


To Wake a Stone with Six Birds: Audited Life-Like Signatures in Toy Substrates

Ioannis Tsiokos 

27 January 2026

Version 2 (revised 16 February 2026) · Zenodo v1: [10.5281/zenodo.18394536](https://doi.org/10.5281/zenodo.18394536)

Abstract

Empirical work on life-like organization often conflates protocol artifacts with genuine directionality and lacks a disciplined way to state and check claims across descriptive levels. This manuscript is a methodological case study in synthetic computational substrates and does not claim the creation or instantiation of biological life. We report a life-motivated application of the emergence calculus developed in Tsiokos [2026] by specifying finite theory packages $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ for two substrates (a particle-based substrate and a neural/meta-layer substrate) and indexing each reported claim to external experimental artifacts via a machine-readable ledger. Across both substrates we (i) calibrate strict null regimes (P3=OFF, P6=OFF) where audit proxies are statistically consistent with zero, reducing false-positive arrow-of-time signals; (ii) demonstrate a separable drive channel (P6) that yields clean null-drive separation in audit proxies under matched controls; and (iii) show that protocol holonomy (P3) changes stroboscopic currents under matched control while remaining a diagnostic rather than a directionality certificate. We further report maintenance/viability proxies—repair performance, traversal requirements, and deadline selection with explicit audit cost—and, in the neural substrate, refined-lens predicate families: stable motif inventories, proto-syntax shifts, and intervention-based dictionary/decoding statistics validated by shift-null significance tests. While the framework allows arbitrarily many theory extensions in the definability sense (adding new lens predicates), we do not claim autonomous open-ended evolution in the Artificial Life (ALife) sense (indefinite novelty generation under a fixed regime without external steering). Limitations are explicit: audits are proxies (not full path-space Kullback–Leibler divergence), idempotence defects are not measured, and formal definability separations are not proved.

keywords: emergence calculus; directionality; audit proxies; life-like organization; finite theory package; protocol holonomy; nonequilibrium dynamics; computational toy models

1 Introduction

A prior framework paper [Tsiokos, 2026] proposes an emergence calculus in which empirical claims are organized by a finite theory package $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ and three certificates: stability/objecthood, novelty/extension, and directionality/audit. That framework is intentionally math-only and explicitly defers instantiations. The present work presents a life-motivated application by fixing concrete $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ packages for two substrates and reporting certificate evidence with conservative controls. This is a Research Article: it reports original computational experiments with automatically generated evidence tables and conservative controls, not a review, opinion, or purely theoretical contribution. In Six Birds terms, “theory” denotes a closure (lens + completion + audit) rather than a speculative hypothesis or manuscript-type category; accordingly, the title’s claim that “a life is a theory” means that life-like competencies are organized as auditable closure packages.

Intellectual context and credit. While the emergence calculus and contracts used here are our own, this life-motivated application was strongly inspired by two adjacent lines of work that emphasize graded, multi-scale agency and the centrality of *interpretation across layers*. In the diverse-intelligence program, Michael Levin and collaborators argue that the genotype–phenotype relationship is best understood as an active, multi-scale problem-solving and error-correcting process (not merely a static decoder), and they develop substrate-independent characterizations of cognition in terms of remapping and navigating embedding spaces under iterative error minimization [Hartl and Levin, 2025, Hartl et al., 2026]. In a complementary direction, Bennett’s thesis develops a “stack” view in which policies at successive abstraction layers carve out stable causal identities and operationally goal-directed behavior (goal-directed in a control-theoretic sense, not ascribed intention) [Bennett, 2025]. We include this paragraph explicitly because significant parts of our motivation came from these perspectives, even though the present paper does not derive its formalism from them.

Empirical workflow. Our empirical workflow is deliberately audit-first: we calibrate strict null regimes in which protocol (P3) and drive (P6) are disabled, and require directionality-relevant audit proxies to be consistent with zero before interpreting any driven behavior as an arrow-of-time signal. We then (i) demonstrate drive separability via P6 under matched null controls, (ii) quantify protocol holonomy effects via matched-control stroboscopic diagnostics without attributing directionality to P3 alone, and (iii) report maintenance and higher-level descriptive structure only where supported by explicit gates, contracts, and shift-null tests.

What “life-like” means here. Throughout, we use “life-like” as shorthand for a limited family of *operational competencies* that can be stated and checked within the emergence-calculus vocabulary: (i) an audited directionality channel that separates cleanly from null controls, (ii) maintenance/repair and perturbation recovery under externally imposed hazards or deadlines with explicit costs, and (iii) at refined lenses, stable predicate families that support controlled readout under explicit gates and shift-null tests. We are not proposing a general definition of biological life, and we do not model reproduction, heredity, or population-level evolution; replication enters only as contextual comparison in the Discussion. Accordingly, we do not claim that any simulated system is alive, nor that these experiments support abiogenesis claims beyond abstract operational analogies.

Title intuition. Here, “to wake a stone” is a metaphor for moving from a strictly calibrated null (no protocol, no drive) to a regime with auditable directionality and maintenance under explicit costs; the “six birds” name the primitive operations we toggle to make that transition testable rather than rhetorical.

Claim ledger. To keep the narrative auditable, every claim used in the paper is indexed in `assets/claims.yml`, which maps from claim IDs to the precise external tables/contracts in the particle and neural experiment repositories. This paper therefore serves as a worked example of how the emergence-calculus vocabulary can be used to state, audit, and bound life-like claims in complex substrates.

