The Binding Invariant

Synthetic Grounding as Cohomological Gradient Flow

Constraint Persistence, Derived Ambiguity, Process-Algebraic Equivalence, and Neural Garbage Collection as Semantic Renormalization

Flyxion

April 2026

Abstract

Every system that undergoes irreversible transformation faces the same question: what, if anything, carries through? This essay introduces the *binding invariant* as the structured property that persists under constraint closure when a system passes from one regime of possibility to another.

We formalize this notion using constraint sheaves over partially ordered time and prove the Binding Invariant Theorem: a system possesses a binding invariant if and only if its constraint history admits a coherent section in the topos of feasible trajectories. This yields precise criteria for uniqueness, stability, and rupture.

We then extend the framework using derived algebraic geometry, modeling ambiguity as higher homotopy in a derived stack of semantic states. In this setting, the binding invariant emerges as a homotopy-stable minimum of a cohomological energy functional, and synthetic grounding is interpreted as a gradient flow that minimizes obstruction across scales.

Bridging to concurrency theory, we show that Linear Logic Concurrent Constraint (LCC) programming provides the operational counterpart of this geometry: constraint consumption corresponds to obstruction reduction, and observational equivalence (bisimulation and barbed congruence) corresponds to equivalence classes of descent trajectories.

We further formalize CLIO as a monoidal endofunctor on a category of semantic states and prove two central results: the associated energy functional is a bisimulation-invariant Lyapunov function, and LCC bisimulation corresponds to homotopy equivalence of descent trajectories.

Finally, we show that Neural Garbage Collection (NGC) provides an empirical realization of the framework: its learned memory eviction dynamics implement constraint-driven descent, its uncertainty signal behaves as a Lyapunov function, and its graph-based reasoning approximates a homotopy colimit. A finite worked example with diagrammatic proof grounds the abstract theory.

Taken together, these results establish a unified view in which identity is not a static property of states, but the Lyapunov-stable equivalence class of trajectories that remain coherent under irreversible constraint accumulation.

Part I: Foundational Framework

1. Introduction

A system changes. Something survives. But *what* survives, and in what sense?

This is not the ancient Ship of Theseus puzzle, though it rhymes with it. That problem asks whether numerical identity persists under material replacement. The question here is structurally deeper: it asks what *constrains* the set of successors that a system can coherently be—not what the system *is* after transformation, but what *binds* its trajectory, making certain transformations continuations and others ruptures.

We call this the *binding invariant*: the property that must be preserved for a transformation to count as development rather than replacement. More precisely, it is the information-theoretic closure of the system's constraint history that any future state must be consistent with in order to inherit the system's trajectory.

This question arises across multiple domains. In concurrency theory, observational equivalence replaces state identity as the primary notion of process sameness [9]. In thermodynamics, irreversibility constrains possible evolutions and places asymmetric costs on state transitions [12, 13, 14]. In modern machine learning, systems that jointly learn reasoning and forgetting force a reconceptualization of what internal state *is* [17]. The binding invariant names what is preserved across all of these domains simultaneously.

The binding invariant arises at the intersection of three lines of work in the RSVP theoretical program. **The Compiled Self** established that personal identity is not a continuous substrate but a fixed point of a constraint-preserving adjunction over an irreversible history topos; the binding invariant names what that fixed point *is*. **The Constraint–Recurrence Principle** showed that systems return not to prior states but to prior constraint configurations, structured by a recurrence operator; the binding invariant is the kernel of that operator. **The Yarncrawler Identifiability Theorem** demonstrated that world-state reconstruction from partial observations is possible if and only if the observation sheaf admits a global section; the binding invariant is the analogue of that section in the context of self-reconstruction across time.

The essay is organized as follows. Sections 2–4 (Part I) develop the sheaf-theoretic foundations and prove the Binding Invariant Theorem. Sections 5–8 (Part II) develop the derived-geometric and variational interpretation. Sections 9–11 (Part III) bridge to Linear Concurrent Constraints. Sections 12–14 (Part IV) formalize CLIO as a monoidal endofunctor with fixed-point algebra and sheaf-theoretic semantics. Sections 15–17 (Part V) prove the Lyapunov and bisimulation theorems. Sections 18–21 (Part VI) absorb NGC and provide the worked example. A bibliography closes the essay.

2. Formal Setting

2.1. The Site of Constraint Histories

Let S be a system evolving through a partially ordered time set $(T, \leq)$.

Definition 2.1 (Site of Constraint Histories)**.** The *constraint site* of S is $\mathcal{J}_S = (T, \tau_{\text{Alex}})$ where τ_{Alex} is the Alexandrov topology on $(T, \leq)$, with coverings given by upward-directed families $\{U_i\}$ such that $U = \bigcup_i U_i$ and $t \in U_i, t \leq t' \Rightarrow t' \in U_i$.

Definition 2.2 (Constraint Sheaf)**.** A *constraint sheaf* for S is a sheaf $\mathcal{C}$ on $\mathcal{J}_S$ such that: (i) the stalk $\mathcal{C}_t$ is the set of constraints active at t; (ii) restriction maps $\rho_{t't} : \mathcal{C}_{t'} \to \mathcal{C}_t$ are monotone for $t \leq t'$ (the *irreversibility condition*: constraint satisfaction is preserved under forward evolution); and (iii) the gluing axiom holds.

Definition 2.3 (Trajectory Sheaf)**.** A trajectory is a section of a sheaf $\mathcal{F}$ on $\mathcal{J}_S$ satisfying:

$$\mathcal{F}(U) \;=\; \bigl\{\, \sigma : U \to \text{States} \;\mid\; \sigma(t) \models \mathcal{C}_t \,\bigr\}.$$

Definition 2.4 (Constraint History)**.** The *constraint history* of S up to time t is $\mathcal{H}_t(S) = \varprojlim_{t' \leq t} \mathcal{C}_{t'}$.

$\mathcal{H}_t(S)$ is the structural sediment of all prior constraint regimes—generally larger than $\mathcal{C}_t$ because it retains the record of past closures even when present constraints do not directly reflect them. Constraint-based models of computation originate in Concurrent Constraint Programming, where systems evolve by monotonic accumulation of information [7]. Linear logic extensions introduce controlled consumption of constraints, breaking monotonicity while preserving logical structure [8, 6].

Definition 2.5 (Closure Event)**.** A *closure event* at time t is a pair (t, γ) where $\gamma : \mathcal{C}_t \to \mathcal{C}_{t^+}$ is injective, generates strictly more constraints, and is irreversible in the sense that no constraint-compatible γ^{-1} exists as a valid restriction map.

2.2. Topos of Feasible Trajectories

Let $\mathbf{Traj}(S)$ be the category of constraint-compatible trajectories of S and constraint-preserving maps between them.

Proposition 2.6. *When $\mathcal{C}$ satisfies the gluing condition,* $\mathbf{Traj}(S)$ *has the structure of a Grothendieck topos.*

Proof. The presheaf category $\widehat{T}$ is a topos. The constraint sheaf condition restricts to the full subcategory of sheaves, which is a reflective subcategory of $\widehat{T}$ and hence a topos by standard results in topos theory [2]. $\square$

The subobject classifier Ω of $\mathbf{Traj}(S)$ classifies constraint-admissible sub-trajectories: the truth values for the proposition "this sub-trajectory is a legitimate continuation of S." The use of Ω aligns the construction with standard topos-theoretic semantics, where truth values classify admissible substructures [2, 1].

3. The Binding Invariant

Definition 3.1 (Binding Invariant)**.** The *binding invariant* of S at time t is

$$\mathcal{B}_t(S) \;=\; \Gamma\big(\mathcal{H}_t(S),\, \Omega_{\mathbf{Traj}(S)}\big),$$

the global sections of the subobject classifier over the constraint history.

