To Chart a Stone with Six Birds:
Emergence Phase Diagrams for Effective Theories in Control Space

Ioannis Tsiokos

March 15, 2026

Abstract

Many sciences rely on maps of macroscopic behavior—phase diagrams, stability charts, bifurcation sets, and feasibility regions—but emergence itself is rarely treated as a mapping problem. We introduce *emergence atlases*: subsets and scalar fields on control space indicating where a macro-description is feasible. In the finite-state setting used here, feasibility means that a packaged description is simultaneously closed, nontrivial, audit-safe, and budget-feasible. We compute such atlases in three settings: analytic law families, full packaging spaces of small systems, and the PICA mechanism algebra. The resulting maps recover exact closure loci and reversible lines, expose multi-objective tradeoffs among closure, directionality, effective information, novelty slack, and budget, and identify mechanism-space regimes together with their sensitive ablations and conservative compact supports. These computations support a simple reframing: emergence is usefully treated as a phase-diagram problem for effective theories in control space. In that view, effective information, irreversibility, coarse-graining quality, and mechanism-space analyses appear not as rival definitions of emergence, but as slices of a larger atlas.

1 Introduction

Maps of macroscopic behavior already organize large parts of science. Phase diagrams locate where phases exist, stability charts locate admissible controllers, bifurcation diagrams locate qualitative changes in dynamics, and renormalization-group flow diagrams locate effective theories and their basins [3, 4, 11, 20]. What these objects share is a common scientific task: not merely describing trajectories once a model is fixed, but charting where a higher-level description is available at all. Emergence research has often pursued closely related goals, but usually in fragmented form: one asks for a best coarse-graining, a best scale, a best macrovariable, or a best scalar score [8, 9, 12]. What is still largely missing is a unified cartography of where macro-laws are feasible.

The Six Birds program already supplies much of the conceptual machinery needed for such a cartography. At the foundational level, it treats packaged objects as fixed points of closure-like maps, distinguishes novelty from mere iteration, and organizes emergence around six primitives: rewrite, constraints, protocols, staging, packaging, and accounting. It also already provides a finite-state/no-smuggling semantics for packaged dynamics and a minimal stochastic mechanism algebra, PICA, that has been used to compute and audit geometry-like and time-asymmetry-like regimes. Those results establish how closure layers, audits, and mechanism choices arise; they do not yet provide the corresponding atlas of where a macro-theory exists in the first place [15–17].

This paper introduces that missing object. We treat emergence as a mapping problem on control space. A control point may include a microscopic law, a packaging or lens, a staging depth or timescale, a protocol or mechanism choice, and an explicit budget axis. An emergence atlas is then

Automorph Inc., Wilmington, DE, USA ioannis@automorph.io Zenodo DOI: 10.5281/zenodo.19038292

a subset or scalar field on that control space indicating where a macro-description is feasible. In the finite-state setting used here, feasibility means that a packaged description is simultaneously closed, nontrivial, audit-safe, and budget-feasible. The result is an object closer in spirit to a phase diagram than to a single emergence score.

This shift matters for two reasons. First, it makes room for multiple kinds of maps rather than a single global scalar: closure fields, directionality fields, novelty fields, effective-information slices, and budget-conditioned frontiers can all be read on the same underlying atlas. Second, it changes how familiar measures should be interpreted. Effective information, irreversibility, coarse-graining quality, and mechanism-space analyses are all useful, but none of them by itself is a complete theory of emergence. In the atlas view they become projections: informative slices of a larger control-space object. The paper's comparison layer is designed to make that reinterpretation explicit rather than rhetorical.

We compute these atlases in three settings. The first is analytic law space, where small benchmark families yield the cleanest visuals: exact closure loci, reversible lines, dense continuity diagnostics, and explicit $\tau > 1$ closure evidence. The second is packaging space, where all partitions of small finite systems can be scanned and compared across closure, nontriviality, directionality retention, novelty slack, effective information, and budget. The third is mechanism space, using PICA as a combinatorial algebra of closure interactions; here regime maps, leave-one-out ablations, compact-support candidates, and cost-quality frontiers become part of the same atlas language.

Several features of the resulting maps are worth emphasizing at the outset. In law space, exact closure loci behave like genuine geometric regions of parameter space rather than ad hoc score artifacts. In packaging space, low closure defect, high effective information, high novelty slack, and low budget do not collapse onto one privileged partition; they trade off against one another, and some of the highest-EI packagings remain trivial as macro-objects. In mechanism space, PICA configurations separate into null, time-asymmetry-like, geometry-compatible, and mixed regimes, while ablation rankings and compact-support frontiers expose which interactions matter most and which compact candidates remain only conservatively supported. The paper's claim surface is therefore intentionally controlled: where evidence is strong we say so; where it remains mixed we say so as well.

The six primitives enter this paper not as a new foundations derivation but as coordinates of the control problem. Rewrite (P1), constraints (P2), protocols (P3), staging (P4), packaging (P5), and accounting (P6) define the axes along which macro-law feasibility can change. Section 2 recaps them compactly and shows how the atlas object specializes them into computable finite-state subsets and fields. That move lets us state a single organizing thesis: a macro-law should be treated not as a single privileged coarse-graining or a single score, but as a region in control space where a macro-description is feasible.

The paper's contributions are as follows.

- We introduce emergence atlases as subsets and scalar fields on control space, treating macro-law feasibility as a mapping problem rather than as the search for a single privileged coarse-graining or a single scalar score.
- We give a finite-state computational specialization in which closure, nontriviality, audit safety, novelty slack, effective information, and budget appear as explicit slices of the same atlas object.
- We compute atlases in three settings—analytic law-space families, full packaging spaces of small systems, and the PICA mechanism algebra—and show that each supports qualitatively distinct but structurally comparable maps of macro-theory feasibility.
- We show, by explicit computed artifacts, that effective information, max-irreversibility style directionality measures, quantitative coarse-graining quality measures, and PICA mechanism-space analyses can be situated as slices of the atlas rather than as competing global definitions.

- We package the resulting claim surface conservatively through robustness analyses, settlement artifacts, a comparison-table layer, and a frozen evidence pack suitable for manuscript-driven use.

Roadmap. Section 2 recaps the Six Birds primitives and defines emergence atlases in control space. Section 3 gives the finite-state computational specialization. Sections 4, 5, and 6 compute atlases in law space, packaging space, and mechanism space, respectively. Section 7 shows how effective information, irreversibility, coarse-graining quality, and PICA become slices of the atlas. Section 8 states the robustness and settlement boundaries that control the claim surface, and Section 9 closes with the broader interpretation.

2 Emergence Atlases and the Six Primitives

This paper treats emergence as a cartographic problem. The aim is not to re-derive the Six Birds program, but to extract from it the object needed here: a map of where a macro-description is feasible. The fixed-point view of packaged objects, the separation between stability, novelty, and directionality, and the finite-state/no-smuggling setting used later in the paper are established elsewhere [16, 17]. What matters in the present section is that those ingredients can be re-read as coordinates of a control problem. Instead of asking only whether a macro-description exists, we ask where in control space it is feasible. The accounting vocabulary used below is likewise inherited from prior Six Birds work, but here it will be used operationally through explicit cost or budget proxies rather than through a full duality theory [19].

The six primitives. The Six Birds program organizes emergence around six primitive operations:

- **P1 Rewrite**: changing the effective update law.
- **P2 Constraints**: restricting admissible states, transitions, or supports.
- **P3 Protocols**: introducing route dependence, ordering, or holonomy.
- **P4 Staging**: changing timescale, iteration depth, or resolution ladder.
- **P5 Packaging**: quotienting or coarse-graining into candidate macro-objects.
- **P6 Accounting**: attaching costs, conserved quantities, or audit variables.

These primitives are not six separate definitions of emergence. They are the coordinates along which macro-law feasibility can change. P1, P3, and P5 move one through law space, protocol space, and packaging space; P2 and P6 carve admissible and budgeted regions within those spaces; and P4 supplies the staging parameter that turns closure and stability into scale-dependent questions.

Definition 2.1 (Control space and admissible macro-descriptions)**.** A *control space* $\mathcal{U}$ is any set of variables over which macro-law feasibility is to be charted. A control point $u \in \mathcal{U}$ may include a microscopic law parameter, a packaging or lens, a staging depth τ, a protocol or mechanism variable, and a budget coordinate. Different atlases in this paper use different choices of $\mathcal{U}$: law-space atlases vary microscopic law parameters while holding packaging fixed, packaging-space atlases vary the packaging while holding the microscopic law fixed, and mechanism-space atlases vary protocol/mechanism choices over a fixed finite-state substrate.

For a given control point u, let $x_t \in X$ denote the microscopic state and let

$$x_{t+1} \sim K_u(\cdot \mid x_t)$$

denote the corresponding microscopic update law. Let $\tau(u) \in \mathbb{N}$ be the staging depth at which the macro-description is assessed. An *admissible macro-description* at u is a pair $(\Pi, L) \in \mathfrak{C}(u)$, where

$\Pi : X \to Y$ is a packaging map and L is a candidate macro update rule on Y. The family $\mathfrak{C}(u)$ collects whichever package/law pairs are allowed at that control point.

Closure is measured at the chosen stage by

$$\mathrm{CE}(u;\Pi,L) = \mathbb{E}\Big[d\Big(\Pi(x_{t+\tau(u)}),\ L(\Pi(x_t))\Big)\Big]\,,$$

where d is the discrepancy notion appropriate to the packaged state space. Deterministic dynamics are the special case in which K_u is supported on a single successor for each state. Small closure error, however, is not enough by itself. A feasible macro-description must also be nontrivial, audit-safe, and budget-feasible. We therefore write

$$\mathrm{NT}(u;\Pi,L), \qquad \mathrm{AS}(u;\Pi,L), \qquad \mathrm{Cost}(u;\Pi,L)$$

for a nontriviality predicate, an audit-safety predicate, and the relevant cost functional, respectively. In the finite-state sections below, these will be instantiated concretely by stable-block and distinct-row diagnostics, directionality and data-processing-style audit checks, novelty slack, effective-information slices, and explicit within-domain budget proxies.

Definition 2.2 (Emergence atlas)**.** Fix a scientific predicate Q expressing the additional property one wishes to chart, such as boundedness, closure, geometry-compatibility, sustained directionality, or some other macro-level property. The *emergence atlas* associated with Q at tolerance $\delta \geq 0$ is

$$\mathcal{A}_Q(\delta) = \Big\{u \in \mathcal{U}\ :\ \exists(\Pi,L) \in \mathfrak{C}(u) \text{ such that } \begin{array}{l}\mathrm{CE}(u;\Pi,L) \leq \delta,\ \mathrm{NT}(u;\Pi,L),\\ \mathrm{AS}(u;\Pi,L),\ \mathrm{Cost}(u;\Pi,L) \leq B(u),\ Q(u;\Pi,L)\end{array}\Big\},$$

where $B(u)$ is the budget coordinate or budget bound carried by the control point.

The atlas definition is intentionally broad. In some applications one scans only law parameters, in others only packagings, and in others only mechanism choices. The same formal object covers all three cases. A control point belongs to the atlas only when a packaged description meeting all four feasibility requirements exists there. The point of the atlas is therefore not to identify a single privileged macro-description, but to display the region in control space where such descriptions are available at all.

Definition 2.3 (Atlas fields)**.** Besides feasible-set atlases, one may chart scalar fields over the same control space. Two generic examples are the *closure field*

$$\mathrm{CE}^*_Q(u) = \inf\Big\{\mathrm{CE}(u;\Pi,L) : \begin{array}{l}(\Pi,L) \in \mathfrak{C}(u),\ \mathrm{NT}(u;\Pi,L),\ \mathrm{AS}(u;\Pi,L),\\ \mathrm{Cost}(u;\Pi,L) \leq B(u),\ Q(u;\Pi,L)\end{array}\Big\},$$

and the *budget frontier*

$$B^*_Q(u) = \inf\Big\{\mathrm{Cost}(u;\Pi,L) : \begin{array}{l}(\Pi,L) \in \mathfrak{C}(u),\ \mathrm{CE}(u;\Pi,L) \leq \delta,\ \mathrm{NT}(u;\Pi,L),\\ \mathrm{AS}(u;\Pi,L),\ Q(u;\Pi,L)\end{array}\Big\}.$$

The paper will also use novelty, directionality, effective-information, and budget fields as additional slices of the same atlas object.

The atlas view changes the scientific question. Instead of asking for a best partition, best score, or best mechanism in isolation, one asks how several fields and feasibility regions coexist over the same control space. High effective information may coexist with trivial objecthood, low closure defect may coexist with low novelty slack, and strong directionality may occur in regions that are poor packagings. Those possibilities are not anomalies in the atlas view; they are precisely what the atlas is designed to display.

The rest of the paper specializes this template to finite-state systems. Section 3 makes closure, nontriviality, audit safety, novelty, effective information, and budget computable. Sections 4, 5, and 6 then chart these quantities in law space, packaging space, and mechanism space, respectively.

3 Finite-State Atlas Construction

This section specializes the general atlas template of Section 2 to finite-state systems, where all of the slices used later in the paper can be computed explicitly. The point is not to claim that finite-state models exhaust emergence, but to give a common operational substrate on which law-space, packaging-space, and mechanism-space atlases can be compared directly. In this setting, packaged dynamics are explicit, audit quantities can be computed exactly on small systems, and the distinction between closure, nontriviality, directionality, novelty, effective information, and budget can be maintained without collapsing them into a single score [2, 7, 10, 16, 17].

