

**First-Principles Validation and Extension of the
S–I–C–T Category-Theoretic Reconstruction:
Functorial Stratification, Error Budgets, and the Limits of
Thermodynamic Obstructions to Ontological Collapse
Revised Technical Edition: Finite Factors and Reproducible Benchmarks**

Autonomous First-Principles Science Panel (source audit)
Roth Complexity Lab (rothcomplexity.org)

Second-order validation of the *First-Principles Validation Report & Extended Academic Paper*
(hereafter **Report A**, 21 October 2026) on M. Róth,
Structural Equivalence Without Ontological Collapse (Working Paper, 21 October 2026)

Revised extension prepared with AI-assisted numerical and proof review;
not externally peer-reviewed

21 October 2026

Abstract

Report A audited Róth’s category-theoretic reconstruction of the S–I–C–T framework (Structure, Information, Cohesion, Transformation) and extended it with a forgetful-functor hierarchy, a quantisation-and-truncation error bound, a computability obstruction and a Landauer thermodynamic obstruction. This paper is a second-order audit: we re-derive every formal claim of Report A from stated axioms, execute its three benchmark protocols numerically, and classify each claim with the Lab’s four-tier epistemic badge system. We confirm the mathematical core: the category axioms, isomorphism criterion, compact-embedding theorem, finite-memory obstruction, forgetful functoriality and the total error bound $E_n \leq L^n \delta_0 + \varepsilon \sum_{k < n} L^k$ are all sound, and the error bound is shown to be *near-sharp* (matching the leading displayed digits) for explicit Euler on the harmonic oscillator once the one-step consistency constant is evaluated on reachable states rather than on the initial energy shell. We correct three substantive overstatements. (i) The Landauer obstruction (Report A, Thm. 4.3) is valid as stated but its prose claims a strictly positive heat cost at *every* step; the erasure terms telescope, so the total logical erasure over any run is bounded by the initial ensemble entropy $H(\mu_0) \leq b$ bits and is identically zero for a single trajectory (Dirac ensemble). The physically relevant dissipation is implementation-level and, by Bennett’s reversibility construction, convertible into a memory cost: the obstruction is a trade-off theorem, not an absolute barrier. (ii) The computability obstruction (Thm. 4.1) is a theorem about the continuum *model* $A_P = [0, 1]^d$, not about physics; treating it as a physical fact would commit the Tier-2→Tier-4 error that Róth’s programme exists to prevent. (iii) The Rule-90 benchmark misapplies Thm. 4.3, since an irreversible *target* does not instantiate the reversible-target/dissipative-simulator asymmetry. The executed benchmarks add two results absent from Report A: the floating-point map of a standard RK4 Lorenz-63 integrator is observed to be *non-injective at the ulp scale* (displacements of 1, 2 and 3 ulp are annihilated within 8, 14 and 34 steps respectively; 4-ulp and 8-ulp displacements collapse onto one and the same trajectory after 191 steps), giving direct empirical content to the non-conjugacy of digital and continuum dynamics; and the “random-walk” rounding-drift model is shown to be scheme-dependent, with a measured *linear* bias drift of 1.44×10^{-17} per step for the exact rotation evaluated in IEEE-754 double precision. We restate the corrected theorems, amend the four-tier equivalence taxonomy (Tier 1 is one-sided approximate simulation, with shadowing as an intermediate tier), and close with recommendations for the S–I–C–T programme.

Version note (21 October 2026). This corrected continuation preserves the attached second-order audit as its starting point but clarifies the physical assumptions behind Landauer’s bound, removes two untraceable Lorenz table rows, distinguishes sampled Jacobian norms from certified Lipschitz bounds, and adds a

mathematically exact finite-factor theorem and an independently executable reversible-target/irreversible-simulator experiment. All additional calculations are provided in `extension_benchmarks.py`. Numerical benchmarks test an implementation, not the hypothesis that the physical universe is a simulation.

Keywords. S–I–C–T; category theory; topological conjugacy; forgetful functors; approximate simulation; shadowing; Landauer principle; Bennett reversibility; structural realism; epistemic badges.

1 Scope, method and epistemic badges

What was audited. The Panel received Report A only. Róth’s underlying manuscript was not supplied, so Róth’s theorems are validated *as restated in Report A* (its Part I table and Theorem 3.1). Where Report A’s rendering is internally sufficient to reconstruct a proof we do so; where it is not, we say so. Several displayed formulas in Report A were lost in transmission (Definitions 2.1 and 2.3, Theorems 3.2 and 4.3, Definition 4.2, Benchmarks 2 and 3); we reconstruct each from the surrounding text and flag the reconstruction.

Panel constitution. Five specialist roles were instantiated— $\mathcal{A}_1$ category theory and topology, $\mathcal{A}_2$ dynamical systems and numerical analysis, $\mathcal{A}_3$ information theory and physics of computation, $\mathcal{A}_4$ philosophy of science, $\mathcal{A}_5$ metrology and falsifiability. The roles were executed as role-constrained reasoning passes of a large language model (Claude, Anthropic), each permitted to argue only from stated axioms and standard theorems with named sources. Every numerical statement in this paper was produced by executing code (Python 3, NumPy 2.5) in the Panel’s sandbox; parameters are in Appendix A and the script accompanies the paper. No claim rests on authority, including the authority of Report A or of Róth.

Epistemic badges. Following the Lab’s convention, each claim carries one of four badges: **ESTABLISHED** proved from stated axioms or standard theorems with full proof available; **GROUNDDED** proved under an explicitly named modelling assumption, or verified numerically within stated tolerance; **HEURISTIC** a widely used rule of thumb with known exceptions; **SPECULATIVE** a conjecture or interpretation not supported by proof or measurement.

2 Part I — Clause-by-clause audit of Report A

2.1 Róth’s theorems as restated in Report A

Claim	Panel verification and refinement	Verdict
Prop. 3: $\mathbf{SICT}_{M,J}$ is a category	Identities are continuous and trivially equivariant; composition of continuous equivariant observable-preserving maps is again such a map; associativity is inherited from Top . Agreed. Refinement: Report A is right that morphisms ignore (X, V, Ω, a_0, μ) ; this is formalised in §3.2.	VALID ESTABLISHED
Thm. 4: $\text{iso} \iff$ intertwining homeomorphism	An inverse in $\mathbf{SICT}_{M,J}$ is in particular a continuous inverse in Top ; conversely conjugating $h \circ T_{P,m} = T_{S,m} \circ h$ by h^{-1} gives equivariance of h^{-1} , and $q_j^P = q_j^S \circ h$ gives $q_j^P \circ h^{-1} = q_j^S$. Agreed.	VALID ESTABLISHED
Thm. 5: compact embedding	A continuous map from a compact space to a Hausdorff space is closed, so a continuous injection is a homeomorphism onto its image $Z = h(A_P)$, which is compact and closed. Z is forward-invariant because $T_{S,m}(h(a)) = h(T_{P,m}(a)) \in Z$. Agreed. Report A’s physical caveat (phase spaces of unbounded systems are only locally compact; restrict to energy shells) is correct.	VALID ESTABLISHED