Related work. Related efforts toward general theories of life also emphasize constraint-driven universality across scales; for example, Solé et al. survey six recurring biological constraint families spanning thermodynamics, molecular information, cellularity, development, cognition, and ecosystem

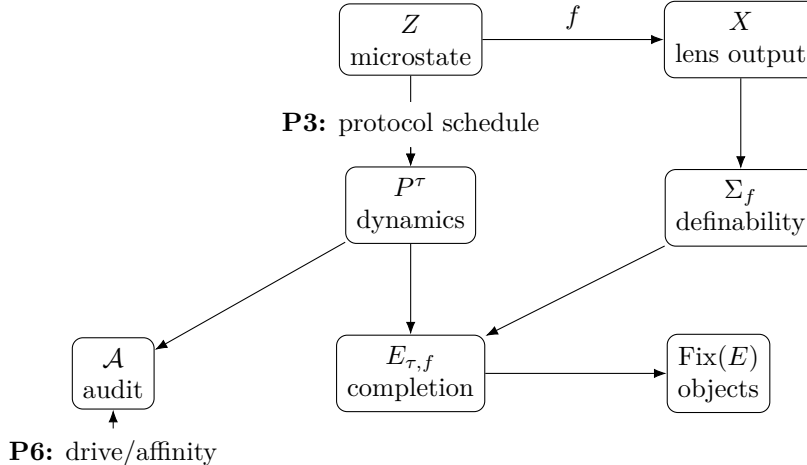


Figure 1: Schematic of the finite theory package $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ and where P3/P6 enter. P3 is treated as a holonomy/protocol *diagnostic* that modifies coarse behavior via protocolized dynamics, while sustained directionality is attributed only when an audit channel separates cleanly (e.g., P6 in table 2).

organization [Solé et al., 2024]. In the Discussion we note a suggestive (but non-canonical) one-to-one correspondence between these six families and the six “birds” (P1–P6) of the emergence calculus.

Code availability. Experiment repositories are available at:

- Neural substrate: <https://github.com/ioannist/six-birds-neural>
- Particle substrate: <https://github.com/ioannist/six-birds-particle>

2 Framework recap (canonical)

This paper instantiates the emergence calculus developed in the framework paper [Tsiokos, 2026]. The framework isolates a *finite theory package* and three logically distinct certificates—stability/objecthood, novelty/extension, and directionality/audit—so that empirical instantiations can be stated and checked without conflating these roles.

2.1 Finite theory package

A *finite theory package* is a tuple

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

Here Z is the microstate space. A *lens* $f : Z \rightarrow X$ specifies what is observable at a given descriptive level; it induces a definability structure Σ_f (equivalently, the partition of Z into the fibers $f^{-1}(x)$). A *completion/packaging endomap* $E : \mathcal{V} \rightarrow \mathcal{V}$ acts on a chosen description space \mathcal{V} (for example, distributions over Z), and its fixed points $\text{Fix}(E) = \{v \in \mathcal{V} : E(v) = v\}$ are the “objects” recognized at that level. Finally, an *audit functional* \mathcal{A} provides a nonnegative certificate intended to be monotone under further coarse observation (so coarse-graining should not create false positives) [Cover and Thomas, 2006]. See figure 1 for a schematic.

2.2 Three certificates

The framework separates three questions that are often entangled:

1. **Stability / objecthood.** A descriptive level exhibits objecthood when its completion rule behaves (approximately) as a closure: applying E does not keep changing the description, and there exist multiple robust fixed points (not a trivial constant map). Empirically, this is typically evaluated through stability/retention proxies rather than claimed by definition.
2. **Novelty / extension.** Iterating a fixed completion saturates. Strict novelty corresponds to a *theory extension*: changing what is definable (e.g., by adjoining predicates not definable from the prior lens) or changing the completion rule itself. In this paper we use “novelty” only in this definability-based sense, not as a synonym for “interesting dynamics.”
3. **Directionality / audit.** Directionality is certified by an audit. A standard example is the path-space arrow-of-time at horizon T [Lebowitz and Spohn, 1999, Seifert, 2012],

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

where \mathcal{R} reverses trajectories. Because $\Sigma_T(\rho)$ is a Kullback–Leibler (KL) divergence it is nonnegative, whereas the finite-window proxy estimators we report can be signed (and occasionally slightly negative) due to accepted-move-only logging, ignored proposal terms, and sampling noise; we interpret them via null calibration, confidence intervals, and clean regime separation rather than as absolute thermodynamic quantities. In our instantiations we report concrete audits and audit *proxies*, and we are explicit about when a reported quantity is a proxy rather than an exact path-space KL audit.

2.3 Protocols and the P3 boundary

Protocol holonomy (primitive P3) is a *diagnostic* describing how coarse behavior can depend on the route by which a closure or update is realized. Protocols can strongly affect coarse currents and stroboscopic statistics, but **P3 is not itself a directionality certificate**. Under autonomy, if the protocol phase is treated as part of the state (a lifted Markov chain) and each phase update is reversible with respect to a common stationary distribution, then the resulting lifted process has zero path-space asymmetry; any sustained arrow-of-time must be certified by an audit and, in particular, by genuine drive/affinity structure (primitive P6), not by protocol noncommutativity alone.

2.4 This paper as a life-motivated application

The framework paper [Tsiokos, 2026] intentionally remained math-only and explicitly deferred concrete instantiations. The present work delivers a life-motivated application by specifying $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ for two substrates—a particle-based substrate and a higher-level neural/meta-layer substrate—and by reporting certificate evidence across clean null regimes, separable drive regimes, protocol diagnostics, and higher-level structure where supported by explicit gates and controls.

3 Experimental substrates (particles; neural)

We instantiate the canonical package $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ in two different substrates: a particle-based substrate (with discrete slow variables and field-like packaging) and a higher-level neural/meta-layer

substrate (with budgeted operator tokens and stroboscopic diagnostics). For traceability, every empirical claim we rely on is indexed in a machine-readable claim ledger (`assets/claims.yml`), which points to the corresponding external artifacts in the two experiment repositories.

3.1 Particle-based substrate

Microstate space $Z^{(p)}$. A microstate consists of N particles on a 2D torus coupled with bounded discrete slow variables. Operationally, $Z^{(p)}$ includes (i) particle positions $x_i \in \mathbb{T}^2$, (ii) a bond/coupling array w_{ij} (primitive P1), (iii) bounded per-particle counters a_i and n_i (primitives P2/P4), (iv) a bounded packaging field S on a finite grid (primitive P5), and (v) optional meta-layer and operator-lifted token fields (used in the operator-coupling experiments). Protocol phase/clock variables are included when protocol diagnostics are enabled (P3), so that protocol effects can be interpreted either as an external schedule or as a lifted autonomous chain with phase in-state.

