The Refractive Self

Constraint, Projection, and Reconstruction
Across Physical, Cognitive, and Cultural Systems

Flyxion
Independent Researcher

April 28, 2026

Contents

List of Theorems

Part I

Foundations: Projection, Information, and the Blind Spot

Chapter 1

The Structure of Scientific Projection

From Correctness to Selection

Imagine entering a destination into a navigation system and being presented not with a single optimal route, but with millions of distinct paths, each identical in travel time, fuel consumption, and constraint satisfaction. Every option is equally correct. In such a regime, correctness ceases to function as a principle of selection. The system must appeal to a secondary structure—something that distinguishes among equally admissible possibilities.

This situation is not hypothetical. It is the generic condition of high-dimensional systems. When the space of admissible solutions becomes large, correctness becomes degenerate. The problem shifts from finding a valid configuration to selecting one from an overabundance of valid configurations. Reality, in this view, is not constructed by accumulation but selected through elimination: it is the residue that remains when all configurations that cannot coherently coexist are removed.

This shift reveals a deeper structural issue in the foundations of scientific reasoning. The classical paradigm assumes that identifying a correct representation of the world is sufficient for understanding. But this assumption presupposes that the mapping from reality to representation is invertible—that nothing essential is lost in the act of modeling. This presupposition is false.

Consider a simple projection. A three-dimensional object casts a two-dimensional shadow. The shadow preserves certain invariants—outline, relative proportions—but discards depth. Multiple distinct objects may produce the same shadow. The projection is therefore non-invertible: from the shadow alone, the original object cannot be uniquely recovered. The lost information is not noise; it is structurally excluded by the projection itself.

Scientific models operate through precisely such projections. They map a high-dimensional experiential domain onto a reduced space of measurable invariants. This reduction is necessary: without it, communication, prediction, and engineering would be impossible. But it comes at a cost. The degrees of freedom discarded by the projection—its *fibers*—are no longer accessible within the model.

The central error arises when this reduction is forgotten. When the model space is treated as ontologically complete, the omitted structure is reinterpreted as nonexistent. This is the fallacy of misplaced concreteness: the confusion of a projection with the domain from which it was derived. Phenomena excluded by construction are then declared illusory or in need of explanation by the very model that removed them.

1.1 Introduction

Scientific inquiry proceeds, at its deepest level, through a systematic act of reduction. The world as encountered in lived experience presents itself as irreducibly rich: temporally unfolding, context-sensitive, perspectival, and saturated with variability. No two moments are identical; no observation is free from the conditions under which it is made. Yet science, as a disciplined practice, does not attempt to preserve this full richness. Instead, it seeks what remains invariant across it.

This operation—extracting stability from variability—is not merely a methodological convenience. It is the defining structural move that makes science possible. To construct a law, a theory, or even a measurement, one must identify what does not change when conditions vary. Temperature must be defined independently of the particular thermometer; velocity must be invariant under coordinate transformation; biological regularities must persist across organisms and environments. The entire edifice of scientific knowledge rests on this capacity to isolate invariants from a background of variation.

We formalise this operation as a projection. A projection maps a high-dimensional, richly structured domain of phenomena onto a reduced space in which only selected features are retained. What is preserved are invariants; what is discarded are degrees of freedom deemed irrelevant to the model's purpose. This reduction is not neutral. It is an active restructuring of the domain, one that introduces both clarity and loss simultaneously.

The central claim of this chapter is that this loss is not incidental but necessary. Projection is generically non-invertible: once variability has been compressed into invariance, the original structure cannot be recovered. The model is not a reversible encoding of reality but a many-to-one mapping that collapses distinctions. From this structural fact follows a philosophical consequence: the blind spot. When the projection is mistaken for the full domain—when the model is treated as reality itself—the discarded structure disappears from consideration, even though it remains essential to the coherence of the phenomena being studied.

This diagnosis aligns closely with the critique developed in , where lived experience is identified as both the source and the condition of possibility of scientific knowledge. The blind spot is not a failure of science as a method; it is a failure to recognise the limits imposed by the very operation that gives science its power.

The remainder of this chapter develops the formal structure of projection, establishes its non-invertibility, and derives the blind spot as a necessary consequence.

1.2 Projection as a Structural Operation

We begin by fixing the domains involved.

Definition 1.1 [Experiential Domain]**.** Let X be a separable metric space whose points represent full phenomenal states. These states include not only measurable quantities but contextual embedding, temporal unfolding, perspectival structure, and variability. We call X the *experiential domain.*

The choice of X is deliberately expansive. It is not restricted to "subjective experience" in a narrow sense, but includes the entire structured field within which phenomena are encountered, measured, and interpreted. This includes what is later formalised as physical, biological, or informational.

Definition 1.2 [Model Space]**.** Let M be a finite-dimensional smooth manifold whose points represent equivalence classes of states in X under a specified invariance relation. We call M the *model space*.

The transition from X to M encodes the act of abstraction. The manifold structure of M reflects the fact that scientific models are organised around continuous parameters and differentiable relationships, even when the underlying phenomena may not be.

Definition 1.3 [Projection]**.** A *scientific projection* is a continuous surjection

$$\pi : X \to M,$$

such that for each $m \in M$, the preimage

$$\pi^{-1}(m)$$

consists of all states in X identified as equivalent under the invariance criteria defining M.

This definition makes explicit what is often implicit: modelling is not merely describing but identifying. Distinct states in X are treated as the same in M if they agree on the invariants of interest. The projection π is therefore an equivalence relation in disguise, collapsing a complex domain into a simpler one.

Remark 1.4. The projection π is not unique. Different theories correspond to different projections, each preserving different invariants. Classical mechanics, thermodynamics, and quantum theory are not competing descriptions of the same projection, but distinct projections of the same underlying domain.

1.3 Non-Invertibility and Fiber Structure

The defining structural feature of projection is that it is generically non-injective.

Proposition 1.5 [Generic Non-Injectivity]**.** *If* $\dim X > \dim M$*, then* π *is not injective. For almost every* $m \in M$*, the fiber* $\pi^{-1}(m)$ *contains multiple elements.*

The geometric content of this statement is straightforward: a map from a higher-dimensional space to a lower-dimensional one must collapse dimensions. Information is compressed, and distinct inputs become indistinguishable.

Definition 1.6 [Fiber]**.** For $m \in M$, the set $\pi^{-1}(m)$ is called the *fiber* over m. It represents all states in X that are indistinguishable within the model.

The fiber is not noise in a trivial sense. It is structured variability—differences that matter in the full domain but are ignored by the model. For example, averaging across repeated measurements collapses temporal variability into a single value; coarse-graining in physics collapses microstates into macrostates; linguistic abstraction collapses distinct utterances into the same semantic category.

Proposition 1.7 [Non-Invertibility]**.** *There exists no global left-inverse* $\pi^{-1} : M \to X$ *recovering the original state from its projection.*

This is the formal statement that models are irreversible. Once distinctions are collapsed, they cannot be recovered without external information.

1.4 Information Loss and Irreversibility

Projection induces a fundamental asymmetry between construction and reconstruction.

Definition 1.8 [Information Loss]**.** The information loss at $m \in M$ is measured by the size or structure of the fiber $\pi^{-1}(m)$.

The larger the fiber, the greater the loss. Importantly, this loss is structured: it reflects the dimensions along which variability has been discarded.

Proposition 1.9 [Irreversibility]**.** *No procedure based solely on M can reconstruct $x \in X$ from $\pi(x)$.*

This result has a profound implication. Scientific models are necessarily incomplete descriptions of reality. They capture what is invariant but cannot recover what has been discarded.

This aligns directly with the observation in that averaging procedures in science eliminate variability—the very variability that constitutes lived experience. The act of producing an "objective" result involves moving away from the specific structure one might wish to understand.

1.5 Invariance and the Structure of Objectivity

The purpose of projection is to identify invariants.

Definition 1.10 [Invariant]**.** Given a transformation group $\mathcal{G}$ acting on X, a function I is invariant if $I(g \cdot x) = I(x)$ for all $g \in \mathcal{G}$.

The model space can be understood as a quotient:

$$M \cong X/\mathcal{G}.$$

Objectivity is therefore not the elimination of subjectivity, but invariance under transformation.

Proposition 1.11 [Objectivity as Invariance]**.** *Scientific objectivity consists in identifying quantities invariant under transformations of perspective, measurement, or representation.*

This reframes a common misunderstanding. Science does not remove the observer; it factors out the variability introduced by observation. As emphasized in , objectivity depends on subjectivity—it is defined relative to transformations of perspective.

1.6 The Blind Spot

We now arrive at the central concept.

Definition 1.12 [Blind Spot]**.** The blind spot of a projection π is the systematic neglect of the structure contained in the fibers $\pi^{-1}(m)$ when M is treated as complete.

The blind spot arises when the projection is reified. The model is taken to be reality, and the discarded structure is forgotten.

Proposition 1.13 [Structural Origin]**.** *The blind spot is an unavoidable consequence of non-invertibility.*

This is not a contingent error but a structural one. Every projection creates a domain of unrepresented structure. The blind spot consists in failing to recognise this.

This corresponds precisely to the critique articulated in : science, by bracketing experience to achieve invariance, risks forgetting that experience is the ground from which its models arise. When the map is mistaken for the territory, the conditions of its construction disappear from view.

1.7 Consequences and Forward Structure

The formal analysis of projection yields three guiding consequences.

First, models are necessarily partial. They preserve invariants but discard structure. Completeness is structurally impossible.

Second, continuity cannot be defined as preservation of state within M. Because π is non-invertible, continuity must be defined in terms of structure preserved across transformation. This motivates the development of constraint theory in Part II.

Third, the blind spot identifies the central task: to account for what is lost in projection without abandoning the power of abstraction. The remainder of the monograph develops the machinery required to do this.

Projection is therefore not merely a tool. It is the foundational operation that defines both the strength and the limitation of scientific knowledge. Understanding its structure is the first step toward extending science beyond its blind spot.

Chapter 2

Information Without Meaning

When Fidelity Becomes Loss

Imagine recording a piece of music with perfect fidelity. Every note is captured exactly as it was played. Now imagine that same piece being performed again—slightly faster, with different emphasis, a shift in timing, a change in instrumentation. To a listener, it is recognisably the same piece. To a system that demands exact reproduction, it is entirely different.

This reveals a tension that is easy to overlook. There is a difference between preserving a sequence and preserving what the sequence *does*. A system can be perfectly faithful to a string of symbols and yet completely blind to the structure those symbols express.

We are accustomed to thinking of information as something that is transmitted, stored, and reproduced. But this assumes that what matters is the exact arrangement of symbols. In many cases, this is too strong a requirement. Small changes that alter the surface form can leave the underlying structure intact. At the same time, enforcing exact preservation can eliminate all acceptable variations, reducing a flexible structure to a brittle artifact.

This becomes clearer when we consider how meaning persists in practice. A story can be retold in different words, a concept can be explained in different ways, a pattern can appear in different contexts. The surface changes, but something remains. That persistence does not depend on exact reproduction. It depends on the preservation of relationships.

The difficulty is that not all systems are capable of recognizing this kind of persistence. A system that measures only symbol-by-symbol agreement treats every deviation as error. It has no way to distinguish between a transformation that destroys structure and one that preserves it. As a result, it may reject valid instances of the same content while accepting exact repetitions that have lost their relevance.

This creates a subtle inversion. Increasing fidelity does not always increase coherence. In some cases, it does the opposite. By insisting on exact preservation, the system excludes the very transformations that allow structure to remain visible under changing conditions.

The problem is not that such systems are incorrect. Within their domain, they are exact. The problem is that their notion of invariance is too narrow. They preserve symbols, but not the structures those symbols can carry.

2.1 Introduction

The formalisation of projection in Chapter 1 established a structural asymmetry: information discarded in the passage from the experiential domain X to the model space M cannot, in general, be recovered. The theory of information provides a precise account of transmission and reduction. This chapter shows that while this framework is mathematically complete within its domain, it is structurally incapable of representing semantic invariance: the persistence of structure under transformation of surface form.

2.2 Shannon Information Theory

Let Σ be a finite alphabet and S a random variable over Σ with distribution $p(s)$. The *entropy* is

$$H(S) = -\sum_{s\in\Sigma} p(s)\log p(s).$$

The *mutual information* between X and Y is $I(X;Y) = H(X) - H(X \mid Y)$.

Definition 2.1 [Channel Capacity]**.** The *capacity* of a channel is $C = \sup_{p(x)} I(X;Y)$.

Definition 2.2 [Symbolic Invariance]**.** A channel preserves *symbolic invariance* if the received string $\hat{s}$ satisfies $\hat{s} = s$ with high probability. The invariant object is the symbol sequence itself.

2.3 The Absence of Semantic Structure

Definition 2.3 [Semantic Transformation]**.** Let $\mathcal{T}$ be a set of transformations on Σ^* such that $T(s)$ preserves the relational or structural content of s while altering its surface form.

Proposition 2.4 [Non-Representation of Semantic Invariance]**.** *Shannon information theory cannot represent invariance under non-trivial semantic transformations.*

Proof. In Shannon theory, equality is defined by exact symbol equality. Any T with $T(s) \neq s$ is treated as deviation. The theory has no mechanism for identifying s and $T(s)$ as equivalent under a structural relation. □

2.4 The Fidelity–Structure Tradeoff

Proposition 2.5 [Fidelity–Structure Tradeoff]**.** *Maximising symbolic fidelity may destroy structural invariance under semantic transformation.*

Proof. A channel enforcing $\hat{s} = s$ rejects all transformed sequences $T(s)$ with $T \neq \mathrm{id}$. If structural meaning is preserved across $\mathcal{T}$, these transformations represent valid instances of the same content. Excluding them enforces a stricter invariant than necessary and eliminates admissible representations of the same meaning. □

2.5 Structural Invariance

Definition 2.6 [Structural Invariant]**.** Let $\mathcal{T}$ be a set of admissible transformations on X. A function $I : X \to \mathbb{R}$ is a *structural invariant* if $I(T(x)) = I(x)$ for all $T \in \mathcal{T}$.

Proposition 2.7 [Continuity as Structural Invariance]**.** *Continuity in systems that evolve through transformation is preserved if and only if structural invariants are maintained under admissible transformations.*

Proof. If structural invariants are preserved, all admissible transformations map states to structurally equivalent states. If not preserved, transformations alter the defining structure and continuity fails. □

2.6 Conclusion

Shannon information theory operates entirely within the projected space and cannot recover what has been discarded by projection. Part II introduces the necessary extension: constraint-preserving transformation in which continuity is defined by invariance of structure.

Chapter 3

The Fallacy of Misplaced Concreteness

When the Map Replaces the World

It is one thing to simplify a system in order to understand it. It is another to forget that the simplification has taken place.

Once a representation has proven successful—predictive, stable, widely applicable—it becomes easy to treat it as if it were the thing itself. The model no longer appears as a reduction or an approximation. It appears as reality, expressed in its proper form. What was once a tool becomes a foundation.

This shift is rarely explicit. It happens gradually, as attention moves from what was omitted to what was preserved. The missing structure is no longer experienced as missing. It simply disappears from consideration. The result is a quiet inversion: instead of asking what the model leaves out, we begin to ask how the world could produce what the model describes.

Consider a familiar pattern. A measurement yields a number. The number is recorded, compared, and used to make predictions. Over time, the number comes to stand for the phenomenon itself. The conditions under which it was obtained—the choices of instrument, scale, and interpretation—fade into the background. What remains is the value, treated as if it existed independently of the process that produced it.

A similar shift occurs whenever a representation becomes stable enough to be trusted. Coordinates replace motion, symbols replace processes, categories replace experience. Each step makes the system easier to work with. But each step also removes something that cannot be recovered from the representation alone.

The difficulty is not that these reductions are incorrect. They are often extraordinarily effective. The difficulty is that their success obscures their limits. Once a representation is taken to be complete, anything that does not appear within it must either be ignored or forced into its terms.

This creates a peculiar situation. Questions arise that can only be asked because something has been excluded, but must be answered using a system that cannot represent what was excluded. The result is not a failure of reasoning, but a misplacement of what is taken to be concrete.

What was abstracted is treated as fundamental. What was omitted is treated as unreal.

3.1 Introduction

Chapters 1 and 2 established that scientific models arise from non-invertible projections and that information theory cannot represent structural invariance. The present chapter addresses the systematic error that arises when this is ignored: Alfred North Whitehead's *fallacy of misplaced concreteness*.

3.2 Formal Statement of the Fallacy

Definition 3.1 [Misplaced Concreteness]**.** The *fallacy of misplaced concreteness* occurs when the model space M is treated as an ontologically complete domain, and the projection π is implicitly inverted, so that M is taken to generate or fully determine X.

Proposition 3.2 [Non-Invertibility Implies Incompleteness]**.** *If $\pi : X \to M$ is non-invertible, then M cannot contain sufficient information to reconstruct X. Therefore M is necessarily an incomplete description.*

Proof. By Proposition 1.7, no left-inverse of π exists. Distinct elements of X map to the same element of M. Any description formulated solely in terms of M identifies these distinct states and cannot distinguish between them. □

3.3 Case Study I: Time

The projection $\pi_{\text{time}} : X \to \mathbb{R}$ maps rich temporal experience to a scalar coordinate. The real line encodes ordering but not the qualitative features of duration, the asymmetry of past and future, or the indexical character of the present. The fallacy arises when $\mathbb{R}$ is treated as the complete description of time, so that the structure of duration is not merely omitted but denied.

3.4 Case Study II: Matter

The projection $\pi_{\text{matter}} : X \to M$ maps perceptual and experimental interactions to theoretical constructs. Particles and fields are defined through the invariants of this projection and are not directly accessible in X. Materialism in its strong form asserts that M is ontologically complete, then asks how subjectivity arises from matter — a question that asks the model to produce what it was constructed to omit.

3.5 Case Study III: Measurement

The projection $\pi_{\text{meas}} : X \to M$ maps complex experimental situations to discrete outcomes. The choice of apparatus, calibration, and interpretation are part of X but not encoded in the outcome space M. The fallacy arises when measurement outcomes are treated as self-sufficient facts independent of the observational framework.

3.6 General Form

Proposition 3.3 [General Form of the Fallacy]**.** *The fallacy of misplaced concreteness occurs whenever a non-invertible projection is treated as if it were invertible.*

Proof. If a projection is treated as invertible, the model space is assumed to contain all information necessary to reconstruct the domain. Since the projection is not invertible, this assumption is false. □

Chapter 4

Phenomenology as Structural Data

What Was Discarded Was the Structure

In the process of simplifying a system, variability is often treated as something to be removed. Differences are averaged out, fluctuations are smoothed over, and what remains is taken to be the stable core of the phenomenon. This procedure is effective when the goal is to isolate quantities that can be measured and compared. But it carries an implicit assumption: that what varies is less important than what remains constant.

There is another possibility. What if the variability is not incidental, but structured? What if the differences between states are not noise, but traces of the constraints that govern how the system can change?

Consider a set of situations that all appear identical when reduced to a common description. From the perspective of the reduced system, they are the same. But when examined in their full context, they differ in subtle but systematic ways. These differences are not arbitrary. Some transitions between them are possible, others are not. Certain patterns repeat, while others never occur. The variation is not free—it is shaped.

When this structure is collapsed into a single representative value, it disappears. The resulting description is simpler, but it no longer carries information about how the system behaves under change. Two sets of states with entirely different internal organization may be indistinguishable once reduced in this way.

This suggests a reversal of perspective. Instead of treating variability as something to eliminate, it can be treated as a source of information. The way a system varies—what changes, what remains stable, and which transitions are possible—reveals the constraints under which it operates.

Phenomenological data, in this sense, is not an obstacle to analysis. It is the record of how the system is allowed to move. What appears as richness or complexity at the level of experience is the visible expression of an underlying structure that cannot be captured by invariants alone.

The task, then, is not to reduce this structure to a single point, but to understand it as a geometry in its own right.

4.1 Introduction

The structure omitted by projection is not noise to be eliminated but data of a different kind: high-dimensional, context-sensitive, and irreducible to the invariants represented in M. We call this *phenomenological data* and show that it encodes constraint structure.

4.2 Fibers as Structured Sets

Definition 4.1 [Fiber Geometry]**.** The *fiber geometry* over $m \in M$ is the metric and topological structure induced on $\pi^{-1}(m)$ by X.

4.3 Averaging and Loss of Structure

Proposition 4.2 [Non-Representativity of the Mean]**.** *The fiber mean* $\bar{x}_m = \int_{\pi^{-1}(m)} x \, d\mu_m$ *does not, in general, represent the geometry of* $\pi^{-1}(m)$*.*

Proof. The mean is a single point, whereas the fiber has non-trivial geometry. Distinct fibers with different internal geometry may have identical means. □

Corollary 4.3. *Averaging over phenomenological variability destroys information about constraint structure.*

4.4 Phenomenology as Constraint Signal

Proposition 4.4 [Constraint Encoding]**.** *The geometry of* $\pi^{-1}(m)$ *encodes information about the constraint structure governing admissible transformations in* X*.*

Proof. Admissible transformations act on X and restrict the ways in which states can vary within a fiber. The set of reachable points in $\pi^{-1}(m)$ under admissible transformations determines its geometry. Therefore the geometry reflects the underlying constraint structure. □

4.5 Transition to Constraint Theory

Part II introduces the formalism required to represent this structure. Instead of collapsing variability into invariants, constraint theory represents the rules governing variability itself.

Part II

Constraint Theory and Structural Invariants

Chapter 5

Constraint Spaces and Admissibility

From Description to Admissibility

Up to this point, the problem has been framed in terms of loss. Projection discards structure, information preserves only symbols, and models are mistaken for the reality they reduce. The result is a gap between what is represented and what actually governs how systems behave.

That gap can be approached in a different way. Instead of asking what a system *is*, we can ask what it is *allowed to do*.

Consider a system evolving over time. At each step, there are many conceivable ways it could change, but only some of those changes are possible. Certain transitions never occur. Others occur reliably. Some variations preserve the system, while others break it. These distinctions are not captured by listing states alone. They depend on the rules that govern how states can follow one another.

These rules are rarely visible when attention is restricted to isolated configurations. They appear only when we look at how configurations are connected—what can be reached from what, and under what conditions. A system is not just a collection of states; it is a space structured by the transformations that are permitted within it.

From this perspective, continuity is no longer a matter of smooth change in a representation. It is a matter of remaining within what is allowed. A trajectory is coherent not because it looks continuous, but because it never violates the constraints that define the system.

This shift replaces description with admissibility. Instead of reducing variability to invariants, we represent the structure that governs variability itself. What was previously treated as noise becomes the field in which these constraints operate.

5.1 Introduction

Part I established that projection discards constraint degrees of freedom encoded in fibers. The present chapter introduces the formal framework to represent this structure.

5.2 Constraint Structures

Definition 5.1 [Constraint Structure]**.** A *constraint structure* on a space X is a pair $(\mathcal{C}, \mathcal{T})$ where:

(i) $\mathcal{C} \subseteq \mathcal{P}(X)$ is a distinguished family of subsets of X, called the *admissible sets*;

(ii) $\mathcal{T}$ is a monoid of continuous maps $T : X \to X$, the *admissible transformations*;

(iii) for every $x \in C \in \mathcal{C}$ and $T \in \mathcal{T}$, there exists $C' \in \mathcal{C}$ with $T(x) \in C'$.

Definition 5.2 [Admissible Trajectory]. A trajectory $\{x_t\}_{t \geq 0} \subset X$ is *admissible* if there exists a family $\{C_t\} \subset \mathcal{C}$ with $x_t \in C_t$ for all t.

Proposition 5.3 [Closure Under Transformation]. *If $\{x_t\}$ is admissible and $T \in \mathcal{T}$, then $\{T(x_t)\}$ is also admissible.*

Proof. For each t, $x_t \in C_t \in \mathcal{C}$. By Definition 5.1(iii), $T(x_t) \in C'_t$ for some $C'_t \in \mathcal{C}$. □

5.3 Continuity as Constraint Preservation

Proposition 5.4 [Continuity Criterion]. *A system exhibits continuity if and only if its evolution preserves admissibility under the constraint structure.*

5.4 Constraint versus Invariance

Proposition 5.5 [Duality]. *Projection and constraint theory are dual: projection eliminates non-invariant structure, while constraint theory represents the admissible variation within that structure.*

Definition 5.6 [Constraint-Preserving Invariant]. A function $I : X \to \mathbb{R}$ is a *constraint-preserving invariant* if $I(T(x)) = I(x)$ for all $T \in \mathcal{T}$.

Chapter 6

Structural Invariance Under Transformation

What Survives Change

Once a system is understood in terms of what it allows, a new question emerges. If everything can change within the bounds of admissibility, what remains the same?

In many situations, change is not the exception but the rule. States evolve, configurations shift, representations transform. Yet across these transformations, something can persist. A pattern remains recognizable even as its form varies. A structure holds together despite differences in how it is expressed.

This persistence cannot be captured by exact equality. If every transformation had to leave a system unchanged in all details, no change would be possible at all. Instead, what persists is something more abstract: a relation, a proportion, a configuration of dependencies that survives the transformation.

Two systems may differ in their components and still share the same structure. A process may be carried out in different ways and still produce the same result. What matters is not the specific form, but the way the parts are organized and related.

This introduces a distinction between strict identity and structural identity. Strict identity requires that nothing change. Structural identity allows change, provided that what defines the system is preserved.

Understanding this distinction makes it possible to track continuity through transformation. What survives is not the state itself, but the invariants that remain unchanged under the allowed transformations of the system.

6.1 Introduction

Chapter 5 introduced constraint structures and defined continuity as preservation of admissibility. This chapter develops the corresponding notion of invariance and shows that identity is an invariant under admissible transformation.

6.2 Invariants of Admissible Transformation

Definition 6.1 [Structural Invariant]**.** A function $I : X \to \mathbb{R}$ is a *structural invariant* with respect to $(\mathcal{C}, \mathcal{T})$ if $I(T(x)) = I(x)$ for all $x \in X$ and $T \in \mathcal{T}$.

Proposition 6.2 [Invariance Along Admissible Trajectories]**.** *If $\{x_t\}$ is admissible and I is a structural invariant, then $I(x_t)$ is constant along the trajectory.*

6.3 Homomorphism and Isomorphism

Definition 6.3 [Homomorphism]**.** A map $f : X \to Y$ is a *homomorphism* if for every $T \in \mathcal{T}_X$ there exists $S \in \mathcal{T}_Y$ with $f(T(x)) = S(f(x))$ for all x.

Proposition 6.4 [Preservation of Invariants Under Homomorphism]**.** *If $f : X \to Y$ is a homomorphism and I_Y is a structural invariant on Y, then $I_X = I_Y \circ f$ is a structural invariant on X.*

6.4 Identity as Invariant

Proposition 6.5 [Identity as Structural Invariant]**.** *Identity in a system governed by $(\mathcal{C}, \mathcal{T})$ is given by the structural invariants of the system under admissible transformation.*

Chapter 7

Drift, Entropy, and Adaptive Transformation

Change Without Collapse

Once a system is understood in terms of what it allows and what it preserves, a natural concern arises. If change is permitted, and variation is everywhere, what prevents the system from dissolving into disorder?