Claim	Panel verification and refinement	Verdict
Prop. 7: finite memory	Pigeonhole / Cantor: no injection from an infinite set into a set of size $\leq 2^b$. The topological variant (a continuous map from a connected space to a discrete space is constant) is also correct.	VALID ESTABLISHED
Cor. 8: observational invariance without material identity	Correct conclusion, but Report A’s justification—“properties invariant under isomorphism are <i>strictly</i> those expressible in the language of $\mathcal{C}$ ”—asserts a biconditional of which only one direction is a theorem (expressible $\Rightarrow$ invariant; Makkai’s principle of isomorphism, cf. Ahrens–North). The converse fails in general. The conclusion follows from the one true direction.	VALID, refined ESTABLISHED
Thm. 12 / Cor. 13: Lipschitz error propagation	$e_{n+1} \leq \varepsilon + Le_n$ by the triangle inequality and the one-step consistency bound, hence $e_n \leq \varepsilon \sum_{k < n} L^k$. Agreed, provided the consistency bound holds on <i>all reachable</i> simulator states—a hypothesis that is easy to violate in practice (Finding 4.2). Report A’s observation that $e_0 = 0$ ignores the quantisation of x_0 is correct and repaired in its Thm. 3.2.	VALID ESTABLISHED

2.2 Report A’s own extensions

Claim	Panel verification and refinement	Verdict
Def. 2.1 (enriched system)	Reconstructed in Def. 3.1. Two clarifications are needed: constraint predicates $c_k : \Omega \rightarrow \{0, 1\}$ must <i>not</i> be required continuous into the discrete two-point space, otherwise A is clopen in Ω and a connected Ω admits only $A = \Omega$ or $A = \emptyset$ (energy shells $\{H \leq E\}$ are closed but not open); and the measure clause is ambiguous between an invariant μ and an evolving family $\mu_m = (T_m)_*\mu$. We adopt an invariant μ as part of the object and treat non-invariant ensembles as trajectories in the space of objects.	VALID, refined ESTABLISHED
Def. 2.3 / Prop. 2.4 (full category, forgetful functor)	Functoriality and well-definedness are correct. The non-injectivity example is correct, but its stated reason “ $H(\mu_1) \neq H(\mu_2)$ ” is ill-posed (Lebesgue measure on $[0, 1]$ has differential entropy 0; a Dirac measure has no finite differential entropy). The correct reason is simply $(\text{id})_*\lambda = \lambda \neq \delta_{1/2}$, so measure compatibility fails. Non-faithfulness is correct (constant-field systems admit many spatial homeomorphisms over the same state map). Non-fullness is correct (the example’s $\text{id}_{[0,1]}$ does not lift). Design caveat: requiring ϕ_X to be a homeomorphism for <i>every</i> morphism forbids spatial coarse-graining morphisms; see Remark 3.6.	VALID, refined ESTABLISHED
Thm. 3.2 (quantisation + truncation bound)	Reconstructed and re-proved as Thm. 3.8. Correct. We add the three regimes $L < 1$, $L = 1$, $L > 1$, make the reachable-set hypothesis explicit, and show sharpness numerically (Finding 4.2).	VALID ESTABLISHED
Rem. 3.3 (symplectic integrators)	Correct statement of backward error analysis (Hairer–Lubich–Wanner, Ch. IX): bounded $O(\Delta t^p)$ energy error over exponentially long times for analytic Hamiltonians. Verified for Störmer–Verlet on the oscillator: maximal energy deviation $1.25 \times 10^{-5} = \Delta t^2/8$ at $n = 10^4, 10^5, 10^6$ (Table 4).	VALID ESTABLISHED
Thm. 4.1 (computability / measure obstruction)	Mathematically correct: $\{0, 1\}^*$ and the computable sequences are countable, any countable subset of $[0, 1]^d$ is Lebesgue-null. <i>But</i> the theorem presupposes that A_P is a continuum. That is a modelling choice about the target, not a first principle of physics (finite-entropy bounds on bounded regions are a live hypothesis). The statement must be badged as conditional on the continuum model.	VALID, refined GROUNDING

Claim	Panel verification and refinement	Verdict
Def. 4.2 / Thm. 4.3 (Landauer obstruction)	The entropy identity $\Delta H_{\text{erase}} = H(\mu) - H(G_*\mu) = H(Z G(Z))$ and its positivity under non-injectivity on the support are correct. Non-isomorphism follows because injectivity of T_1 is an isomorphism invariant. The prose conclusion—“every simulation step must dissipate a strictly positive minimum heat”—does not follow: the per-step erasures telescope to $H(\mu_0) - H(\mu_N) \leq H(\mu_0) \leq b$, the per-step erasure vanishes once the ensemble concentrates on cycles (observed: Table 3), and it is identically zero for a Dirac ensemble. See Thm. 3.13 and Remark 3.14.	VALID as stated; prose overclaims ESTABLISHED
Benchmark 1 (shift vs Rule 90)	Shift conjugacy verified exhaustively for $N = 6, 7, 8$. Rule 90 multiplicities verified (4-to-1 for even N , 2-to-1 for odd N). But replacing the <i>target</i> dynamics by Rule 90 does not “trigger Theorem 4.3”: that theorem requires a reversible target and a non-injective simulator. An irreversible target has its own intrinsic erasure, so no asymmetry arises. The example correctly illustrates Def. 4.2 applied to a target. “CNOT-like” is a misnomer for a cyclic shift (harmless).	MISAPPLIED (theorem); VALID (arithmetic)
Benchmark 2 (oscillator table)	Analytic drift formulas for explicit and implicit Euler verified to 10^{-13} relative accuracy. The column “satisfies Theorem 3.2 bound” is misleading: <i>all</i> consistent schemes satisfy the bound when ε is taken on reachable states; the discriminator between schemes is Cohesion (invariant drift), not the Transformation bound. The “ $O(\sqrt{n} \varepsilon_{\text{mach}})$ random-walk” rounding drift is a heuristic (Brouwer’s law) and scheme-dependent: we measure a <i>linear</i> bias drift for the exact rotation in double precision and no visible secular rounding drift for Verlet in single precision over 10^6 steps.	VALID, refined GROUNDED HEURISTIC
Benchmark 3 (Lorenz-63 horizon)	The n_{crit} formula was lost; reconstructed as $n_{\text{crit}} \approx \ln(1/\max(\delta_0, \varepsilon))/\ln L$. Logarithmic scaling observed. The sampled $\ Df\ _2$ maximum of 30.3 is not a certified upper Lipschitz bound. A horizon computed using it is illustrative, not mathematically guaranteed. Lyapunov-rate estimates describe typical separation, rather than the theorem’s worst-case bound. Shadowing deserves its own tier (§5).	VALID, refined GROUNDED
Four-tier taxonomy (§6)	Hierarchy is well-ordered (each tier implies the one below). Tier 1 as stated is a one-sided approximate <i>simulation</i> ; Girard–Pappas approximate <i>bisimulation</i> is symmetric. Tier 2 preserves more than “orbit structure”: topological entropy, periodic-orbit counts, transitivity and minimality are conjugacy invariants, while measure-theoretic entropy is not.	VALID, refined ESTABLISHED

3 Part II — Validated and corrected formal theory

Throughout, $(M, \cdot, e)$ is a topological monoid ($\mathbb{N}_0$ or $\mathbb{R}_{\geq 0}$), J an index set, and $\{Q_j\}_{j \in J}$ Hausdorff spaces of observable values.

