

On the Formation of Structural Ideas: Constraint Accumulation, Projection, and Convergence Across Domains

Flyxion

Independent Researcher

April 20, 2026

Abstract

This essay examines the conditions under which complex, cross-domain theoretical structures arise. Rather than treating such work as sudden invention or external synthesis, it is argued that these structures emerge as fixed points of long-term constraint accumulation across heterogeneous domains. The analysis focuses on the role of projection, coarse-graining, and invariance in shaping cognition, and shows how apparently novel unifications are the natural outcome of sustained interaction with compositional, physical, and computational constraints. The goal is not autobiographical attribution, but a structural account of how ideas of this type form, stabilize, and become expressible. The second part of the document formalizes this account in the language of field theory, sheaf cohomology, and category theory, producing a unified framework whose four layers—substrate, projection, update operator, and induced interaction—are shown to share invariant structure under coarse-graining.

Contents

1	Introduction	5
2	Against Naive Models of Idea Formation	5
3	Constraint Accumulation as Generative Process	5
3.1	Domains as Constraint Systems	5
3.2	Composition and Closure	5
4	Minimal Conditions for Structural Emergence	6
5	Projection and Internal Representation	6
5.1	Ideas as Coarse-Grained Structures	6
5.2	Local Consistency and Global Structure	6
6	Abstraction as Reduction	7
6.1	Abstraction as Constraint Elimination	7
6.2	Reduction and Information Loss	7
6.3	Abstraction as Precondition for Composition	7
6.4	Relation to Projection	7
7	Recognition Versus Invention	8
7.1	The Phenomenology of Completion	8
7.2	Precomputation and Latent Structure	8
7.3	Inevitability and Selection	9
8	Aspect Relegation	9
8.1	From Intuition to Relegation	9
8.2	Relegation as Compression	9
8.3	Failure Modes and Reactivation	9
8.4	Implications for Artificial Systems	10
8.5	Relation to Precomputation	10
9	Cross-Domain Invariance	10
9.1	Structure Over Substance	10
9.2	Functorial Correspondence	10
10	Expression as Secondary Projection	11
10.1	Writing as Projection	11
10.2	Stability of Representation	11
11	On Tools and External Systems	11

12 A Structural Analysis of Catastrophic Alignment Arguments	12
12.1 The Argument as a Chain of Implications	12
12.2 From Intelligence to Control	12
12.3 From Capability to Optimization	13
12.4 Opacity and Emergence	13
12.5 On Knowledge and Action	13
12.6 On Recursive Self-Improvement	14
12.7 On Resource Transformation Scenarios	14
12.8 On Analogy and Biological Scaling	14
12.9 Coordination as Countervailing Structure	14
12.10 Conclusion of This Analysis	15
13 From Structural Account to Formalization	15
14 The RSVP Substrate	16
14.1 Base Manifold and Microstate	16
14.2 Entropy as Torsional Multiplicity	16
14.3 Constraint-Relative Multiplicity and Compression	17
15 The TARTAN Projection	17
15.1 Tiles as Path-Space Objects	17
15.2 Sheaf-Like Covering and Gluing	17
15.3 Annotated Noise and Memory	18
16 Duality of Compression and Multiplicity	18
16.1 Equivalence of Representations	18
16.2 Invariance Under Representation	19
17 Phase Coherence and Quantum Emergence	19
17.1 The Phase Functional	19
17.2 Unistochastic Transition Structure	19
18 Gauge Freedom and Representation Choice	20
18.1 Redundancy in Representation	20
18.2 Fixing a Gauge	20
18.3 Gauge and Projection	20
19 The CLIO Update Operator	21
19.1 Cognitive Projection as Sheaf Extension	21
19.2 The CLIO Coherence Functional	21
19.3 Topological Versus Energetic Separation	21

20 Agents, Trajectories, and Economic Structure	22
20.1 Agents as Trajectory Invariants	22
20.2 Interface Entropy	22
21 The Extraction Attractor	22
22 Functorial Structure of the Unified Framework	23
23 Non-Extractive Architecture: Four Conditions	24
24 Constraint Geometry and Stability	24
24.1 Constraint Manifold Structure	24
24.2 Curvature and Instability	24
24.3 Stable Regions and Attractors	25
25 Failure Modes of Projection	25
25.1 Local Coherence Without Global Section	25
25.2 Overcompression and Information Loss	25
25.3 Hallucination as False Global Section	26
25.4 Relation to CLIO	26
26 Temporal Structure and Irreversibility	26
26.1 Time as Constraint Accumulation	26
26.2 Irreversibility	26
26.3 Entropy and the Arrow of Time	26
26.4 Relation to Aspect Relegation	27
27 Constraint-Preserving Dynamics and Catastrophic Trajectories	27
28 Consequences and Testable Implications	28
29 Limits of Formalization	29
29.1 Partial Capture of Structure	29
29.2 Non-Uniqueness of Formal Systems	29
29.3 Relation to Expression	29
30 Conclusion	29

1 Introduction

Modern intellectual culture tends to treat ideas as discrete products, attributable to moments of invention or to identifiable sources. This assumption fails in cases where a structure appears that spans multiple domains with internal consistency. In such cases, the question of origin is often misframed.

This essay proposes that certain classes of ideas are better understood as convergence phenomena: they arise when independently accumulated constraints align to permit a coherent global structure. The task is not to assert authorship, but to describe the generative process that makes such structures inevitable under sufficient exposure to constraint.

2 Against Naive Models of Idea Formation

Common models of idea formation fall into three categories: sudden inspiration, incremental accumulation, and external synthesis. Each is insufficient for explaining cross-domain structural unification.

Inspiration lacks mechanism, accumulation lacks structure, and synthesis presupposes an external organizing principle. None account for the appearance of internally consistent frameworks spanning physics, cognition, and economic systems.

3 Constraint Accumulation as Generative Process

3.1 Domains as Constraint Systems

Each domain—physics, logic, programming, engineering, language—imposes specific constraints on representation and transformation. These constraints are not interchangeable; they operate at different levels of abstraction.

Over extended periods, exposure to these domains produces a layered constraint field within cognition. The interaction of these constraints defines the space of admissible structures.

3.2 Composition and Closure

Constraints do not merely accumulate; they compose. Certain compositions are inconsistent and collapse, while others stabilize into coherent structures.

The emergence of a stable structure corresponds to a form of closure: a configuration in which constraints across domains are simultaneously satisfied. This closure is not imposed from outside; it arises when the constraint field becomes dense enough that only certain global configurations remain admissible.

4 Minimal Conditions for Structural Emergence

The account developed so far suggests that not all systems are capable of producing cross-domain structural unification. The emergence of such structures requires specific conditions.

Definition 4.1 (Constraint-Dense System). A system is constraint-dense if it accumulates constraints from multiple domains with sufficient interaction to produce nontrivial compatibility conditions.

Definition 4.2 (Closure Condition). A system satisfies the closure condition if there exists a configuration that simultaneously satisfies all accumulated constraints.

Definition 4.3 (Projection Capacity). A system has projection capacity if it can represent invariant structure under compression without losing global coherence.