$\mathcal{B}_t(S)$ is not a single property but the structured space of coherent judgments about what counts as S continuing. An element $b \in \mathcal{B}_t(S)$ may be understood as a rule assigning to each admissible sub-trajectory $U \subseteq T$ a truth value indicating whether U extends to a coherent global trajectory compatible with $\mathcal{H}_t(S)$. Thus $\mathcal{B}_t(S)$ *classifies admissible continuations* rather than describing states.

Definition 3.2 (Binding Section)**.** A *binding section* at (t, γ) is a section $b \in \Gamma(\mathcal{H}_t(S), \mathcal{C})$ such that $b(t') \in \mathcal{C}_{t'}$ for all $t' \leq t$ and $\{b(t')\}$ is coherent in the inverse system.

Theorem 3.3 (Binding Invariant Theorem)**.** *Let S have irreversible constraint sheaf $\mathcal{C}$ and closure sequence $E = \{(t_i, \gamma_i)\}_{i=1}^{n}$ with effective descent for trajectories. The following are equivalent:*

- *(i)* $\mathcal{B}_{t_n}(S) \neq \varnothing$.
- *(ii)* $\mathbf{Traj}(S)$ *admits a global section compatible with* $\mathcal{C}$ *after the last closure event.*
- *(iii)* $\mathcal{H}_{t_n}(S)$ *admits a binding section compatible with all events in* E.
- *(iv)* E *has a coherent limit in* $\mathbf{Pro}(\mathbf{Traj}(S))$.

Proof. The argument relies on standard sheaf-theoretic gluing and limit constructions [1, 4]. (i)⇒(ii): By Yoneda, a global section of Ω over $\mathcal{H}_{t_n}(S)$ corresponds to a non-empty object in $\mathbf{Traj}(S)$, which is a global trajectory section. (ii)⇒(iii): The constraint components $b(t') = \mathsf{traj}(\sigma(t')) \cap \mathcal{C}_{t'}$ of any global section σ form a binding section, compatible with all closure events via compatibility of σ with the restriction maps of $\mathcal{C}$. (iii)⇒(iv): A binding section defines a compatible family of morphisms in $\mathbf{Traj}(S)$; the pro-limit exists by the universal property. (iv)⇒(i): The pro-limit's projection maps define an element of $\varprojlim_i \mathbf{Traj}(S)_{t_i}$; by the sheaf condition on $\mathcal{C}$ this yields a global section of Ω. $\square$

Corollary 3.4 (Uniqueness Criterion)**.** *$\mathcal{B}_t(S)$ is a singleton iff $\mathbf{Traj}(S)$ is Boolean at t.*

Corollary 3.5 (Rupture Criterion)**.** *A transformation at* t *is a* rupture *(replacement, not continuation) iff* $\mathcal{B}_t(S) = \varnothing$*.*

Corollary 3.6 (Binding Stability)**.** *If* $\mathcal{B}_t(S) \neq \varnothing$ *and all subsequent closure events are conservative, then* $\mathcal{B}_{t'}(S) \neq \varnothing$ *for all* $t' > t$*.*

Remark 3.7 (Degrees of Rupture)*.* Rupture may occur gradually. A weakening of the binding invariant corresponds to loss of coherence in portions of $\mathcal{H}_t(S)$, reflected in the non-existence of sections over certain subdomains. Total rupture is $\mathcal{B}_t(S) = \varnothing$; partial rupture is failure of gluing on a covering. Identity is therefore not an intrinsic property of states, but a fixed point in the category of constraint-consistent trajectories.

3.1. Dynamical Interpretation

The binding invariant admits a dynamical interpretation that bridges to Parts II–VI. Let $\mathcal{C} : X \to X$ be an update operator on semantic states, and let $\mathcal{E}$ be an energy functional measuring constraint inconsistency. A trajectory $x_{t+1} = \mathcal{C}(x_t)$ induces a descent process $\mathcal{E}(x_{t+1}) \leq \mathcal{E}(x_t)$. Then $\mathcal{B}_t(S)$ corresponds to the set of equivalence classes of trajectories converging to the same minimal energy configuration. The binding invariant is the fixed-point structure of a constraint-induced descent flow.

4. Applications and Philosophical Consequences

Cognitive identity. The binding invariant gives the Compiled Self thesis a precise interpretation: the self is the structured space of coherent judgments about what would count as *this person* continuing. Identity is not a substance but a section; it is not determined by present state alone; rupture is formally possible; and amnesia is structural—severe disruption of the binding section, not merely data loss. This perspective aligns with modern machine learning systems in which reasoning and memory are jointly optimized, and internal state is actively pruned during inference [17].

Computation. A computation's binding invariant is the record of what it has committed to. Naïve checkpointing fails when side effects alter the external constraint environment in ways not captured by the checkpoint. System migration (not mere update) is required when an upgrade destroys the binding section. In process calculi, equivalence is defined observationally: two processes are identical if no environment can distinguish them [10, 9].

Physical processes. First-order phase transitions are closure events; absence of a binding section across the transition implies the post-transition system is a distinct entity. The thermodynamic arrow of time is the direction along which binding invariants accumulate: entropy production is the physical signature of irreversible closure [12, 13, 14, 15].

Philosophical summary. The binding invariant supports structural persistence (what continues is a coherence condition, not a substance), refutes presentism (the past is constitutive, not merely historical), and formalizes the distinction between development and rupture. This position is structurally aligned with categorical and structuralist approaches to identity, where objects are determined by their relations rather than intrinsic properties [5, 4].

Part II: Synthetic Grounding as Cohomological Gradient Flow

Central claim: A self, or coherent world-state, is a homotopy-stable minimum of a cohomological energy functional defined over a derived moduli space of candidate interpretations. The binding invariant is the structure that remains fixed at such minima.

5. Ambiguity as Higher Homotopy in Derived Stacks

5.1. The Derived Stack of Semantic States

Let Ω denote the Spherepop event-log base. We model the space of candidate interpretations as a derived stack:

$$\mathfrak{M} \;:\; \Omega^{\mathrm{op}} \;\longrightarrow\; \infty\text{-}\mathbf{Grpd},$$

assigning to each open event-history $U \subseteq \Omega$ an ∞-groupoid $\mathfrak{M}(U)$ of candidate interpretations. A "thought"—a globally coherent interpretation—is a section:

$$x \;\in\; \Gamma(\Omega, \mathfrak{M}) \;\simeq\; \mathrm{holim}_{U\subseteq\Omega}\, \mathfrak{M}(U).$$

Instead of assigning to each U_i a set of local sections, we assign an ∞-groupoid whose objects are candidate local thoughts, whose 1-morphisms are identifications between candidate thoughts, whose 2-morphisms are homotopies between identifications, and so on. Ambiguity is not an external imperfection but intrinsic higher structure. A global thought is a point in the homotopy limit $\Gamma(\Omega, \mathfrak{M}) \simeq \mathrm{holim}_{U_i \in \mathcal{U}}\, \mathfrak{M}(U_i)$.

5.2. Mapping Spaces and Ambiguity

Two conflicting interpretations $x_1, x_2 \in \mathfrak{M}(U)$ are classified by the mapping space $\mathrm{Map}_{\mathfrak{M}(U)}(x_1, x_2)$:

- *Empty*: hard contradiction.
- *Nontrivial*: ambiguity—a path exists but is underdetermined.
- *Contractible*: equivalence—no meaningful distinction survives.

Higher-order ambiguity corresponds to higher homotopy. When paths between x_1 and x_2 are themselves ambiguous, we ascend to 2-morphisms; ambiguity about 2-morphisms yields

3-morphisms, and so on. The system carries this tower forward without forcing premature resolution. Two candidate global thoughts need not be strictly identical in their overlap data; they may be related by a path in $\mathrm{Map}_{\mathfrak{M}(U_i \cap U_j)}(x_i|_{U_i \cap U_j}, x_j|_{U_i \cap U_j})$.

