

Six Birds: Foundations of Emergence Calculus

Ioannis Tsiokos
`ioannis@automorph.io`

25 January 2026

© 2026 Automorph Inc.

Abstract

We present a math-only framework in which *theories* arise as fixed points of idempotent operators and *open-endedness* requires changing the operator rather than iterating it. Under minimal hypotheses of composability plus limited interface access with bounded observation, we show that the operator vocabulary P1–P6 arises canonically as unavoidable closure mechanics (Theorem 7). Order-theoretic closure operators and dynamics-induced endomaps $E_{\tau,f}$ (built from a Markov kernel, a coarse-graining lens, and a timescale) are treated uniformly via the abstraction of idempotent endomaps and their fixed points. We define a computable total-variation idempotence defect for $E_{\tau,f}$ and show that small retention error implies approximate idempotence, yielding robust packaged “objects” at a given time-scale τ within a fixed lens/theory. To prevent spurious directionality, we define arrow-of-time as path-space KL divergence and prove that coarse-graining cannot increase it (data processing); we further formalize the protocol-trap principle showing that autonomous protocol holonomy does not generate sustained asymmetry unless a genuine affinity is present in the lifted dynamics. Finally, we give a finite forcing-style counting lemma: definable predicates are exponentially rare relative to a partition-based theory, so generic predicate extensions are overwhelmingly non-definable, providing a mathematically clean anti-saturation mechanism for strict ladder climbing.

Contents

1	Introduction	4
1.1	The organizing picture: a three-certificate loop	6
2	Related work	7
3	Conventions and notation	9
3.1	Finite state spaces, distributions, and kernels	9
3.2	Paths, time reversal, and relative entropy	9
3.3	Coarse-graining lenses	10
3.4	A unified theory package viewpoint	10
3.5	A minimal instantiation recipe	11
3.6	Support graphs and discrete 1-forms	13
3.7	Assumption bundles	13
4	Order-closure and closure ladders	14
4.1	Order-theoretic closure and fixed points	14
4.2	Closure ladders and saturation	15
5	Idempotent endomaps and induced closures	16
5.1	Idempotent endomaps	16
5.2	Dynamics-induced empirical endomaps	17
6	AUT + REV + ACC regime and graph 1-forms	21
6.1	Bidirected support and the log-ratio 1-form	21
6.2	Cycle integrals, exactness, and the null regime	21
6.3	Accounting as coordinates on cycle space	23
7	Arrow-of-time as path reversal asymmetry; data processing; protocol trap	23
7.1	Data processing: coarse-graining cannot create asymmetry	24
7.2	Protocol trap: apparent stroboscopic arrows and the “clock audit”	25
8	Generic extension and the finite forcing lemma	27
8.1	Theories as partitions and definability	27
8.2	Counting lemma: definable predicates are rare	27
8.3	Finite forcing: generic extensions are non-definable	28
9	Why the primitives are unavoidable	29

10 Primitives P1–P6 as closure-changing operations	31
10.1 Definitions of P1–P6	31
10.2 How the primitives compose to generate theory growth	32
10.3 Downward influence across theories	34
10.4 Two load-bearing propositions	35
11 Examples	36
11.1 A biased three-cycle: a non-exact graph 1-form	36
11.2 Protocol trap: external schedule vs autonomous lifted model	36
11.3 Finite forcing count: definability is exponentially rare	37
12 Discussion and scope boundary	38
12.1 What the theory does and does not claim	38
12.2 Outlook: forthcoming instantiations	38
A Appendix A: Reproducibility checklist	39
B Appendix B: Lean formalization map	40
C Appendix C: Python evidence map	41
D Appendix D: Zeno cascades and depth	44
D.1 Setup: frontier and Zeno criterion	44
D.2 Bridges, ports, and accounting	44
D.3 Storage-based activity and the WORK quantum (Option B) .	44
D.4 Integrated throughput (ICAP) and feasibility \Rightarrow latency bound	45
D.5 No-Zeno criterion via divergence	46
D.6 Hard lemma slots (WORK/CAP/route)	47
D.7 Checkable divergence criteria	48
D.8 Toy model families (necessity witnesses)	48
D.9 Decision tree (settlement frontier)	49
E Appendix E: Toolkit theory—defects	49
E.1 Defects as quantitative relaxations of exact laws	49
E.2 Bridge objects: ports, passivity, and integrated throughput .	51
E.3 Emergent coercivity template via sector compression	54
E.3.1 Summary: slots and divergence consequence	54
E.3.2 Dissipative atoms and semigroup decay	55
E.3.3 Coercivity from feasibility gating (P2)	59
E.3.4 Mode compression from sectorization (P4) and mini- mality (P5)	60

1 Introduction

This paper isolates a finite, discipline-agnostic *emergence calculus*: a minimal operator toolkit for when a coarse description becomes a *theory* with stable objects, how strict theory extension remains possible, and how directionality audits stay honest under observation. A new structural result makes this toolkit non-ad hoc: Theorem 7 shows that once processes are composable but interface-limited, the operator vocabulary P1–P6 arises canonically as unavoidable closure mechanics of description (we refer to these as the “six birds”). The organizing principle is to distinguish *completion* (idempotent endomaps with fixed points) from *extension* (refinements of what is definable), and to keep directionality certificates strictly tied to audit functionals that are monotone under coarse observation. Within that scope we prove a small spine of claims: closure saturation, data processing for path reversal asymmetry, a protocol-trap audit for hidden schedules, and a finite forcing-count lemma that quantifies definability rarity.

The manuscript is structured to keep assumptions explicit and finite. After conventions, we develop order-closure and idempotent endomaps, connect them to dynamics-induced operators on the simplex, and then turn to accounting and arrow-of-time diagnostics. We treat definability/forcing in finite theory terms and collect primitives P1–P6 as closure-changing operations. The appendices record reproducibility checks, a Lean map, a Python evidence map, a Zeno decision frontier, and a toolkit appendix that packages defect calculus, route mismatch, and bridge objects for reuse. A compact interaction diagram for how the six primitives fit into the rest of the calculus appears in Section 10.2. We then make explicit how the same toolkit expresses *downward influence* (macro-to-micro constraint within a coupled theory package) in Section 10.3.

Remark 1 (Interoperability of certificates (toolkit safety)). This manuscript uses three logically distinct “certificate families”: *stability* (existence/robustness of fixed points of a completion rule), *novelty* (strict theory extension by adding non-definable predicates), and *directionality* (nontrivial audit values such as affinities or path reversal asymmetry). They are intentionally separated to prevent accidental overinterpretation. The three certificate families (Stability, Novelty, Directionality) are logically independent. Each can hold or fail independently of the others. The only designed interactions are: (C1) completion defines stable objects, (C2) extension changes definability, (C3) audit is monotone under coarse-graining. No other implications are intended or valid.

Non-implications.

- Stability (small idempotence defect / approximate closure) does *not* imply novelty: a system may stabilize within a fixed definability structure (Section 5.1) without any strict theory extension (Section 8).
- Novelty does *not* imply directionality: adding a generic predicate (Section 8) can change what is expressible while leaving the audit certificates unchanged.
- Stability and novelty do *not* imply directionality: directionality is certified only by the audit functional and its monotonicity laws (Sections 6 and 7).

No false positives for directionality under observation. Coarse-graining cannot create arrow-of-time: the path reversal asymmetry contracts under pushforward (data processing in Section 7). Apparent irreversibility caused by hiding an internal protocol/phase is diagnosed by the protocol-trap/clock-audit pattern (Section 7).

Where coupling actually occurs. The only intended coupling mechanisms are explicit operators: (i) *completion* changes fixed-point sets (Sections 4 and 5.1), (ii) *extension* refines definability (Section 8), and (iii) the audit functional is functorial/monotone under lenses (Sections 6 and 7). Appendix E’s ECT template is an optional coupling that links completion/accounting to a throughput/capacity certificate under explicit slots; it still does not certify novelty or directionality. Appendix E’s balanced-atom route (Definitions 32–33 and Corollary 6) is one optional way to couple packaging/accounting to a throughput certificate; it still does not certify novelty (Section 8) or directionality (Sections 6–7). Route dependence (holonomy) is treated separately as a reduction-order effect and is not, by itself, a directionality certificate.

1.1 The organizing picture: a three-certificate loop

The results of this paper can be read as a minimal calculus for building hierarchical structure while keeping directionality honest. The core observation is that three notions often conflated in informal discussions are logically distinct and have distinct mathematical *certificates*:

- **Stability / objects (completion).** A theory level is specified by an idempotent endomap (exact or approximate); its fixed points are the “objects” at that theory level. In order theory this is a closure operator; in dynamics it is the induced endomap $E_{\tau,f}$. The certificate is (approximate) idempotence, quantified by the TV defect $\delta_{\tau,f}$.
- **Novelty / open-endedness (theory extension).** Iterating a fixed completion rule saturates by idempotence. Strict ladder climbing therefore requires changing the theory/completion rule. The certificate is *non-definability*: a new predicate is not measurable with respect to the old partition-based theory. Our finite forcing lemma makes this generic: when there is hidden volume ($N - K > 0$), uniformly random predicate extensions are overwhelmingly non-definable.
- **Directionality / irreversibility (audit).** Arrow-of-time is defined on path space as $\Sigma_T(\rho) = D_{\text{KL}}(\mathbb{P}_{\rho,T} \| \mathcal{R}_* \mathbb{P}_{\rho,T})$. The certificate is either positive path-space KL or, under the accounted regime, a non-exact log-ratio 1-form (nonzero cycle affinities). Data processing ensures coarse-graining cannot create this certificate, and the protocol-trap theorem ensures it is not an artifact of hiding an internal clock.

These certificates compose into a simple loop: choose a theory (lens) f , induce a packaging endomap $E_{\tau,f}$ from dynamics and timescale, extract its (approximate) fixed points as objects, then adjoin a generic non-definable predicate to obtain a strict theory extension and repeat—while auditing any apparent arrow-of-time via path-space KL and affinity. The main text develops each piece of this loop in a representation-independent way and keeps all nonessential computational evidence in the appendices.

Executive map (emergence calculus). A *theory* in this paper is a package of (i) a lens/definability structure, (ii) a completion rule, and (iii) an audit. The six primitives are the minimal moves that act on these components: P5 selects objects as fixed points of a completion, P6 supplies audits/monotones that remain honest under observation, P4 provides a

bounded theory index for coherent refinement, P2 carves feasibility and representability, P1 rewrites dynamics when closure fails to descend. P3 diagnoses noncommuting reduction routes (route mismatch); this is *not* a directionality certificate. Emergence means the appearance of stable fixed points within a theory. Open-endedness means strict theory extension rather than iteration of a fixed closure. Section 10.2 makes this interaction loop explicit.

Roadmap. Section 4 develops the order-closure backbone and the saturation principle. Section 5.1 isolates idempotent endomaps as the unifying abstraction and introduces the induced empirical endomap $E_{\tau,f}$ with a computable idempotence defect and a qualified theory-birth criterion. Section 6 formulates **A_AUT+A_REV+A_ACC** as a graph 1-form exactness statement and defines affinities. Section 7 defines arrow-of-time as path reversal asymmetry, proves data processing, and formalizes the protocol-trap/clock-audit principle (including the corrected autonomous reading of protocol holonomy). Section 8 gives the finite forcing lemma as an anti-saturation move for strict theory extension. Section 10 summarizes the primitives P1–P6 as closure-changing operations, and Section 11 gives a few surgical examples. Section 10.2 ties the primitives into the closure/forcing/audit spine as an explicit emergence calculus.

2 Related work

This paper is written to be largely self-contained; we briefly situate standard ingredients that we use as mathematical tools.

Closure operators, reflections, and idempotents. Closure operators on posets/lattices are classical objects in order theory: extensive, monotone, and idempotent endomaps whose fixed points form a “closed” substructure. In categorical language, (order-theoretic) closure operators are the thin-category shadow of reflective subcategories and idempotent monads: a reflector is left adjoint to the inclusion of the full subcategory of fixed points, and the universal property characterizes the “least closed object above” a given one. Our main text uses the order-theoretic formulation, and the Lean scaffold mechanizes the thin (poset) instance. See, e.g., [DP02, ML98, Bor94].

Coarse-graining of Markov dynamics and lumpability. State aggregation and lumpability for finite Markov chains provide exact hypotheses under which a coarse observation of a Markov chain is itself Markov (or approximately so after subsampling). We use coarse-graining primarily as a measurable map (a “lens”) and emphasize statements that are valid without assuming exact lumpability (notably, data-processing for relative entropy). When we do speak about induced macro-dynamics, the relevant classical references are the standard Markov chain texts that treat aggregation/lumpability and time reversal; see [KS60, Nor97, RP81]. For approximate lumpability and metastable coarse objects in finite-state settings, see e.g. [KADK19, NJ10].

Path-space KL and DPI. Relative entropy/KL divergence and the data processing inequality are standard in information theory: pushforward under a measurable map cannot increase KL. We apply this at the level of path measures, so that coarse observation cannot create time-reversal asymmetry. The use of path-space relative entropy between forward and time-reversed Markov path measures as a measure of irreversibility also appears in the literature on reversible Markov chains and (more recently) stochastic thermodynamics; we cite these only as orientation since our arguments are finite and measure-theoretic. See, e.g., [CT06, CK11, Kel79, Sei12, Mer11, PPRV10] and recent coarse-graining/hidden-drive analyses such as [DvdMS23, BSvdM25].

Graph cycles, affinities, and nonequilibrium network structure. The cycle/cocycle decomposition viewpoint for Markov networks and the role of affinities as cycle integrals are standard in nonequilibrium statistical mechanics; see, e.g., [Pol12, AGH⁺12, DCLP23].

Protocol geometry and stochastic pumps. Geometric-phase currents and stochastic-pump effects in driven Markov systems provide a related route-dependence perspective on nonequilibrium protocols; see [SN07].

Definability rarity and the forcing analogy. In our finite setting, “definable from a theory” is treated as measurability with respect to a finite σ -algebra (equivalently: being constant on the blocks of a partition). Our use of “theory” is intentionally modest: it is a finite definability structure (partition / finite σ -algebra), not an infinitary logical theory. Our “finite forcing lemma” is a counting statement: when blocks have nontrivial internal

volume, predicates that are definable from the coarse theory form an exponentially small subset of all predicates, so a random predicate extension is generically non-definable. The term “forcing” is used only as an analogy to the role of generic extensions in set theory; our results are finite and do not invoke infinitary semantics. Background references include standard finite model theory and set-theory texts, e.g. [Lib04, Coh63, Kun11, MS11, Hyt15].

3 Conventions and notation

3.1 Finite state spaces, distributions, and kernels

Throughout, Z denotes a finite set (the microstate space). We write $\Delta(Z)$ for the probability simplex on Z . We represent distributions as row vectors $\mu \in \Delta(Z)$ and Markov kernels as row-stochastic matrices P so that the time-1 update is $\mu \mapsto \mu P$. All logarithms are natural logs. We use *theory* (and *theory level*) as the primary term throughout.

3.2 Paths, time reversal, and relative entropy

Fix a horizon $T \in \mathbb{N}$. A path is $z_{0:T} := (z_0, z_1, \dots, z_T) \in Z^{T+1}$. Given an initial distribution $\rho \in \Delta(Z)$ and a time-homogeneous kernel P , the forward path law is

$$\mathbb{P}_{\rho, T}(z_{0:T}) = \rho(z_0) \prod_{t=0}^{T-1} P(z_t, z_{t+1}).$$

Let $\mathcal{R} : Z^{T+1} \rightarrow Z^{T+1}$ be path reversal, $\mathcal{R}(z_0, \dots, z_T) = (z_T, \dots, z_0)$, and let $\mathcal{R}_* \mathbb{P}_{\rho, T}$ denote the pushforward law.

For probability measures p, q on a finite set, we use the KL divergence

$$D_{\text{KL}}(p||q) := \sum_x p(x) \log \frac{p(x)}{q(x)},$$

with the conventions $0 \log(0/q) := 0$ and $p(x) > 0, q(x) = 0 \Rightarrow D_{\text{KL}}(p||q) = +\infty$.

The finite-horizon path reversal asymmetry $\Sigma_T(\rho)$ is defined in Section 7 as the KL divergence between $\mathbb{P}_{\rho, T}$ and its time-reversed pushforward $\mathcal{R}_* \mathbb{P}_{\rho, T}$.

Remark 2 (Stationarity is not needed for Σ_T or data processing). The definition in Section 7 makes sense for any initial distribution ρ (it may be $+\infty$ if reverse-support is missing). The data processing inequality for KL is

purely measure-theoretic and therefore applies to $\Sigma_T(\rho)$ without any stationarity assumption. Stationarity enters only when one wants a steady-state interpretation (e.g. relating $\Sigma_T(\pi)$ to entropy production and/or considering long-time rates).

3.3 Coarse-graining lenses

A (deterministic) coarse-graining or *lens* is a map $f : Z \rightarrow X$ to a finite set X . The pushforward of $\mu \in \Delta(Z)$ is $(f_*\mu)(x) := \sum_{z:f(z)=x} \mu(z)$. On path space we use the coordinatewise lens $f^{T+1} : Z^{T+1} \rightarrow X^{T+1}$, and write $f_*\mathbb{P}_{\rho,T}$ for the induced macro path law.