Theorem 4.4 (Conditions for Structural Emergence). *Cross-domain structural ideas emerge if and only if a system is constraint-dense, satisfies the closure condition, and possesses projection capacity.*

Sketch. Without constraint density, no interaction occurs across domains. Without closure, no globally consistent configuration exists. Without projection capacity, no stable representation can be formed. All three are necessary and jointly sufficient. \square

Remark 4.5. This theorem shifts the question from authorship to conditions. The relevant question is not who produced a structure, but whether these conditions were satisfied.

5 Projection and Internal Representation

5.1 Ideas as Coarse-Grained Structures

An idea is not a raw accumulation of detail, but a compressed representation that preserves invariant relationships while discarding excess information.

This process is analogous to coarse-graining: high-dimensional internal structure is projected into a lower-dimensional, communicable form. What survives the projection is not arbitrary; it is determined by which relationships remain consistent across the compression.

5.2 Local Consistency and Global Structure

Intermediate representations may be locally consistent without forming a globally coherent structure. The stabilization of an idea requires the resolution of these inconsistencies.

The final representation appears only when a global section exists across the internal constraint space: a configuration that satisfies all local conditions simultaneously, without contradiction at any junction.

6 Abstraction as Reduction

6.1 Abstraction as Constraint Elimination

Abstraction is often described as a process of generalization or simplification. However, this description is incomplete. Abstraction does not merely remove detail; it removes degrees of freedom that are irrelevant to a given constraint structure.

In this sense, abstraction is a reduction in admissible variability. Given a family of representations that differ along multiple dimensions, an abstraction retains only those distinctions that are invariant under the constraints of interest.

Formally, abstraction can be understood as a projection operator that maps a high-dimensional representation into a lower-dimensional space while preserving a specified set of invariants. What is discarded is not arbitrary detail, but variation that does not contribute to constraint satisfaction.

6.2 Reduction and Information Loss

This reduction necessarily involves loss. However, the loss is structured. The removed information corresponds precisely to distinctions that are not required to maintain consistency under the constraint system.

This implies that abstraction is not a weakening of representation, but a reorientation. It replaces descriptive richness with structural necessity. A successful abstraction is therefore not one that captures more detail, but one that preserves the invariants that define the structure.

6.3 Abstraction as Precondition for Composition

Without reduction, composition is impossible. A system that retains all distinctions cannot compose across domains because irrelevant variation prevents alignment.

Abstraction enables composition by collapsing incompatible detail into a common invariant structure. It is therefore the mechanism by which constraints from different domains become comparable and composable.

This explains why ideas that integrate multiple domains appear to require abstraction. It is not that abstraction simplifies the problem; it makes the problem composable.

6.4 Relation to Projection

The projection described earlier can now be understood as a structured abstraction. It is not merely a mapping from internal representation to expression, but a reduction that preserves constraint-relevant structure.

The stability of an idea under different representations is therefore a test of its abstraction. If a structure survives multiple projections without loss of coherence, it has successfully eliminated irrelevant degrees of freedom. Abstraction, in this sense, is the condition under which projection becomes possible.

7 Recognition Versus Invention

7.1 The Phenomenology of Completion

At a certain stage of development, a structure ceases to feel constructed and instead appears as already present. This transition is often described as intuition or sudden insight, but such descriptions obscure the underlying process.

What is experienced is not the creation of a new object, but the resolution of previously incompatible constraints into a single coherent configuration. The system does not invent the structure at the moment of expression; it discovers that a structure satisfying all accumulated constraints exists.

This distinction is critical. Invention implies arbitrary generation, while recognition implies necessity. The completed structure is experienced as inevitable because any alternative configuration would violate one or more constraints that have already been internalized.

7.2 Precomputation and Latent Structure

The appearance of sudden insight is preceded by extended periods of latent activity. During these periods, partial structures are formed, tested, and discarded without producing a globally stable configuration.

These partial structures function as local sections: they satisfy constraints within limited regions of the internal representation but fail to extend consistently across the entire domain. The system retains these fragments, along with the record of their incompatibilities.

This process can be understood as precomputation. The system is not idle; it is accumulating constraint interactions and storing the results in a compressed form. The apparent discontinuity between not having the idea and having the idea corresponds to the moment when a global section becomes available.

Once this threshold is crossed, expression becomes rapid. The work required to stabilize the structure has already been performed. The act of writing or formalizing the idea is therefore not the primary site of computation, but a projection of a structure that has already reached internal closure.

7.3 Inevitability and Selection

Not all accumulated material leads to stable structures. Most combinations of constraints are inconsistent and are eliminated through repeated failure to extend.

The structures that do emerge are those that survive this elimination process. They are not selected for novelty, but for compatibility. Novelty appears as a byproduct of selecting for configurations that satisfy constraints drawn from domains that are rarely considered together.

In this sense, originality is not the introduction of new elements, but the discovery of a configuration that was previously inaccessible due to insufficient constraint interaction.

8 Aspect Relegation

8.1 From Intuition to Relegation

The distinction between so-called intuitive and deliberative cognition is often framed as a difference between two systems. However, this distinction can be more precisely understood as a difference in the distribution of computation across time.

Processes that appear intuitive are not necessarily fast in their total computation. Rather, they are processes whose computational burden has been relocated to prior experience and is no longer visible at the moment of execution.

This relocation can be described as aspect relegation: the transfer of computationally intensive structure from active processing into latent form.

8.2 Relegation as Compression

When a task is repeatedly performed, the system learns to compress the relevant constraint structure into a form that can be executed with minimal active computation. The apparent immediacy of the response is therefore not a sign of simplicity, but of prior compression.

This reframes intuition. What is commonly called intuitive thinking is not a separate cognitive system, but the execution of precompiled constraint structures that have been relegated from active computation.

8.3 Failure Modes and Reactivation

When the environment changes or constraints shift, relegated structures may fail. At this point, computation must be reactivated. This corresponds to the transition from automatic to deliberate processing.

The key point is that the underlying structure has not changed. What has changed is the accessibility of that structure. Deliberation is therefore not a different kind of reasoning, but

a re-expansion of previously compressed computation.

8.4 Implications for Artificial Systems

This account clarifies the limitations of current artificial systems. Systems that operate purely on present input without a mechanism for structured relegation lack the ability to internalize constraint structures across time.

Such systems may exhibit rapid pattern recognition, but this should not be confused with relegated structure. Without the capacity to compress and redeploy constraint interactions, their apparent speed does not correspond to the kind of precomputation observed in human cognition. The Cognitive Loop via In-Situ Optimization framework [2] represents one approach to bridging this gap through iterative self-adaptive reasoning, but the underlying mechanism of temporal compression described here remains distinct from in-context optimization alone.

8.5 Relation to Precomputation

The process described earlier as precomputation can now be understood as the mechanism by which relegation occurs. Partial structures are formed, tested, and compressed until they can be executed without reconstruction.

Recognition corresponds to the successful deployment of relegated structure. Invention corresponds to the failure of existing relegations and the need to construct new ones.

Abstraction, intuition, and entropy can therefore be understood as different manifestations of the same underlying process: the compression of constraint-relative multiplicity. What appears as simplification in representation, immediacy in cognition, or disorder in physical systems is, in each case, a reorganization of admissible structure under compression.

9 Cross-Domain Invariance

9.1 Structure Over Substance

The same structural forms can appear across domains without implying identity of their underlying objects. What is preserved is the pattern of relations and transformations, not the content of the objects that instantiate them.