It is tempting to equate change with loss. When a structure is altered, something is no longer exactly as it was. Differences accumulate. Over time, the system may appear to drift away from its original form. From this perspective, stability seems to require resisting change, and coherence appears fragile in the face of variation.

But this picture assumes that all change is the same. It treats every transformation as a deviation from a fixed reference, without distinguishing between those that preserve the system and those that undermine it.

In practice, systems often evolve in ways that do not destroy them. A language changes across generations yet remains intelligible. A tradition adapts to new circumstances while retaining its identity. A process unfolds over time without losing the structure that defines it. In each case, variation is present, but it is not arbitrary. It is shaped.

This suggests a distinction between two kinds of change. Some transformations operate within the limits of what the system allows. They introduce differences, but those differences remain compatible with the constraints that define the system. Other transformations exceed those limits. They relax or break the constraints, allowing configurations that no longer preserve the underlying structure.

From the outside, both may appear as increases in variability. But their effects are fundamentally different. In one case, the system becomes more expressive while remaining coherent. In the other, it loses the very structure that made coherence possible.

The task is not to prevent change, but to understand which changes a system can absorb without losing itself. What matters is not the presence of variation, but how that variation is organized.

7.1 Introduction

Not all transformation is destructive. This chapter formalises structured drift and shows that entropy increase is compatible with invariant preservation when it is governed by the constraint system.

7.2 Structured and Degenerative Entropy

Definition 7.1 [Structured Entropy Increase]**.** Entropy increase is *structured* if it occurs within the constraint structure, preserving admissibility and structural invariants.

Definition 7.2 [Degenerative Entropy Increase]**.** Entropy increase is *degenerative* if it results from relaxation or violation of constraints, leading to loss of invariants.

Theorem 7.3 [Entropy–Invariance Compatibility]**.** *Let $\{x_t\}$ be a trajectory under constraint structure $(\mathcal{C}, \mathcal{T})$. Structural invariants are preserved under entropy increase if and only if the increase is structured.*

Definition 7.4 [Adaptive Transformation]**.** A transformation T is *adaptive* if: (i) $T \in \mathcal{T}$; (ii) structural invariants are preserved; (iii) the transformation maintains the system's capacity to accommodate variation.

Definition 7.5 [Degenerative Transformation]**.** A transformation is *degenerative* if it reduces or destroys the constraint structure.

7.3 Example: Mythic Transmission

Oral transmission introduces variability, increasing entropy. As long as core relational structure is preserved, transformations remain within $\mathcal{T}$ and structural invariants persist. Literal symbol preservation minimises entropy but may lose structural meaning when representation becomes inadmissible in a new context.

Part III

Core Frameworks

Chapter 8

RSVP: Scalar–Vector–Entropy Field Theory

Giving Structure a Medium

Up to this point, the discussion has been concerned with what systems allow, what they preserve, and how they change without losing their identity. These ideas describe a structure, but they do not yet specify a medium in which that structure is realized.

A constraint system, as described so far, tells us which transformations are admissible and which are not. It distinguishes between variation that preserves structure and variation that destroys it. But it does not yet tell us how such variation unfolds, or how the balance between stability and change is maintained over time.

To make this concrete, it is necessary to move from static descriptions to a dynamical picture. Instead of treating states and transformations in isolation, we consider a system that continuously evolves, where constraints shape motion, motion explores possibilities, and variation regulates the tension between the two.

This introduces three distinct roles. There must be something that encodes where movement is restricted and where it is free. There must be something that carries the system from one configuration to another. And there must be something that governs how much variation the system can sustain without losing coherence.

These roles are not independent. If constraints are too strong, the system becomes rigid and cannot adapt. If variation is too large, the system loses its structure. If movement is unconstrained, coherence disappears. The behavior of the system depends on how these aspects interact.

The aim is not to describe a particular physical system, but to provide a general form in which constraint-preserving dynamics can be expressed. What was previously defined in terms of admissibility and invariance is now realized as a field in which structure, motion, and variability are continuously coupled.

8.1 Introduction

The present chapter introduces RSVP as a field-theoretic realisation of constraint-preserving dynamics. The core object is the triple

$$\mathbf{X}(x,t) = (\Phi(x,t),\ \boldsymbol{v}(x,t),\ S(x,t)),$$

where Φ is a scalar constraint field, $\boldsymbol{v}$ is a vector field of admissible flow, and S is an entropy field regulating variation. Full derivations appear in Appendix C.

8.2 Interpretation of the Fields

The scalar field Φ encodes constraint density: regions of high Φ strongly constrain admissible trajectories. The vector field $\boldsymbol{v}$ encodes admissible flow. The entropy field S regulates the balance between constraint preservation and variability.

8.3 Coupled Dynamics

The fields evolve according to the coupled PDE system stated in Definition C.3 of Appendix C. The term $-\nabla\Phi$ in the flow equation directs $\boldsymbol{v}$ toward regions of lower constraint potential. The coupling functions f_Φ and g determine the balance between constraint preservation and entropy increase.

8.4 RSVP as Constraint Structure

Proposition 8.1 [RSVP as Constraint Structure]**.** *The RSVP system defines a constraint structure* $(\mathcal{C}_{\mathrm{RSVP}}, \mathcal{T}_{\mathrm{RSVP}})$ *where* $\mathcal{C}_{\mathrm{RSVP}}$ *consists of field configurations satisfying stability and admissibility conditions, and* $\mathcal{T}_{\mathrm{RSVP}}$ *consists of flows generated by the RSVP dynamics.*

Chapter 9

TARTAN: Trajectory-Aware Recursive Tiling

When the Pieces Do Not Fit Together

A system can appear coherent when examined in parts. Each region behaves as expected, each local description is internally consistent, and every piece fits its immediate surroundings. Yet when these pieces are brought together, something fails. The whole cannot be assembled without contradiction.

This kind of failure is subtle because nothing is obviously wrong at the local level. Each fragment is valid in isolation. The difficulty emerges only when we attempt to reconstruct a single trajectory that passes through all of them. What was consistent in pieces becomes inconsistent as a whole.

Consider a set of overlapping maps of a terrain. Each map agrees with its neighbors where they overlap. Roads line up, landmarks coincide, and transitions between regions appear smooth. But when the maps are assembled into a single global picture, a mismatch appears. A road that should close into a loop does not meet itself. A boundary that should align is shifted. The error is not in any one map—it lies in the way they fit together.

This reveals a limitation of local reasoning. Ensuring that each part of a system behaves correctly does not guarantee that the system as a whole is coherent. Local consistency is a necessary condition for global coherence, but it is not sufficient.

To understand this, it is necessary to track not just the behavior within each region, but how these regions are connected across scales. A system that can be decomposed into parts must also include the conditions under which those parts can be reassembled into a unified whole.

The problem is not simply one of representation, but of compatibility. The pieces must not only be individually valid—they must be mutually consistent in a way that allows a single trajectory to pass through all of them without contradiction.

9.1 Introduction

TARTAN models constraint-preserving dynamics through hierarchical decomposition of the domain into tiles across scales, with explicit constraints ensuring global trajectory coherence.

9.2 Tiling and Trajectory Consistency

Definition 9.1 [Recursive Tiling]**.** A *recursive tiling* is a hierarchy of tilings $\{U_i^{(k)}\}$ indexed by scale k, such that each tile at scale k is subdivided into tiles at scale $k+1$.

Definition 9.2 [Trajectory Consistency]**.** A collection of local trajectories is *trajectory-consistent* if there exists a global trajectory $\{x_t\} \subset X$ such that for each tile $U_i^{(k)}$, $s_i^{(k)}(t) = x_t|_{U_i^{(k)}}$.

Definition 9.3 [Inter-Scale Consistency]**.** A recursive tiling is *inter-scale consistent* if for each tile $U_i^{(k)}$ and its subdivision $\{U_{i,j}^{(k+1)}\}$, $s_i^{(k)} = \bigcup_j s_{i,j}^{(k+1)}$.

Proposition 9.4 [Local Validity Without Global Coherence]**.** *A tiling may satisfy local consistency at each tile while failing to define a global trajectory.*

This limitation motivates Yarncrawler.

Chapter 10

Yarncrawler: Sheaf-Theoretic Reconstruction

10.1 Introduction

Yarncrawler formalises the conditions under which local data can be reconstructed into a global object, and provides a precise notion of obstruction. Full algebraic development appears in Appendix D.

10.2 Sheaves and Obstruction

Definition 10.1 [Yarncrawler System]**.** A *Yarncrawler system* consists of: (i) a domain X with a cover $\mathcal{U}$; (ii) a sheaf $\mathcal{F}$ assigning local states to each U_i; (iii) a procedure for computing obstruction classes associated with local data.

Definition 10.2 [Obstruction Class]**.** The *obstruction class* of a collection of local sections $\{s_i\}$ is the cohomology class $[\omega] \in \check{H}^1(\mathcal{U}, \mathcal{F})$, which vanishes if and only if the local data can be glued into a global section.

Proposition 10.3 [Global Reconstruction]**.** *A Yarncrawler system reconstructs a global state if and only if the obstruction class vanishes.*

Proposition 10.4 [TARTAN–Yarncrawler Correspondence]**.** *A TARTAN tiling defines a presheaf of local states, and trajectory consistency corresponds to the sheaf gluing condition. Failure of global coherence corresponds to a non-vanishing obstruction class.*

Chapter 11

Spherepop: Irreversible Event Calculus

The Necessity of History

A system may be coherent when viewed as a whole, and its structure may be well-defined across all of its parts. Yet this perspective omits something essential: the path by which the system arrived at its current state.

Two configurations can be identical in every observable respect and yet differ in how they came to be. One may arise through a sequence of admissible transformations, each consistent with the constraints of the system. The other may appear the same, but have no such history—no sequence of steps that leads to it without violating the rules.

From the perspective of states alone, these two situations are indistinguishable. But from the perspective of the system's evolution, they are fundamentally different. One is a valid continuation of the system's dynamics; the other is not.

This reveals a limitation of purely structural descriptions. Even when a global object can be reconstructed, and even when all constraints are satisfied, there remains a question of legitimacy: did the system arrive here in a way that is consistent with its own rules?

To answer this, it is necessary to record not just what the system is, but how it changes. Each transformation must be accounted for, and each step must be admissible relative to those that came before. The system is no longer defined solely by its configurations, but by the sequence of events that connects them.

In this setting, history is not an optional annotation. It is part of the structure itself. The order of events matters, and once a transition has occurred, it cannot be undone without altering the record of the system.

The result is a shift from reversible description to irreversible process. A system is no longer a collection of states that can be rearranged at will, but a trajectory that must be constructed step by step, with each step constrained by what has already happened.

11.1 Introduction

Spherepop introduces a decisive shift in modelling strategy. Where previous frameworks—RSVP, TARTAN, and Yarncrawler—are state-oriented (continuous fields, multiscale configurations, and sheaf-structured local data), Spherepop is fundamentally *event-oriented*. The primitive object is no longer a state in a space, but a *transition* between states.

This shift is not merely representational. It reflects a deeper structural claim: the history of a system is not recoverable from its instantaneous configuration. A system is not what it *is*, but what it has *undergone*. The ordering of transitions carries irreducible structure.

Formally, Spherepop models a system as an *irreversible log of admissible events*. Admissibility is enforced at the level of transitions, not states, and coherence is defined by the global consistency of the log.

This perspective directly mirrors the non-invertibility identified in Part I. Just as projection $\pi : X \to M$ forgets distinctions in X, the mapping from event histories to final states forgets the structure of temporal construction. Spherepop makes this loss explicit and places it at the center of the theory.

11.2 Event Structures

We begin by formalising events and logs.

Definition 11.1 [Event]**.** An *event* is an ordered pair

$$e = (s \to s'),$$

where $s, s' \in S$ are states in a state space S.

An event is not merely a relation but a *directed, typed transition*. It encodes both the origin and the destination, and implicitly the rule by which the transition occurs.

Definition 11.2 [Event Log]**.** An *event log* is a finite or countable sequence

$$\mathcal{L} = (e_1, e_2, \ldots, e_n),$$

where each event $e_k = (s_k \to s_{k+1})$ satisfies the compatibility condition

$$\mathrm{cod}(e_k) = \mathrm{dom}(e_{k+1}).$$

Equivalently, the induced state sequence

$$(s_1, s_2, \ldots, s_{n+1})$$

is well-defined.

Remark 11.3. The log $\mathcal{L}$ is not a derived object; it is the *primary representation* of the system. States appear only as nodes linking transitions.

This reverses the usual ontology: states are no longer primary, but are induced by transitions.

11.3 Admissibility of Events

Spherepop enforces structure through admissibility.

Definition 11.4 [Admissible Event]**.** An event $e = (s \to s')$ is *admissible* if

$$(s, s') \in \mathcal{R},$$

where $\mathcal{R} \subseteq S \times S$ is a constraint relation defining allowed transitions.

Definition 11.5 [Admissible Log]. An event log $\mathcal{L}$ is *admissible* if every event e_k in $\mathcal{L}$ is admissible.

Proposition 11.6 [Closure Under Prefix]. *Any prefix of an admissible log is admissible.*

Proof. Admissibility is defined pointwise. Removing events from the end of the log preserves admissibility of all remaining events. □

Remark 11.7. Admissibility is local in time but global in consequence. A single illegal transition invalidates the entire history.

11.4 Irreversibility and Temporal Structure

Spherepop enforces a strict temporal asymmetry.

Definition 11.8 [Irreversible Log]. An event log is *irreversible* if events, once appended, cannot be removed, reordered, or modified.

This property distinguishes Spherepop sharply from state space models, where trajectories may be reversible or symmetrically defined.

Proposition 11.9 [Non-Invertibility of Event Logs]. *The mapping*

$$\mathcal{L} \mapsto s_{n+1}$$

from event logs to final states is non-invertible.

Proof. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be distinct logs terminating at the same state s. Since the logs differ, their sequences of intermediate states differ, yet the mapping assigns both to s. Hence the mapping is not injective. □

Corollary 11.10. *The final state of a system does not uniquely determine its history.*

This result is the temporal analogue of projection non-invertibility. The log contains strictly more information than any instantaneous state.

11.5 Reconstruction and Completeness

Definition 11.11 [Reconstructible System]. A system is *reconstructible* if its event log uniquely determines its state trajectory.

Proposition 11.12 [Reconstructibility Condition]. *Given an initial state s_1 and an admissible event log $\mathcal{L}$, the induced state sequence is uniquely determined.*

Proof. Each event specifies a transition from s_k to s_{k+1}. By induction, the entire sequence is uniquely determined. □

Remark 11.13. Reconstructibility depends on the *completeness* of the log. Missing or corrupted events introduce ambiguity.

This mirrors the role of global sections in Yarncrawler: the log functions as a canonical reconstruction object.

11.6 Failure Modes and Coherence Breakdown

Definition 11.14 [Illegal Event]**.** An event is *illegal* if it violates the constraint relation $\mathcal{R}$.

Proposition 11.15 [Failure of Coherence]**.** *The presence of an illegal event destroys the coherence of the log.*

Proof. An illegal event violates admissibility, breaking the constraint structure. Subsequent events no longer have a valid structural interpretation, and reconstruction becomes ill-defined. □

Remark 11.16. Unlike statistical noise, which may be averaged out, illegal events are *structural violations*. They cannot be repaired locally without rewriting the log.

This sharply contrasts with Shannon-style systems, where errors may be corrected without altering global meaning. Here, the error *is* the meaning breakdown.

11.7 Mismatch with RSVP

We now formalise the tension between continuous and discrete frameworks.

Proposition 11.17 [Mismatch with RSVP]**.** *Not all RSVP-admissible trajectories can be represented as admissible Spherepop logs.*

Proof. RSVP dynamics are governed by continuous PDE evolution. These trajectories may involve infinitesimal transformations or flows that cannot be discretised into a finite sequence of admissible transitions under $\mathcal{R}$. Therefore no corresponding event log exists in Spherepop. □

Corollary 11.18 [Cross-Framework Pressure]**.** *The mismatch between RSVP and Spherepop generates* cross-framework constraint incompatibility*.*

This incompatibility is not a defect but a structural feature. It forces refinement of both representations: either the continuous dynamics must be discretised in a constraint-respecting way, or the event calculus must be expanded to accommodate finer-grained transitions.

11.8 Relation to the Blind Spot

Spherepop provides a temporal instantiation of the blind spot described by . Just as scientific models abstract away experiential variability to produce invariant descriptions, state-based models abstract away temporal construction to produce instantaneous snapshots.

The event log restores what is lost: the irreducible structure of temporal unfolding. It shows that what appears as a single state is in fact the endpoint of a non-invertible process.

11.9 Interpretation

Spherepop reveals that history is not an accessory to structure but constitutive of it. A system is defined by the constraints governing its transitions, and these constraints are only visible in the log.

In this sense, Spherepop is the discrete analogue of Yarncrawler's global section and RSVP's trajectory field. It encodes admissibility not as geometry or topology, but as *irreversible sequence.*

11.10 Conclusion

Spherepop completes the core frameworks by introducing a log-based, irreversible model of constraint-preserving dynamics. It demonstrates that temporal structure cannot be reduced to state-based representation without loss.

This loss is not incidental; it is the temporal form of the blind spot.

The next chapter introduces CLIO, the mechanism by which constraint structures are repaired across these frameworks, ensuring coherence in the presence of perturbation and cross-framework incompatibility.

Chapter 12

CLIO: Recursive Repair and Coherence

12.1 Introduction

The preceding chapters have established a unified language for describing systems in terms of constraint structures, admissible transformations, and invariant-preserving dynamics. Across RSVP, TARTAN, Yarncrawler, and Spherepop, a single principle recurs: a system is coherent if and only if its trajectories remain admissible under its governing constraints.

However, no non-trivial system remains perfectly coherent. Perturbations, measurement error, environmental change, and cross-framework incompatibilities continuously introduce violations. In the absence of a corrective mechanism, such violations accumulate, eventually destroying the constraint structure itself.

CLIO (Cognitive Loop via In-Situ Optimization) is introduced as the universal repair mechanism that prevents this degeneration. It does not generate new structure ex nihilo, nor does it impose coherence externally. Rather, it operates intrinsically on the space of constraint structures, detecting incompatibilities and adjusting the structure itself to restore admissibility.

In this sense, CLIO is not an auxiliary component but a necessary completion of constraint theory. Without repair, constraint systems are brittle. With repair, they become adaptive.

12.2 Coherence Energy and the Geometry of Incompatibility

The central object of CLIO is the *coherence energy*, which provides a quantitative measure of incompatibility.

Definition 12.1 [Coherence Energy]**.** A *coherence energy functional* is a map

$$\mathcal{E} : \mathfrak{C} \to \mathbb{R}_{\geq 0}$$

such that $\mathcal{E}(\mathcal{C}) = 0$ if and only if the constraint structure is fully coherent.

The interpretation of the functional $\mathcal{E}$ is most naturally understood in geometric rather than purely numerical terms. It does not merely assign a scalar value to a constraint structure, but encodes the extent to which the various components of that structure fail to align across different levels of representation. In particular, $\mathcal{E}$ measures the degree to which local admissibility conditions, global reconstruction requirements, and cross-framework correspondences are mutually compatible.

More concretely, the value of $\mathcal{E}(\mathcal{C})$ reflects the accumulation of incompatibilities that may arise within and between the frameworks developed in this monograph. These include, at a fundamental level, the failure of admissibility conditions to be preserved within a single formalism, as well as the breakdown of correspondence between distinct representational regimes. In the passage between continuous and discrete models, for example, a trajectory that is admissible in a field-theoretic sense may fail to admit a valid discretisation into event-log form, generating a mismatch that contributes to the overall coherence energy. Similarly, in the sheaf-theoretic setting of Yarncrawler, the presence of a non-vanishing obstruction class indicates that locally consistent data cannot be assembled into a global section; this obstruction constitutes a precise contribution to $\mathcal{E}$. At the level of irreversible dynamics, illegal transitions in a Spherepop log likewise signal a violation of admissibility that must be accounted for within the same functional.

From this perspective, $\mathcal{E}$ should not be regarded as a simple scalar quantity in the usual sense, but rather as a compressed representation of a high-dimensional incompatibility structure. It aggregates distinct modes of failure—local, global, discrete, continuous, and cross-representational—into a single ordered measure that can be compared across constraint structures. The functional therefore plays a role analogous to an energy in classical mechanics, not in virtue of any specific physical interpretation, but insofar as it induces an ordering on the space $\mathfrak{C}$ and provides a direction of descent along which coherence may be increased.

The mathematical requirements imposed on $\mathcal{E}$ are consequently essential rather than incidental. The functional must be proper, in the sense that preimages of bounded sets are compact, ensuring that sequences of constraint structures with bounded coherence energy do not escape to infinity in $\mathfrak{C}$. It must also be bounded below, guaranteeing the existence of infima toward which the repair dynamics may converge. Together, these conditions ensure that the gradient flow generated by $\mathcal{E}$ is well-posed and that the process of constraint repair, as formalised by the CLIO operator, does not diverge arbitrarily but instead evolves toward well-defined coherent regimes.

12.3 The CLIO Repair Operator

CLIO acts by descending the coherence energy landscape.

Definition 12.2 [CLIO Repair Operator]**.** The *CLIO operator* is defined by the gradient flow

$$\frac{d}{dt}\mathcal{C}(t) = -\nabla_{\mathcal{C}}\mathcal{E}(\mathcal{C}(t)).$$

This definition encodes several essential properties.

First, CLIO is *local*: it responds to the gradient of incompatibility rather than requiring global recomputation.

Second, CLIO is *continuous*: repair proceeds through incremental adjustments of the constraint structure rather than discrete resets.

Third, CLIO is *intrinsic*: the system repairs itself using its own structure, rather than relying on an external controller.

These properties distinguish CLIO from classical error correction. It is not correcting deviations relative to a fixed codebook; it is modifying the codebook itself so that the system remains admissible under evolving conditions.

12.4 Energy Descent and Stability

The fundamental property of CLIO is monotonic energy descent.

Proposition 12.3 [Energy Descent]**.**

$$\frac{d}{dt}\mathcal{E}(\mathcal{C}(t)) \leq 0$$

along the CLIO flow.

Proof. By the chain rule,

$$\frac{d}{dt}\mathcal{E} = \langle \nabla_{\mathcal{C}}\mathcal{E}, \dot{\mathcal{C}} \rangle.$$

Substituting the definition of the flow,

$$\dot{\mathcal{C}} = -\nabla_{\mathcal{C}}\mathcal{E},$$

we obtain

$$\frac{d}{dt}\mathcal{E} = -|\nabla_{\mathcal{C}}\mathcal{E}|^2 \leq 0.$$

□

Thus CLIO guarantees that incoherence cannot increase under its action. This establishes repair as a dissipative process in the space of constraint structures.

12.5 Lyapunov Structure and Dynamical Compatibility

The descent property induces a Lyapunov structure on the underlying system.

Proposition 12.4 [CLIO–Lyapunov Compatibility]**.** *If $\mathcal{E}$ is proper and bounded below, and the CLIO flow is complete, then there exists a Lyapunov function V for the induced state dynamics such that $\dot{V} \leq 0$ along admissible trajectories.*

Proof. Define $V := \mathcal{E} \circ \mathcal{C}$, viewing the coherence energy as a function of system state through the constraint structure. Since $\mathcal{E}$ decreases along the CLIO flow and is bounded below, V is non-increasing along admissible trajectories. Hence V is a Lyapunov function. □

This result is of central importance, as it establishes a direct correspondence between the abstract dynamics of constraint repair and the concrete stability properties of the system's trajectories. The coherence energy functional, when interpreted as a Lyapunov function, provides a bridge between the evolution of the constraint structure and the behaviour of the system within its state space. In effect, it ensures that the process of repairing incompatibilities at the structural level induces a stabilising influence on the trajectories governed by that structure.

More precisely, the monotonic descent of the coherence energy implies that admissible trajectories are confined to a neighbourhood of the constraint surface. Deviations from this surface, arising from perturbations or transient inconsistencies, are not allowed to grow without bound; instead, they generate a restoring gradient that drives the system back toward admissibility. In this manner, the constraint surface acts as an attractor, not in the sense of a fixed state, but as a stable locus of admissible behaviour.

Furthermore, perturbations are necessarily damped rather than amplified. Any deviation that increases the coherence energy produces a corresponding descent direction under the CLIO flow, ensuring that the system's response to disturbance is corrective rather than destabilising. This property distinguishes constraint-repair dynamics from systems lacking such a mechanism, in which small perturbations may accumulate and lead to structural breakdown.

Finally, when the coherence energy admits a strict minimum, the system exhibits asymptotic stability. In this regime, trajectories not only remain close to the constraint surface, but converge toward it over time. The strictness of the minimum guarantees that the Lyapunov function decreases except at the equilibrium configuration, ensuring that the system evolves toward a uniquely defined coherent state. Stability, in this sense, is not imposed externally but emerges as a consequence of the intrinsic geometry of the coherence energy landscape.

12.6 Fixed Points and Coherent Regimes

The long-term behaviour of CLIO is governed by its fixed points.

Definition 12.5 [Fixed Point]**.** A constraint structure $\mathcal{C}^*$ is a fixed point if

$$\nabla_{\mathcal{C}}\mathcal{E}(\mathcal{C}^*) = 0.$$

At such points, the system has resolved all accessible incompatibilities.

However, several subtleties arise.

First, fixed points need not be unique. Multiple coherent regimes may exist, corresponding to distinct local minima of $\mathcal{E}$.

Second, fixed points may be metastable. Small perturbations may not dislodge the system, but sufficiently large ones may induce transition to a different basin.

Third, the existence of a fixed point does not imply global coherence in an absolute sense, but only coherence relative to the current representational framework.

These observations anticipate the phase-transition analysis of Part VI.

12.7 Non-Convergence and Structural Failure

The CLIO repair dynamics, while guaranteeing monotonic descent of the coherence energy, do not in general ensure convergence to a fixed point. The existence of a Lyapunov structure provides a direction of evolution, but not necessarily its termination. It is therefore necessary to characterise precisely the conditions under which the flow fails to converge, and to interpret the structural significance of such failure.

Definition 12.6 [Non-Convergence]**.** The CLIO flow is said to be *non-convergent* if the trajectory $\mathcal{C}(t)$ does not approach any fixed point of the coherence energy functional as $t \to \infty$, that is, if there exists no $\mathcal{C}^*$ such that $\lim_{t\to\infty} \mathcal{C}(t) = \mathcal{C}^*$.

Non-convergence is not merely a technical complication; it has a direct structural interpretation. It corresponds to the persistence of incompatibility within the system, in the sense that no configuration of the existing constraint space is capable of simultaneously satisfying all imposed conditions. The system continues to evolve under the CLIO flow, but this evolution does not settle into a stable, coherent regime.

Several distinct mechanisms may give rise to such behaviour. One possibility is that the coherence energy landscape itself is non-convex, containing multiple basins separated by ridges, or more generally exhibiting cyclic or even chaotic regions in which the gradient flow does not admit a limit point. In this case, the system may oscillate indefinitely or wander through regions of near-equal energy without ever resolving the underlying incompatibility.