Remark 3. WLOG, we replace X by the image $f(Z)$ so that f is surjective and every fiber $B_x = f^{-1}(x)$ is nonempty. This does not change the induced σ -algebra/partition or any pushforward laws. We will often refer to a lens $f : Z \rightarrow X$ (equivalently, the induced finite partition / σ -algebra) as a *finite observational theory* on Z : it specifies which predicates are expressible and which distinctions are invisible. Refining f by adjoining a predicate h will be called a *theory extension*.

3.4 A unified theory package viewpoint

Remark 4 (Emergence and open-endedness (terminology)). Throughout, *emergence* means the appearance of stable fixed points of a completion rule within a theory package, while *open-endedness* means strict theory extension (a change in definability or completion), not iteration of a fixed closure. These are structural notions; no claim is made that either occurs in all systems.

Definition 1 (Finite theory package). *Assume **A_FIN**. A finite theory package is a tuple*

$$\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$$

where:

- Z is a finite carrier set (microstate space).
- $f : Z \rightarrow X$ is a lens (coarse description). The induced definability structure Σ_f is the finite σ -algebra (equivalently, partition) on Z generated by f , i.e. the family of predicates measurable with respect to f .

- $E : \mathcal{V} \rightarrow \mathcal{V}$ is a completion/packaging endomap on a chosen space of descriptions \mathcal{V} associated to Z . Typically E is idempotent ($E \circ E = E$) or approximately idempotent in a specified norm. In the empirical/dynamical instantiation $\mathcal{V} = \Delta(Z)$, while in the order-theoretic instantiation $\mathcal{V} = Z$ (with E a closure operator).
- \mathcal{A} is an audit functional (a certificate of directionality or drive) defined on the relevant objects in context (e.g. path laws for a Markov dynamics, or antisymmetric edge 1-forms on a support graph), together with a designated monotonicity principle under coarse observation: for any further coarsening $g : X \rightarrow Y$, the induced pushforward should not increase the audit, schematically

$$\mathcal{A}((g \circ f)_{\#}(\cdot)) \leq \mathcal{A}(f_{\#}(\cdot)),$$

whenever both sides are defined.

We call Σ_f the expressive content of the theory and E its completion rule; the fixed points of E are the internally complete descriptions/objects recognized by the theory at the chosen scale and lens.

Remark 5. We use the symbol Σ in two unrelated roles: Σ_f denotes the definability σ -algebra induced by a lens f , while $\Sigma_T(\rho)$ denotes the finite-horizon path reversal asymmetry defined in Section 7. Context (and subscripts) distinguish these usages.

Remark 6 (Instantiations in this manuscript). The components of \mathcal{T} appear throughout the paper in specialized forms. In Section 4, an order-closure operator c on a poset is an idempotent completion rule; fixed points of c are the “objects” at that rung of a closure ladder. In Section 5.1, we isolate idempotent endomaps as the minimal abstraction behind this fixed-point viewpoint. In Section 5.1 we also construct, from a Markov kernel, a lens f , and a timescale, a dynamics-induced completion endomap $E_{\tau,f}$ on $\Delta(Z)$. In Section 8, adjoining a new predicate corresponds to refining Σ_f to a strictly stronger definability structure, i.e. a theory extension. Finally, Sections 6 and 7 provide two audit families used here: graph 1-form exactness/affinities (ACC) and path reversal asymmetry (arrow-of-time), the latter enjoying a precise monotonicity principle under coarse-graining (data processing).

3.5 A minimal instantiation recipe

To apply the emergence calculus to an unrelated multiscale system, one supplies (at minimum) the following five pieces of data.

1. **Micro description space.** A finite (or otherwise well-posed) space Z of microstates, histories, or happenings, together with the ambient space of descriptions \mathcal{V} used in the domain (e.g. Z itself with an order, or $\Delta(Z)$ for distributional descriptions).
2. **Lens/refinement family.** A lens $f : Z \rightarrow X$ (or a refinement family $f_j : Z \rightarrow X_j$) specifying which distinctions are observable/definable at the current theory. This induces the definability structure (written Σ_f in the finite setting).
3. **Completion/packaging rule.** A completion endomap $E : \mathcal{V} \rightarrow \mathcal{V}$ (often idempotent or approximately idempotent) whose fixed points are the theory's *objects* (stable packaged descriptions).
4. **Audit functional.** An audit \mathcal{A} that is functorial/monotone under lenses (so it does not generate false positives under coarse observation). Examples in this paper include the ACC 1-form affinities (Section 6) and finite-horizon path reversal asymmetry (Section 7).
5. **Extension/update move.** A rule that changes the theory package (typically by extending definability, changing the completion rule, or rewriting admissible dynamics). Open-endedness is never obtained by iterating a fixed completion; it requires such strict extension (Section 8).

Given these inputs, the three core certificates used throughout the paper can be read as pass/fail tests:

Emergence (stability/objecthood). Pass when the completion has multiple robust fixed points (objects) and the idempotence defect is small at the intended scale (Section 5.1 and the empirical endomap construction).

Open-endedness (novelty). Pass when a theory extension is *strict* in the definability sense (new predicates are not definable from the old lens), so that closing again can yield genuinely new objects; the finite forcing lemma quantifies that such strict extensions are generic (Section 8).

Directionality (audit). Pass when the chosen audit is nontrivial in the full model (e.g. nonzero affinity class or nonzero path reversal asymmetry) and remains consistent under observation (data processing). Apparent directionality caused by hiding protocol/state variables is ruled out by the protocol-trap analysis (Section 7).

Operationally, one applies the toolkit by iterating: choose (f, E, \mathcal{A}) for a candidate theory, identify its objects (fixed points), audit consistency under coarse observation, and then perform an explicit extension/rewrite step to obtain the next theory package. Appendices collect reproducibility artifacts and sanity checks for the finite Markov and finite definability instantiations used as running examples.

3.6 Support graphs and discrete 1-forms

Given a kernel P on Z , its directed support has an edge $z \rightarrow z'$ whenever $P(z, z') > 0$. When microreversibility holds (support is bidirected), we also use the induced undirected support graph with edge $\{z, z'\}$ whenever $P(z, z') > 0$ and $P(z', z) > 0$.

On bidirected support edges we use the antisymmetric edge 1-form

$$a(z, z') := \log \frac{P(z, z')}{P(z', z)}.$$

For a directed cycle $\gamma = (z_0 \rightarrow z_1 \rightarrow \dots \rightarrow z_m = z_0)$ in the bidirected support, its *cycle integral* is

$$\oint_{\gamma} a := \sum_{i=0}^{m-1} a(z_i, z_{i+1}).$$

If a is exact, i.e. there exists $\Phi : Z \rightarrow \mathbb{R}$ with $a(z, z') = \Phi(z') - \Phi(z)$ on all bidirected edges, then $\oint_{\gamma} a = 0$ for every cycle γ .

3.7 Assumption bundles

We will cite standing hypotheses by short tags (e.g. **A_FIN+A_LENS**) rather than repeating prose. Use the following as the canonical assumption bundle list in this manuscript (tag spellings are aligned with `docs/spec/theorem-inventory.md`):

A_FIN Finite combinatorial setting: underlying state/index sets are finite (so sums/KL on finite alphabets are unambiguous), and when we treat completion/packaging maps as linear operators they are finite-dimensional matrices.

A_AUT Autonomy: time-homogeneous dynamics; phase variables are part of the state (no external schedule).

A_REV Microreversibility on support: $P(z, z') > 0 \Rightarrow P(z', z) > 0$ (bidirected support).

A_ACC Accounting/1-form decomposition: the antisymmetric log-ratio 1-form admits a decomposition into an exact part plus fixed affinity components (graph H^1 viewpoint).

A_NULL Null regime: all affinity coordinates / cycle integrals vanish (e.g. all basis affinities $A_r = 0$).

A_STAT A stationary distribution π exists (full support when required by context).

A_LENS A specified coarse-graining lens $f : Z \rightarrow X$ is part of the statement.

4 Order-closure and closure ladders

4.1 Order-theoretic closure and fixed points

We begin with the minimal order-theoretic notion of closure. This section is intentionally agnostic about dynamics, probability, or measurement; those enter later. The only idea we need is that a closure operator defines a notion of “completion,” and its fixed points constitute the “objects” of the theory determined by that completion rule.

Definition 2 (Closure operator). *Let (L, \leq) be a poset. A map $c : L \rightarrow L$ is a closure operator if for all $x, y \in L$:*

1. (*Extensive*) $x \leq c(x)$.
2. (*Monotone*) $x \leq y \Rightarrow c(x) \leq c(y)$.
3. (*Idempotent*) $c(c(x)) = c(x)$.

For a preorder, the same discussion applies after quotienting by the induced equivalence relation $x \sim y \iff x \leq y \leq x$. We use posets to identify fixed points by equality.

Definition 3 (Closed points / objects). *Given a closure operator $c : L \rightarrow L$, its closed points (or objects of the theory determined by c) are the fixed points*

$$\text{Fix}(c) := \{x \in L : c(x) = x\}.$$

The order on L induces a natural order on closure operators: we write $c \leq d$ if $c(x) \leq d(x)$ for all $x \in L$ (pointwise order). Intuitively, d is “stronger” than c if it closes more aggressively.

Lemma 1 (Closed points are antitone in closure strength). *If $c \leq d$ are closure operators on L , then $\text{Fix}(d) \subseteq \text{Fix}(c)$. Equivalently: if $d(x) = x$ then $c(x) = x$.*

Proof. Let $x \in \text{Fix}(d)$, so $d(x) = x$. By extensiveness, $x \leq c(x)$. By $c \leq d$ we have $c(x) \leq d(x) = x$. Hence $x \leq c(x) \leq x$, so $c(x) = x$. \square

4.2 Closure ladders and saturation

The Box is the Thing. Idempotence implies that repeatedly applying a *fixed* completion rule cannot yield unbounded novelty.

Lemma 2 (One-step stabilization under iteration). *Let c be a closure operator on L . Define the iterate $c^{(n)}$ by $c^{(0)} = \text{id}$ and $c^{(n+1)} = c \circ c^{(n)}$. Then for all $x \in L$ and all integers $n \geq 1$,*

$$c^{(n)}(x) = c(x).$$

Proof. For $n = 1$ the statement is trivial. If $c^{(n)}(x) = c(x)$ for some $n \geq 1$, then $c^{(n+1)}(x) = c(c^{(n)}(x)) = c(c(x)) = c(x)$ by idempotence. \square

Corollary 1 (Closure saturates). *Fix a closure operator c . For any $x \in L$, the sequence $x, c(x), c(c(x)), c(c(c(x))), \dots$ stabilizes after one closure step: it is constant from the second term onward.*

Remark 7. Viewed abstractly, a closure operator acts like a completion rule. Iterating a fixed completion rule stabilizes. Thus strict ladder growth requires changing the completion rule (equivalently, passing to a stronger theory).

Thus, any notion of open-endedness based purely on repeated application of a single completion rule is structurally limited. To obtain a strict hierarchy of “new objects,” one must change the completion rule.

Definition 4 (Strictly stronger closure). *For closure operators c, d on L , write $c \prec d$ (“ d is strictly stronger than c ”) if $c \leq d$ and there exists an $x \in L$ such that $d(x) \neq c(x)$.*

Definition 5 (Closure ladder). A closure ladder is a sequence of closure operators $(c_n)_{n \in \mathbb{N}}$ on L such that $c_n \prec c_{n+1}$ for all n . The fixed-point sets form a nested chain

$$\text{Fix}(c_0) \supseteq \text{Fix}(c_1) \supseteq \text{Fix}(c_2) \supseteq \dots$$

by Lemma 1.

Remark 8. The strictness relation \prec is irreflexive and transitive on closure operators (by the pointwise order), and Lemma 1 records the antitone behavior of fixed points. On a fixed finite lattice, any \prec -chain is finite; the open-endedness modeled here comes from changing the theory/closure package or enlarging domains, not from a single finite lattice. Viewed abstractly, a closure operator acts like a completion rule. Iterating a fixed completion rule stabilizes. In a poset, closure operators correspond to reflections onto their fixed points (see Remark 9).

Remark 9 (Lean formalization note). The order-theoretic closure operator used here is formalized in the repository using mathlib’s `ClosureOperator`. The basic closure and ladder API lives in `formal/ClosureLadder/Basic.lean`. In particular, the iterate-stabilization statement Lemma 2 corresponds to Lean lemmas `closure*iterate*succ` and `closure*iterate*ge*one`, and ladder monotonicity corresponds to `ladder*mono`.

5 Idempotent endomaps and induced closures

5.1 Idempotent endomaps

The common mathematical nucleus behind both order-theoretic closures and the dynamics-induced operators used later is *idempotence*. In this section we package that nucleus in a way that avoids conflating two different notions: *order-closures* (which require an ambient poset) versus *idempotent endomaps* (which require no order at all).

Definition 6 (Idempotent endomap). Let A be a set (or type). An idempotent endomap on A is a function $e : A \rightarrow A$ such that $e \circ e = e$, i.e. $e(e(a)) = e(a)$ for all $a \in A$.

Definition 7 (Fixed points). For an endomap $e : A \rightarrow A$ define its fixed-point set

$$\text{Fix}(e) := \{a \in A : e(a) = a\}.$$

We interpret $\text{Fix}(e)$ as the collection of “objects” completed by e .

Lemma 3 (Idempotents split). *Let $e : A \rightarrow A$ be an idempotent endomap. Let $i : \text{Fix}(e) \hookrightarrow A$ denote inclusion, and define $r : A \rightarrow \text{Fix}(e)$ by $r(a) := e(a)$ (viewed as an element of $\text{Fix}(e)$). Then $r \circ i = \text{id}_{\text{Fix}(e)}$ and $i \circ r = e$. In particular, e restricts to the identity on $\text{Fix}(e)$ and $\text{Fix}(e)$ canonically identifies with the image of e .*

Proof. For $a \in A$, idempotence gives $e(e(a)) = e(a)$, so $e(a) \in \text{Fix}(e)$ and r is well-defined. For $x \in \text{Fix}(e)$ we have $r(i(x)) = e(x) = x$, hence $r \circ i = \text{id}$. Finally, for $a \in A$, $(i \circ r)(a) = i(e(a)) = e(a)$ by definition of i , so $i \circ r = e$. \square

Remark 10 (Order-closures as a special case). If L is a poset and $c : L \rightarrow L$ is an order-closure (extensive, monotone, idempotent), then forgetting the order yields an idempotent endomap on the underlying set. The converse is false: an idempotent endomap need not be monotone or extensive with respect to any nontrivial order. In particular, the dynamics-induced operators $E_{\tau, f}$ (defined later on probability simplices) are treated as (*approximate*) *idempotent endomaps* with fixed points and idempotence defects measured in a chosen metric, rather than as order-closures. We will sometimes call an idempotent endomap a *semantic completion map* for a theory: its fixed points are the internally complete descriptions with respect to that completion.

Remark 11 (Formalization note). The constructions in this subsection are formalized in Lean:

- `formal/ClosureLadder/IdempotentEndo.lean`
- Definitions: `ClosureLadder.IdempotentEndo`
`ClosureLadder.IdempotentEndo.FixedPoints.`
- Bridge lemma:
`ClosureLadder.ClosureOperator.toIdempotentEndo.`

5.2 Dynamics-induced empirical endomaps

Throughout this section we assume **A_FIN** and fix a finite state space Z . Let P be a (row-stochastic) Markov kernel on Z , and let $\Delta(Z)$ denote the simplex of probability distributions on Z (row vectors). Fix a *lens* $f : Z \rightarrow X$ (assumption **A_LENS**), with fibers (blocks) $B_x := f^{-1}(x)$.

Definition 8 (Coarse map and canonical lift; **D-IC-01**). Let $Q_f : \Delta(Z) \rightarrow \Delta(X)$ be the pushforward (coarse-graining) map

$$(Q_f \mu)(x) := \sum_{z \in B_x} \mu(z).$$

Fix a canonical lift $U_f : \Delta(X) \rightarrow \Delta(Z)$ specified by a choice of prototype distributions $(u_x)_{x \in X}$ with $\text{supp}(u_x) \subseteq B_x$ and $\sum_{z \in B_x} u_x(z) = 1$, by setting

$$(U_f \nu)(z) := \sum_{x \in X} \nu(x) u_x(z).$$

Equivalently, $U_f \delta_x = u_x$ and $Q_f U_f = \text{id}_{\Delta(X)}$.

Existence Requires Choosing a Scale. For an integer time scale $\tau \geq 1$, define the induced empirical endomap

$$E_{\tau, f} : \Delta(Z) \rightarrow \Delta(Z), \quad E_{\tau, f}(\mu) := U_f(Q_f(\mu P^\tau)).$$

The pair $(f, E_{\tau, f})$ may be viewed as an *empirical theory at scale τ* : f specifies the expressive content (definability), while $E_{\tau, f}$ specifies a semantics/completion rule whose fixed points are the objects recognized by the theory.

Remark 12 (Not an order-closure). The map $E_{\tau, f}$ is an affine endomap of $\Delta(Z)$ (indeed a Markov kernel on Z), and it is always idempotent as an *endomap* only in special regimes. We treat $E_{\tau, f}$ within the “fixed points of (approx) idempotents” framework (§5.1), not as an order-closure operator. See the repo spec note `docs/spec/empirical-vs-order-closure.md`.