Definition 5.1 (Derived Binding Invariant)**.** The *derived binding invariant* of S is

$$\mathbf{R}\mathcal{B}(S) \;=\; \mathrm{Map}_{\mathbf{dSt}}(\mathcal{H}_t(S),\; \mathfrak{M}).$$

The classical $\mathcal{B}_t(S)$ is recovered as $\pi_0(\mathbf{R}\mathcal{B}(S))$.

5.3. The Self as n-Stack

The binding invariant in the derived setting is the homotopy limit of the system's homotopy-coherent diagram of candidate interpretations:

$$\mathcal{B}_\infty(S) \;=\; \operatorname{holim}\big(\text{homotopy-coherent diagram of candidate sections}\big).$$

Theorem 5.2 (Derived Binding Invariant Theorem)**.** $\mathbf{R}\mathcal{B}(S) \neq \varnothing$ *iff the homotopy-coherent diagram of candidate world-states over* Ω *has a non-trivial homotopy limit. When it exists,* $\mathcal{B}_\infty(S)$ *is an* n*-stack—a derived geometric object that persists despite lower-dimensional fluctuations.*

A point $x \in \mathfrak{M}$ has a tangent complex $T_x\mathfrak{M}$ encoding infinitesimal deformations, while its higher cohomology encodes obstructions to extending local deformations globally. A candidate thought is not merely present or absent: it has nearby deformations, infinitesimal continuations, and obstruction classes preventing extension into a stable global realization.

6. The Cohomological Energy Functional

6.1. Obstruction Classes

Let $H^\bullet(\Omega, \mathfrak{M})$ be the cohomology of $\mathfrak{M}$ over Ω. The degrees carry distinct interpretations: H^0 recovers global sections; H^1 classifies local inconsistencies (gluing failures, contradictions between adjacent event windows); H^k for $k \geq 2$ classifies k-fold coherence failures (higher-order structural ambiguity). Let $\mathrm{Obs}^k(x) \in H^k(\Omega, \mathfrak{M})$ denote the k-th obstruction class of state x.

A local version over a cover $\{U_i\}$ provides operational form:

$$\mathcal{E}_{\mathcal{U}}(x) = \sum_{i<j} \alpha_{ij}\, d(x_i|_{U_i \cap U_j}, x_j|_{U_i \cap U_j})^2 + \sum_{i<j<k} \beta_{ijk} \|\omega_{ijk}\|^2 + \mu S(x) + \nu C(x),$$

where the first term measures pairwise mismatch on overlaps and the second measures higher cocycle defect on triple overlaps.

6.2. The Global Functional

Let $S(x)$ be semantic dispersion (entropy: the degree to which x spreads probability mass over incompatible interpretations), $C(x)$ be substrate cost (memory, metabolic, or computational), $\mathcal{A}(x)$ the set of pending unsatisfied constraints, and $r(a, x)$ the residual cost of pending constraint a under state x.

Definition 6.1 (Constraint Energy Functional)**.** With parameters $\lambda_k, \mu, \nu, \eta > 0$:

$$\mathcal{E}(x) \;=\; \sum_{k \geq 1} \lambda_k \left\| \mathrm{Obs}^k(x) \right\|^2 \;+\; \mu\, S(x) \;+\; \nu\, C(x) \;+\; \eta \sum_{a \in \mathcal{A}(x)} r(a, x).$$

The four terms penalize distinct failure modes: failure of constraint closure at each scale; semantic indeterminacy; unsustainable resource use; and residual unfulfilled constraints (the LCC "ask" term, which drives the connection to process algebra in Part III). Low energy corresponds to vanishing obstruction, concentrated probability, efficient representation, and satisfied constraints.

7. Synthetic Grounding as Gradient Flow

7.1. CLIO Dynamics as Descent

In the CLIO framework the system evolves via $x_{t+1} = \mathcal{C}(x_t)$, where each step restricts (prunes incompatible sections), infers (extends by constraint propagation), and stabilizes. Geometrically: project → flow → stabilize.

Definition 7.1 (Synthetic Grounding)**.** *Synthetic grounding* is the stochastic gradient flow:

$$\frac{dx}{dt} \;=\; -\nabla_x \mathcal{E}(x), \qquad \mathbb{E}\big[\mathcal{E}(x_{t+1})\big] \;\leq\; \mathbb{E}\big[\mathcal{E}(x_t)\big],$$

with strict decrease away from fixed points.

Synthetic grounding is *descent on cohomological inconsistency under resource constraint*: the system prunes incompatible sections (reduces $\|\mathrm{Obs}^k\|^2$), merges compatible ones (reduces $S(x)$), compresses redundancy (reduces $C(x)$), and satisfies outstanding constraints (reduces the residual ask-cost term).

7.2. Binding Invariants as Critical Points

Theorem 7.2 (Binding Invariants as Critical Points)**.** *The critical points of* $\partial_t x = -\nabla\mathcal{E}(x)$ *are precisely the binding invariants: states* x^* *such that* $[x^*] \in \mathcal{B}_t(S)$*.*

Proof. At a critical point, $\nabla\mathcal{E}(x^*) = 0$. The gradient of the obstruction terms vanishes iff $\mathrm{Obs}^k(x^*) = 0$ for all $k \geq 1$—x^* is a globally coherent section. The gradient of $\mu S(x)$ vanishes at entropy-minimum; $\nu C(x)$ at the resource-efficient point; and $\eta \sum_a r(a,x)$ when all constraints are satisfied. Together these are equivalent to x^* being a global section of $\mathfrak{M}$ compatible with the full event-history—a binding invariant. □

A binding invariant is:

$$x^* \in \operatorname*{argmin}_{x\in\mathfrak{M}} \mathcal{E}(x) \quad \text{such that} \quad x^* \simeq \mathcal{C}(x^*).$$

The self is a *cohomologically minimal, homotopy-stable fixed point of CLIO dynamics.*

Theorem 7.3 (Convergence of Synthetic Grounding)**.** *If* $\mathcal{E}$ *is convex on* $H^\bullet(\Omega, \mathfrak{M})$ *and the event-log satisfies bounded total variation, then* $x(t) \to x^* \in \mathcal{B}_t(S)$ *as* $t \to \infty$*, uniquely when* $\mathcal{E}$ *is strictly convex.*

Remark 7.4 (Non-convex case and basin structure)*.* When $\mathcal{E}$ is non-convex—as is typical for cognitive and social systems—the flow may converge to any local minimum. Different constraint histories $\mathcal{H}_t(S)$ define different energy landscapes and hence different basins of attraction. CLIO is *locally contractive* within a basin: $d(\mathcal{C}(x), \mathcal{C}(y)) \leq \kappa\, d(x,y)$ for $\kappa < 1$ when x, y lie in the same basin. Return Under Constraint is therefore guaranteed within a basin, not globally. This explains identity fragmentation, allows multiple stable selves, and aligns with LCC non-determinism.

8. Return Under Constraint as Geometric Stability

Proposition 8.1. *Every binding invariant* $x^* \in \mathcal{B}_t(S)$ *satisfies the return property:* $\mathcal{R}_T(x^*) = x^*$ *for all* $T \geq 0$*, where* $\mathcal{R}_T = \Phi_T$ *is the time-*T *map of the gradient flow.*

Theorem 8.2 (Geometric Persistence Theorem)**.** *If* S' *is obtained from* S *by a conservative closure event* τ*, then: (i)* S' *has a binding invariant; (ii) it is homotopic to that of* S*; (iii) the gradient flows of* $\mathcal{E}$ *on* S *and* S' *are conjugate via an isomorphism induced by* τ*.*

Theorem 8.3 (The Self as Cohomological Fixed Point)**.** *Under the convergence conditions of Theorem 7.3: (i) the self is* $x^* = \lim_{t\to\infty} \Phi_t(x_0)$ *for* x_0 *in the basin of attraction; (ii)* x^* *is the fixed point of the return operator; (iii)* x^* *persists under* τ *iff* τ *is conservative; (iv) rupture occurs iff* τ *moves* x_0 *outside the basin of attraction.*

The self compiles itself continuously: each incoming event perturbs x, and synthetic grounding returns it to x^*. Identity is *dynamical equilibrium*, not static persistence.