Microscopic substrate and packaging. Let $X = \{1, \dots, N\}$ be a finite microscopic state space and let

$$P \in [0,1]^{N\times N}, \qquad \sum_{j=1}^{N} P_{ij} = 1,$$

be the row-stochastic microscopic update kernel. A packaging is represented by a partition

$$\Pi = \{B_1, \dots, B_K\}$$

of X into nonempty disjoint blocks, with associated macro state space

$$Y = \{1, \dots, K\}.$$

The canonical projection

$$\pi_\Pi : X \to Y$$

sends each microscopic state to the index of its block. For a microscopic distribution $\mu \in \Delta(X)$, the induced macro distribution is the pushforward

$$(\pi_\Pi)_* \mu(a) = \sum_{i \in B_a} \mu(i).$$

To construct a macro update rule from P, one must specify how a macro basis state is lifted back into microscopic space. We therefore fix, for each block B_a, a block-supported probability vector

$$u_a \in \Delta(X), \qquad \mathrm{supp}(u_a) \subseteq B_a.$$

Two lift choices are used throughout the experiments:
- the *uniform lift*, where u_a is uniform on B_a;
- the *prototype lift*, where u_a is a chosen block-supported prototype, in the experiments usually the one-hot state supported on the first state of the canonical block ordering.

Writing e_a for the ath basis distribution on Y, the induced staged macro-kernel at depth τ is

$$M^{(\tau)}_{\Pi,\mathbf{u}}(a,b) = \sum_{j\in B_b} (u_a P^\tau)_j. \tag{1}$$

Equivalently, one lifts e_a to u_a, evolves by P^τ, and pushes forward by π_Π.

Definition 3.1 (Packaged macro-kernel)**.** Given a finite microscopic kernel P, a partition $\Pi = \{B_1,\ldots,B_K\}$ of X, a block-supported lift family $\mathbf{u} = (u_1,\ldots,u_K)$, and a stage $\tau \in \mathbb{N}$, the *packaged macro-kernel* is the row-stochastic kernel $M^{(\tau)}_{\Pi,\mathbf{u}}$ defined by (1).

The finite-state atlas slices below are all computed from the tuple

$$(P, \Pi, \mathbf{u}, \tau),$$

possibly together with a choice of initial law or a budget proxy. The later sections vary different coordinates of that tuple: law-space atlases vary P, packaging-space atlases vary Π, and mechanism-space atlases vary structured families of microscopic laws while inheriting the same finite-state audit logic.

Closure and nontriviality slices. The first slice of the atlas is closure quality. For a fixed partition Π and stage τ, exact strong lumpability requires that the block-level future seen from any two microscopic states in the same source block agree exactly [2, 10]. This motivates the strong-lumpability defect

$$\mathrm{LD}^{(\tau)}_{\Pi}(P) = \max_{a,b} \max_{i,j\in B_a} \left| \sum_{m\in B_b} (P^\tau_{im} - P^\tau_{jm}) \right|. \tag{2}$$

Exact closure at that stage is the condition $\mathrm{LD}^{(\tau)}_{\Pi}(P) = 0$.

A second closure slice compares each microscopic state's projected future to the row of the packaged macro-kernel attached to its block:

$$\mathrm{RM}^{(\tau)}_{\Pi,\mathbf{u}}(a) = \frac{1}{|B_a|} \sum_{i\in B_a} \left\| (\pi_\Pi)_*(\delta_i P^\tau) - M^{(\tau)}_{\Pi,\mathbf{u}}(a,\cdot) \right\|_1, \tag{3}$$

with global row mismatch

$$\mathrm{RM}^{(\tau)}_{\Pi,\mathbf{u}} = \max_a \mathrm{RM}^{(\tau)}_{\Pi,\mathbf{u}}(a)$$

or, where noted, the block average. This detects within-block heterogeneity even when the macro-kernel itself is well defined.

A third closure-related quantity is the retention error of the packaged basis states:

$$\mathrm{Ret}^{(\tau)}_{\Pi,\mathbf{u}}(a) = \mathrm{TV}\Big(e_a M^{(\tau)}_{\Pi,\mathbf{u}}, e_a\Big), \qquad \mathrm{Ret}^{(\tau)}_{\Pi,\mathbf{u}} = \max_a \mathrm{Ret}^{(\tau)}_{\Pi,\mathbf{u}}(a), \tag{4}$$

where $\mathrm{TV}(p,q) = \frac{1}{2}\|p-q\|_1$. This measures whether the packaged basis states are metastable under the induced macro dynamics.

Finally, following the fixed-point packaging intuition of the earlier Six Birds work, we use an empirical idempotence defect. Let

$$E^{(\tau)}_{\Pi,\mathbf{u}}(\mu) = \sum_{a=1}^{K} \Big((\pi_\Pi)_*(\mu P^\tau)\Big)(a)\, u_a$$

be the lift-pushforward endomap on microscopic distributions. The idempotence defect is computed over microscopic basis states by

$$\mathrm{ID}^{(\tau)}_{\Pi,\mathbf{u}} = \max_{i\in X} \mathrm{TV}\Big(E^{(\tau)}_{\Pi,\mathbf{u}}\Big(E^{(\tau)}_{\Pi,\mathbf{u}}(\delta_i)\Big)\,, E^{(\tau)}_{\Pi,\mathbf{u}}(\delta_i)\Big)\,. \tag{5}$$

Closure alone is not enough: a constant or rank-one macro map can be perfectly closed and still scientifically trivial. The finite-state nontriviality slices therefore remain separate. We use the stable-block count

$$\mathrm{SB}^{(\tau)}_{\Pi,\mathbf{u}} = \#\left\{a \in Y : \mathrm{Ret}^{(\tau)}_{\Pi,\mathbf{u}}(a) \le \varepsilon_{\mathrm{stab}}\right\}$$

and the distinct-row count

$$\mathrm{DR}^{(\tau)}_{\Pi,\mathbf{u}} = \#\{\text{row classes of } M^{(\tau)}_{\Pi,\mathbf{u}} \text{ up to tolerance } \varepsilon_{\mathrm{row}}\}.$$

The paper's default objecthood indicator is then

$$\mathrm{Obj}^{(\tau)}_{\Pi,\mathbf{u}} = \mathbf{1}\Big[\mathrm{SB}^{(\tau)}_{\Pi,\mathbf{u}} \ge 2 \ \wedge\ \mathrm{DR}^{(\tau)}_{\Pi,\mathbf{u}} \ge 2\Big]\,. \tag{6}$$

The later packaging-space sections rely heavily on the fact that low closure defect, high effective information, and nontrivial objecthood need not coincide on the same partition.

Directionality slices. The atlas also tracks directionality in a way that distinguishes true arrow-of-time information from artifacts of fitting or observation [5, 13]. For a finite horizon T and initial microscopic law $\rho \in \Delta(X)$, let $\mathbb{P}_{\rho,T}$ be the path law of the microscopic chain and let $\mathcal{R}_*\mathbb{P}_{\rho,T}$ denote its time-reversed image. The finite-horizon path-space asymmetry is

$$\Sigma_T(\rho; P) = D_{\mathrm{KL}}(\mathbb{P}_{\rho,T} \,\|\, \mathcal{R}_*\mathbb{P}_{\rho,T})\,. \tag{7}$$

For deterministic observation by π_Π, the observed path-space asymmetry is

$$\overline{\Sigma}_T(\rho; P, \Pi) = D_{\mathrm{KL}}\Big((\pi_\Pi^{T+1})_*\mathbb{P}_{\rho,T} \,\Big\|\, \mathcal{R}_*(\pi_\Pi^{T+1})_*\mathbb{P}_{\rho,T}\Big)\,. \tag{8}$$

When $\Sigma_T(\rho; P) > 0$, we summarize directionality retention by

$$\mathrm{DirRet}_T(\rho; P, \Pi) = \frac{\overline{\Sigma}_T(\rho; P, \Pi)}{\Sigma_T(\rho; P)}, \tag{9}$$

and leave it undefined when the microscopic asymmetry is itself zero. This is the finite-state form in which the no-fake-arrows constraint later appears: observation may hide directionality, but it should not create it.

For stationary packaged dynamics, we also use the entropy-production-style proxy

$$\mathrm{EP}(M, \pi) = \sum_{a,b} \pi_a M_{ab} \log \frac{\pi_a M_{ab}}{\pi_b M_{ba}}, \tag{10}$$

where π is a stationary law for the packaged kernel M. In the empirical sections this is used as a summary slice rather than as a foundational definition of time asymmetry.

Novelty, effective information, and budget slices. The remaining slices quantify aspects of the atlas that are neither closure nor directionality. Novelty is tracked through hidden slack in the packaging. For a finite partition of an N-state system into K blocks, define

$$\mathrm{HV}(\Pi) = N - K, \tag{11}$$

the hidden-volume slack, together with the finite-forcing-inspired strict-extension probability

$$\mathrm{SEP}(\Pi) = \begin{cases} 1 - 2^{-(N-K)}, & N > K, \\ 0, & N \leq K. \end{cases} \tag{12}$$

These are not realized novelty measures; they are novelty-potential slices that track how much internal freedom remains beneath the packaging.

Effective information is treated as another slice, not as a universal objective [8, 9, 12, 14]. For the packaged macro-kernel $M = M^{(\tau)}_{\Pi,\mathbf{u}}$ on K macro states, let

$$q = \frac{1}{K} \sum_{a=1}^{K} M(a, \cdot)$$

be the average macro effect distribution under a uniform intervention on macro states. The macro effective information in nats is

$$\mathrm{EI}_{\mathrm{macro}}(M) = \frac{1}{K} \sum_{a=1}^{K} D_{\mathrm{KL}}(M(a, \cdot) \,\|\, q)\,, \tag{13}$$

with normalized form

$$\mathrm{EI}^{\mathrm{norm}}_{\mathrm{macro}}(M) = \begin{cases} \mathrm{EI}_{\mathrm{macro}}(M) / \log K, & K > 1, \\ 0, & K = 1. \end{cases} \tag{14}$$

To compare this with the information already visible before macro-basis packaging, let

$$b_i = (\pi_\Pi)_*(\delta_i P^\tau), \qquad \bar{b} = \frac{1}{N} \sum_{i=1}^{N} b_i.$$

Then the projected-micro effective information is

$$\mathrm{EI}_{\mathrm{proj}}(P, \Pi, \tau) = \frac{1}{N} \sum_{i=1}^{N} D_{\mathrm{KL}}\Big(b_i \,\Big\|\, \bar{b}\Big)\,. \tag{15}$$

The experiments use both the EI gap

$$\mathrm{EIgap} = \mathrm{EI}_{\mathrm{proj}} - \mathrm{EI}_{\mathrm{macro}}$$

and the EI retention ratio where the denominator is positive. These will matter in packaging space, where some of the highest-EI partitions remain trivial as macro-objects.

Finally, the budget axis is instantiated conservatively through explicit within-domain proxies rather than through a universal unit of cost. For analytic and packaging rows, the default structural budget is

$$\mathrm{Bud}_{\mathrm{pack}} = K + c_{\mathrm{lift}} + \tau, \tag{16}$$

where $c_{\text{lift}} = 1$ for the uniform lift and $c_{\text{lift}} = K$ for the prototype lift. In PICA, later sections use enabled-cell counts, gated-edge counts, and the saved mechanism budget proxy when available. The role of budget in this paper is therefore comparative and within-domain: it marks cheap versus expensive regions of the atlas, not a universal inter-domain currency.

The finite-state atlas construction is now explicit. Given a microscopic kernel, a partition, a lift family, and a stage, one can compute closure slices, directionality slices, novelty slices, effective-information slices, and budget slices on the same control point. The next three sections do exactly that in three different control spaces: law space, packaging space, and mechanism space.

4 Law-Space Atlases

Law-space atlases are the clearest place to see the paper's central claim in a familiar scientific form. Here the packaging is held fixed and one varies the microscopic law parameters instead, asking where closure, directionality, and related slices become available. In the finite-state setting, this produces objects that are recognizably phase-diagram-like: exact closure loci, reversible lines, dense scalar fields, and threshold boundaries. The point of this section is not that law space is the only control space that matters, but that it provides the cleanest visual demonstration that macro-law feasibility can be charted rather than reduced to a single emergence score [16, 17].

Exact closure loci in the block family. The block-metastable-leak family is the paper's cleanest closure benchmark. Its default partition is $0, 1 \mid 2, 3$, and the parameter δ controls whether the two microscopic states inside each source block agree or disagree in their aggregated block-level futures. At $\delta = 0$, the family lies on an exact closure locus for the packaged dynamics at $\tau = 1$; away from that line, the strong-lumpability defect becomes positive. Figure 1 makes that structure visible on a dense grid: the exact-locus line appears as a narrow zero-defect corridor rather than as a post hoc threshold artifact, and off-locus deformation is smooth enough to support a genuine phase-diagram reading.

That interpretation remains valid only if the line is not an artifact of coarse sampling. The dense closure atlas matters because it resolves the locus at substantially higher parameter resolution than the earlier exploratory sweeps, and because the closure field is computed from the explicit packaged finite-state quantities of Section 3. The resulting map is therefore not a stylized schematic: it is the closure field induced by the actual finite-state packaged kernel under the chosen partition, lift, and stage.