Definition 9 (TV idempotence defect; **D-IC-02**). Let $\|\cdot\|_{\text{TV}}$ denote total variation on $\Delta(Z)$, i.e. $\|\mu - \nu\|_{\text{TV}} = \frac{1}{2}\|\mu - \nu\|_1$. The idempotence defect of $E_{\tau, f}$ is

$$\delta_{\tau, f} := \sup_{\mu \in \Delta(Z)} \|E_{\tau, f}(E_{\tau, f}(\mu)) - E_{\tau, f}(\mu)\|_{\text{TV}}.$$

Guardrail. A small idempotence defect $\delta(E)$ is a saturation diagnostic (reapplying E changes $E(\mu)$ only slightly), but it does not by itself certify non-trivial emergence or multiplicity (e.g., a constant map has $\delta(E) = 0$ but only one fixed point). Any “multiple objects” conclusion therefore requires a separate nontriviality condition (for example: at least two distinct stable labels/prototypes at the chosen tolerance). Since $\mu \mapsto \|\mu M\|_1$ is convex

for fixed M and $\Delta(Z)$ is the convex hull of the Dirac masses $(\delta_z)_{z \in Z}$, the supremum is attained on extreme points:

$$\delta_{\tau,f} = \max_{z \in Z} \|\delta_z(E_{\tau,f}^2 - E_{\tau,f})\|_{\text{TV}} = \frac{1}{2} \max_{z \in Z} \sum_{z' \in Z} |(E_{\tau,f}^2 - E_{\tau,f})(z, z')|.$$

Definition 10 (Prototype stability at scale τ). For each $x \in X$ define the (scale- τ) stability of the prototype u_x by

$$\begin{aligned} s_{\tau,f}(x) &:= \|E_{\tau,f}(u_x) - u_x\|_{\text{TV}} \\ &= \|U_f(Q_f(u_x P^\tau)) - U_f(\delta_x)\|_{\text{TV}}. \end{aligned}$$

We say u_x is ε -stable if $s_{\tau,f}(x) \leq \varepsilon$.

Theorem 1 (Approximate idempotence from retention). **T-IC-02.**

Assume **A_FIN+A_LENS** and fix $E_{\tau,f}$ as in Definition 8. Define a retention error

$$\varepsilon_{\tau,f} := \max_{x \in X} \|Q_f(u_x P^\tau) - \delta_x\|_{\text{TV}}.$$

Then:

1. (Prototype stability bound) For all $x \in X$, $s_{\tau,f}(x) \leq \varepsilon_{\tau,f}$.
2. (Idempotence defect bound) $\delta_{\tau,f} \leq \varepsilon_{\tau,f}$. In particular, if $\varepsilon_{\tau,f} = 0$ then $E_{\tau,f}$ is exactly idempotent.

Proof. Since U_f is a stochastic map, it is a contraction in total variation. Thus for each x ,

$$\begin{aligned} \|E_{\tau,f}(u_x) - u_x\|_{\text{TV}} &= \|U_f(Q_f(u_x P^\tau)) - U_f(\delta_x)\|_{\text{TV}} \\ &\leq \|Q_f(u_x P^\tau) - \delta_x\|_{\text{TV}} \leq \varepsilon_{\tau,f}. \end{aligned}$$

This proves (1).

For (2), write any $\mu \in \Delta(Z)$ as $\mu_1 := E_{\tau,f}(\mu) = \sum_{x \in X} w_x u_x$ where $w := Q_f(\mu P^\tau) \in \Delta(X)$ (so $w_x \geq 0$ and $\sum_x w_x = 1$). Then by linearity,

$$E_{\tau,f}(\mu_1) - \mu_1 = \sum_{x \in X} w_x (E_{\tau,f}(u_x) - u_x).$$

Taking TV norms and using the bound from (1),

$$\begin{aligned} \|E_{\tau,f}(\mu_1) - \mu_1\|_{\text{TV}} &\leq \sum_{x \in X} w_x \|E_{\tau,f}(u_x) - u_x\|_{\text{TV}} \\ &\leq \sum_{x \in X} w_x \varepsilon_{\tau,f} = \varepsilon_{\tau,f}. \end{aligned}$$

Now take the supremum over μ to obtain $\delta_{\tau,f} \leq \varepsilon_{\tau,f}$. □

Refinement is not monotone-good in general: it can reveal additional stable objects when aligned with metastable structure, but it can also destroy stability when it misaligns with mixing/leakage scales.

Theorem 2 (Refinement can help or hurt). *T-IC-01.*

Assume **A_FIN+A_LENS**. There exist finite Markov kernels P and times τ together with two lenses

$$f_{\text{fine}} : Z \rightarrow X_{\text{fine}}, \quad f_{\text{coarse}} : Z \rightarrow X_{\text{coarse}},$$

such that f_{coarse} is a coarsening of f_{fine} (the fine partition refines the coarse), and for the canonical lift U_f defining $E_{\tau,f}$ (Definition 8), the number of ε -stable prototypes can either increase or decrease under refinement for a fixed stability threshold $\varepsilon > 0$.

More precisely, both behaviors occur:

1. (Refinement reveals objects) in one example, refinement increases the number of ε -stable prototypes.
2. (Refinement destroys objects) in another example, the stable count decreases.

Proof idea by explicit two-block constructions. The mechanism is controlled by the retention errors $\varepsilon_{\tau,f}$ of Theorem 1.

Consider a state space split into two blocks B_0, B_1 and a lift U_f whose prototypes u_0, u_1 are supported on the blocks. If the dynamics jump from B_0 to B_1 with probability p each step (and otherwise remain in the current block), then starting from u_0 , the probability of exiting B_0 within τ steps is $1 - (1 - p)^\tau$. Hence $\|Q_f(u_0 P^\tau) - \delta_0\|_{\text{TV}}$ is on the order of $1 - (1 - p)^\tau$, and similarly for u_1 .

Now take a coarse lens that merges two metastable regions into one block, and a refinement that splits them. If p is very small (relative to τ), the refined prototypes have small retention error and are stable, so refinement increases the stable count. If p is not small (relative to τ), the refined prototypes rapidly mix across the split boundary, their retention errors exceed the stability threshold, and refinement decreases the stable count. \square

Remark 13 (Qualified “theory birth”). The induced endomap $E_{\tau,f}$ packages a distribution by evolving for τ steps and then forgetting micro-detail below the lens f , re-instantiating it via the chosen prototypes. Theorem 1 shows that a small retention error implies small idempotence defect. In this regime, the (approximately) stable prototypes u_x behave as emergent “objects” at scale τ in the theory f . Refinement changes the lens and can therefore reveal or destroy such objects (Theorem 2).

6 AUT + REV + ACC regime and graph 1-forms

This section isolates the representation-independent content behind the admissibility regime **A_AUT+A_REV+A_ACC** that we use throughout: the obstruction to “no-drive” behavior is a cohomological (cycle) obstruction on the bidirected support graph. Our presentation is discrete-time and purely graph-theoretic; any continuous-time (CTMC) reading is optional and out of scope for the main text.

6.1 Bidirected support and the log-ratio 1-form

Assume **A_FIN** and let P be a time-homogeneous Markov kernel on a finite state space Z . Write $P_{ij} := P(i, j)$. Under **A_REV** (microreversibility on support), we have

$$P_{ij} > 0 \implies P_{ji} > 0,$$

so every allowed transition comes with a reverse transition.

Define the (undirected) *support graph* $G = (Z, E)$ by declaring $\{i, j\} \in E$ iff $P_{ij} > 0$ (equivalently $P_{ji} > 0$ by **A_REV**). Let \vec{E} be the corresponding set of *oriented edges*: for each $\{i, j\} \in E$ we include both (i, j) and (j, i) .

Definition 11 (Edge log-ratio 1-form). *Assume **A_FIN+A_REV**. The antisymmetric edge 1-form associated to P is the function $a : \vec{E} \rightarrow \mathbb{R}$ defined by*

$$a_{ij} := a(i, j) := \log \frac{P_{ij}}{P_{ji}} \quad \text{for } (i, j) \in \vec{E}.$$

Then $a_{ji} = -a_{ij}$.

The cycle/affinity viewpoint for Markov networks and its cohomological structure is standard; see e.g. [Sch76, Pol12, AGH⁺12, DCLP23].

Remark 14 (Why **A_REV** matters). If **A_REV** fails, then ratios such as P_{ij}/P_{ji} may be undefined, and many time-reversal quantities become infinite rather than merely nonzero. In this paper we treat **A_REV** as the basic finiteness/admissibility condition for edge log-ratios.

6.2 Cycle integrals, exactness, and the null regime

A (simple) cycle in G is a sequence $(v_0, v_1, \dots, v_m = v_0)$. We require $\{v_k, v_{k+1}\} \in E$ for all k .

Definition 12 (Cycle integral / affinity along a cycle). *Given an antisymmetric edge 1-form a and an oriented cycle $\gamma = (v_0, \dots, v_m = v_0)$, define*

$$\mathcal{A}(\gamma) := \sum_{k=0}^{m-1} a_{v_k v_{k+1}}.$$

Reversing the orientation negates $\mathcal{A}(\gamma)$.

Definition 13 (Exact 1-forms). *An antisymmetric 1-form a on \vec{E} is exact if there exists a potential $\Phi : Z \rightarrow \mathbb{R}$ such that for every $(i, j) \in \vec{E}$,*

$$a_{ij} = \Phi(j) - \Phi(i).$$

We write $a = d\Phi$ in this case.

Lemma 4 (Exact forms have zero cycle integrals).

*Assume **A_FIN+A_REV**. If $a = d\Phi$ is exact, then $\mathcal{A}(\gamma) = 0$ for every cycle γ in G .*

Proof. Along a cycle $\gamma = (v_0, \dots, v_m = v_0)$ we have

$$\mathcal{A}(\gamma) = \sum_{k=0}^{m-1} (\Phi(v_{k+1}) - \Phi(v_k)) = \Phi(v_m) - \Phi(v_0) = 0.$$

□

Force Lives on Loops.

Theorem 3 (Cycle criterion for exactness). *Assume **A_FIN+A_REV** and let a be an antisymmetric 1-form on the oriented edge set \vec{E} of G . The following are equivalent:*

1. *a is exact: $a = d\Phi$ for some $\Phi : Z \rightarrow \mathbb{R}$.*
2. *$\mathcal{A}(\gamma) = 0$ for every cycle γ in G .*
3. *$\mathcal{A}(\gamma) = 0$ for every cycle γ in some fixed cycle basis of G (e.g. a fundamental cycle basis from a spanning tree).*

Proof sketch. (1) \Rightarrow (2) is Lemma 4. (2) \Rightarrow (3) is immediate.

For (3) \Rightarrow (1), fix a spanning tree T of each connected component of G . Choose a root r in each component and define $\Phi(r) = 0$; for any vertex v in that component, let $\Phi(v)$ be the sum of a along the unique tree path from r

to v (with signs determined by orientations). This makes $a_{ij} = \Phi(j) - \Phi(i)$ hold for every tree edge (i, j) . For any chord edge $e = (u, v) \notin T$, the fundamental cycle γ_e is the tree path from u to v together with e ; the hypothesis $\mathcal{A}(\gamma_e) = 0$ forces $a_{uv} = \Phi(v) - \Phi(u)$ on the chord as well. Thus $a = d\Phi$ on all of \vec{E} . \square

Corollary 2 (Null regime as “no-drive” / detailed-balance-like). *Assume **A_FIN**+**A_REV**+**A_NULL** and let a be the log-ratio 1-form on the bidirected support graph. Then a is exact. Equivalently, there exists $\Phi : Z \rightarrow \mathbb{R}$ such that*

$$\log \frac{P_{ij}}{P_{ji}} = \Phi(j) - \Phi(i) \quad \text{for all } (i, j) \in \vec{E}.$$

In particular, defining $\pi(i) \propto e^{\Phi(i)}$ yields the detailed balance relations $\pi(i)P_{ij} = \pi(j)P_{ji}$ on support.

6.3 Accounting as coordinates on cycle space

Drive Is Coordinate-Free. Let $\beta_1(G)$ be the cycle rank of the undirected support graph. Choosing a cycle basis $(\gamma_1, \dots, \gamma_{\beta_1})$ defines a coordinate vector of *affinities*

$$A_r := \mathcal{A}(\gamma_r), \quad r = 1, \dots, \beta_1(G).$$

Changing the cycle basis changes coordinates but not the underlying obstruction: $A = 0$ in one basis iff $A = 0$ in every basis (Theorem 3).

We regard **A_ACC** as the decision to work with this decomposition viewpoint: the antisymmetric log-ratio 1-form a splits into an exact (gauge) part plus a cycle/affinity part, and the affinity part vanishes exactly in the null regime **A_NULL**.

Remark 15 (Implementation note for Appendix C). The repository includes a small graph-theoretic routine that computes cycle affinities and tests exactness on finite Markov kernels (reversible vs biased 3-cycle examples). We relegate all such numerical evidence to Appendix C per the paper contract.

7 Arrow-of-time as path reversal asymmetry; data processing; protocol trap

Throughout this section we assume **A_FIN**. Let Z be a finite state space and let $P : Z \times Z \rightarrow [0, 1]$ be a time-homogeneous Markov kernel.

Definition 14 (Forward path law and reversal (D-AOT-01)). *Fix a horizon $T \in \mathbb{N}$ with $T \geq 1$ and an initial distribution $\rho \in \Delta(Z)$. The (forward) length- T path law $\mathbb{P}_{\rho,T}$ is the probability measure on Z^{T+1} defined by*

$$\mathbb{P}_{\rho,T}(z_0, \dots, z_T) := \rho(z_0) \prod_{t=0}^{T-1} P(z_t, z_{t+1}).$$

Let $\mathcal{R} : Z^{T+1} \rightarrow Z^{T+1}$ be the reversal map $\mathcal{R}(z_0, \dots, z_T) = (z_T, \dots, z_0)$, and write $\mathcal{R}_*\mu$ for pushforward of a measure μ . Define the (finite-horizon) path reversal asymmetry

$$\Sigma_T(\rho) := D_{\text{KL}}(\mathbb{P}_{\rho,T} \parallel \mathcal{R}_*\mathbb{P}_{\rho,T}) \in [0, \infty].$$

We use the standard convention $D_{\text{KL}}(p \parallel q) = +\infty$ if p is not absolutely continuous with respect to q .

Classical path-space fluctuation formulations of time-reversal asymmetry appear in [LS99].

Remark 16 (Stationarity is not required (L-AOT-STAT-01)). The quantity $\Sigma_T(\rho)$ is defined for any initial distribution ρ ; it can be strictly positive even in a detailed-balance chain when ρ is not stationary (a boundary-term effect). Stationarity is only needed when we want a *steady-state* interpretation (entropy production rate), not for the definition of Σ_T or for the data-processing inequality below.

Definition 15 (Steady-state entropy production (D-AOT-EP-01)). *Assume additionally **A_STAT** and let π be a stationary distribution for P . The (finite-horizon) steady-state entropy production is $\text{EP}_T := \Sigma_T(\pi)$. When a per-step rate is desired, we write $\text{ep} := \frac{1}{T}\Sigma_T(\pi)$ (and, under standard ergodicity conditions, one can pass to the limit $T \rightarrow \infty$).*

7.1 Data processing: coarse-graining cannot create asymmetry

No Fake Arrows. Now assume **A_LENS** and fix a coarse-graining (lens) $f : Z \rightarrow X$. Let $f^{T+1} : Z^{T+1} \rightarrow X^{T+1}$ be the coordinatewise map

$$f^{T+1}(z_0, \dots, z_T) = (f(z_0), \dots, f(z_T)).$$

Define the observed path law $\mathbb{Q}_{\rho,T} := (f^{T+1})_*\mathbb{P}_{\rho,T}$. Note that the observed process on X need not be Markov; the DPI statement below is purely about path measures.

Theorem 4 (Data processing for path reversal asymmetry (T-AOT-01)).
Under **A_FIN+A_LENS**, for every $T \geq 1$ and every initial distribution ρ ,

$$D_{\text{KL}}(\mathbb{Q}_{\rho,T} \parallel \mathcal{R}_* \mathbb{Q}_{\rho,T}) \leq D_{\text{KL}}(\mathbb{P}_{\rho,T} \parallel \mathcal{R}_* \mathbb{P}_{\rho,T}) = \Sigma_T(\rho).$$

Proof sketch. The reversal map \mathcal{R} commutes with coordinatewise coarse-graining:

$$\mathcal{R} \circ f^{T+1} = f^{T+1} \circ \mathcal{R}.$$

Hence $\mathcal{R}_* \mathbb{Q}_{\rho,T} = (f^{T+1})_*(\mathcal{R}_* \mathbb{P}_{\rho,T})$. The claim is then the usual contraction of KL under measurable maps: for any measurable g , $D_{\text{KL}}(g_* p \parallel g_* q) \leq D_{\text{KL}}(p \parallel q)$. \square

Importantly, this monotonicity rules out *false positives*: coarse observation can hide irreversibility but cannot create it.

Corollary 3 (No false positives). *Assume **A_FIN+A_LENS**. If the lifted system has zero asymmetry at horizon T , then every coarse observation has zero asymmetry at the same horizon. There are no false positives from forgetting variables. Coarse-graining can hide irreversibility (false negatives), but it cannot create it.*

Related coarse-graining and hidden-drive discussions in stochastic thermodynamics include [PPRV10, DvdMS23, BSvdM25, Mer11].

Equivalently: passing to a coarser theory (via a lens pushforward) cannot create time-reversal asymmetry.