Cognitive / epistemic concept	Derived-geometric realization
Ambiguity	Nontrivial homotopy in $\mathfrak{M}(U)$
Contradiction	Empty mapping space
Reasoning	Descent on $\mathcal{E}$
Identity	Stable minimum of $\mathcal{E}$
Memory	Retained structure along gradient flow
Binding	Stability under recursive projection
Rupture	Exit from basin of attraction

Part III: Process-Algebraic Correspondences

9. Linear Concurrent Constraints

9.1. Background

Linear Concurrent Constraints (LCC), following Haemmerlé [9], is a process algebra built on a constraint store that evolves by *consumption*: the `ask` operation reads a constraint from the store and destructs it, rewriting the state dynamically. This breaks the monotonicity of classical constraint systems [7] in precisely the way that CLIO dynamics breaks monotonicity via eviction and projection. The key notions of observational equivalence in LCC are: logical equivalence ($P^\dagger \dashv\vdash Q^\dagger$), may-testing, barbed congruence ($\forall C,\ C[P] \sim C[Q]$), and labeled bisimulation. Haemmerlé proves that labeled bisimilarity and barbed congruence coincide [9].

9.2. Accessible Constraints as Binding Residue

In LCC, the observable of a process P is not its full state but

$$\mathcal{O}_D(P) \;=\; \{\, c \mid P \;\Rightarrow\; P' \parallel c \,\}:$$

the set of constraints that remain *accessible* after evolution.

Proposition 9.1 (Accessible Constraints = Binding Residue)**.** *$\mathcal{O}_D(P)$ is the operational shadow of the binding invariant: it corresponds to the support of $\mathcal{B}_t(P)$ in the constraint lattice.*

The identification is structural, not analogical. The binding invariant is the constraint-closed residue of the system's trajectory; $\mathcal{O}_D(P)$ is the constraint-closed residue of the process's execution. They are the same object viewed from different levels of description.

9.3. Constraint Consumption as Obstruction Projection

Proposition 9.2. *Constraint consumption in LCC corresponds to the projection*

$$H^k(\Omega, \mathfrak{M}) \;\longrightarrow\; H^k(\Omega, \mathfrak{M}) \,/\, \langle \textit{consumed constraints} \rangle.$$

Each $\mathsf{ask}/\mathsf{consume}$ *step modifies the cocycle data, reducing* $\|\mathrm{Obs}^k\|^2$ *in expectation:* $\omega \mapsto \omega'$ *with* $\|\omega'\| \leq \|\omega\|$.

Constraint consumption is *cohomological simplification*: it reduces obstruction data, driving the gradient flow toward lower energy.

LCC operation	Cohomological interpretation
$\mathsf{tell}(c)$	Add local section; may increase H^1
$\mathsf{ask}(c)$	Consume constraint; reduce H^k obstruction
Parallel composition	Gluing of local sections
Commitment / choice	Irreversible closure event
Non-determinism	Branching of derived stack $\mathfrak{M}$

10. Equivalences as Geometric Relations

10.1. Bisimulation as Gradient-Flow Equivalence

Definition 10.1 (Flow Bisimulation)**.** $P \approx_{\mathcal{E}} Q$ iff for all flows Φ_t, the trajectories $\Phi_t(P)$ and $\Phi_t(Q)$ are homotopic in $H^\bullet(\Omega, \mathfrak{M})$ at all times t.

10.2. Barbed Congruence as Descent Congruence

Definition 10.2 (Descent Congruence)**.** $P \cong_{\mathcal{E}} Q$ iff for all perturbations δ (context embeddings), $\mathcal{C}(\delta(P)) \sim \mathcal{C}(\delta(Q))$ under flow bisimulation.

Barbed congruence is invariance under all gradient flows: no perturbation can separate the binding invariants of P and Q. This is equivalence under irreversible descent.

11. CLIO-Observational Equivalence

Definition 11.1 (CLIO-Observational Equivalence)**.**

$$x \ \sim_{\mathcal{C}} \ y \quad \Longleftrightarrow \quad \lim_{t\to\infty} \mathcal{C}^t(x) \ = \ \lim_{t\to\infty} \mathcal{C}^t(y).$$

Two states are equivalent if they converge to the same binding invariant.

The refined asymptotic version avoids requiring a unique attractor:

$$x \sim_{\mathcal{C}} y \iff \lim_{t\to\infty} \operatorname{dist}\big(\mathcal{C}^t(x), \mathcal{C}^t(y)\big) = 0.$$

Definition 11.2 (Strong CLIO-Observational Equivalence)**.** $x \approx_{\mathcal{C}} y$ iff $\forall t,\ \mathcal{C}^t(x) \sim \mathcal{C}^t(y)$. This is bisimulation lifted to gradient flow: indistinguishability at every stage of descent.

LCC concept	Geometric realization	Formal object
Constraints	Local sections	$\mathfrak{M}(U)$
Accessible constraints	Binding residue	$\operatorname{supp}(\mathcal{B}_t)$
ask/consume	Obstruction projection	$H^k \to H^k/\langle c\rangle$
Bisimulation	Flow equivalence	$\Phi_t(P) \simeq \Phi_t(Q)$
Barbed congruence	Invariance under irreversible descent	$\mathcal{C}(\delta P) \sim \mathcal{C}(\delta Q)$
Logical equivalence	Derived stack equivalence	$\mathfrak{M}_P \simeq \mathfrak{M}_Q$
May-testing	Shared basin of attraction	$\exists x \in \operatorname{holim} \mathfrak{M}$ reachable from both

Theorem 11.3 (Triple Characterization of the Binding Invariant)**.** *The binding invariant of S admits three equivalent characterizations: (i) global sections of Ω over $\mathcal{H}_t(S)$ (Part I); (ii) homotopy-stable minima of $\mathcal{E}$ (Part II); (iii) the CLIO-observational equivalence class: $\mathcal{B}_t(S) \cong \{x \mid x \sim_{\mathcal{C}} S\}/\simeq$ (Part III).*

Part IV: CLIO as Monoidal Endofunctor

12. Semantic Category and Monoidal Structure

Let **Sem** be a symmetric monoidal category whose objects are semantic states $X \in \mathcal{X}$, whose morphisms $\phi : X \to Y$ are admissibility-preserving transformations, and whose tensor product $\otimes$ represents compositional aggregation of independent semantic components. We interpret $X = (\Phi, \mathbf{v}, S)$ as a structured RSVP field object in **Sem**, where Φ is a scalar semantic potential, $\mathbf{v}$ is a directional inferential flow, and S is an entropy field capturing unresolved structure.

Definition 12.1 (CLIO Operator)**.** The *CLIO operator* is a map $\mathcal{C} : \mathcal{X} \to \mathcal{X}$:

$$\mathcal{C}(X) \;=\; \mathbb{E}_{o \sim \pi_\theta(\cdot|E(X))}\big[\mathcal{F}(E(X), o)\big],$$

composing: (i) restriction $X \mapsto E(X)$; (ii) inference $o \sim \pi_\theta$; (iii) state update $\mathcal{F}$.

Proposition 12.2 (Endofunctor Structure)**.** *$\mathcal{C}$ is an endofunctor on* **Sem***: for any morphism $\phi : X \to Y$, $\mathcal{C}(\phi) : \mathcal{C}(X) \to \mathcal{C}(Y)$ is induced by functoriality of $\mathcal{F}$ and invariance of E over admissible maps.*

Proposition 12.3 (Lax Monoidality)**.** *There exists a natural transformation $\mu_{X,Y} : \mathcal{C}(X) \otimes \mathcal{C}(Y) \to \mathcal{C}(X \otimes Y)$ making $\mathcal{C}$ lax monoidal. Independent reasoning processes can be composed before or after constraint closure, but the two are not strictly equivalent: interaction terms arise through shared resource constraints.*

13. Fixed-Point Algebra of CLIO

Definition 13.1 (CLIO Algebra)**.** An algebra over $\mathcal{C}$ is a pair (X, α) where $\alpha : \mathcal{C}(X) \to X$.