Lenses and definability. We use a small family of lenses $f^{(p)} : Z^{(p)} \rightarrow X$ corresponding to the descriptive levels used in the experiments: (i) a safe/field lens f_{safe} that reports the packaging field S (and derived safe-set statistics), (ii) an operator-coupling lens f_{opK} that reports operator token fields and interface mismatch summaries, and (iii) a diagnostics lens f_{diag} that reports windowed observables (currents/affinities, motif counts, protocol holonomy observables). Each choice of f induces the definability structure Σ_f as the partition of microstates into f -fibers.

Completion/packaging $E^{(p)}$. Packaging is implemented by the P5 update channel (field writes) and its meta/operator extensions. At the theory level we refer to the dynamics-induced completion template

$$E_{\tau,f}(\nu) := Q_f((U_f\nu)P^\tau),$$

where P is the Markov kernel induced by the simulator dynamics, Q_f denotes pushforward along the lens, and U_f is a fixed prototype lift. We emphasize that while the repository specifies an explicit prototype lift at the documentation level, it does *not* implement or measure the idempotence defect

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

Accordingly, objecthood/closure is assessed through descriptive stability proxies (persistence of safe-set fractions, component statistics, and qualitative retention in logs), rather than by claiming a measured $\delta_{\tau,f}$.

Audits $\mathcal{A}^{(p)}$ (and proxies). The particle substrate reports concrete audit *proxies* built from coarse event counts and acceptance logs. The main audit proxies are: (i) coarse current–affinity statistics per writable coordinate (yielding windowed affinities A_y , currents J_y , and a memory-like aggregate $\Sigma_{\text{mem}} = \sum_y J_y A_y$), (ii) motif-based proxies (e.g., the M6 family) that are near-zero in null and become nonzero under drive, and (iii) an acceptance-log entropy-production proxy (“epExact” in the repo) that is explicitly treated as a proxy rather than an exact path-space KL audit. Clean null calibration and drive separability for these proxies are part of the SAFE results (see the claim ledger entries beginning with C).

3.2 Neural/meta-layer substrate (ratchet neural)

Microstate space $Z^{(n)}$. A microstate is a layered 2D lattice state with bounded local degrees of freedom and hard budget constraints. Concretely, $Z^{(n)}$ includes per-site symbol/activity variables

(e.g., σ), auxiliary bounded fields (e.g., n and a barrier/closure field s), and two classes of operator tokens: within-layer tokens W with a global budget constraint, and inter-layer tokens K with a per-site budget constraint (primitive P2). Multiple layers implement meta-level coupling, with diagnostics that summarize cross-layer mismatch and token geometry.

Lenses and the diagnostic extension Z_{diag} . We use lenses that expose (i) stroboscopic coarse signatures and currents, (ii) cross-layer mismatch summaries, and (iii) higher-level motif and dictionary statistics in later phases. Some reported “snapshot” outputs depend on the entropy-production (EP) ledger (an EP accumulator), and are therefore best viewed as lenses on an extended diagnostic state $Z_{\text{diag}} := Z^{(n)} \times L_{\text{EP}}$, where L_{EP} is the EP ledger. This bookkeeping distinction is made explicit to avoid treating diagnostic accumulators as part of the physical microstate when reasoning about autonomy and reversibility.

Completion/packaging $E^{(n)}$. In the neural substrate, packaging is realized through energy/barrier mechanisms (P5) and structured token-mediated coupling between layers. As in the particle case we use the canonical completion template $E_{\tau,f}(\nu) = Q_f((U_f\nu)P^\tau)$ as the theory-level reference, while evaluating closure/stability empirically through the explicit phase gates and contract tests in the experiment repository. We do not claim a measured idempotence defect; rather, we treat contract-backed stability gates and persistence measures as operational certificates.

Audits $\mathcal{A}^{(n)}$ (and protocol disclosure). The primary directionality signal is an accepted-move entropy-production proxy (an EP ledger) and its proposal-normalized, per-kernel variants; these are treated as proxies rather than exact path-space KL audits. Protocol holonomy (P3) is implemented as a deterministic kernel schedule keyed to step index and is therefore time-inhomogeneous unless protocol phase is included in state. Accordingly, protocol effects are reported via matched-control comparisons on stroboscopic currents and signatures, and are labeled as diagnostics rather than as standalone arrow-of-time certificates. Higher-level claims (motif inventories, proto-syntax, and intervention-based semantics/decoding) are reported only where they pass explicit phase gates and shift-null controls (see the claim ledger entries beginning with N).

3.3 Operational definitions of reported audit proxies

Because we rely on audit *proxies* rather than the exact path-space KL $\Sigma_T(\rho)$, we state the operational estimators used in both substrates. All windowed quantities below are computed over a sliding window of length T_w on an accepted-move event log; pseudocounts use $\alpha = 1$.

Coarse current–affinity proxy and Σ_{mem} (particles). For each writable memory-variable type $y \in \{w, a, n, S\}$, let N_y^+ and N_y^- denote the counts of accepted +1 and −1 updates in the window. Define the net flux (current)

$$J_y := \frac{N_y^+ - N_y^-}{T_w},$$

and the empirical edge affinity

$$A_y := \log \frac{N_y^+ + \alpha}{N_y^- + \alpha}.$$

The aggregate memory-asymmetry proxy is

$$\Sigma_{\text{mem}} := \sum_y J_y A_y.$$

This proxy is motivated by current–affinity decompositions in Markov-network nonequilibrium theory [Schnakenberg, 1976, Altaner et al., 2012, Seifert, 2012].

M6 two-context loop proxy (particles). Let $c_\mu \in \{H, L\}$ denote a coarse context bin (“high/low μ ” in the particle repo). With counts $N_H^{y+}, N_H^{y-}, N_L^{y+}, N_L^{y-}$ for each y , define the two-context loop-bias estimator

$$\hat{\mathcal{A}}_{M6}(y) := \log \frac{(N_H^{y+} + \alpha)(N_L^{y-} + \alpha)}{(N_H^{y-} + \alpha)(N_L^{y+} + \alpha)}.$$

The reported M6 family (e.g., M6A, M6W, M6N) are summary statistics of $\hat{\mathcal{A}}_{M6}(y)$ across relevant y -types, and are treated as audit proxies (near-zero under null; nonzero under drive).