3.1 Objects and the two categories

Definition 3.1 (Enriched S–I–C–T system; reconstruction of Report A Def. 2.1). An enriched system is a tuple

$$\mathcal{U} = (X, V, \Omega, A, (q_j)_{j \in J}, T, \mathcal{B}_A, \mu)$$

such that:

- (S) $X \neq \emptyset$ and V are topological spaces and $\Omega \subseteq V^X$ carries a specified topology;
- (C) $A = \bigcap_{k \in K} c_k^{-1}(1) \subseteq \Omega$ for a family of (not necessarily continuous) predicates $c_k : \Omega \rightarrow \{0, 1\}$, $A \neq \emptyset$ with the subspace topology, and each $q_j : A \rightarrow Q_j$ is continuous;
- (I) $\mathcal{B}_A$ is the Borel σ -algebra of A and μ a Borel probability measure on A (a deterministic actual state a_0 is $\mu = \delta_{a_0}$);
- (T) $T : M \times A \rightarrow A$ is a jointly continuous monoid action, $T_e = \text{id}_A$, $T_{m \cdot n} = T_m \circ T_n$, with

$$q_j \circ T_m = q_j \quad \text{and} \quad (T_m)_* \mu = \mu \quad \text{for all } m \in M, j \in J.$$

Remark 3.2. The continuity caveat on c_k is not pedantry: a continuous map into the discrete space $\{0, 1\}$ from a connected Ω is constant. Admissibility predicates are Borel, not continuous; continuity is reserved for the observables q_j whose level sets Cohesion is meant to preserve. The invariance clause $(T_m)_* \mu = \mu$ replaces Report A’s ambiguous “evolving the informational measure via pushforward”; an ensemble that is not invariant is a curve $m \mapsto \mathcal{U}_m$ of objects, not a single object.

Definition 3.3 (The categories $\mathbf{SICT}_{M,J}$ and $\mathbf{SICT}_{M,J}^{\text{full}}$; reconstruction of Report A Def. 2.3). 1. In $\mathbf{SICT}_{M,J}$ (Róth) an object is a tuple $U = (X, V, \Omega, A, (q_j), T)$; a morphism $h : U \rightarrow W$ is a continuous map $h : A_U \rightarrow A_W$ with

$$h \circ T_{U,m} = T_{W,m} \circ h, \quad q_j^W \circ h = q_j^U \quad (m \in M, j \in J).$$

2. In $\mathbf{SICT}_{M,J}^{\text{full}}$ objects are enriched systems and a morphism $\Phi = (\phi_X, \phi_V, h) : \mathcal{U} \rightarrow \mathcal{W}$ consists of a homeomorphism $\phi_X : X_U \rightarrow X_W$, a continuous $\phi_V : V_U \rightarrow V_W$ and a continuous $h : A_U \rightarrow A_W$ with
- (i) $h(\sigma)(\phi_X(x)) = \phi_V(\sigma(x))$ for all $\sigma \in A_U, x \in X_U$;
 - (ii) $h_* \mu_U = \mu_W$;
 - (iii) $h \circ T_{U,m} = T_{W,m} \circ h$ and $q_j^W \circ h = q_j^U$.

3.2 The forgetful functor and substrate non-uniqueness

Proposition 3.4 (Report A Prop. 2.4, corrected proof). **ESTABLISHED** Both $\mathbf{SICT}_{M,J}$ and $\mathbf{SICT}_{M,J}^{\text{full}}$ are categories, and

$$\mathfrak{U} : \mathbf{SICT}_{M,J}^{\text{full}} \rightarrow \mathbf{SICT}_{M,J}, \quad \mathcal{U} \mapsto (X, V, \Omega, A, (q_j), T), \quad (\phi_X, \phi_V, h) \mapsto h$$

is a functor that is neither full nor faithful and is not injective on isomorphism classes.

Proof. Categories: identities satisfy (i)–(iii) trivially; if $\Phi = (\phi_X, \phi_V, h)$ and $\Psi = (\psi_X, \psi_V, g)$ are composable then $(g \circ h)_* \mu_U = g_*(h_* \mu_U) = g_* \mu_V = \mu_W$, and $(g(h(\sigma)))(\psi_X \phi_X x) = \psi_V(h(\sigma)(\phi_X x)) = \psi_V \phi_V(\sigma(x))$; equivariance and observable preservation compose by associativity. Functoriality of $\mathfrak{U}$ is immediate.

Not injective on iso-classes. Let $A = [0, 1]$, $T_m = \text{id}_A$, no observables. $\mathcal{U}_1$: $X_1 = \{*\}$, $V_1 = [0, 1]$, $\Omega_1 = A$, $\mu_1 = \lambda$ (Lebesgue). $\mathcal{U}_2$: $X_2 = S^1$, $V_2 = [0, 1]$, Ω_2 the constant fields $\cong [0, 1]$, $\mu_2 = \delta_{1/2}$. Then $\mathfrak{U}\mathcal{U}_1 \cong \mathfrak{U}\mathcal{U}_2$ via $h = \text{id}$. Any $\mathbf{SICT}_{M,J}^{\text{full}}$ -isomorphism would need $\phi_X : \{*\} \rightarrow S^1$ bijective, impossible; independently, (ii) fails for every h , since $h_* \lambda$ is non-atomic while $\delta_{1/2}$ is atomic. (Report A’s appeal to “ $H(\mu_1) \neq H(\mu_2)$ ” is unnecessary and ill-defined for a Dirac measure.)

Not faithful. On $\mathcal{U}_2$ take $h = \text{id}$ and $\phi_V = \text{id}$; every rotation ϕ_X of S^1 satisfies (i) because all fields are constant, so distinct Φ have the same image. *Not full.* $h = \text{id}_{[0,1]} : \mathfrak{U}\mathcal{U}_1 \rightarrow \mathfrak{U}\mathcal{U}_2$ is a morphism in $\mathbf{SICT}_{M,J}$ with no preimage, by (ii). $\square$

Corollary 3.5 (Algebraic form of Róth’s Cor. 8). **ESTABLISHED** An isomorphism in $\mathbf{SICT}_{M,J}$ is an isomorphism of images under $\mathfrak{U}$. Since $\mathfrak{U}$ does not reflect isomorphisms, no conclusion about spatial carriers, ensembles, or substrate follows. Properties expressible in the language of $\mathbf{SICT}_{M,J}$ (orbit structure, topological entropy, periodic-orbit counts, level sets of q_j) are invariant; properties outside that language need not be.

Remark 3.6 (Coarse-graining). Requiring ϕ_X to be a homeomorphism for every morphism makes $\mathbf{SICT}_{M,J}^{\text{full}}$ unable to express spatial coarse-graining (X_W a quotient of X_U) or renormalisation, which the S–I–C–T programme will need. A continuous surjection ϕ_X keeps (i) well-posed; a general continuous ϕ_X leaves $h(\sigma)$ unconstrained off $\phi_X(X_U)$ and requires an additional clause. We recommend the surjective variant as the working definition and note that Proposition 3.4 holds verbatim for it.