Definition 13.2 (Fixed-Point Algebra)**.** A *fixed-point algebra* satisfies $\alpha \circ \mathcal{C}(\alpha) = \alpha$, and in the special case $\alpha = \mathrm{Id}_X$, we obtain $\mathcal{C}(X) = X$, or in the derived homotopy-correct form, $X \simeq \mathcal{C}(X)$.

Theorem 13.3 (Cognitive Fixed Point)**.** *Cognitive states correspond to fixed-point algebras of $\mathcal{C}$: states invariant under restriction (memory constraint), inference (generation), and update (state evolution). A* thought *is not a sequence but a fixed point of a recursive constraint-inference operator.*

14. CLIO as Sheaf Endofunctor

We lift $\mathcal{C}$ from an endofunctor on semantic states to an endofunctor on presheaves:

$$\mathcal{C} : \mathbf{PSh}(\Omega) \;\longrightarrow\; \mathbf{PSh}(\Omega),$$

acting on sections by $\mathcal{C}(s_U) = \mathbb{E}_o[\mathcal{F}(E(s_U), o)]$.

Proposition 14.1. *$\mathcal{C}$ preserves restriction maps up to stochastic equivalence, making it a weak endofunctor on the category of presheaves.*

Definition 14.2 (Obstruction Class (Sheaf Form)). Given a cover $\{U_i\}$ of Ω, the obstruction to gluing local sections $\{s_i\}$ is $[\omega] \in \check{H}^1(\Omega, \mathcal{F})$.

Theorem 14.3 (Constraint Closure Criterion). *A reasoning trajectory achieves constraint closure if and only if $\check{H}^1(\Omega, \mathcal{F}) = 0$. Successful reasoning corresponds to the existence of a global section: all surviving local computations are mutually compatible.*

Theorem 14.4 (Obstruction Descent). *Under the entropy-reduction condition, $\mathbb{E}[\|\check{H}^1(s^{(n+1)})\|] \leq \mathbb{E}[\|\check{H}^1(s^{(n)})\|]$. Thus CLIO induces a descent process on cohomological obstruction.*

Proposition 14.5 (Functoriality under Admissible Maps). *If $\phi : \mathcal{X} \to \mathcal{X}$ preserves admissibility ($\phi(\mathcal{A}_B) \subseteq \mathcal{A}_B$), then $\mathcal{C} \circ \phi \approx \phi \circ \mathcal{C}$ up to stochastic equivalence.*

We obtain a unified picture of the KV-cache model: the cache defines a presheaf of local semantic sections; reasoning generates local sections; eviction prunes sections to reduce obstruction; CLIO iterates this as a functor; and cognition converges to fixed points where obstruction vanishes.

Part V: Lyapunov Stability and Bisimulation Theorems

15. Setup and Assumptions

Let $\mathcal{X}$ be a semantic state space equipped with CLIO update $\mathcal{C} : X \to X$, energy functional $\mathcal{E} : X \to \mathbb{R}_{\geq 0}$ of the form in Definition 6.1, and a realization map $F : \mathbf{LCC} \to X$ sending LCC processes to semantic states. Let $\sim$ denote the observational equivalence relation induced by LCC.

We operate under the following assumptions:

(A1) **Observational invariance**: $P \sim Q \Rightarrow F(P)$ and $F(Q)$ have the same accessible constraint image under $\mathcal{O}_D$.

(A2) **Bisimulation-compatibility of CLIO**: $x \sim y \Rightarrow \mathcal{C}(x) \sim \mathcal{C}(y)$.

(A3) **Class-function property**: $x \sim y \Rightarrow \mathcal{E}(x) = \mathcal{E}(y)$.

(A4) **Descent property**: $\mathcal{E}(\mathcal{C}(x)) \leq \mathcal{E}(x)$ for all x, with strict inequality unless x is in the stationary set.

(A5) **Homotopy-coherent realization**: LCC transitions relevant for bisimulation yield paths $\gamma_{P,\alpha} : [0, 1] \to X$ well-defined up to homotopy.

Definition 15.1 (Admissible Fixed Point). $X^* \in \mathcal{A}_B$ is an admissible fixed point if there exists π_θ such that $\mathbb{E}[X_{t+1} \mid X_t = X^*] = X^*$ and $R(X^*)$ is locally maximal under perturbations within $\mathcal{A}_B$.

Theorem 15.2 (Constrained Convergence)**.** *Assume: (1)* $R : \mathcal{X} \to \mathbb{R}$ *is bounded and Lipschitz on* $\mathcal{A}_B$*; (2)* E_t *satisfies* $\mathbb{E}[d(E_t(X), E_t(Y))] \leq \kappa\, d(X,Y)$ *for* $\kappa < 1$*; (3)* π_θ *is updated via an unbiased policy gradient with diminishing step size. Then* $\{X_t\}$ *converges in distribution to a set of admissible fixed points.*

Proof. The contraction condition prevents unbounded growth of state variability. Boundedness and Lipschitz continuity of R ensure stability under restriction. Standard stochastic approximation results imply convergence of π_θ to a stationary policy. Together, these imply that the coupled process (X_t, π_θ) converges to stationary points where both state and policy are invariant in expectation, corresponding to admissible fixed points. □

Proposition 15.3 (Entropy Reduction under Restriction)**.** *If* E_t *removes components whose expected contribution to future reward is non-positive, then* $\mathbb{E}[S_{t+1}] \leq \mathbb{E}[S_t]$*, with strict inequality whenever non-contributing components are present.*

16. Theorem 1: Bisimulation-Invariant Lyapunov Function

Theorem 16.1 (CLIO Induces a Bisimulation-Invariant Lyapunov Function)**.** *Under assumptions A2–A4,* $\mathcal{E}$ *is a Lyapunov function on the quotient space* $\mathcal{X}/{\sim}$*, and therefore a bisimulation-invariant Lyapunov function for CLIO.*

Proof. **Step 1.** By A3, $\mathcal{E}$ descends to a well-defined function $\bar{\mathcal{E}} : \mathcal{X}/{\sim} \to \mathbb{R}_{\geq 0}$, setting $\bar{\mathcal{E}}([x]) := \mathcal{E}(x)$. This is well-defined because $x \sim y$ implies $\mathcal{E}(x) = \mathcal{E}(y)$.

Step 2. By A2, $\mathcal{C}$ induces a quotient map $\bar{\mathcal{C}} : \mathcal{X}/{\sim} \to \mathcal{X}/{\sim}$, setting $\bar{\mathcal{C}}([x]) := [\mathcal{C}(x)]$, which is well-defined because $x \sim y$ implies $\mathcal{C}(x) \sim \mathcal{C}(y)$.

Step 3. Apply A4: $\bar{\mathcal{E}}(\bar{\mathcal{C}}([x])) = \mathcal{E}(\mathcal{C}(x)) \leq \mathcal{E}(x) = \bar{\mathcal{E}}([x])$. So $\bar{\mathcal{E}}$ is non-increasing along CLIO iterates on equivalence classes. If $[x]$ is not stationary, the decrease is strict, so $\bar{\mathcal{E}}$ is a strict Lyapunov function away from the stationary set.