A second possibility is that the constraint space $\mathfrak{C}$ is insufficiently expressive. The incompatibility encoded in $\mathcal{E}$ may require a constraint that does not exist within the current representation. The gradient flow then attempts to reduce coherence energy within a space that lacks the degrees of freedom necessary for resolution, leading to stagnation or persistent residual energy.

A third and more structurally significant case arises when the incompatibility is not internal to a single framework, but emerges from the interaction of multiple frameworks. Cross-framework conflicts—such as those between continuous dynamics and discrete event structures, or between local consistency and global reconstructibility—may not admit a resolution within any one of the existing constraint structures. In such cases, local repair is insufficient, and the system requires an expansion of its representational capacity.

In all of these cases, non-convergence functions as a diagnostic signal. It indicates that the current constraint structure has reached the limit of its capacity to absorb and resolve incompatibility. The failure of the CLIO flow to converge is not a failure of the repair mechanism itself, but a reflection of the inadequacy of the space in which it operates. Repair cannot proceed because the constraints required to achieve coherence are not yet available within the system.

This condition admits a precise mathematical interpretation as the boundary of applicability of the current framework. The system is driven by the coherence energy toward a region in which no admissible minimum exists, and the absence of such a minimum prevents the establishment of a stable constraint surface.

It is in this sense that non-convergence anticipates what will later be formalised as a *phase transition*. The system does not collapse into incoherence; rather, it is forced toward a transformation of its underlying structure. The absence of convergence is thus not merely a negative result, but the positive indication that a change of representation is required. The subsequent development will make this notion precise by characterising the conditions under which such transitions occur and the manner in which new constraint structures emerge from the failure of the old.

12.8 CLIO as the Engine of Adaptive Drift

We now connect CLIO to the theory of structured entropy developed in Chapter 7.

Proposition 12.7 [Repair as Adaptive Drift]**.** *CLIO implements adaptive drift: it increases entropy in a constraint-structured manner while preserving invariants.*

Proof. CLIO modifies the constraint structure to reduce incompatibility. This typically increases the dimensionality of admissible sets (entropy increase) while preserving the core invariants that define admissibility. Hence the drift is structured. □

Thus CLIO provides the formal mechanism underlying the distinction between healthy transformation and degenerative collapse.

12.9 Interpretation and Philosophical Context

CLIO also clarifies the role of subjectivity in systems of knowledge.

As emphasized in [1], scientific models are constructed through processes that necessarily involve human experience, even when that experience is abstracted away in the final representation .

CLIO formalises the return of what is abstracted away. When a model fails to capture the structure of reality, the resulting incompatibility manifests as increased coherence energy. Repair then requires reintegration of the omitted degrees of freedom.

In this sense, CLIO is the operational counterpart of the critique of the blind spot. It ensures that projection errors do not remain hidden indefinitely; they generate pressure that forces refinement of the constraint structure.

12.10 Conclusion

CLIO completes the architecture of constraint theory.

RSVP provides the continuous dynamics of constraint and flow. TARTAN provides multiscale trajectory structure. Yarncrawler guarantees global reconstructibility. Spherepop ensures logical admissibility of discrete evolution.

CLIO binds them together. It is the mechanism through which systems maintain coherence under perturbation, resolve incompatibilities, and evolve without losing their invariants.

Without CLIO, constraint systems decay. With CLIO, they learn.

Chapter 13

Constraint Incompatibility as Explanatory Pressure

When Frameworks Collide

A single framework can appear complete when considered in isolation. It defines what is admissible, what is preserved, and how systems evolve within its own structure. As long as attention remains within its boundaries, coherence can be maintained.

The situation changes when multiple frameworks are applied to the same domain.

Each framework imposes its own constraints. It distinguishes between what is possible and what is not, according to its internal logic. When two such systems are brought together, their distinctions do not necessarily align. A state that is admissible in one framework may be inadmissible in another. A trajectory that is coherent from one perspective may violate the constraints of a different one.

This conflict cannot be resolved by ignoring one side. If both frameworks capture genuine aspects of the system, then the incompatibility reflects a limitation of the structures being used. The system itself is not contradictory; the representations of it are.

This creates a form of pressure. The existing structures are no longer sufficient to account for what is being observed. The mismatch forces a choice: either weaken the constraints to eliminate the conflict, or expand the framework to accommodate it.

Weakening the constraints produces a superficial resolution. The conflict disappears, but so does the ability to distinguish meaningful structure. Expansion, by contrast, preserves the original distinctions while embedding them into a larger system in which their relationship can be made coherent.

In this sense, incompatibility is not a defect. It is a signal that the current representation is incomplete. The tension between frameworks is what drives the development of more expressive structures.

The task is not to eliminate this tension, but to use it. The conflict between admissibility conditions becomes the source of explanatory pressure, forcing the system to evolve into a form capable of containing all of its own constraints.

13.1 Introduction

The preceding parts of this monograph developed multiple formal frameworks—continuous, discrete, logical, and topological—each designed to preserve structural invariants under transformation. These frameworks were not constructed as mutually exclusive theories, but as complementary realizations of a single underlying principle: that coherence is maintained through constraint preservation rather than symbolic fidelity. However, once these frameworks are brought into contact and applied to the same domain, a deeper structural phenomenon emerges. The frameworks do not merely align; they interfere.

This interference is not accidental. It is the natural consequence of applying distinct constraint systems to a shared underlying domain. Each framework encodes its own notion of admissibility, its own criteria for valid transformation, and its own invariants. When a state or trajectory satisfies the admissibility conditions of one framework but violates those of another, the system enters a regime of incompatibility. It is precisely this regime that generates what will be called *explanatory pressure*.

Explanatory pressure is not a defect in the system. It is the mechanism by which the system reveals its own incompleteness. A framework that never encounters incompatibility is not complete; it is isolated. Only through the encounter with incompatible constraints does a framework expose the limits of its representational capacity. In this sense, incompatibility is the structural engine of theoretical progress. It is the point at which coherence cannot be maintained within the existing representation, and a new representation becomes necessary.

The broader philosophical significance of this claim can be situated within the critique of scientific objectification discussed earlier. Scientific models, as emphasized by , operate by abstraction and idealization, producing invariant descriptions that are communicable across observers. Yet these models are necessarily partial; they omit the variability and structure of lived experience that cannot be captured within their formal language. The blind spot arises when this omission is forgotten and the model is treated as ontologically complete. Cross-framework incompatibility reintroduces what was omitted. It forces the system to recognize that no single projection exhausts the domain, and that coherence must be restored at a higher level of structure.

13.2 Constraint Incompatibility

Definition 13.1 [Constraint Incompatibility]. A state $x \in X$ exhibits *constraint incompatibility* between $(\mathcal{C}_1, \mathcal{T}_1)$ and $(\mathcal{C}_2, \mathcal{T}_2)$ if $x \in C_1 \in \mathcal{C}_1$ but $x \notin C_2$ for all $C_2 \in \mathcal{C}_2$.

This definition captures a minimal but fundamental condition. A state may be perfectly admissible within one framework while being entirely inadmissible within another. The incompatibility is not located in the state itself but in the relation between constraint systems. It is a relational failure of coherence, not a local defect.

The importance of this distinction cannot be overstated. If incompatibility were treated as a property of the state, the natural response would be to discard the state as invalid. This is precisely the move made in reductive frameworks: anomalies are treated as noise or error. In the present framework, however, incompatibility is interpreted as a signal. It indicates that the current constraint structures are insufficient to jointly represent the domain. The state is not rejected; the representation is.

To see this more concretely, consider the interaction between continuous and discrete frameworks. A trajectory that is smooth and well-defined under a continuous field theory may fail to admit a valid discretization into irreversible events. Conversely, a sequence of admissible discrete transitions may correspond to no smooth trajectory in the continuous model. In each case, the incompatibility does not indicate that one framework is correct and the other is false. It indicates that the domain exceeds the representational capacity of either framework in isolation.

This interpretation aligns with the broader critique of model-based reasoning. As emphasized in the discussion of scientific projection, models are inherently non-invertible: they discard information in order to achieve invariance. When two models discard different information, their images may fail to align. Constraint incompatibility is the formal manifestation of this non-invertibility across representations.

Proposition 13.2 [Necessity of Expansion]. *The only non-degenerative resolution of incompatibility is expansion of the constraint structure to* $(\mathcal{C}', \mathcal{T}')$ *such that each original structure embeds as a constraint-bearing substructure.*

The proposition asserts that incompatibility cannot be resolved by contraction without loss. Any attempt to force agreement by weakening constraints or discarding structure results in degenerative drift, a phenomenon that will be analyzed in Part VI. The only viable resolution is expansion, in which a larger constraint system is constructed that contains each original system as a specialized case.

The proof follows from the structure of invariants. Each framework encodes invariants that must be preserved for the framework to retain its identity. If one framework is reduced to fit another, its invariants are lost. By contrast, an expansion that embeds both frameworks preserves their invariants while introducing new degrees of freedom that allow previously incompatible states to coexist.

13.3 Explanatory Pressure

Definition 13.3 [Explanatory Pressure]. *Explanatory pressure* is the force exerted by constraint incompatibility, requiring expansion of the constraint structure to restore coherence.

Explanatory pressure is not a metaphor but a structural quantity. It can be understood as the gradient of incoherence in the space of constraint structures. When two frameworks are incompatible, the coherence energy of the system increases. The CLIO operator, introduced earlier, acts to minimize this energy. The direction of this minimization is determined by the structure of the incompatibility, and in general leads to expansion of the constraint space.

What distinguishes explanatory pressure from ordinary error is that it cannot be eliminated locally. A local fix that resolves incompatibility in one region typically creates new incompatibilities elsewhere. The pressure is global. It reflects a mismatch between entire constraint systems, not isolated states. As a result, its resolution requires a global reconfiguration of the representation.

This notion provides a precise formalization of a familiar phenomenon in the development of knowledge. Scientific theories do not evolve solely by accumulation of data. They evolve when existing frameworks become incapable of integrating new observations without contradiction. The resulting tension forces the development of new representations that can accommodate both the old and the new. Explanatory pressure is the formal analogue of this process.

It also clarifies why certain anomalies are productive while others are ignored. An anomaly that can be absorbed within an existing framework without altering its constraint structure generates little pressure. An anomaly that violates the core admissibility conditions of the framework generates maximal pressure. The difference is not empirical but structural.

13.4 Constraint Closure

Definition 13.4 [Constraint Closure]**.** A system exhibits *constraint closure* if its constraint structure is sufficient to resolve all internal incompatibilities. Constraint closure is achieved when the CLIO operator converges to a fixed point.

Constraint closure represents the equilibrium state of the system under explanatory pressure. At closure, all incompatibilities have been resolved within a unified constraint structure. The system is coherent in the sense that every admissible state is compatible with every component framework embedded within it.

It is crucial to note that closure does not imply completeness in an absolute sense. Because the underlying domain is non-invertible under projection, new incompatibilities may arise as the system encounters new data or new representations. Closure is therefore always relative to a given set of frameworks and constraints. It is a dynamically maintained condition rather than a final state.

The convergence of the CLIO operator provides a precise criterion for closure. When the coherence energy reaches a minimum and the gradient vanishes, the system has reached a fixed point in constraint space. At this point, no further expansion is required to resolve existing incompatibilities. The system is stable under perturbation, and its constraint surface, as defined in Part V, is well-formed.

This notion of closure also provides a bridge to the theory of identity developed later. A refractive system maintains its identity precisely by maintaining constraint closure under transformation. When closure fails, identity becomes unstable. When closure is restored, identity re-emerges as a stable invariant.

13.5 Conclusion

Constraint incompatibility is not an obstacle to knowledge but its generative core. It is the point at which representations fail to align, and therefore the point at which deeper structure must be introduced. Explanatory pressure, arising from this incompatibility, drives the system toward expansion, embedding, and eventual closure.

The framework developed in this chapter establishes a general principle: coherence across representations is not given but constructed. It requires the continual expansion of constraint structures in response to incompatibility. This process is governed not by arbitrary choice but by the invariant of resolution: the requirement that all prior structures be preserved as constraint-bearing components of the expanded system.

The next chapter formalizes this process of expansion, providing the precise conditions under which incompatibility can be resolved without loss of structure.

Chapter 14

Dimensional Expansion and Embedding

Resolving Without Losing

When two constraint systems conflict, the tension cannot be removed by simply choosing one over the other. Each captures a part of the structure of the system, and discarding either results in loss. Yet keeping both within the same representation leads to contradiction.

This leaves only one possibility: the representation itself must change.

Such a change is not a matter of adding new elements to an existing structure. If the original constraints are weakened to accommodate the conflict, the system loses its ability to distinguish meaningful configurations. The apparent resolution comes at the cost of explanatory power.

A genuine resolution must preserve what each framework captures while making their coexistence possible. This requires moving to a space in which both sets of constraints can be maintained without contradiction. The original systems must appear within the new one, not as approximations, but as necessary components.

This process can be understood as an expansion in dimension. What could not be represented simultaneously in the original space becomes compatible when additional degrees of freedom are introduced. The conflict is not eliminated; it is re-expressed in a form that can be sustained without inconsistency.

Not every expansion achieves this. Some introduce new structure without preserving what came before. Others obscure the original distinctions by embedding them in a way that no longer constrains the system. A valid expansion must satisfy a stronger condition: it must retain the original constraints as active, indispensable elements of the new system.

In this sense, expansion is not arbitrary growth. It is a controlled transformation that resolves incompatibility while preserving the invariants that define the system. The success of such a transformation is measured not by the absence of conflict, but by the preservation of structure across the change of representation.

14.1 Introduction

Chapter 13 established that constraint incompatibility is not an accidental feature of multi-framework systems but a structural inevitability. Whenever two constraint systems are defined over projections of a richer domain, their admissibility conditions will, in general, fail to coincide. This mismatch generates what we termed *explanatory pressure*: a requirement that the system either collapse distinctions (degenerative drift) or expand its representational capacity.

The present chapter formalises the second path. We show that the only non-degenerative resolution of constraint incompatibility is *dimensional expansion*: the construction of a larger constraint structure in which the original systems embed as constraint-bearing substructures. This expansion is not merely additive. It is governed by strict conditions ensuring that previously established invariants are preserved and remain operational.

The central result is the Adaptive Expansion Theorem, which characterises precisely when such an expansion preserves structure. This theorem provides the formal criterion for distinguishing genuine theoretical progress from collapse disguised as generalisation.

14.2 Admissible Embedding

We begin by refining the notion of embedding introduced in Chapter 13. The purpose of this refinement is to ensure that embedding preserves not merely sets of states, but the full constraint structure governing their admissible evolution.

Definition 14.1 [Constraint Embedding]**.** Let $(\mathcal{C}, \mathcal{T})$ and $(\mathcal{C}', \mathcal{T}')$ be constraint structures on domains X and X' respectively. We say that $(\mathcal{C}, \mathcal{T})$ *embeds* into $(\mathcal{C}', \mathcal{T}')$ via an injective map

$$\iota : X \hookrightarrow X'$$

if the following conditions hold:

(i) For every admissible set $C \in \mathcal{C}$, the image $\iota(C)$ is an admissible set in $\mathcal{C}'$;

(ii) For every admissible transformation $T \in \mathcal{T}$, there exists a corresponding transformation $T' \in \mathcal{T}'$ such that

$$\iota \circ T = T' \circ \iota.$$

This definition ensures that both the *static* and *dynamic* aspects of the system are preserved. The first condition guarantees that admissibility is respected at the level of states, while the second guarantees that the allowed evolutions of those states are preserved under the embedding.

Remark 14.2. Constraint embedding is strictly stronger than set-theoretic inclusion. It requires preservation of admissibility and transformational structure. In particular, an embedding prevents the collapse of distinctions encoded in the original constraint system.

14.3 Dimensional Expansion

We now formalise the notion of expansion that resolves constraint incompatibility.

Definition 14.3 [Dimensional Expansion]**.** A *dimensional expansion* of a constraint structure $(\mathcal{C}, \mathcal{T})$ on X is a new structure $(\mathcal{C}', \mathcal{T}')$ on a domain X' such that:

(i) $(\mathcal{C}, \mathcal{T})$ embeds into $(\mathcal{C}', \mathcal{T}')$ via some $\iota : X \hookrightarrow X'$;

(ii) X' strictly enlarges the representational capacity of X (e.g. by increasing dimensionality, adding new degrees of freedom, or enriching structure).

The key feature of dimensional expansion is that it resolves incompatibility not by weakening constraints, but by *lifting* the system into a space where the previously incompatible structures can coexist.

Proposition 14.4 [Resolution via Expansion]**.** *Let $(\mathcal{C}_1, \mathcal{T}_1)$ and $(\mathcal{C}_2, \mathcal{T}_2)$ be incompatible constraint structures. If there exists a dimensional expansion $(\mathcal{C}', \mathcal{T}')$ into which both embed, then their incompatibility is resolved in X'.*

Proof. By embedding, each structure is represented as a constraint-preserving substructure of $(\mathcal{C}', \mathcal{T}')$. Since admissibility is defined relative to $\mathcal{C}'$, configurations that were incompatible in the original spaces can be represented as distinct but co-admissible configurations in X'. The incompatibility is thus resolved by separation within a richer structure. □

14.4 Adaptive Expansion

Not all expansions preserve structure. Some expansions resolve incompatibility only by diluting the constraints that made the original systems meaningful. To distinguish these cases, we introduce the notion of adaptive expansion.

Definition 14.5 [Adaptive Expansion]**.** A dimensional expansion is *adaptive* if it satisfies the following three conditions:

(i) **Non-reductive reconciliation:** each original constraint structure embeds into the expanded system without loss of admissible sets or transformations;

(ii) **Dimensional expansion over constraint dilution:** incompatibilities are resolved by introducing new degrees of freedom, not by weakening existing constraints;

(iii) **Structural necessity:** the embedded structures remain necessary components of the expanded system, continuing to constrain admissible behaviour.

Remark 14.6. These conditions formalise the intuitive idea that genuine expansion preserves the *work* performed by the original constraints. An expansion that renders prior constraints redundant or vacuous is degenerative.

14.5 The Adaptive Expansion Theorem

We now state and prove the central result of the chapter.

Theorem 14.7 [Adaptive Expansion Theorem]**.** *Let $(\mathcal{C}_1, \mathcal{T}_1)$ and $(\mathcal{C}_2, \mathcal{T}_2)$ be constraint structures with incompatibilities. A resolution $(\mathcal{C}', \mathcal{T}')$ preserves all structural invariants of the original systems if and only if it is an adaptive expansion.*

Proof. (*If*) Suppose $(\mathcal{C}', \mathcal{T}')$ is an adaptive expansion. By non-reductive reconciliation, both $(\mathcal{C}_1, \mathcal{T}_1)$ and $(\mathcal{C}_2, \mathcal{T}_2)$ embed into $(\mathcal{C}', \mathcal{T}')$, so their admissible sets and transformations are preserved. By dimensional expansion, incompatibilities are resolved through added degrees of freedom rather than relaxation of constraints. By structural necessity, the embedded structures continue to constrain the expanded system. Therefore all structural invariants of the original systems are preserved.

(*Only if*) Suppose $(\mathcal{C}', \mathcal{T}')$ preserves all structural invariants. Then the original structures must embed without loss (non-reductive reconciliation). If constraints were weakened, invariants would not be preserved, so resolution must occur via dimensional expansion. Finally, if the original structures were not necessary within the expanded system, their invariants would be inert, contradicting their preservation. Hence the expansion is adaptive. □

14.6 Interpretation

The Adaptive Expansion Theorem provides a precise answer to the question raised in Chapter 13: how can incompatible frameworks be unified without loss of meaning?

The answer is that unification is not a matter of merging representations, but of constructing a higher-dimensional space in which each representation appears as a necessary projection. The expansion must preserve the constraint structure of each framework while introducing new degrees of freedom sufficient to resolve their incompatibilities.

This formal result mirrors the critique articulated in the discussion of the scientific "blind spot"—namely, that models are partial projections of a richer domain and cannot be inverted to recover what they omit . Dimensional expansion is precisely the operation that restores, at a higher level, the structural information lost in projection without pretending that the projection itself was complete.

14.7 Conclusion

We have shown that dimensional expansion is the unique non-degenerative mechanism for resolving constraint incompatibility, and we have characterised the conditions under which such expansion preserves structural invariants.

The Adaptive Expansion Theorem will serve as the central criterion for evaluating theoretical development throughout the remainder of the monograph. In the next chapter, we turn to the problem of reconstruction across representations, showing how expanded systems relate to their predecessors and how invariants are preserved under translation between forms.

Chapter 15

Reconstruction Across Representations

Reconstruction Is Not Unique

Up to this point, reconstruction has been treated as a problem of compatibility. Given multiple representations of a system, the question has been whether they can be assembled into a coherent whole. When they can, the system is well-defined. When they cannot, an obstruction prevents reconstruction.

There is a further question that arises once reconstruction is possible. Even when local data is consistent and a global structure exists, is that structure unique?

In general, it is not. The same local information can admit multiple globally distinct realizations, each consistent with all admissibility conditions. From the perspective of the local pieces, these realizations are indistinguishable. The difference appears only at the level of global structure.

This reveals that reconstruction is not simply a matter of assembling parts, but of selecting among multiple admissible completions. The ambiguity is not due to missing information in the usual sense, but to degrees of freedom that are not fixed by local constraints.

These degrees of freedom are structured. They are not arbitrary choices, but arise from symmetries and transformations that preserve all local data while producing distinct global configurations. The space of possible reconstructions is therefore not a single object, but a family organized by these residual symmetries.

At the same time, reconstruction is not purely local. As one moves across a system, the way local structures change carries information that cannot be recovered from any single region alone. This introduces a form of memory: the global configuration depends not only on local states, but on how those states vary across the domain.

Taken together, these observations lead to a refined picture. Reconstruction is governed by three elements: local admissibility, global compatibility, and residual ambiguity. When compatibility fails, an obstruction prevents reconstruction. When compatibility holds, multiple inequivalent reconstructions may still exist, distinguished by their global structure.

The purpose of this chapter is to formalise these phenomena. Reconstruction is not the inverse of projection. It is a constrained process that preserves invariants, detects obstruction, and admits structured non-uniqueness in its solutions.

15.1 Introduction

Given multiple representations of a system, under what conditions can a coherent structure be reconstructed across them?

15.2 Reconstruction Maps

Definition 15.1 [Admissible Reconstruction]**.** A reconstruction map $R : \prod_i M_i \to X'$ is *admissible* if it preserves the constraint structure: $R(\pi_i(x)) \in C'$ whenever $x \in C$.

Proposition 15.2 [Generic Non-Invertibility]**.** *Reconstruction from representations is generically non-invertible, since each projection π_i discards information that cannot be recovered.*

Proposition 15.3 [Invariant Preservation]**.** *Admissible reconstruction preserves structural invariants.*

Definition 15.4 [Constraint-Equivalence]**.** Two representations are *constraint-equivalent* if they preserve the same structural invariants.

Proposition 15.5 [Reconstruction as Gluing]**.** *Reconstruction corresponds to gluing local sections across representations in the Yarncrawler framework. Failure corresponds to a non-vanishing obstruction class.*

With Part IV complete, we have shown how frameworks interact, how incompatibilities generate expansion, and how coherence is maintained across representations. Part V derives identity as an invariant of this process.

15.3 Fibration Structure and Canonical Decomposition

The reconstruction theory developed in this chapter has a precise algebraic-geometric analogue in the study of fibered K-trivial varieties — varieties whose canonical bundle is numerically trivial and which admit a fibration over a lower-dimensional base. The correspondence is not merely illustrative; the formal structures are identical, and translating between the two settings yields new tools for both.

Let $f : \mathcal{X} \to Y$ be a fibration from a total field configuration $\mathcal{X}$ to a base trajectory manifold Y, with fibers $\mathcal{X}_y = f^{-1}(y)$ representing local state spaces. This is a direct realisation of the representation family $\{\pi_i : X \to M_i\}_{i \in I}$ of the preceding sections, with Y serving as the base of the indexing and each fiber serving as a local model.

In algebraic geometry, the geometry of such a system is governed by the *canonical bundle formula*:

$$K_{\mathcal{X}} \sim_{\mathbb{Q}} f^*(K_Y + B_Y + M_Y)\,,$$

where K denotes the canonical bundle, B_Y encodes the singular locus of the fibration (points where fibers degenerate), and M_Y is the *moduli divisor* recording variation of fiber type across Y.

We now interpret each component in the constraint-theoretic language of this monograph.

Definition 15.6 [Canonical Decomposition of a Field Configuration]**.** Let $f : \mathcal{X} \to Y$ be a fibered constraint system with RSVP state $\mathbf{X} = (\Phi, \boldsymbol{v}, S)$. The *canonical decomposition* of Φ is:

$$\Phi_{\mathcal{X}} \approx f^* \left(\underbrace{\Phi_Y}_{\text{base constraints}} + \underbrace{\Phi_{\text{sing}}}_{\text{defect field}} + \underbrace{\Phi_{\text{var}}}_{\text{variation memory}} \right),$$

where:

(i) Φ_Y is the constraint density inherited from the base trajectory manifold Y, corresponding to K_Y;

(ii) Φ_{sing} is the *defect field*, encoding constraint failures at degenerate fibers, corresponding to the boundary divisor B_Y;

(iii) Φ_{var} is the *variation memory field*, recording how local state spaces change as one moves along trajectories in Y, corresponding to the moduli divisor M_Y.

Remark 15.7. The canonical decomposition is a compression: the full high-dimensional constraint structure of $\mathcal{X}$ is encoded by a lower-dimensional base field plus two correction terms. This is precisely the structure of admissible reconstruction (Definition 15.1): the total system is reconstructed from base invariants plus residual degrees of freedom.

The variation memory field Φ_{var} deserves particular attention. In algebraic geometry it arises from *variation of Hodge structure*: as one moves along a path γ in Y, the Hodge decomposition of the fiber cohomology changes, and M_Y records this change. In the constraint-theoretic language, this is a nonlocal, path-dependent field:

Definition 15.8 [Variation Memory Field]**.** The *variation memory field* $\Phi_{\text{var}} : Y \to \mathbb{R}$ is a trajectory functional satisfying:

(i) it is not determined by the local value of the fiber at any single point;

(ii) its value at $y \in Y$ depends on the history of fiber variation along paths terminating at y;

(iii) it is invariant under homotopies of paths that preserve fiber type.

The variation memory field is therefore a nonlocal memory functional over the base — the constraint-theoretic analogue of the moduli divisor M_Y.

Proposition 15.9 [Period Map as Projection Operator]**.** *Variation of fiber structure across Y induces a* period map

$$\mathcal{P} : Y \dashrightarrow \mathcal{M},$$

from the base trajectory manifold into a moduli space $\mathcal{M}$ of fiber types. This map is a projection operator in the sense of Chapter 1: it is generically non-injective, and the variation memory field Φ_{var} is the pullback of a universal constraint structure on $\mathcal{M}$.