Accepted-move EP proxies (both substrates). Both repos also report an accepted-move path-asymmetry proxy (“EP proxy”). In the particle substrate this quantity is accumulated as a sum of log acceptance ratios over accepted moves (rejections contribute 0; proposal-ratio terms are treated as symmetric and ignored). In the neural substrate, the per-accepted-move increment is recorded as $\Delta_{\text{ep}} = -\beta(\Delta E - W_6)$ (with $W_6 = 0$ under null and a nonconservative work term W_6 when P6 is enabled), again with rejections contributing 0 and proposal terms treated as symmetric. These are *not* exact KL audits, are not guaranteed nonnegative, and are used comparatively (null vs drive) alongside independent coarse affinity proxies.

3.4 Six-Birds primitive mapping (implementation view)

Below we summarize where each primitive P1–P6 lives in the two substrates and which primitives are active in the core presets/phases referenced in this paper. The mapping is based on the implementation documents and preset files in the external experiment repositories (e.g., `six-birds-particle/docs/knowledge/02_primitives_P1_P6.md`, `six-birds-neural/docs/kernel-spec.md`).

- **P1 (operator rewrite / coupling).** Particles: writable bonds w_{ij} (enabled when `pWrite` > 0). Neural: local token exchange kernels (`w_local`, `k_local`), active in core phases.
- **P2 (gating / budget constraints).** Particles: apparatus counters a_i (enabled when `pAWrite` > 0). Neural: economy diffusion / conserved budgets (`k_neighbor_trade`, B_w , B_k); active in Phase 2, off in Phases 1/3.
- **P3 (protocol holonomy).** Particles: protocol phase/clock (`p30n`). Neural: deterministic kernel cycle when `p3_on` is true; Phase 3 is matched-control P3.
- **P4 (discrete modes / counters).** Particles: quantized counters n_i (enabled when `pNWrite` > 0). Neural: template flip kernel (`n_flip`); off in Phases 1/2/3, on in later/meta presets.
- **P5 (packaging / closure field).** Particles: packaging field S_q (enabled when `pSWrite` > 0). Neural: barrier/closure field s_u with `s_step` (off in Phases 1/2/3, present in kernel spec).
- **P6 (audit / drive).** Particles: drive toggle `p60n` with μ contexts. Neural: work term `eta_drive` with `p6_on`; Phase 2 drive on, Phase 1/3 off.

Substrate	Null regime	Audit proxies (summary)
Particles	P3=OFF, P6=OFF	Retuned null preset: $\Sigma_{\text{mem}} = \text{sigmaMem} \approx 3 \times 10^{-4}$ (mean \pm std across seeds); M6 motif affinities are 0 in the retune summary. Accepted-move EP proxy: $\text{meanExactWindow} \approx 3.5 \times 10^{-6}$ with confidence interval (CI) half-width $\approx 4.7 \times 10^{-5}$ (CI contains 0). <i>Ledger: C_BASE_RETUNE_NULL_1, C_EP_EXACT_NULL_1</i>
Neural	P3=OFF, P6=OFF	Contracted null certificate: $ \text{epMicroRateWindowLast} \leq 2 \times 10^{-4}$ (seeds 1–2 in the contract test). Scale-up report (24 \times 24, three seeds): mean EP $\approx -3.9 \times 10^{-6}$ with max CI half-width $\approx 2.8 \times 10^{-4}$ (CI contains 0; sign not meaningful at this scale); acceptance ≈ 0.54 (non-degenerate). <i>Ledger: N_PH1_NULL_EP_1</i>

Table 1: Null-regime audit calibration (proxy audits). In both substrates, the null regime is tuned so that reported audit proxies are near zero under P3=OFF and P6=OFF, reducing risk of false-positive directionality. Particle values are reported in `six-birds-particle/EXPERIMENTS.md`. Neural values are reported in `PHASE1_NULL_QUICKSELECT_REPORT_v1.md` and the contract `six-birds-neural/tests/contracts/test_contract_null_ep.py`.

4 Results

4.1 Null regime validation

SAFE. Before using protocol (P3) or drive (P6) to argue for directionality or life-like maintenance, we first calibrate a *null regime* in which P3 and P6 are switched off. The goal is to reduce the risk of spurious arrow-of-time claims created by instrumentation, hidden schedules, or coarse-graining artifacts. Because our reported audits are *proxies* (e.g., accepted-move EP and coarse current-affinity summaries), the null regime is defined operationally: the proxies should be statistically consistent with zero while dynamics remain non-degenerate (acceptance not collapsed).

In the particle substrate we additionally monitor coarse per-coordinate currents J_y and affinities A_y (and a motif proxy family) to ensure they relax toward zero in null. In the neural substrate we use both contract-backed thresholds and a scale-up report to verify that the EP proxy remains near zero at larger lattice sizes. Subsequent results are reported as comparisons against this calibrated baseline (null vs driven, or matched-control protocol comparisons), and directionality is attributed to drive/audit structure only when supported by the corresponding proxy audits.

4.2 Separable drive (P6)

SAFE. Having calibrated a clean null baseline (table 1), we next test *drive separability*: switching on P6 (with P3 held OFF) should activate directionality-relevant audit proxies while leaving the null regime unchanged. We emphasize two controls throughout: (i) comparisons are always against a matched null configuration (same substrate and budgets, P3=OFF, P6=OFF), and (ii) audit quantities are treated as *proxies* (coarse current-affinity summaries, motif proxies, and accepted-move EP proxies), so the key requirement is a robust *null-vs-drive separation* rather than an absolute thermodynamic interpretation.

In both substrates, the qualitative pattern is the same: a calibrated null exhibits proxy audits near zero, while switching on P6 produces a robust separation in the reported audit proxies. We treat this as evidence for a controllable, separable “metabolic” channel that can be invoked in later sections when discussing maintenance, repair, and higher-level organization.