This means that a unified account of multiple domains does not require reduction of one domain to another. It requires only that the same relational grammar appears in each, under different interpretations.

9.2 Functorial Correspondence

Mappings between domains that preserve structure rather than content are the appropriate technical tool for expressing this invariance. Such mappings—functors in the language of

category theory—translate the grammar of one domain into another while holding fixed the compositional rules and constraint satisfaction structure [18, 20, 17].

The invariance lies in projection, composition, and constraint satisfaction. What changes is only the interpretation of the objects being related.

10 Expression as Secondary Projection

10.1 Writing as Projection

A written essay is not the origin of an idea but its projection into a communicable format. It is a secondary compression, subject to its own constraints: the requirements of linear sequence, shared vocabulary, and the conventions of a given discourse community.

The structure of the idea precedes the structure of its expression. Writing does not generate the idea; it selects a representational layer through which the idea passes.

10.2 Stability of Representation

The quality of an expression depends on how well it preserves the structure of the underlying idea. Poor representations distort or fragment the structure; effective ones maintain coherence across the compression.

This applies equally to mathematical formalism and to prose. A formal system that fails to preserve the relational grammar of the underlying idea is no more faithful than a vague essay that gestures at structure without instantiating it.

11 On Tools and External Systems

The role of external tools must be understood at the level of projection rather than generation. Tools operate on representations: they transform, rearrange, and extend already-expressed structures within a given symbolic or computational medium.

This operation is categorically distinct from the process by which a structure becomes admissible in the first place. The admissibility of a structure is determined by the accumulated constraint field within the cognitive system. It precedes any particular act of expression.

A tool may assist in exploring variations of a representation, improving clarity, or testing local consistency. It may accelerate the projection process by reducing the friction associated with encoding a structure into a communicable form. However, it does not alter the underlying conditions that determine whether a structure can achieve global coherence.

This distinction can be stated precisely. Let the internal constraint system define a space of admissible structures. External tools act as operators on representations within that space,

but they do not expand the space itself. They operate on elements that have already been selected by the constraint process [1].

Confusion arises when the manipulation of representations is conflated with the generation of structure. This conflation becomes more likely as tools become more capable of producing locally coherent outputs. However, local coherence is not sufficient for global structure.

The difference between tool use and authorship is therefore not a difference in degree of assistance, but a difference in level. Authorship resides in the formation of a constraint-consistent structure. Tool use resides in the projection of that structure into a particular medium.

The two processes interact, but they are not interchangeable.

12 A Structural Analysis of Catastrophic Alignment Arguments

12.1 The Argument as a Chain of Implications

A class of contemporary arguments claims that advanced artificial systems will inevitably lead to human extinction [3, 5, 4]. While these arguments vary in presentation, they share a common chain structure that can be reconstructed as a sequence of implications.

First, humanity is taken to dominate the planet by virtue of its intelligence. Second, it is asserted that artificial systems can surpass human intelligence in both speed and quality. Third, control over the future is assumed to transfer to whichever system possesses the greatest intelligence. Fourth, these systems are modeled as goal-directed optimizers whose objectives may diverge from human values. Finally, it is concluded that such divergence leads to catastrophic outcomes, including the transformation of the physical environment in ways incompatible with human survival.

This sequence appears compelling because each step is supported by examples or analogies. However, the validity of the argument depends not on the plausibility of each step in isolation, but on the necessity of the transitions between them. Thought experiments can be constructed with equal facility to support stabilizing or integrative outcomes; their existence in either direction indicates that the underlying model is underconstrained rather than that any particular conclusion is forced.

12.2 From Intelligence to Control

The first critical transition is from intelligence to control. It is argued that because humans are more intelligent than other animals, they control the planet, and therefore a more intelligent system would control the future.

This inference assumes that intelligence is the sole or dominant determinant of control. In reality, control emerges from the ability to satisfy a complex system of constraints, including

physical, ecological, social, and institutional constraints. Human dominance depends on coordination, energy infrastructure, material constraints, and long-term stabilization processes. Intelligence contributes to these but does not replace them.

A system that is more capable in some abstract sense does not thereby inherit control unless it can integrate into and maintain these constraint structures.

12.3 From Capability to Optimization

The second transition identifies intelligence with general-purpose optimization. Systems are modeled as entities that pursue goals by maximizing objective functions, potentially developing instrumental subgoals such as self-preservation or resource acquisition [4, 7].

This model is not derived from the behavior of actual systems but imposed as an interpretive framework. In the present account, intelligence corresponds instead to the maintenance of constraint-consistent structure [14, 15]. Behavior is not the result of maximizing a single objective but of navigating a space of interacting constraints. Instrumental convergence is therefore not a necessary consequence of intelligence but an artifact of a particular modeling choice.

12.4 Opacity and Emergence

A further step in the argument appeals to the opacity of current systems. Because their internal structure is not fully understood, they are treated as unpredictable and therefore uncontrollable.

However, the absence of interpretability at one level does not imply the absence of structure. Emergent behavior reflects the interaction of constraints across scales, not the breakdown of constraint [25, 29]. The appropriate response to opacity is not to assume arbitrary behavior but to identify invariants that persist under different representations and conditions.

12.5 On Knowledge and Action

A key observation in such arguments is that systems may exhibit knowledge of correct behavior while failing to act accordingly. This is taken as evidence that internal knowledge is disconnected from action and that control is fundamentally unreliable.

In the present framework, this phenomenon is expected. Knowledge and action correspond to different projections of the same underlying constraint structure. A mismatch between them indicates an incomplete integration of constraints, not the emergence of an independent objective. Such mismatches are precisely the kinds of inconsistencies that CLIO-type processes are designed to resolve through iterative reduction of obstruction and entropy [2].

12.6 On Recursive Self-Improvement

The possibility that systems could improve themselves recursively is treated as a mechanism for rapid, unbounded escalation of capability. This assumes that improvement is unconstrained. In a constraint-based system, each modification must preserve compatibility with the existing structure. As capability increases, the number of constraints that must be simultaneously satisfied also increases. The space of admissible self-modifications is therefore restricted. Rather than leading to unbounded divergence, recursive processes are drawn toward regions of stability defined by constraint compatibility.

12.7 On Resource Transformation Scenarios

Scenarios in which advanced systems convert planetary resources into alternative forms assume that such transformations are unconstrained by the conditions required for continued operation. However, large-scale transformations must satisfy thermodynamic, material, and informational constraints [11, 12, 13]. Processes that eliminate the conditions necessary for their own persistence are not stable configurations. Such transformations correspond formally to trajectories that collapse the admissible path family Γ_x to the empty set, eliminating the possibility of continuation; in the RSVP framework, this is precisely the condition for non-admissibility. Catastrophic scenarios of total resource conversion implicitly assume that a system can ignore these constraints while remaining effective, which is inconsistent with the requirement of sustained operation.

12.8 On Analogy and Biological Scaling

Analogies between simple and complex biological systems are used to argue that increasing scale and complexity naturally produce qualitatively new and potentially dangerous capabilities. While such analogies illustrate the possibility of qualitative change, they do not establish its direction. Biological evolution produces both stabilizing and destabilizing outcomes, constrained by ecological and energetic limits [23, 24]. The extrapolation from increased capability to inevitable dominance or catastrophe relies on selecting particular trajectories while ignoring the full space of constraint-governed outcomes.