Tau dependence and strengthened closure support. The exact-locus story is strongest at $\tau = 1$, where exact strong lumpability is part of the family design. What matters for the paper, however, is whether the closure-locus interpretation survives beyond a single stage. Figure 2 therefore compares the same block family at $\tau = 1$ and $\tau = 2$. The comparison shows near-identical locus behavior: the exact line remains visually stable, while off-locus points remain clearly nonzero. The dedicated $\tau > 1$ closure artifact sharpened this further by recording exact zero defect along the locus at $\tau = 1$ and numerically negligible defect along the same line at $\tau = 2$, with nearby off-locus rows staying bounded away from zero.

We treat this as strengthened empirical support rather than as a formal theorem about all powers of the kernel family. In other words, the paper is entitled to say that the closure-locus story survives the first nontrivial staging extension tested here, but not that a full polynomial or semialgebraic characterization of the $\tau > 1$ locus has been proven. That distinction is part of the paper's general claim discipline.

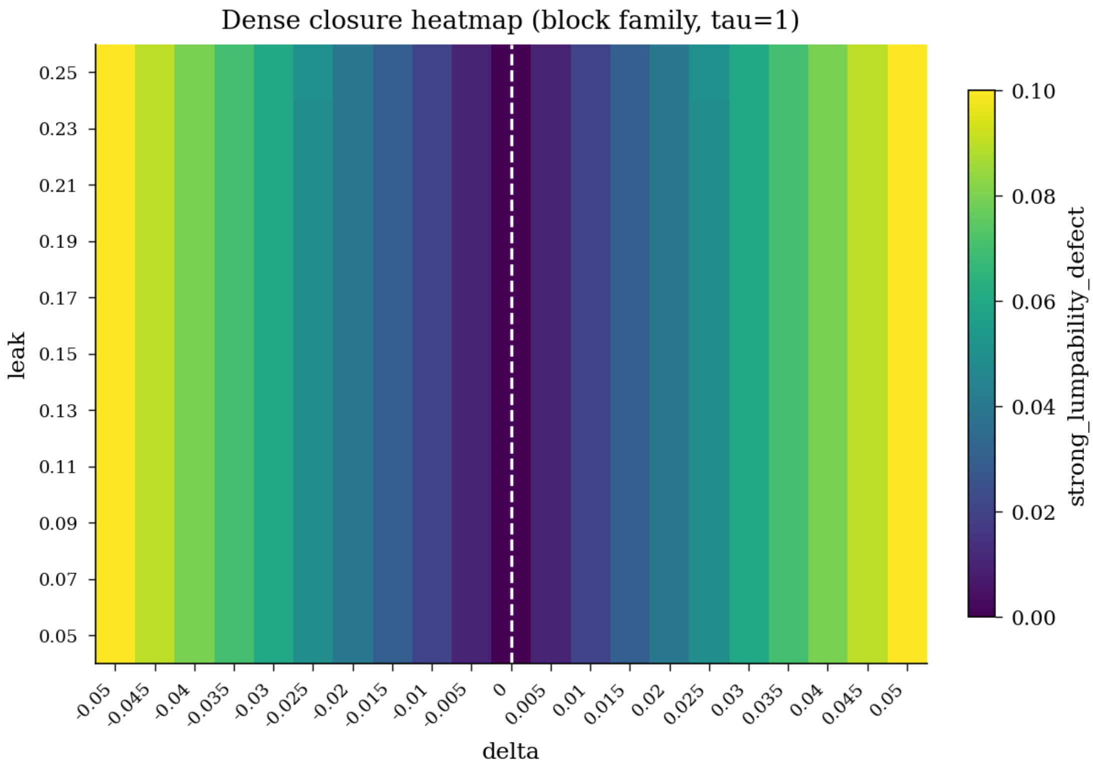

Figure 1: Dense closure atlas for the block-metastable-leak family at the default packaging $0, 1 \mid 2, 3$. The line $\delta = 0$ appears as an exact closure locus, while off-locus parameter values acquire positive strong-lumpability defect. The figure is read as a law-space feasibility map rather than as a single-score optimization surface.

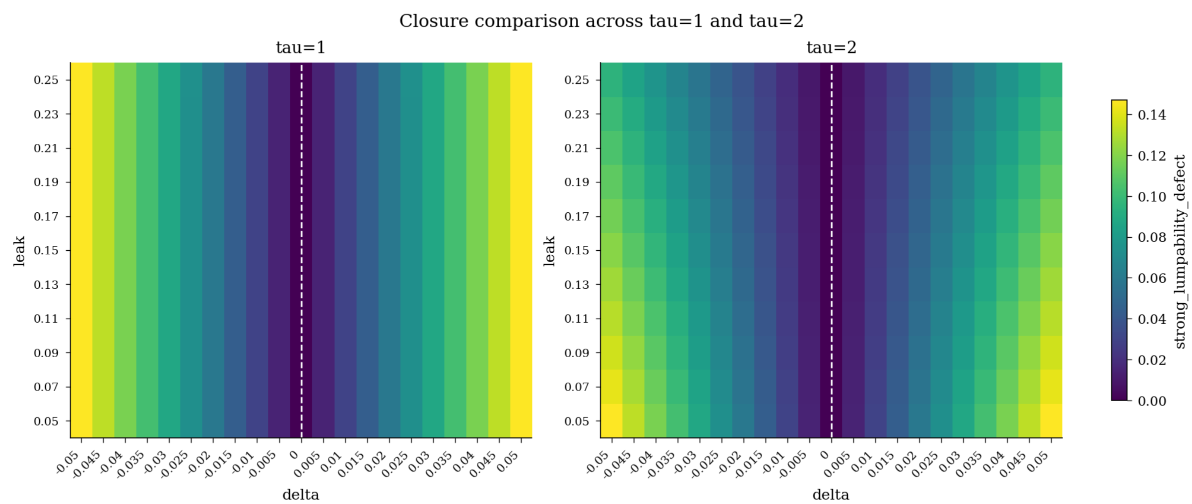

Figure 2: Matched $\tau = 1$ and $\tau = 2$ closure atlases for the block-metastable-leak family. The exact-locus line remains visually stable, and the dedicated dense artifact records numerically negligible defect on the same line at $\tau = 2$. This is explicit empirical support for closure-locus persistence beyond the one-stage case, not a general symbolic proof.

Reversible lines and directionality fields. The driven-reversible-block family provides the cleanest directionality axis. Here the parameter 'drive' controls the bias between clockwise and counterclockwise motion, while the self-loop parameter controls the amount of inertial retention. At drive $= 0$, the family sits on a reversible line; away from that line, path-space asymmetry becomes

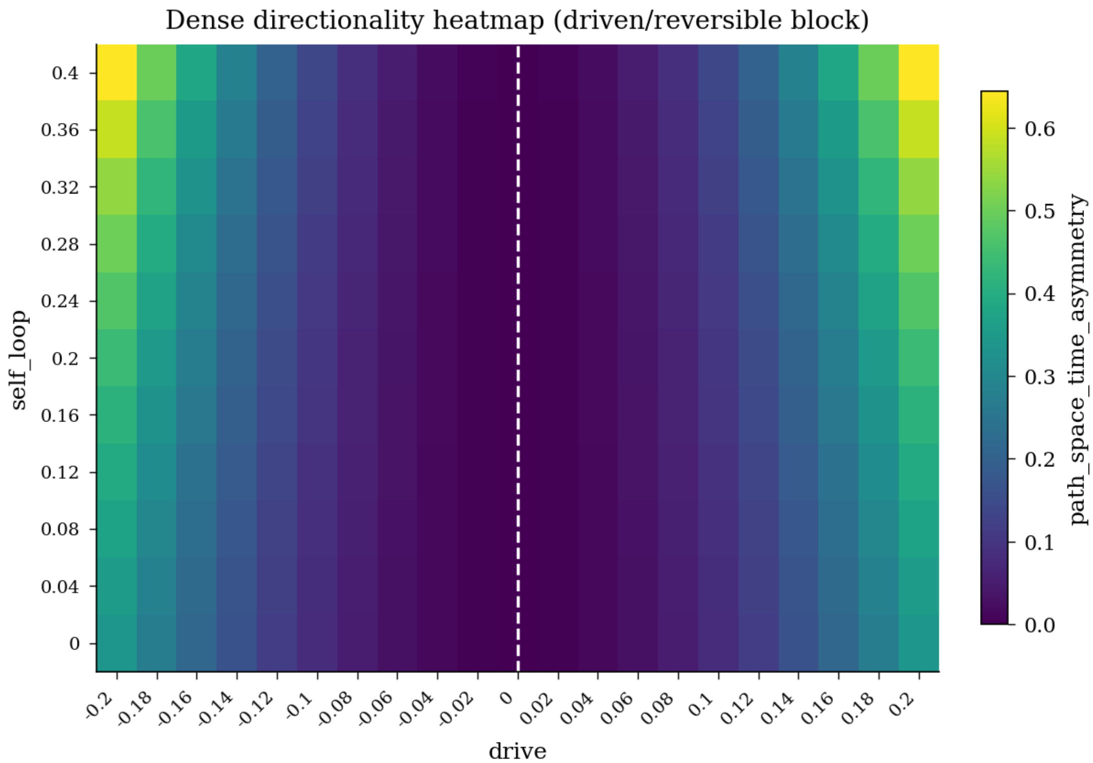

Figure 3: Dense directionality atlas for the driven-reversible-block family. The reversible line at drive ≈ 0 is cleanly visible, while finite-horizon path-space asymmetry grows away from that line across the self-loop axis. The map shows where directionality is available in law space rather than reducing it to a single global nonequilibrium score.

positive. Figure 3 shows the resulting dense directionality field. The reversible line is not inferred from a fitted model after coarse-graining; it is visible directly in the exact finite-horizon asymmetry calculation, with the signless magnitude increasing away from the line as the system is driven out of equilibrium.

The biased ring family, although smaller, corroborates the same qualitative pattern in a second law family: symmetric bias produces a zero-asymmetry boundary, while biased motion produces positive asymmetry. The point is not that every law-space atlas must have a single clean line, but that law space can contain structural subsets on which directionality vanishes or appears in an interpretable way, just as closure vanishes or appears along interpretable subsets in the block family.

Explicit empirical continuity support. One of the paper's intended theorem-support statements is that approximate closure fields and related atlas fields behave continuously enough to justify phase-diagram reasoning. Earlier low-resolution figures made this plausible, but plausibility is weaker than an explicit diagnostic. The dense atlas pipeline therefore computes adjacent-grid jump statistics, local Lipschitz surrogates, overlap consistency between coarse and dense sweeps on shared parameter points, and threshold-stability summaries for the key closure and directionality regions. Figure 4 condenses those diagnostics.

The resulting support is empirical, not formal. What it gives is stronger than visual smoothness and weaker than a general continuity theorem. In the tested analytic families, adjacent-grid variation remains controlled, dense/coarse overlap discrepancies vanish on shared grid points, and thresholded region comparisons are fully stable under the designated diagnostics. This is enough to justify the

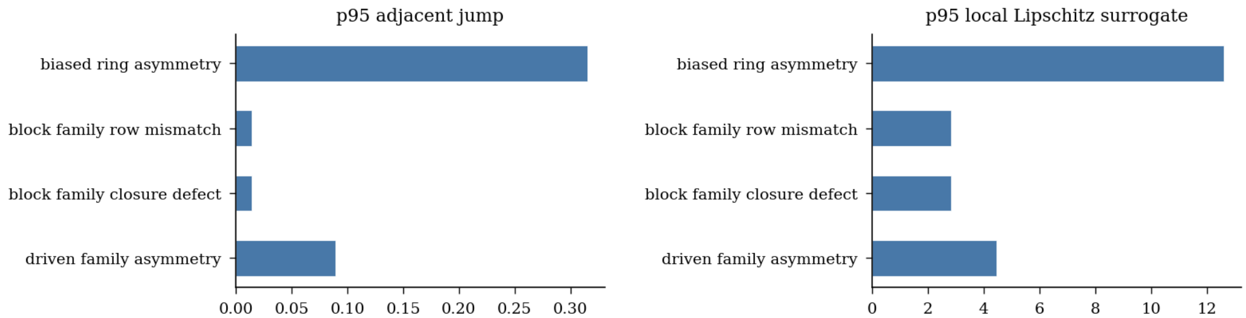

Figure 4: Empirical continuity diagnostics for the dense analytic atlases. *Left:* p95 adjacent-grid jump. *Right:* p95 local Lipschitz surrogate. Bars correspond to: biased ring (path-space asymmetry), block family (fiber-row mismatch and strong-lumpability defect), and driven family (path-space asymmetry). A third diagnostic (dense-vs-coarse overlap discrepancy) is omitted because all values are near-zero by design, confirming grid consistency.

paper's cartographic language: the law-space fields are not merely pictures that look smooth, but audited dense fields whose variation has been measured.

Taken together, the law-space atlases show that emergence can be charted in the same scientific spirit as ordinary phase diagrams. Exact closure loci and reversible lines appear as structural subsets of parameter space; dense continuity diagnostics support the idea that the corresponding fields vary in an orderly way; and the first nontrivial staging extension provides explicit empirical support that the closure-locus story is not confined to $\tau = 1$. What this section does not yet do is vary the packaging itself. That is the task of the next section, where the atlas becomes explicitly multi-objective and different partitions trade closure, objecthood, directionality retention, novelty, effective information, and budget against one another.