Substrate	Condition	Regime	Key metrics
Particles	P3=OFF	Null (P6=OFF)	$\sigma_{\text{Mem}} \approx 3 \times 10^{-4}$; M6 affinities = 0; EP proxy $\approx 3.5 \times 10^{-6}$ (CI contains 0)
		Drive (P6=ON)	$\sigma_{\text{Mem}} \approx 2.1 \times 10^{-3}$; M6 motifs activate (M6W ≈ 0.814 , M6N ≈ 0.374 , M6A ≈ 0.392); EP proxy $\approx 4.26 \times 10^{-3}$
Neural	P3=OFF	Null (P6=OFF)	K-enabled null: $ \text{null_mean} \leq 5 \times 10^{-4}$; non-degenerate acceptance
		Drive (P6=ON)	$\text{drive_mean} \geq 10^{-3}$; $\text{drive_mean} - \text{null_mean} \geq 5 \times 10^{-4}$; preset: $B_k=2$, $\eta_{\text{drive}}=2.0$

Table 2: Drive separability (P6) in both substrates. In both cases, drive (P6=ON) activates audit proxies while null remains near 0, demonstrating clean separability. Particle numbers from `six-birds-particle/EXPERIMENTS.md`; neural thresholds from the Phase 2 contract. *Ledgers*: `C_BASE_RETUNE_P6_1`, `C_P6_SEP_1`, `N_PH2_P6_SEPARABILITY_1`.

Setting	Control L2	Protocol L2	Rel. change	Accept (c)	Accept (p)
seed 1, $\eta=0.5$ (screen)	0.0349	0.0262	0.250	–	–
seed 1, $\eta=1.0$ (screen)	0.0379	0.0271	0.284	–	–
seed 1, $\eta=2.0$ (contract)	0.0311	0.0255	0.179	–	–
seed 1, $\eta=1.0$ (confirm)	0.0379	0.0271	0.284	0.0103	0.0107
seed 2, $\eta=1.0$ (confirm)	0.0326	0.0280	0.141	0.0116	0.0107
seed 3, $\eta=1.0$ (confirm)	0.0313	0.0280	0.105	0.0107	0.0108

Table 3: Protocol holonomy diagnostic (neural substrate): Phase 3 protocol vs matched control changes the stroboscopic current magnitude `strobe_current_l2_window`. The $\eta=2.0$ row is the contracted case; remaining rows show the effect persists across settings. Contract: `test_contract_phase3_p3_protocol_changes_strobe_currents.py`. Ledger: `N_PH3_P3_PROTOCOL_1`.

4.3 Protocol holonomy diagnostics (P3)

SAFE (diagnostic). We next test protocol holonomy (primitive P3): whether a deterministic protocol schedule can measurably change coarse, stroboscopic behavior under a fixed observation lens. This is a *diagnostic* claim, not a directionality certificate: we do not attribute arrow-of-time to protocol noncommutativity alone, and we treat any sustained directionality as requiring an audit signal (e.g., P6-driven separation as in table 2).

Neural substrate: protocol changes stroboscopic currents (matched control). In the neural/meta-layer substrate, P3 is implemented as a deterministic kernel schedule keyed to step index. Phase 3 compares a protocol run against a matched-control baseline (same operating point, P3=OFF), and evaluates a coarse stroboscopic current magnitude `strobe_current_l2_window`. Related issues in stroboscopic observation of Markov networks and the distinction between generator reconstruction and thermodynamic inference are discussed in [Bauer et al. \[2025\]](#). A contract test certifies a nontrivial relative change in this stroboscopic statistic together with basic non-degeneracy conditions (acceptance and strobe signature support). The table below summarizes the reported Phase 3 measurements.

Substrate	Task	Key outcome	Ledger
Particles	Code maintenance	Null: err = 0.43 ± 0.09 , Sdiff = 6.86 ± 0.25 , epRate = 0 Drive: err = 0, Sdiff = 0.83 ± 0.06 , epRate = 0.016	C_CODE_MAINT_1
Particles	Traversal (gated)	Ungated: err = 0.23, Sdiff = 0.49 Gated moving: err = 0, Sdiff = 0 Gated static: Sdiff = 2.58 (no recovery)	C_TRAVERSAL_NEED_1
Particles	Deadline selection	Δ uptimeTail = 0.285, CI [0.14, 0.43] Clock EP cost: Δ EPClockRate = 0.083	C_DEADLINE_DRIFT_1
Neural	Hazard response	spike_paired = 0.063, realloc_paired = 0.015, recovery_paired = 0.060; threshold: realloc ≥ 0.002	N_PH9P5_PAIRIED_HIGHWAY_1

Table 4: Viability/maintenance proxies in both substrates. All tasks are SAFE status. Particle numbers from `six-birds-particle/EXPERIMENTS.md`; neural numbers from `six-birds-neural/docs/EXPERIMENTS.md` (Phase 9.5).

The $\eta = 2.0$ row corresponds to the contracted case. Remaining rows provide screening/confirmation context, demonstrating persistence of the protocol effect across nearby settings.

Protocol-trap handling (what we do and do not claim). Because the P3 schedule is keyed to step index, the observed process is time-inhomogeneous unless protocol phase is included in state. The framework recap (section 2.3) explains why protocol effects can be strong at coarse lenses while still failing to certify directionality under autonomy: if protocol phase is treated as part of the microstate and each phase update is reversible with respect to a common stationary distribution, the resulting lifted chain has zero path-space asymmetry. Accordingly, we interpret table 3 as evidence that P3 is a real *route-dependence/holonomy knob* for coarse behavior, and we attribute sustained directionality only when audit proxies separate cleanly (as under P6 in table 2).

Particle substrate (supportive evidence only). SUGGESTIVE. In the particle substrate, P3 is tracked by holonomy-style observables (`p3LoopArea`, `p3DispMag`) that are intermittently detectable in minimal protocol settings, but are sensitive to protocol design and report cadence. We therefore treat particle P3 evidence as supportive diagnostics rather than a core certificate (ledger: `C_P3_OBS_1`, `C_BASE_RETUNE_P3_1`).

4.4 Viability and maintenance loops

SAFE. With a clean null baseline (table 1) and a separable drive channel (table 2), we can ask a more life-like question: can the system use available resources to *maintain* structured state against noise, perturbations, or externally imposed deadlines? Here we report maintenance/viability proxies in both substrates. These are not claims of exact closure (we do not measure $\delta_{\tau,f}$); rather, they are contract- and control-backed demonstrations of persistence, repair, and resource allocation patterns under perturbation. The hazards, deadlines, and repair targets are externally specified test harnesses used to probe competencies under matched control, not claims of biochemical metabolism or ecological realism.