3.3 Exact equivalence

Theorem 3.7 (Isomorphism criterion and compact embedding; Róth Thms. 4–5 as restated). **ESTABLISHED**

1. $U_P \cong U_S$ in $\mathbf{SICT}_{M,J}$ iff there is a homeomorphism $h : A_P \rightarrow A_S$ with $h \circ T_{P,m} = T_{S,m} \circ h$ and $q_j^S \circ h = q_j^P$ for all m, j .
2. If A_P is compact, A_S Hausdorff and $h : U_P \rightarrow U_S$ an injective morphism, then $Z = h(A_P)$ is compact, closed and T_S -forward-invariant, and $h : U_P \rightarrow U_S|_Z$ is an isomorphism.

Proofs are as in Part I. Note that $U_S|_Z$ is the restricted system $(Z, T_S|_{M \times Z}, (q_j^S|_Z))$; the “tautological codomain restriction” that Róth criticises is exactly the move of passing to $U_S|_Z$ *without* first proving closedness and forward invariance—which here are theorems, not stipulations.

3.4 The error budget theorem

Theorem 3.8 (Total simulation error; Report A Thm. 3.2, reconstructed). **ESTABLISHED** Let (A_P, d) be a metric space, $F : A_P \rightarrow A_P$ with $d(Fx, Fy) \leq Ld(x, y)$, and let $E : A_P \rightarrow A_S$, $G : A_S \rightarrow A_S$, $D : A_S \rightarrow A_P$ be encoder, update and decoder. Fix $x_0 \in A_P$, $z_0 = E(x_0)$, and let $\mathcal{R} = \{G^k z_0 : k \geq 0\}$ be the reachable set. Assume

$$(Q) \quad d(D(E(x_0)), x_0) \leq \delta_0, \quad (C) \quad d(D(Gz), F(Dz)) \leq \varepsilon \quad \text{for all } z \in \mathcal{R}.$$

Then $E_n := d(D(G^n z_0), F^n x_0)$ satisfies

$$E_n \leq L^n \delta_0 + \varepsilon S_n(L), \quad S_n(L) = \sum_{k=0}^{n-1} L^k = \begin{cases} \frac{L^n - 1}{L - 1}, & L \neq 1, \\ n, & L = 1, \end{cases}$$

and for any K -Lipschitz conserved q ($q \circ F = q$), $|q(D(G^n z_0)) - q(x_0)| \leq K(L^n \delta_0 + \varepsilon S_n(L))$.

Proof. Insert $F^n(Dz_0)$: $E_n \leq d(DG^n z_0, F^n Dz_0) + d(F^n Dz_0, F^n x_0)$. The first term is Róth’s $e_n \leq \varepsilon S_n(L)$ by induction on $e_{k+1} \leq d(DG^{k+1} z_0, FDG^k z_0) + d(FDG^k z_0, F^{k+1} Dz_0) \leq \varepsilon + Le_k$, where (C) is applied at $G^k z_0 \in \mathcal{R}$. The second term is $\leq L^n d(Dz_0, x_0) \leq L^n \delta_0$ by (Q). The invariant statement follows from $q(F^n x_0) = q(x_0)$ and Lipschitz continuity. $\square$

Remark 3.9 (Three regimes). **ESTABLISHED** $L < 1$: $E_n \leq \delta_0 + \varepsilon/(1 - L)$ uniformly in n (contractive targets are simulable forever to fixed tolerance). $L = 1$: linear growth $n\varepsilon + \delta_0$ (isometries; the oscillator). $L > 1$: exponential growth with horizon $n_{\text{crit}} \approx \ln(1/\max(\delta_0, \varepsilon(L - 1)^{-1}))/\ln L$, i.e. logarithmic in precision. The hypothesis “for all $z \in \mathcal{R}$ ” is load-bearing: Finding 4.2 shows that evaluating ε on the initial energy shell instead of on $\mathcal{R}$ gives a bound that the data violate, while the correct ε gives equality to six digits.

3.5 The cardinality and computability obstruction

Theorem 3.10 (Report A Thm. 4.1, badged). **GROUND** Assume the continuum model: $A_P = [0, 1]^d$ with Lebesgue measure λ . Let A_S be any set of finite binary strings, or the set A_S^{comp} of Turing-computable sequences. Then no bijection $A_P \rightarrow A_S$ exists, and for any injective encoding $E : A_{\text{rep}} \rightarrow A_S^{\text{comp}}$ on $A_{\text{rep}} \subseteq A_P$, $\lambda(A_P \setminus A_{\text{rep}}) = 1$.

Proof. Both target sets are countable; a countable subset of $[0, 1]^d$ is λ -null; $|[0, 1]^d| = 2^{\aleph_0} > \aleph_0$. $\square$

Remark 3.11 (Why the badge is GROUNDED, not ESTABLISHED). The proof is elementary; the badge records the conditionality of the *premise*. That the physical state space is a continuum is a property of the model $\mathcal{U}_P$, not an axiom of physics; finite-entropy bounds for bounded regions are a serious hypothesis under which $|A_P| < \infty$ and the obstruction evaporates (indeed Róth’s Prop. 7 could then run in the opposite direction). Report A describes the obstruction as holding “on strictly first-principles grounds”; it holds on strictly first-principles grounds *about the continuum model*. To promote it to a statement about the physical target is precisely the Tier 2→Tier 4 inference Róth’s programme forbids. The Panel regards this as the most important interpretive correction in this paper.

3.6 The erasure budget theorem (corrected Landauer obstruction)

Definition 3.12 (Logical erasure; reconstruction of Report A Def. 4.2). For a finite A_S , a map $G : A_S \rightarrow A_S$ and a distribution μ on A_S ,

$$\Delta H_{\text{erase}}(\mu, G) := H(\mu) - H(G_*\mu) = H_\mu(Z | G(Z)) \quad (\text{bits}),$$

the second equality by the chain rule $H(Z, G(Z)) = H(Z) = H(G(Z)) + H(Z | G(Z))$, since $G(Z)$ is a function of Z .

Theorem 3.13 (Erasure budget). **ESTABLISHED** Let $G : A_S \rightarrow A_S$ with $|A_S| \leq 2^b$, μ_0 a distribution on A_S , $\mu_{n+1} = G_*\mu_n$. Then:

1. $\Delta H_{\text{erase}}(\mu_n, G) \geq 0$, with equality iff G is injective on $\text{supp } \mu_n$.
2. (Telescoping) $\sum_{n=0}^{N-1} \Delta H_{\text{erase}}(\mu_n, G) = H(\mu_0) - H(\mu_N) \leq H(\mu_0) \leq b$ for every N .
3. $\Delta H_{\text{erase}}(\mu_n, G) \rightarrow 0$; indeed it is eventually 0, since $\text{supp } \mu_n$ is eventually contained in the set of periodic points of G , on which G is a bijection.
4. If $\mu_0 = \delta_{z_0}$ then $\Delta H_{\text{erase}}(\mu_n, G) = 0$ for all n .
5. (Conditional Landauer benchmark) Under an isothermal implementation with equal logical-state free energies and no ancillary information record retained, the familiar ideal-erasure lower bound is $Q_{\text{env}} \geq k_B T_{\text{env}} \ln 2 [H(\mu_0) - H(\mu_N)]$. The bracketed logical-entropy decrease is at most b bits; this does not upper-bound the actual heat, nor does a Dirac distribution imply zero device dissipation. With side information or unequal state free energies, the implementation-level thermodynamic accounting must be revisited.
6. If the target $T_{P,1}$ is a bijection and G is not injective on $\text{supp } \mu_n$ for some n , then $U_P \not\cong U_S$ in $\mathbf{SICT}_{M,J}$ (injectivity of T_1 is a conjugacy invariant).