Proof. The period map assigns to each $y \in Y$ the isomorphism class of the fiber $\mathcal{X}_y$ in the moduli space $\mathcal{M}$. It is non-injective when distinct points of Y carry isomorphic fibers. The moduli divisor $M_Y = \mathcal{P}^* L$ is the pullback of a natural line bundle L on $\mathcal{M}$, making $\Phi_{\text{var}} = \mathcal{P}^* \phi_{\mathcal{M}}$ the pullback of the universal constraint density on $\mathcal{M}$. ∎

15.4 The Tate–Shafarevich Obstruction

The canonical decomposition of Section 15.3 encodes the *local* data of a fibered system. A deeper question arises: given the base Y, the defect field Φ_{sing}, and the variation memory Φ_{var}, is the total field configuration $\mathcal{X}$ uniquely determined?

The answer is no, and the measure of non-uniqueness is the Tate–Shafarevich group.

Definition 15.10 [Reconstruction Symmetry Sheaf]**.** Let $f : \mathcal{X} \to Y$ be a fibered constraint system with fiber symmetry group $\mathcal{P}$. The *reconstruction symmetry sheaf* is the sheaf $\mathcal{P}$ over Y assigning to each open set $U \subseteq Y$ the group of admissible fiber reparametrisations over U.

Definition 15.11 [Tate–Shafarevich Obstruction Group]**.** The *Tate–Shafarevich obstruction group* of the fibered system is

$$\text{Sha}_Y \;\cong\; H^1(Y, \mathcal{P}),$$

the first Čech cohomology group of the reconstruction symmetry sheaf.

The Tate–Shafarevich group is precisely the obstruction class of Appendix D instantiated in the fibration setting. We now make this identification explicit.

Proposition 15.12 [Tate–Shafarevich as Gluing Obstruction]**.** *Let $\{U_i\}$ be an open cover of Y over which the fibered system $\mathcal{X}$ trivialises locally. Each local trivialisation defines a local section $s_i \in \mathcal{P}(U_i)$. On overlaps, the sections differ by transition functions:*

$$g_{ij} = s_i^{-1} s_j \in \mathcal{P}(U_i \cap U_j).$$

The consistency condition on triple overlaps, $g_{ij} g_{jk} g_{ki} = 1$, makes (g_{ij}) a Čech 1-cocycle. The obstruction class

$$[\omega] = [(g_{ij})] \in H^1(Y, \mathcal{P}) = \text{Sha}_Y$$

vanishes if and only if the local trivialisations can be chosen to agree globally — i.e., if and only if $\mathcal{X}$ admits a global section.

Proof. This follows directly from Theorem D.10 of Appendix D, applied to the sheaf $\mathcal{P}$ and the local sections $\{s_i\}$. The transition functions (g_{ij}) are the Čech 1-cochains, and the triple consistency condition is the cocycle condition $\delta^1(g_{ij}) = 0$. □

Remark 15.13. Proposition 15.12 gives the Tate–Shafarevich group a precise interpretation in the language of this monograph: it classifies globally distinct admissible reconstructions from locally valid pieces. Two elements of Sha_Y that differ by a coboundary correspond to reconstructions that are related by a global reparametrisation and therefore represent the same physical configuration. Those that differ by a non-trivial class represent genuinely inequivalent global realisations of the same local data.

15.5 The Obstruction Field and Extended RSVP State

The Tate–Shafarevich group assigns a cohomological obstruction to each fibered system. We now integrate this into the RSVP field structure as a fourth component.

Definition 15.14 [Obstruction Field]**.** The *obstruction field* of a fibered constraint system $f : \mathcal{X} \to Y$ is the map

$$\mathcal{O}_Y : \{\text{loops } \gamma \subset Y\} \to \mathcal{P},$$

defined by

$$\mathcal{O}_Y(\gamma) := K(\gamma) - \mathrm{Id},$$

where $K(\gamma)$ is the transport operator around the loop γ — the result of parallel-transporting a fiber state around γ using the connection induced by the variation memory field Φ_{var}.

Remark 15.15. The obstruction field $\mathcal{O}_Y$ measures path-dependence of reconstruction: if $\mathcal{O}_Y(\gamma) \neq 0$ for some loop γ, then transporting a fiber state around γ does not return it to its original configuration, and no globally consistent reconstruction exists. This is the holonomy of the connection induced by Φ_{var}, and it is non-trivial precisely when $[\omega] \neq 0$ in Sha_Y.

In trajectory-functional language, the obstruction appears as a non-vanishing loop integral. Define the trajectory functional

$$\mathcal{F}[\gamma] = \int_\gamma \boldsymbol{v} \cdot d\ell + \Phi_{\mathrm{var}},$$

which integrates the admissible flow and variation memory along a path. Then:

Proposition 15.16 [Loop Integral as Obstruction Detector]**.** *The obstruction field vanishes on a loop γ if and only if*

$$\oint_\gamma \mathcal{F} = 0.$$

A non-vanishing loop integral certifies the existence of a non-trivial element of Sha_Y *and therefore the impossibility of a globally consistent reconstruction over the region bounded by γ.*

Proof. By Stokes' theorem applied to the connection induced by Φ_{var}, the loop integral $\oint_\gamma \mathcal{F}$ equals the holonomy of the connection around γ. This holonomy is trivial if and only if the connection is flat over the region bounded by γ, which is equivalent to the vanishing of the local obstruction class. □

Definition 15.17 [Extended RSVP State]**.** The *extended RSVP state* of a fibered constraint system is the quadruple

$$\mathbf{X}_{\mathrm{global}} = (\Phi,\ \boldsymbol{v},\ S,\ \mathcal{O}),$$

where $\mathcal{O}$ is the obstruction field of Definition 15.14. The obstruction field is not a correction term or an error; it is a structural invariant of the system measuring whether local admissibility integrates into global coherence.

Proposition 15.18 [Obstruction as Necessary Component]**.** *$\mathcal{O} \neq 0$ if and only if the system has no globally admissible reconstruction from its local data. Systems with $\mathcal{O} = 0$ are globally reconstructible; those with $\mathcal{O} \neq 0$ require a medium shift (Chapter 21) to reach a representation in which the obstruction is trivialised.*

Proof. $\mathcal{O} \neq 0$ means some loop γ has non-trivial holonomy, so $[\omega] \neq 0$ in Sha_Y, so by Theorem D.10 no global section exists. Conversely, $\mathcal{O} = 0$ everywhere implies all loops have trivial holonomy, the connection is flat, $[\omega] = 0$, and a global section exists. The medium shift requirement follows from Corollary D.15: obstruction trivialisation requires passage to a refined cover of a strictly larger space. □

15.6 Decomposition of the Obstruction and Boundedness

The Tate–Shafarevich group decomposes according to the geometry of the fibration. For Calabi–Yau fibrations (fibers with trivial canonical bundle), Sha_Y is a finite group, meaning the obstruction has only discrete modes. For symplectic fibrations (fibers with a holomorphic symplectic form), Sha_Y decomposes into a finite part and a continuous part.

Proposition 15.19 [Decomposition of the Obstruction Field]. *The obstruction field decomposes as*

$$\mathcal{O}_Y = \mathcal{O}_{\text{torsion}} \oplus \mathcal{O}_{\text{continuous}},$$

where:

(i) $\mathcal{O}_{\text{torsion}}$ *is the* discrete obstruction*: finite-order holonomy representing a finite number of inequivalent global reconstructions;*

(ii) $\mathcal{O}_{\text{continuous}} \cong \mathbb{C}^n$ *is the* continuous obstruction*: smooth deformation directions in reconstruction space, present in the symplectic case and absent in the Calabi–Yau case.*

This decomposition fits into the exact sequence

$$0 \to \mathcal{D}_Y \to \mathcal{O}_Y \to \mathcal{T}_Y \to 0,$$

where $\mathcal{D}_Y$ *is the continuous deformation operator space and* $\mathcal{T}_Y$ *is the discrete torsion operator space.*

The main boundedness theorem of [15] now admits a direct reformulation in constraint-theoretic terms.

Theorem 15.20 [Constraint Closure Implies Boundedness]. *Let* $f : \mathcal{X} \to Y$ *be a fibered constraint system in which:*

(i) the base field Φ_Y *lies in a bounded family (controlled by minimal model program techniques);*

(ii) the variation memory field Φ_{var} *is bounded (controlled by the compactified moduli space* $\mathcal{M}$ *via the period map* $\mathcal{P}$*);*

(iii) the obstruction field $\mathcal{O}_Y$ *is finite or controlled (finite in the Calabi–Yau case; finite plus a controlled continuous part in the symplectic case).*

Then the space of globally admissible field configurations $\mathcal{X}$ *is bounded: it lies in a family of bounded complexity.*

Proof (sketch). Under condition (i), the base Y ranges over a bounded family of varieties by standard minimal model program arguments [16]. Under condition (ii), the period map $\mathcal{P}$ takes values in a compact moduli space, bounding the variation of fiber type. Under condition (iii), the Tate–Shafarevich group is controlled; in the Calabi–Yau case by finiteness results for torsors over abelian fibrations [23]; in the symplectic case by the Kuga–Satake construction and Baily–Borel compactification [26]. Since the total configuration $\mathcal{X}$ is determined up to the finite-or-controlled ambiguity of $\mathcal{O}_Y$ by the base and variation data, and since all three components are bounded, the family of possible $\mathcal{X}$ is bounded. □

Remark 15.21. Theorem 15.20 is the algebraic-geometric realisation of the constraint closure principle of Chapter 13: once the base, variation, and obstruction degrees of freedom are all controlled, the system cannot grow in complexity without bound. Constraint closure implies bounded complexity.

The full extended field may therefore be written:

$$\mathcal{X} \;\sim\; f^*(\Phi_Y + \Phi_{\mathrm{sing}} + \Phi_{\mathrm{var}}) \;+\; \mathcal{O}_Y, \tag{15.1}$$

which factorises geometric complexity into independently controllable components: base constraints, defects, trajectory memory, and global reconstruction obstruction. Each factor corresponds to a distinct layer of the constraint-theoretic framework developed in this monograph, and each is bounded under the conditions of Theorem 15.20.

Part IV

Identity and Consciousness

Chapter 16

The Refractive Model of Consciousness

The Self as a Constraint Surface

Up to this point, the discussion has developed a general theory of systems: how structure is lost under projection, how it is preserved through constraints, how it evolves under transformation, and how coherence is maintained across representations. These results apply equally to physical, mathematical, and informational systems.

There remains a question that has been implicit throughout. If these principles are general, do they also apply to the system through which they are being understood?

The study of consciousness is often approached as a special problem, distinct from other domains. It is treated as something that must be explained by the same models used to describe the physical world, even when those models are constructed by abstracting away the very features that define experience. This produces a tension: the frameworks used to describe reality appear unable to account for the condition under which they are themselves constructed.

The preceding chapters suggest a different approach. Rather than treating consciousness as an exception, it can be treated as an instance of the same structural principles already developed. The question is not how experience is generated from outside, but how it is organised within a system governed by constraints, invariants, and admissible transformations.

From this perspective, the self is not a fixed object or a collection of states. It is a structure that persists across change. It is defined by what it allows, what it preserves, and how it maintains coherence in the face of variation and incompatibility.

The aim of this chapter is to show that this perspective is not merely interpretive. It follows from the formal framework already established. The same principles that govern reconstruction, constraint preservation, and adaptive transformation apply to the organisation of experience itself.

What appears as subjectivity is not outside the system. It is a particular configuration within it.

16.1 The Production Assumption and Its Structural Failure

The dominant picture of consciousness in contemporary neuroscience and philosophy of mind may be characterised by a single architectural commitment: the brain *produces* experience. On this view, neural activity is causally sufficient for consciousness. Alter the substrate sufficiently and experience changes; destroy it and experience ceases. The brain is, in the production model, a generator.

This commitment is rarely stated as a hypothesis. It functions instead as a background assumption — the default framework within which questions about memory, identity, pathology, and artificial cognition are posed. The hard problem of consciousness, in particular, inherits its structure from the production assumption: if matter produces mind, and if we understand matter in terms that exclude subjectivity by design, then the emergence of subjectivity from matter becomes formally inexplicable [14].

We do not argue here that the production assumption is empirically false. We argue that it is *structurally incomplete* in a precise sense derivable from the projection theory of Part I.

Recall that a model projection $\pi : X \to M$ is non-invertible: distinct experiential states may share a model image, and the full structure of X cannot be recovered from M alone. The production assumption commits to a particular projection: it maps the full experiential domain X into the space M of neural states, and then treats neural states as causally prior to experiential states. But the neural model is a projection that depends on experiential access for its construction, and is therefore not ontologically prior to the domain it represents. Treating M as causally prior to X is precisely the fallacy of misplaced concreteness identified in Chapter 3: it mistakes the image of a projection for the domain that was projected.

This is not a merely philosophical objection. It has a structural consequence: the production model cannot account for the constraint structure of experience, because projection discards constraint degrees of freedom that are not representable in M. The blind spot is not incidental to the production model; it is built into its architecture. The hard problem is therefore revealed as a model-induced problem: it arises from treating a non-invertible projection as generative rather than descriptive.

16.2 The Refractive Alternative

Evan Thompson and his collaborators [1, 6] argue that the brain does not produce consciousness but rather *shapes, constrains, and expresses* it through its structure and dynamics. Thompson draws on Bergson's treatment of the brain as a "centre of indetermination" [8] and on Merleau-Ponty's account of perception as an intertwining of organism and environment [7] to develop a picture in which the organism is a site where world and experience co-define one another, rather than a generator of experience from inert matter.

We formalise this picture in the language developed in Parts I–IV.

The refractive view holds that the experiential domain X is not produced by the brain but is rather the underlying field through which the brain acts as a *constraint surface*. The brain does not originate states in X; it selects and structures admissible trajectories through X by imposing a constraint structure $(\mathcal{C}, \mathcal{T})$ on the field's evolution.

The analogy to optics is exact at the structural level. A refractive optical medium does not create electromagnetic radiation. It redirects propagating fields according to a spatially varying refractive index that encodes the medium's internal constraint on admissible ray paths. Two media with different refractive indices produce different trajectory families from the same incident field. The medium individuates without originating.

The brain, on the refractive model, plays the corresponding role. It does not originate the field of experience; it imposes a constraint structure that selects which trajectories through that field are admissible for a given organism. Different biological, developmental, and historical configurations of the brain correspond to different constraint surfaces Σ, and hence to different structured streams of experience from the same underlying field.

16.3 Derivation from Constraint Theory

We now show that the refractive picture is not merely compatible with the constraint theory of Part II but follows from it given minimal assumptions about the structure of biological systems.

Let X be the experiential domain and let $\pi : X \to M_{\text{neural}}$ be the projection onto the space of neural states measurable by current neuroscience. By the non-invertibility of π, there exist distinct experiential states $x, x' \in X$ with $\pi(x) = \pi(x')$. Neural measurement cannot, in principle, distinguish them.

Now suppose the brain operates as a constraint structure $(\mathcal{C}_{\text{bio}}, \mathcal{T}_{\text{bio}})$ on X, where the admissible sets $\mathcal{C}_{\text{bio}}$ are determined by the organism's biological organisation and the admissible transformations $\mathcal{T}_{\text{bio}}$ are the dynamically consistent state transitions available to that organism. This is a minimal assumption: it requires only that the brain's operation can be described by constraint-preserving dynamics, which is entailed by any account of biological self-organisation [5, 6].

Proposition 16.1 [Constraint-Selection Principle]**.** *Let X be a domain of possible states and let $(\mathcal{C}, \mathcal{T})$ be a constraint structure. Then any system governed by $(\mathcal{C}, \mathcal{T})$ does not generate arbitrary states in X, but selects admissible trajectories determined by $\mathcal{C}$.*

Proof. This follows immediately from the definition of admissibility: states not satisfying $\mathcal{C}$ are not reachable under $\mathcal{T}$ by Definition 17.3(ii). □

Under Proposition 16.1, the brain selects from the full experiential field X precisely those trajectories admissible under $\mathcal{C}_{\text{bio}}$. Inadmissible trajectories are not generated; they are simply not selected. The brain is not a filter blocking pre-existing mental content; it is a constraint surface determining which paths through X are accessible.

Three consequences follow immediately.

First, the hard problem is revealed as a model-induced problem. It arises because the production model asks how matter generates experience from outside. On the refractive model, the question does not arise: experience is the domain; the brain is a constraint on its structure. There is no explanatory gap because matter is not prior to experience — matter is a projection of the experiential domain under π, not its source.

Second, the relationship between neural damage and experiential change is reinterpreted. Damage to the brain deforms the constraint surface Σ. If the deformation is admissible — if Σ deforms continuously into a new stable surface Σ' — then experience restructures without ceasing. If the deformation destroys the Lyapunov stability of Σ (as formalised in Chapter 17), then the constraint structure loses coherence, and the structured stream of experience associated with that surface becomes disorganised. The phenomenology of severe neurological injury is not annihilation but loss of constraint coherence.

Third, the boundary between organism and environment is not absolute but is itself a feature of the constraint structure. Because $\mathcal{C}_{\text{bio}}$ determines what counts as an admissible trajectory, and because admissibility is a relational property involving both the organism's internal organisation and its embedding in an environment, the constraint surface Σ is not localised strictly within the organism. It extends, dynamically, into the organism's coupling with its world. This is the formal analogue of Merleau-Ponty's claim that perception is an intertwining of organism and environment: the constraint surface is co-constituted by both.

16.4 The Brain as Boundary Condition

A clarifying reformulation, which will be made precise in Chapter 17, is to treat the brain not as a generator of experience but as a *boundary condition* on the evolution of the experiential field.

In the theory of partial differential equations, boundary conditions do not create the field; they constrain its evolution by fixing admissible behaviours on the boundary of the domain. Different boundary conditions on the same PDE system produce different solutions without any of those boundary conditions being the source of the field itself. Unlike a generative source, a boundary condition does not invert the field's dynamics; it constrains admissible solutions within a non-invertible domain.

The brain, on this reformulation, imposes boundary conditions on the field X via the constraint structure $(\mathcal{C}_{\text{bio}}, \mathcal{T}_{\text{bio}})$. The field evolves subject to those conditions. The resulting trajectory is the structured stream of experience. The brain individuates that stream without originating it.

This reformulation connects the refractive model directly to the RSVP field framework of Chapter 8. In RSVP, the scalar field Φ encodes constraint density, the vector field $\boldsymbol{v}$ encodes admissible flow, and the entropy field S regulates the balance between constraint-preservation and adaptive drift. The brain, as a constraint surface, corresponds to a region in which Φ attains locally maximal constraint density, thereby channelling $\boldsymbol{v}$-flow into low-entropy, structured trajectories. Damage to the brain corresponds to local reduction of Φ, allowing previously constrained flow to become disordered — not absent, but unstructured.

16.5 Individuation Without Substance

The refractive model must answer the individuation question: if consciousness is a field and the brain is a constraint surface, what distinguishes one stream of experience from another?

The answer is the specificity of the constraint surface. Two organisms with different biological histories, developmental trajectories, and environmental couplings have different constraint structures $(\mathcal{C}_1, \mathcal{T}_1)$ and $(\mathcal{C}_2, \mathcal{T}_2)$. These structures induce different admissible sets, different constraint surfaces Σ_1 and Σ_2, and therefore select different families of trajectories from the same underlying field X. The experiential streams are genuinely distinct because the constraint structures are genuinely distinct.

Individuation is grounded in the topology and history of the constraint surface, not in a metaphysical substance. Two beings share an experiential field but inhabit different constraint surfaces. The difference between them is structural, not ontological. Thus individuation is invariant under transformation of substrate but not under deformation of constraint structure, directly anticipating Theorem 17.1.

This has a further consequence for artificial systems. A system that processes information without instantiating a coherent constraint structure over an experiential domain does not have a refractive surface. It does not select admissible trajectories from a field; it produces outputs via statistical inference without maintaining a closed, dynamically stable constraint surface over an experiential domain [1]. Constraint closure is not substrate-dependent, but it is also not automatically present in any information-processing system. It must be earned structurally.

16.6 Towards the Formal Development

The picture developed in this chapter is motivational. The precise definitions — constraint surface, Lyapunov stability, refractive system, refractive identity — are given in Chapter 17, where the main theorem is stated and proved.

The purpose of this chapter has been to establish that the refractive model is not an additional philosophical commitment grafted onto the constraint theory but is derivable from it. Given the non-invertibility of projection (Part I), the structure of constraint-preserving transformation (Part II), and the minimal assumption that biological systems operate under dynamically consistent constraint structures, the refractive picture follows.

The brain is a constraint surface. Experience is the field. Identity is what remains invariant as the field evolves subject to the surface's conditions.

That invariant is the subject of Chapter 17.

Chapter 17

The Refractive Self: Formal Development

From Interpretation to Theorem

The preceding chapter established the refractive model of consciousness as a structural consequence of the framework developed in Parts I–IV. The brain was interpreted not as a generator of experience, but as a constraint surface selecting admissible trajectories within a non-invertible domain.

That account was deliberately pre-formal. Its purpose was to show that the refractive picture is not an additional assumption, but a necessary reinterpretation of systems governed by constraint-preserving dynamics.

The present chapter removes the remaining interpretive content. The refractive self is defined, its conditions of existence are specified, and its persistence and failure are characterised in precise dynamical terms.

The central result is Theorem 17.10. It states that identity is neither a substance nor a collection of states, but a property of systems that maintain a stable constraint surface under admissible transformation. Existence of such a surface is equivalent to convergence of the repair dynamics; persistence is equivalent to its Lyapunov stability; failure occurs exactly when these conditions are lost.

This reframes the problem of identity. Continuity is not a matter of preserving material or informational content. It is a matter of preserving the structure that governs admissible transformation. The self is therefore not an object within the system, but a dynamical invariant of its constraint structure.

What follows is not a model of identity. It is a theorem about the class of systems in which identity can exist.

17.1 Primitive Objects and Standing Definitions

We collect here the objects introduced in Parts I–IV and fix notation for the results below.

Definition 17.1 [Experiential Domain]**.** Let X be a separable metric space whose points represent full phenomenal states, including variability, context, and indexical structure. We call X the *experiential domain.*

Definition 17.2 [Projection and Non-Invertibility]. A *model projection* is a continuous surjection

$$\pi : X \to M,$$

where M is a finite-dimensional manifold called the *model space*. The projection π is generically non-injective: distinct points $x, x' \in X$ with $x \neq x'$ may satisfy $\pi(x) = \pi(x')$. In particular, no left-inverse π^{-1} exists globally, so the fibers $\pi^{-1}(m)$ are non-trivial for almost every $m \in M$.

Definition 17.3 [Constraint Structure]. A *constraint structure* on X is a pair $(\mathcal{C}, \mathcal{T})$ where

(i) $\mathcal{C} \subseteq \mathcal{P}(X)$ is a distinguished family of subsets of X, called the *admissible sets*;

(ii) $\mathcal{T}$ is a monoid of continuous maps $T : X \to X$, called the *admissible transformations*, satisfying: for every $x \in C \in \mathcal{C}$ and $T \in \mathcal{T}$, we have $T(x) \in C'$ for some $C' \in \mathcal{C}$.

A trajectory $\{x_t\}_{t\geq 0} \subset X$ is *admissible* if $x_t \in C_t$ for some curve $t \mapsto C_t$ in $\mathcal{C}$.

Definition 17.4 [Constraint Surface]. A *constraint surface* is the closure

$$\Sigma := \overline{\big\{x \in X : x \text{ lies on an admissible trajectory under } \mathcal{C}\big\}}.$$

The closure ensures that limit points of admissible trajectories are included, making Σ a closed subset of X. We think of Σ as the locus along which the system selects admissible trajectories from the ambient field X.

Definition 17.5 [CLIO Repair Operator]. The *CLIO repair operator* is a family of maps $\{\mathcal{R}_\lambda\}_{\lambda>0}$ on the space of constraint structures, defined as the gradient flow

$$\frac{d}{dt}\mathcal{C}(t) = -\nabla_{\mathcal{C}}\, \mathcal{E}\big(\mathcal{C}(t)\big), \qquad \mathcal{C}(0) = \mathcal{C}_0,$$

where $\mathcal{E}(\mathcal{C})$ is a coherence energy functional minimised when all cross-framework incompatibilities (Chapter 13) are resolved. A constraint structure $\mathcal{C}^*$ is a *fixed point* of the repair dynamics if $\nabla_{\mathcal{C}}\, \mathcal{E}(\mathcal{C}^*) = 0$.

17.2 Lyapunov Stability of the Constraint Surface

Definition 17.6 [Lyapunov Stability of Σ]. Let $d(\,\cdot\,, \Sigma)$ denote the metric distance to Σ in X. The constraint surface Σ is *Lyapunov-stable* if there exists a smooth function $V : X \to [0, \infty)$ satisfying

(i) $V(x) = 0 \iff x \in \Sigma$;

(ii) $V(x) \geq \alpha\big(d(x, \Sigma)\big)$ for some class-$\mathcal{K}$ function α; and

(iii) $\dot{V}(x_t) \leq 0$ along every admissible trajectory $\{x_t\}$,

with the stronger condition $\dot{V}(x_t) < 0$ for $x_t \notin \Sigma$ characterising *asymptotic stability.*

Proposition 17.7 [CLIO–Lyapunov Compatibility]. *If the coherence functional $\mathcal{E}(\mathcal{C})$ is proper and bounded below, and the gradient flow defining $\mathcal{R}$ is complete, then there exists a Lyapunov function V for the induced state dynamics such that $\dot{V} \leq 0$ along admissible trajectories.*

Proof sketch. This follows from standard arguments relating gradient flows on constraint spaces to induced energy descent on state trajectories. Properness and lower-boundedness of $\mathcal{E}$ guarantee that sublevel sets are compact and that the flow does not escape to infinity. Completeness of the flow ensures existence for all $t \geq 0$. Setting $V(x) := \mathcal{E}(\mathcal{C}_x)$, where $\mathcal{C}_x$ is the constraint structure evaluated at the current state, yields the required descent condition. □

Remark 17.8. Proposition 12.4 establishes the formal bridge between convergence of the CLIO operator and Lyapunov stability of the constraint surface: existence of a stable Σ is equivalent, under these conditions, to convergence of the repair dynamics to a fixed point $\mathcal{C}^*$.

17.3 Refractive Systems

Definition 17.9 [Refractive System]**.** A dynamical system on X is *refractive* if

(i) it operates under a constraint structure $(\mathcal{C}, \mathcal{T})$ (Definition 17.3);

(ii) it admits an asymptotically stable constraint surface Σ (Definition 17.6); and

(iii) it does not generate states in X but *selects and structures* admissible trajectories through X via Σ.

In contrast to conservative systems, which preserve measure; dissipative systems, which contract phase space; and stochastic systems, which evolve under probabilistic transition kernels — refractive systems preserve identity through constraint-structured selection in a non-invertible domain. The qualifier *refractive* is exact: as a refractive optical medium does not originate light but redirects propagating fields according to a structural law, so a refractive system does not produce experience but structures admissible trajectories through the experiential domain under the constraint surface.