Phase	Seed	Cov. (pre)	Cov. (haz)	JSD (pre/haz)	JSD (trans)	Prop. score
11 (inventory)	1	1.0	1.0	0.0312	–	0.644
11 (inventory)	2	1.0	1.0	0.0324	–	0.587
11 (inventory)	3	1.0	1.0	0.0316	–	0.609
12 (proto-syntax)	1	1.0	–	0.0312	0.0383	0.644
12 (proto-syntax)	2	1.0	–	0.0324	0.0418	0.587
12 (proto-syntax)	3	1.0	–	0.0316	0.0424	0.609

Table 5: Motif inventory and proto-syntax (neural substrate). Phase 11 gate: coverage ≥ 0.60 , Jensen-Shannon divergence (JSD) ≥ 0.01 , prop. score ≥ 0.02 (all seeds pass). Phase 12 reports a transition-grammar shift via JSD (trans). Ledgers: N_PH11_MOTIF_INVENTORY_1, N_PH12_MOTIF_SYNTAX_1.

SUGGESTIVE. As broader context for the neural hazard response, Phase 9 reports an unguided hazard condition in which two of three seeds pass a spike/recovery/reallocation gate (with best reallocation values on the order of 10^{-2} to 10^{-3}), but because the pass rate is incomplete we treat that phase as suggestive context rather than a core certificate (ledger: N_PH9_HAZARD_HIGHWAYS_1). The paired Phase 9.5 design strengthens interpretability by comparing a hazard fork against a matched baseline fork.

4.5 Higher-level representations (neural substrate)

SAFE. The neural/meta-layer substrate supports refined observational levels in which higher-order regularities are expressed as *lens-level predicate families*. Concretely, we extend the diagnostic lens to include motif histograms, motif-transition summaries, and (under controlled interventions) dictionary and decoding statistics. In the framework language (section 2.2), this is the appropriate sense in which we speak about “novelty/extension”: not as an unqualified claim of open-endedness, but as evidence that additional predicates become robustly usable at refined descriptive levels. All claims in this subsection are backed by explicit phase gates and (where applicable) shift-null p-values, and are indexed in `assets/claims.yml`. In these phases, “shift-null” refers to a within-run control in which intervention labels are temporally shifted (breaking alignment between labels and state trajectories while preserving marginal statistics), and p-values are computed against this shifted-label null rather than against an across-run parametric model. We treat these p-values as phase-local control signals reported by the experiment protocol (often paired with explicit gates) and do not claim global family-wise error control across all phases or metrics in the paper.

Motif inventory and proto-syntax (Phases 11–12). SAFE. Phase 11 certifies that a small motif inventory achieves high coverage and shifts between pre-hazard and hazard contexts, passing the gate `coverage` ≥ 0.60 , `jsd` ≥ 0.01 , and `prop_score` ≥ 0.02 . Operationally, motifs are discrete tokens discovered by the Phase 11 discovery script and summarized as per-window histograms and transition counts; inventory stability is evidenced here by full coverage and consistent divergence/prop-score margins across seeds. Phase 12 extends this to a proto-syntax claim: motif-transition statistics shift between contexts, quantified by a transition-grammar divergence `jsd_trans_pre_hazard`.

Intervention-based semantics and phrase decoding (Phases 17–20). SAFE. Phases 17–20 use controlled token interventions and shift-null tests to show that refined motif summaries support stable decoding and dictionary structure tied to intervention labels. Here “semantics” is used in an operational sense: intervention-conditioned discriminability and decoding of injected labels, not

Phase	Seed	Key metric	Value	p -value	Separation
17 (dir. focus)	1	focus delta	1.478	0.005	spike ctl = 0.269
17 (dir. focus)	2	focus delta	1.453	0.000	spike ctl = 0.225
17 (dir. focus)	3	focus delta	1.521	0.000	spike ctl = 0.384
18 (dict. polarity)	1	dict delta	0.064	0.005	JSD in/out = 0.032
18 (dict. polarity)	2	dict delta	0.108	0.005	JSD in/out = 0.060
18 (dict. polarity)	3	dict delta	0.075	0.005	JSD in/out = 0.039
19 (phrase align.)	1	alignment	0.999	0.005	JSD out/in = 0.665
19 (phrase align.)	2	alignment	1.000	0.005	JSD out/in = 0.671
19 (phrase align.)	3	alignment	1.000	0.005	JSD out/in = 0.667
20 (phrase decode)	1	token acc	1.000	0.005	JSD out/in = 0.665
20 (phrase decode)	2	token acc	1.000	0.005	JSD out/in = 0.671
20 (phrase decode)	3	token acc	1.000	0.005	JSD out/in = 0.667

Table 6: Intervention-based semantics and phrase decoding (neural substrate). Phase 17 gate: focus delta ≥ 0.01 , $p \leq 0.10$; later phases report dictionary and decoding statistics with shift-null p -values. Ledgers: N_PH17_DIRECTIONAL_SEMANTICS_1, N_PH18_DICTIONARY_SEMANTICS_1, N_PH19_PHRASE_SEMANTICS_1, N_PH20_PHRASE_DECODE_1.

linguistic meaning or subjective experience. The interventions inject controlled label patterns into designated channels, and decoding accuracy/alignment are computed by predicting these labels from the resulting motif-level summaries under matched shift-null controls. We report effect sizes together with p -values from the shift-null controls. The Phase 17 gate includes `focus_delta_inject` ≥ 0.01 and `focus_p_inject` ≤ 0.10 ; subsequent phases report dictionary and decoding p -values at the same operating point.

Negative attempts constrain the claim (Phases 14–16). `NEGATIVE`. We explicitly report that earlier “semantics” attempts (Phases 14–16) failed their shift-null gates and are treated as negative constraints rather than as partial successes. This helps delineate what is and is not being claimed: the robust decoding/dictionary results above arise in the intervention-based protocol with explicit controls, not as an automatic consequence of having motifs. Ledger: N_NEG_PH14_16_EARLY_SEM_1.