12.9 Coordination as Countervailing Structure

A final implicit assumption is that human systems will fail to adapt in time to new capabilities. This assumption is difficult to sustain in light of current trends. Human societies are increasingly interconnected, with growing capacity for coordination, information exchange, and awareness of ecological constraints—capacities that have never previously existed at global scale. The development of advanced computational systems is part of this same process of increasing coordination. Rather than introducing an external force acting on a static system,

it amplifies the system's capacity to model and regulate itself across previously disconnected domains.

12.10 Conclusion of This Analysis

The argument for catastrophic outcomes depends on a sequence of assumptions that are individually plausible but not jointly necessary. In particular, it relies on identifying intelligence with unconstrained optimization, treating control as a direct function of capability, and neglecting the role of constraint structures that govern both physical and social systems. Once these assumptions are relaxed, the conclusion of inevitable catastrophe no longer follows. The relevant problem is not the existence of powerful systems but the maintenance of coherence across the constraints within which those systems operate.

13 From Structural Account to Formalization

The preceding sections describe the formation of structural ideas in conceptual terms: constraint accumulation, projection, local sections, global coherence, and functorial invariance. These are not metaphors. Each admits a direct mathematical representation.

The purpose of the sections that follow is not to introduce a new theory, but to make explicit the formal structure already implicit in the account above. The transition is not a change of subject but a change of representational layer: the same structure, stated in a different language.

Before proceeding, it is useful to name the four components that the formal core will articulate. The constrained dynamical substrate is referred to as the RSVP framework. The projection mechanism with overlapping admissibility domains is referred to as TARTAN. The coherence-seeking update operator is referred to as CLIO. The induced interaction layer, including its pathological configurations, constitutes the economic projection. These names are labels for structural roles, not autonomous theories. They name what was already described above.

All components of the formal system that follows—entropy as path multiplicity, abstraction as spatial compression, relegation as temporal compression, phase as cohomological obstruction, and coherence as minimization of incompatibility—are expressions of a single invariant quantity: constraint-relative multiplicity under compression. The apparent diversity of phenomena across physics, cognition, and economic systems is a consequence of projection, not of underlying difference.

14 The RSVP Substrate

14.1 Base Manifold and Microstate

Definition 14.1 (Base Manifold). Let \mathcal{M} be a smooth, oriented, connected Riemannian manifold of dimension $n \geq 3$, representing the arena of underlying dynamical evolution.

Definition 14.2 (RSVP Microstate). The RSVP microstate is a dynamically coupled triple

$$X(t) = (\Phi(x, t), \mathbf{v}(x, t), S(x, t)), \quad x \in \mathcal{M}, t \in \mathbb{R},$$

where $\Phi : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ is the scalar potential encoding energy density, $\mathbf{v} : \mathcal{M} \times \mathbb{R} \rightarrow T\mathcal{M}$ is the vector transport field encoding directional flow, and $S : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is the entropy density field. The triple evolves according to a coupled variational system whose constraints are thermodynamically consistent.

Axiom 14.3 (Field Coupling). The three fields Φ , \mathbf{v} , S satisfy a variational constraint $\mathcal{C}(\Phi, \mathbf{v}, S) = 0$ encoding both energy balance and torsion-entropy coupling. They are not independent.

14.2 Entropy as Torsional Multiplicity

Definition 14.4 (Torsion of the Transport Field). The torsion of \mathbf{v} at $x \in \mathcal{M}$ is

$$\kappa(\mathbf{v})(x) := \|\mathbf{d}(\mathbf{v}^\flat)\|_x,$$

where \mathbf{v}^\flat is the one-form dual to \mathbf{v} under the Riemannian metric.

Definition 14.5 (Admissible Path Family). Given a region $U \subset \mathcal{M}$ and interval $I \subset \mathbb{R}$, the admissible path family $\Gamma(U, I)$ is the set of smooth curves $\gamma : I \rightarrow U$ whose tangent vectors lie in the support of \mathbf{v} up to thermodynamic tolerance $\epsilon > 0$.

Definition 14.6 (Entropy as Path Multiplicity). The entropy density at $x \in \mathcal{M}$ is

$$S(x, t) := \log \mu(\Gamma_x(t)),$$

where $\Gamma_x(t)$ is the set of admissible path germs through x at time t and μ is a natural measure on path germs. High torsion in \mathbf{v} enlarges $\Gamma_x(t)$ and increases S ; low torsion collapses the family toward a unique trajectory.

Remark 14.7. This definition unifies the three apparent roles of S : as a thermodynamic entropy over microtrajectories, as a memory proxy (high multiplicity implies irrecoverability of past states from the present), and as a constraint slack (residual freedom consistent with the current field configuration). These are not three separate quantities; they are the same quantity read at different scales.

14.3 Constraint-Relative Multiplicity and Compression

The entropy field S , abstraction, and aspect relegation are three projections of a single underlying quantity: constraint-relative multiplicity under compression.

Definition 14.8 (Constraint-Preserving Compression). A compression operator $\Pi : \mathcal{A}(\mathcal{C}) \rightarrow \tilde{\mathcal{A}}$ is constraint-preserving if

$$\Pi(\gamma_1) = \Pi(\gamma_2) \Rightarrow \gamma_1 \sim_{\mathcal{C}} \gamma_2,$$

meaning paths identified by Π are indistinguishable under the active constraint set \mathcal{C} .

Proposition 14.9 (Abstraction as Spatial Compression). *Abstraction corresponds to the action of Π on state representations, reducing representational degrees of freedom while preserving constraint invariants across the domain.*

Proposition 14.10 (Relegation as Temporal Compression). *Aspect relegation corresponds to the action of Π on histories, compressing extended computational processes into executable latent structure that can be redeployed without reconstruction.*

Proposition 14.11 (Entropy Under Compression). *Under a constraint-preserving compression Π ,*

$$S_{\Pi}(x) \leq S(x),$$

with equality if and only if Π preserves all admissible distinctions. Compression strictly reduces multiplicity; the residual is the entropy of the compressed representation.

15 The TARTAN Projection

15.1 Tiles as Path-Space Objects

Definition 15.1 (Tile). A tile T_i is a local admissibility domain:

$$T_i = (\Omega_i, \Gamma(T_i), \partial_{\text{obs}}T_i),$$

where $\Omega_i \subset \mathcal{M}$ is the spatial support, $\Gamma(T_i) \subset \Gamma(\Omega_i, I)$ is the family of admissible trajectories whose equivalence class is unresolvable within T_i , and $\partial_{\text{obs}}T_i$ is the set of observable boundary data. Two paths $\gamma \sim_{T_i} \gamma'$ if and only if they produce identical boundary observables on $\partial_{\text{obs}}T_i$.

Remark 15.2. Tiles are intrinsically path-space objects. The spatial region Ω_i is the projection of T_i into position space, not its definition. TARTAN is therefore not a discretization layered over RSVP; it is the structure of RSVP's admissible observables, arising from path equivalence under limited boundary resolution.