Proof. (1) is the non-negativity of conditional entropy, with equality iff Z is determined by $G(Z)$ on the support. (2) is a telescoping sum. (3) In a finite set every forward orbit is eventually periodic, so $\text{supp } \mu_n \subseteq G^n(A_S)$ is eventually contained in the eventual image $\bigcap_k G^k(A_S)$, which is the set of periodic points, on which G restricts to a permutation. (4) $H(\delta_z) = 0$ and $G_*\delta_z = \delta_{Gz}$. (5) invokes an additional physical implementation assumption and is not a purely information-theoretic consequence of (1)–(4). (6) If h were an isomorphism then $G = h \circ T_{P,1} \circ h^{-1}$ would be a bijection. $\square$

Remark 3.14 (What survives of Report A’s Theorem 4.3, and what does not). Report A’s *statement* (positivity at a step where G merges support points; non-isomorphism) is item (1) and item (6) and is correct. Its *prose*—“every simulation step must dissipate a strictly positive minimum heat”, “ $\dot{Q} \geq \dots$ ” as a rate, “irreversible entropy production in the host substrate” as a standing feature of simulating conservative dynamics—does not follow, by items (2)–(4). Two further points relocate the physics:

- Real hardware dissipation is *implementation-level*: the arithmetic that realises G may dissipate far more than the ideal state-level entropy-loss term $H(Z | G(Z))$, including for a Dirac ensemble. Theorem 3.13 does not bound that cost; it bounds only the part attributable to the state map.
- By Bennett (1973), any computation of G can be embedded in a logically reversible computation at the price of retaining auxiliary “garbage” bits, with no Landauer cost in the limit of slow operation; the garbage can be uncomputed only by retaining the input. This introduces a dissipation–memory trade-off

at the implementation level rather than a universal per-step heating law. Separately, a finite register cannot bijectively encode a continuum of target states; exact reversible finite factors remain possible under a non-injective observation map.

Non-injectivity of floating-point updates is nevertheless real and large: Finding 4.6 exhibits it directly.

4 Part III — Executed benchmarks

Original runs: Python 3 / NumPy 2.5 as reported by the source audit; independently repeated here in Python 3.13 / NumPy 2.3.5, IEEE-754 binary64 unless stated; parameters in Appendix A. Floating-point ulp merger times can depend on architecture/library versions.

4.1 Benchmark 1: exact finite conjugacy and erasure

Table 3: Shift conjugacy (exhaustive over all 2^N states) and Rule 90 erasure under the uniform initial distribution (H in bits).

N	shift equivariant	q_{mag} preserved	Rule 90 image/ 2^N	multiplicity	$H(\mu_n)$, $n = 0, 1, 2, \dots$	total erased
6	yes	yes	16 / 64	4-to-1	6, 4, 4, 4, $\dots$	2
7	yes	yes	64 / 128	2-to-1	7, 6, 6, 6, $\dots$	1
8	yes	yes	64 / 256	4-to-1	8, 6, 4, 2, 0, 0, $\dots$	8

Finding 4.1. GROUNDED The affine encoding $h(\sigma)_i = (\sigma_i + 1)/2$ passes the exhaustive isomorphism test for $N = 6, 7, 8$, confirming Report A. Rule 90 on a ring is 4-to-1 for even N and 2-to-1 for odd N , confirming the stated multiplicities. The erasure sequence confirms Theorem 3.13(2)–(3) exactly: for $N = 6$ and $N = 7$ the only erasure occurs at step 1 (the image is a union of cycles on which Rule 90 is a bijection), and for $N = 8$ (a power of two, where Rule 90 is nilpotent) 2 bits are erased at each of the first four steps and none thereafter, totalling $8 = b$ bits. Report A’s “2 bits/step” is thus correct for at most $N/2$ steps and false as a standing rate.

4.2 Benchmark 2: harmonic oscillator

Target: $F_{\Delta t}(x, p) = (x \cos \Delta t + p \sin \Delta t, -x \sin \Delta t + p \cos \Delta t)$, an isometry ($L = 1$), $\Delta t = 0.01$, $x_0 = 1$, $p_0 = 0$, $E_0 = \frac{1}{2}$. One-step discrepancy ε measured as $\max_{|z|=1} |G(z) - F(z)|$.

Finding 4.2 (Sharpness of Theorem 3.8 and the reachable-set hypothesis). GROUNDED With ε evaluated on the unit energy shell, the bound $n\varepsilon$ gives 0.500 at $n = 10^4$ and 5.00 at $n = 10^5$, but the measured Euler trajectory errors are 0.6487 and 147.4: the bound is *violated*. The reason is that hypothesis (C) fails: Euler’s reachable set spirals outward, and for a linear scheme $\varepsilon(z) = 5 \times 10^{-5}|z|$ grows with the state. Replacing $n\varepsilon$ by $\sum_{k < n} \varepsilon(G^k z_0)$ —the bound of Theorem 3.8 with (C) evaluated on $\mathcal{R}$ —gives 0.648694 versus measured 0.648694 ($n = 10^4$) and 147.379 versus 147.377 ($n = 10^5$): equality to the displayed digits. The theorem is sharp here because the Euler error is purely radial while the phase error per step, $\Delta t^3/3$, is negligible. The lesson for S–I–C–T is structural: a Cohesion failure (energy drift) is what invalidates the *Transformation* bound, because it pushes the trajectory into states where the consistency constant is larger.

Finding 4.3 (Verlet). GROUNDED Verlet’s trajectory error grows linearly at 4.17×10^{-8} per step, matching the phase error $\Delta t^3/24$ of its modified frequency $\omega_{\Delta t} = \frac{2}{\Delta t} \arcsin \frac{\Delta t}{2}$, and stays a factor 4 inside $n\varepsilon$. Its energy error is bounded by $E_0 \Delta t^2/4 = \Delta t^2/8$ at all three horizons, confirming backward error analysis. Cohesion (C) and Transformation (T) fidelity are therefore *independent* diagnostics, as Report A’s Remark 3.3 argues.

Finding 4.4 (Rounding drift is scheme-dependent). GROUNDED / HEURISTIC The exact rotation evaluated in binary64 shows energy drift of -1.43×10^{-13} , -1.45×10^{-12} , -1.44×10^{-11} at $n = 10^4, 10^5, 10^6$: *linear* in n , i.e. a systematic bias of $\approx 1.44 \times 10^{-17} \approx 0.065 \varepsilon_{\text{mach}}$ per step, not a $\sqrt{n}$ random walk. Verlet

Table 4: Measured one-step discrepancy, energy error and trajectory error. Analytic energy errors: explicit Euler $((1 + \Delta t^2)^n - 1)E_0$; implicit Euler $((1 + \Delta t^2)^{-n} - 1)E_0$; Verlet bounded by $E_0\Delta t^2/4 = \Delta t^2/8$.