Therefore $\mathcal{E}$ is Lyapunov on $\mathcal{X}/{\sim}$ modulo bisimulation. □ □

Corollary 16.2 (Stationary Classes are Binding Invariants)**.** *If* $\bar{\mathcal{C}}([x^*]) = [x^*]$ *and* $\bar{\mathcal{E}}$ *is minimized at* $[x^*]$*, then* $[x^*]$ *is a binding invariant in the strong sense: a stable observational class that no admissible CLIO update can distinguish from itself.*

Interpretation 16.3. Identity is not assigned at the level of raw presentation but at the level of observational class. $\mathcal{E}$ does not care which representative is chosen, and CLIO dissipates it monotonically. What stabilizes is not a state snapshot but an equivalence class under interaction—in Haemmerlé's language, a bisimulation class; in ours, a binding invariant.

17. Theorem 2: Bisimulation as Homotopy Equivalence of Descent Trajectories

Theorem 17.1 (LCC Bisimulation = Homotopy Equivalence of Descent Trajectories)**.** *Assume A1, A2, A5. Let P, Q be LCC processes such that $P \approx Q$ (labeled bisimilarity in Haemmerlé's sense). Then the CLIO descent trajectories issued from $F(P)$ and $F(Q)$ are homotopy equivalent in the quotient dynamics: for every bisimulation-matching transition sequence from P, there is a matching sequence from Q whose realized geometric trajectory is homotopic relative endpoints in X.*

Proof. Labeled bisimulation means that if $P \xrightarrow{\alpha} P_1$, then there exists Q_1 such that $Q \overset{\alpha}{\Longrightarrow} Q_1$ and $P_1 \approx Q_1$ (Haemmerlé [9], Definition of DE-bisimulation). Apply F. By A5, the transition $P \xrightarrow{\alpha} P_1$ yields a path $\gamma_{P,\alpha} : [0,1] \to X$ from $F(P)$ to $F(P_1)$; the matched transition yields $\gamma_{Q,\alpha}$ from $F(Q)$ to $F(Q_1)$. By A1 and bisimulation, $F(P) \sim F(Q)$ and $F(P_1) \sim F(Q_1)$. Therefore in the quotient space $\mathcal{X}/{\sim}$, both paths represent the same arrow $[F(P)] \to [F(P_1)]$ up to homotopy. Iterating over all matched transitions yields a ladder diagram whose squares commute up to homotopy. By standard pasting, the full descent trajectories are homotopy equivalent relative endpoints. ☐ ☐

Corollary 17.2 (Barbed Congruence and Homotopy)**.** *Since Haemmerlé proves that labeled bisimilarity and barbed congruence coincide [9], the homotopy classification above applies equally to barbed congruence: $P \cong Q \iff F(P), F(Q)$ determine the same homotopy class of descent trajectories. This is especially important because barbed congruence is the irreversible, committed-choice notion—exactly matching the RSVP program's emphasis on projection, path dependence, and return under constraint.*

Theorem 17.3 (Quotient Descent on Observational Classes)**.** *Under A1–A5, CLIO defines a descent dynamics on the homotopy category of observational equivalence classes, and the induced $\mathcal{E}$ is a Lyapunov function on that category.*

Proof. From Theorem 16.1, $\mathcal{E}$ descends to a Lyapunov function on $\mathcal{X}/{\sim}$. From Theorem 17.1, the morphisms generated by CLIO trajectories in $\mathcal{X}/{\sim}$ are precisely homotopy classes of operationally equivalent LCC evolutions. The quotient dynamics is simultaneously operationally meaningful and geometrically coherent. ☐

Final compression of Part V. Two claims are now fully established:

> *(1) $\mathcal{E}(\mathcal{C}(x)) \le \mathcal{E}(x)$ is not merely an energy drop on raw states—it is an energy drop on bisimulation classes. The Lyapunov object is identity-respecting.*

(2) $P \approx Q$ does not merely mean two programs cannot be told apart by an observer. Under realization into the derived semantic space, it means their descent dynamics belong to the same homotopy class.

Part VI: Neural Garbage Collection as Constraint-Driven State Reduction

18. From Monotonic Context Growth to Adaptive State Restriction

Standard autoregressive language models operate under a monotonic growth assumption: the internal state X_t expands strictly with each inference step:

$$X_{t+1} \;=\; \mathcal{F}(X_t, o_t), \qquad X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots$$

This presupposes that all previously generated information remains equally available and relevant.

Neural Garbage Collection (NGC) [17] introduces a fundamental deviation. At discrete intervals, the model applies a learned restriction operator $E_t : X_t \to X_t'$, yielding:

$$X_{t+1} \;=\; \mathcal{F}(E_t(X_t), o_t),$$

breaking monotonicity. The model jointly learns *what to think* and *what to forget* under a single reward signal. The eviction decisions are not heuristic (like recency or attention thresholds) but are learned via reinforcement learning using only final task success. The most important sentence in the paper is essentially:

> *"If a model can learn to reason, why can't it learn to forget?"*

This is a conceptual inversion: previously, reasoning was learned and memory was engineered; now reasoning and memory are co-learned.

18.1. RSVP Field Representation

We reinterpret the internal state as an RSVP field:

$$X_t \;=\; (\Phi_t, \mathbf{v}_t, S_t),$$

where Φ_t is the scalar semantic potential, $\mathbf{v}_t$ is the directional inferential flow, and S_t is the unresolved entropy field. The KV cache is a discrete support over a domain Ω_t:

$$X_t \;\approx\; X|_{\Omega_t}.$$

Eviction corresponds to selection of a subdomain $K_t \subset \Omega_t$, yielding a restricted field:

$$X'_t \;=\; X|_{K_t}, \qquad X_{t+1} \;=\; \mathcal{F}(X|_{K_t}, o_t).$$

Reasoning proceeds not on the full accumulated field, but on a dynamically selected support: inference is no longer transport on a fixed field history, but transport on a *self-pruned field manifold*.

18.2. Variational Formulation

Let $R(X_T)$ denote terminal task realizability and $C(K_t)$ retained computational substrate. The joint optimization over reasoning and eviction is:

$$\max_{\{o_t, K_t\}} R(X_T) \quad \text{subject to} \quad \sum_t C(K_t) \;\leq\; B.$$

Introducing a Lagrange multiplier λ, we obtain the unconstrained functional:

$$\mathcal{J} \;=\; R(X_T) \;-\; \lambda \sum_t C(K_t).$$

In RSVP language, this is a constrained variational principle: reasoning and forgetting are coupled controls minimizing future work while preserving task closure. This is precisely the "selection determines reality" thesis: the realized trajectory is the residue of a decimation process.

18.3. Eviction as Entropic Projection

Definition 18.1 (Admissible Projection)**.** Eviction is interpreted as an approximate projection

$$E_t : X_t \;\mapsto\; \Pi_{\mathcal{A}}(X_t),$$

where $\mathcal{A}$ denotes the admissible subspace of states preserving downstream realizability under resource constraints.

NGC implements *entropic decimation*: components of the state that do not contribute to future constraint closure are removed. The realized trajectory is not the accumulation of all intermediate steps, but the residue of selective pruning.

18.4. Sheaf-Theoretic Interpretation

Let $\{U_i\}$ be a cover of the reasoning trajectory, where each U_i corresponds to a contiguous block of reasoning tokens. The KV cache encodes local sections $s_i \in \mathcal{F}(U_i)$. NGC performs a learned pruning $\{s_i\} \mapsto \{s_i\}_{i \in I_t}$ of sections that are likely to admit extension to a global section under future constraints. Eviction functions as a coarse, learned filter on the gluing problem, discarding sections unlikely to contribute to eventual consistency. This interpretation is especially natural because NGC performs eviction at the level of contiguous token blocks, which correspond to semantically coherent subcomputations occupying contiguous spans.

More precisely, eviction is an *approximate obstruction-sensitive restriction functor*: the model is learning, however crudely, which local sections need not be preserved for later consistency. The stronger formulation is that a reasoning process should preserve precisely those local sections whose omission would change the obstruction class of the evolving global solution.