7.2 Protocol trap: apparent stroboscopic arrows and the “clock audit”

P3 Loves P6 Law. We now pin down a class of autonomous protocol models and the sense in which “P3 requires P6_drive”. Let X be a microstate space and let $\Phi = \{0, 1, \dots, m-1\}$ be a finite phase space. A *protocol family* is a collection of kernels $\{K_\phi\}_{\phi \in \Phi}$ on X together with a phase kernel S on Φ .

Definition 16 (Autonomous m -phase protocol). **D-PROT-02.**

Fix $\alpha \in [0, 1]$ and define an autonomous Markov chain on the product space $Z := X \times \Phi$ by

$$P((x, \phi), (x', \phi')) := \alpha \mathbf{1}_{x'=x} S(\phi, \phi') + (1 - \alpha) \mathbf{1}_{\phi'=\phi} K_\phi(x, x').$$

We call this the lifted (phase-included) protocol model.

Theorem 5 (Protocol trap / hidden-drive principle (T-AOT-02)). *Assume **A_FIN+A_AUT**. Suppose: (i) the phase chain S is reversible with stationary distribution s ; (ii) each K_ϕ is reversible with respect to a common stationary distribution π on X . Then the product measure $\mu := \pi \times s$ is stationary for the lifted chain on $Z = X \times \Phi$ and the lifted chain is reversible (hence $\Sigma_T(\mu) = 0$ for all T). Consequently, by Theorem 4 applied to the projection $X \times \Phi \rightarrow X$, the observed process on X has zero path reversal asymmetry at every finite horizon.*

*In particular, any apparent arrow-of-time obtained by analysing a stroboscopic composition of kernels on X must be traced to either: (a) an external schedule (violating **A_AUT**), or (b) a driven/bias phase dynamics (nonzero affinity, cf. Section 6).*

Proof sketch. Let P_{phase} and P_{state} be the phase-update-only and state-update-only kernels defined by

$$\begin{aligned} P_{\text{phase}}((x, \phi), (x', \phi')) &= \mathbf{1}_{x'=x} S(\phi, \phi'), \\ P_{\text{state}}((x, \phi), (x', \phi')) &= \mathbf{1}_{\phi'=\phi} K_\phi(x, x'). \end{aligned}$$

By reversibility of S with respect to s and of each K_ϕ with respect to the common π , both P_{phase} and P_{state} are reversible with respect to $\mu = \pi \times s$. The lifted kernel is their convex combination $P = \alpha P_{\text{phase}} + (1 - \alpha) P_{\text{state}}$, hence is also reversible with respect to μ . Therefore $\Sigma_T(\mu) = 0$ for all T , and the DPI argument yields zero asymmetry on X . \square

Corollary 4 (“P3 needs P6_drive” under autonomy (C-ACC-01)). *Under **A_FIN+A_REV+A_AUT+A_ACC+A_NULL** (no affinities anywhere in the lifted state space), protocol holonomy alone does not yield sustained arrow-of-time in the autonomous accounted model. For any stationary distribution of the lifted autonomous chain, nonzero steady-state entropy production requires a nontrivial affinity component (P6_drive) in the lifted dynamics (e.g. a biased phase cycle) or an external schedule.*

Remark 17 (What the “protocol trap” actually warns about). The theorem above is about *true* path reversal asymmetry. In practice, one can still obtain a spurious “arrow” by fitting a one-step Markov model to a non-Markov observed process (for example, by replacing the hidden-phase model by a stroboscopic kernel on X). The cure is the clock audit: include the phase variable in the state, or work directly with path-space quantities that do not assume Markovianity at the observed level.

8 Generic extension and the finite forcing lemma

8.1 Theories as partitions and definability

Assume **A_FIN** and fix a finite microstate set Z . A *theory* (or *lens*) is a finite-valued map $f : Z \rightarrow X$ (assumption **A_LENS**), inducing the partition

$$\Pi_f := \{B_x\}_{x \in X}, \quad B_x := f^{-1}(x).$$

We interpret Π_f as the collection of distinctions expressible in the current theory.

Definition 17 (Predicates and definability). *A (Boolean) predicate is a function $h : Z \rightarrow \{0, 1\}$. We say that h is definable from the theory f (equivalently: Π_f -measurable) if it is constant on each block:*

$$\forall x \in X, \forall z, z' \in B_x, \quad h(z) = h(z').$$

Equivalently, there exists $\tilde{h} : X \rightarrow \{0, 1\}$ with $h = \tilde{h} \circ f$.

Finite-model-theoretic context for definability in finite structures and counting viewpoints appears in, e.g., [MS11, Hyt15].

Remark 18. In a fixed finite Z there is a hard cap on how many strict refinements are possible; open-endedness is modeled by repeated extensions across growing problems/domains (each step finite). The “finite forcing lemma” is a combinatorial proxy and does not claim set-theoretic independence.

Adjoining a new predicate h refines the theory by considering the joint lens

$$(f, h) : Z \rightarrow X \times \{0, 1\}.$$

This extension is *trivial* precisely when h was already definable from f .

8.2 Counting lemma: definable predicates are rare

Let $N := |Z|$ and $K := |X|$.

Lemma 5 (Counting definable predicates). *Assume **A_FIN+A_LENS**. The number of predicates definable from $f : Z \rightarrow X$ is exactly 2^K . Consequently, if h is sampled uniformly from $\{0, 1\}^Z$, then*

$$\mathbb{P}(h \text{ is definable from } f) = \frac{2^K}{2^N} = 2^{-(N-K)}.$$

Proof. A definable predicate is determined by choosing one bit per block B_x (the common value on that block), hence 2^K choices. There are 2^N total predicates on Z , so the stated probability follows. \square

8.3 Finite forcing: generic extensions are non-definable

The preceding lemma is the finite analogue of the forcing slogan: *generic extensions add non-definable predicates*. To emphasize the “genericity” aspect, we also record a simple sufficient condition for a predicate to be *strongly novel*.

Almost Nothing Is Definable.

Theorem 6 (Finite forcing lemma / physical forcing lemma). *Assume **A_FIN+A_LENS**. Let $f : Z \rightarrow X$ be a theory with blocks $\Pi_f = \{B_x\}_{x \in X}$, and sample $h : Z \rightarrow \{0, 1\}$ uniformly from $\{0, 1\}^Z$ (i.i.d. fair bits across microstates). Then **non-definability is generic**:*

$$\mathbb{P}(h \text{ is not definable from } f) = 1 - 2^{-(N-K)}.$$

Proof. This is Lemma 5. □

The “Nothing Stays Constant” Lemma.

Lemma 6. *Assume **A_FIN+A_LENS** and let $f : Z \rightarrow X$ have blocks $\Pi_f = \{B_x\}_{x \in X}$. For each block B_x ,*

$$\mathbb{P}(h \text{ is constant on } B_x) = 2^{1-|B_x|},$$

and by a union bound,

$$\mathbb{P}(\exists x \in X \text{ such that } h \text{ is constant on } B_x) \leq \sum_{x \in X} 2^{1-|B_x|}.$$

In particular, if every block has size at least $m \geq 2$, then

$$\mathbb{P}(h \text{ splits every block}) \geq 1 - K 2^{1-m}.$$

Proof. On a block of size $|B_x|$, the only constant labelings are “all zeros” and “all ones”, each with probability $2^{-|B_x|}$, hence $2^{1-|B_x|}$ total. The union bound gives the stated inequality, and the final display uses $|B_x| \geq m$ for all x . □

Corollary 5 (Generic extension is strict).

*Assume **A_FIN+A_LENS**. With probability $1 - 2^{-(N-K)}$, the refined lens $(f, h) : Z \rightarrow X \times \{0, 1\}$ is a strict refinement of f . Equivalently, it distinguishes at least one pair of microstates previously indistinguishable under f .*

Proof. If the refinement were not strict, then h would be constant on each block of f , i.e. definable from f . The complement event has probability $1 - 2^{-(N-K)}$ by Theorem 6(1). \square

Remark 19 (Why this is the anti-saturation move). Closure (or packaging) at a fixed theory saturates by idempotence: repeated closure steps do not create an infinite strict chain. Corollary 5 exhibits a mathematically clean *extension step*: adjoining a generic predicate changes the theory itself in a way that is overwhelmingly not reconstructible from the old distinctions. This is the finite “forcing-style” mechanism we use to justify strict ladder climbing via theory/completion change.

9 Why the primitives are unavoidable

Definition 18 (Process soup (D-META-PROC-01)). *A process soup is a set \mathcal{P} equipped with a partially defined associative composition \circ (a partial semigroup), with identities where relevant.*

Definition 19 (Interface lens (D-META-LENS-01)). *An interface lens is a map $f : \mathcal{P} \rightarrow X$ to a set of observables. It induces a kernel equivalence $x \sim_f y$ iff $f(x) = f(y)$. Limited access means $\exists x \neq y$ with $x \sim_f y$. This abstracts the notion of a coarse-graining lens $f : Z \rightarrow X$ used elsewhere: here the lens acts on processes/happenings rather than on state spaces.*

Definition 20 (Refinement family (D-META-REF-01)). *A refinement family is a chain of equivalence relations $(\sim_j)_{j \geq 0}$ with $x \sim_{j+1} y \Rightarrow x \sim_j y$; equivalently, lenses f_j with $f_j = g_j \circ f_{j+1}$ for some maps g_j . As j increases, observables become more discriminating (partition refinement), hence the induced saturation operators on predicates become weaker (descending in the pointwise order).*

Definition 21 (Bounded interface (D-META-BND-01)). *There exists $C_0 \geq 1$ such that the quotient index satisfies $|\mathcal{P}/\sim_j| \leq C_0(j+1)$ for all j . This is an explicit bounded-interface hypothesis on the lens hierarchy (a modeling assumption encoding bounded observable signature); it is not derived from refinement alone.*

For example, refinement alone permits exponential growth (e.g. $|X_j| = 2^j$ with $X_j = \{0, 1\}^j$), so the linear bound must be assumed or verified in an instantiation.

Theorem 7 (Self-generated primitives). *Assume D-META-PROC-01, D-META-LENS-01, D-META-REF-01, and D-META-BND-01. Then the six primitives P1–P6 appear canonically as closure mechanics. Here P5, P6, P4, and P2 are by construction from the lens/refinement data, while P1 and P3 isolate the only nontrivial compatibility requirements as named hypothesis slots (HL-META-1/2/3).*

1. **P5 (Packaging)**. *Each equivalence \sim_j yields a quotient/packaging map $\Pi_j : \mathcal{P} \rightarrow \mathcal{P}/\sim_j$ and induces an idempotent saturation on predicates/events via $\text{cl}_j(A) := \Pi_j^{-1}(\Pi_j(A))$ (cf. Sections 4 and 5.1).*
2. **P6 (Accounting)**. *The refinement order on quotients induces a pre-order and monotone “audit” quantities (e.g. $j \mapsto |X_j|$); this is an information/feasibility order only (not, by itself, a thermodynamic arrow-of-time).*
3. **P4 (Staging)**. *The refinement chain generates a depth index j , and bounded interfaces ensure $|X_j| \lesssim (j + 1)$. Nontrivial staging corresponds to strict refinement at some scale (HL-META-2: $\sim_{j+1} \subsetneq \sim_j$ for some j).*
4. **P2 (Constraints)**. *Feasible macrostates are those representable as images under Π_j , yielding a constraint set carved by interface compatibility.*
5. **P1 (Operator rewrite)**. *Fix a micro-update $F : \mathcal{P} \rightarrow \mathcal{P}$. The induced macro-update $F^\sharp([p]) := [F(p)]$ on \mathcal{P}/\sim_j is well-defined iff $p \sim_j q \Rightarrow F(p) \sim_j F(q)$ (HL-META-1). When (HL-META-1) fails, one must either (i) refine the lens so stability holds, (ii) modify/extend F to respect \sim_j , or (iii) accept that no closed macro dynamics exists at depth j .*
6. **P3 (Holonomy)**. *Assuming there exist two admissible closure/evolution routes with the same input/output type (HL-META-3) (e.g. $\Pi_j \circ F$ versus $F^\sharp \circ \Pi_j$ when F^\sharp is well-defined), their discrepancy defines a holonomy-like diagnostic. In normed operator instantiations we measure this by $\text{RM}(j)$ (Appendix E, Eq. (D-TK-ROUTE-01)); this is a discrepancy diagnostic, not itself a directionality certificate (see Section 7).*

Proof sketch. P5 is the quotient/packaging map induced by the equivalence, which is idempotent by construction. P6 follows from the refinement pre-

order and yields information/feasibility monotones only (not a thermodynamic arrow-of-time). P4 is the stage index implicit in the refinement chain (nontrivial if refinements are strict, HL-META-2), which also connects to closure ladders (Section 4). P2 is the feasible image under Π_j , i.e. representability at the interface. For P1, given a micro-update F , the expression $F^\sharp([p]) = [F(p)]$ descends to a map on \mathcal{P}/\sim_j exactly when (HL-META-1) holds; when it fails, closed macro dynamics at depth j requires rewrite/extension, refinement, or must be abandoned. For P3, whenever two admissible routes exist (HL-META-3), their noncommutativity produces a route-mismatch/holonomy defect, formalized by $\text{RM}(j)$ in Appendix E; this defect is agnostic to directionality without additional structure (Section 7). \square

10 Primitives P1–P6 as closure-changing operations

The primitives are presented below as a minimal operator vocabulary, but they are not axioms: Theorem 7 (Section 9) shows they are structurally forced once composability, limited access, and bounded interfaces are assumed. This framing remains compatible with the corrected stance on P3: protocol/phase is part of the autonomous state (no external schedule), and route dependence alone does not certify arrow-of-time. Primitives P1–P6 form a vocabulary of operations that change (i) the admissible transition structure (support graph and cycle space), (ii) the available idempotent endomaps (packaging and fixed points), and (iii) the affinity data encoded by the graph 1-form. These are not ratchets by themselves; directionality is certified by path reversal asymmetry and/or non-exactness of the log-ratio 1-form under the **A_AUT**+**A_REV**+**A_ACC** regime.

10.1 Definitions of P1–P6

Definition 22 (P1: operator rewrite). *P1 is a rewrite of the substrate operator: replace a Markov kernel P by a new kernel P' (or a finite family $\{P^{(m)}\}$), thereby changing the endomap $\mu \mapsto \mu P^\tau$ and the induced empirical endomap $E_{\tau,f}$. This is a change in the operator itself, not an external schedule.*

Definition 23 (P2: gating / constraints). *P2 restricts the support graph by deleting edges (setting selected $P_{ij} = 0$) and renormalizing rows, or equiva-*

lently restricting to a subgraph on which the kernel lives. Cycle-rank claims are interpreted on the undirected support graph justified by **A_REV**.

Definition 24 (P3: autonomous protocol holonomy). *P3 is modeled in the autonomous lifted form on $Z := X \times \Phi$. The phase $\phi \in \Phi$ evolves by an internal kernel \mathcal{S} , and conditioned on ϕ the microstate updates by K_ϕ . No external schedule is assumed in the main text; externally scheduled stroboscopic protocols are non-autonomous and fall outside **A_AUT**. P3 (route mismatch/holonomy) is a protocol-geometry diagnostic: by itself it does not certify directionality, and any arrow-of-time claim must be supported by an audit functional (e.g. affinity or path reversal asymmetry) in a model where the protocol state is included.*

Definition 25 (P4: sectors / invariants). *P4 is a conserved sector label: the support decomposes into disconnected components (block structure), or equivalently P is block diagonal up to permutation, so evolution preserves a sector coordinate.*

Definition 26 (P5: packaging). *P5 is an idempotent endomap e whose fixed points $\text{Fix}(e)$ are the packaged objects of a given theory. This is an endomap notion (Section 5.1). It is not an order-closure.*

Definition 27 (P6: accounting / audit). *P6 is an accounting/audit structure: a certificate or functional that is monotone under coarse maps or packaging. Canonical instantiations in this manuscript include:*

1. *the information/feasibility order induced by limited access (META),*
2. *path-space KL asymmetry with data processing (Section 7),*
3. *the ACC graph 1-form/cycle-integral audit (Section 6).*

Remark 20. We write *P6_drive* for the ACC specialization: a non-exact log-ratio 1-form (equivalently, a nonzero cycle integral). This is the condition used in the slogan “P3 needs P6_drive” under autonomy.

10.2 How the primitives compose to generate theory growth

In this manuscript a *theory* is the package (lens/definability, completion, audit) formalized in Section 3.4. Emergence means stable fixed points of the completion rule inside a theory, while open-endedness means strict theory extension rather than iteration of a fixed closure (Section 8).

1. **Limited access \Rightarrow P5 packaging.** A lens collapses distinctions, so a completion endomap packages microstates into fixed-point objects (Sections 4 and 5.1).
2. **Lossy packaging \Rightarrow P6 accounting.** Audits/monotones track what survives under lenses (Sections 6 and 7); coarse-graining cannot create directionality by DPI.
3. **Saturation \Rightarrow extension.** Iterating a fixed completion saturates (Section 4), so strict growth requires theory extension (Section 8).
4. **P4 staging.** A refinement family supplies a theory index (Section 9); bounded interfaces/quantized indices make repeated refinement coherent.
5. **P2 gating.** Feasibility curves representable macrostates and restricts support/cycle space (see the P2 theorem below and Section 6).
6. **P1 rewrite.** When induced macro-dynamics fails to descend, an operator rewrite/refinement is forced to restore closure on packaged objects (Section 5.1).
7. **P3 route mismatch.** Noncommuting reductions yield route mismatch $RM(j)$ (Appendix E, Eq. (D-TK-ROUTE-01)); this is route dependence, not a directionality certificate (Section 7).
8. **Iteration.** The updated (lens, completion, audit) defines the next theory; objects and audits change and the loop repeats.