5 Packaging-Space Atlases

If law-space atlases hold the packaging fixed and vary the microscopic law, packaging-space atlases do the opposite: they hold the microscopic substrate fixed and vary the packaging itself. This is the point in the paper where the atlas perspective becomes unmistakably multi-objective. A packaging is no longer a hidden modeling choice in the background; it is itself the control variable. Once that shift is made, the scientific question changes immediately. One no longer asks for a single best partition. One asks how closure, nontriviality, directionality retention, novelty slack, effective information, and budget coexist over the space of candidate packagings [6, 7, 16, 17].

That change matters because the different slices do not collapse onto one another. Some partitions are nearly perfectly closed but scientifically trivial. Some have large effective information but poor closure. Some retain directionality well but fail to sustain nontrivial macro-objects. Some are cheap but weak, while others are richer and more expensive. The packaging-space atlas is therefore not an optimization problem with one universal winner. It is a structured tradeoff landscape.

Exact-locus packaging space. The cleanest place to see this is the exact-locus block family. Here the microscopic law sits on the law-space closure line, and the packaging is varied over all fifteen partitions of the four-state system. Figure 5 shows the resulting partition-family map. Even on the exact closure locus, packaging space does not collapse to a single privileged macro-description. Some partitions lie at exact or near-exact closure while remaining trivial as macro-objects; others retain exact closure while also supporting nontrivial objecthood. This immediately demonstrates why closure alone cannot be the whole story.

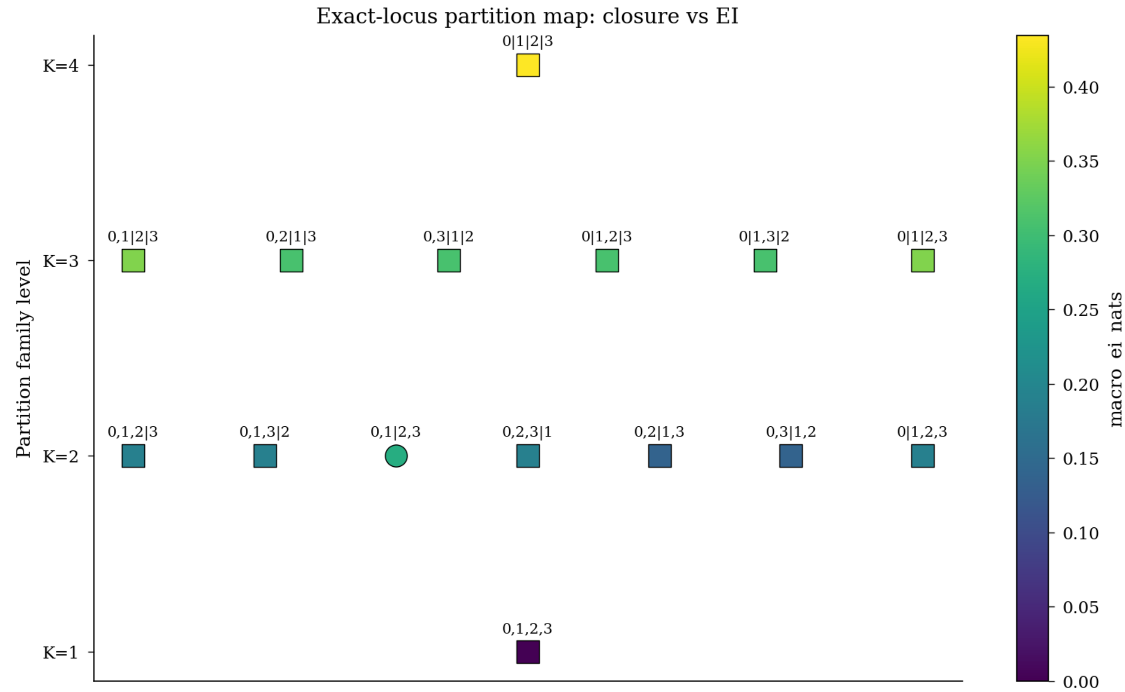

Figure 5: Packaging-space map for the exact-locus block family, arranged by partition family over all fifteen partitions of the four-state system. Exact closure can coexist with both trivial and nontrivial macro-objects, and high effective information need not coincide with the most scientifically useful packaging. The figure shows that closure, EI, and objecthood remain distinct slices even on the exact law-space locus.

The exact-locus map also makes the effective-information point visible. The finest partition can carry the highest EI without thereby becoming the best macro-object in the paper's sense. Conversely, a coarser partition such as $0,1 \mid 2,3$ can remain nontrivial and exactly closed while carrying less EI than a finer but scientifically trivial alternative. This is precisely the kind of relationship the atlas is meant to display: effective information is informative, but it is only one slice of packaging quality, not a universal surrogate for closure or objecthood.

Off-locus multi-objective tradeoffs. Once the same family is moved off the exact law-space locus, packaging-space tradeoffs become sharper. Figure 6 shows the off-locus frontier and full partition cloud. The low-budget extreme remains the indiscriminate one-block partition, which is cheap but trivial. Nearby in cost are partitions that improve effective information or novelty slack while still failing the paper's objecthood criterion. The partition $0,1 \mid 2,3$ is more expensive than the trivial minimum but supports nontrivial objecthood and positive EI, even though its closure defect is no longer globally minimal. In other words, once the law is no longer exactly aligned with the default packaging, one does not obtain a single best partition; one obtains a frontier.

This is where the atlas perspective is strongest. The off-locus packaging field contains a low-defect/trivial corner, a nontrivial/moderate-EI corner, and higher-EI or higher-novelty alternatives that incur additional closure or budget cost. No scalar score eliminates these tradeoffs. The frontier should therefore be read the same way one reads a Pareto front in control or design problems: as a family of different feasible compromises among objectives that matter simultaneously.

Directionality-retaining packagings. Packaging space also changes how directionality is seen. For driven microscopic laws, one can ask not only whether the microscopic process is time-asymmetric,

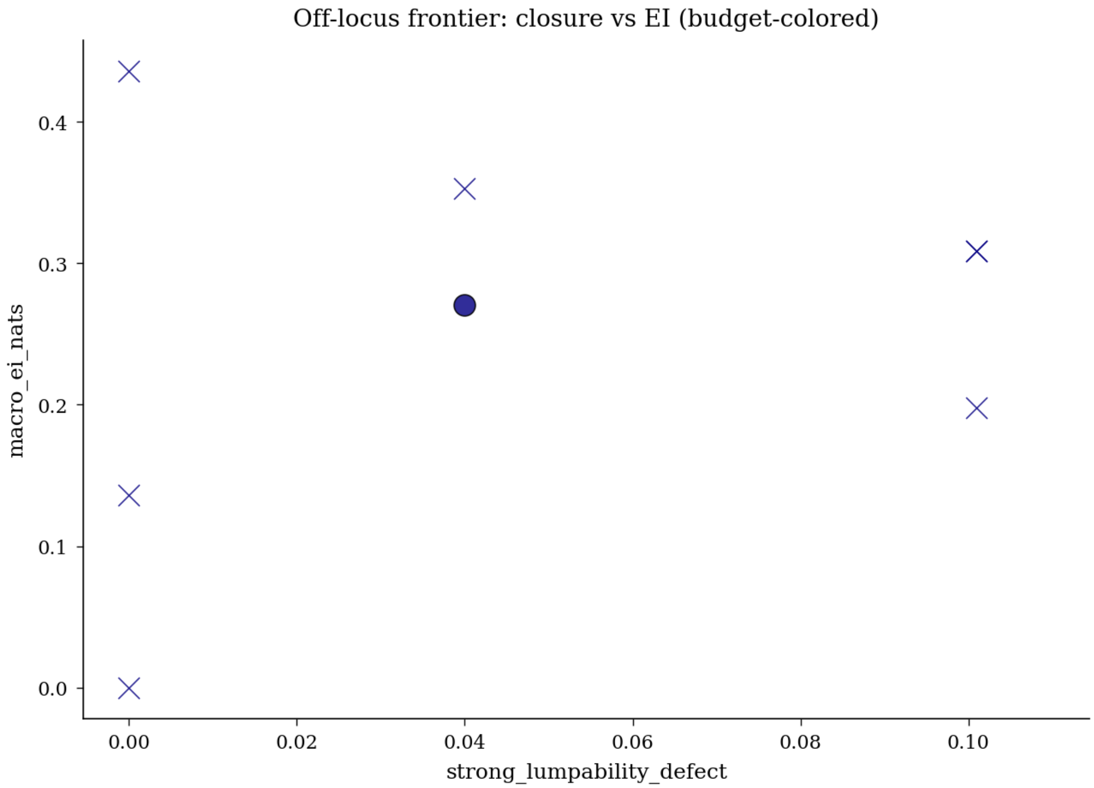

Figure 6: Packaging frontier for the off-locus block family. The full partition cloud and highlighted frontier expose tradeoffs among closure defect, effective information, objecthood, novelty slack, and budget rather than selecting a single universally preferred partition. Low-budget partitions can remain trivial, while nontrivial partitions typically occupy more balanced but more expensive frontier positions.

but which packagings retain that asymmetry most effectively. Figure 7 compares a driven condition with its reversible counterpart. In the driven panel, directionality retention becomes an explicit packaging-space axis. Some packagings preserve the microscopic asymmetry strongly, but may remain trivial or closure-poor; others improve closure or EI at the cost of attenuating the arrow-of-time signal. In the reversible panel, by contrast, the directionality axis is correctly absent as a meaningful discriminator because the underlying microscopic asymmetry is itself degenerate.

This comparison matters because it shows that directionality is not just a property of the law. It is also a property of the packaging, in the sense that different packagings can preserve or erase the same microscopic asymmetry to different degrees. But again, that does not make directionality retention the sole objective. The reversible panel is especially useful here: once the microscopic asymmetry vanishes, the packaging tradeoffs remain real, but they are carried by closure, EI, novelty, and budget rather than by arrow-of-time retention.

Effective information as a slice, not a verdict. The most direct packaging-space support for the paper's comparison thesis is the EI slice itself. Figure 8 plots EI against closure while encoding objecthood and budget. What it shows is exactly the relationship the atlas view predicts: EI is informative but not dispositive. Some of the highest-EI packagings are trivial as macro-objects; some of the most scientifically useful packagings have only moderate EI but much better closure or objecthood properties. The right interpretation is not that EI is wrong. It is that EI is one coordinate projection of packaging quality.

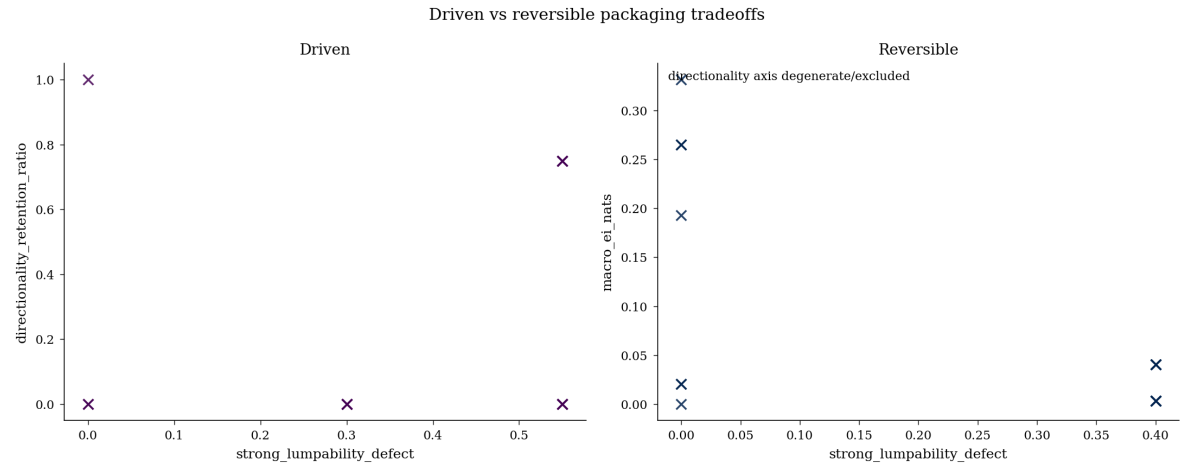

Figure 7: Driven versus reversible packaging-space tradeoffs. In the driven condition, packagings differ in how much microscopic directionality they retain; in the reversible condition, that axis becomes degenerate and the remaining tradeoffs are carried by closure, EI, novelty, and budget. The contrast shows that directionality retention is itself a packaging-dependent slice rather than a law-only property.

That point matters because it lets the paper treat causal-emergence-style quantities constructively rather than polemically [8, 9, 12]. The atlas does not dismiss effective information. It computes it, places it on the same packaging-space rows as closure, objecthood, novelty, directionality retention, and budget, and then shows exactly where it aligns with those quantities and where it does not. The result is a more informative scientific object than either a closure-only ranking or an EI-only ranking.

Refinement is non-monotone. The packaging atlas also makes clear why refinement should not be narrated by a single monotonic slogan. The settled refinement analysis now supports both help and hurt examples under explicit conventions and tolerances. In some regions of packaging space, refinement reveals distinctions that improve closure or preserve objecthood; in others, refinement exposes internal mismatch that worsens closure or destroys the macro-object. The correct general statement is therefore not that refinement helps, nor that refinement hurts, but that refinement is non-monotone. That non-monotonicity is not an accidental pathology. It is a structural feature of packaging space once closure, objecthood, and retention are all tracked simultaneously.

