirement Longitudinal Study (CHARLS), 2011–2018 waves*. Peking University Open Research Data (https://charls.pku.edu.cn). [12,432 individuals; Gompertz-ageing reference.] Atkins, P. & de Paula, J. *Atkins' Physical Chemistry*. 10th edn. (Oxford Univ. Press, 2014). [Arrhenius rate-constant calibration: $E_a = 40/60/70/50$ kJ·mol^{-1}, $T = 500$ K.] Mishkin, F. S. *The Economics of Money, Banking, and Financial Markets*. 11th edn. (Pearson, 2016). [Money multiplier $m = 1/r$ and bank-leverage reference values.] Poltorak, A. et al. Defective LPS signaling in C3H/HeJ and C57BL/10ScCr mice: mutations in the *Tlr4* gene. *Science* **282**, 2085–2088 (1998). doi:10.1126/science.282.5396.2085. [TLR4 inflammatory positive feedback, amplification $T \approx 25$.] Furness, J. B. *The Enteric Nervous System*. (Wiley-Blackwell, 2006). [ENS–LPS–TT physiological coupling parameters / mucosal-barrier loop.] Liu, Z., Li, Y., Han, L. et al. PDB-wide collection of binding data: the PDBbind database (v2020 refined set, 5,316 complexes). *Nucleic Acids Research* (Database issue). https://www.pdbbind.org.cn. [Independent measured $-\log K_d/K_i$ distribution for the four-step catalytic cross-check: median 6.42, IQR 4.93–7.74, spanning ≈ 10 orders of magnitude.] Batista, C. et al. Role of endogenous lipopolysaccharides in neurological disorders. *Cells* **11**, 1770 (2022). doi:10.3390/cells11111770. PMID 36552802. [Barrier

dysfunction → LPS influx → systemic endotoxaemia; a_1 range anchor.] Long, H. et al. The TLR4 D299G and T399I SNPs are constitutively active to up-regulate expression of Trif-related adaptor molecule. *PLoS One* **9**, e105076 (2014). doi:10.1371/journal.pone.0105076. PMID 25365308. [WT vs D299G: LPS-induced NF-κB reporter activation $\approx 12\times$ higher; a_2 range anchor.] Jo, S. et al. Age-dependent shift in macrophage polarisation causes inflammation-mediated degeneration of enteric neurons. *Gut* **67**, 1081–1090 (2018). doi:10.1136/gutjnl-2017-314165. PMID 28228489. [Age-dependent inflammation-mediated enteric neurodegeneration; a_3 range anchor.] Lindahl, M. et al. Cerebral dopamine neurotrophic factor-deficiency leads to degeneration of enteric neurons. *Neurobiol. Dis.* **144**, 105080 (2020). doi:10.1016/j.nbd.2020.105080. PMID 31783118. [Age-dependent enteric neuron loss; a_3 anchor.] Buergel, T. et al. Metabolomic profiles predict individual multidisease outcomes. *Nat. Med.* **28**, 2309–2320 (2022). doi:10.1038/s41591-022-01980-3. PMID 36289334. [UK Biobank $n = 117{,}981$; 168 circulating metabolites; 24 diseases; 10-yr individual-level C-index/AUC improvement (existence anchor).] UK Biobank Metabolomics Consortium (You, J. et al.). Metabolome Phenome Atlas: circulating metabolites and future disease across 274,241 UK Biobank participants. *Nat. Metab.* **7**, 2366–2384 (2025). doi:10.1038/s42255-025-01371-1. [$n = 274{,}241$; 317 metabolites; 97.8% change > 10 yr before ≥ 1 disease, 96.2% 5–10 yr before; time-scale separation anchor.]

Declarations

Data availability. All data used in this study are publicly available. Public databases are read locally by the pipeline: PDBbind v2020 refined set , CHARLS (China Health and Retirement Longitudinal Study) , and the FDIC bank-failure record . Quantitative results taken from published studies (Buergel 2022 , You 2025 , and the literature-anchored parameter ranges in SI §1) are entered as explicit constants from the cited papers; the source citations are given in situ and in the References. No privately held or otherwise unavailable dataset is required to reproduce any figure in the paper.

Code availability. The four scripts that reproduce the paper-derived calibrations and the individual-level empirical anchor of this paper (SI §1 and §2.7) are publicly available at https://github.com/genlinjiao/cross-layer-invariant: `c001_bio_L_paper_derived.py`, `c002_chem_L_pdbbind_calibration.py`, `c004_cross_layer_amplification_scan.py`, and `cross_layer_empirical_anchor.py`. These run against publicly available data (PDBbind v2020 , CHARLS , and published values entered as explicit constants from and the SI §1 ranges); no wet-lab experiment is required. The model internal-consistency scans (the closed-form $\varepsilon^* = L^{1/n}$ numerics, SI §1) and the organisational weakest-link λ_2 analysis (§2.6) operate on internal project data (a knowledge-carrier graph and four in-house code subsystems), which are not released in this repository; their method is described in Methods §4, and the invariant itself is a closed-form linear-algebraic identity verifiable on paper. A single entry point (`run_all.py`) reruns the four released public-data scripts; intermediate outputs are written to `md_p1_work/` as JSON, with filenames cited in the text.

Competing interests. The author declares no competing interests.

Author contributions. G.J. conceived the framework, derived the identities, wrote the scripts, ran the calibrations, and wrote the manuscript. (Single author.)

Acknowledgements. The manuscript was prepared with editorial and computational assistance from an AI assistant (drafting, independent calibration calculations), disclosed in accordance with the journal's AI policy; no AI tool is listed as an author.