18.5. Policy Unification and Single-Signal Optimization

Let $\tau = \{(o_t, K_t)\}_{t=1}^{T}$ be a trajectory. The policy gradient objective is:

$$\nabla_\theta \, \mathbb{E}_{\tau \sim \pi_\theta}[R(\tau)].$$

There is no decomposition into separate objectives for reasoning and memory. Consequently, the model learns a joint strategy in which retained state determines future reasoning, reasoning determines reward, and reward updates both retention and reasoning. This *closes the loop* between inference and state formation—the model becomes a self-editing dynamical system.

18.6. Budget-Aware Interoception as Field-Awareness

NGC prepends the eviction rate to the prompt so the model has explicit awareness of its own resource constraint before reasoning. In RSVP terms, this is not a superficial prompt trick but a crude form of *endogenous field-awareness*: the system is told about its own thermodynamic boundary condition before transport begins. This moves it one step closer to a model that reasons as a function of its own internal constraint state.

18.7. The Replay Mask and Path-Faithful Admissibility

The replay-mask mechanism ensures that once the model has evicted parts of its history, training respects the exact visibility structure induced by those evictions. This is very close to the insistence that truth is not local plausibility but fixed-point realizability under the actual constraints of the path taken. A reasoning trace cannot be judged by a context it never possessed. The replay mask is a path-faithful reconstruction of the admissibility conditions under which each token was produced.

18.8. NGC's Empirical Lyapunov Signal

NGC's internal uncertainty functional behaves empirically as a Lyapunov function. The paper observes that the gradient of uncertainty with respect to time is negative for correct trajectories and positive for incorrect ones. This is not merely a heuristic observation: it is the defining property of a Lyapunov function. We can therefore identify

$$\text{uncertainty}(t) \;\approx\; \mathcal{E}(x_t),$$

and the empirical result becomes

$$\frac{d}{dt}\mathcal{E}(x_t) \;<\; 0 \quad \text{on successful trajectories.}$$

Proposition 18.2 (Empirical Lyapunov Principle)**.** *The internal uncertainty functional measured by NGC behaves as a Lyapunov function for successful reasoning trajectories, providing empirical grounding for Theorem 16.1.*

18.9. Graph Aggregation as Computational Homotopy Colimit

NGC's graph-based aggregation constructs a graph of reasoning paths, clusters them, and merges them into a final structure. This is structurally:

$$\text{tree expansion} \to \text{selection} \to \text{collapse.}$$

In the homotopy-theoretic picture:

- Different paths → different local trajectories (local sections).
- Clustering → identification of equivalent reasoning (homotopy identification).
- Summarization → quotienting (homotopy colimit).

NGC's graph reduction is therefore a computational version of taking a homotopy colimit:

$$\text{stable CLIO reasoning graph} \;\equiv\; \text{bisimulation equivalence class.}$$

18.10. Toward Semantic Renormalization

While NGC operates via binary retention decisions, the formalism suggests a broader class of operators. Instead of selecting subsets K_t, one may consider *semantic renormalization operators*:

$$M_t : X_t \;\mapsto\; \tilde{X}_t$$

that perform semantic compression, merging, or coarse-graining while preserving invariants relevant to realizability:

$$\mathcal{I}(M_t(X)) \;=\; \mathcal{I}(X).$$

The generalized CLIO operator is:

$$\mathcal{C}_M(X) \;=\; \mathbb{E}_{o\sim\pi_\theta(\cdot|M(X))}[\mathcal{F}(M(X),o)],$$

supporting compression, merging, and abstraction rather than binary deletion. This defines a renormalization flow over semantic states, in which cognition operates by preserving invariants while reducing representational complexity:

$$\text{evict} \;\subset\; \text{restrict} \;\subset\; \text{quotient} \;\subset\; \text{renormalize}.$$

Binary garbage collection is only the crudest member of this family. A more mature cognitive architecture would perform *homotopy-aware renormalization*, preserving invariants while collapsing redundant or obstruction-producing structure.

The deepest way to state the significance of NGC inside the present framework is this: it is an early empirical example of a system in which cognition is being forced to acknowledge that persistence is costly and that not all locally generated structure deserves continued ontological status. What survives is what remains useful under constraint. That is a weak form of the broader claim that realization is the residue left after incompatible structure has been decimated.

19. Worked Example: Binding and Rupture in a Finite Constraint System

We present a minimal example illustrating the Binding Invariant Theorem concretely, exhibiting a system with a non-trivial binding invariant, a closure event that preserves it, and a closure event that destroys it.

19.1. The Constraint System

Let $T = \{t_0 \leq t_1 \leq t_2\}$. Define:

$$\mathcal{C}_{t_0} = \{A\}, \quad \mathcal{C}_{t_1} = \{A, B\}, \quad \mathcal{C}_{t_2} = \{A, B, C\},$$

with restriction maps $\rho_{t_1 t_0}(A, B) = A$ and $\rho_{t_2 t_1}(A, B, C) = (A, B)$. The constraint history at t_2 is:

$$\mathcal{H}_{t_2}(S) \;=\; \varprojlim\{\mathcal{C}_{t_0} \leftarrow \mathcal{C}_{t_1} \leftarrow \mathcal{C}_{t_2}\} \;=\; \{(A,\ (A,B),\ (A,B,C))\}.$$

A valid trajectory $\sigma(t_i) = \bigwedge_{j \le i} \mathcal{C}_{t_j}$ exists: $\Gamma(T, \mathcal{F}) \neq \varnothing$, hence $\mathcal{B}_{t_2}(S) \neq \varnothing$.

19.2. Preserving Closure Event

Adjoining D consistent with (A, B, C): $\mathcal{C}_{t_2^+} = \{A, B, C, D\}$. The binding section extends, $\mathcal{B}_{t_2^+}(S) \neq \varnothing$, and the system continues.

19.3. Rupture Event

Imposing $\mathcal{C}_{t_2'} = \{A, B, \neg A\}$: since $A \wedge \neg A = \bot$, no consistent state exists, no trajectory extends to t_2', and $\mathcal{B}_{t_2'}(S) = \varnothing$. By the Rupture Criterion (Corollary 3.5), this is replacement, not continuation.

19.4. Partial Rupture / Ambiguity

Imposing $\mathcal{C}_{t_2''} = \{A, B, C \vee E\}$: multiple extensions exist, global sections exist but are non-unique, and $\mathcal{B}_{t_2''}(S)$ has multiple elements. This is *indeterminate continuation*: the system persists, but its future identity is underdetermined.

19.5. Diagrammatic Form

The three cases correspond to the existence, non-existence, and multiplicity of compatible cones over the inverse system.

Binding (normal continuation). The inverse limit $\mathcal{H}_{t_2}(S)$ is the apex of a compatible cone:

$$\begin{tikzcd}
\mathcal{H}_{t_2}(S) \arrow[d, dashed, "\pi_2"'] \arrow[rd, dashed, "\pi_1"] & & \\
\{A, B, C\} \arrow[r, "\rho_{t_2 t_1}"'] & \{A, B\} \arrow[r, "\rho_{t_1 t_0}"'] & \{A\}
\end{tikzcd}$$

A compatible binding section exists; $\mathcal{B}_{t_2}(S) \neq \varnothing$.

Rupture. Imposing $\neg A$ alongside A destroys compatibility:

$$\begin{tikzcd}
\{A, B, \neg A\} \arrow[d, "\rho_{t_2' t_1}"'] \arrow[rd, "\rho_{t_2' t_0}"] & \\
\{A, B\} \arrow[r, "\rho_{t_1 t_0}"'] & \{A\}
\end{tikzcd}$$

No object $\mathcal{H}_{t_2'}(S)$ completes this as an inverse limit: any putative section would require $A \wedge \neg A = \bot$. Hence $\mathcal{B}_{t_2'}(S) = \varnothing$.