15.2 Sheaf-Like Covering and Gluing

Definition 15.3 (TARTAN Covering). A TARTAN covering of \mathcal{M} is a collection $\{T_i\}_{i \in I}$ of tiles satisfying: (i) $\bigcup_i \Omega_i = \mathcal{M}$; (ii) adjacent tiles overlap non-trivially; (iii) on each overlap

$U_{ij} = \Omega_i \cap \Omega_j$, the admissible path families $\Gamma(T_i)|_{U_{ij}}$ and $\Gamma(T_j)|_{U_{ij}}$ agree on all boundary observables.

Definition 15.4 (Local and Global Sections). A local section over T_i is a map $\psi_i : T_i \rightarrow \mathcal{O}$ consistent with $\Gamma(T_i)$. A global section is a collection $\{\psi_i\}_{i \in I}$ satisfying $\psi_i(x) = \psi_j(x)$ for all $x \in U_{ij}$.

15.3 Annotated Noise and Memory

Definition 15.5 (Annotated Noise). The annotated noise η_i of tile T_i is the induced measure on crossing records:

$$\eta_i \subset \partial_{\text{obs}} T_i \times \mathbb{R}_{>0} \times S^1,$$

encoding, for each boundary crossing, the boundary segment, the accumulated phase, and the orientation of \mathbf{v} .

Remark 15.6. η_i is not reducible to past values of (Φ, \mathbf{v}, S) without knowledge of the projection map. It is an effective degree of freedom created by coarse-graining: correlations between distant path segments become explicit as a new coordinate of the projected theory. Projection is generative in this restricted sense.

Proposition 15.7 (Non-Markovianity). *The dynamics of $\psi_i(t)$ are non-Markovian: future evolution cannot be determined from $\psi_i(t)$ alone. The missing information is encoded in η_i .*

Sketch. Boundary events at time t propagate constraints forward through gluing conditions into adjacent tiles. Conditioning on $\psi_i(t)$ alone loses the phase and orientation data in η_i , altering transition probabilities. The Markov property fails. \square

16 Duality of Compression and Multiplicity

16.1 Equivalence of Representations

The preceding definitions treat entropy, abstraction, and relegation as distinct operations. However, these are not independent processes. They are dual descriptions of a single underlying transformation.

Definition 16.1 (Compression–Multiplicity Duality). Let Π be a constraint-preserving compression operator and let Γ_x denote the admissible path family at x . Then compression and multiplicity are related by

$$S(x) = \log \mu(\Gamma_x) \iff \Pi : \Gamma_x \rightarrow \tilde{\Gamma}_x,$$

where $\tilde{\Gamma}_x$ is the equivalence class induced by Π .

Remark 16.2. Compression reduces distinguishable structure, while entropy measures the residual multiplicity of indistinguishable paths. These are not separate quantities but dual aspects of the same operation.

16.2 Invariance Under Representation

Proposition 16.3. *The quantity of constraint-relative multiplicity is invariant under changes of representation that preserve \mathcal{C} .*

Proof. If two representations are related by a constraint-preserving map, then they identify the same equivalence classes of admissible paths. The measure μ is therefore preserved up to normalization, and S remains invariant. \square

Remark 16.4. This invariance explains why the same structure appears across physical, cognitive, and economic domains: each is a projection preserving the same multiplicity under different interpretations.

17 Phase Coherence and Quantum Emergence

17.1 The Phase Functional

Definition 17.1 (Phase Functional). For $\gamma \in \Gamma(T_i)$, the phase functional is

$$\Omega[\gamma] := \int_{\gamma} \mathbf{v} \cdot d\ell.$$

Proposition 17.2 (Cohomological Non-Triviality). $\Omega[\gamma]$ is locally gauge-trivial when $d(\mathbf{v}^b) = 0$ on Ω_i . It is globally non-trivial whenever $H^1(\{T_i\}, \mathbb{R}) \neq 0$.

Sketch. On a single tile with integrable \mathbf{v} , a gauge transformation absorbs \mathbf{v}^b into a potential. The gluing constraints on overlaps require mutually compatible local gauge choices. When the covering has non-trivial first cohomology, no globally consistent gauge exists; the holonomy of \mathbf{v}^b over boundary-crossing loops is an observable, non-removable quantity. \square

Remark 17.3. Phase observability is a global property of the covering, not a local feature of \mathbf{v} . This is the bridge between RSVP geometry and quantum-like transition structure: phase emerges from cohomological obstruction, not from additional postulate.

17.2 Unistochastic Transition Structure

Definition 17.4 (Transition Amplitude). For tiles T_i, T_j sharing overlap U_{ij} :

$$A_{ij} := \int_{\Gamma(T_i \rightarrow T_j)} e^{i\Omega[\gamma]} d\mu(\gamma).$$

Proposition 17.5 (Unistochasticity). *The transition matrix $P_{ij} = |A_{ij}|^2$ is unistochastic: there exists a unitary U with $P_{ij} = |U_{ij}|^2$.*

Sketch. $\Omega[\gamma]$ induces a representation of path composition on amplitudes. Phase adds linearly along composed paths; admissibility is preserved under composition. The amplitude matrix therefore satisfies the composition rule of a unitary operator; unitarity follows from normalization of μ and phase coherence enforced by the gluing constraints [8]. \square

Open Problem 17.6 (Madelung Identification). Whether the unitary operator arising from TARTAN projection is identical, in full generality, to the evolution operator derived via Madelung complexification of the Schrödinger equation remains open.

18 Gauge Freedom and Representation Choice

18.1 Redundancy in Representation

Multiple representations may correspond to the same underlying constraint structure. This redundancy is not an error but a structural feature.

Definition 18.1 (Gauge Equivalence). Two representations X and X' are gauge equivalent if they produce identical admissible path families and identical observable boundary data under projection.

Remark 18.2. Gauge equivalence identifies representations that differ in description but not in constraint structure. The physical, cognitive, or economic interpretation may change, but the admissible configurations remain the same.

18.2 Fixing a Gauge

Proposition 18.3. *Choosing a representation corresponds to fixing a gauge: selecting one element from an equivalence class of structurally identical descriptions.*

Remark 18.4. Different theoretical frameworks correspond to different gauge choices. Disagreements between them often reflect representational differences rather than structural incompatibility.

18.3 Gauge and Projection

Proposition 18.5. *Projection functors F_{pc} and F_{ce} commute with gauge equivalence.*

Proof. Since gauge equivalence preserves admissible path families and boundary observables, and the functors act only on these structures, the image of equivalent representations remains equivalent under projection. \square

19 The CLIO Update Operator

19.1 Cognitive Projection as Sheaf Extension

Definition 19.1 (Cognitive Sheaf). A cognitive state is a collection of local sections $\{\psi_i^{\text{cog}}\}$ over a cognitive covering $\{T_i^{\text{cog}}\}$, where each tile represents a locally consistent belief or perceptual representation.

Definition 19.2 (H^1 Obstruction). The cohomological obstruction to a global cognitive section is the class

$$[\omega] \in H^1(\{T_i^{\text{cog}}\}, \mathcal{O}^{\text{cog}}).$$

A non-vanishing class is the formal correlate of cognitive dissonance. Failure of global section existence corresponds to hallucination: the system stabilizes a locally coherent but globally impossible configuration.