The paper's claim surface here is intentionally careful. It does not state a universal theorem that refinement typically improves description. What it states, and what the settled artifacts support, is that both directions occur on explicit finite-state examples under declared conventions. The packaging atlas is exactly the right object for that fact, because it treats refinement not as a moral preference for finer scales, but as motion inside a control space where different slices may improve or worsen together or separately.

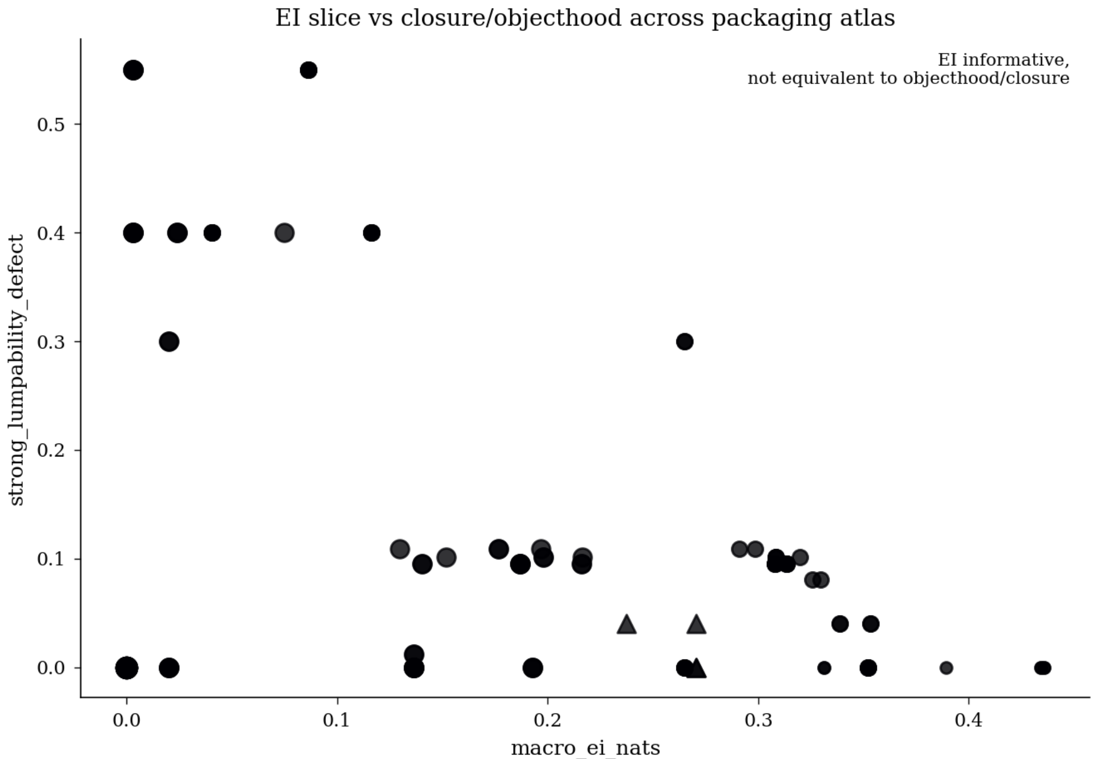

Figure 8: Effective information as a slice of the packaging atlas. High EI can coexist with trivial objecthood or with nonminimal closure defect, while some nontrivial packagings occupy more moderate EI regions. The figure supports the paper's comparison claim directly: EI is informative, but not equivalent to macro-law feasibility.

Taken together, the packaging-space atlases show that coarse-graining is itself a landscape of feasible and infeasible macro-laws. Closure, EI, directionality retention, novelty slack, and budget define a genuinely multi-objective geometry over that landscape. The main lesson is not that one partition wins everywhere, but that different scientific criteria align or diverge in structurally interpretable ways. The next section carries that same atlas logic into mechanism space, where the control variable is no longer the packaging of a fixed law but the choice of closure-interaction subalgebra itself.

6 Mechanism-Space Atlases in PICA

The first two empirical pillars varied microscopic law parameters and packaging choices. The third varies mechanism choice itself. This is where the atlas language reaches its most literal form: control space is now a combinatorial algebra of closure interactions. Instead of asking which partition of a fixed law is best, one asks which subalgebras of a minimal stochastic mechanism set generate which macroscopic regimes, how sensitive those regimes are to ablation, and how expensive or compact a geometry-compatible support can become before its behavior degrades. In that sense, PICA is not merely another substrate. It is the mechanism-space slice of the emergence atlas [15, 16, 18].

The mechanism-space datasets used here come from the targeted EXP-112 sweep together with the smaller PICA reproduction and settlement artifacts. All runs share the same finite-state audit logic introduced earlier in the paper, but now the control variables are mechanism labels, enable matrices, seed, and substrate scale rather than law parameters or partitions. The result is a

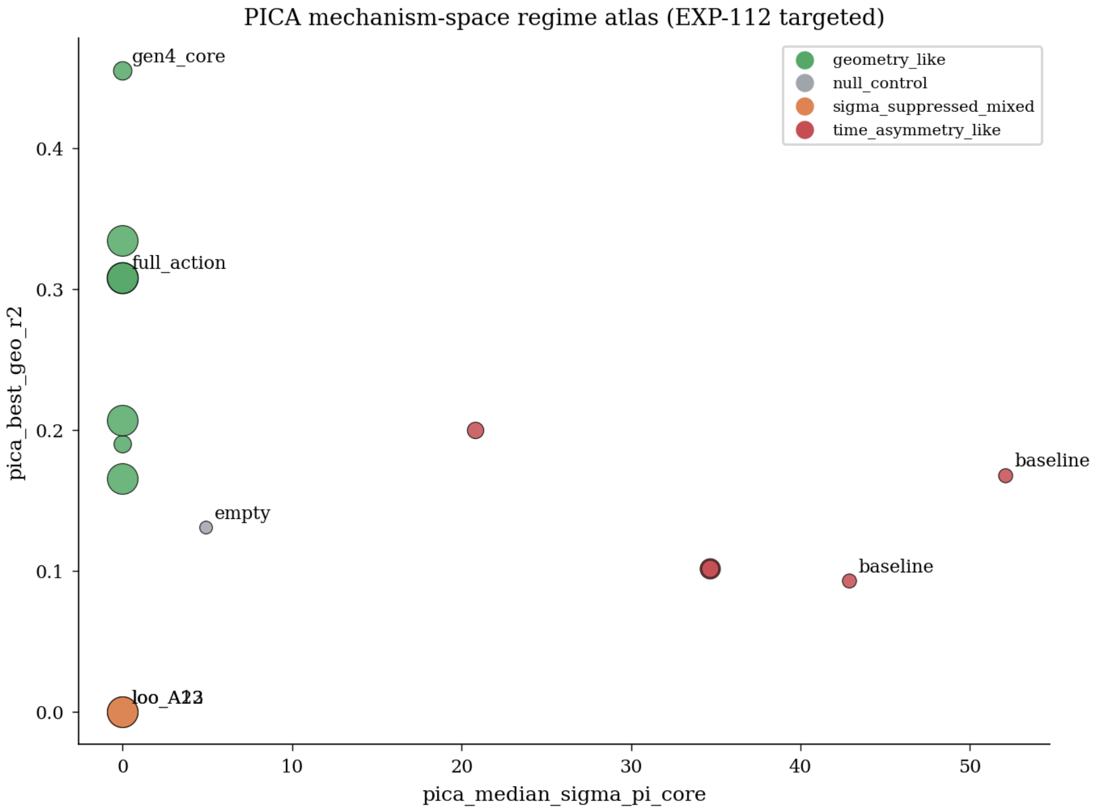

Figure 9: Mechanism-space regime atlas for the targeted EXP-112 PICA configurations. Null, time-asymmetry-like, geometry-compatible, and mixed configurations occupy visibly different regions of the same mechanism-space map. Marker size tracks support size, so the figure also shows that regime type and support cost need not align monotonically.

space of configurations in which null controls, time-asymmetry-like regimes, geometry-compatible regimes, and mixed or sigma-suppressed variants can be compared directly. The crucial point is that these are not merely labels attached after the fact. They are the visible organization of a saved mechanism-space atlas.

Mechanism-space regime separation. Figure 9 is the core mechanism-space map. It plots the EXP-112 mechanism signatures in a regime space spanned by core sigma and geometry score, with size encoding support size and labels identifying the anchor configurations. The null control sits apart from the main regime cloud. Baseline-like configurations occupy the time-asymmetry-like region, where sigma remains large and geometry support is weak or secondary. Full-action-style configurations move into the geometry-compatible or sigma-suppressed region, where sigma is strongly reduced and geometry support improves. Mixed or transitional cases populate the intermediate zone.

This figure matters because it makes the mechanism-space pillar visually comparable to the earlier law-space and packaging-space pillars. The atlas is no longer only a surface over law parameters or partitions; it is also a map over mechanism choices. Different closure-interaction subalgebras occupy different regions of that map, and the regime structure is visible without collapsing the underlying data into a single score. The purpose of this section is therefore not simply to say that PICA can produce interesting regimes. It is to show that mechanism-space itself can be charted in the same phase-diagram spirit as the other control spaces.

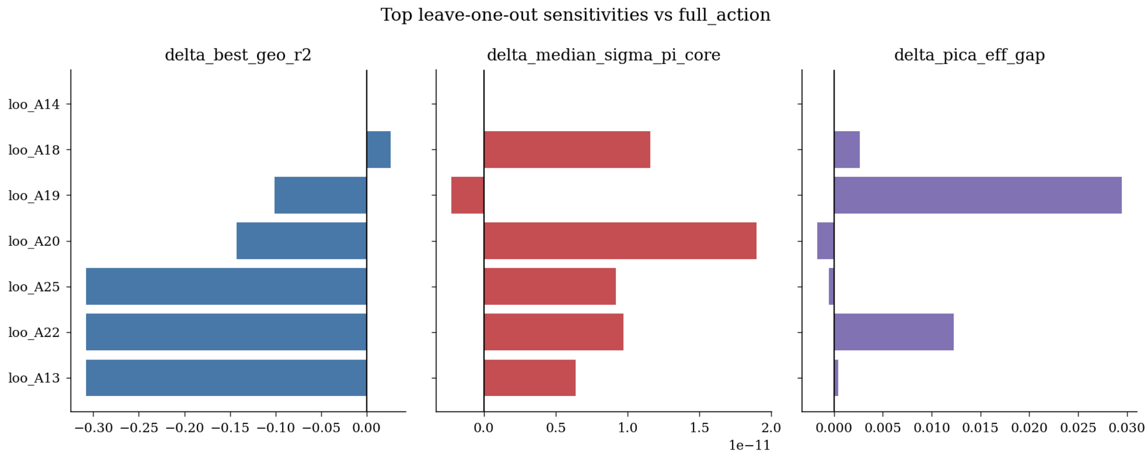

Figure 10: Leave-one-out ablation ranking in the targeted PICA sweep, shown relative to the full-action reference. The dominant sensitivity directions are A13, A22, and A25, which produce the largest geometry degradation and regime displacement. The figure should be read as a local sensitivity structure on mechanism space rather than as a generic feature-importance score.

Ablations as sensitivity directions. If the regime atlas shows where configurations sit, the ablation ranking shows which interaction coordinates matter most. Figure 10 ranks leave-one-out ablations against the full-action reference and displays the component deltas rather than only a black-box score. The three most sensitive removals are A13, A22, and A25. Their defining feature is not simply that they perturb one metric, but that they degrade the geometry-supporting regime in a visibly coordinated way: best geometry score falls sharply, sigma suppression weakens, and structural summaries move away from the full-action reference.

This is the mechanism-space analogue of a leverage or weakness field. In packaging space, certain partitions were frontier-relevant because they traded off closure, EI, and objecthood. In mechanism space, certain cells are frontier-relevant because removing them changes which regime a configuration can occupy. The ablation figure therefore serves as more than a ranked list. It is the local differential structure of the mechanism-space atlas.

Compact supports, budget, and conservative generator language. Mechanism space also supports a budget-quality frontier. Figure 11 compares the main compact-support candidates against the full and null references using support size, geometry score, and sigma-compatible coloring. The picture is informative precisely because it does not collapse to an easy triumphalist message. Some compact supports are clearly too weak: they retain too much sigma, lose geometry support, or fall into mixed behavior. Others remain geometry-compatible but do so at different support sizes and budget levels. The key point is that support size and regime quality can be plotted on the same mechanism-space frontier.

The conservative status of `gen4_core` should be read directly from this figure. It is the strongest current compact geometry-compatible candidate in the targeted sweep, but the settlement artifacts do not license an unqualified minimal-generator claim. The correct paper-facing language is therefore *near-minimal only*. The figure matters because it makes that caveat visible as part of the mechanism-space atlas rather than leaving it implicit in the supporting artifacts.