Mapping to the spine.

- P5 \leftrightarrow completion/idempotents (Sections 4, 5.1).
- P6 \leftrightarrow audits (Sections 6, 7).
- P1/P2/P4 \leftrightarrow closure-changing structure (this section and the P1/P2 theorems below).
- P3 \leftrightarrow route mismatch (Appendix E).
- Open-endedness \leftrightarrow strict theory extension (Section 8).

We do not claim every system self-organizes; the toolkit identifies the minimal operations and certifies a system must instantiate if it exhibits autonomous theory growth and stable emergent objects.

10.3 Downward influence across theories

We make explicit how the emergence calculus represents *downward influence* (macro-to-micro constraint) within a coupled theory package. In this paper, “downward influence” is not a seventh primitive and not a metaphysical claim: it is a *structural* phenomenon that appears whenever a coarse description is fed back into a finer one through the existing primitives (Section 3.4).

Three canonical downward mechanisms already in P1–P6. At the abstraction level of this paper, every downward influence path factors through one of the following constructions (the rest are compositions/specializations):

1. **Representative selection via lift/completion (P5).** A lens $f : Z \rightarrow X$ induces a pushforward Q_f and, once prototypes are fixed, a canonical lift $U_f : \Delta(X) \rightarrow \Delta(Z)$ (Definition 8). The map $x \mapsto u_x := U_f(\delta_x)$ is an explicit macro-to-micro channel: a coarse description selects a representative micro-description. This is exactly the “completion” move used to define dynamics-induced packaging maps such as $E_{\tau,f}$.
2. **Feasibility gating and budgets (P2/P6).** A coarse constraint $C \subseteq X$ restricts the finer theory to the preimage $f^{-1}(C) \subseteq Z$ (and similarly for path constraints). Equivalently: feasibility/accounting at the coarse theory can *exclude* micro directions by shrinking the admissible set of states/histories, without modifying the underlying micro update rule.
3. **Internal protocol/phase as state (P3 with the autonomy guardrail).** A coarse context variable (“phase”) can be included as part of the state and can index a family of micro updates. This yields genuine context-dependent micro dynamics *without* smuggling in an external schedule; see the protocol-trap discussion in Section 7. Route mismatch/holonomy (P3) then measures noncommutativity of routes, but remains *agnostic* about directionality unless coupled to a driven/asymmetric accounting structure (P6).

Non-claim. This subsection does not assert that coarse descriptions “override” micro laws. It only records that once a coarse theory is fed back into the finer one via (i) completion/lift, (ii) feasibility gating, or (iii) an internalized protocol variable, the resulting coupled theory package has bona

vide macro-to-micro influence *in the sense of constrained admissibility and context-indexed operators*.

10.4 Two load-bearing propositions

Constraints Kill Engines.

Theorem 8 (P2 gating shrinks cycle space). *T-P2-01*.

*Assume **A_FIN+A_REV**.*

Let $G = (V, E)$ be an undirected support graph and let $G' = (V, E')$ be obtained by deleting edges. Then

$$\begin{aligned}\beta_1(G') &\leq \beta_1(G), \\ \beta_1(G) &:= |E| - |V| + c(G),\end{aligned}$$

where $c(G)$ is the number of connected components.

Proof. Deleting a single edge reduces $|E|$ by one and increases $c(G)$ by at most one. Hence $|E| - |V| + c(G)$ cannot increase. Iterating over all deleted edges yields the claim. \square

Theorem 9 (P1 can change cycle rank and spectral gap (T-P1-01)). *Assume **A_FIN+A_REV**. P1 rewrites can increase or decrease cycle rank and can increase or decrease a metastability proxy (the spectral gap of a reversible lazy random walk).*

Proof sketch. (Cycle rank.) Adding an edge within an existing connected component leaves $c(G)$ unchanged and increases $|E|$ by one, so $\beta_1(G)$ increases by one. More general rewires can change $c(G)$ and thus increase or decrease β_1 .

(Spectral gap.) By rewriting edge weights or topology, one can create or remove bottlenecks. Highly connected graphs have large gaps, while nearly decomposable graphs have small gaps; thus P1 can increase or decrease the gap by explicit constructions (see Appendix C for examples). \square

Remark 21 (Evidence pointers). Appendix C records finite graph and kernel witnesses for the P1 and P2 behaviors described above, including explicit cycle-rank changes and spectral-gap increases/decreases.

11 Examples

This section gives a few finite examples that instantiate the preceding definitions and results. They introduce no new machinery and are included purely for intuition and for checking that the assumption tags match the intended use-cases.

11.1 A biased three-cycle: a non-exact graph 1-form

Let $Z = \{0, 1, 2\}$ and consider the Markov kernel P defined by the parameters (p, q, s) with $p, q, s > 0$ and $p + q + s = 1$, where from each $i \in Z$,

$$P(i, i) = s, \quad P(i, i + 1 \bmod 3) = p, \quad P(i, i - 1 \bmod 3) = q.$$

This satisfies **A_FIN**+**A_REV**. The antisymmetric edge 1-form $a_{ij} := \log(P_{ij}/P_{ji})$ is therefore defined on each bidirected edge.

Along the directed cycle $\gamma = (0 \rightarrow 1 \rightarrow 2 \rightarrow 0)$, the cycle integral (affinity) is

$$\oint_{\gamma} a = \log \frac{P_{01}}{P_{10}} + \log \frac{P_{12}}{P_{21}} + \log \frac{P_{20}}{P_{02}} = 3 \log \frac{p}{q}.$$

For instance with $(p, q, s) = (0.7, 0.2, 0.1)$ we get $|\oint_{\gamma} a| \approx 3.7583 \neq 0$, hence the 1-form is not exact and no global potential Φ can satisfy $a_{ij} = \Phi_j - \Phi_i$ on all edges (cf. Section 6).

11.2 Protocol trap: external schedule vs autonomous lifted model

We illustrate the difference between an externally scheduled protocol and an autonomous lifted model. Let $Z = \{0, 1, 2\}$ and let K_0, K_1 be two kernels that are each reversible w.r.t. the uniform distribution (on Z), but do not commute (so the composition depends on order). One concrete instance used in our reproducibility artifacts is:

$$K_0 = \begin{pmatrix} 0.573333 & 0.393333 & 0.033333 \\ 0.393333 & 0.573333 & 0.033333 \\ 0.033333 & 0.033333 & 0.933333 \end{pmatrix},$$

$$K_1 = \begin{pmatrix} 0.933333 & 0.033333 & 0.033333 \\ 0.033333 & 0.573333 & 0.393333 \\ 0.033333 & 0.393333 & 0.573333 \end{pmatrix}.$$

for which $\max_{i,j} |(K_0K_1 - K_1K_0)_{ij}| \approx 1.296 \times 10^{-1}$.

Externally scheduled (not autonomous): if one applies K_0 then K_1 in a fixed periodic order but hides the schedule variable, the stroboscopic one-step kernel on Z is $K := K_1K_0$, which can exhibit positive steady-state entropy production (positive time-reversal asymmetry in the one-step Markov description).

Autonomous lifted model (clock included in state): define an augmented state space $Z \times \Phi$ where $\Phi = \{0, 1\}$ is a phase variable, and define a time-homogeneous random-scan Markov chain that, at each step, (i) updates ϕ using S while keeping z fixed with probability α , and (ii) updates z using K_ϕ while keeping ϕ fixed with probability $1 - \alpha$. If S is unbiased/reversible (no phase affinity) and both K_0, K_1 are reversible w.r.t. a common π , then the lifted model exhibits vanishing steady-state entropy production (and hence vanishing path reversal asymmetry at stationarity), whereas introducing a phase bias (nontrivial affinity in the Φ -subsystem) yields strictly positive entropy production. This is the clean mathematical form of the slogan “P3 needs P6_drive” in autonomous models: noncommuting actions alone do not create sustained directionality unless the protocol degree of freedom itself is driven.

11.3 Finite forcing count: definability is exponentially rare

Let Z be a finite set with $|Z| = N$, and let a “theory” be represented by a partition Π of Z into K blocks (equivalently a lens $f : Z \rightarrow X$ with $|X| = K$). A predicate $h : Z \rightarrow \{0, 1\}$ is definable from Π iff it is constant on each block of Π , so the number of definable predicates is exactly 2^K , while the total number of predicates is 2^N . Therefore, if h is sampled uniformly from $\{0, 1\}^Z$,

$$\mathbb{P}(h \text{ is definable from } \Pi) = \frac{2^K}{2^N} = 2^{-(N-K)}.$$

For example, when $(N, K) = (16, 4)$, this gives $2^{-(16-4)} = 2^{-12} \approx 2.44 \times 10^{-4}$, i.e. a uniformly random new predicate is overwhelmingly likely to be non-definable from the old theory. This is the finite counting core behind the “generic extension” step in Section 8.

12 Discussion and scope boundary

12.1 What the theory does and does not claim

The organizing loop in Section 1.1 separates three roles that are often conflated: packaging (idempotence and fixed points) and extension (non-definability), versus directionality (path-space KL and affinities). The paper’s contribution is to make these roles explicit and composable, with certificates that are stable under coarse observation (DPI) and robust to hidden-clock artifacts (protocol audit).

In this paper, “theory” is used in the minimal mathematical sense of a definability structure (lens/partition) together with its completion/packaging map. Later interpretations are intentionally out of scope here. Viewed through Theorem 7, the primitives are not hypotheses about the world but structural consequences of description under bounded interfaces. Instantiating the framework amounts to specifying a lens/refinement family, a completion rule (packaging), and an audit functional. One then verifies that the domain’s interfaces realize the standing hypotheses. None of these domain choices are fixed by the abstract theorem.

Equally important are the non-claims. We do not claim that closure ladders arise automatically in arbitrary dynamics, nor that any particular choice of lens f or timescale τ is preferred. We also do not claim that protocol holonomy by itself yields sustained directionality under autonomy; on the contrary, the lifted-model statements formalize that sustained asymmetry requires a genuine affinity (a non-exact 1-form component) or an externally imposed schedule (which lies outside **A_AUT**). Finally, the finite forcing lemma is used only as a *finite proxy* for generic extension: it justifies that strict theory extension is cheap and typical when there is hidden volume, not that all forms of novelty reduce to random predicates.

This paper keeps its scope finite, structural, and diagnostic: it does not advance claims about continuous-time stochastic thermodynamics, empirical estimation guarantees, or domain-specific applications. Appendix D frames the Zeno decision frontier as a research outlook for when Zeno can or cannot be ruled out in the abstract bridge setting.

12.2 Outlook: forthcoming instantiations

This paper isolates a math-only emergence calculus. It uses a *theory package* (lens/definability, completion, audit). It is paired with the six primitive moves that change what theories can exist and what objects they recognize. A series of companion papers will instantiate this calculus in multiple

domains (with domain-specific choices of lenses, bridges, and audits) while keeping the present framework fixed:

- mathematics theories,
- physics theories as families of effective closures across scales,
- life and cognition as theory stacks,
- conscious entities as theories,
- societies and civilizations,
- the universe as a “theorist”.

We intentionally do not include details here; the purpose of this paragraph is only to indicate scope and to delineate the boundary between the general calculus proved in this paper and its forthcoming instantiations.

A Appendix A: Reproducibility checklist

This repository is intended to be mechanically checkable at three levels: (i) a claim/dependency registry, (ii) a small Lean formalization core, and (iii) a deterministic Python evidence harness.

Repository: <https://github.com/ioannist/six-birds-theory>.

Arweave ID: 1kWCQsvq071tmOBx-uOVQfnsfsJZeHXhK7AWx-P-pII

A.1 Repository integrity checks (from repo root)

Run the following commands from the repository root:

```
python3 scripts/check_deps_dag.py
python3 scripts/check_kb_pointers.py
python3 scripts/check_paper_contract.py
./check_lean.sh
```

A.2 Python evidence harness (deterministic tests)

The Python code lives under `python/` and is run in an isolated virtual environment. From the repository root:

```
python3 -m venv python/.venv
python/.venv/bin/pip install -q pytest numpy
python/.venv/bin/pip install -e python
cd python && .venv/bin/python -m pytest -q
```

A.3 Notes on tool availability

A \LaTeX toolchain is not required to run the checks above. If you wish to compile the PDF locally, install one of `latexmk`, `pdflatex`, or `tectonic` and run:

```
./scripts/build_paper.sh
```

B Appendix B: Lean formalization map

The Lean development provides a minimal mechanized backbone for the order-theoretic and “thin packaging” parts of the theory. It is intentionally small: no probability theory or Markov chain results are mechanized in Lean for this paper.

B.1 What is mechanized

The following items are implemented in Lean under `formal/`:

- **Order-closure and closure ladders** (order-theoretic closure operators, fixed points, strict ladder relation, and iterate stabilization).
- **Thin packaging / reflection** for closure operators, expressed via the Galois insertion associated to a closure operator.
- **Idempotent endomaps** and fixed-point subtypes.
Bridge lemma from order-closures to idempotent endomaps.

B.2 File map and key declarations

- `formal/ClosureLadder/Basic.lean`:
closure operators, closed subtype, strict ladder relation, and iterate stabilization.
Key lemmas used in the manuscript include `ladder*mono`, `closure*iterate*succ`, and `closure*iterate*ge*one`.
- `formal/ClosureLadder/Packaging.lean`:
reflection-style lemmas for the closed subtype (thin-category packaging anchor).
- `formal/ClosureLadder/IdempotentEndo.lean`:

the structure `IdempotentEndo` and its fixed-point subtype.
Bridge lemma: `ClosureOperator.toIdempotentEndo`.

- `formal/ClosureLadder.lean`:
umbrella import module.

B.3 How to build

From the repository root:

```
./check_lean.sh
```

This script performs a cache fetch (when available) and then runs `lake build`.

C Appendix C: Python evidence map

The repository contains a small deterministic Python harness (finite-state Markov kernels, path laws, KL divergence, coarse-graining pushforwards, graph cycle utilities, and related diagnostics). These computations are *evidence and sanity checks* for the mathematical definitions and theorem statements; they are not used as premises in any proof.

C.1 How to run

See Appendix A.2 for a one-shot test run via `pytest`. All tests are deterministic (fixed seeds or no randomness).

C.2 Evidence by theme (tests and scripts)

The following tests and scripts witness the main diagnostic phenomena discussed in the paper:

Path reversal asymmetry and data processing (DPI). Tests include randomized DPI checks, strict hiding examples, and nonstationary initial distributions. Representative files:

```
python/tests/test_paths_kl_dpi.py  
python/tests/test_paths_dpi_randomized.py  
python/tests/test_paths_dpi_nonstationary_init.py  
python/tests/test_paths_strict_hiding.py  
python/scripts/dpi_sweep.py
```

Protocol trap and the “P3 needs P6_drive” correction under autonomy. Tests separate externally scheduled noncommuting kernels from autonomous phase-included models, including m-phase protocols and the “protocol clock audit”. Representative files:

```
python/tests/test_protocol_trap_external_vs_autonomous.py
python/tests/test_protocol_clock_audit_phi_included.py
python/tests/test_protocol_mphase_random_scan.py
python/tests/test_protocol_stroboscopic.py
python/scripts/protocol_trap_demo.py
```

Graph topology effects of P2 (edge deletion) and P1 (rewrites). Tests verify cycle-rank monotonicity under edge deletion, show cycle-rank increases under edge addition, and demonstrate spectral-gap changes (including a decisive gap-decrease rewrite). Representative files:

```
python/tests/test_graph_p2_cycle_rank_monotone.py
python/tests/test_graph_p1_edge_add_beta1_delta.py
python/tests/test_p1_spectral_gap_decrease.py
python/scripts/graph_p1_p2_demo.py
python/scripts/p1_gap_decrease_demo.py
```

Dynamics-induced empirical endomap and idempotence defect. Tests implement the induced empirical endomap, the total-variation idempotence defect, and show that refinement can help or hurt depending on leakage and timescale. Representative files:

- python/tests/
test_empirical_closure_
idempotence_defect_tv.py
- python/tests/
test_empirical_closure_
refinement_reveals_objects.py
- python/tests/
test_empirical_closure_
refinement_can_hurt.py
- python/scripts/empirical_closure_demo.py
- python/scripts/refinement_help_hurt_sweep.py

Finite forcing / definability rarity. Tests include exact enumeration for small cases, Monte Carlo checks, and a contrast with fixed-weight sampling. Representative files:

```
python/tests/test_forcing_exact_enumeration.py
python/tests/test_forcing_monte_carlo_matches_theory.py
python/tests/test_forcing_fixed_weight_differs.py
python/scripts/forcing_sweep.py
```

Finite-memory ICAP witnesses. Tests implement the discrete-time finite-memory mass bound for truncated inputs. Representative files:

```
python/src/clt/bridge_kernels.py
python/tests/test_bridge_kernels_fmем_icap.py
```

Balanced-atom route (definitions + kernel-mass hinge). These files provide the finite-memory ICAP and toy compression witnesses used by the balanced-atom route. Representative files:

```
python/src/clt/bridge_kernels.py
python/tests/test_bridge_kernels_fmем_icap.py
python/src/clt/balanced_atoms.py
python/tests/test_balanced_atoms_icap.py
python/scripts/balanced_atoms_witness.py
python/src/clt/ect_toy.py
python/tests/test_ect_compression_witness.py
python/scripts/ect_compression_witness.py
```

ECT compression and capacity witnesses. Tests and scripts witness linear mode count and linear ICAP capacity growth in a toy family. Representative files:

```
python/src/clt/ect_toy.py
python/tests/test_ect_compression_witness.py
python/scripts/ect_compression_witness.py
```

ACC as graph 1-form exactness and cycle affinities. Tests compute the antisymmetric log-ratio 1-form, cycle integrals, and potential reconstruction when exact. Representative files:

```
python/tests/test_acc_graph_one_form.py
```

D Appendix D: Zeno cascades and depth

D.1 Setup: frontier and Zeno criterion

Fix a threshold $\theta > 0$ and a scale index $j \in \mathbb{N}$. Let $E_j : [0, \infty) \rightarrow [0, \infty)$ be an activity observable at level j , and define the frontier

$$j_*(t) := \max\{j : E_j(t) \geq \theta\}.$$

Define crossing times and step durations by

$$t_j := \inf\{t : j_*(t) \geq j\}, \quad \Delta t_j := t_{j+1} - t_j.$$

A *Zeno cascade* is the event that (t_j) remains bounded as $j \rightarrow \infty$, equivalently $\sum_{j \geq j_0} \Delta t_j < \infty$.