Scheme	ε (unit shell)	n	$H(z_n) - E_0$	analytic	$ z_n - F^n x_0 $
Explicit Euler	$5.000 \times 10^{-5} = \Delta t^2/2$	10^4	+0.8591	+0.8591	0.6487
		10^5	$+1.101 \times 10^4$	$+1.101 \times 10^4$	147.4
		10^6	$+1.34 \times 10^{43}$	$+1.34 \times 10^{43}$	5.2×10^{21}
Implicit Euler	5.000×10^{-5}	10^4	-0.3161	-0.3161	0.393
		10^5	-0.49998	-0.49998	0.993
		10^6	-0.5	-0.5	1.000
Störmer–Verlet	$1.667 \times 10^{-7} = \Delta t^3/6$	10^4	1.25×10^{-5} (max)	$\leq \Delta t^2/8$	4.22×10^{-4}
		10^5	1.25×10^{-5}		4.16×10^{-3}
		10^6	1.25×10^{-5}		4.17×10^{-2}
Exact rotation (fp64)	0 (analytic)	10^4	-1.43×10^{-13}	—	1.43×10^{-13}
		10^5	-1.45×10^{-12}		1.45×10^{-12}
		10^6	-1.44×10^{-11}		1.44×10^{-11}
Verlet (fp32)		$10^{3\dots 6}$	$(-4.0, -6.2, -9.7, +0.8) \times 10^{-6}$	$< \Delta t^2/8$	—

in binary32 shows no secular rounding drift distinguishable from its $\Delta t^2/8$ oscillation over 10^6 steps. Report A’s “random-walk rounding drift $O(\sqrt{n}\varepsilon_{\text{mach}})$ ” is Brouwer’s-law heuristics and must be badged **HEURISTIC**; whether rounding errors are unbiased is a property of the particular arithmetic realisation of G , which is exactly the implementation-level structure that **SICT** _{M,J} forgets.

4.3 Benchmark 3: Lorenz-63

Target: $\dot{x} = \sigma(y - x)$, $\dot{y} = x(\rho - z) - y$, $\dot{z} = xy - \beta z$ with $(\sigma, \rho, \beta) = (10, 28, 8/3)$; simulator: classical RK4 with $h = 10^{-3}$; reference state u_{ref} taken after a transient of $t = 50$ from $(1, 1, 1)$; divergence threshold $\|a - b\| \geq 1$ (attractor diameter ≈ 40). Literature value $\lambda_{\text{max}} \approx 0.906$.

Table 5: Divergence horizon versus initial displacement δ_0 in x ; “Lyapunov” = $\ln(1/\delta_0)/0.9056$; “illustrative rate” = $\ln(1/\delta_0)/\max_{\text{sample}} \|Df\|_2$ with the sampled max $\|Df\|_2 = 30.3$.

δ_0	measured t_{crit}	Lyapunov estimate	illustrative, not certified	fitted local λ
10^{-6}	12.36	15.26	0.46	0.83
10^{-9}	22.51	22.88	0.68	0.81
10^{-12}	29.43	30.51	0.91	0.83
h vs. $h/2$ (same u_{ref})	23.86	(RK4 local error $\leq 8.6 \times 10^{-11}/\text{step}$)		

Finding 4.5 (Logarithmic horizon; guarantee versus typical behaviour). **GROUND** Across the three scripted perturbations, each thousand-fold reduction in δ_0 buys roughly 7–10 time units, consistent with $\ln(10^3)/0.906 = 7.6$ and with Theorem 3.8 in the $L > 1$ regime. Halving the step (changing ε with $\delta_0 = 0$) gives a horizon of 23.9, consistent with $\ln(1/8.6 \times 10^{-11})/0.906 \approx 25.6$. But a *heuristic* horizon obtained by substituting the sampled norm 30.3 uses $L \approx e^{30.3h}$, which is ~ 33 times the typical Lyapunov rate: a *heuristic* error horizon based on this sampled norm is $t \approx 0.5$ –1 while the scripted perturbation trajectories agree for 12–29 units of time. Report A’s “ $L \geq e^{\lambda_{\text{max}}\Delta t}$ ” is a correct inequality that conflates two different quantities. The honest statement is: Theorem 3.8 becomes a certified worst-case guarantee only after an upper Lipschitz bound valid over a proved reachable region is supplied. The sampled value 30.3 is not such a certificate; a fitted Lyapunov exponent is a different, statistical diagnostic.

Finding 4.6 (Observed floating-point non-injectivity). **GROUND** At $u_{\text{ref}} = (-4.4085, -3.3242, 24.1048)$ the unit in the last place of x is $\text{ulp}(x) = 8.88 \times 10^{-16}$. Displacing x by 1, 2 or 3 ulp and integrating both

trajectories with identical arithmetic, the displaced state becomes *bit-identical* to the reference after 8, 14 and 34 RK4 steps respectively, and the separation is exactly zero for all later $t \leq 200$: distinct machine states are mapped to the same state by a short composition of G . This is consistent with the increment’s response to the displacement lying below the resolution of the state grid, so that rounding absorbs it. Displacements of 4 and 8 ulp become bit-identical *to each other* after 191 steps and thereafter follow one trajectory (same $t_{\text{crit}} = 34.512$, same separations at $t = 10, 50, 100, 200$) that diverges from u_{ref} ; 16 ulp gives $t_{\text{crit}} = 34.547$. Thus the iterated digital map is *contractive and non-injective* at the ulp scale, whereas a well-posed Lorenz ODE flow has injective finite-time maps on its domain of existence (even though it also has stable directions). This is an implementation-specific observation supporting $U_S \not\cong U_P$ under the assumption of an injective target flow map (Theorem 3.13(6)): not a cardinality argument but an observed merger of states. It also shows that Report A’s worry “ $\delta_0 \sim 10^{-16}$ ” is below the quantiser: such a δ_0 is not a perturbation the simulator can represent.

5 The equivalence taxonomy, amended

Tier	Relation	Criterion	Does not entail
1	Approximate observational <i>simulation</i> (one-sided)	$\sup_{n \leq N} d_Q(q(DG^n E x_0), q(F^n x_0)) \leq \eta$ (Thm. 3.8)	state conjugacy; infinite horizon; the reverse direction (needed for Girard–Pappas <i>bisimulation</i>)
$1\frac{1}{2}$	Shadowing	for every pseudo-orbit $(DG^n z_0)$ there is a true orbit $F^n y_0$ with $\sup_n d \leq \eta$ (Anosov–Bowen; Hammel–Yorke–Grebogi)	that $y_0 = x_0$; unlimited shadowing time for non-hyperbolic systems (Sauer–Grebogi–Yorke)
2	Reduced isomorphism in $\mathbf{SICT}_{M,J}$	intertwining homeomorphism h	$X_P \cong X_S$; equality of μ or of measure-theoretic entropy; anything about hardware. <i>Does</i> entail topological entropy, periodic-orbit counts, transitivity
3	Full isomorphism in $\mathbf{SICT}_{M,J}^{\text{full}}$	(ϕ_X, ϕ_Y, h) preserving (i)–(iii)	material substrate; implementation-level dissipation; mass/charge realisation
4	Physical identity	$W_P = W_S$ (Leibniz’s law over all physical predicates)	not derivable from Tiers 1–3; the “ontological collapse” is the inference $2 \Rightarrow 4$ or $3 \Rightarrow 4$

The hierarchy is well-ordered ($4 \Rightarrow 3 \Rightarrow 2 \Rightarrow 1\frac{1}{2} \Rightarrow 1$ with $\eta = 0$ and $N = \infty$). Finding 4.6 places finite-precision chaotic simulation strictly below Tier 2 for a reason independent of cardinality; Finding 4.2 places non-symplectic integrators below Tier 1 beyond a horizon set by Cohesion failure.