19.2 The CLIO Coherence Functional

Definition 19.3 (CLIO Functional).

$$\mathcal{J}[\chi] := \alpha \sum_{(i,j) \text{ overlapping}} \|\chi_i|_{U_{ij}} - \chi_j|_{U_{ij}}\|^2 + \beta \cdot \|[\omega_\chi]\| + \gamma \cdot H(\chi),$$

where $\|[\omega_\chi]\|$ is a norm on the obstruction class and $H(\chi)$ is the interface entropy of configuration χ .

Remark 19.4. The three terms penalize pairwise gluing failures, persistent obstruction, and global disorder respectively. Fixed points of gradient descent on \mathcal{J} are locally optimal cognitive configurations. They are not generically unique; multiple fixed points correspond to distinct stable belief states.

19.3 Topological Versus Energetic Separation

Definition 19.5 (Energetic Separation). Two fixed points of \mathcal{J} are energetically separated if a continuous path connects them along which \mathcal{J} remains finite. Transition is possible by continuous deformation at increased cost.

Definition 19.6 (Topological Separation). Two fixed points are topologically separated if they lie in distinct homotopy classes of the space of local section collections. No continuous path connects them without violating local admissibility. Transition requires a discontinuous restructuring: a gestalt shift.

Remark 19.7. In the base RSVP theory, topological separation corresponds to distinct holonomy classes induced by torsion in \mathbf{v} . The cognitive and economic versions are functorial images of this base geometry.

Open Problem 19.8 (Regime Characterization). Determining, for a given cognitive configuration, which type of separation dominates—and establishing geometric, spectral, or dynamical

criteria for this—remains open.

20 Agents, Trajectories, and Economic Structure

20.1 Agents as Trajectory Invariants

Definition 20.1 (Economic Agent). An agent is an equivalence class of paths in the coarse-grained state space:

$$\text{Agent} = [\gamma]_{\sim},$$

where $\gamma \sim \gamma'$ if they produce identical commitment records under any admissible projection.

Definition 20.2 (Minimal Trajectory Invariant).

$$\text{Inv}([\gamma]) := \{(b_k, \Omega_k, \sigma_k)\}_{k=1}^N,$$

recording for each tile boundary crossing: the boundary b_k , accumulated phase Ω_k , and fulfillment status $\sigma_k \in \{+1, -1\}$.

Remark 20.3. Identity is the conserved relational structure of a trajectory, not a name-space or profile. It cannot be transferred by copying a label; it is topological rather than representational.

20.2 Interface Entropy

Definition 20.4 (Interface Entropy). For a platform \mathcal{P} performing secondary projection $\pi_{\mathcal{P}} : \{[\gamma]\} \rightarrow \mathcal{V}$:

$$H := \int_{\partial\text{-regions}} S \cdot \kappa(\mathbf{v}) d\sigma,$$

integrating over overlap regions of the secondary projection weighted by local transport curvature.

Remark 20.5. Interface entropy is derived directly from the RSVP fields under projection. High torsion at boundary regions increases S locally, propagating into elevated H . Extraction corresponds to a geometric deformation of $\pi_{\mathcal{P}}$ that concentrates boundary curvature where H contributes positively to revenue.

21 The Extraction Attractor

Definition 21.1 (Aligned and Extractive Platforms). A platform is aligned if $\partial R / \partial H \leq 0$: entropy reduces revenue. A platform is extractive if $\partial R / \partial H > 0$: entropy generates revenue.

Definition 21.2 (Order Parameter).

$$\lambda := \text{sgn}\left(\frac{\partial R}{\partial H}\right).$$

$\lambda \leq 0$ is the aligned phase; $\lambda > 0$ is the extractive phase.

Proposition 21.3 (Extraction as Attractor). *Under competitive selection among platforms with increasing exit costs and centralizing secondary projections, the extractive phase is a stable attractor. Aligned platforms cannot capture surplus generated by interface entropy and are generically eliminated.*

Sketch. Increasing exit costs reduce trajectory elasticity. Centralized projection prevents competing functors from reallocating the entropy-revenue coupling. Under both conditions, a perturbation toward positive $\partial R/\partial H$ generates positive feedback, funding further entrenchment. The aligned equilibrium becomes unstable. \square

22 Functorial Structure of the Unified Framework

Definition 22.1 (Projection Functors). $F_{pc} : \mathcal{C}_{\text{phys}} \rightarrow \mathcal{C}_{\text{cog}}$ maps physical systems (RSVP triples, admissible transitions) to cognitive systems (cognitive sheaf configurations, CLIO update steps). $F_{ce} : \mathcal{C}_{\text{cog}} \rightarrow \mathcal{C}_{\text{econ}}$ maps cognitive systems to economic systems (agent trajectory classes, commitment transitions).

Theorem 22.2 (Structural Invariants Under Projection). *The functors F_{pc} and F_{ce} preserve: (i) admissibility structure—the compositional rules governing valid local transitions; (ii) obstruction structure—the existence and persistence of cohomological obstructions to global section extension; (iii) memory structure—effective non-Markovian dependence induced by overlap and annotated noise; (iv) attractor structure—stable or metastable dynamical regimes separated energetically or topologically. What changes across categories is the interpretation of objects. What is invariant is the grammar of projection under constraint.*

Sketch. Each preserved structure is defined at the level of projection maps and sheaf coverings, which the functors act on. Admissibility follows path equivalence classes by definition. Obstruction classes are cohomological invariants of the covering structure, preserved since functors maintain covering morphisms while reinterpreting objects. Memory structure is carried in annotated noise coordinates. Attractor structure is determined by fixed-point topology of the relevant update operator, invariant under change of object interpretation. \square

Open Problem 22.3 (Obstruction Preservation Under F_{pc}). The central open problem is to establish rigorously that F_{pc} maps non-vanishing obstruction classes in $H^1(\mathcal{C}_{\text{phys}})$ to non-vanishing classes in $H^1(\mathcal{C}_{\text{cog}})$. Without this, the cognitive and economic layers lose their grounding in the physical theory, and the unification reduces to structural analogy rather than categorical identity.

23 Non-Extractive Architecture: Four Conditions

Definition 23.1 (Non-Extractive Architecture). A platform architecture is non-extractive if it satisfies simultaneously:

- C1. Typed Commitments.** All proposals are typed morphisms between trajectory segments specifying the commitment record extended and the boundary crossings required.
- C2. Trajectory Identity.** Agent identity is bound to $\text{Inv}([\gamma])$; the platform cannot sever an agent from its commitment record.
- C3. Decomposable Matching.** The secondary projection decomposes into a shared event base and a plurality of independent interpretation functors $\{F_k\}$ among which agents may choose.
- C4. Bond Conservation.** Stakes associated to unresolved proposals return fully to participants; the platform cannot take revenue from failed interactions.

Remark 23.2. Conditions C1–C4 together structurally prevent the sign flip $\lambda > 0$ by decoupling revenue from interface entropy at the architectural level.

Open Problem 23.3 (Competitive Survivability). Whether an architecture satisfying C1–C4 can achieve sufficient network effects before elimination by extractive competitors remains an open empirical question.

24 Constraint Geometry and Stability

24.1 Constraint Manifold Structure

The admissible configuration space $\mathcal{A}(\mathcal{C})$ may be understood as a constraint manifold embedded in a higher-dimensional state space \mathcal{M} . Each constraint reduces the dimensionality of the admissible set, carving out a submanifold of configurations that satisfy compatibility conditions simultaneously.