D.2 Bridges, ports, and accounting

For each boundary j introduce a real Hilbert space U_j of port variables. A port input is a signal $u_j : [0, \infty) \rightarrow U_j$ (measurable), and the eliminated sector is packaged as a causal bridge operator Z_j producing a port response $y_j = Z_j[u_j]$. Define instantaneous interface power and positive supply by

$$p_j(t) := \langle u_j(t), y_j(t) \rangle_{U_j}, \quad p_j^+(t) := \max\{p_j(t), 0\},$$

and the supplied work over an interval by

$$W_j^+[s, t] := \int_s^t p_j^+(\tau) d\tau.$$

Using W^+ rather than net work avoids cancellations from backscatter and matches the use of “forward supply” in the latency argument.

D.3 Storage-based activity and the WORK quantum (Option B)

A bridge Z_j is *passive with storage* if there exists a nonnegative storage functional $S_j(\cdot) \geq 0$ on an internal state $x_j(t)$ such that for all $s < t$,

$$S_j(x_j(t)) - S_j(x_j(s)) \leq \int_s^t p_j(\tau) d\tau \leq W_j^+[s, t]. \quad (*)$$

Define the (storage-based) activity at level $j + 1$ relative to the activation time t_j by

$$E_{j+1}(t) := \begin{cases} 0, & t < t_j, \\ \sup_{s \in [t_j, t]} (S_j(x_j(s)) - S_j(x_j(t_j))), & t \geq t_j. \end{cases}$$

Define the next crossing time by

$$t_{j+1} := \inf\{t \geq t_j : E_{j+1}(t) \geq \theta\}.$$

This definition is robust to “pre-loading” before t_j because the baseline resets at t_j .

Lemma 7 (WORK from storage growth). *Assume that the passive-storage inequality (*) holds at boundary j , and that $t \mapsto E_{j+1}(t)$ is right-continuous. If $t_{j+1} < \infty$, then the supplied work over the crossing interval obeys*

$$W_j^+[t_j, t_{j+1}] \geq \theta.$$

Proof. By right-continuity and the definition of t_{j+1} , $E_{j+1}(t_{j+1}) \geq \theta$. For any $\varepsilon > 0$, choose $s_\varepsilon \in [t_j, t_{j+1}]$ with $S_j(x_j(s_\varepsilon)) - S_j(x_j(t_j)) \geq \theta - \varepsilon$. Apply (*) on $[t_j, s_\varepsilon]$ to get $W_j^+[t_j, s_\varepsilon] \geq \theta - \varepsilon$, and by monotonicity of W_j^+ in the interval, $W_j^+[t_j, t_{j+1}] \geq \theta - \varepsilon$. Letting $\varepsilon \rightarrow 0$ yields $W_j^+[t_j, t_{j+1}] \geq \theta$. \square

D.4 Integrated throughput (ICAP) and feasibility \Rightarrow latency bound

The correct throughput notion for bridges with memory is an *integrated* cap. For all $s < t$, with truncated input $u_j^{[s,t]} := u_j \mathbf{1}_{[s,t]}$ and truncated output $y_j^{[s,t]} := Z_j[u_j^{[s,t]}]$, define the windowed positive work

$$W_{j,\text{win}}^+[s, t] := \int_s^t \max\{\langle u_j(\tau), y_j^{[s,t]}(\tau) \rangle, 0\} d\tau.$$

We say boundary j satisfies an integrated capacity inequality (ICAP) if there exists $\Lambda(j) \in (0, \infty)$ such that

$$W_{j,\text{win}}^+[s, t] \leq \Lambda(j) \int_s^t \|u_j(\tau)\|_{U_j}^2 d\tau. \quad (\text{ICAP})$$

To convert this into a time lower bound we also isolate a primitive feasibility constraint: there exists $\bar{B}(j) \in (0, \infty)$ such that on each crossing interval,

$$\int_{t_j}^{t_{j+1}} \|u_j(\tau)\|_{U_j}^2 d\tau \leq \bar{B}(j) \Delta t_j. \quad (\text{FEAS})$$

This is an L^2 *energy-density cap* (average squared port amplitude), which is compatible with memory and allows short spikes.

Lemma 8 (Latency lower bound from WORK+ICAP+FEAS). *If Lemma 7 holds at boundary j , (ICAP)+(FEAS) hold on $[t_j, t_{j+1}]$, and the following fresh-start alignment holds on the crossing interval,*

$$W_j^+[t_j, t_{j+1}] \leq W_{j,\text{win}}^+[t_j, t_{j+1}], \quad (\text{ALIGN})$$

then

$$\Delta t_j \geq \frac{\theta}{\Lambda(j) \bar{B}(j)}.$$

Proof. By Lemma 7, $W_j^+[t_j, t_{j+1}] \geq \theta$. By (ALIGN), $W_{j,\text{win}}^+[t_j, t_{j+1}] \geq W_j^+[t_j, t_{j+1}] \geq \theta$. By ICAP and FEAS,

$$W_{j,\text{win}}^+[t_j, t_{j+1}] \leq \Lambda(j) \int_{t_j}^{t_{j+1}} \|u_j\|^2 \leq \Lambda(j) \bar{B}(j) \Delta t_j.$$

Rearrange. □

Remark 22 (When the alignment hypothesis holds). Condition (ALIGN) is automatic for memoryless bridges and holds for memoryful bridges when the boundary is "fresh" at activation (e.g. the port is inactive before t_j or the bridge state is reset at t_j).

D.5 No-Zeno criterion via divergence

Define the effective per-boundary capacity $\text{Cap}(j) := \Lambda(j) \bar{B}(j)$. Combining Lemma 8 with the elementary summability criterion yields a clean decision rule.

Theorem 10 (No-Zeno from divergent reciprocal capacity). *Assume that for all sufficiently large j , Lemma 8 applies with the same threshold $\theta > 0$. If*

$$\sum_{j \geq j_0} \frac{1}{\text{Cap}(j)} = \infty,$$

then $\sum_{j \geq j_0} \Delta t_j = \infty$ and hence no Zeno cascade occurs.

Proof. From Lemma 8, $\Delta t_j \geq \theta/\text{Cap}(j)$, and the sum diverges by assumption. \square

Appendix E (Section E.3) provides one concrete route to certify DIV via mode compression and per-atom ICAP, summarized in Theorem 11. This is optional evidence and does not change the formal criterion above.

D.6 Hard lemma slots (WORK/CAP/route)

The No-Zeno criterion above is elementary once three ingredients are available: (i) a storage-based work quantum (Lemma 7), (ii) an integrated throughput cap (ICAP), and (iii) a feasibility/energy-density cap (FEAS). In general systems, (ii)–(iii) are the true load-bearing frontier, and divergence is tied to how these capabilities scale with depth; we isolate them as hard lemma slots.

HL-WORK (storage existence). For each j , there exists a storage functional $S_j \geq 0$ and internal state $x_j(t)$ such that $(*)$ holds. *Failure mode:* the packaged bridge is effectively active (storage can grow without supplied work).

HL-CAP-X1 (bounded dissipation density). Assuming a linear time-invariant (LTI) bridge, the positive-real constraint from passivity is supplemented by a uniform upper bound on the dissipative part, yielding ICAP (ICAP). *Failure mode:* arbitrarily sharp resonant dissipation (unbounded dissipative density) destroys any uniform $\Lambda(j)$.

HL-CAP-X2 (finite memory / kernel mass). Assuming a causal convolution form, a bound $\int_0^\infty \|k_j(t)\| dt < \infty$ yields an L^2 operator bound and hence ICAP. *Failure mode:* long memory tails (nonintegrable kernel mass) or singular instantaneous response.

HL-CAP-X3 (finite stable passive realization). Assuming a finite-dimensional passive state realization with exponential stability and uniformly bounded parameters, ICAP follows with an explicit $\Lambda(j)$ (up to an additive term depending on initial storage). *Failure mode:* marginally stable modes or scale-dependent complexity producing near-resonances.

HL-ROUTE (route mismatch controls gain growth). Let $\text{RM}(j)$ quantify the mismatch between one-shot and two-step packaging across two levels. A representative target inequality is

$$\text{Gain}(j \leftarrow j+2) \leq \text{Gain}(j \leftarrow j+1) \text{Gain}(j+1 \leftarrow j+2) + \text{RM}(j),$$

so small (e.g. summable) route mismatch prevents hidden super-gain from shortcut reductions. *Failure mode:* large route mismatch allows a one-shot packaged bridge to have gain far larger than sequential composition predicts.

D.7 Checkable divergence criteria

Write $\text{Cap}(j) = \Lambda(j)\bar{B}(j)$. The No-Zeno condition in Theorem 10 is checkable from asymptotics. For example, if there exist $C, c > 0$ and exponents $\alpha, \beta \in \mathbb{R}$ such that for all large j ,

$$\Lambda(j) \leq C(j+1)^\alpha, \quad \bar{B}(j) \leq C(j+1)^\beta,$$

then $\text{Cap}(j) \leq C'(j+1)^{\alpha+\beta}$. In particular,

$$\sum_j \frac{1}{\text{Cap}(j)} = \infty \quad \text{whenever } \alpha + \beta \leq 1,$$

by comparison with the p -series. More refined tests can be stated in terms of dyadic subsequences when $1/\text{Cap}(j)$ is eventually nonincreasing.

D.8 Toy model families (necessity witnesses)

The criterion above is sharp in the following minimal sense. If

$$\sum_j \frac{1}{\text{Cap}(j)} < \infty,$$

then Zeno is compatible with WORK+ICAP+FEAS.

Zeno by fast capacity growth (DIV fails). Fix $\theta > 0$ and define $\text{Cap}(j) = 2^j$. Set $\Delta t_j := \theta/\text{Cap}(j) = \theta 2^{-j}$ and $t_{j+1} := t_j + \Delta t_j$. Then $\sum_j \Delta t_j < \infty$. Define $W_j^+[t_j, t_{j+1}] := \theta$ and choose any u_j with $\int_{t_j}^{t_{j+1}} \|u_j\|^2 = \Delta t_j$ so that FEAS holds with $\bar{B}(j) = 1$. ICAP holds with $\Lambda(j) = \text{Cap}(j)$ by construction: $W_j^+ = \Lambda(j) \int \|u_j\|^2$. This realizes Zeno while keeping the WORK quantum fixed.

Zeno by vanishing work quantum (WORK fails). Alternatively, even with a uniform capacity $\text{Cap}(j) \equiv 1$, if the required work per level decays (e.g. $\theta_j := 2^{-j}$), then taking $\Delta t_j = \theta_j$ yields $\sum_j \Delta t_j < \infty$. This shows that some form of uniform work quantization is genuinely load-bearing for No-Zeno.

D.9 Decision tree (settlement frontier)

The No-Zeno engine decomposes into three independent settlement points:

- **WORK (storage-based):** does the packaged bridge admit a passive storage inequality (*)?
- **CAP (integrated):** can one derive an ICAP bound (ICAP) (e.g. from passivity plus an additional structural condition)?
- **Divergence:** does $\sum_j 1/\text{Cap}(j)$ diverge?

If all three hold (together with FEAS), Theorem 10 rules out Zeno. If any one fails, Zeno is not excluded in this abstract setting, and the failure mode indicates what structural assumption must break (e.g. active bridges, uncontrolled resonance/memory, or rapidly growing effective capacity across scales).

E Appendix E: Toolkit theory—defects

E.1 Defects as quantitative relaxations of exact laws

A recurring pattern in this manuscript is that exact algebraic laws (idempotence, commutativity of reductions, audit monotonicity) hold only approximately in realistic multiscale settings. We therefore standardize a *defect calculus*: each desired law is paired with a nonnegative scalar measuring its violation, and simple norm inequalities propagate these defects.

Idempotence defect (TV/L1). Let Z be finite and let $\Delta(Z)$ denote the simplex of distributions on Z . For a linear map $E : \Delta(Z) \rightarrow \Delta(Z)$ represented by a matrix acting on row distributions (so $\mu \mapsto \mu E$), define the total-variation idempotence defect by

$$\delta_{\text{TV}}(E) := \sup_{\mu \in \Delta(Z)} \|\mu(E^2 - E)\|_{\text{TV}}. \quad (\text{D-TK-DEF-02})$$

For finite Z , $\|\nu\|_{\text{TV}} = \frac{1}{2}\|\nu\|_1$ for signed measures ν , hence $\delta_{\text{TV}}(E)$ is computable.

Remark 23 (Extreme points attain the TV idempotence defect). Because $\mu \mapsto \|\mu(E^2 - E)\|_{\text{TV}}$ is convex in μ , the supremum in (D-TK-DEF-02) is attained at extreme points:

$$\delta_{\text{TV}}(E) = \max_{z \in Z} \|\delta_z(E^2 - E)\|_{\text{TV}} = \frac{1}{2} \max_i \sum_j |(E^2 - E)_{ij}|.$$

This is the norm used by the reproducible diagnostics in the accompanying code (row-vector convention).

Route mismatch defect (abstract). Many constructions in multiscale reduction depend on an elimination *route*. Let $\text{Pack}_{j \leftarrow k}$ denote the effective packaging map from depth k to depth j (with $k > j$), and let $d(\cdot, \cdot)$ be a metric on the space of packaged objects (e.g. an operator norm in linear settings). Define the route mismatch defect across two steps by

$$\text{RM}(j) := d(\text{Pack}_{j \leftarrow j+2}, \text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}). \quad (\text{D-TK-ROUTE-01})$$

This quantity is purely definitional here; in a later toolkit step we relate it to controlled growth of effective capacities.

Route dependence, capacity, and a one-step control inequality

To make (D-TK-ROUTE-01) operational we isolate a minimal setting in which packaged objects are linear operators between normed port spaces and the metric d is induced by an operator norm. This covers, for example, linearized bridges and bounded-gain summaries of more complex components.

Abstract packaging maps. For integers $k > j$, let $\text{Pack}_{j \leftarrow k}$ denote the packaged object obtained by eliminating intermediate structure between depths k and j . We do not assume functoriality or strict associativity: route dependence is precisely the possibility that $\text{Pack}_{j \leftarrow j+2}$ differs from $\text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}$.

Linear-operator specialization.

Assume each $\text{Pack}_{j \leftarrow k}$ is a bounded linear operator between normed spaces, with operator norm $\|\cdot\|$.

We take the metric in (D-TK-ROUTE-01) to be

$$d(A, B) := \|A - B\|.$$

Define the induced *gain* of a packaged object by

$$\text{Gain}(j \leftarrow k) := \|\text{Pack}_{j \leftarrow k}\|.$$

Lemma 9 (Route mismatch controls gain growth). *In the default linear-operator setting above, the two-step route mismatch*

$$\begin{aligned} \text{RM}(j) = & \|\text{Pack}_{j \leftarrow j+2} \\ & - \text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}\| \end{aligned}$$

implies the gain bound

$$\begin{aligned} \text{Gain}(j \leftarrow j+2) \leq & \text{Gain}(j \leftarrow j+1) \text{Gain}(j+1 \leftarrow j+2) \\ & + \text{RM}(j). \end{aligned} \tag{L-TK-ROUTE-01}$$

Proof. By the triangle inequality,

$$\begin{aligned} \|\text{Pack}_{j \leftarrow j+2}\| \leq & \|\text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}\| \\ & + \|\text{Pack}_{j \leftarrow j+2} - \text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}\|. \end{aligned}$$

Submultiplicativity of the operator norm gives

$$\|\text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}\| \leq \|\text{Pack}_{j \leftarrow j+1}\| \|\text{Pack}_{j+1 \leftarrow j+2}\|.$$

The final term is $\text{RM}(j)$ by definition. \square

Remark 24 (Holonomy without teleology). The purpose of $\text{RM}(j)$ is to quantify *route dependence* (a holonomy-like effect) of reduction/packaging. This is logically distinct from directionality: a system may exhibit nontrivial route mismatch while having zero audit (e.g. zero affinities or vanishing path reversal asymmetry under the stated assumptions). The arrow-of-time results in Section 7 remain certified by the audit functional and its monotonicity, not by route dependence alone.