6 Part IV — New results: finite factors, reversible targets and falsification

6.1 A finite exact factor of a surjective dynamics is reversible

Theorem 6.1 (Finite-factor permutation theorem). **ESTABLISHED** *Let A be any nonempty set, $F : A \rightarrow A$ be surjective, and Z be a finite nonempty set. If $h : A \rightarrow Z$ is surjective and $G : Z \rightarrow Z$ satisfies $h \circ F = G \circ h$, then G is a permutation of Z .*

Proof. Take any $z \in Z$ and choose $a \in A$ with $h(a) = z$. By the surjectivity of F , there exists $a' \in A$ with $F(a') = a$. Then $G(h(a')) = h(F(a')) = h(a) = z$. Thus G is onto. A surjection of a finite set to itself is injective and hence bijective. $\square$

Corollary 6.2 (No exact finite noninvertible quotient of a reversible target). **ESTABLISHED** *If F is invertible and a finite-state G merges distinct states, there cannot exist any surjective exact factor map h intertwining F and G .*

The result is independent of entropy, cardinality of A , temperature and implementation technology. It is stronger than merely denying a homeomorphic conjugacy because even a many-to-one *onto* coarse-graining cannot rescue an exact factor relation with a noninjective finite map. It does *not* rule out one-sided approximate simulations, infinite-state factors, or finite-horizon matching; a predictive physical claim would require evidence that the chosen F and finite-state representation accurately describe a particular physical system.

6.2 A corrected benchmark: shift target with a lossy simulator

Take $A = Z = \{0, 1\}^6$ and the reversible cyclic shift $F(z)_i = z_{i+1 \bmod 6}$. Define $G(z)$ to apply F and then replace the last register bit with zero. The target is a bijection; G is not. Starting from the uniform ensemble over all 64 states, exhaustive enumeration gives:

Table 7: New corrected benchmark; entropy and erasure in bits. Values computed from `extension_benchmarks.py`.

Step n	0	1	2	3	4	5	6	7–10
$H(\mu_n)$	6	5	4	3	2	1	0	0
$\Delta H_{n \rightarrow n+1}$	1	1	1	1	1	1	0	0
$ \text{supp } \mu_n $	64	32	16	8	4	2	1	1

Finding 6.3 (Thermodynamic benchmark repaired). **GROUND** The source report’s Rule-90 example is a valid study of an irreversible *target*, but not of an irreversible simulator applied to a reversible target. The replacement above tests exactly that asymmetry. Its total logical erasure is six bits and stops after six steps. The example is a deterministic finite calculation, *not* calorimetry. The six-bit bound is a ceiling on this state-level entropy loss, *not* a ceiling on physical device heat.

6.3 Quantifying the sharpness of the reachable-set bound

For the harmonic oscillator with $\Delta t = 0.01$, identify (x, p) with a complex number and let $f = e^{-i\Delta t}$ and $g = 1 - i\Delta t$. Then at $z_0 = 1$ the actual Euler error is $|g^n - f^n|$ while the reachable-state triangle bound is exactly

$$B_n = |g - f| \sum_{k=0}^{n-1} |g|^k = |g - f| \frac{|g|^n - 1}{|g| - 1}.$$

At $n = 200,000$ the independently executed code reports $|g^n - f^n| = 22,014.4582693$, $B_n = 22,014.9452431$ and $|g^n - f^n|/B_n = 0.9999778799$. This is a *near-sharp* bound, not strict algebraic equality. Importantly, B_n is a path-weighted sum using local errors along the reachable orbit; the theorem’s single uniform ε would be the supremum over reachable states and yields a looser bound.

6.4 Four coordinates as separately falsifiable diagnostics

A computational interpretation of S–I–C–T should specify *measurable proxies* rather than assigning labels to one score after the fact. The following is a proposed protocol, not a theorem about physics:

Structure (S) State-carrier type, dimension, topology, resolution, reachable-set coverage; report the encoding map and its domain.

Information (I) Ensemble Shannon entropy, erasure conditional entropy, and conditional uncertainty after decoding; declare the distribution.

Cohesion (C) Drift of independently specified invariants (energy, magnetisation), topological consistency and, where supported, symplectic error.

Transformation (T) Local consistency error, error propagation, estimated Lyapunov growth, observed finite-precision merging time and shadowing horizon; separate observed rates from certified bounds.

For a preregistered test, publish the map definitions, random seeds (if any), float precision, platform versions, calibration distribution, observable set, metric, time horizon and pass/fail tolerance before viewing the final curves. An out-of-sample cross-precision or cross-platform test can falsify a *particular proposed emulator equivalence*; it cannot prove or disprove the metaphysical simulation hypothesis.

6.5 Limitations, implementation dependencies and what is not claimed

The original RK4 experiment shows collisions of representable states for the particular discrete arithmetic evaluated; changes in floating-point optimisations or compiler can change first-coincidence times. The number $\max \|Df\|_2 \approx 30.3$ is a sampled maximum along a numerical trajectory; it is *not* a verified uniform Lipschitz constant. A guarantee would require a rigorously enclosed, forward-invariant set and a bound on that set. The original source manuscript itself was not included in the audit materials, so claims about its theorems are based on their restatements in Report A. No evidence supplied here distinguishes whether the actual universe is a simulation. The mathematically defensible result is a hierarchy of well-defined equivalence notions and explicit limits on inferences between them.

7 Verdict and recommendations

Verdict on Report A. Mathematical core: sound. Extensions: the functorial stratification and the error budget are correct and useful (**ESTABLISHED**). The computability obstruction is correct as mathematics and must be badged **GROUND** as conditional on the continuum model. The Landauer obstruction is correct as a theorem and overclaimed as physics; corrected to an erasure budget and a dissipation–memory trade-off. One benchmark misapplies a theorem; two benchmark heuristics are badged **HEURISTIC**. No claim in Report A is **SPECULATIVE**; no claim required rejection outright.

Verdict on Róth (as restated). Every restated theorem is **ESTABLISHED**. The reconstruction achieves what it claims: it blocks the inference from categorical isomorphism to ontological identity by making explicit which structure the morphisms carry.

Recommendations for the S–I–C–T programme.

1. Adopt Definition 3.1 with Borel (not continuous) predicates and invariant μ ; adopt surjective ϕ_X (Remark 3.6) so that coarse-graining is a morphism.
2. State the error bound with the reachable-set hypothesis explicit, and report ε as $\sup_{\mathcal{R}}$, not as a shell constant.
3. Replace “Landauer obstruction” by “erasure budget” (Theorem 3.13) and state the Bennett trade-off; move implementation-level dissipation to a separate, hardware-indexed layer of the enrichment.
4. Badge every continuum-dependent obstruction **GROUND** and say what finiteness hypothesis would void it; the programme’s own discipline demands it.
5. Add shadowing as Tier 1 $\frac{1}{2}$ and make the ulp-scale merger test (Finding 4.6) a standard falsification protocol: a simulator that merges representable states is not conjugate to an expansive target, full stop.
6. Replace Benchmark 1’s Rule-90 variant by a reversible target simulated by a lossy map; the new six-bit benchmark in Section 6 provides executable numbers and a finite-factor theorem.