Ambiguity. The disjunctive constraint $C \vee E$ admits multiple extensions:

$$\begin{tikzcd}
\{A, B, C \vee E\} \arrow[d, "\rho_{t_2''t_1}"'] \arrow[dr, "\rho_{t_2''t_0}"] & \\
\{A, B\} \arrow[r, "\rho_{t_1t_0}"'] & \{A\}
\end{tikzcd}$$

Compatible cones exist, but not uniquely: $\mathcal{B}_{t_2''}(S)$ is non-empty and non-singleton. The system persists with underdetermined continuation.

Remark 19.1. These diagrams make visible the core distinction of the theory. Preservation corresponds to extendability of cones. Rupture corresponds to non-existence of a compatible cone. Ambiguity corresponds to multiplicity of cones. The binding invariant is therefore best understood not as a substance carried through time, but as the existence and structure of a limiting cone over the system's constraint history. This finite example is the discrete analogue of cohomological obstruction: incompatibility of constraints corresponds to non-vanishing obstruction classes, while compatible extensions correspond to trivialization.

20. Open Problems

Three problems would convert the framework's correspondences into fully tight theorems.

Problem 1 (Lyapunov stability). Prove that CLIO updates induce a bisimulation-invariant Lyapunov function $V(x) = \mathcal{E}(x) - \mathcal{E}(x^*)$ satisfying $V(\mathcal{C}(x)) \le V(x)$ for all x, with equality iff $x = x^*$, without requiring the assumptions A2–A4 as inputs—i.e., derive them from RSVP field dynamics.

Problem 2 (Fisher metric). Define a natural Fisher information metric on $\mathfrak{M}$ making the gradient flow of $\mathcal{E}$ explicit as a natural gradient flow in the sense of Amari. The entropy term $\mu S(x)$ suggests $\mathfrak{M}$ carries natural information geometry; the Fisher metric would make the gradient flow $\partial_t x = -\nabla \mathcal{E}(x)$ the most efficient descent with respect to the geometry of distributions.

Problem 3 (LCC bisimulation as homotopy—converse). The forward direction of Theorem 17.1 is proved. The converse—that homotopy equivalence of descent trajectories implies LCC bisimulation—requires a realizability result connecting the derived-stack semantics of $\mathfrak{M}$ to the operational semantics of LCC.

A solution to all three would unify process algebra, derived algebraic geometry, and learning dynamics into a single formal object: the metric space of CLIO-observational equivalence classes under Fisher-information gradient flow.

21. Conclusion

We have developed the binding invariant through six interlocking formalisms.

Part I (sheaf and topos theory): the binding invariant is the global section of the subobject classifier over the constraint history topos. The Binding Invariant Theorem establishes four equivalent conditions; three corollaries characterize uniqueness, rupture, and stability.

Part II (derived algebraic geometry and variational dynamics): ambiguity is higher homotopy in a derived semantic stack; the self is the minimum-energy homotopy class of a multi-scale cohomological energy functional; synthetic grounding is gradient descent; return-under-constraint is the geometric stability of the fixed point. Identity is dynamical equilibrium, not static persistence.

Part III (process algebra): LCC constraint consumption is cohomological obstruction projection; bisimulation is gradient-flow equivalence; barbed congruence is invariance under irreversible descent. The Triple Characterization Theorem (Theorem 11.3) proves equivalence of all three levels.

Part IV (monoidal category theory): CLIO is a lax monoidal endofunctor on **Sem**; cognitive states are fixed-point algebras of CLIO; the sheaf-theoretic lift makes CLIO an obstruction-descending functor on $\mathbf{PSh}(\Omega)$.

Part V (formal theorems): $\mathcal{E}$ is a bisimulation-invariant Lyapunov function (Theorem 16.1); LCC bisimulation corresponds to homotopy equivalence of descent trajectories (Theorem 17.1); quotient descent on observational classes is simultaneously operationally meaningful and geometrically coherent (Theorem 17.3).

Part VI (concrete instantiation): NGC is a concrete empirical case of CLIO dynamics, its "grow-then-evict" mechanism is admissible-subspace projection in RSVP field coordinates, its internal uncertainty is an empirical Lyapunov signal, its graph aggregation is a computational homotopy colimit, and its limitation—binary retention—points toward semantic renormalization as the next architecture.

The deepest result is the convergence of all six levels on the same object. The binding invariant is not an artifact of any one formal choice but the invariant structure that appears at every level of description of a system that undergoes irreversible change. The framework is not speculative: it is the correct mathematical language for what NGC is already doing. And the framework is not merely interpretive: it generates three open problems whose resolution would close the loop between process algebra, derived geometry, and machine learning.

> *The binding invariant is the minimum-energy homotopy class that survives irreversible projection—the observational equivalence class of a system under cohomological descent on its constraint structure. A system is "itself" if it lies in a homotopy class invariant under descent on a cohomological energy functional defined by its constraints.*

Acknowledgments. This essay develops ideas first articulated in *The Compiled Self*, *Return Under Constraint*, and the Yarncrawler framework. Parts III–V draw on Haemmerlé's formalization of Linear Concurrent Constraints. Part VI engages with Li, Hamid, Fox, and Goodman's Neural Garbage Collection.

Keywords. binding invariant, constraint closure, irreversibility, sheaf theory, topos theory, derived stacks, ∞-groupoids, cohomological energy functional, gradient flow, synthetic grounding, CLIO, RSVP field theory, Yarncrawler, linear concurrent constraints, bisimulation, barbed congruence, observational equivalence, neural garbage collection, semantic renormalization, monoidal endofunctor, fixed-point algebra.

References

[1] S. Mac Lane, *Categories for the Working Mathematician*, Springer, 2nd ed., 1998.

[2] P. T. Johnstone, *Sketches of an Elephant: A Topos Theory Compendium*, Oxford University Press, 2002.

[3] J. Lurie, *Higher Topos Theory*, Princeton University Press, 2009.

[4] E. Riehl, *Category Theory in Context*, Dover, 2016.

[5] D. I. Spivak, *Category Theory for the Sciences*, MIT Press, 2014.

[6] J.-Y. Girard, Linear Logic, *Theoretical Computer Science*, 50 (1987), 1–102.

[7] V. A. Saraswat, *Concurrent Constraint Programming*, MIT Press, 1993.

[8] V. A. Saraswat and P. Lincoln, Higher-Order, Linear, Concurrent Constraint Programming, Technical Report, Xerox PARC, 1992.

[9] R. Haemmerlé, Observational Equivalences for Linear Logic Concurrent Constraint Languages, *Theory and Practice of Logic Programming*, 11(4–5): 469–485, 2011. Also available as arXiv:1108.0329.

[10] R. Milner, *Communication and Concurrency*, Prentice Hall, 1989.

[11] J. Parrow and B. Victor, The Fusion Calculus: Expressiveness and Symmetry in Mobile Processes, *Proceedings of LICS*, 1998.

[12] R. Landauer, Irreversibility and Heat Generation in the Computing Process, *IBM Journal of Research and Development*, 1961.

[13] C. H. Bennett, The Thermodynamics of Computation, *International Journal of Theoretical Physics*, 1982.

[14] S. Lloyd, Ultimate Physical Limits to Computation, *Nature*, 406 (2000), 1047–1054.

[15] T. Jacobson, Thermodynamics of Spacetime: The Einstein Equation of State, *Physical Review Letters*, 75 (1995), 1260–1263.

[16] E. Verlinde, On the Origin of Gravity and the Laws of Newton, *Journal of High Energy Physics*, 2011.

[17] M. Y. Li, J. I. Hamid, E. B. Fox, and N. D. Goodman, Neural Garbage Collection: Learning to Forget while Learning to Reason, arXiv:2604.18002, 2026.

[18] S. Abramsky and N. Tzevelekos, Introduction to Categories and Categorical Logic, in *New Structures for Physics*, Springer, 2011.

[19] R. S. Sutton, The Bitter Lesson, 2019.