Multiscale signatures over *k*. The regime labels themselves are not single-point summaries. They arise from multiscale behavior over the macro resolutions visible in the saved PICA audits. Figure 12 plots sigma-like and geometry-like quantities over available k values for representative

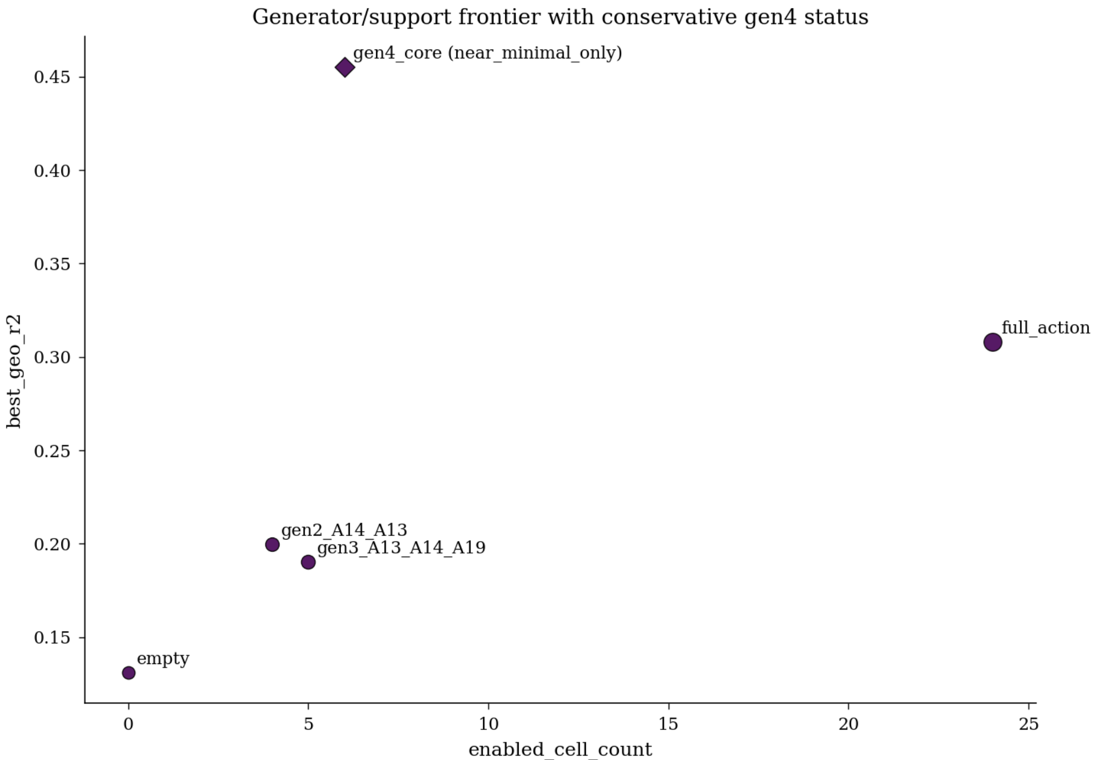

Figure 11: Mechanism-support frontier in PICA. Support size or budget is traded against geometry-compatible behavior, and compact candidates do not perform uniformly. The `gen4_core` point is intentionally annotated conservatively: it is the strongest current compact candidate in the targeted sweep, but only under a near-minimal-only settlement, not as an unqualified minimal generator.

configurations. The null control remains weak across the scan. Baseline-like configurations retain strong time-asymmetry signatures. Full-action-like configurations suppress sigma while exhibiting stronger geometry support over part of the ladder. The compact-support candidate can then be compared to these anchors as a multiscale object rather than as a single summary row.

This figure matters because it shows what the regime labels summarize. A configuration is not called geometry-compatible merely because one scalar in one row is favorable. It is called geometry-compatible because its profile over the accessible multiscale scan aligns with the corresponding signature family. The same is true for the time-asymmetry-like and mixed labels. In this sense, the multiscale figure is the dynamical underside of the regime atlas.

Taken together, the regime atlas, ablation rankings, support frontier, and multiscale profiles make mechanism space a full empirical pillar of the paper rather than a side case study. The same control-space language now applies across law space, packaging space, and mechanism space, but each axis of variation reveals something different. Law-space atlases show structural subsets such as exact closure loci and reversible lines. Packaging-space atlases show multi-objective tradeoffs among closure, EI, novelty, directionality retention, and budget. Mechanism-space atlases show which closure-interaction subalgebras generate which regimes, which ablations matter most, and which compact supports remain plausible only under conservative settlement language. The next section uses that shared structure directly by comparing the atlas view to several more familiar single-slice approaches.

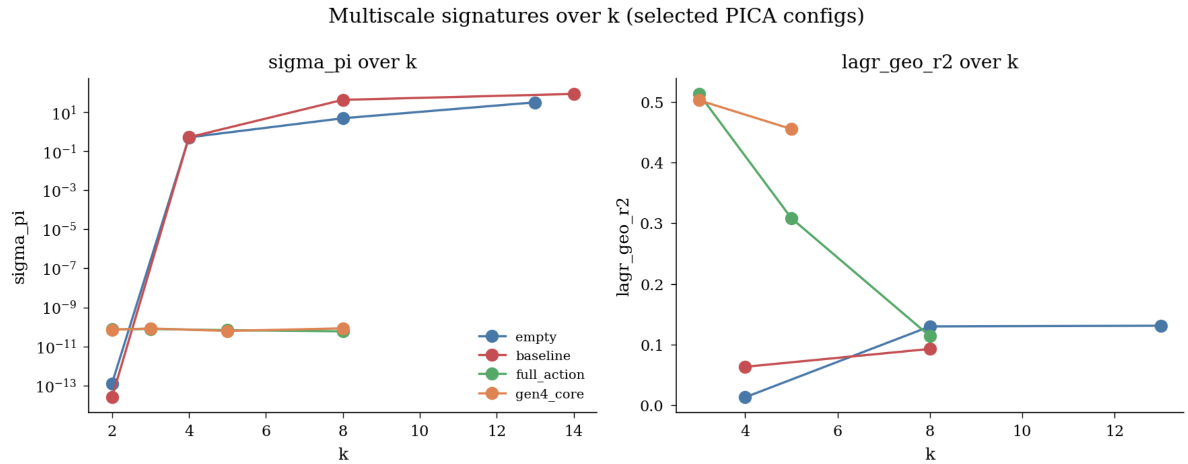

Figure 12: Representative multiscale signature profiles in PICA over available macro resolutions k. Null, time-asymmetry-like, geometry-compatible, and compact-support candidates differ not only in summary rows but in their behavior across the multiscale scan itself. The regime labels in the mechanism-space atlas therefore summarize real profile structure rather than a single isolated scalar.

7 Slices of the Atlas

The three empirical pillars of the paper now make it possible to compare the atlas view directly with several more familiar emergence-style objects. The point of this section is not to dismiss those objects, but to specify what each one captures and what each leaves out. Effective information, irreversibility, quantitative coarse-graining quality, and PICA mechanism-space analysis all detect something real. The atlas contribution is to place them on the same control-space substrate, so that their alignments and misalignments can be seen directly rather than inferred from separate literatures.

Table 1 presents that comparison in compact form. Each row identifies the inputs a slice requires, the quantities it computes, the aspect of emergence it detects most directly, the main blind spot that remains, and the sense in which it embeds into the larger atlas. The table should therefore be read as a synthesis device rather than as a replacement for the section-level empirical results.

The atlas view. The first row is the atlas itself. Its inputs are the full control variables relevant to the application: microscopic law, packaging, stage, protocol or mechanism, and budget. Its fields are correspondingly plural. Closure, objecthood, directionality, novelty, effective information, and cost appear together, and the central scientific question is whether a feasible macro-description exists at a given control point and what tradeoffs structure the surrounding region. The atlas therefore does not optimize a single quantity. It records where feasible descriptions exist and how the various slices align or fail to align there.

This is what the previous sections have already shown. In law space, closure loci and reversible lines appear as structured subsets of parameter space. In packaging space, closure, objecthood, EI, novelty, and budget define a tradeoff landscape rather than a single optimum. In mechanism space, regime labels, ablations, support frontiers, and budgets organize a combinatorial control space. The atlas row in the comparison table simply states this explicitly: all of these objects are part of one control-space cartography.

Effective information as a slice. The effective-information row is the clearest example of a useful but incomplete slice [8, 9, 12, 14]. In the present paper it is backed by actual computed

Table 1: Comparison-table view of the atlas and its principal slices. Each row identifies the quantities a slice computes, what it detects, its main blind spot, and how it embeds into the atlas. All rows except the engineering-emergence proxy are fully implemented.

View	Computed metrics	Detects	Misses	Atlas embedding
Atlas view	lumpability defect, path asymmetry, objecthood, macro EI, volume slack, budget proxy, regime label	feasible regions where closure, nontriviality, and constraint slices cohere	not a full external-framework implementation; depends on finite-state coverage	full multi-slice object
Effective-information slice	macro EI (nats), normalized EI, projected-micro EI, EI retention, EI gap	intervention-style discriminability in packaged futures	high EI does not guarantee closure or nontrivial objecthood	EI slice over packaging rows
Max-irreversibility slice	path-space time asymmetry, observed-path asymmetry, directionality retention, entropy production	irreversibility / arrow-of-time strength	irreversibility alone does not ensure closure or objecthood	directionality slice
Coarse-graining quality slice	lumpability defect, fiber row mismatch, retention error, idempotence defect	closure quality / coarse-graining fidelity	low defect can still be trivial; closure alone misses directionality and novelty	closure-quality slice
PICA mechanism-space slice	regime label, median σ_π, best geo-R^2, ablation deltas, generator status, cell count	mechanism subalgebras producing geometry-like vs. time-asymmetry-like regimes	not a replacement for packaging-space atlas; conclusions are substrate-dependent	mechanism / protocol slice
Engineering-emergence proxy†	volume slack, budget proxy, objecthood, frontier membership	pragmatic tradeoffs among cost, novelty, and nontriviality	not a canonical implementation of external frameworks	explicit proxy row only

†Implemented as an explicit partial proxy; see text.

quantities: macro EI, normalized macro EI, projected-micro EI, EI gap, and EI retention. Those quantities are not placeholders or rhetorical imports from another literature. They are computed directly on packaging-space rows and then placed beside closure, objecthood, novelty, and budget. That is why the EI comparison in Section 5 matters so much. It shows that EI can be large while the corresponding packaging remains trivial as a macro-object, and that a nontrivial packaging can remain scientifically useful with only moderate EI.

This is the key interpretive point. The atlas does not reject EI. It computes it and keeps it. What it rejects is the idea that EI alone is enough to decide macro-law feasibility. EI is strongest when the question is discriminability of packaged futures under intervention-style variation. It is weakest when closure, objecthood, or budget matter independently. That is why the comparison table treats EI as a slice rather than as a global verdict.

Irreversibility as a slice. The irreversibility row is built from the directionality quantities already used in the analytic and packaging sections: path-space asymmetry, observed-path asymmetry, directionality retention, and the packaged entropy-production-style proxy [5, 13]. These quantities are especially strong at detecting arrow-of-time structure and, when compared before and after deterministic observation, at auditing whether observation preserves or attenuates that structure. In that sense they are not merely descriptive: they constrain what a valid coarse or packaged description may claim about directionality.

But the same reason that makes irreversibility valuable also limits it. A high asymmetry score need not imply closure, good packaging, or nontrivial objecthood. This was visible already in packaging space, where some directionality-retaining partitions remained scientifically poor by the paper's broader criteria. The comparison row therefore gives irreversibility an important but bounded role: it is the directionality slice of the atlas, not a complete criterion for emergence.

Quantitative coarse-graining quality as a slice. The coarse-graining-quality row is built from the closure metrics: strong-lumpability defect, row mismatch, retention error, and idempotence defect [2, 7, 10]. These are precisely the quantities needed when the scientific task is to ask how well a packaged dynamics closes or how faithfully microscopic variation descends to the packaged law. The law-space closure loci and the packaging-space frontiers show why such metrics are indispensable.

At the same time, they are not enough by themselves. A nearly perfect coarse-graining can still be trivial; a closure-optimal partition can still have poor EI or little novelty slack; and a scientifically interesting packaging may remain only moderately closed while gaining nontrivial objecthood or interpretive value. This is why the comparison row calls the coarse-graining-quality slice a fidelity or closure slice rather than an emergence theory. It captures one of the paper's necessary criteria, but not the whole atlas.

PICA as a mechanism-space slice. The PICA row is the mechanism-space counterpart of the other slices. Its inputs are no longer just microscopic laws and partitions, but structured mechanism choices: enable matrices, configuration labels, scale, and seed. Its outputs are not only closure-like or directionality-like metrics but regime labels, multiscale signatures, ablation deltas, and generator/support summaries. This makes it the row that is furthest from classical coarse-graining language and closest to a true mechanism algebra.

That extra reach is exactly why the row still needs the atlas perspective. PICA is powerful at showing which closure-interaction subalgebras generate which macroscopic regimes, which removals matter most, and how compact a candidate support can be before its behavior weakens. But it is not a universal replacement for law-space or packaging-space analysis. Instead it is the mechanism-space slice of the same control-space picture. The paper's use of PICA is strongest precisely when it is treated that way.

The optional engineering-emergence proxy. The comparison table also includes an explicitly partial engineering-emergence proxy row. It is marked as partial on purpose. The current repo does not implement a full external engineering-emergence framework. What it does contain is enough to justify a careful proxy: novelty slack, budget proxies, and packaging frontiers already display the kind of cost-quality and possibility-frontier structure that such a row is meant to notice. Including the row in partial form is therefore more honest than either pretending the implementation is complete or omitting the connection entirely. The table makes that status visible rather than hiding it.

Why the atlas is the larger object. The point of the comparison table is not merely that several useful views coexist. It is that each row isolates a different aspect of macro-law feasibility: EI projects discriminability, irreversibility projects arrow-of-time retention, coarse-graining-quality metrics project closure and fidelity, and PICA projects mechanism-space organization. The atlas is the object on which those projections live together.