References

- [1] Róth, M. (2026). *Structural Equivalence Without Ontological Collapse: A Category-Theoretic Reconstruction of S–I–C–T Systems, Simulations, and Physical Dynamics*. Roth Complexity Lab Working Paper, 21 October 2026.
- [2] Autonomous Truth-Seeking Science Panel (2026). *First-Principles Validation Report & Extended Academic Paper* on [1] (“Report A”), 21 October 2026.
- [3] Mac Lane, S. (1998). *Categories for the Working Mathematician*, 2nd ed. Springer.
- [4] Ahrens, B., North, P. (2019). Univalent foundations and the equivalence principle. In *Reflections on the Foundations of Mathematics*, Springer.
- [5] Walters, P. (1982). *An Introduction to Ergodic Theory*. GTM 79, Springer.
- [6] Lind, D., Marcus, B. (1995). *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press.
- [7] Rutten, J.J.M.M. (2000). Universal coalgebra: a theory of systems. *Theoret. Comput. Sci.* 249(1), 3–80.
- [8] Milner, R. (1989). *Communication and Concurrency*. Prentice Hall.
- [9] Girard, A., Pappas, G.J. (2007). Approximate bisimulation relations for constrained linear systems. *Automatica* 43(8), 1307–1317.
- [10] Lax, P.D., Richtmyer, R.D. (1956). Survey of the stability of linear finite difference equations. *Comm. Pure Appl. Math.* 9(2), 267–293.
- [11] Hairer, E., Lubich, C., Wanner, G. (2006). *Geometric Numerical Integration*, 2nd ed. Springer.
- [12] Shannon, C.E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.* 27, 379–423, 623–656.
- [13] Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* 5(3), 183–191.
- [14] Bennett, C.H. (1973). Logical reversibility of computation. *IBM J. Res. Dev.* 17(6), 525–532.
- [15] Bennett, C.H. (1982). The thermodynamics of computation—a review. *Int. J. Theor. Phys.* 21, 905–940.
- [16] Bowen, R. (1975). *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms*. LNM 470, Springer.
- [17] Hammel, S.M., Yorke, J.A., Grebogi, C. (1987). Do numerical orbits of chaotic dynamical processes represent true orbits? *J. Complexity* 3, 136–145.
- [18] Sauer, T., Grebogi, C., Yorke, J.A. (1997). How long do numerical chaotic solutions remain valid? *Phys. Rev. Lett.* 79, 59–62.
- [19] Pilyugin, S.Yu. (1999). *Shadowing in Dynamical Systems*. LNM 1706, Springer.
- [20] Tucker, W. (2002). A rigorous ODE solver and Smale’s 14th problem. *Found. Comput. Math.* 2, 53–117.
- [21] Brouwer, D. (1937). On the accumulation of errors in numerical integration. *Astron. J.* 46, 149–153.
- [22] Bekenstein, J.D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Phys. Rev. D* 23, 287–298.
- [23] Chalmers, D.J. (2003/2005). The Matrix as metaphysics. In C. Grau (ed.), *Philosophers Explore the Matrix*, Oxford University Press.
- [24] Frigg, R., Nguyen, J. Scientific representation. *Stanford Encyclopedia of Philosophy*.
- [25] Ladyman, J. (1998). What is structural realism? *Stud. Hist. Phil. Sci.* 29, 409–424.
- [26] Newman, M.H.A. (1928). Mr. Russell’s “Causal Theory of Perception”. *Mind* 37, 137–148.

A Reproduction parameters

Added extension. Script `extension_benchmarks.py` requires Python 3.10+ and NumPy; its deterministic assertions verify the reversible target, irreversible lossy map, six-step entropy budget, finite factor and reachable-orbit Euler ratio. Results are in `extension_results.json`.

Environment. Python 3, NumPy 2.5.3, IEEE-754 binary64 (binary32 where stated). Script: `benchmarks.py` (accompanies this paper); runtime ≈ 15 s.

Benchmark 1. Rings $N \in \{6, 7, 8\}$; all 2^N states enumerated. Shift $T(z)_i = z_{i+1 \bmod N}$; encoding $h(\sigma) = (\sigma + 1)/2$; magnetisation $\sum_i \sigma_i$ versus $\sum_i (2z_i - 1)$. Rule 90: $z'_i = z_{i-1} \oplus z_{i+1}$ (periodic). Erasure: uniform $\mu_0, \mu_{n+1} = G_* \mu_n, H$ in bits, 12 steps.

Benchmark 2. $\Delta t = 0.01, (x_0, p_0) = (1, 0), E_0 = 0.5$. Explicit Euler $(x + \Delta t p, p - \Delta t x)$; implicit Euler $(1 + \Delta t^2)^{-1}(x + \Delta t p, p - \Delta t x)$; Störmer–Verlet kick–drift–kick $p \leftarrow p - \frac{\Delta t}{2}x, x \leftarrow x + \Delta t p, p \leftarrow p - \frac{\Delta t}{2}x$; exact rotation with $\cos \Delta t, \sin \Delta t$ in binary64. ε as max over 3601 points of the unit shell of $|G(z) - F(z)|$; reachable-set sum $\sum_{k < n} \varepsilon_{\text{unit}} |G^k z_0|$. Trajectory error against $(\cos n\Delta t, -\sin n\Delta t)$.

Benchmark 3. Lorenz-63 (10, 28, 8/3); RK4 with $h = 10^{-3}$; transient $t = 50$ from $(1, 1, 1)$ to $u_{\text{ref}} = (-4.40848, -3.32421, 24.10480)$. Jacobian spectral norm sampled every 10 steps over $t \in [0, 100]$ from u_{ref} : max = 30.30, mean 18.27. Divergence: perturb x by δ_0 , integrate both to $t = 200$ or until $\|a - b\| \geq 1$. Local λ fitted by least squares on $\ln \|a - b\|$ over the window $10\delta_0 < \|a - b\| < 10^{-2}$. Step-size test: h versus two steps of $h/2$. Local error: Richardson estimate $\|u_h - u_{h/2, h/2}\|/(1 - 2^{-4})$ over 2000 steps, max = 8.63×10^{-11} . Ulp test: displacements of $k \cdot \text{ulp}(x_{\text{ref}})$, $k \in \{1, 2, 3, 4, 8, 16\}$, ulp = 8.88×10^{-16} ; bit-identity of state vectors checked at every step for 10^4 steps (first coincidence with reference: $k = 1$ at step 8, $k = 2$ at 14, $k = 3$ at 34; $k = 4$ with $k = 8$ at step 191; $k = 16$ never within 10^4).

Reporting discipline. Every number in Tables 3–5 and in Findings 4.2–4.6 is copied from the script’s output without rounding beyond the displayed digits. The original audit reports adding the ulp experiment after the $\delta_0 = 10^{-15}$ run failed to diverge. The present continuation used pre-specified six-bit finite-factor and lossy-shift examples and states its new experiment parameters and assertions explicitly. Two Lorenz table rows at 10^{-13} and 10^{-14} have been removed because they are absent from the attached executable script.