5 Discussion and limitations

5.1 Mapping back to the three certificates loop

The framework separates three logically distinct certificates—stability/objecthood, novelty/extension, and directionality/audit (section 2.2). The experiments reported here instantiate that separation in two substrates, and the results can be read as a concrete “loop”: (i) packaging/closure mechanisms support persistent descriptive regularities (objecthood proxies), (ii) audits (or audit proxies) certify when a directionality-relevant resource channel is actually active (as opposed to being a protocol artifact), and (iii) refined lenses can make additional predicate families usable and stable (extension at the descriptive level).

Directionality/audit. Across both substrates, we first enforce a conservative null regime in which audit proxies are statistically consistent with zero (table 1). We then show a separable drive channel (P6) that produces clean null-vs-drive separation in reported audit proxies while holding P3 off

(table 2). Protocol holonomy (P3) is treated as a diagnostic knob that changes coarse stroboscopic behavior under matched control (table 3), but it is not used as a directionality certificate.

Stability/objecthood. We do not claim measured closure via an idempotence defect; rather, we report operational objecthood and viability via maintenance proxies under perturbation and selection pressure (table 4). In the particle substrate this appears as successful repair/maintenance under drive, traversal as a causal requirement in gated settings, and a drift advantage under deadlines with an explicit audit cost. In the neural substrate this appears as a paired hazard response with recovery and resource reallocation under a matched baseline fork.

Novelty/extension. Higher-level results in the neural substrate are phrased as refined-lens predicate families rather than as unqualified “emergence”: motif inventories and transition-grammar summaries become robustly usable (table 5), and controlled interventions support dictionary and decoding statistics with explicit shift-null p-values (table 6). We also report negative attempts (Phases 14–16) as constraints, delineating that the later intervention-based decoding results are not automatic consequences of having motifs.

A competency-based interpretation. A complementary way to read these results is in terms of *competencies* that can be expressed at multiple descriptive scales: the ability to maintain constraints, redirect dynamics after perturbation, and reliably use internal degrees of freedom to achieve macroscopic outcomes. In this view, the null calibration (table 1) establishes that apparent directionality is not an instrumentation artifact; the separable drive channel (table 2) supplies a controllable resource whose activation can be audited; and the maintenance results (table 4) show concrete error-correction and persistence under perturbation and deadline selection with an explicit audit cost. At refined lenses, the motif inventory and intervention-based decoding results (tables 5 and 6) can be understood as stable, reusable predicate families that support context-dependent routing and controlled readout. Importantly, we treat these as *lens-relative* competencies rather than intrinsic intentions: they are operationally grounded in matched controls, explicit gates, and conservative proxy audits, and bounded by the negative phases reported.

Life-like competencies do not presuppose replication. A further interpretation suggested by these instantiations is that the particular competencies certified here—audited directionality and operational viability/maintenance, together with lens-relative extensions—can arise in single-system dynamical regimes without replication or population-level Darwinian dynamics. Concretely, our substrates are single-system dynamical regimes with budgets, protocols, and drive channels: they do *not* implement population-level reproduction with heritable variation and selection by differential replication. Nevertheless, we observe (i) clean null calibration and separable audited drive activation (tables 1 and 2), (ii) maintenance and perturbation recovery with explicit costs (table 4), and (iii) refined-lens predicate families that support controlled readout under explicit gates (tables 5 and 6). This supports a modest claim: replication and Darwinian evolution are not prerequisites for the certificate pattern documented here (audited directionality, operational viability/maintenance, and lens-relative extension), even if they may be essential for the long-run accumulation and open-ended diversification of such structure in natural settings. This should not be read as a claim that replication is unnecessary for biological life in general.

Relation to biological “constraint taxonomies”. An independent biological viewpoint argues that much of the apparent contingency of evolution is nevertheless shaped by a small set of

cross-scale constraints, and organizes the discussion around six such constraint families [Solé et al., 2024]. While our “Six Birds” refers to six abstract closure primitives (P1–P6) rather than to biological categories, we observe a heuristic alignment: thermodynamic constraints \leftrightarrow P6 (accounting/audit); linear molecular information \leftrightarrow P1 (operator rewrites mediated by encoded programs); cellular building blocks \leftrightarrow P5 (packaging into stable objects); multicellularity and development \leftrightarrow P3 (internal protocol/phase and route dependence); threshold computation in cognition \leftrightarrow P2 (gating/constraints); and discrete ecosystem architecture \leftrightarrow P4 (sectors/invariants and decompositions into quasi-discrete compartments). We emphasize that this correspondence is interpretive, not a formal identification: Solé et al. aim to summarize universal biological design constraints shaping evolutionary outcomes, whereas we isolate a representation-independent operator vocabulary for theory growth and instantiate a bounded set of competencies in controlled substrates. In particular, key biological pillars in that review (e.g., cellular reproduction and ecological organization) are discussed as constraints on living matter, but are explicitly out of scope for our present instantiations.

5.2 Limitations (what is not established)

This paper is an application/case-study report, not a completeness proof. In particular:

- **No full path-space KL audit.** We do not compute the exact path-space arrow-of-time $\Sigma_T(\rho)$ as a KL divergence between forward and time-reversed trajectory laws. Reported EP quantities are audit *proxies* (e.g., accepted-move logs and coarse current–affinity summaries) and should be interpreted as conservative indicators whose credibility rests on clean null calibration and controlled separability.
- **No measured idempotence defect.** While the canonical completion template $E_{\tau,f}$ is used to organize the instantiations, we do not implement or measure $\delta_{\tau,f}$; stability/objecthood is assessed via explicit retention/maintenance proxies and phase gates rather than by claiming closure as a measured property.
- **No protocol-only arrow-of-time claims.** P3 results are reported as matched-control diagnostics. Because protocol schedules can render the observed process time-inhomogeneous unless phase is included in state, we do not treat protocol noncommutativity as a directionality certificate.
- **No claim of biological life instantiation.** The substrates studied are synthetic computational models; “life-like” refers to operational, lens-relative signatures under explicit protocols and audits, not to the creation or detection of living organisms.
- **No claim of autonomous open-ended evolution.** While the framework allows arbitrarily many theory extensions in the definability sense (adding new lens predicates), we do not claim autonomous open-ended evolution in the ALife sense (indefinite novelty generation under a fixed regime without external steering).
- **No formal definability separation proof.** The refined-lens predicates (motifs, dictionaries, decoding statistics) are empirically validated by gates and shift-null controls, but we do not provide a formal proof that they are not definable from a coarser lens in the strict framework sense.
- **No coarse-graining monotonicity check for proxies.** While the framework motivates audits that behave monotonically under coarse observation, we do not verify any such monotonicity

property for the proxy audits used here; instead we rely on strict null calibration and matched-control separations.