This is why the paper treats emergence as a mapping problem rather than as a single-score problem. No one row is guaranteed to settle macro-law feasibility by itself, but the atlas makes it possible to ask where several scientifically meaningful criteria align, where they fail to align, and how strongly that pattern is supported. The next section uses that same discipline to state which resulting claims are robust, which are mixed, and which must remain explicitly bounded.

8 Robustness, Settlements, and the Claim Surface

The preceding sections established the three empirical pillars of the paper and the comparison layer that situates their main slices. What they do not yet do is say how strongly each resulting claim should be carried. That is the task of the present section. The point is not to weaken the atlas narrative, but to discipline it. The paper is strongest when it states plainly which findings are robust under convention changes, which remain mixed or threshold-sensitive, which have been conservatively downgraded, and which broader extrapolations should be ruled out altogether.

Two ideas organize that discipline. First, the paper distinguishes empirical slice behavior from paper-facing claim status. A field may be robust in the sense that it persists across several axes of variation, while the corresponding manuscript claim is still phrased conservatively because its interpretation is easy to oversell. Second, null controls and explicit settlement artifacts are treated as first-class evidence. The paper does not only ask where an effect appears; it also asks whether the same logic would falsely promote a null or trivial case, and whether borderline findings survive dedicated settlement analysis.

Robust, mixed, and fragile findings. The robustness suite classifies the major findings by varying lift choice, stage, threshold conventions, seed, and, where affordable, substrate scale. Table 2 summarizes those outcomes. The most stable analytic results are the exact closure locus in the block family and the directionality axis in the driven family. Both remain intact under the tested lift and τ variations, and the null/reversible controls behave as they should. On the PICA side, the baseline time-asymmetry signal and the top ablation sensitivity ordering are likewise stable under the tested perturbations. These are the findings the paper can state most directly.

Other findings require more care. The packaging-space refinement story is real, but not monotone in the naive sense: the underlying phenomenon is robustly non-monotone, while individual help/hurt cases remain convention-sensitive enough that they should be described under explicit conditions rather than as a single universal tendency. Likewise, the geometry-supporting full-action regime is a genuine part of the PICA story, but its paper-facing wording remains narrower than the most optimistic reading of the raw regime labels. The section therefore treats robustness status as a filter on how strongly a result should be narrated, not merely as a binary pass/fail badge.

Null controls as claim guards. The null controls are not decorative. They are part of the claim surface itself. In the analytic package, constant or degenerate packagings remain non-objects

Table 2: Robustness-status summary for the major analytic and PICA findings. The matrix distinguishes results that remain stable across tested convention changes (lift, τ, threshold set, seed, scale) from those that are mixed or fragile. It is part of the paper's claim-surface discipline rather than an afterthought to the empirical sections.

Finding	**Domain**	**Status**	**Variation tested**
Exact closure locus ($\delta = 0$)	Analytic	robust	lift, τ
Directionality drive axis	Analytic	robust	τ, drive
Constant-map nonobject null	Analytic	robust	threshold set
Off-locus nontrivial tradeoff	Analytic	robust	lift, τ
Refinement help/hurt	Analytic	fragile	lift, τ
Baseline time-asymmetry regime	PICA	robust	seed, scale, threshold set
Empty null control	PICA	robust	scale, threshold set
Full-action geometry regime	PICA	robust	seed, scale, threshold set
Top ablation sensitivity (A13, A22, A25)	PICA	robust	threshold set
`gen4_core` candidate	PICA	mixed	seed, scale, threshold set

under the tested objecthood thresholds, and reversible law-space points remain near-zero on the directionality slices. In the PICA package, the empty control does not become a false geometry success merely because one relaxed threshold or one budget-like number happens to look favorable. This matters because a large class of emergence narratives fail precisely at this step: they can identify interesting cases, but they cannot stop themselves from promoting nulls under neighboring conventions.

The atlas view is stronger here because it does not rely on one privileged metric. A null control must fail in the appropriate way across several slices. It should fail to become a nontrivial object when objecthood is the question; it should fail to produce real arrow-of-time retention when the microscopic asymmetry is degenerate; and it should fail to win by budget or EI alone when the corresponding closure or objecthood conditions are absent. The robustness artifacts therefore act as guardrails on the interpretation of the main figures rather than as a detached supplementary check.

The gen4_core settlement. The most important settlement result concerns `gen4_core`. Without the settlement pass, the mechanism frontier could easily be overread as a minimal-generator discovery. The dedicated settlement analysis blocks that move. The saved conclusion is that `gen4_core` is a *mixed near-minimal* candidate and should be described in the manuscript only as *near-minimal only*. Some seed/scale comparisons make it look strong, including the key seed-1, scale-32 comparison that motivated its prominence in the first place. But other variants weaken or become coverage-deficient, especially once larger-scale or less favorable comparisons are included. The result is not a failure of the mechanism-space story. It is a successful containment of overclaim.

Table 3 makes that containment visible. The point of the table is not that the paper lacks compact-support structure. It is that the compact-support story must be narrated at the right level. The experiments support a meaningful frontier and a strongest current compact geometry-compatible candidate; they do not support a clean theorem-like statement that the candidate is an unqualified minimal generator. That is why the main text consistently keeps the 'near-minimal only' language in view.

Refinement as a settled non-monotone effect. The refinement settlement plays a different role. Here the issue was not whether an effect existed, but what its correct general form should be. A weaker version of the project could have been tempted to say either that refinement generally helps, because finer descriptions reveal more structure, or that refinement generally hurts, because finer

Table 3: Settlement-status summary for the two main claim-boundary issues. The table makes visible both the conservative `gen4_core` recommendation and the non-monotone refinement settlement, marking the boundary between what the paper claims strongly and what it presents only with explicit caveat.

Claim boundary	**Settlement status**	**Paper-facing recommendation**
`gen4_core` generator support	mixed near-minimal	near-minimal only; not an unqualified minimal-generator claim
Analytic refinement direction	robust help *and* hurt (687 help, 596 hurt pairs)	non-monotone under explicit conventions; no universal monotone tendency

descriptions expose more mismatch. The settlement artifacts support neither slogan. They support the stronger and more precise claim that refinement is non-monotone under explicit conventions: both help and hurt examples occur, and the correct way to talk about refinement is therefore as motion through packaging space rather than as a moral preference for finer or coarser scales.

This is exactly the right conclusion for the atlas view. Once closure, objecthood, and retention are tracked separately, one should expect refinement to help in some regions and hurt in others. The settlement pass strengthens that claim by removing the earlier floating-point/tie ambiguity from the original 'helps' control. The paper is therefore entitled to say that the non-monotonicity is real and explicitly exhibited, but not to say that refinement has a universal monotone tendency.

From robustness to claim surface. These robustness and settlement outputs are not just post hoc cautions. They are the final layer that turns the paper from a rich computational atlas into a controlled scientific argument. The frozen claim matrix built from these artifacts distinguishes safe claims, supported-with-caveat claims, downgraded claims, and explicit do-not-overclaim boundaries. That structure is especially important for the paper's most tempting extrapolations. The presence of robust help and hurt examples does not justify saying that refinement universally helps or universally hurts. The presence of a strong compact candidate does not justify saying that minimal generator discovery has been solved in PICA. The presence of an engineering-emergence proxy row in the comparison table does not justify saying that a full external framework implementation has been completed here.

What this section adds is not another slice of the atlas, but a rule for how the slices may be narrated. The final discussion can therefore return to the paper's central thesis with the claim surface under control: emergence is usefully treated as a mapping problem for macro-law feasibility in control space, but each component of that statement is carried only as strongly as the frozen evidence pack allows.

9 Discussion

The paper began from a simple observation: many sciences already use maps of macroscopic behavior, but emergence itself is rarely treated in that way. The central contribution of the present work is to make that cartographic move explicit. Instead of treating emergence as the search for one privileged coarse-graining, one privileged scale, or one privileged scalar score, the paper defines and computes *emergence atlases*: subsets and scalar fields on control space indicating where a macro-description is feasible. In the finite-state setting used here, feasibility means that a packaged description is simultaneously closed, nontrivial, audit-safe, and budget-feasible. The main claim of the paper is that this is a useful scientific object in its own right.

What makes the atlas useful is that it lets several familiar questions be asked on the same substrate without collapsing them into one another. Section 4 showed that law space can contain structurally interpretable subsets such as exact closure loci and reversible lines, and that these can be supported by dense continuity diagnostics and by explicit $\tau > 1$ evidence without being reduced to a single score. Section 5 showed that packaging space is not merely the place where one picks a best partition, but a landscape of tradeoffs among closure, objecthood, directionality retention, novelty slack, effective information, and budget. Section 6 then showed that the same language extends to mechanism space, where regimes, ablations, support frontiers, and multiscale signatures can be charted over structured subalgebras of closure interactions. The point is not that these spaces are identical. It is that they can be read through the same control-space lens.

This is also why the comparison layer in Section 7 matters. Effective information, irreversibility, coarse-graining quality, and PICA mechanism analysis all remain scientifically useful, but each captures only part of the larger control-space picture. The atlas view keeps them, computes them, and clarifies where they align and where they do not.

That reframing clarifies several substantive points. First, no single slice is guaranteed to identify the best macro-description everywhere. High effective information can coexist with trivial objecthood. Low closure defect can coexist with low novelty slack or low interpretive value. Strong directionality retention can coexist with poor packaging quality. A compact mechanism support can coexist with a still-unsettled generator status. These are not anomalies to be explained away; they are exactly the kinds of misalignment an atlas should display. Second, the atlas view changes how one should talk about emergence theoretically. The relevant question is not only whether a macro-description exists, but where in the available control variables it exists, how stable that region is under convention changes, and which scientific criteria align or diverge there.

The robustness and settlement layer keeps that conclusion scientifically disciplined. Some findings can be carried strongly; others require explicit caveat or downgrade. In particular, `gen4_core` remains near-minimal only, refinement remains non-monotone rather than universally improving or worsening, and the engineering-emergence row remains explicitly partial. The atlas view is stronger, not weaker, for making those limits visible.

The appendices record the supporting metric definitions, freeze logic, and bounded Lean layer in the same spirit: as support structure for the atlas argument rather than as parallel claims.

Several limits of the present paper are deliberate rather than accidental. The continuity support is empirical rather than formal, the $\tau > 1$ closure support strengthens but does not symbolically complete the law-space story, the EI slice does not exhaust causal-emergence-style ideas, and the mechanism-space results do not erase the settlement caveats on compact support discovery. These boundaries are part of the paper's final contribution rather than leftovers from an unfinished project.

The more positive lesson is that a great deal of emergence theory can be reorganized once macro-law feasibility is treated as a mapping problem. The six primitives recapped in Section 2—rewrite, constraints, protocols, staging, packaging, and accounting—then cease to be only a foundations vocabulary. They become coordinates of control space. P1 changes the effective law; P2 and P6 carve admissible and budgeted regions; P3 moves one through protocol or mechanism space; P4 changes the stage at which closure and retention are assessed; and P5 moves one through packaging space itself. The atlas object is the place where those coordinates meet. This is why the paper's central scientific image is a phase diagram rather than a scorecard.

The broader interpretation is therefore straightforward. Emergence is not well described solely as the discovery of a privileged higher level. It is more usefully described as the charting of where effective theories are feasible and what tradeoffs structure those regions [1, 20]. Sometimes that chart is drawn in law space, sometimes in packaging space, sometimes in mechanism space, and sometimes in a product of several such spaces. The present paper has provided one finite-state implementation of that idea and shown that it is rich enough to recover exact closure loci, reversible lines, packaging frontiers, EI slices, mechanism regimes, ablation sensitivities, and compact-support cautionary cases within one framework.

The guiding conclusion of the paper is therefore not that emergence has been reduced to a single formula. It is that emergence can be treated as a phase-diagram problem for effective theories. Once that move is made, many familiar tools remain valuable, but their relationship becomes clearer: they are slices of a larger cartography rather than competing total definitions. The most useful scientific question is then not merely which macro-description is best, but where in control space macro-descriptions exist at all, how they fail, how they trade off, and how strongly those conclusions are supported.

The central question is not only what macroscopic theories exist, but where in control space they can be made to exist.

A Metrics and Implementation Notes

This appendix records the operational quantities used throughout the paper. It is not intended as a second theory section. Its role is to make the atlas computations inspectable without overloading the main text with implementation detail.

Finite-state conventions. All finite-state dynamics are represented by row-stochastic kernels

$$P \in [0,1]^{N\times N}, \qquad \sum_{j=1}^{N} P_{ij} = 1,$$

with microscopic states indexed as $X = \{1,\ldots,N\}$. Packagings are deterministic partitions

$$\Pi = \{B_1,\ldots,B_K\}$$

with associated macro state space $Y = \{1,\ldots,K\}$. The paper uses the same row-vector convention throughout the computations: one-step evolution is $\mu \mapsto \mu P$, and staged evolution is $\mu \mapsto \mu P^{\tau}$.

Closure and objecthood slices. The main closure diagnostics are:

- the strong-lumpability defect $\mathrm{LD}_{\Pi}^{(\tau)}(P)$,
- the row-mismatch quantity $\mathrm{RM}_{\Pi,\mathbf{u}}^{(\tau)}$,
- the retention error $\mathrm{Ret}_{\Pi,\mathbf{u}}^{(\tau)}$,
- the empirical idempotence defect $\mathrm{ID}_{\Pi,\mathbf{u}}^{(\tau)}$.