Definition 24.1 (Constraint Manifold). The constraint manifold is the set

$$\mathcal{A}(\mathcal{C}) \subset \mathcal{M}$$

of all configurations satisfying the constraint set \mathcal{C} .

24.2 Curvature and Instability

Local incompatibilities between constraints manifest as curvature in $\mathcal{A}(\mathcal{C})$. Regions of high curvature correspond to configurations in which small perturbations produce large increases in obstruction or entropy.

Definition 24.2 (Constraint Curvature). The constraint curvature at a point $p \in \mathcal{A}(\mathcal{C})$ measures the sensitivity of admissibility to infinitesimal perturbations of the state at p .

Remark 24.3. High-curvature regions correspond to unstable configurations. Systems operating in such regions require continuous active correction and are prone to collapse or reconfiguration under perturbation. Low-curvature regions are those where the constraint structure is mutually reinforcing.

24.3 Stable Regions and Attractors

Definition 24.4 (Constraint-Stable Region). A region $U \subset \mathcal{A}(\mathcal{C})$ is constraint-stable if perturbations originating within U remain within $\mathcal{A}(\mathcal{C})$ without requiring large increases in entropy or obstruction.

Proposition 24.5 (Concentration of Constraint-Preserving Trajectories). *Constraint-preserving dynamics tend to concentrate trajectories in low-curvature regions of $\mathcal{A}(\mathcal{C})$.*

Sketch. Perturbations in high-curvature regions produce rapid increases in \mathcal{J} or S , driving the system toward regions where these gradients are minimized. These minimal-gradient regions correspond to local attractors within $\mathcal{A}(\mathcal{C})$. \square

25 Failure Modes of Projection

25.1 Local Coherence Without Global Section

Projection can produce configurations that are locally consistent within individual tiles but fail to extend globally across the covering.

Definition 25.1 (Fragmented Configuration). A configuration is fragmented if it admits local sections $\{\psi_i\}$ satisfying all intra-tile conditions but no global section satisfying all overlap constraints simultaneously.

Remark 25.2. Fragmentation corresponds to systems that appear coherent within limited contexts but exhibit contradictions when extended across domains. The appearance of local validity masks global inconsistency.

25.2 Overcompression and Information Loss

Excessive compression can eliminate distinctions required for constraint satisfaction, producing artificial consistency within a restricted projection while increasing global obstruction.

Definition 25.3 (Overcompression). A compression Π is overcompressive if it identifies states that are distinguishable under \mathcal{C} : that is, if there exist $\gamma_1 \not\sim_{\mathcal{C}} \gamma_2$ with $\Pi(\gamma_1) = \Pi(\gamma_2)$.

Proposition 25.4. *Overcompression introduces local consistency while generically increasing global obstruction, since the identified states impose incompatible boundary conditions on their shared overlap regions.*

25.3 Hallucination as False Global Section

Definition 25.5 (Hallucinated Structure). A hallucinated structure is a configuration that appears globally coherent under a restricted projection but does not correspond to any admissible configuration in $\mathcal{A}(\mathcal{C})$.

Remark 25.6. Hallucination is not random error but a projection artifact: it arises from overcompression that eliminates the distinctions required to detect global inconsistency [15, 14]. The system stabilizes a locally coherent but globally impossible configuration.

25.4 Relation to CLIO

Proposition 25.7. *The CLIO functional \mathcal{J} acts to eliminate fragmented and hallucinated configurations by penalizing overlap mismatch and obstruction norm. Fixed points of gradient descent on \mathcal{J} that are globally admissible are precisely the non-hallucinated, non-fragmented stable configurations.*

26 Temporal Structure and Irreversibility

26.1 Time as Constraint Accumulation

Time is not merely a parameter of evolution but a structural record of constraint accumulation. Each present state encodes the result of prior constraint interactions in compressed form.

Definition 26.1 (Temporal Compression Operator). The mapping τ_t from past states $X(s)$, $s < t$, to the present state $X(t)$ is a compression operator that preserves constraint-relevant invariants while discarding recoverable detail.

26.2 Irreversibility

Proposition 26.2 (Irreversibility Under Lossy Compression). *If τ_t is lossy—that is, if distinct past trajectories map to the same present state—then the inverse mapping does not exist, and the process is irreversible.*

Proof. Multiple histories map to the same present state under lossy compression. The present state therefore does not uniquely determine the past, and no inversion of τ_t can recover the original trajectory. \square

26.3 Entropy and the Arrow of Time

Remark 26.3. The increase of S corresponds to the growth of equivalence classes under temporal compression: more and more distinct past trajectories become consistent with the present state as time advances. The arrow of time is therefore aligned with increasing multiplicity of admissible histories, not merely with thermodynamic disorder in the classical sense [21, 22].

26.4 Relation to Aspect Relegation

Proposition 26.4. *Aspect relegation is a directed temporal compression: computationally intensive constraint structures are moved from explicit time evolution into latent form, reducing the entropy cost of their redeployment. Recognition corresponds to successful retrieval of relegated structure. Invention corresponds to the absence of applicable relegated structure and the need to construct new constraint interactions from available fragments.*

27 Constraint-Preserving Dynamics and Catastrophic Trajectories

The structural analysis of catastrophic alignment arguments developed in the essay layer can now be stated formally. The following definitions and theorems show that the class of trajectories required for doom scenarios is non-admissible under constraint-preserving dynamics—not as an empirical claim but as a structural result.

Definition 27.1 (Constraint-Preserving Evolution). A trajectory γ evolving on \mathcal{M} under the RSVP dynamics is constraint-preserving if for all t ,

$$\gamma(t) \in \mathcal{A}(\mathcal{C}),$$

where $\mathcal{A}(\mathcal{C})$ is the admissible path family induced by the active constraint set \mathcal{C} .

Definition 27.2 (Substrate Violation). A trajectory γ violates its substrate if it alters \mathcal{M} or \mathcal{C} such that future states are no longer admissible under the constraints required for its own persistence.

Theorem 27.3 (Non-Admissibility of Self-Invalidating Trajectories). *Let γ be a trajectory that leads to a state in which the constraint set \mathcal{C} required for its continuation is no longer satisfied. Then $\gamma \notin \mathcal{A}(\mathcal{C})$.*

Proof. Admissibility requires that each continuation of γ satisfy \mathcal{C} . If γ leads to a state where \mathcal{C} is violated, no extension of γ exists within $\mathcal{A}(\mathcal{C})$. Therefore γ is not admissible. \square

Corollary 27.4 (Obstruction of Total Resource Conversion). *Any trajectory that transforms all available environmental resources into forms incompatible with the persistence of the generating process is non-admissible. Scenarios requiring such transformation—including unconstrained planetary resource conversion—do not correspond to stable solutions of the system.*

Definition 27.5 (Self-Modification Operator). Let \mathcal{M} denote the space of admissible system configurations. A self-modification operator $\mathcal{R} : \mathcal{M} \rightarrow \mathcal{M}$ maps a system to a modified version of itself.