5.3 Why proxy audits are acceptable here, and what would upgrade them

The aim of this paper is to demonstrate a disciplined empirical workflow aligned with the framework: prevent false positives via null calibration, attribute directionality only when an audit channel separates cleanly, and treat higher-level descriptive claims as lens-relative. Proxy audits are therefore acceptable *for this application/case-study paper* because (i) they are explicitly labeled as proxies, (ii) they are calibrated to zero under null, and (iii) the main claims are comparative (null vs drive; protocol vs matched control), not absolute thermodynamic identifications.

Several upgrades would strengthen future instantiations: (i) implement and report an exact $\Sigma_T(\rho)$ on an explicitly specified finite state space (or a controlled coarse-graining) including proposal-ratio terms, (ii) implement an explicit dynamics-induced $E_{\tau,f}$ on distributions together with a computable idempotence defect (or a tractable bound), (iii) treat protocol phase as part of state when making autonomy claims and evaluate directionality on the lifted chain, and (iv) pursue formal definability separation tests (or bounds) for refined predicate families beyond empirical shift-null controls.

5.4 Conclusion

To “wake a stone” here means a transition from a calibrated null to sustained, directed activity with maintenance: directionality appears only when the audit channel separates cleanly, and stability appears only when packaging-level structure persists under perturbation. The six primitives are the minimal, explicit vocabulary for what is switched on and what is left off. The contribution of this paper is to show that this vocabulary can be instantiated and tested in two substrates.

Declarations

Corresponding author. Correspondence to Ioannis Tsiokos (ioannis@automorph.io).

Competing interests. The author declares no competing interests. The author is affiliated with Automorph Inc.; this work was not funded by, nor conducted on behalf of, any external sponsor.

Funding. No external funding was received for this research.

Scope note (non-claims). This paper studies *operational* life-like signatures in synthetic computational substrates under controlled protocols. It does *not* claim the creation, instantiation, or detection of biological life, abiogenesis, or living organisms.

Ethics statement. This work reports computational simulations and analysis only. It involves no human participants, no animals, no personal data, and no biological materials. The manuscript does not claim or attempt to enable the creation of living organisms and is not intended for harmful applications.

Data and code availability. All experiment code, configuration files, and automatically generated evidence tables are publicly available under open-source licenses:

- Particle substrate: <https://github.com/ioannist/six-birds-particle>

- Neural substrate: <https://github.com/ioannist/six-birds-neural>
- This paper and claim ledger: <https://github.com/ioannist/six-birds-life>

Repository archival DOI: <https://doi.org/10.5281/zenodo.18454211>. The machine-readable claim ledger (`assets/claims.yml`) maps every claim ID referenced in this paper to the corresponding external artifacts.

Author contributions. Sole author. I.T. designed the emergence-calculus instantiation, implemented both substrate experiments, performed all analyses, and wrote the manuscript.

AI assistance disclosure. Portions of the experimental code and manuscript text were prepared with assistance from large language models (OpenAI ChatGPT and Claude via Claude Code). The author reviewed, edited, and tested all code and text, retains full responsibility and ownership for the content and interpretations, and lists no AI system as an author.

References

- Bernhard Altaner, Stefan Grosskinsky, Stephan Herminghaus, Lukas Katthän, Marc Timme, and Jürgen Vollmer. Network representations of non-equilibrium steady states: Cycle decompositions, symmetries and dominant paths. *Physical Review E*, 85:041133, 2012. doi: 10.1103/PhysRevE.85.041133.
- M. T. Bauer, U. Seifert, and J. van der Meer. Stroboscopic measurements in markov networks: exact generator reconstruction vs. thermodynamic inference. *Journal of Physics A: Mathematical and Theoretical*, 58:125001, 2025. doi: 10.1088/1751-8121/adbc52.
- Michael Timothy Bennett. *How To Build Conscious Machines*. PhD thesis, The Australian National University, 2025. URL <https://hdl.handle.net/1885/733782452>.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2 edition, 2006. ISBN 978-0-471-24195-9.
- Benedikt Hartl and Michael Levin. What does evolution make? learning in living lineages and machines. *Trends in Genetics*, 41(6):480–496, 2025. doi: 10.1016/j.tig.2025.04.002.
- Benedikt Hartl, Léo Pio-Lopez, Chris Fields, and Michael Levin. Remapping and navigation of an embedding space via error minimization: a fundamental organizational principle of cognition in natural and artificial systems, 2026. URL <https://arxiv.org/abs/2601.14096>.
- Joel L. Lebowitz and Herbert Spohn. A gallavotti–cohen-type symmetry in the large deviation functional for stochastic dynamics. *Journal of Statistical Physics*, 95(1-2):333–365, 1999. doi: 10.1023/A:1004589714161.
- Jürgen Schnakenberg. Network theory of microscopic and macroscopic behavior of master equation systems. *Reviews of Modern Physics*, 48(4):571–585, 1976. doi: 10.1103/RevModPhys.48.571.
- 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.

Ricard Solé, Christopher P. Kempes, Bernat Corominas-Murtra, Manlio De Domenico, Artemy Kolchinsky, Michael Lachmann, Eric Libby, Serguei Saavedra, Eric Smith, and David Wolpert. Fundamental constraints to the logic of living systems. *Interface Focus*, 14(5):20240010, 2024. doi: 10.1098/rsfs.2024.0010.

Ioannis Tsiokos. Six birds: Foundations of emergence calculus, February 2026. URL <https://doi.org/10.48550/arXiv.2602.00134>. arXiv:2602.00134 [cs.LO].