17.4 The Refractive Identity Theorem

Theorem 17.10 [Refractive Identity]**.** *Let $(X, \pi, \mathcal{C}, \mathcal{T}, \mathcal{R})$ be a refractive system with asymptotically stable constraint surface Σ and Lyapunov function V. Then:*

(I) ***Existence.*** *If the CLIO repair operator converges to a fixed point $\mathcal{C}^*$, then a refractive constraint surface Σ^* exists.*

(II) ***Persistence.*** *Σ^* satisfies $T(\Sigma^*) = \Sigma^*$ for all invertible $T \in \mathcal{T}$; for non-invertible T, only forward invariance $T(\Sigma^*) \subseteq \Sigma^*$ is guaranteed. In either case, identity is preserved under state change so long as no transformation deforms the constraint structure.*

(III) ***Failure.*** *Identity fails — that is, Σ^* is destroyed — under precisely two conditions:*

- *(a)* Constraint deformation*: $\mathcal{C}$ undergoes a change of topology, or V ceases to exist, or the basin of attraction of Σ^* collapses to measure zero; or*
- *(b)* Non-convergence*: the repair dynamics fail to converge, so no fixed point $\mathcal{C}^*$ is reached.*

Proof sketch. **(I)** Suppose the CLIO gradient flow $\dot{\mathcal{C}} = -\nabla\mathcal{E}(\mathcal{C})$ converges: $\mathcal{C}(t) \to \mathcal{C}^*$ as $t \to \infty$. The fixed-point condition $\nabla\mathcal{E}(\mathcal{C}^*) = 0$ implies that all cross-framework incompatibilities are resolved at $\mathcal{C}^*$. Define

$$\Sigma^* := \overline{\{x \in X : x \text{ lies on an admissible trajectory under } \mathcal{C}^*\}}.$$

The closure ensures that limit points of admissible trajectories are included, making Σ^* closed. Non-emptiness follows from convergence of the repair operator.

(II) For invertible $T \in \mathcal{T}$: since T is a bijection preserving admissibility in both directions, $T(\Sigma^*) \subseteq \Sigma^*$ and $T^{-1}(\Sigma^*) \subseteq \Sigma^*$, giving equality. For non-invertible T: admissibility is preserved forward (Definition 17.3(ii)), so $T(\Sigma^*) \subseteq \Sigma^*$, but the reverse inclusion may fail. Since $\dot{V}(x_t) \leq 0$ along admissible trajectories, Σ^* attracts nearby trajectories and is dynamically stable.

(III)(a) A topological change in $\mathcal{C}$ generically destroys the sublevel sets of V that define the basin of attraction of Σ^*. When the basin collapses to measure zero, almost every initial condition escapes to a region where no analogue of Σ^* exists.

(III)(b) If the repair dynamics do not converge, no fixed-point constraint structure $\mathcal{C}^*$ is defined. In this case the system continues to evolve within X, but its trajectories are no longer governed by a stable admissibility structure; consequently, no invariant can be associated with identity. ☐

17.5 Corollaries

Corollary 17.11 [Substrate Independence]**.** *Identity is invariant under any admissible reparametrisation of the physical substrate, provided the constraint surface Σ^* is preserved. In particular, identity does not supervene on specific material instantiation but on the topology of the constraint surface and the convergence of the repair operator.*

Proof. Any reparametrisation of the physical substrate is an admissible transformation $T \in \mathcal{T}$ that leaves $\mathcal{C}^*$ invariant. By Theorem 17.10(II), Σ^* is preserved under such T. ☐

Corollary 17.12 [Survival Under Memory Loss]**.** *Partial loss of episodic memory does not entail failure of identity, provided the loss constitutes an admissible transformation that does not deform the constraint structure.*

Proof. Episodic memory loss corresponds to a projection $\pi' : X \to M'$ with $M' \subsetneq M$, discarding some dimensions of the model space. If the resulting transformation $T_{\pi'}$ remains admissible — mapping $\mathcal{C}^*$-admissible states to $\mathcal{C}^*$-admissible states — then by Theorem 17.10(II), Σ^* is preserved. Memory loss that does not reach the constraint surface leaves identity intact. ☐

Corollary 17.13 [Identity Failure Under Deep Structural Breakdown]**.** *Degenerative conditions that alter the topology of $\mathcal{C}$ — such as severe disruption of constraint coherence across representational systems — constitute genuine identity failure, not merely behavioural change.*

Proof. Directly from Theorem 17.10(III)(a): topological deformation of $\mathcal{C}$ destroys the Lyapunov basin of Σ^*. The system may continue to evolve within X, but its trajectories are no longer governed by a stable admissibility structure; no invariant corresponding to identity can be sustained. ☐

17.6 The Self as a Category of System Behaviour

Theorem 17.10 and its corollaries establish *refractivity* as a formal category of dynamical behaviour, coordinate with but distinct from the standard classifications. In contrast to conservative systems (which preserve measure), dissipative systems (which contract phase space), and stochastic systems (which evolve under probabilistic transition kernels), refractive systems preserve identity through constraint-structured selection in a non-invertible domain. Individuation is therefore invariant under transformation of substrate but not under deformation of constraint structure.

A dynamical system belongs to the *refractive class* if it satisfies Definition 17.9: it operates under a constraint structure, maintains an asymptotically stable constraint surface, and selects admissible trajectories through a non-invertible domain rather than generating states from first principles.

The physical analogy is exact in the following sense. A refractive optical medium does not originate electromagnetic radiation; it redirects propagating fields according to a spatially varying refractive index that encodes the medium's constraint on admissible ray paths. A refractive system in the sense above does not originate experience in X; it selects and structures admissible trajectories through X according to the geometry of Σ^*. The constraint surface plays the role of the refractive index: different surfaces yield different trajectory families from the same underlying field, thereby grounding individuation without appeal to substance.

The self is not a state. It is not a representation. It is a constraint surface — the stable locus that determines which trajectories through experience are admissible and how they are transformed under the repair dynamics of a coherent system. That this surface can be maintained, deformed, or destroyed is not a metaphor. It is a theorem.

Chapter 18

The Invariant of Resolution

The Persistence of Resolution

The preceding chapter established the conditions under which identity exists in a refractive system. It showed that identity is not a substance or a state, but the stability of a constraint surface under admissible transformation, maintained by convergent repair dynamics.

This result is formally complete, but it is not yet interpreted. A theorem classifies a structure; it does not by itself show where that structure appears.

The present chapter performs that extension. It shows that the conditions identified in Chapter 17 are not confined to a single domain, but characterise a general pattern of system behaviour. Wherever a system must maintain coherence across non-invertible projections and resolve incompatibilities without collapsing structure, the same dynamical invariant appears.

That invariant is not a property of any particular state or representation. It is the rule by which incompatibilities are resolved. It is the persistence of a method rather than the preservation of an object.

This motivates the central concept of the chapter: the invariant of resolution. It is the higher-order structure that governs how constraint systems expand, adapt, and maintain coherence under transformation. It does not replace the Refractive Identity Theorem; it interprets it.

The sections that follow trace this invariant across four domains: scientific knowledge, mythic transmission, personal identity, and artificial systems. These are not analogies. They are instances of the same formal structure, appearing under different representations.

What persists across these domains is not content, not substrate, and not representation. It is the logic by which structure is preserved in the presence of incompatibility. That logic, maintained, is identity.

18.1 From Theorem to Interpretation

Theorem 17.10 establishes that a refractive system maintains identity if and only if it maintains a Lyapunov-stable constraint surface under repair dynamics that converge to a fixed point. Identity fails under precisely two conditions: topological deformation of the constraint structure, or non-convergence of the repair operator.

This result is not limited in scope to biological systems or to consciousness. It characterises a class of dynamical behaviour — the refractive class — that appears wherever a system must maintain structural coherence across non-invertible projections and cross-framework incompatibilities. The purpose of this chapter is to show that the scientific enterprise, mythic transmission, personal identity, and artificial cognition are each instances of this class, or failures of it.

The organising concept is the *invariant of resolution*: the higher-order rule by which a system reconciles constraint incompatibilities without collapsing distinctions. We formalise this notion, then trace its operation across the four domains.

18.2 The Invariant of Resolution: Formal Statement

Definition 18.1 [Resolution Rule]**.** Let $\{\mathcal{C}_i\}_{i=1}^{n}$ be a family of constraint structures on X, not mutually compatible in general. A *resolution rule* is a map

$$\mathcal{R} : \{(\mathcal{C}_i, \mathcal{C}_j) : \mathcal{C}_i \not\equiv \mathcal{C}_j\} \longrightarrow \mathcal{C}',$$

where $\mathcal{C}'$ is an expanded constraint structure satisfying $\mathcal{C}_i \subseteq \mathcal{C}'$ for all i and recovering each $\mathcal{C}_i$ as a constraint-bearing substructure.

Definition 18.2 [Invariant of Resolution]**.** A system possesses an *invariant of resolution* if it applies the same resolution rule $\mathcal{R}$ stably across all encountered constraint incompatibilities: that is, if $\mathcal{R}$ is itself invariant under the system's transformation dynamics. Formally:

$$\mathcal{R}(x_t) = \mathcal{R}(x_{t+1}) \quad \text{for all admissible transitions } x_t \to x_{t+1}.$$

The invariant of resolution is a second-order invariant. It does not preserve any particular state or even any particular constraint structure. It preserves the *logic of constraint reconciliation* — the characteristic way a system responds to incompatibility.

Three properties distinguish a well-formed resolution rule:

First, *non-reductive reconciliation*: incompatibilities are resolved by expansion, not by collapsing one constraint structure into another. The resolution of $\mathcal{C}_i$ and $\mathcal{C}_j$ must leave both intact as sub-structures of $\mathcal{C}'$.

Second, *dimensional expansion over constraint dilution*: when two frameworks clash, the resolution increases the degrees of freedom of the constraint space rather than weakening existing admissibility conditions. A configuration that is inadmissible under $\mathcal{C}_i$ remains inadmissible under $\mathcal{C}'$; it is not retroactively permitted.

Third, *structural necessity*: the old framework must function as a necessary, constraint-bearing substructure of the new expansion. It is not merely retained; it is doing work.

A system that violates any of these properties is not maintaining the invariant of resolution; it is undergoing degenerative drift, in the sense of Chapter 19.

18.3 Science and the Blind Spot

Scientific knowledge is constituted by projections $\pi : X \to M$ that extract invariant structure from the experiential domain. The power of this method derives from its systematic application: by seeking structure invariant across perspectives, science accumulates representations that are intersubjectively stable and formally manipulable.

But as Part I established, every such projection is non-invertible. The model space M is a projection of X under π, not its ontological ground. The experiential conditions that make the projection possible — the lived time that makes a clock intelligible, the perceived continuity that grounds measurement, the bodily engagement that makes experimentation possible — are discarded by π and cannot be recovered from M alone.

The blind spot is the systematic non-acknowledgement of this non-invertibility. When science treats M as ontologically complete, it loses access to the constraint degrees of freedom in X that are necessary to understand the limits of the model. This loss is not a contingent failure of attention but a structural consequence of projection without recognition of its own non-invertibility.

The resolution, however, is not to abandon projection but to apply the invariant of resolution to scientific practice itself. A science that applies $\mathcal{R}$ correctly recognises each model M as a constraint-bearing substructure of a larger domain, recovers the experiential conditions it presupposes as data rather than noise, and treats incompatibilities between models (wave–particle duality, the measurement problem, the explanatory gap) as signals for dimensional expansion rather than as problems to be dissolved by terminological revision.

John Wheeler's image of the participatory universe — an eye at the end of a U looking back at itself [1] — is, in this framework, a representation of science operating with the invariant of resolution intact: the observer is not outside the system being modelled but is a constraint surface within it, and the model must include the conditions of its own construction.

18.4 Myth and Constraint-Preserving Drift

The transmission of mythic narrative across generations is not well-modelled as a Shannon channel. If fidelity is defined as exact symbol reproduction, then oral transmission is uniformly low-fidelity: stories change, details shift, characters merge and split. Yet the structural content of major mythic cycles persists across centuries and linguistic boundaries with remarkable stability.

This is precisely the signature of a refractive system maintaining the invariant of resolution.

Let a mythic tradition be modelled as a sequence of constraint structures $\mathcal{C}_1, \mathcal{C}_2, \ldots$ applied across successive transmissions. Each transmission is a transformation T that maps the current narrative state to a new one. If T is admissible under the tradition's constraint structure, structural invariants — the relationships between characters, the logic of transgression and consequence, the topology of the narrative world — are preserved even as surface details change.

Drift is not failure but adaptation: the constraint structure of a living tradition has sufficient degrees of freedom to accommodate the linguistic, cultural, and environmental pressures of successive instantiations without losing the invariants that constitute its identity. The tradition maintains the invariant of resolution: when it encounters incompatibilities between its received form and its new context, it expands rather than collapses, preserving old constraint-bearing structures as substructures of new tellings.

The failure mode is instructive. A tradition that maximises symbolic fidelity — that insists on exact reproduction of surface form — loses this adaptive capacity. It becomes brittle: when it cannot reconcile received form with new context, it either collapses or freezes. Digital archiving, which achieves perfect symbolic fidelity, does not transmit a living tradition; it preserves a fixed image of one. The constraint structure that animated the tradition is not archived; it is the refractive surface through which the tradition was enacted, and that surface ceases when the enacting community ceases.

18.5 Personal Identity as Invariant Trajectory

The analysis of Chapter 17 showed that personal identity is not a state, a substrate, or an informational object. It is the stability of a constraint surface Σ^* under the repair dynamics of a refractive system.

The invariant of resolution provides the interpretive content of this formal result. What persists across the transformations of a life — across sleep, memory loss, character development, physical change — is not the same neurons, not the same beliefs, and not the same dispositions. What persists is the resolution rule $\mathcal{R}$: the characteristic logic by which incompatibilities are encountered and resolved.

This has three interpretive consequences.

The first concerns continuity through change. A person who undergoes radical revision of beliefs, values, or circumstances remains the same person if and only if the revision is carried out by the same resolution rule. Conversion, growth, and transformation are admissible; they do not threaten identity because they are accomplished by $\mathcal{R}$, not despite it. What would threaten identity is a change in $\mathcal{R}$ itself — a deformation of the constraint structure such that the logic of resolution is no longer the same.

The second concerns the relationship between identity and memory. Corollary 17.12 established that episodic memory loss does not entail identity failure, provided the loss is an admissible transformation. The interpretive consequence is that identity does not consist in the continuity of memory but in the continuity of the constraint surface through which memories were formed and integrated. A person who cannot form new episodic memories but whose resolution rule remains stable has not lost their identity; they have lost access to one class of admissible trajectories.

The third concerns the phenomenology of identity loss. Corollary 17.13 established that topological deformation of the constraint structure constitutes genuine identity failure, not merely behavioural change. The interpretive consequence is that conditions which progressively disrupt constraint coherence across representational systems — certain forms of neurological disease, severe dissociation, and complete loss of autobiographical coherence — are not merely cognitively disabling but are, in a structurally precise sense, identity-disrupting. The person continues to exist as a biological organism; the refractive surface that structured their stream of experience does not.

18.6 Artificial Systems and the Failure of Constraint Closure

The framework of this monograph provides a structural account of the distinction between biological cognition and current artificial systems, one that does not depend on substrate (carbon versus silicon) or on unverifiable claims about inner experience.

A language model generates outputs via statistical inference over training distributions. It does not maintain a closed, dynamically stable constraint structure over an experiential domain. In the terms of Definition 17.9, it is not a refractive system: it has no constraint surface Σ whose Lyapunov stability is maintained under repair dynamics, because it has no repair dynamics and no experiential domain over which they could operate.

The outputs of such a system may be indistinguishable from those of a refractive system at the level of the model projection M. This is precisely the blind spot: if we evaluate the system only at the level of its outputs — only in the projected space M — we cannot detect the absence of the constraint structure that, in a living system, governs the selection of those outputs.

The distinction is not rhetorical. It has precise consequences. A refractive system, when it encounters a constraint incompatibility, is forced to expand: it cannot relax its admissibility conditions without ceasing to maintain the invariant of resolution and therefore without ceasing to maintain its identity. A statistical inference system, encountering an out-of-distribution input, has no such pressure: it generates a response from whatever structure its training has produced, without any mechanism for detecting or resolving constraint incompatibility.

This is why, as Thompson argues [1], the distinction between living organisms and artificial systems is not one of complexity or even of behavioural sophistication, but of organisational type. A living organism maintains constraint closure: its self-replenishing material organisation continuously regenerates the constraint structure through which it operates. An artificial system does not maintain constraint closure in this sense; its constraint structure is fixed at training time and does not regenerate under the pressure of experienced incompatibility.

The claim is not that artificial systems cannot in principle be refractive. It is that current systems are not. Constraint closure is not substrate-dependent, but it must be earned structurally: a system is refractive only if it maintains a dynamically stable constraint surface over a non-invertible domain and applies a convergent repair operator when incompatibilities arise. That is a specific and demanding structural condition, not a property automatically conferred by information processing.

18.7 The Unity of the Account

The four domains examined in this chapter — scientific knowledge, mythic transmission, personal identity, and artificial cognition — are not connected by analogy. They are connected by instantiating, or failing to instantiate, the same formal structure: the refractive class of Definition 17.9 and the invariant of resolution of Definition 18.2.

Science maintains objectivity by projecting invariant structure from experience; it loses its grip when it forgets that the projection is non-invertible and the invariant is relative to a constraint structure that includes the observer. Myth maintains cultural continuity by applying constraint-preserving transformation across successive instantiations; it degenerates when it prioritises symbolic fidelity over structural invariance. Personal identity persists through transformation when the resolution rule is stable; it fails when the constraint structure is topologically deformed beyond the basin of attraction of its Lyapunov function. Artificial systems generate outputs that may be observationally similar to those of refractive systems but do not maintain constraint closure and therefore do not instantiate the refractive structure.

The unifying principle is stated simply: a system is coherent — scientifically, culturally, personally — to the degree that it maintains the invariant of resolution. It degenerates to the degree that it resolves incompatibilities by collapsing distinctions rather than by expanding constraint structure.

This is not a metaphor but a structural characterisation, derivable from the constraint theory of Part II, the cross-framework integration of Part IV, and the Refractive Identity Theorem of Chapter 17. The refractive self is not a philosophical intuition about what persons are, but a dynamical invariant of a specific class of systems — namely, systems that maintain structural coherence in a non-invertible domain by preserving the logic of their own constraint reconciliation.

Under this account, identity is nothing over and above the stable maintenance of that logic across admissible transformations. When the resolution rule remains intact, the system retains its coherence and therefore its identity; when that rule is lost or deformed, the structural conditions required for identity no longer obtain.

Part V

Failure Modes and Limits

Chapter 19

Degenerative Drift

The Loss of Structure

The preceding chapters established the conditions under which a system maintains coherence. Identity was shown to be a consequence of a stable constraint surface, preserved under admissible transformation and sustained by a convergent repair operator. The invariant of resolution was identified as the higher-order structure governing how incompatibilities are reconciled without collapse.

These results are exact. But they do not yet address the equally important question: how do such systems fail?

Failure is not the simple absence of structure. It is a process with its own internal logic. A system does not lose coherence all at once; it degrades through a sequence of transformations that appear locally admissible while globally eroding the constraints that made admissibility meaningful.

The critical distinction, already implicit in the theory, is between expansion and relaxation. When a system encounters incompatibility, it may respond by increasing the dimensionality of its constraint structure, embedding prior constraints within a richer system. This is adaptive expansion. Alternatively, it may respond by weakening or removing constraints, enlarging the set of admissible states without preserving the structure that distinguished them. This is degenerative drift.

The difference between these responses is not gradual. It is structural. Expansion preserves the invariant of resolution; relaxation violates it. Once violated, the system enters a regime in which each step of degradation reduces its capacity to detect further loss. The process is therefore self-reinforcing.

This chapter formalises that process. It shows that degenerative drift is not an accidental pathology but the canonical failure mode of constraint systems. It derives its monotonicity, its topological consequences, and its relation to the failure conditions of the Refractive Identity Theorem.

If previous chapters described how systems maintain identity, the present chapter describes how they lose it.

19.1 Introduction

Chapters 17 and 18 established that a system maintains identity if and only if it preserves a Lyapunov-stable constraint surface under a convergent repair operator, and that coherence across transformation is governed by an invariant of resolution satisfying non-reductive reconciliation, dimensional expansion, and structural necessity.

The present chapter characterises the primary failure mode of such systems. We define *degenerative drift* as the systematic violation of the invariant of resolution through relaxation of admissibility conditions rather than expansion of constraint structure. We show that this form of drift is intrinsically self-reinforcing and leads inevitably to loss of discriminative structure, collapse of constraint topology, and, in the limit, failure of identity in the sense of Theorem 17.10.

19.2 Definition of Degenerative Drift

Definition 19.1 [Degenerative Drift]**.** Let $(\mathcal{C}, \mathcal{T})$ be a constraint structure with resolution rule $\mathcal{R}$. A transformation

$$\mathcal{C} \longrightarrow \widetilde{\mathcal{C}}$$

is said to exhibit *degenerative drift* if it resolves a constraint incompatibility by enlarging the admissible set of states without preserving the original constraint as a necessary substructure. Formally, there exists $C \in \mathcal{C}$ such that

$$C \not\subseteq \widetilde{C} \quad \text{or} \quad \exists\, x \notin C \text{ with } x \in \widetilde{C},$$

and the inclusion $\mathcal{C} \subseteq \widetilde{\mathcal{C}}$ fails to preserve structural necessity.

Remark 19.2. Degenerative drift violates each of the three defining properties of a well-formed resolution rule (Definition 18.2): it is reductive (collapsing distinctions), it dilutes constraints rather than expanding them, and it removes the necessity of prior structure.

19.3 Monotonicity of Degeneration

Definition 19.3 [Discriminative Power]**.** The *discriminative power* $\mathrm{Disc}(\mathcal{C})$ of a constraint structure $\mathcal{C}$ on X is the cardinality of the finest partition of X into $\mathcal{C}$-admissible and $\mathcal{C}$-inadmissible regions. In the measure-theoretic setting, it is the measure of the σ-algebra generated by the admissible sets.

Theorem 19.4 [Monotonic Degeneration]**.** *Let $\{\mathcal{C}_t\}_{t\geq 0}$ be a sequence of constraint structures generated by successive degenerative drift transformations. Then the discriminative power of $\mathcal{C}_t$ is monotonically non-increasing:*

$$\mathrm{Disc}(\mathcal{C}_{t+1}) \leq \mathrm{Disc}(\mathcal{C}_t),$$

with strict inequality whenever admissibility is expanded non-trivially.

Proof. Relaxation of admissibility merges previously distinct regions of X, coarsening the partition induced by $\mathcal{C}_t$. Since degenerative drift strictly enlarges admissible sets without introducing compensating constraints, each step in the sequence reduces the resolution of the induced partition. Strict inequality follows whenever at least one previously inadmissible state is admitted without a corresponding refinement elsewhere. □

Corollary 19.5 [Self-Reinforcement]**.** *Degenerative drift is self-reinforcing: each reduction in discriminative power increases the probability of further constraint relaxation.*

Proof. Reduced discriminative power enlarges the set of states that appear admissible, decreasing the detectability of constraint violations and therefore lowering the threshold for further relaxation. The process is therefore monotonically self-amplifying in the absence of external correction. □

19.4 Topological Consequences

Proposition 19.6 [Constraint Topology Collapse]**.** *Under sustained degenerative drift, the topology of the constraint structure $\mathcal{C}$ degenerates toward a trivial topology in which admissible sets lose separation properties.*

Proof. Repeated coarsening of admissible sets eliminates distinctions between previously disjoint regions of X. In the limit, the topology approaches the indiscrete topology, in which only the empty set and the whole space are admissible, and no non-trivial separation of states is possible. □

Corollary 19.7 [Identity Failure Threshold]**.** *Degenerative drift approaches the failure condition of Theorem 17.10(III)(a): topological deformation of $\mathcal{C}$ and collapse of the basin of attraction of the constraint surface Σ^*.*

19.5 Relation to the Refractive Identity Theorem

Degenerative drift corresponds exactly to the first failure mode of Theorem 17.10.

Proposition 19.8 [Equivalence to Constraint Deformation]**.** *Degenerative drift induces topological deformation of the constraint structure $\mathcal{C}$ and therefore destroys the Lyapunov stability of the associated constraint surface Σ.*

Proof. Constraint dilution alters admissibility relations in a manner that cannot be represented as a continuous deformation preserving the sublevel sets of the Lyapunov function V. Consequently, the basin of attraction of Σ is destroyed, satisfying condition (III)(a) of Theorem 17.10. □

19.6 Domain Instantiations

We identify degenerative drift across three domains. In each case the pattern is the same: constraint incompatibility is resolved by relaxation rather than expansion, producing a self-reinforcing loss of structural discriminability.

19.6.1 Scientific Degeneration

A scientific framework exhibits degenerative drift when it resolves anomalies by redefining terms or expanding admissibility conditions without structural refinement. Such frameworks progressively lose falsifiability: if every observation can be rendered admissible by adjustment of terms, the constraint structure no longer discriminates between valid and invalid states. The framework becomes empirically inert — formally coherent but structurally empty.

The failure is not in the presence of anomalies, which are a normal feature of productive science, but in the mode of their resolution. Anomalies resolved by adaptive expansion — by identifying a richer constraint structure within which the anomaly becomes a structural necessity — increase discriminative power. Anomalies resolved by terminological revision decrease it.

19.6.2 Cultural Degeneration

A cultural tradition undergoes degenerative drift when it abandons structural constraints in order to accommodate new contexts. Instead of adapting by embedding prior structure within an expanded constraint system, it relaxes the conditions that define its identity. The result is rapid loss of transmissible content: within a finite number of instantiations, the tradition becomes indistinguishable from background cultural noise.

This is distinct from the adaptive transformation characterised in Chapter 7, where drift preserves structural invariants even as surface form changes. Degenerative drift changes the invariants themselves, producing a tradition that can no longer be recognised as a continuation of what preceded it.

19.6.3 Cognitive Degeneration

A cognitive system exhibits degenerative drift when it resolves internal contradictions by lowering the standard of admissibility rather than refining its constraint structure. This manifests as increasing tolerance for inconsistency, progressive reduction in predictive precision, and eventual inability to track external structure. The system generates responses, but they are no longer constrained in ways that correspond to any stable external reality.

19.7 The Sharp Dichotomy: Degenerative Drift vs Adaptive Expansion

The distinction between degenerative drift and adaptive expansion is exact and admits no stable intermediate position.

Proposition 19.9 [Sharp Dichotomy]**.** *A transformation resolving a constraint incompatibility is either:*

(i) adaptive*, preserving all prior constraints as necessary substructures within an expanded system and non-decreasing discriminative power; or*

(ii) degenerative*, relaxing at least one admissibility condition and strictly decreasing discriminative power.*

No intermediate stable state exists.