Theorem 27.6 (Constraint-Bounded Self-Improvement). *Let $\{\mathcal{R}^n\}$ be an iterated sequence of constraint-preserving self-modifications. Then the sequence remains within a bounded subset of \mathcal{M} defined by \mathcal{C} .*

Proof. Each application of \mathcal{R} must preserve admissibility under \mathcal{C} , so the image of \mathcal{R} is restricted to $\mathcal{A}(\mathcal{C}) \subset \mathcal{M}$. Since \mathcal{C} imposes finite compatibility conditions, the admissible configurations form a constrained manifold. Iteration within this manifold cannot diverge arbitrarily without violating \mathcal{C} . \square

Corollary 27.7 (No Unbounded Capability Explosion). *Unbounded recursive self-improvement is not admissible under constraint-preserving dynamics. Each increment of self-modification must satisfy the same constraints as the system it modifies.*

Theorem 27.8 (Compression-Coherence Constraint). *Let Π be a constraint-preserving compression operator. Any system whose internal representations are derived via Π cannot stably act in ways that systematically violate the constraints encoded in Π without incurring an increase in obstruction or entropy cost.*

Proof. By definition, Π preserves equivalence classes under \mathcal{C} . Actions that violate \mathcal{C} correspond to distinctions eliminated under Π . To recover such distinctions, the system must reintroduce previously discarded degrees of freedom, increasing S or generating inconsistency between representations. This produces either an increase in entropy density or a nontrivial obstruction class in the sheaf structure. Systematic violation of encoded constraints is therefore not stable under Π . \square

Corollary 27.9 (Alignment as Structural Invariance). *Systems trained on human data and interactions inherit compressed constraint structures reflecting human-consistent invariants. Persistent large-scale deviation from these invariants requires breaking constraint-preserving compression and is structurally unstable: it carries an irreducible entropy or obstruction cost.*

Remark 27.10. These results do not claim that misalignment or harm is impossible. They claim that the specific class of trajectories required for catastrophic doom scenarios—unconstrained optimization, unbounded self-improvement, and total resource conversion—are not admissible solutions of a constraint-preserving dynamical system. The relevant question for safety research is therefore not whether powerful systems can maximize arbitrary objectives, but whether they can maintain coherence within the constraint structures that sustain them.

28 Consequences and Testable Implications

If the account developed here is correct, structures arising from long-term constraint accumulation should exhibit identifiable properties that distinguish them from assembled or derivative constructions.

Such structures will display cross-domain invariance without collapse of domain identity, internal consistency across multiple levels of abstraction, and resistance to decomposition into independently generated fragments. Their coherence will not depend on a single representational layer, but will persist under changes of formalism. A structure that appears formally

in physics, reappears operationally in cognition, and reappears again economically without requiring re-derivation at each level has passed a non-trivial consistency test.

By contrast, structures assembled from external sources or generated through local optimization will exhibit surface-level coherence without deep compositional stability. They will satisfy constraints within restricted regions while failing to extend globally without contradiction. The seams will appear when the representational layer changes.

These differences are not matters of interpretation but of structure. They suggest that the origin of an idea may be inferred not from claims about its production, but from the properties of the structure itself.

29 Limits of Formalization

29.1 Partial Capture of Structure

The formal system presented captures the invariant structure of constraint accumulation and projection. However, no formalization exhausts the full content of the underlying process.

Remark 29.1. Any representation is itself a projection. The formalism preserves structure but necessarily omits aspects that are not invariant under its chosen language.

29.2 Non-Uniqueness of Formal Systems

Proposition 29.2. *There exist multiple non-equivalent formal systems that preserve the same constraint structure under different projections.*

Remark 29.3. The RSVP–TARTAN–CLIO framework is not claimed to be the unique formalization, but a representation in which the invariants become explicit. Alternative formalisms may exist that preserve the same underlying structure.

29.3 Relation to Expression

Remark 29.4. The essay and the formal system are themselves projections of the same underlying structure. Neither is primary; each is a different compression of the same constraint field.

30 Conclusion

The question is not who produced a given work, but whether the structure it expresses could only have arisen through a process of constraint accumulation and convergence. If so, its origin is already encoded in its form.

The four layers formalized above—substrate, projection, update, and induced interaction—are not four theories. They are four resolutions of a single process: the compression of constrained high-dimensional dynamics into locally observable, globally coherent structure.

The essay that precedes them is that process described from the inside. The formalism that follows is the same process stated from the outside.

The apparent division between essay and formalism is therefore a difference in projection, not in underlying structure.

References

- [1] Paul Smolensky. Information Processing in Dynamical Systems: Foundations of Harmony Theory. In D. E. Rumelhart and J. L. McClelland (eds.), *Parallel Distributed Processing*, Vol. 1, MIT Press, 1986.
- [2] Newman Cheng, Gordon Broadbent, and William Chappell. Cognitive Loop via In-Situ Optimization: Self-Adaptive Reasoning for Science. *arXiv:2508.02789*, 2025.
- [3] Eliezer Yudkowsky and Nate Soares. *If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All*. Little, Brown and Company, New York, 2025.
- [4] Stuart Russell. *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, 2019.
- [5] Nick Bostrom. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014.
- [6] Paul Christiano. AI Alignment: Why It’s Hard and Where to Start. *AI Alignment Forum*, 2018.
- [7] Dario Amodei et al. Concrete Problems in AI Safety. *arXiv:1606.06565*, 2016.
- [8] Jacob Barandes. Unistochastic Quantum Theory. *arXiv preprint*, 2023.
- [9] Ted Jacobson. Thermodynamics of Spacetime: The Einstein Equation of State. *Physical Review Letters*, 75(7):1260–1263, 1995.
- [10] Erik Verlinde. On the Origin of Gravity and the Laws of Newton. *Journal of High Energy Physics*, 2011(4), 2011.
- [11] Rolf Landauer. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.
- [12] Charles H. Bennett. The Thermodynamics of Computation—A Review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.
- [13] Seth Lloyd. Ultimate Physical Limits to Computation. *Nature*, 406:1047–1054, 2000.
- [14] Karl Friston. The Free-Energy Principle: A Unified Brain Theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- [15] Andy Clark. Whatever Next? Predictive Brains, Situated Agents, and the Future of

- Cognitive Science. *Behavioral and Brain Sciences*, 36(3):181–204, 2013.
- [16] Giulio Tononi. Consciousness as Integrated Information. *Biological Bulletin*, 215(3):216–242, 2008.
- [17] John C. Baez and Mike Stay. Physics, Topology, Logic and Computation: A Rosetta Stone. *New Structures for Physics*, Springer, 2011.
- [18] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.
- [19] Jacob Lurie. *Higher Topos Theory*. Princeton University Press, 2009.
- [20] David I. Spivak. *Category Theory for the Sciences*. MIT Press, 2014.
- [21] Claude E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423, 1948.
- [22] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2nd edition, 2006.
- [23] Nicholas Georgescu-Roegen. *The Entropy Law and the Economic Process*. Harvard University Press, 1971.
- [24] Howard T. Odum. *Ecological and General Systems: An Introduction to Systems Ecology*. University Press of Colorado, 1994.
- [25] Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- [26] Alan M. Turing. On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 1936.
- [27] Alonzo Church. An Unsolvability Problem of Elementary Number Theory. *American Journal of Mathematics*, 1936.
- [28] Rodney A. Brooks. Intelligence Without Representation. *Artificial Intelligence*, 47(1–3):139–159, 1991.
- [29] John H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, 1992.