E.2 Bridge objects: ports, passivity, and integrated throughput

We isolate a minimal, representation-agnostic interface for *bridge objects*: a packaged eliminated sector acting causally on port variables of the retained sector. The purpose is not to specialize to any particular physics, but to standardize what must be provided to make WORK/CAP-style arguments and accounting statements well-posed.

Definition 28 (Ports, causal bridge, and accounting). A port space is a finite-dimensional real Hilbert space U with inner product $\langle \cdot, \cdot \rangle_U$ and norm $\| \cdot \|_U$. Fix a finite time horizon $T > 0$ and write $\mathcal{U} := L^2([0, T]; U)$ for square-integrable U -valued signals. A (deterministic) bridge operator is a map $Z : \mathcal{U} \rightarrow \mathcal{U}$. It is causal if for all $t \in [0, T]$ and all $u, v \in \mathcal{U}$,

$$u(\tau) = v(\tau) \text{ for a.e. } \tau \in [0, t] \implies (Zu)(\tau) = (Zv)(\tau) \text{ for a.e. } \tau \in [0, t].$$

Given an input $u \in \mathcal{U}$ and output $y = Zu$, define the instantaneous supply rate

$$p(t) := \langle u(t), y(t) \rangle_U, \quad p^+(t) := \max\{p(t), 0\},$$

and the positive supplied work over $[s, t] \subseteq [0, T]$ by

$$W^+[s, t] := \int_s^t p^+(\tau) d\tau.$$

Definition 29 (Passivity with storage). A causal bridge $Z : \mathcal{U} \rightarrow \mathcal{U}$ is passive with storage (cf. [Wil72]) if there exist: (i) a state space X , (ii) a nonnegative storage functional $S : X \rightarrow [0, \infty)$, and (iii) for each input $u \in \mathcal{U}$ a state trajectory $x_u : [0, T] \rightarrow X$, such that for all $0 \leq s \leq t \leq T$,

$$S(x_u(t)) - S(x_u(s)) \leq \int_s^t \langle u(\tau), (Zu)(\tau) \rangle_U d\tau. \quad (\text{D-TK-BRG-02})$$

Definition 30 (Integrated capacity inequality (ICAP)). A causal bridge $Z : \mathcal{U} \rightarrow \mathcal{U}$ satisfies an integrated capacity inequality (ICAP) if there exists $\Lambda \in (0, \infty)$ such that for all $0 \leq s \leq t \leq T$ and all inputs $u \in \mathcal{U}$, with truncated input $u^{[s, t]} := u \mathbf{1}_{[s, t]}$ and output $y^{[s, t]} := Z(u^{[s, t]})$,

$$W^+[s, t] \leq \Lambda \int_s^t \|u(\tau)\|_U^2 d\tau. \quad (\text{D-TK-BRG-03})$$

ICAP is the natural throughput notion for bridges with memory: it allows short spikes in $p(t)$ while bounding cumulative positive supply by an L^2 budget.

Remark 25. The incremental (activated) form above avoids prehistory contamination in memory systems; a full-output per-subinterval ICAP can fail even for simple delays.

Lemma 10 (Causality is closed under composition). Assume the setup of Definition 28. If $Z_1 : \mathcal{U} \rightarrow \mathcal{U}$ and $Z_2 : \mathcal{U} \rightarrow \mathcal{U}$ are causal, then $Z_2 \circ Z_1$ is causal.

Proof. If u and v agree a.e. on $[0, t]$, then by causality of Z_1 the signals $Z_1 u$ and $Z_1 v$ agree a.e. on $[0, t]$. Applying causality of Z_2 yields $(Z_2 \circ Z_1)u = (Z_2 \circ Z_1)v$ a.e. on $[0, t]$. \square

Lemma 11 (Parallel sum: passivity and ICAP constants add). *Assume the setup of Definition 28. Let $Z_1, Z_2 : \mathcal{U} \rightarrow \mathcal{U}$ be causal bridges on the same port space U , and define $Z := Z_1 + Z_2$ by $(Zu)(t) := (Z_1 u)(t) + (Z_2 u)(t)$.*

1. *If Z_1 and Z_2 are passive with storage (Definition 29), then Z is passive with storage (with storage equal to the sum of storages on the product state space).*
2. *If Z_1 and Z_2 satisfy ICAP with constants Λ_1 and Λ_2 , then Z satisfies ICAP with constant $\Lambda_1 + \Lambda_2$.*

Proof. For (2), fix u and $[s, t]$, and set $u^{[s,t]} := u \mathbf{1}_{[s,t]}$. Let $y_i = Z_i u^{[s,t]}$. Then $p = p_1 + p_2$ pointwise. Since $(a + b)^+ \leq a^+ + b^+$ for real a, b , we have $W^+ \leq W_1^+ + W_2^+$, and applying ICAP to each term yields

$$W^+[s, t] \leq (\Lambda_1 + \Lambda_2) \int_s^t \|u(\tau)\|_U^2 d\tau.$$

For (1), add the two passivity inequalities (D-TK-BRG-02) on a product state space. Use linearity of Z in the supply rate. \square

Remark 26 (Serial composition: gain submultiplicativity is model-free; passivity needs port-compatibility). For bounded linear bridges on \mathcal{U} one can define an operator-norm gain $\text{Gain}(Z) := \|Z\|_{\mathcal{U} \rightarrow \mathcal{U}}$, which satisfies $\text{Gain}(Z_2 \circ Z_1) \leq \text{Gain}(Z_2)\text{Gain}(Z_1)$. Such a bound yields an ICAP constant via Cauchy–Schwarz (so $\text{Cap}(Z) \leq \text{Gain}(Z)$), hence the resulting ICAP constants are submultiplicative under serial composition in this default linear setting. By contrast, passivity under serial/feedback interconnection depends on how port variables are paired and how the external supply decomposes into internal supplies; we therefore record the parallel-sum lemma above as the universally valid closure property in the present one-port notation.

Remark 27 (Relation to primitives P5 and P6). In the main text, **P5** (packaging) is expressed as a completion endomap E with fixed points, while the present appendix expresses packaging at a boundary as a *causal bridge* Z acting on port signals. Both viewpoints represent “the eliminated sector as an effective operator”; one acts on descriptions (e.g. distributions), the other acts on boundary interactions. The audit components used in this

manuscript—path reversal asymmetry (Section 7) and graph 1-form affinities (Section 6)—are instances of **P6**-style accounting: they certify non-trivial drive/directionality and obey monotonicity principles under coarse observation.

E.3 Emergent coercivity template via sector compression

E.3.1 Summary: slots and divergence consequence

Theorem 11 (ECT summary: linear ICAP capacity and DIV from three slots). *Fix a depth-indexed family of packaged bridges $(Z_j)_j$ with port spaces U_j and a horizon $T > 0$. Assume the following slot conditions hold for all sufficiently large j :*

HL-ECT-1 (Mode compression) Z_j admits an atomic/sector decomposition with mode count $m_j \leq C_0(j + 1)$. A sufficient route in this appendix is: sectorization + P5-min + quantized index, which yields $m_j \lesssim (j + 1)$ (see Corollary 7).

HL-ECT-3 (Uniform per-atom ICAP) each atom $Z_{j,r}$ in the decomposition satisfies ICAP with a uniform constant Λ_0 (independent of j, r). A sufficient condition is bounded kernel mass (Lemma 13). In particular, dissipative atoms with semigroup decay (Definition 32) and balanced coupling (Definition 33) satisfy uniform ICAP via Lemma 12 and Corollary 6.

HL-ECT-2 (Feasibility removes lossless directions, if strictness is needed) feasibility gating excludes lossless directions so that a coercive norm is induced on feasible inputs (see Remark 32, Lemma 15, and the depth-scaling lemma Lemma 16).

Then the ICAP capacity constant obeys a linear bound

$$\text{Cap}(Z_j) \leq \Lambda_0 C_0 (j + 1),$$

and consequently

$$\sum_{j \geq 0} \frac{1}{\text{Cap}(Z_j)} = \infty.$$

In particular, this verifies the DIV ingredient used in the No-Zeno criterion (Theorem 10 in Appendix D). WORK quantization and a feasibility/energy-density cap are supplied there.

Proof. Under **HL-ECT-1** and **HL-ECT-3**, Lemma 14 gives $\text{Cap}(\mathbf{Z}_j) \leq m_j \Lambda_0 \leq \Lambda_0 C_0(j+1)$, and the divergence of $\sum_j 1/\text{Cap}(\mathbf{Z}_j)$ follows by comparison with the harmonic series. \square

Remark 28 (Non-implication: why ECT is an add-on, not automatic). The ECT conclusion is *not* a consequence of passivity alone. Passivity (existence of a storage functional with $S(t) - S(s) \leq \int_s^t \langle u, \mathbf{Z}u \rangle$) does not by itself control ICAP constants, does not force mode compression, and does not eliminate lossless feasible directions. Even within convolution/finite-memory classes, passivity alone allows arbitrarily large ICAP constants via high- Q slowly decaying internal directions; the balanced coupling condition (Definition 33) is one explicit way to exclude this. The template therefore isolates explicit structural sources of coercivity/capacity control: **P4**-style sector compression (bounded sector complexity), **P5** packaging outcomes (e.g. atomization with finite memory), and **P2** feasibility that removes null directions. These conditions couple packaging+accounting to a through-put certificate, but they do not by themselves imply novelty (Section 8) or directionality (Sections 6–7).

Definition 31 (Atomic/sector decompositions of a packaged bridge (ID: D-ECT-ATOM-01)). *Fix a port space U and a finite horizon $T > 0$, and write $\mathcal{U} := L^2([0, T]; U)$. Let $\mathbf{Z} : \mathcal{U} \rightarrow \mathcal{U}$ be a causal bridge (Definition 28). An atomic (sector) decomposition of \mathbf{Z} is a representation*

$$\mathbf{Z} = \sum_{r=1}^m \mathbf{Z}_r$$

as a parallel sum of causal bridges $\mathbf{Z}_r : \mathcal{U} \rightarrow \mathcal{U}$ on the same port space (so $(\mathbf{Z}u)(t) = \sum_r (\mathbf{Z}_r u)(t)$). We call m the mode count of the chosen decomposition. In depth-indexed settings we write $\mathbf{Z} = \mathbf{Z}_j$ and denote the mode count by m_j .

E.3.2 Dissipative atoms and semigroup decay

Definition 32 (Dissipative atom representation (ID: D-ECT-ATOM-SS-01)). *Let \mathbf{Z}_r be an atom in an atomic decomposition (Definition 31) acting on $\mathcal{U} := L^2([0, T]; U)$. We say \mathbf{Z}_r admits a dissipative atom representation if it has a kernel*

$$K_r(\tau) = C_r e^{A_r \tau} B_r, \quad \tau \geq 0,$$

with an internal Hilbert space H_r , bounded linear maps $A_r : H_r \rightarrow H_r$, $B_r : U \rightarrow H_r$, $C_r : H_r \rightarrow U$, and decay

$$\|e^{A_r \tau}\|_{\text{op}} \leq e^{-\lambda_r \tau} \quad \text{for some } \lambda_r > 0.$$

All windowed/interval inequalities in this appendix are evaluated using truncated inputs $u^{[s,t]} := u \mathbf{1}_{[s,t]}$, as in Definition 34.

Definition 33 (Balanced coupling (ID: D-ECT-BAL-01)).

In the setting of Definition 32, fix a constant $\Lambda_0 > 0$. We say the atom Z_r satisfies the balanced coupling condition with constant Λ_0 if

$$\|C_r\|_{\text{op}} \|B_r\|_{\text{op}} \leq \Lambda_0 \lambda_r,$$

where $\lambda_r > 0$ is the semigroup decay rate from Definition 32.

Definition 34 (ICAP capacity functional (ID: D-ECT-CAP-01)). For a causal bridge $Z : \mathcal{U} \rightarrow \mathcal{U}$, define its ICAP capacity constant

$$\text{Cap}(Z) := \inf \left\{ \Lambda \in (0, \infty) : \begin{array}{l} W^+[s,t] \leq \Lambda \int_s^t \|u(\tau)\|_U^2 d\tau \\ \text{for all } 0 \leq s \leq t \leq T \text{ and } u \in \mathcal{U} \end{array} \right\},$$

where $u^{[s,t]} := u \mathbf{1}_{[s,t]}$, $y^{[s,t]} := Z(u^{[s,t]})$, and $W^+[s,t]$ is computed from $u^{[s,t]}$, $y^{[s,t]}$ as in Definition 28. (Equivalently, $\text{Cap}(Z)$ is the best constant in (D-TK-BRG-03) when it holds.)

Lemma 12 (Balanced decay implies a uniform kernel-mass bound (ID: L-ECT-BAL-KMASS-01)). Assume Z_r admits a dissipative atom representation (Definition 32) and satisfies the balanced coupling condition with constant Λ_0 (Definition 33). Then the operator-valued kernel $K_r(\tau) = C_r e^{A_r \tau} B_r$ satisfies

$$\int_0^\infty \|K_r(\tau)\|_{\text{op}} d\tau \leq \Lambda_0.$$

Proof. By submultiplicativity and decay,

$$\|K_r(\tau)\|_{\text{op}} \leq \|C_r\|_{\text{op}} \|e^{A_r \tau}\|_{\text{op}} \|B_r\|_{\text{op}} \leq \|C_r\|_{\text{op}} \|B_r\|_{\text{op}} e^{-\lambda_r \tau}.$$

Integrating gives

$$\int_0^\infty \|K_r(\tau)\|_{\text{op}} d\tau \leq \frac{\|C_r\|_{\text{op}} \|B_r\|_{\text{op}}}{\lambda_r} \leq \Lambda_0,$$

by Definition 33. □

Lemma 13 (Finite kernel-mass convolution bridges satisfy ICAP for truncated inputs (ID: L-ECT-FMEM-01)). *Let U be a finite-dimensional port space and $\mathcal{U} := L^2([0, T]; U)$. Consider a causal convolution bridge $Z : \mathcal{U} \rightarrow \mathcal{U}$ of the form*

$$(Zu)(t) := \int_0^t K(t - \tau) u(\tau) d\tau,$$

where $K : [0, \infty) \rightarrow \text{End}(U)$ is measurable and $k(\tau) := \|K(\tau)\|_{\text{op}}$ satisfies $M := \int_0^\infty k(\tau) d\tau < \infty$. For $0 \leq s \leq t \leq T$ define the truncated input $u^{[s,t]} := u \mathbf{1}_{[s,t]}$ and output $y^{[s,t]} := Z(u^{[s,t]})$. Then the positive supplied work computed from $y^{[s,t]}$ satisfies

$$W^+[s, t] \leq M \int_s^t \|u(\tau)\|_U^2 d\tau.$$

Proof. For $\tau \in [s, t]$ we have

$$\|y^{[s,t]}(\tau)\|_U \leq \int_s^\tau k(\tau - \sigma) \|u(\sigma)\|_U d\sigma,$$

so by Young's inequality, $\|y^{[s,t]}\|_{L^2([s,t])} \leq M \|u\|_{L^2([s,t])}$. Then

$$\begin{aligned} W^+[s, t] &\leq \int_s^t \|u(\tau)\|_U \|y^{[s,t]}(\tau)\|_U d\tau \\ &\leq \|u\|_{L^2([s,t])} \|y^{[s,t]}\|_{L^2([s,t])} \\ &\leq M \int_s^t \|u(\tau)\|_U^2 d\tau. \end{aligned}$$

□

Corollary 6 (Balanced atoms have uniform ICAP).

C-ECT-BAL-ICAP-01.

Assume Z_r admits a dissipative atom representation (Definition 32) and satisfies the balanced coupling condition with constant Λ_0 (Definition 33). Then $\text{Cap}(Z_r) \leq \Lambda_0$ in the sense of Definition 34.

That is, the ICAP inequality holds with constant Λ_0 for truncated inputs $u^{[s,t]}$ and positive work $W^+[s, t]$.

Proof. Lemma 12 gives a kernel-mass bound $M \leq \Lambda_0$. Applying Lemma 13 with this M yields

$$W^+[s, t] \leq \Lambda_0 \int_s^t \|u(\tau)\|_U^2 d\tau$$

for all windows. Hence $\text{Cap}(Z_r) \leq \Lambda_0$ by Definition 34. □

Remark 29 (Finite-memory bridges as a P5 packaging target). Lemma 13 provides a concrete sufficient condition for the uniform per-atom ICAP hypothesis used in Lemma 14. Accordingly, one natural packaging target for **P5** is to return atom bridges that admit a convolution representation with bounded kernel mass M , so that $\text{Cap}(\mathbf{Z}_r) \leq M$ is explicit and checkable. This is only a model class and not a universal consequence of packaging.