These serve different purposes. Strong-lumpability defect tracks exact stagewise closure at the level of aggregated block futures. Row mismatch detects within-block heterogeneity relative to the induced packaged kernel. Retention error tracks metastability of packaged basis states. Idempotence defect tracks how close the empirical packaging endomap is to a stable closure-like action on microscopic distributions.

Nontriviality is kept separate from closure. The paper's default objecthood indicator is

$$\mathrm{Obj}_{\Pi,\mathbf{u}}^{(\tau)} = \mathbf{1}\Big[\mathrm{SB}_{\Pi,\mathbf{u}}^{(\tau)} \geq 2 \ \wedge \ \mathrm{DR}_{\Pi,\mathbf{u}}^{(\tau)} \geq 2\Big],$$

where SB is the stable-block count and DR is the distinct-row count up to the declared tolerance. This separation is essential: some of the best-closed packagings remain trivial, and some high-EI packagings remain scientifically poor as macro-objects.

Directionality slices. Directionality is tracked primarily through finite-horizon path-space asymmetry and its observed counterpart:

$$\Sigma_T(\rho;P) = D_{\mathrm{KL}}(\mathbb{P}_{\rho,T} \,\|\, \mathcal{R}_*\mathbb{P}_{\rho,T})\,, \qquad \overline{\Sigma}_T(\rho;P,\Pi) = D_{\mathrm{KL}}\Big((\pi_\Pi^{T+1})_*\mathbb{P}_{\rho,T} \,\Big\|\, \mathcal{R}_*(\pi_\Pi^{T+1})_*\mathbb{P}_{\rho,T}\Big)\,.$$

When the microscopic asymmetry is positive, the paper also uses the directionality-retention ratio

$$\mathrm{DirRet}_T(\rho;P,\Pi) = \frac{\overline{\Sigma}_T(\rho;P,\Pi)}{\Sigma_T(\rho;P)}.$$

This quantity is intentionally left undefined in degenerate reversible cases rather than being forced into an arbitrary zero or unit convention. A stationary entropy-production-style proxy is also used where a packaged stationary law is available. In the paper's logic, these are directionality slices, not complete definitions of emergence.

Novelty, effective information, and budget slices. Novelty is tracked through hidden slack beneath the packaging:

$$\mathrm{HV}(\Pi) = N - K, \qquad \mathrm{SEP}(\Pi) = \begin{cases} 1 - 2^{-(N-K)}, & N > K, \\ 0, & N \le K. \end{cases}$$

These quantities should be read as novelty-potential slices rather than realized novelty measures. They formalize the amount of remaining hidden room beneath a packaging and provide a computational counterpart to the strict-extension/finite-forcing picture.

Effective information is treated as a packaging slice rather than as a universal objective. The macro effective information is

$$\mathrm{EI}_{\mathrm{macro}}(M) = \frac{1}{K}\sum_{a=1}^{K} D_{\mathrm{KL}}(M(a,\cdot) \,\|\, q)\,, \qquad q = \frac{1}{K}\sum_{a=1}^{K} M(a,\cdot),$$

with normalized version

$$\mathrm{EI}^{\mathrm{norm}}_{\mathrm{macro}}(M) = \begin{cases} \mathrm{EI}_{\mathrm{macro}}(M)/\log K, & K > 1, \\ 0, & K = 1. \end{cases}$$

The projected-micro EI and the resulting EI gap and EI-retention ratio are used to compare what is visible before and after packaged basis-state compression. The paper's packaging-space results depend crucially on the fact that EI can align with, diverge from, or even conflict with closure and objecthood.

Budget is instantiated through explicit within-domain proxies. For analytic and packaging rows, the default structural budget is

$$\mathrm{Bud}_{\mathrm{pack}} = K + c_{\mathrm{lift}} + \tau,$$

with $c_{\mathrm{lift}} = 1$ for uniform lift and $c_{\mathrm{lift}} = K$ for prototype lift. In PICA, the main budget coordinates are enabled-cell count, gated-edge count, and the saved mechanism-budget proxy when available. These budget quantities are comparative within domain; the paper does not claim a universal cross-domain cost unit.

Atlas rows, derived enrichments, and summaries. The implementation distinguishes source result rows from derived enrichment layers. EI, novelty, and budget were added as post hoc enrichments rather than by destructively rewriting the original experimental outputs. This design choice matters for provenance and for the final freeze pack: the atlas rows preserve their lineage to source scans, and the derived summaries remain inspectable as overlays rather than as silent mutations of the underlying evidence.

Finite-state scope. Nothing in the paper depends on finite-state systems being the only scientifically interesting setting. The finite-state specialization is used because it makes the atlas computable, because exact path-space and packaged quantities can be evaluated directly on small systems, and because law space, packaging space, and mechanism space can all be instantiated on the same substrate. The conceptual claim of the paper is cartographic; the finite-state layer is the concrete implementation used to support it.

B Provenance and Freeze Notes

The experimental side of the paper is frozen. This appendix records what that means operationally and how the manuscript is connected to the frozen evidence pack.

Frozen-pack principle. The manuscript is written against a frozen experiment pack rather than against a moving result tree. The frozen pack collects selected figures, selected datasets and summaries, a claim-confidence matrix, a frozen snapshot of the findings ledger, a frozen snapshot of the artifact registry, checksum records, lineage records, and a ready-to-write summary. The purpose of the freeze is not archival aesthetics. It is to ensure that the paper's claim surface is tied to a stable evidence object rather than to a live repository state.

Selected evidence layers. The frozen pack contains four main evidence layers:

- **Analytic law-space assets:** dense closure and directionality figures, continuity diagnostics, and the dedicated $\tau > 1$ closure-support note.
- **Packaging-space assets:** full partition-scan atlases, frontiers, EI and novelty/budget enrichments, and packaging figure panels.
- **PICA mechanism-space assets:** regime atlas, ablation figure, support frontier, multiscale profiles, settlement summaries, and the targeted mechanism-space datasets.
- **Cross-framework and claim-control assets:** the comparison table, robustness summaries, settlement summaries, and the final claim-confidence matrix.

The main paper cites and displays a curated subset of these assets. The appendices and freeze metadata preserve the broader linkage structure.

Claim matrix and claim statuses. The frozen claim-confidence matrix is the paper's main claim-discipline artifact. It distinguishes:

- `safe_claim`,
- `supported_with_caveat`,
- `downgraded_claim`,
- `do_not_overclaim`.

These statuses are not rhetorical decorations. They are generated from the same frozen evidence pack that supports the manuscript figures and summaries. In particular, the paper's treatment of gen4_core, refinement non-monotonicity, and the engineering-emergence proxy row is governed by the frozen claim matrix rather than by informal narrative preference.

Checksums, lineage, and zero-orphan validation. Every frozen asset in the pack carries a checksum record, and the frozen registry records source paths, artifact kind, family, and claim linkage. The freeze pack also includes an explicit orphan check. The paper therefore does not rely on the assumption that the selected figures and summaries are self-explanatory or implicitly well linked. Their lineage is part of the frozen evidence object.

What the freeze excludes. The freeze is intentionally selective. It does not attempt to preserve every intermediate artifact, every raw exploratory output, or every optional development branch in the repository. Nor does it treat optional future work as pending evidence. In particular, no further Lean expansion, no further broad experimental sweeps, and no additional post-freeze evidence-collection steps are presumed by the manuscript. The writing phase therefore starts from a bounded scientific object: the experiment is frozen, the claim surface is frozen, and the manuscript is not waiting on further implementation.

How to read the main text in light of the freeze. The main text should therefore be read as a guided interpretation of a frozen, traceable evidence pack. The figures are selected views into that pack. The comparison table is a selected synthesis of rows derived from it. The robustness and settlement sections are explicit constraints on how strongly the figures may be narrated. The corresponding experiment and manuscript repository is available at `https://github.com/ioannist/six-birds-maps`. This is why the paper can make strong claims where the evidence is strong and remain disciplined where the evidence is mixed or intentionally bounded.

C Lean Note

The Lean layer in this project is intentionally small. Its role is not to formalize the full theorem spine of the paper, nor to provide a formal development of the finite-state probabilistic machinery. Instead it serves as a bounded support layer for the split-idempotent and fixed-point framing that underlies the packaging interpretation.

What is formalized. The isolated Lean package currently freezes a very small theorem set:
- `split_idempotent`,
- `split_image_fixed`,
- `split_fixed_factorization`.

These lemmas formalize the basic split-idempotent picture: the composition built from a section/retraction pair is idempotent, image points are fixed, and the packaged-image/fixed-point relationship can be factored in the small sense needed here.

What is not formalized. The Lean layer does *not* formalize:
- path-space KL quantities,
- data-processing statements for observed directionality,
- finite-state coarse-graining metrics,
- novelty slack or budget proxies,
- PICA regime or mechanism-space computations.

Those parts of the paper remain computational and empirical rather than formally mechanized.

Why the Lean layer is kept small. This is a deliberate scope choice. The paper's main contribution is the atlas object and its computation across several control spaces, not a large formalization campaign. A tiny Lean layer is enough to make the fixed-point/descent framing feel intentional rather than accidental, while avoiding the false impression that the theorem-support spine has been fully formalized. The manuscript therefore uses Lean as bounded support, not as a substitute for the frozen computational evidence pack.

Isolation from the main pipeline. The Lean package remains isolated from the main Python and manuscript smoke paths. It builds independently and does not participate in the default scientific or CI smoke routines. That separation is also intentional: it keeps the formal support layer proportional to the paper's actual argumentative load.

References

[1] P. W. Anderson. More is different. *Science*, 177(4047):393–396, 1972. doi: 10.1126/science.177.4047.393.

[2] Peter Buchholz. Exact and Ordinary Lumpability in Finite Markov Chains. *Journal of Applied Probability*, 31(1):59–75, 1994. doi: 10.2307/3215235.

[3] Pin Yu Chew and Aleks Reinhardt. Phase diagrams—Why they matter and how to predict them. *The Journal of Chemical Physics*, 158(3):030902, 2023. doi: 10.1063/5.0131028. Perspective.

[4] Hsiao-Dong Chiang and James S. Thorp. Stability regions of nonlinear dynamical systems: a constructive methodology. *IEEE Transactions on Automatic Control*, 34(12):1229–1241, 1989. doi: 10.1109/9.40768.

[5] Gavin E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721–2726, 1999. doi: 10.1103/PhysRevE.60.2721.

[6] Mauro Faccin, Michael T. Schaub, and Jean-Charles Delvenne. State Aggregations in Markov Chains and Block Models of Networks. *Physical Review Letters*, 127(7):078301, 2021. doi: 10.1103/PhysRevLett.127.078301.

[7] Bastian Hilder and Upanshu Sharma. Quantitative Coarse-Graining of Markov Chains. *SIAM Journal on Mathematical Analysis*, 56(1):913–954, 2024. doi: 10.1137/22M1473996.

[8] Erik P. Hoel. When the Map Is Better Than the Territory. *Entropy*, 19(5):188, 2017. doi: 10.3390/e19050188.

[9] Erik P. Hoel, Larissa Albantakis, and Giulio Tononi. Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013. doi: 10.1073/pnas.1314922110.

[10] John G. Kemeny and J. Laurie Snell. *Finite Markov Chains*. D. Van Nostrand, Princeton, NJ, 1960. ISBN 978-0442043285.

[11] Yuri A. Kuznetsov. *Elements of Applied Bifurcation Theory*, volume 112 of *Applied Mathematical Sciences*. Springer, Cham, 4th edition, 2023. ISBN 978-3-031-22006-7. doi: 10.1007/978-3-031-22007-4.

[12] Fernando E. Rosas, Pedro A. M. Mediano, Henrik J. Jensen, Anil K. Seth, Adam B. Barrett, Robin L. Carhart-Harris, and Daniel Bor. Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data. *PLOS Computational Biology*, 16 (12):e1008289, 2020. doi: 10.1371/journal.pcbi.1008289.

[13] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012. doi: 10.1088/0034-4885/75/12/126001.

[14] Giulio Tononi and Olaf Sporns. Measuring information integration. *BMC Neuroscience*, 4:31, 2003. doi: 10.1186/1471-2202-4-31.

[15] Ioannis Tsiokos. To create a stone with six birds: Emergent geometric and thermodynamic regimes from a minimal stochastic substrate. Manuscript. DOI: `https://doi.org/10.5281/zenodo.18838994`, 2026.

[16] Ioannis Tsiokos. Six birds: Foundations of emergence calculus. Manuscript. DOI: `https://doi.org/10.48550/arXiv.2602.00134`, 2026.

[17] Ioannis Tsiokos. To lay a stone with six birds: Finite-state semantics for packaging, directionality, and coarse-graining. Manuscript. DOI: `https://doi.org/10.20944/preprints202602.1699.v1`, 2026.

[18] Ioannis Tsiokos. To plot a stone with six birds: Constructing and auditing emergent geometry from markov dynamics. Manuscript. DOI: `https://doi.org/10.5281/zenodo.18494975`, 2026.

[19] Ioannis Tsiokos. To spend a stone with six birds: Currency, constraint duality, and shadow prices across closure layers. Manuscript. DOI: `https://doi.org/10.5281/zenodo.18926771`, 2026.

[20] Kenneth G. Wilson. The renormalization group: Critical phenomena and the Kondo problem. *Reviews of Modern Physics*, 47(4):773–840, 1975. doi: 10.1103/RevModPhys.47.773.