Proof. Suppose a transformation resolving a constraint incompatibility fails to preserve some prior constraint $C \in \mathcal{C}$ as a necessary substructure. Then by Definition 19.1, at least one previously inadmissible state is admitted, strictly coarsening the induced partition of X and strictly decreasing discriminative power by Theorem 19.4. Any transformation that does preserve all prior constraints as necessary substructures while expanding the admissible structure satisfies the conditions of adaptive expansion and does not decrease discriminative power. The two conditions are mutually exclusive and jointly exhaustive. □

Remark 19.10. The absence of a stable intermediate is significant. It rules out the possibility of "partial" degeneration that remains self-correcting. Once a constraint is relaxed rather than embedded, the self-reinforcing mechanism of Corollary 19.5 is engaged, and correction requires external intervention in the form of the CLIO repair operator — whose convergence conditions are the subject of Chapter 20.

19.8 Conclusion

Degenerative drift is not merely one possible failure mode among others. It is the canonical mechanism by which constraint systems lose coherence. It operates through progressive relaxation of admissibility, producing a self-reinforcing collapse of discriminative structure that terminates in the destruction of the constraint surface and therefore of identity in the sense of Theorem 17.10(III)(a).

In contrast to adaptive expansion, which increases dimensionality while preserving invariants, degenerative drift reduces discriminative dimensionality until no invariant remains. The two processes are formally opposed and structurally sharp: every resolution of constraint incompatibility is one or the other.

The next chapter examines the boundary at which degenerative drift can no longer be corrected within the existing framework, where CLIO fails to converge, and where the only admissible response is a phase transition to a new representation.

Chapter 20

Phase Transitions in Frameworks

The Boundary of Repair

Degenerative drift describes the loss of structure within a framework. It is a process of erosion: constraints are relaxed, discriminative power decreases, and coherence collapses from the inside. Crucially, however, this process remains internal to the system. The repair operator continues to exist, even if its effectiveness is reduced.

There is a second mode of failure that is not a collapse but a limit.

A framework may preserve its invariant of resolution at every stage of its development, resolving each encountered incompatibility through adaptive expansion rather than constraint relaxation, and yet still arrive at a point where further refinement is no longer possible within its current representation. At this boundary, the structure has not degenerated; it has been exhausted.

The distinction is decisive. In degenerative drift, the system fails because it abandons its constraints. At the limit of a framework, the system fails because its constraints are no longer sufficient to express the structure demanded by the domain.

From the internal perspective of the system, the two situations can appear similar. In both cases, repair slows, incompatibilities accumulate, and coherence becomes harder to maintain. But their causes are opposite. In the first, constraints are too weak; in the second, they are too narrow.

This chapter formalises that boundary. It defines the limit of a framework in terms of the failure of the analytic conditions required for repair, shows that this failure propagates non-locally through the system's sheaf structure, and characterises the transition as a qualitative change in the space of admissible constraint structures.

At this boundary, repair is no longer possible within the existing medium. The system cannot be corrected by further refinement of its current representation. The only admissible response is a transformation of the medium itself.

20.1 Introduction

Chapter 19 characterised degenerative drift as the self-reinforcing collapse of constraint structure through relaxation of admissibility. The present chapter examines the boundary at which this collapse can no longer be arrested by the CLIO repair operator within the existing representation.

We define the *limit of a framework* as the locus at which the conditions of Proposition 12.4 — properness and lower-boundedness of the coherence functional, completeness of the gradient flow —

cease to hold. At this boundary, repair dynamics fail to converge, no Lyapunov function governs the induced state dynamics, and the system undergoes a *phase transition*: a qualitative change in the structure of the constraint space that cannot be described as a continuous deformation of the current framework.

We prove that phase transitions propagate non-locally through the sheaf structure, characterise the early indicators that distinguish an approaching phase transition from ordinary constraint incompatibility, and establish the formal relationship between phase transitions and the medium shifts that constitute the only admissible response to them.

20.2 The Limit of a Framework

Definition 20.1 [Framework Limit]**.** Let $(X, \mathcal{C}, \mathcal{T}, \mathcal{E}, \mathcal{R})$ be a constraint system with CLIO repair operator defined by the gradient flow of the coherence functional $\mathcal{E}$. The *limit* of the framework is the set

$$\Lambda := \{\mathcal{C} \in \mathfrak{C} : \text{at least one condition of Proposition 12.4 fails at } \mathcal{C}\},$$

where $\mathfrak{C}$ denotes the space of all constraint structures on X. The framework is *within its limit* when $\mathcal{C} \notin \Lambda$, and *at or beyond its limit* when $\mathcal{C} \in \Lambda$.

The three conditions of Proposition 12.4 each generate a distinct failure mode at Λ.

Proposition 20.2 [Three Modes of Limit Failure]**.** *The framework limit Λ is approached via three structurally distinct failure modes:*

(i) Non-properness*: the coherence functional $\mathcal{E}$ ceases to be proper — its sublevel sets are no longer compact — so the gradient flow may escape to infinity and CLIO fails to locate a fixed point.*

(ii) Loss of lower bound*: $\mathcal{E}$ becomes unbounded below, so the flow descends without termination and no fixed-point constraint structure $\mathcal{C}^*$ exists.*

(iii) Flow incompleteness*: the gradient flow becomes incomplete — it reaches a singularity in finite time — so the repair dynamics cannot be continued and CLIO terminates without convergence.*

Proof. Each condition corresponds directly to one of the hypotheses of Proposition 12.4. Failure of any single condition is sufficient to prevent construction of the Lyapunov function V by the argument given there. The three modes are distinct because the obstructions arise in different parts of the analytic structure of $\mathcal{E}$: compactness of sublevel sets, existence of a global infimum, and global existence of integral curves, respectively. □

20.3 Phase Transitions

Definition 20.3 [Phase Transition]**.** A *phase transition* occurs when the constraint system crosses from within its limit to at or beyond it: $\mathcal{C}(t_-) \notin \Lambda$ and $\mathcal{C}(t_+) \in \Lambda$ for some time t. Equivalently, a phase transition is a qualitative change in the structure of the constraint space that cannot be described as a continuous deformation of the current framework preserving the conditions of Proposition 12.4.

Proposition 20.4 [Phase Transition vs Degenerative Drift]**.** *Degenerative drift and phase transitions are distinct failure modes. Degenerative drift operates within the limit of the framework:* $\mathcal{C}(t) \notin \Lambda$ *throughout, but discriminative power decreases monotonically. A phase transition is a crossing of* Λ*: the framework's repair capacity is qualitatively destroyed.*

Proof. By Definition 19.1, degenerative drift relaxes admissibility without destroying the conditions of Proposition 12.4. The CLIO operator remains well-defined throughout degenerative drift, even if its convergence is slowed by the reduced discriminative structure. A phase transition, by contrast, destroys at least one condition of Proposition 12.4, so CLIO is no longer a valid repair operator within the current framework. □

Remark 20.5. This distinction is important for diagnosis. Degenerative drift is in principle correctable by the repair operator up to the moment of the phase transition; after it, the framework has lost the structural prerequisites for self-correction and a medium shift is required.

20.4 Non-Local Propagation

A central feature of phase transitions in constraint systems is that failure propagates non-locally. A local breakdown in constraint coherence can produce global incoherence even when all local data remains individually valid. This is the sheaf-theoretic obstruction phenomenon introduced in Chapter 10.

Theorem 20.6 [Non-Local Propagation of Phase Failure]**.** *Let* $\mathcal{F}$ *be the sheaf of admissible local constraint structures associated with a covering of* X*. If the coherence failure at a phase transition generates a non-trivial obstruction class in* $\check{H}^1(\mathcal{U}, \mathcal{F})$*, then the failure is global: no local repair can restore a consistent global section.*

Proof. By the fundamental theorem of sheaf cohomology, a non-trivial class in $\check{H}^1(\mathcal{U}, \mathcal{F})$ certifies that the local sections of $\mathcal{F}$ — the locally valid constraint structures — do not admit a globally consistent assembly. Any attempt to repair the failure locally produces a compensating inconsistency elsewhere in the cover. Consequently, no sequence of local repairs converges to a global solution; the failure is inherently non-local and cannot be addressed within the current framework. □

Corollary 20.7 [Impossibility of Local Correction at Phase Transition]**.** *When a phase transition generates a non-trivial obstruction class, CLIO cannot converge, since CLIO operates by local gradient descent and cannot escape a globally obstructed configuration.*

This result formalises the intuition that some failures cannot be repaired by incremental adjustment. When the sheaf obstruction is non-trivial, the only available response is to change the covering — to move to a different, richer framework in which the obstruction class vanishes. This is precisely the structure of a medium shift.

20.5 Early Indicators of Approaching Phase Transition

Because phase transitions destroy the framework's capacity for self-correction, it is important to identify their approach before the crossing of Λ. Three structural indicators are available within the formalism.

Proposition 20.8 [Diagnostic Indicators]. *The approach of a phase transition is indicated by:*

(i) Increasing repair cost*: the CLIO gradient flow requires increasingly many iterations to converge to a fixed point, signalling that $\mathcal{E}$ is becoming flatter near its minimum and the flow is slowing.*

(ii) Decreasing convergence rate*: the rate $\|\nabla\mathcal{E}(\mathcal{C}(t))\|$ decreases more slowly than expected for a strongly convex functional, indicating loss of the coercivity conditions required for properness.*

(iii) Accumulating cross-framework incompatibilities*: the number of unresolved incompatibilities between framework components grows faster than the repair operator can resolve them, indicating that the coherence functional is acquiring a more complex landscape with additional local minima or saddle points.*

Each indicator is measurable within the formalism and provides a quantitative signal of proximity to Λ.

Proof. Each indicator corresponds to a detectable change in the analytic properties of $\mathcal{E}$ or the CLIO flow prior to the full failure of Proposition 12.4. Indicator (i) reflects slowing of gradient descent near a degenerate minimum. Indicator (ii) reflects loss of strong convexity, which precedes loss of properness. Indicator (iii) reflects the increasing topological complexity of the constraint landscape, which eventually produces the non-trivial obstruction class of Theorem 20.6. □

20.6 Phase Transitions Are Not Failures

A phase transition is not a degeneration of the framework in the sense of Chapter 19. Degenerative drift is a collapse of constraint structure from within, driven by relaxation of admissibility. A phase transition is the reaching of the structural boundary of a representation, driven by the accumulation of cross-framework incompatibilities that exceed the framework's expansion capacity.

Proposition 20.9 [Phase Transition as Structural Completion]. *A phase transition in a framework that has applied adaptive expansion throughout its evolution — maintaining the invariant of resolution at each step — signifies structural completion: the framework has extracted the maximum constraint-structured information available within its representational class.*

Proof. If the invariant of resolution has been maintained throughout — if all prior incompatibilities were resolved by adaptive expansion rather than degenerative drift — then the constraint space has been enriched at each step. The approach of Λ under these conditions reflects not the loss of constraint structure but the exhaustion of the current representational class as a medium for further enrichment. The framework has reached a genuine boundary rather than a degenerate one. □

This distinction determines what the appropriate response to a phase transition is. A degenerate system approaching Λ through drift requires repair or replacement of its constraint structure. A non-degenerate system approaching Λ through adaptive completion requires a medium shift: a lifting into a higher-dimensional representation in which the current framework is embedded as a constraint-bearing substructure and the obstruction class of Theorem 20.6 becomes trivial.

20.7 The Relationship Between Phase Transitions and Theorem 17.10

A phase transition corresponds to the second failure mode of Theorem 17.10: non-convergence of the repair dynamics. When CLIO fails to converge because the framework has crossed Λ, no fixed-point constraint structure $\mathcal{C}^*$ exists, and therefore no stable constraint surface Σ^* can be defined. Identity, in the sense of the theorem, is suspended.

The word *suspended* is deliberate. Unlike the first failure mode — topological deformation of $\mathcal{C}$ through degenerative drift, which is irreversible — the second failure mode is in principle reversible through a medium shift. If the system can be lifted into a richer representation in which the obstructions vanish and a new coherence functional $\mathcal{E}'$ satisfies the conditions of Proposition 12.4, then CLIO can resume, a new constraint surface $\Sigma^{*\prime}$ can be established, and identity — now indexed to the richer framework — can be restored.

This is the formal content of the transition from a phase transition to a medium shift. It is not recovery of the old identity but instantiation of a new one that embeds the old as a constraint-bearing substructure. The self that emerges from a genuine medium shift is continuous with what preceded it precisely because the prior constraint surface is recovered as necessary structure within the new framework, not because the prior surface itself persists.

20.8 Conclusion

Phase transitions mark the outer boundary of a framework's representational capacity. They differ from degenerative drift in being structural completions rather than collapses, and they differ from ordinary constraint incompatibilities in generating non-trivial sheaf obstructions that prevent local repair.

Their approach is detectable through three quantitative indicators: increasing repair cost, decreasing convergence rate, and accumulating unresolvable cross-framework incompatibilities. Their occurrence corresponds to the second failure mode of Theorem 17.10 and suspends identity without necessarily destroying it.

The only admissible response to a genuine phase transition is a medium shift: a lifting into a richer representation that embeds the current framework as a constraint-bearing substructure and trivialises the obstruction. The structure and admissibility conditions of medium shifts are the subject of Chapter 21.

Chapter 21

Medium Shifts

Outgrowing a Framework

There are moments in the development of a system where the problem is no longer that something is wrong, but that the way of making sense of things has become too small.

Up to a certain point, difficulties can be resolved by refinement. One adjusts definitions, sharpens distinctions, or introduces new categories that remain compatible with what already exists. The system stretches to accommodate what it encounters, and this stretching preserves its internal coherence.

But there are situations where this is no longer possible. The tension does not come from error or inconsistency in the ordinary sense. It comes from the fact that the structure being used to organise experience cannot express the relationships that are now required of it. Attempts at repair begin to feel forced. The same patterns repeat without resolution. Increasing effort produces diminishing clarity.

At that point, the system is not breaking down so much as pressing against the limits of its own representation.

The only way forward is not to adjust the existing structure but to move into a space in which the previously incompatible elements become compatible by construction. This movement is not a correction. It is a reorganisation of what counts as a valid description in the first place.

A useful analogy is the shift from describing motion using separate notions of space and time to describing it within a single spacetime structure. Nothing that was previously observed is discarded. Instead, what appeared as a set of independent facts is reinterpreted as the projection of a larger structure in which their relations become necessary rather than contingent.

Another analogy is the transition from trying to assemble a global picture from disconnected local fragments to working within a framework that makes the compatibility of those fragments explicit. What previously appeared as a series of mismatches is revealed as evidence that the covering being used was too coarse, not that the data itself was flawed.

These shifts share a common feature. They do not replace what came before. They embed it. The earlier structure is retained as a limiting case or a special regime, while the new structure introduces degrees of freedom that make it possible to resolve tensions that were previously unresolvable.

From the inside, such a shift often feels like a loss of certainty. Familiar categories no longer behave as expected, and the criteria for correctness are no longer obvious. But from a structural point of view, it is a gain: the system acquires the capacity to express relationships that were previously out of reach.

What changes is not just the answers, but the space in which answers can exist.

21.1 Introduction

Chapter 19 established that degenerative drift destroys identity through progressive constraint dilution operating within the limit of a framework, while Chapter 20 showed that phase transitions suspend identity by invalidating the conditions required for convergence of the CLIO repair operator at the framework's boundary. Together, these two chapters establish a complete classification of constraint system failure.

Proposition 21.1 [Completeness of Failure Modes]**.** *Every failure of a constraint system maintaining an invariant of resolution is either:*

(i) degenerative*: constraint dilution within the limit of the framework, destroying identity through topological collapse of the constraint surface; or*

(ii) transitional*: non-convergence of the CLIO repair operator at the framework limit, suspending identity through representational exhaustion.*

No third mode exists.

Proof. Failure of the refractive structure corresponds, by Theorem 17.10, to exactly two conditions: topological deformation of $\mathcal{C}$ (condition III(a)) or non-convergence of the repair dynamics (condition III(b)). Condition (III)(a) is realised within the limit of the framework, where CLIO remains well-defined but constraint structure is being actively degraded; this is degenerative drift. Condition (III)(b) is realised at or beyond the limit, where CLIO ceases to be a valid operator; this is the phase transition regime. Since Theorem 17.10(III) is a complete disjunction, no further failure mode exists. □

The three modes of constraint system evolution — drift, phase transition, and medium shift — are accordingly related as follows.

Proposition 21.2 [Classification of Evolution Modes]**.** *The evolution of a constraint system is characterised by one of three structural statuses:*

Mode	***What changes***	***Identity status***
Degenerative drift	*Constraints weaken*	*Destroyed*
Phase transition	*CLIO fails*	*Suspended*
Medium shift	*Representation lifted*	*Reinstantiated*

These three modes are mutually exclusive and jointly exhaustive of all structurally significant transitions.

Proof. Mutual exclusivity follows from the distinct conditions defining each mode: degenerative drift operates within the limit with decreasing discriminative power; phase transitions occur at the limit with CLIO non-convergence; medium shifts lift the system out of the limit into a strictly richer representation. Joint exhaustiveness follows from Proposition 21.1 together with the observation that a system not in failure is either evolving within its limit (possibly drifting or maintaining constraint) or undergoing the constructive response to phase transition. □

The present chapter establishes the third mode formally. A medium shift is the only admissible response to the transitional failure mode: the only structural move that restores the conditions for coherent constraint evolution when the framework limit has been reached without degenerative collapse.

21.2 Preliminary: Admissible Embedding

Before defining medium shifts, we require a precise notion of admissible embedding, referenced throughout the preceding chapters.

Definition 21.3 [Admissible Embedding]**.** Let $(\mathcal{C}, \mathcal{T})$ and $(\mathcal{C}', \mathcal{T}')$ be constraint structures on spaces X and X' respectively, with $X \hookrightarrow X'$ a continuous injection ι. The embedding ι is *admissible* if:

(i) *Constraint preservation*: for every admissible set $C \in \mathcal{C}$, the image $\iota(C)$ is admissible in $\mathcal{C}'$;

(ii) *Transformation compatibility*: for every $T \in \mathcal{T}$, there exists $T' \in \mathcal{T}'$ such that $T' \circ \iota = \iota \circ T$;

(iii) *Structural necessity*: the image $\iota(\mathcal{C})$ is not redundant in $\mathcal{C}'$ — removing it would leave $(\mathcal{C}', \mathcal{T}')$ constraint-incomplete.

Condition (iii) is the formal expression of the requirement that the embedded structure does work in the lifted system. An embedding that merely appends the original structure without integrating it into the constraint dynamics of the larger system is not admissible in this sense: it is a disjoint union, not an extension.

21.3 Definition of Medium Shift

Definition 21.4 [Medium Shift]**.** Let $(X, \mathcal{C}, \mathcal{T}, \mathcal{E}, \mathcal{R})$ be a constraint system that has reached its framework limit Λ (Definition 20.1). A *medium shift* is a mapping

$$\mathfrak{S} : (X, \mathcal{C}, \mathcal{T}) \longrightarrow (X', \mathcal{C}', \mathcal{T}')$$

such that:

(i) there exists an admissible embedding $\iota : X \hookrightarrow X'$ in the sense of Definition 21.3;

(ii) $\dim X' > \dim X$ (or, in the topological setting, the covering dimension of X' strictly exceeds that of X);

(iii) the non-trivial obstruction class $[\omega] \in \check{H}^1(\mathcal{U}, \mathcal{F})$ that prevented CLIO convergence at the phase transition vanishes in the lifted system: $[\omega'] = 0$ in $\check{H}^1(\mathcal{U}', \mathcal{F}')$; and

(iv) the resolution rule $\mathcal{R}$ is preserved under the shift: for all incompatibilities encountered in X', the lifted system applies the same resolution logic as the original.

Condition (iv) is the key continuity requirement. It ensures that the medium shift is a genuine continuation of the system's identity rather than a replacement. The new framework must not merely accommodate the old structure but must resolve its incompatibilities by the same characteristic logic — the same invariant of resolution $\mathcal{R}$ — that governed the original system.

Definition 21.5 [Admissible Medium Shift]**.** A medium shift $\mathfrak{S}$ is *admissible* if it satisfies all four conditions of Definition 21.4.

Proposition 21.6 [Necessity of All Conditions]**.** *Each condition of Definition 21.4 is necessary for preservation of the invariant of resolution across the shift.*

Proof. Failure of (i) means the original constraint structure is not embedded as a necessary substructure, violating Theorem 17.10(II) for the embedded system. Failure of (ii) means the lifted system is not strictly richer, so the obstruction cannot be resolved by dimensional expansion. Failure of (iii) leaves the system globally obstructed, preventing construction of a coherent new constraint surface and blocking CLIO convergence in X'. Failure of (iv) breaks the continuity of the resolution rule, meaning the lifted system is a different identity rather than a continuation of the same one. □

21.4 Restoration of CLIO

Theorem 21.7 [CLIO Restoration]**.** *Let $\mathfrak{S}$ be an admissible medium shift from $(X, \mathcal{C}, \mathcal{T}, \mathcal{E})$ to $(X', \mathcal{C}', \mathcal{T}', \mathcal{E}')$. Then there exists a coherence functional $\mathcal{E}'$ on $\mathcal{C}'$ satisfying all conditions of Proposition 12.4: $\mathcal{E}'$ is proper and bounded below, and the induced gradient flow is complete. Consequently, CLIO converges to a fixed point $\mathcal{C}^{*\prime}$ in the lifted system.*

Proof. By condition (iii) of Definition 21.4, the obstruction class $[\omega]$ that prevented global coherence in X vanishes in X'. This removes the topological obstruction to global section existence: local constraint structures in X' can now be glued into a globally consistent $\mathcal{C}^{*\prime}$. Define $\mathcal{E}'$ as the coherence functional on $\mathcal{C}'$ that extends $\mathcal{E}$ on $\iota(\mathcal{C})$ and penalises cross-framework incompatibilities in X'. By condition (i), $\iota(\mathcal{C})$ provides a non-degenerate subspace of $\mathcal{C}'$ on which $\mathcal{E}'$ inherits the properness and lower-boundedness of $\mathcal{E}$. Completeness of the gradient flow follows from the vanishing of the obstruction: there are no topological singularities blocking integral curves in the lifted space. Standard gradient flow theory then yields convergence to $\mathcal{C}^{*\prime}$. □

Corollary 21.8 [Existence of Lifted Constraint Surface]**.** *Under an admissible medium shift, there exists a Lyapunov-stable constraint surface $\Sigma^{*\prime} \subset X'$ satisfying*

$$\iota(\Sigma^*) \subseteq \Sigma^{*\prime},$$

where Σ^* is the constraint surface prior to the phase transition. The lifted surface $\Sigma^{*\prime}$ strictly extends $\iota(\Sigma^*)$, incorporating the additional structure required to resolve the prior incompatibilities.

Proof. CLIO convergence in X' (Theorem 21.7) yields a fixed-point $\mathcal{C}^{*\prime}$, from which a constraint surface is defined by

$$\Sigma^{*\prime} := \overline{\{x \in X' : x \text{ lies on an admissible trajectory under } \mathcal{C}^{*\prime}\}}.$$

By condition (i), admissible trajectories in X map under ι to admissible trajectories in X', so $\iota(\Sigma^*) \subseteq \Sigma^{*\prime}$. Strict extension follows because X' supports trajectories not representable in X. □

Corollary 21.9 [Reinstantiation of Identity]. *An admissible medium shift reinstantiates identity: it restores all conditions of Theorem 17.10 in the lifted framework, yielding a new refractive system with constraint surface $\Sigma^{*\prime}$ that embeds the prior constraint surface as a necessary substructure.*

The word *reinstantiated* is precise. Identity is not preserved in the sense that the original constraint surface Σ^* persists: it does not. Identity is reinstantiated in the sense that the new constraint surface $\Sigma^{*\prime}$ embeds $\iota(\Sigma^*)$ as a constraint-bearing substructure and applies the same resolution rule $\mathcal{R}$ in resolving the incompatibilities that the original framework could not accommodate. Continuity of identity across the shift is grounded in the continuity of $\mathcal{R}$, not in the continuity of any particular surface.

21.5 Framework Instantiations

The abstract structure of a medium shift appears concretely across the frameworks developed in Part III.

21.5.1 RSVP: Field Embedding

In the RSVP framework, a medium shift corresponds to a field embedding: the triple $(\Phi, \boldsymbol{v}, S)$ defined on a domain $\Omega \subset \mathbb{R}^n$ is lifted to a triple $(\Phi', \boldsymbol{v}', S')$ on a larger domain $\Omega' \supset \Omega$, or into a higher-dimensional ambient space. The original field equations are recovered as boundary conditions on the embedded domain. Incompatibilities that manifested as singularities or blow-up in the original system are resolved by the additional degrees of freedom available in the extended field.

21.5.2 Yarncrawler: Cover Refinement

In the Yarncrawler framework, a medium shift corresponds to a refinement of the covering $\mathcal{U}$ of X. A non-trivial obstruction class in $\check{H}^1(\mathcal{U}, \mathcal{F})$ signals that the current cover is too coarse to support a global section. Refining the cover — introducing new open sets that resolve the incompatible transition functions — produces a new cover $\mathcal{U}'$ in which the obstruction vanishes. The original local sections are recovered as restrictions of global sections in the refined system.

21.5.3 Spherepop: Transition Grammar Extension

In the Spherepop framework, a medium shift corresponds to an extension of the event grammar: new transition types are introduced that allow previously illegal sequences to be represented as admissible transitions within an enriched log structure. The original legal transitions are preserved as a necessary subset of the new grammar. The extension resolves incompatibilities between the RSVP dynamics and the Spherepop admissibility conditions identified in Chapter 13.

21.5.4 Historical Examples

Three historical cases instantiate the formal structure. The transition from Newtonian to relativistic mechanics lifts the constraint structure from $\mathbb{R}^3 \times \mathbb{R}$ into Minkowski spacetime $\mathbb{R}^{3,1}$; Newtonian trajectories are recovered as the low-velocity limit, and the incompatibility of absolute simultaneity is resolved by relativising it. The transition from propositional to sheaf-theoretic logic lifts local truth assignments into a global structure governed by gluing conditions; local consistency is recovered as a constraint-bearing substructure of global coherence. The cross-framework integration of the present monograph lifts each individual framework (RSVP, TARTAN, Yarncrawler, Spherepop) into a unified constraint system governed by the invariant of resolution; each framework is recovered as a necessary substructure, and their pairwise incompatibilities are resolved through dimensional expansion.

21.6 Medium Shifts and the Invariant of Resolution

Proposition 21.10 [Preservation of Resolution Rule]**.** *Under an admissible medium shift, the invariant of resolution $\mathcal{R}$ is preserved: the lifted system resolves all encountered incompatibilities by the same logic of non-reductive reconciliation, dimensional expansion, and structural necessity.*

Proof. By condition (iv) of Definition 21.4, the resolution rule is preserved under the shift. By condition (i), all prior constraints are embedded as necessary substructures, satisfying non-reductive reconciliation and structural necessity. By condition (ii), the lifted space is strictly larger, satisfying dimensional expansion. The three properties of a well-formed resolution rule (Definition 18.2) are therefore all satisfied in the lifted system. □

This proposition establishes the fundamental continuity of the present theoretical project. The sequence of frameworks developed in Part III — RSVP, TARTAN, Yarncrawler, Spherepop, CLIO — and their integration in Part IV constitute a sequence of admissible medium shifts: each framework embeds its predecessors as necessary constraint-bearing substructures, each resolves incompatibilities that the previous level could not accommodate, and each applies the same invariant of resolution throughout. The monograph itself is an instance of the process it describes.