Lemma 14 (Mechanical ECT: sector ICAP + mode count \Rightarrow linear capacity (ID: L-ECT-MECH-01)). *Let $\mathbf{Z} = \sum_{r=1}^m \mathbf{Z}_r$ be an atomic decomposition in the sense of Definition 31. Assume each atom satisfies ICAP with a uniform constant Λ_0 , i.e. for all r and all $0 \leq s \leq t \leq T$,*

$$W_r^+[s, t] \leq \Lambda_0 \int_s^t \|u(\tau)\|_{\mathcal{U}}^2 d\tau,$$

where $u^{[s,t]} := u \mathbf{1}_{[s,t]}$ and $W_r^+[s, t]$ is computed from $y_r = \mathbf{Z}_r u^{[s,t]}$. Then \mathbf{Z} satisfies ICAP with constant $m\Lambda_0$, hence

$$\text{Cap}(\mathbf{Z}) \leq m\Lambda_0.$$

In particular, in a depth-indexed family $(\mathbf{Z}_j)_j$ with mode counts $m_j \leq C_0(j+1)$ we obtain

$$\text{Cap}(\mathbf{Z}_j) \leq \Lambda_0 C_0(j+1), \quad \text{and consequently} \quad \sum_{j \geq 0} \frac{1}{\text{Cap}(\mathbf{Z}_j)} = \infty.$$

Proof. Apply Lemma 11(2) iteratively to the parallel sum $\mathbf{Z} = \mathbf{Z}_1 + \dots + \mathbf{Z}_m$ to conclude that the ICAP constants add. Thus $W^+[s, t] \leq \sum_{r=1}^m W_r^+[s, t] \leq m\Lambda_0 \int_s^t \|u(\tau)\|_{\mathcal{U}}^2 d\tau$. The bound on $\text{Cap}(\mathbf{Z})$ is immediate from Definition 34. If $\text{Cap}(\mathbf{Z}_j) \leq \Lambda_0 C_0(j+1)$, then $\sum_{j \geq 0} 1/\text{Cap}(\mathbf{Z}_j)$ diverges by comparison with the harmonic series. \square

Remark 30 (DIV-ready chain). Fix a depth j and interval $[t_j, t_{j+1}]$ with input u_j and output $y_j = \mathbf{Z}_j u_j$. If the positive supplied work satisfies $W_j^+[t_j, t_{j+1}] \geq \theta$ for some $\theta > 0$, then by ICAP,

$$\int_{t_j}^{t_{j+1}} \|u_j(\tau)\|_{\mathcal{U}_j}^2 d\tau \geq \frac{\theta}{\text{Cap}(\mathbf{Z}_j)}.$$

If, in addition, there is a bound

$$\int_{t_j}^{t_{j+1}} \|u_j(\tau)\|_{\mathcal{U}_j}^2 d\tau \leq B \Delta t_j, \quad \Delta t_j := t_{j+1} - t_j,$$

then $\Delta t_j \geq \theta/(B \text{Cap}(\mathbf{Z}_j))$. Under a linear capacity bound $\text{Cap}(\mathbf{Z}_j) \leq C_1(j+1)$, the sum $\sum_j \Delta t_j$ diverges by comparison with the harmonic series.

Remark 31 (Route mismatch as an additive gain error). In the default linear-operator specialization of Appendix E, Lemma 9 and (L-TK-ROUTE-01) show that route mismatch contributes an additive term to two-step gain growth. Thus, when estimating gains along successive reductions, $\text{RM}(j)$ appears as an additive perturbation term. When the mismatch operator itself satisfies ICAP with a comparable constant, this additive term can be absorbed into the same $\text{Cap}(\cdot)$ calculus; otherwise it remains an explicit robustness correction.

E.3.3 Coercivity from feasibility gating (P2)

Definition 35 (Feasibility gate and induced coercive norm (ID: D-ECT-GATE-01)). *Fix a finite horizon $T > 0$, a port Hilbert space U_j , and write $\mathcal{U}_j := L^2([0, T]; U_j)$. Let $G_j : U_j \rightarrow U_j$ be a self-adjoint positive semidefinite operator. Define the feasible input class*

$$\mathcal{U}_j^{\text{feas}} := \{u \in \mathcal{U}_j : u(t) \in \text{Ran}(G_j) \text{ for a.e. } t\},$$

the pointwise seminorm on U_j ,

$$|v|_{X_j}^2 := \langle v, G_j v \rangle \quad (v \in U_j),$$

and the associated integrated energy on a window $[s, t]$,

$$|u|_{L_{X_j}^2(s,t)}^2 := \int_s^t |u(\tau)|_{X_j}^2 d\tau.$$

On $\mathcal{U}_j^{\text{feas}}$ this is a genuine norm, since the feasibility constraint removes the kernel of G_j .

Lemma 15 (Coercivity on the feasible set (ID: L-ECT-COER-01)). *Let $Z_j : \mathcal{U}_j \rightarrow \mathcal{U}_j$ be a causal bridge. Assume that for every feasible input $u \in \mathcal{U}_j^{\text{feas}}$ and every $0 \leq s \leq t \leq T$, with output $y = Z_j u$,*

$$\int_s^t \langle u(\tau), y(\tau) \rangle d\tau \geq \int_s^t \langle u(\tau), G_j u(\tau) \rangle d\tau.$$

Then, on $\mathcal{U}_j^{\text{feas}}$,

$$\int_s^t \langle u(\tau), (Z_j u)(\tau) \rangle d\tau \geq |u|_{L_{X_j}^2(s,t)}^2.$$

Proof. The right-hand side is exactly $\int_s^t \langle u(\tau), G_j u(\tau) \rangle d\tau$ by definition of $|u|_{L_{X_j}^2(s,t)}^2$. \square

Remark 32 (Hard slot HL-P2-CORE (P2 must eliminate lossless feasible directions)). *Slot (HL-P2-CORE)*. There exists a self-adjoint $G_j \succeq 0$ whose kernel coincides with the lossless directions of the packaged bridge at depth j (e.g. in a linear specialization, $\ker(G_j) = \ker(\text{Sym } Z_j)$), and **P2** feasibility implies $u(t) \in \text{Ran}(G_j)$ for a.e. t . Once this slot holds, Lemma 15 turns the gating statement into the strict/coercive inequality on feasible inputs. All difficulty is concentrated in establishing HL-P2-CORE for the particular packaging construction at hand.

Lemma 16 (Feasibility shrinkage \Rightarrow depth-scaled coercivity (ID: L-ECT-SHRINK-01)). *Assume the setup of Definition 35 and Lemma 15. Suppose there exists a constant $a_j > 0$ such that for every feasible input $u \in \mathcal{U}_j^{\text{feas}}$ and every interval $0 \leq s \leq t \leq T$,*

$$\|u\|_{L^2_X(s,t)}^2 \geq a_j^2 \int_s^t \|u(\tau)\|_{U_j}^2 d\tau.$$

Then for every feasible u and every $0 \leq s \leq t \leq T$ we have the depth-scaled coercive lower bound

$$\int_s^t \langle u(\tau), (Z_j u)(\tau) \rangle d\tau \geq a_j^2 \int_s^t \|u(\tau)\|_{U_j}^2 d\tau.$$

Equivalently, relative to the baseline $L^2([0, T]; U_j)$ norm the effective coercivity coefficient is $c_j = a_j^2$.

Proof. Combine Lemma 15 with the stated shrinkage inequality. \square

Remark 33 (Examples of depth scaling). If $a_j \gtrsim (j+1)$, then $c_j = a_j^2 \gtrsim (j+1)^2$. If $a_j \gtrsim 2^{sj}$, then $c_j = a_j^2 \gtrsim 2^{2sj}$. These are purely algebraic consequences of the shrinkage hypothesis.

E.3.4 Mode compression from sectorization (P4) and minimality (P5)

Definition 36 (Sectorization of the port space (ID: D-ECT-SECT-01)). *Fix a depth j and a port Hilbert space U_j . A sectorization is the choice of a finite index set Q_j and a family of orthogonal projections $\{\Pi_{j,q} : U_j \rightarrow U_j\}_{q \in Q_j}$ such that $\Pi_{j,q} \Pi_{j,q'} = 0$ for $q \neq q'$ and $\sum_{q \in Q_j} \Pi_{j,q} = I$. We extend $\Pi_{j,q}$ pointwise to signals $u \in \mathcal{U}_j := L^2([0, T]; U_j)$ by $(\Pi_{j,q} u)(t) := \Pi_{j,q}(u(t))$. A causal bridge $Z_j : \mathcal{U}_j \rightarrow \mathcal{U}_j$ is sector-respecting if*

$$Z_j = \sum_{q \in Q_j} \Pi_{j,q} Z_j \Pi_{j,q},$$

i.e. the output on each sector depends only on the input restricted to that sector (no cross-sector coupling).

Definition 37 (Sector-respecting minimal decompositions (P5-min) (ID: D-ECT-P5MIN-01)). *Assume a sectorization as in Definition 36. An atomic decomposition $Z_j = \sum_{r=1}^m Z_{j,r}$ (Definition 31) is sector-respecting if for each atom r there exists a sector label $\sigma(r) \in Q_j$ such that*

$$Z_{j,r} = \Pi_{j,\sigma(r)} Z_{j,r} \Pi_{j,\sigma(r)}.$$

*Define the sector-minimal mode count m_j^{\min} to be the smallest m for which Z_j admits a sector-respecting decomposition into m nonzero atoms. We say packaging satisfies **P5-min** at depth j if it returns a sector-respecting decomposition achieving m_j^{\min} .*

Lemma 17 (Sectorization bounds the minimal mode count (ID: L-ECT-MODE-UB-01)). *Assume a sectorization as in Definition 36. Then any sector-respecting atomic decomposition of Z_j can be merged within each sector to produce a sector-respecting decomposition with at most $|Q_j|$ nonzero atoms. In particular,*

$$m_j^{\min} \leq |Q_j|.$$

*Consequently, under **P5-min** packaging (Definition 37) the returned mode count satisfies $m_j \leq |Q_j|$.*

Proof. Let $Z_j = \sum_{r=1}^m Z_{j,r}$ be sector-respecting with labels $\sigma(r) \in Q_j$. For each $q \in Q_j$ define the merged atom $\tilde{Z}_{j,q} := \sum_{\{r:\sigma(r)=q\}} Z_{j,r}$. Each $\tilde{Z}_{j,q}$ is sector-respecting and supported on q , and $Z_j = \sum_{q \in Q_j} \tilde{Z}_{j,q}$. Discarding any zero $\tilde{Z}_{j,q}$ yields a sector-respecting decomposition with at most $|Q_j|$ nonzero atoms, proving $m_j^{\min} \leq |Q_j|$. The final statement follows from the definition of **P5-min**. \square

Definition 38 (Quantized index hypothesis (ID: D-ECT-P4IDX-01)).

Fix depth j and a sectorization as in Definition 36. We say the sectorization has a quantized index if there exists a constant $C_0 \geq 1$ and an injective map

$$\iota_j : Q_j \hookrightarrow \{0, 1, 2, \dots, \lceil C_0(j+1) \rceil - 1\}.$$

Equivalently, each sector admits a distinct integer label from a set of size $O(j)$. This is an explicit modeling/design hypothesis encoding bounded interface complexity; it is not implied by the other primitives and must be verified in each intended instantiation.

Lemma 18 (Quantized index implies linear sector count (ID: L-ECT-Q-SIZE-01)). *Under the quantized index hypothesis of Definition 38, we have*

$$|Q_j| \leq \lceil C_0(j+1) \rceil.$$

Proof. Since ι_j is injective, $|Q_j| \leq |\{0, 1, \dots, \lceil C_0(j+1) \rceil - 1\}| = \lceil C_0(j+1) \rceil$. \square

Corollary 7 (Linear mode bound from P4 index (ID: C-ECT-HL1-01)).

Assume a sectorization (Definition 36) and P5-min packaging (Definition 37). Assume also the quantized index hypothesis (Definition 38). Then the packaged mode count satisfies

$$m_j \leq |Q_j| \leq \lceil C_0(j+1) \rceil.$$

Proof. The inequality $m_j \leq |Q_j|$ follows from Lemma 17 under P5-min packaging. The bound on $|Q_j|$ is Lemma 18. \square

Remark 34 (Hard lemma slots behind coercivity emergence). Lemma 14 is purely algebraic. In applications, the substantive work is to justify: (i) a *sector/atom model* produced by packaging (P5) in which a uniform per-atom ICAP bound holds, (ii) a *mode-count compression* estimate $m_j \leq C_0(j+1)$ (often a P4-style sectorization statement; cf. Lemma 17 and Corollary 7), and (iii) a feasibility/gating principle (P2) strong enough to rule out lossless feasible directions (if strict coercivity is required beyond ICAP).

Audit monotonicity defect (optional). Let \mathcal{A} be an audit functional equipped with a target monotonicity principle under coarse observation (e.g. the data-processing inequality for path reversal asymmetry). Given a refinement $f : Z \rightarrow X$ and a further coarsening $g : X \rightarrow Y$, define the monotonicity defect on an argument (\cdot) by

$$\text{Def}_{\mathcal{A}}(g \mid f; \cdot) := \max\{0, \mathcal{A}((g \circ f)_{\#}(\cdot)) - \mathcal{A}(f_{\#}(\cdot))\}. \quad (\text{D-TK-DEF-01})$$

When DPI holds in the stated setting, this defect vanishes identically.

Defect propagation rules (toolkit)

The following inequalities are repeatedly useful. They require no system-specific structure beyond norms and triangle inequalities.

1. **Composition perturbation (triangle inequality).** For composable maps A, B, A', B' between normed spaces with an operator norm $\|\cdot\|$,

$$\|A \circ B - A' \circ B'\| \leq \|A - A'\| \|B\| + \|A'\| \|B - B'\|.$$

2. **Idempotence factorization.** For any map E with a submultiplicative operator norm,

$$E^2 - E = E \circ (E - \text{Id}) \quad \Rightarrow \quad \|E^2 - E\| \leq \|E\| \|E - \text{Id}\|.$$

3. **Audit monotonicity as a zero-defect rule.** In contexts where \mathcal{A} is known to contract under pushforward (as in Section 7), $\text{Def}_{\mathcal{A}}(g | f; \cdot) = 0$ provides a standardized certificate that coarse observation does not create spurious directionality.

References

- [AGH⁺12] Bernhard Altaner, Stefan Grosskinsky, Stephan Herminghaus, Lukas Katthän, Marc Timme, and Jürgen Vollmer. Network representations of non-equilibrium steady states: Cycle decompositions, symmetries and dominant paths. *Physical Review E*, 85:041133, 2012.
- [Bor94] Francis Borceux. *Handbook of Categorical Algebra 1: Basic Category Theory*. Cambridge University Press, 1994.
- [BSvdM25] M. T. Bauer, U. Seifert, and J. van der Meer. Stroboscopic measurements in markov networks: exact generator reconstruction vs. thermodynamic inference. *Journal of Physics A: Mathematical and Theoretical*, 58:125001, 2025.
- [CK11] Imre Csiszár and János Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2 edition, 2011.
- [Coh63] Paul J. Cohen. The independence of the continuum hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 50(6):1143–1148, 1963.
- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2 edition, 2006.

- [DCLP23] S. Dal Cengio, V. Lecomte, and M. Polettini. Geometry of nonequilibrium reaction networks. *Physical Review X*, 13:021040, 2023.
- [DP02] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2 edition, 2002.
- [DvdMS23] J. Degünther, J. van der Meer, and U. Seifert. Coarse-grained entropy production in markov processes. *arXiv preprint*, 2023.
- [Hyt15] Tapani Hyttinen. Counting measure and forking in finite models. In Åsa Hirvonen, Juha Juutinen, Juha Kontinen, Kimmo Räsänen, and Jouko Väänänen, editors, *Logic Without Borders: Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics*, volume 44 of *Trends in Logic*. De Gruyter, 2015.
- [KADK19] Wasiur R. KhudaBukhsh, Arnab Auddy, Yann Disser, and Heinz Koepl. Approximate lumpability for markovian agent-based models using local symmetries. *Journal of Applied Probability*, 56(3):647–671, 2019.
- [Kel79] Frank P. Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 1979.
- [KS60] John G. Kemeny and J. Laurie Snell. *Finite Markov Chains*. Van Nostrand, 1960.
- [Kun11] Kenneth Kunen. *Set Theory*. College Publications, 2011.
- [Lib04] Leonid Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [LS99] Joel L. Lebowitz and Herbert Spohn. A gallavotti–cohen type symmetry in the large deviation functional for stochastic dynamics. *Journal of Statistical Physics*, 95:333–365, 1999.
- [Mer11] Neri Merhav. Data processing theorems and the second law of thermodynamics. *IEEE Transactions on Information Theory*, 57(8):4926–4939, 2011.
- [ML98] Saunders Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer, 2 edition, 1998.

- [MS11] Dugald Macpherson and Charles Steinhorn. Definability in classes of finite structures. In *Finite and Algorithmic Model Theory*, pages 140–176. Cambridge University Press, 2011.
- [NJ10] Martin Nilsson Jacobi. A robust spectral method for finding lumpings and metastable states of non-reversible markov chains. *Electronic Transactions on Numerical Analysis*, 37:296–306, 2010.
- [Nor97] J. R. Norris. *Markov Chains*. Cambridge University Press, 1997.
- [Pol12] Matteo Poletti. *Geometric and Combinatorial Aspects of Nonequilibrium Statistical Mechanics*. Phd thesis, Università di Bologna, 2012.
- [PPRV10] A. Puglisi, S. Pigolotti, L. Rondoni, and A. Vulpiani. Entropy production and coarse-graining in markov processes. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(05):P05015, 2010.
- [RP81] L. C. G. Rogers and J. W. Pitman. Markov functions. *The Annals of Probability*, 9(4), 1981.
- [Sch76] Jürgen Schnakenberg. Network theory of microscopic and macroscopic behavior of master equation systems. *Reviews of Modern Physics*, 48:571–585, 1976.
- [Sei12] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012.
- [SN07] N. A. Sinitsyn and I. Nemenman. The berry phase and the pump flux in stochastic chemical kinetics. *Europhysics Letters*, 77(5):58001, 2007.
- [Wil72] J. C. Willems. Dissipative dynamical systems part i: General theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.