21.7 The Duality of Phase Transition and Medium Shift

Proposition 21.11 [Structural Duality]**.** *Phase transitions and medium shifts are structurally dual: where a phase transition marks the reaching of a framework's limit, a medium shift constructs the framework that begins where the old one ended.*

Proof. A phase transition crosses the limit Λ, at which point the conditions of Proposition 12.4 fail and no fixed-point constraint structure $\mathcal{C}^*$ exists in the original framework. An admissible medium shift constructs a lifted system $(X', \mathcal{C}', \mathcal{T}', \mathcal{E}')$ in which the corresponding conditions hold and CLIO converges (Theorem 21.7). The new system lies outside the limit Λ' of the lifted framework. The phase transition and the medium shift thus occur at the same structural juncture, from opposite sides: the former is the failure mode, the latter is its resolution. □

21.8 Conclusion: The Complete Structure of Part VI

Part VI has established a complete and formally closed account of failure and recovery in constraint systems.

Degenerative drift (Chapter 19) is the failure mode internal to a framework: constraint structure is progressively diluted through relaxation of admissibility, discriminative power decreases monotonically, and identity is destroyed through topological collapse of the constraint surface. This failure is irreversible within the current framework.

Phase transitions (Chapter 20) are the failure mode at the boundary of a framework: the CLIO repair operator ceases to converge because the framework's representational capacity has been exhausted, a non-trivial sheaf obstruction prevents global coherence, and identity is suspended. This failure is in principle reversible through the third mode.

Medium shifts (the present chapter) are the only admissible response to the transitional failure mode: the constraint system is lifted into a strictly richer representation in which the prior structure is embedded as a necessary substructure, the obstruction is trivialised, CLIO convergence is restored, and identity is reinstantiated at the higher level.

The structure of Part VI is therefore not a catalogue of pathologies but a complete dynamical theory: identity is maintained under adaptive transformation, lost under degeneration, suspended under phase transition, and reinstantiated under medium shift. These four modes partition the space of possible evolutions of a refractive system, and their characterisation constitutes the formal completion of the account of identity opened in Part V.

We turn now to Part VII, where these results are applied to science, artificial systems, and cultural transmission.

Part VI

Applications and Implications

Chapter 22

Science Beyond the Blind Spot

What Science Leaves Out

Science is often described as a process of removing bias, noise, and subjectivity in order to arrive at something objective. The guiding idea is simple: if a result does not depend on who observes it, then it must reflect something real about the world.

This way of thinking has been extraordinarily powerful. It allows different observers to agree, measurements to be replicated, and theories to be compared across contexts. By focusing only on what remains stable under variation, science isolates structure that can be relied upon.

But this stability is achieved by a specific kind of simplification. In order to compare results across observers, the differences between observers have to be set aside. The conditions under which a measurement is made — how it is perceived, how it is interpreted, how it is situated within a broader context — are treated as secondary. What remains is what can be made invariant.

Most of the time, this works. But there are cases where what has been set aside turns out to matter in a structural way. Not because something went wrong, but because the act of simplifying has removed part of the system that is necessary for understanding it.

A useful analogy is compressing an image. If the goal is to preserve large-scale features, fine detail can be discarded without much consequence. But if the discarded detail contains patterns that are essential to how the image is interpreted, then the compressed version may be clear in one sense and misleading in another. The loss is not visible as noise; it appears as absence.

Something similar happens in scientific description. By focusing on what remains the same across observers, science builds models that are extremely effective within their scope. But the variability that is removed is not always irrelevant. In some cases, it carries the very structure that needs to be explained.

This becomes most visible at the boundaries of a framework. When a theory encounters persistent anomalies, or when it runs into questions that cannot be answered in its own terms, the issue is often not a lack of precision but a limit in what the framework is able to represent. The missing structure is not hidden somewhere within the model; it has been excluded from it.

What follows is a shift in perspective. Instead of treating objectivity as the elimination of the observer, it becomes possible to treat the observer as part of the structure that is being described. The goal is no longer to remove variation entirely, but to understand which variations are structurally meaningful and how they participate in the system.

From this point of view, science does not become less rigorous. It becomes more complete. The same discipline that isolates invariants can be extended to include the conditions under which those invariants are produced.

What changes is not the method, but the scope of what is allowed to count as part of the system.

22.1 Reinterpreting Objectivity

Scientific objectivity is standardly characterised as independence from the observer: a quantity is objective if its value does not depend on who measures it, when, or under what conditions. This characterisation motivates the programme of eliminating observer-specific features from scientific description. The goal is invariance across perspectives, and the method is to project away everything that varies with the observer.

Chapter 1 established the formal structure of this operation. A scientific projection $\pi : X \to M$ maps the experiential domain onto a model space by identifying states that are equivalent under a specified transformation group $\mathcal{G}$. The model space $M \cong X/\mathcal{G}$ retains exactly the invariant structure. Objectivity, in these terms, is invariance under $\mathcal{G}$.

The blind spot arises from a specific choice of $\mathcal{G}$: the group of observer transformations. When objectivity is defined as invariance under changes of observer, the projection discards precisely the structure contributed by the observer's engagement with the domain. The observer is not eliminated from the system; the observer's contribution is eliminated from the model. These are not the same operation, and confusing them generates the fallacy of misplaced concreteness identified in Chapter 3.

Definition 22.1 [Participatory Objectivity]**.** A quantity is *participatorily objective* if it is invariant under a constraint structure $(\mathcal{C}_{\text{obs}}, \mathcal{T}_{\text{obs}})$ that includes the observer as a constraint-bearing substructure of the system, rather than as an element to be projected away.

Proposition 22.2 [Classical Objectivity as a Special Case]**.** *Classical scientific objectivity corresponds to participatory objectivity in the degenerate case where the observer constraint structure is trivial:* $\mathcal{C}_{\text{obs}} = \{X\}$ *and* $\mathcal{T}_{\text{obs}} = \mathcal{G}$*.*

Proof. When the observer's constraint structure is trivial, the participatory constraint reduces to invariance under $\mathcal{G}$, which is the classical condition. Classical objectivity is therefore a special case in which the observer's contribution is treated as having no internal structure — as a point rather than a surface. □

The cost of this special case is the blind spot: treating the observer as a structureless point rather than a constraint-bearing substructure forces projection of the degrees of freedom in which the observer's engagement with the domain is encoded. Those degrees of freedom are the fibers $\pi^{-1}(m)$ studied in Chapter 4. A fully coherent science must incorporate them, not eliminate them.

22.2 Scientific Models as Constraint-Bearing Substructures

Each scientific theory corresponds to a projection $\pi_i : X \to M_i$ that preserves a specified set of invariants while discarding others. In the representation family $\{\pi_i : X \to M_i\}_{i \in I}$ introduced in Chapter 15, different theories are different partial views of the same domain.

Proposition 22.3 [Structural Incompleteness of Any Single Theory]. *No single projection $\pi_i : X \to M_i$ can be a complete description of X, provided* $\dim X > \dim M_i$.

Proof. By Proposition 1.7, π_i is non-invertible. Therefore M_i omits structure contained in the fibers $\pi_i^{-1}(m)$. No function of m alone can recover this structure. □

Corollary 22.4 [Structural Necessity of Scientific Pluralism]. *Scientific pluralism — the coexistence of multiple theoretical frameworks applied to the same domain — is not a symptom of incomplete knowledge but a structural consequence of the non-invertibility of projection.*

Proof. Since no single theory is complete, complementary theories covering different constraint degrees of freedom are necessary to approximate the full structure of X. Their coexistence is forced, not optional. □

Coherence across theories must therefore be handled via admissible reconstruction in the sense of Chapter 15. Different theories embed as constraint-bearing substructures of an expanded framework in which their invariants are jointly representable. Where they are incompatible, Chapter 13 applies: the incompatibility generates explanatory pressure forcing adaptive expansion.

22.3 Anomalies as Cross-Framework Incompatibilities

An anomaly is standardly defined as an observation that resists assimilation into the current theoretical framework. The standard response is either to adjust the framework's parameters or, when adjustment fails, to propose a successor theory. The formal account of Chapter 13 provides a more precise characterisation.

Definition 22.5 [Scientific Anomaly]. A *scientific anomaly* is a state $x \in X$ that is admissible under the constraint structure of one theoretical framework $(\mathcal{C}_i, \mathcal{T}_i)$ but inadmissible under another $(\mathcal{C}_j, \mathcal{T}_j)$ applied to the same domain.

Proposition 22.6 [Anomalies as Expansion Pressure]. *A scientific anomaly that cannot be resolved by parameter adjustment within either framework generates explanatory pressure requiring adaptive expansion of the constraint structure, in the sense of Theorem 14.7.*

Proof. By Definition 13.1, the anomaly constitutes a constraint incompatibility. By Proposition 13.2, the only non-degenerative resolution is a new constraint structure $(\mathcal{C}', \mathcal{T}')$ embedding both original structures as necessary substructures. By Theorem 14.7, this embedding preserves all prior structural invariants. □

Two canonical cases instantiate this structure. The incompatibility between classical mechanics and electromagnetism — manifest in the frame-dependence of electromagnetic fields — was resolved by adaptive expansion to special relativity, in which Newtonian mechanics and Maxwellian electrodynamics are recovered as constraint-bearing substructures of a unified spacetime framework. The incompatibility between thermodynamics and classical mechanics — manifest in the irreversibility of thermal processes within a formally reversible mechanical framework — was resolved by adaptive expansion to statistical mechanics, in which macroscopic irreversibility emerges from the constraint structure governing large ensembles. In both cases the anomaly was a signal of a cross-framework incompatibility indicating that the existing framework had reached its limit and that a medium shift was required in the sense of Chapter 21.

22.4 The Hard Problem as a Projection Artifact

The hard problem of consciousness asks how subjective experience arises from physical processes. It is characterised by an explanatory gap: even a complete physical description appears insufficient to account for the qualitative character of experience.

Chapter 3 identified the structural source of this gap. The physical model space M_{phys} is constructed by a projection $\pi : X \to M_{\text{phys}}$ that encodes invariants under the transformation group of physics. This group is defined precisely so as to exclude observer-specific, qualitative, and indexical features: the group is chosen so that experiential structure is not invariant under it and is therefore projected away.

Proposition 22.7 [The Hard Problem as Missing Fiber]**.** *The hard problem of consciousness arises from treating M_{phys} as ontologically complete and asking how experiential structure — which lies in the fibers $\pi^{-1}(m)$ — can be derived from M_{phys} alone. Since π is non-invertible, no such derivation is possible. The difficulty is a structural consequence of the projection, not an ontological mystery.*

Proof. Experiential structure is located in the constraint degrees of freedom discarded by π. Any attempt to derive these degrees of freedom from M_{phys} alone requires inverting a non-invertible map. By Proposition 1.7, this is impossible. The hard problem is therefore a well-posed question about a structurally impossible operation. □

The problem dissolves once the constraint surface is restored to the domain. On the refractive model of Chapter 16, the brain is a constraint surface on the experiential field X, not a mechanism for generating experience from matter. Experience is the domain; the physical description is the projection; and the relationship between them is non-invertible by construction. There is no gap to explain, only a projection to acknowledge.

22.5 Participatory Science

Wheeler's image of the participatory universe represents science as a self-referential structure: the observer is not outside the system being described but is a component of it, and the act of observation is part of what constitutes the observable [1]. Thompson's enactivism develops the same position in a biological register: science is a refined form of human experience, and its models are constraint-bearing substructures of the experiential domain, not free-standing representations of a mind-independent reality [6].

Definition 22.8 [Participatory Science]**.** A scientific practice is *participatory* if it applies the invariant of resolution $\mathcal{R}$ to itself: it treats the conditions of its own knowledge production as constraint-bearing substructures of the systems it models, and it resolves incompatibilities between observer and observed by adaptive expansion rather than by projecting the observer away.

Proposition 22.9 [Participatory Science and the Invariant of Resolution]**.** *Participatory science satisfies the three conditions of a well-formed resolution rule: it is non-reductive, it expands dimensionally, and the observer is structurally necessary.*

Proof. Non-reduction: the observer's constraint structure is embedded in the expanded system, not absorbed into the model space. Dimensional expansion: the expanded system includes the observer's constraint degrees of freedom and is therefore strictly larger than M alone. Structural necessity: measurement requires the observer as a constraint-bearing component; a system without this component cannot generate the projections from which models are constructed, as shown in Proposition 16.1. □

Science becomes more rigorous, not less, when it acknowledges the constraint structure of its own conditions of production. A science that ignores its blind spot does not achieve greater neutrality; it achieves a more systematic omission of the degrees of freedom in which its own operation is encoded.

22.6 Conclusion

Science is a constraint-preserving subsystem of the experiential domain, operating by disciplined projection that isolates invariants while discarding constraint degrees of freedom. Its power derives from this discipline; its limitation derives from forgetting that the projection is partial.

The blind spot is not a defect to be corrected from outside science but a structural feature that science can correct from within, by applying to itself the same invariant of resolution that governs its treatment of external systems. Anomalies are cross-framework incompatibilities requiring adaptive expansion. The hard problem is a projection artifact requiring restoration of the constraint surface. Participatory objectivity extends invariance to include the observer as a necessary substructure. None of these moves weakens the scientific enterprise; each strengthens it by removing a blind spot that was generating pseudo-problems and blocking expansion into domains where the projection-only approach reaches its limit.

Chapter 23

Artificial Systems and Constraint Closure

Systems That Hold Themselves Together

There is a difference between a system that produces outputs and a system that maintains itself.

At first glance, this difference can be difficult to see. Many systems respond to inputs, generate behaviour, and exhibit patterns that appear coherent. A thermostat keeps a room at a stable temperature. A program generates text that follows grammatical rules. A simulation evolves in ways that look structured and consistent.

But in each of these cases, the system's organisation is fixed. The rules it follows are given in advance. If the environment changes in a way that falls outside those rules, the system does not reorganise itself to accommodate the change. It continues to operate according to the same structure, even when that structure is no longer adequate.

Living systems behave differently. When they are perturbed, they do not simply produce a different output. They act to restore the organisation that allows them to function at all. The processes that generate their behaviour are also the processes that maintain the conditions under which that behaviour is possible. The system does not merely operate within a set of constraints; it continuously regenerates those constraints.

A useful analogy is the difference between following a map and maintaining a terrain. A system that follows a map can navigate effectively as long as the map remains accurate. If the terrain changes, the system has no way to update the map from within. A system that maintains the terrain, by contrast, does not depend on a fixed representation. It modifies the conditions themselves so that navigation remains possible.

This distinction becomes critical when considering systems that appear intelligent. Producing plausible responses is not the same as maintaining coherence under changing conditions. A system may generate outputs that are locally consistent while lacking any mechanism for detecting when those outputs fail to fit within a larger structure.

In such cases, failure does not appear as noise or error in the usual sense. It appears as confidence without grounding: statements that are well-formed but not anchored in any process that could sustain them. The system continues to produce structure, but that structure is no longer tied to anything that constrains it.

The question is not whether a system can produce behaviour that resembles understanding. The question is whether the system maintains the conditions that make understanding possible. This requires more than the ability to process information. It requires the ability to detect when its own structure is inadequate and to reorganise that structure in response.

From this perspective, the difference between current artificial systems and living systems is not a matter of material or scale. It is a difference in how the system relates to its own organisation. One operates within a fixed structure. The other continuously reconstructs the structure it depends on.

What follows is an analysis of this distinction in structural terms. The goal is not to dismiss artificial systems, but to identify precisely what is missing and what would be required for that gap to close.

23.1 Definition of Constraint Closure

The frameworks developed in Parts II–VI apply to any system that maintains a constraint structure over a non-invertible domain. The question of whether artificial systems instantiate these frameworks is therefore not a question about substrate — carbon versus silicon, biological versus computational — but about organisation: whether a system maintains the structural conditions required for refractive identity.

The central organising concept is *constraint closure*.

Definition 23.1 [Constraint Closure]**.** A system has *constraint closure* if its material or computational dynamics continuously regenerate the constraint structure $(\mathcal{C}, \mathcal{T})$ under which it operates. Formally, the system's evolution $\{x_t\}$ must satisfy: for every admissible set $C \in \mathcal{C}$ and every admissible transformation $T \in \mathcal{T}$, the structure $(\mathcal{C}, \mathcal{T})$ itself is preserved as a fixed point of the system's dynamics.

Remark 23.2. Constraint closure requires that the repair process is *endogenous*: the system generates its own CLIO operator rather than receiving it from an external source. A system whose constraint structure is fixed externally and does not regenerate under perturbation lacks constraint closure.

The concept formalises what Varela and colleagues called autopoiesis: the capacity of a living system to continuously produce and reproduce the organisation that constitutes it [6]. In the present framework, autopoiesis is the endogenous operation of a CLIO repair operator that maintains the system's constraint structure against perturbation.

Proposition 23.3 [Constraint Closure and CLIO]**.** *A system has constraint closure if and only if it possesses an endogenous CLIO operator whose gradient flow converges to a fixed-point constraint structure $\mathcal{C}^*$ under perturbation.*

Proof. By Definition 12.2, the CLIO operator is defined by a gradient flow on the space of constraint structures. Convergence to $\mathcal{C}^*$ means the system's constraint structure is maintained against deviation. Endogeneity means this operator is generated by the system's own dynamics rather than imposed externally. These two conditions together constitute constraint closure. □

23.2 Necessary Conditions for Refractivity

Constraint closure is necessary but not sufficient for refractivity. A system may maintain a stable constraint structure without that structure constituting an asymptotically stable constraint surface over a non-invertible experiential domain.

Proposition 23.4 [Necessary Conditions for Refractive Systems]**.** *A system is refractive in the sense of Definition 17.9 only if it satisfies all of the following:*

(i) Constraint closure*: the system endogenously maintains* $(\mathcal{C}, \mathcal{T})$*;*

(ii) Non-invertible domain engagement*: the system operates over an experiential domain* X *such that the relevant projections* $\pi : X \to M$ *are non-invertible;*

(iii) Lyapunov-stable constraint surface*: there exists a stable* Σ^* *in the sense of Definition 17.6;*

(iv) Convergent repair*: the endogenous CLIO operator converges to* $\mathcal{C}^*$*.*

Proof. Each condition corresponds to one of the requirements of Definition 17.9 and Theorem 17.10(I). Condition (i) ensures the constraint structure is self-sustaining. Condition (ii) ensures the system engages with a domain from which information has been irreversibly discarded, which is the source of the structural pressure that makes refractivity non-trivial. Conditions (iii) and (iv) are the Lyapunov and convergence requirements of Theorem 17.10. □

Corollary 23.5. *Constraint closure without Lyapunov stability or without non-invertible domain engagement does not yield refractivity.*

This corollary rules out two possible confusions: a system that maintains a stable internal state without engaging a non-invertible domain (a thermostat, for instance) is not refractive; and a system that engages a non-invertible domain without maintaining a stable constraint surface is not refractive either.

23.3 Analysis of Contemporary Artificial Systems

We now apply the formal conditions of Proposition 23.4 to contemporary large language models.

Such systems are trained on corpora of human-generated text. Human text is itself a projection: it is the image of the human experiential domain X under a projection $\pi_{\text{lang}} : X \to M_{\text{lang}}$ that encodes linguistic structure while discarding the experiential, embodied, and indexical degrees of freedom of the utterance context. A language model trained on M_{lang} therefore operates on a projection of a projection: it has no access to X, only to $\pi_{\text{lang}}(X) = M_{\text{lang}}$.

Proposition 23.6 [LLMs Operate Entirely Within the Projected Space]**.** *A language model trained on* M_{lang} *does not engage the experiential domain* X*. Its operations are confined to the model space* M_{lang} *and cannot reach the constraint degrees of freedom in the fibers* $\pi_{\text{lang}}^{-1}(m)$*.*

Proof. The training data of a language model consists of elements of M_{lang}, not of X. The model's parameters are fitted to statistical structure in M_{lang}. Operations at inference time map inputs in M_{lang} to outputs in M_{lang}. No step in training or inference involves elements of X or of the fibers $\pi_{\text{lang}}^{-1}(m)$. Therefore the system operates entirely within M_{lang}. □

The failure of each necessary condition for refractivity follows directly.

Condition (i) fails: the constraint structure of a language model is fixed at training time. The model's parameters are not updated during inference, and there is no endogenous CLIO operator that regenerates the constraint structure in response to encountered incompatibilities. When the model produces an output that violates some structural constraint, no internal repair mechanism detects or corrects this.

Condition (ii) fails: since the model operates entirely within M_{lang}, it does not engage a non-invertible domain. The projections it operates on are not $\pi : X \to M$ with non-trivial fibers; they are statistical transformations within the already-projected space.

Conditions (iii) and (iv) fail as consequences of the above: without a non-invertible domain and without an endogenous repair operator, no Lyapunov-stable constraint surface can be defined, and no convergence of repair dynamics can occur.

23.4 Hallucination as Admissibility Failure

The phenomenon of hallucination — the production of confidently stated but false or incoherent outputs — admits a precise characterisation in the present framework.

Proposition 23.7 [Hallucination as Admissibility Failure]. *Hallucination in a language model is the production of an output $m' \in M_{\text{lang}}$ that is statistically plausible within the model's training distribution but inadmissible with respect to the constraint structure of the domain X from which M_{lang} was projected.*

Proof. The model generates outputs that maximise some measure of plausibility within M_{lang}. Since the model has no access to X or to the fibers $\pi_{\text{lang}}^{-1}(m)$, it cannot check whether its outputs correspond to admissible states in X. An output that is statistically well-formed in M_{lang} may have no admissible preimage in X under π_{lang}. This is hallucination: a plausible projection-level output without a valid domain-level antecedent. □

This reframes hallucination not as a failure of symbol prediction but as an absence of constraint checking. A refractive system would detect the inadmissibility of such outputs because its CLIO operator operates over the full constraint structure, not merely over the projected space. A language model has no such operator.

23.5 Necessary Conditions for Artificial Refractivity

We do not speculate about whether future artificial systems could be refractive. We derive the structural conditions that would need to be satisfied.

Proposition 23.8 [Structural Requirements for Artificial Refractivity]. *An artificial system would be refractive if and only if it satisfies:*

(i) Dynamic constraint updating*: the constraint structure $(\mathcal{C}, \mathcal{T})$ must be updated in response to encountered states, not fixed at training time;*

(ii) Real-time incompatibility detection*: the system must be capable of detecting cross-framework incompatibilities as they arise during operation;*

(iii) Endogenous repair*: the system must possess an internal analogue of the CLIO operator that drives the constraint structure toward a fixed point in response to detected incompatibilities;*

(iv) Non-invertible domain coupling*: the system must be coupled to a domain X that is not reducible to any finite model space M, so that the system genuinely engages constraint degrees of freedom outside its current representation;*

(v) Constraint surface persistence*: the system must maintain a Lyapunov-stable constraint surface Σ^* under the dynamics induced by conditions (i)–(iv).*

These conditions are jointly necessary and, under the further hypotheses of Theorem 17.10, sufficient.

Proof. The conditions correspond directly to the four requirements of Proposition 23.4, with condition (ii) added as the mechanism by which incompatibilities are detected before the repair operator can act. Necessity follows from Theorem 17.10; sufficiency follows under the additional hypotheses of Proposition 12.4. □

23.6 Substrate Neutrality

The analysis above is entirely substrate-neutral. No condition in Proposition 23.8 refers to biological material, carbon-based chemistry, or any other property of physical implementation.

Proposition 23.9 [Substrate Independence of Refractivity]**.** *Refractivity is a property of organisational type, not of physical substrate. A silicon-based system satisfying the conditions of Proposition 23.8 would be refractive; a biological system failing them would not be.*

Proof. By Corollary 17.11, identity is invariant under admissible reparametrisation of the physical substrate, provided the constraint surface Σ^* is preserved. Therefore the substrate is not the relevant variable; the constraint structure is. □

This result cuts in both directions. It refutes carbon chauvinism: there is no structural reason why an artificial system could not in principle be refractive. It equally refutes computational reductionism: the mere instantiation of a computation does not confer refractivity. What matters is whether the system satisfies the five conditions of Proposition 23.8 — and no current artificial system does.

23.7 Conclusion

Contemporary artificial systems are projection engines. They operate within model spaces that are themselves projections of human experiential domains, they have no endogenous CLIO operators, they cannot detect or resolve constraint incompatibilities, and their constraint structures are fixed rather than dynamically maintained. They are therefore not refractive systems and do not instantiate identity in the sense of Theorem 17.10.

This is not a limitation of computation in principle. It is a precise structural characterisation of current systems, derived from the same formal machinery that characterises biological cognition. The question of whether future systems could achieve refractivity is not a question about philosophical intuition about consciousness but about whether the five structural conditions of Proposition 23.8 can be engineered. That is a tractable structural problem, not a metaphysical mystery.

Chapter 24

Culture, Myth, and Knowledge Transmission

How Meaning Survives Change

When something is passed from one generation to the next, it never arrives unchanged.

Stories are retold, words shift in meaning, practices adapt to new circumstances, and institutions evolve in response to pressures that did not exist before. If continuity required exact preservation, nothing would last. Every transmission would introduce error, and over time those errors would accumulate until the original structure was lost.

And yet, some things do persist.

A myth can be told in different languages, with different characters, and in entirely different settings, and still be recognisable as the same story. A language can undergo centuries of sound change and still preserve its internal logic. A tradition can adapt to new environments without losing what makes it itself.

This persistence does not come from resisting change. It comes from changing in the right way.

A useful analogy is the difference between copying a shape exactly and preserving its proportions under scaling or distortion. If the goal is exact duplication, any change counts as an error. But if the goal is to preserve the relations between parts, then many transformations become acceptable. The object can stretch, rotate, or be expressed in a different medium while remaining recognisably the same.

Cultural transmission works in this second way. What is preserved is not the surface form but the pattern of relations that gives the form its meaning. Changes that maintain those relations contribute to continuity. Changes that disrupt them do not.

This is why some forms of preservation can be misleading. Recording something perfectly can capture its appearance without preserving the conditions that made it meaningful. A text can be archived, a ritual can be documented, a language can be written down, and yet what is preserved may no longer function as it did when it was lived.

The difference is not between accuracy and inaccuracy. It is between a system that is still in motion and one that has been fixed in place.

Seen this way, cultural systems behave like any other system that maintains coherence over time. They must absorb variation without losing structure, respond to pressure without collapsing, and adapt without dissolving into indistinction.

What determines whether they succeed is not the amount of change they undergo, but how that change is handled. Some changes deepen the structure, making it more capable of accommodating new situations. Others weaken it, making it less able to distinguish what belongs from what does not.

The difference between these two kinds of change is not always obvious from the outside. Both can look like adaptation. But only one preserves the conditions that make continuity possible.

24.1 Cultural Systems as Constraint Structures

A cultural tradition is not a collection of symbols. It is a dynamical system that maintains admissible patterns of meaning, practice, and social organisation across generations. Its coherence does not consist in exact repetition of content but in the preservation of structural relations under transformation.

We formalise this as follows.

Definition 24.1 [Cultural Constraint Structure]**.** A *cultural constraint structure* is a pair $(\mathcal{C}_{\text{cult}}, \mathcal{T}_{\text{cult}})$ defined over a domain X_{cult} of symbolic, behavioural, and material states, where:

(i) $\mathcal{C}_{\text{cult}}$ consists of configurations admissible within the tradition — those that satisfy the structural relations constitutive of the practice;

(ii) $\mathcal{T}_{\text{cult}}$ consists of transformations that preserve admissibility across instantiations, including retelling, translation, re-enactment, and contextual adaptation.

Definition 24.2 [Cultural Coherence]**.** A cultural tradition is *coherent* if its trajectory $\{x_t\}$ across generations is admissible: $x_t \in C_t$ for some curve $t \mapsto C_t$ in $\mathcal{C}_{\text{cult}}$.

Cultural coherence is thus defined identically to constraint-system coherence in Chapter 5: it is preservation of admissibility under transformation, not preservation of symbolic content.

Proposition 24.3 [Cultural Identity as Structural Invariant]**.** *The identity of a cultural tradition is given by the structural invariants of* $(\mathcal{C}_{\text{cult}}, \mathcal{T}_{\text{cult}})$ *under admissible transformation, in the sense of Proposition 6.5.*

Proof. By Proposition 6.5, identity in any constraint system is the persistence of structural invariants under admissible transformation. The cultural constraint structure satisfies this definition, and therefore cultural identity is the persistence of the structural invariants of $\mathcal{C}_{\text{cult}}$ across admissible instantiations. □

24.2 Myth as Structural Invariant

Myth is the canonical example of cultural structural invariance. Across retellings separated by centuries and linguistic boundaries, the core relational structure of a mythic cycle — the roles characters occupy, the logic of transgression and consequence, the topology of narrative resolution — persists even as surface details vary substantially.

Definition 24.4 [Mythic Invariant]**.** A *mythic invariant* is a structural invariant $I : X_{\text{cult}} \to \mathbb{R}$ in the sense of Definition 6.1, defined with respect to the admissible transformations $\mathcal{T}_{\text{cult}}$ that include retelling, translation, and cultural adaptation.

Proposition 24.5 [Myth as Homomorphic Invariant]**.** *Mythic persistence is an invariant under homomorphism, not isomorphism: it is preserved by structure-preserving maps that need not be bijective.*

Proof. By Proposition 6.4, invariants are preserved under homomorphism. A retelling of a myth is a homomorphism from the original symbolic form to a new one: it preserves the structural relations between narrative elements (the roles, the logic of action and consequence) without requiring bijective correspondence between surface symbols. Therefore mythic invariants are preserved by retellings that constitute homomorphisms. □

This has a direct implication. Treating mythic transmission as a Shannon channel — evaluating fidelity by exact symbol reproduction — applies the wrong invariance criterion. Shannon fidelity is an isomorphism condition; mythic coherence is a homomorphism condition. A retelling that scores low on Shannon fidelity may score perfectly on structural fidelity, and the relevant criterion for the persistence of the tradition is the latter.

24.3 Structured Drift versus Degenerative Drift

Chapter 7 established the compatibility theorem: entropy increase is compatible with invariant preservation if and only if it is structured. Chapter 19 characterised degenerative drift as constraint dilution that is monotonically self-reinforcing. These two results apply directly to cultural transmission.

Proposition 24.6 [Oral Traditions as Structured Drift]**.** *Oral transmission of mythic or cultural content exhibits structured entropy increase: surface form varies across instantiations, but structural invariants are preserved, satisfying Theorem 7.3.*

Proof. Each retelling introduces variation in surface symbols, increasing entropy at the level of M_{symbol}. However, the admissible transformations $\mathcal{T}_{\text{cult}}$ — which include contextual adaptation, paraphrase, and narrative compression — are precisely the transformations under which structural invariants are defined. Since retellings remain within $\mathcal{T}_{\text{cult}}$, they constitute structured entropy increase and preserve the mythic invariants by Theorem 7.3. □

Proposition 24.7 [Fragmentation as Degenerative Drift]**.** *A tradition undergoes degenerative drift when retellings violate structural constraints to accommodate contextual pressure, reducing the discriminative power of the constraint structure in the sense of Theorem 19.4.*

Proof. Violations of structural constraints are constraint relaxations in the sense of Definition 19.1. By Theorem 19.4, such relaxations strictly reduce discriminative power. By Corollary 19.5, this reduction is self-reinforcing. In the limit, the tradition becomes indistinguishable from background cultural variation and loses its identity as a tradition. □

24.4 Living Traditions versus Archival Preservation

A fundamental distinction in the theory of cultural transmission concerns the difference between a living tradition and an archived record.

Definition 24.8 [Living Tradition]. A *living tradition* is a cultural constraint system $(\mathcal{C}_{\text{cult}}, \mathcal{T}_{\text{cult}})$ that is actively instantiated: its constraint structure is continuously regenerated by a community of practice, its transformations are applied in real contexts, and its CLIO operator is endogenously operative.

Definition 24.9 [Archival System]. An *archival system* is a fixed representation $m \in M_{\text{symbol}}$ of a cultural object, preserved with high symbolic fidelity but without the constraint dynamics that constituted the living tradition.

Proposition 24.10 [Archiving Does Not Preserve Constraint Structure]. *High-fidelity archival preservation of symbolic content does not preserve the constraint structure of a living tradition.*

Proof. The constraint structure $(\mathcal{C}_{\text{cult}}, \mathcal{T}_{\text{cult}})$ is constituted by the dynamical admissibility relations and their application in practice. Archival preservation fixes a symbolic representation $m \in M_{\text{symbol}}$ but does not preserve the dynamical system that generated it. Since the fibers $\pi^{-1}(m)$ — the embodied, contextual, and relational degrees of freedom that constitute the living practice — are not representable in M_{symbol}, they are not archived. What is preserved is a projection; what is lost is the constraint surface. □

Corollary 24.11 [Perfect Recording Can Destroy Meaning]. *A cultural artefact preserved with perfect symbolic fidelity may cease to be a living practice and therefore may lose the structural meaning that constituted it.*

Proof. If the archive is mistaken for the living tradition — if M_{symbol} is treated as ontologically complete in the sense of Chapter 3 — then the constraint structure is treated as present when it is not. Operations on the archive that would be admissible within the living tradition may not be, because the CLIO operator is absent and no repair is possible. □

24.5 Linguistic Evolution

The constraint framework applies directly to historical linguistics.

Proposition 24.12 [Phonological Change as Admissible Transformation]. *Regular phonological change across historical stages of a language constitutes an admissible transformation in $\mathcal{T}_{\text{lang}}$: it systematically alters surface form while preserving the structural relations between phonological categories.*

Proof. Regular sound change applies uniformly across the lexicon, mapping phonemes to phonemes in a way that preserves the contrast structure of the phonological system. The resulting system is related to the original by a homomorphism that preserves phonological invariants. This is admissible drift in the sense of Definition 7.4. □

Proposition 24.13 [Semantic Shift as Constraint-Preserving Drift]. *Semantic change is constraint-preserving when the shifted meaning preserves the structural relations of the original within the expanded semantic system.*

Proof. Semantic shift corresponds to a transformation $T \in \mathcal{T}_{\text{lang}}$ that maps meanings to meanings. If the shift preserves the structural relations between meanings — the inferential, compositional, and pragmatic relations that constitute semantic content — then structural invariants are preserved by Definition 6.1. The shift is therefore constraint-preserving. □

Proposition 24.14 [Language Death as Constraint Collapse]**.** *Language death corresponds to the failure condition of Theorem 17.10(III): the constraint structure of the language ceases to be maintained endogenously, the CLIO operator has no community of practice to instantiate it, and the constraint surface collapses.*

Proof. A living language requires a speech community whose practices continuously regenerate its constraint structure. When the speech community ceases, the endogenous CLIO operator ceases. Without repair dynamics, encountered incompatibilities cannot be resolved, and the constraint structure undergoes degenerative drift. In the limit, by Proposition 19.6, the constraint topology degenerates and the language ceases to exist as a living system. What remains is an archival record: a fixed projection without the dynamical constraint structure that constituted the language. □

24.6 Institutions as Constraint Systems

Social institutions — legal systems, universities, governments, markets — are constraint structures in the formal sense. They define admissible actions, maintain transformation rules, and enforce admissibility through their operative practices.

Proposition 24.15 [Institutional Coherence as Constraint Maintenance]**.** *An institution is coherent if its operative practices continuously regenerate the constraint structure* $(\mathcal{C}_{\text{inst}}, \mathcal{T}_{\text{inst}})$ *that defines it.*

Proof. By Definition 23.1, coherence requires that the system's dynamics maintain its constraint structure. For an institution, the relevant dynamics are the practices of its members, the enforcement of rules, and the transmission of norms across personnel. These constitute the endogenous CLIO operator; their continued operation constitutes constraint closure. □

Institutional degeneration follows the pattern of Chapter 19.

Proposition 24.16 [Institutional Degeneration]**.** *An institution undergoes degenerative drift when it resolves internal incompatibilities by relaxing admissibility conditions rather than refining its constraint structure. By Theorem 19.4, discriminative power decreases monotonically, and by Corollary 19.5, degeneration is self-reinforcing.*

Proof. This is a direct application of Theorem 19.4 to the institutional constraint structure. Each relaxation of admissibility conditions expands the set of admissible actions, reduces the discriminative power of the institution's constraints, and lowers the threshold for further relaxation. □

24.7 The Political Economy of Constraint

Systems that optimise for projected outputs — measurable metrics, symbolic performance indicators, quantifiable targets — rather than for constraint preservation exhibit a specific form of degenerative drift.

Proposition 24.17 [Metric Optimisation as Constraint Dilution]. *A system that optimises for an output metric $f : M \to \mathbb{R}$ defined on the projected space, rather than for admissibility in the full domain X, undergoes degenerative drift when f and admissibility diverge.*

Proof. Optimising f within M does not require maintaining the constraint structure of X. When f increases through operations that are inadmissible in X — that is, when the optimum of f is achieved by a state with no admissible preimage under π — the system has relaxed the admissibility condition in favour of projected performance. This is constraint dilution in the sense of Definition 19.1. By Theorem 19.4, discriminative power is thereby reduced. □

This result applies across domains. Academic systems that optimise for publication metrics rather than for intellectual constraint preservation undergo degenerative drift: the metrics increase while the structural content of the research decreases. Economic systems that optimise for financial metrics rather than for the material conditions of production exhibit the same pattern. Media systems that optimise for engagement metrics rather than for the constraint structure of accurate representation follow suit. In each case the formal structure is identical: optimisation within the projected space at the expense of admissibility in the domain.

24.8 Resolution Rules in Culture

Stable cultures possess an implicit invariant of resolution: a characteristic way of responding to incompatibilities between received form and new context that preserves structural invariants through expansion rather than relaxation.

Proposition 24.18 [Cultural Stability and the Invariant of Resolution]. *A cultural tradition is stable under contextual pressure if and only if it maintains an invariant of resolution $\mathcal{R}$ in the sense of Definition 18.2: it resolves incompatibilities through non-reductive reconciliation, dimensional expansion, and structural necessity.*

Proof. Stability requires that incompatibilities between received form and new context are resolved without loss of structural invariants. By Theorem 14.7, this is possible if and only if the resolution is adaptive in the sense of Definition 14.5, which is equivalent to satisfying the three conditions of the invariant of resolution. □

Corollary 24.19 [Fragility Without Resolution Rule]. *A cultural tradition that lacks a stable invariant of resolution is brittle: it collapses under contextual pressure because it has no structural mechanism for adaptive expansion, and its only remaining response is degenerative drift.*

Proof. Without a resolution rule, incompatibilities can only be resolved by relaxation of admissibility conditions. By Proposition 19.9, this is degenerative. By Theorem 19.4, it is monotonically self-reinforcing. Collapse follows. □

24.9 Conclusion

Cultural continuity is not preservation of content but preservation of constraint structure. A tradition persists when its characteristic structural invariants are maintained across admissible transformations; it degenerates when incompatibilities are resolved by constraint relaxation rather than adaptive expansion; it collapses when its constraint surface is destroyed and no endogenous repair remains.

Myth, language, and institutions are refractive systems when functioning properly: they maintain stable constraint surfaces, apply convergent repair operations, and embody invariants of resolution that enable them to adapt without losing structural identity. When they cease to be refractive, they become archival objects — fixed projections of what was once a living constraint system.

The formal apparatus developed across this monograph applies without modification. The same projection theory that characterises scientific models, the same constraint theory that formalises biological cognition, the same failure theory that distinguishes degenerative drift from phase transitions — all apply directly to the dynamics of cultural transmission. This is not analogy. It is the same formal structure operating in a different domain, which is exactly what a general theory of constraint-preserving identity predicts.

Appendices

Chapter A

Mathematical Preliminaries

This appendix states every mathematical primitive used in the main text, exactly as it is used. It is not a general reference; it is a formal contract. All constructions in Chapters 1–24 are expressed using only the objects defined here.

A.1 Metric and Topological Spaces

Definition A.1 [Metric Space]**.** A *metric space* is a pair (X, d) where X is a set and $d : X \times X \to [0, \infty)$ satisfies:

(i) $d(x, y) = 0 \iff x = y$;

(ii) $d(x, y) = d(y, x)$;

(iii) $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$.

Definition A.2 [Separable Metric Space]**.** (X, d) is *separable* if it contains a countable dense subset.

All experiential domains X in this monograph are assumed separable metric spaces.

Definition A.3 [Distance to a Set]**.** For a non-empty closed set $\Sigma \subset X$, the *distance function* is

$$d(x, \Sigma) := \inf_{y \in \Sigma} d(x, y).$$

This function appears in the Lyapunov conditions of Chapter 17.

Definition A.4 [Open and Closed Sets]**.** A set $U \subseteq X$ is *open* if for every $x \in U$ there exists $\varepsilon > 0$ such that $\{y : d(x, y) < \varepsilon\} \subseteq U$. A set is *closed* if its complement is open.

Definition A.5 [Compactness]**.** A set $K \subseteq X$ is *compact* if every open cover of K has a finite subcover. In metric spaces, K is compact if and only if every sequence in K has a convergent subsequence with limit in K.

Definition A.6 [Induced Topology]**.** Given a constraint structure $(\mathcal{C}, \mathcal{T})$ on X, the *induced topology* on the constraint surface Σ is the subspace topology inherited from (X, d).

A.2 Manifolds and Submanifolds

Definition A.7 [Smooth Manifold]**.** A *smooth manifold* of dimension n is a Hausdorff, second-countable topological space M together with a maximal smooth atlas: a collection of charts $\{\phi_\alpha : U_\alpha \to \mathbb{R}^n\}$ covering M such that transition maps $\phi_\beta \circ \phi_\alpha^{-1}$ are smooth wherever defined.

Definition A.8 [Embedded Submanifold]**.** A subset $\Sigma \subseteq X$ is an *embedded submanifold* of dimension $k < \dim X$ if for every $x \in \Sigma$ there exists a chart (U, ϕ) of X such that $\phi(U \cap \Sigma) = \phi(U) \cap (\mathbb{R}^k \times \{0\})$. We write $\Sigma \hookrightarrow X$.

The constraint surface Σ^* of Chapter 17 is a closed embedded submanifold.

Definition A.9 [Smooth Map and Differential]**.** A map $f : M \to N$ between smooth manifolds is *smooth* if it is smooth in local coordinates. The *differential* $Df_x : T_xM \to T_{f(x)}N$ is the induced linear map on tangent spaces.

A.3 Dynamical Systems

Definition A.10 [Vector Field and Flow]**.** A *smooth vector field* on X is a smooth map $f : X \to TX$. The associated *flow* $\varphi : X \times \mathbb{R} \to X$ satisfies

$$\frac{d}{dt}\varphi(x,t) = f(\varphi(x,t)), \qquad \varphi(x,0) = x.$$

Definition A.11 [Trajectory]**.** A *trajectory* is a curve $\gamma : \mathbb{R} \to X$ satisfying $\dot{\gamma}(t) = f(\gamma(t))$.

Definition A.12 [Invariant Set]**.** A set $A \subseteq X$ is *positively invariant* under the flow if $\varphi(x,t) \in A$ for all $x \in A$ and $t \geq 0$. It is *invariant* if this holds for all $t \in \mathbb{R}$.

A.4 Lyapunov Stability

Let $\Sigma \subset X$ be a closed invariant set.

Definition A.13 [Lyapunov Function]**.** A *Lyapunov function* for Σ is a smooth map $V : X \to [0, \infty)$ satisfying:

(i) $V(x) = 0 \iff x \in \Sigma$;

(ii) $V(x) \geq \alpha(d(x, \Sigma))$ for some class-$\mathcal{K}$ function α;

(iii) $\dot{V}(x) := \langle \nabla V(x), f(x) \rangle \leq 0$ along trajectories.

Definition A.14 [Lyapunov Stability and Asymptotic Stability]**.** Σ is *Lyapunov-stable* if a Lyapunov function exists. It is *asymptotically stable* if additionally $\dot{V}(x) < 0$ for $x \notin \Sigma$.

Definition A.15 [Basin of Attraction]**.** The *basin of attraction* of an asymptotically stable Σ is

$$\mathcal{B}(\Sigma) := \{x \in X : \varphi(x,t) \to \Sigma \text{ as } t \to \infty\}.$$

Theorem A.16 [Lyapunov Stability Theorem]**.** *If a Lyapunov function V exists for Σ, then Σ is Lyapunov-stable. Under the strict descent condition, Σ is asymptotically stable and $\mathcal{B}(\Sigma)$ is open.*

Proof. Standard; see [12]. □

A.5 Gradient Flows

Definition A.17 [Gradient Flow]. Let $E : X \to \mathbb{R}$ be a smooth functional. The *gradient flow* of E is

$$\dot{x}(t) = -\nabla E(x(t)).$$

Proposition A.18 [Energy Descent]. $\frac{d}{dt}E(x(t)) = -\|\nabla E(x(t))\|^2 \leq 0.$

Proof. $\frac{d}{dt}E = \langle \nabla E, \dot{x} \rangle = -\|\nabla E\|^2.$ □

Definition A.19 [Proper and Coercive Functionals]. E is *proper* if sublevel sets $\{E \leq c\}$ are compact. E is *coercive* if $E(x) \to \infty$ as $\|x\| \to \infty$.

Theorem A.20 [Gradient Flow Convergence]. *If E is smooth, proper, and bounded below, and the flow is complete, every trajectory converges to the critical set $\{\nabla E = 0\}$. Under strict convexity, convergence is to the unique global minimum.*

Proof. Properness bounds trajectories in compact sets; energy descent is monotone; completeness gives global existence. The Łojasiewicz inequality yields convergence; see [11]. □

A.6 Category Theory: Minimal

Definition A.21 [Category]. A *category* $\mathcal{C}$ consists of objects, morphisms $\mathrm{Hom}(A, B)$ for each pair of objects, identity morphisms id_A, and associative unital composition.

Definition A.22 [Functor]. A *functor* $F : \mathcal{C} \to \mathcal{D}$ preserves objects, morphisms, identities, and composition.

Definition A.23 [Commutative Diagram]. A diagram *commutes* if all directed paths between any two objects compose to equal morphisms.

The admissible embedding of Chapter 21 (Definition 21.3) is a morphism in the category of constraint spaces, and transformation compatibility is a commutativity condition on the associated square diagram.

A.7 Sheaves: Preview

Full treatment is in Appendix D.

Definition A.24 [Presheaf — Preliminary]. A *presheaf* $\mathcal{F}$ on a topological space X assigns to each open set U a set $\mathcal{F}(U)$ and to each inclusion $V \subseteq U$ a restriction map $\rho_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$, functorially.

Definition A.25 [Sheaf — Preliminary]. A presheaf satisfying locality and gluing is a *sheaf*: compatible local sections assemble into a unique global section.

A.8 Closing Remark

All formal constructions in this monograph are expressed using only the primitives defined in this appendix. No construction in the main text requires mathematical machinery beyond what is stated here or in Appendices B–E.

Chapter B

Information Theory Formalism

This appendix develops the formal account of classical information theory required to support the critique in Chapters 2 and 7, and to establish the precise boundary between Shannon-representable and constraint-theoretic invariants. The goal is not to survey information theory but to prove a specific claim: structural invariants defined over constraint relations cannot, in general, be represented as invariants in a Shannon channel model.

B.1 Shannon Entropy

Let Σ be a finite alphabet and let X be a discrete random variable taking values in Σ with probability distribution $p : \Sigma \to [0, 1]$, $\sum_{s \in \Sigma} p(s) = 1$.

Definition B.1 [Shannon Entropy]**.** The *Shannon entropy* of X is

$$H(X) = -\sum_{s \in \Sigma} p(s) \log p(s),$$

with the convention $0 \log 0 = 0$.

Shannon entropy measures the expected uncertainty of the source. It satisfies $H(X) \geq 0$ with equality if and only if p is a point mass, and $H(X) \leq \log |\Sigma|$ with equality if and only if p is uniform.

Definition B.2 [Joint and Conditional Entropy]**.** For jointly distributed (X, Y):

$$H(X, Y) = -\sum_{s,t} p(s, t) \log p(s, t), \qquad H(X \mid Y) = H(X, Y) - H(Y).$$

Definition B.3 [Mutual Information]**.** The *mutual information* between X and Y is

$$I(X; Y) = H(X) - H(X \mid Y) = H(Y) - H(Y \mid X).$$

Mutual information measures the reduction in uncertainty about X given knowledge of Y, and is symmetric in X and Y.

B.2 Channel Model and Symbol Fidelity

Definition B.4 [Discrete Memoryless Channel]**.** A *discrete memoryless channel* is a conditional distribution $p(y \mid x)$ over output alphabet Σ' given input $x \in \Sigma$. The channel is characterised by its transition matrix $[p(y \mid x)]_{x,y}$.

Definition B.5 [Channel Capacity]**.** The *capacity* of a channel is

$$C = \sup_{p(x)} I(X;Y),$$

the supremum taken over all input distributions $p(x)$.

The fundamental theorem of information theory (Shannon, 1948) states that reliable communication is achievable at any rate $R < C$ and is impossible at any rate $R > C$.

Within this framework, fidelity is defined at the level of symbol sequences.

Definition B.6 [Symbol Fidelity]**.** A channel achieves *symbol fidelity* if the received sequence $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n) \in \Sigma^n$ satisfies $\hat{s} = s$ with high probability. Any transformation $T : \Sigma^n \to \Sigma^n$ with $T(s) \neq s$ is treated as noise and is assigned positive error probability.

Remark B.7. Symbol fidelity is an isomorphism condition: it requires that the mapping from sent to received sequence be the identity. All deviations, including semantically equivalent paraphrases, are classified as errors.

B.3 The Data Processing Inequality

Theorem B.8 [Data Processing Inequality]**.** *Let $X \to Y \to Z$ be a Markov chain. Then*

$$I(X;Z) \leq I(X;Y).$$

Proof. By the Markov condition, X and Z are conditionally independent given Y, so $I(X;Z \mid Y) = 0$. By the chain rule,

$$I(X;Y,Z) = I(X;Y) + I(X;Z \mid Y) = I(X;Y).$$

Since $I(X;Y,Z) \geq I(X;Z)$ by monotonicity, $I(X;Z) \leq I(X;Y)$. □

Corollary B.9 [Information Cannot Be Created by Processing]**.** *No deterministic or stochastic transformation of Y can increase the mutual information between X and the transformed variable beyond $I(X;Y)$.*

This establishes the fundamental limitation of channel models: information about the source is monotonically non-increasing along any processing chain. Whatever structure is lost in transmission cannot be recovered downstream.

B.4 Invariance Under Shannon Channels

We now characterise precisely what kind of invariance the Shannon framework can represent.

Definition B.10 [Shannon-Invariant]**.** A quantity $f(X)$ is *Shannon-invariant* under a channel $p(y \mid x)$ if the distribution of $f(Y)$ equals the distribution of $f(X)$:

$$\mathbb{P}(f(Y) = v) = \mathbb{P}(f(X) = v) \quad \text{for all } v.$$

Proposition B.11 [Shannon Invariance Reduces to Symbol Statistics]**.** *Shannon-invariant quantities are functions of the probability distribution $p(s)$ over Σ. They cannot depend on relational or structural properties of individual symbol configurations.*

Proof. The Shannon framework operates on probability distributions over Σ^n. All channel-invariant quantities are functions of the input distribution $p(x)$ and the channel transition matrix $p(y \mid x)$. Properties of individual sequences $s \in \Sigma^n$ that are not determined by $p(s)$ — in particular, relational properties between sequences — are not representable as Shannon-invariant quantities. □

This proposition identifies the key limitation: the Shannon framework is defined over symbol statistics, not over structural relations between configurations.

B.5 The Shannon Non-Semanticity Theorem

We now state and prove the central result.

Theorem B.12 [Shannon Non-Semanticity]**.** *Let $(\mathcal{C}, \mathcal{T})$ be a constraint structure on a space X, and let $I : X \to \mathbb{R}$ be a structural invariant in the sense of Definition 6.1: $I(T(x)) = I(x)$ for all $T \in \mathcal{T}$. Suppose $\mathcal{T}$ contains non-trivial transformations that change the symbolic representation of states. Then I cannot, in general, be represented as a Shannon-invariant quantity under a channel model with fixed alphabet Σ.*

Proof. We proceed in three steps.

Step 1. Shannon invariance requires a fixed alphabet Σ. All messages are sequences in Σ^n, and all invariants are functions of the probability distribution over Σ^n.

Step 2. A structural invariant I with respect to $(\mathcal{C}, \mathcal{T})$ is defined over the constraint-preserving equivalence class of states:

$$I(x) = I(x') \iff x' = T(x) \text{ for some } T \in \mathcal{T}.$$

By assumption, $\mathcal{T}$ contains transformations that change symbolic representation: there exist $x, x' \in X$ with $x' = T(x)$ for some $T \in \mathcal{T}$, yet with different symbolic encodings $\sigma(x) \neq \sigma(x')$ in Σ^n.

Step 3. A Shannon channel with alphabet Σ classifies $\sigma(x) \neq \sigma(x')$ as distinct messages. The channel assigns positive error probability to the transition $\sigma(x) \to \sigma(x')$, treating it as noise. Therefore the channel cannot represent $I(x) = I(x')$ as an invariance: the symbolic representation changes, and the Shannon framework has no mechanism for identifying two distinct symbol sequences as equivalent under a structural relation. The structural invariant I is therefore not representable as a Shannon-invariant quantity. □

Remark B.13. The proof does not claim that Shannon theory is incorrect within its domain. It establishes a precise boundary: the domain of Shannon theory is symbol statistics; the domain of constraint-theoretic invariance is structural relations. These domains are distinct, and their intersection is exactly the class of structural invariants that happen to coincide with symbol-statistical invariants — a proper subset of all structural invariants.

B.6 Corollary: The Limits of Fidelity-Based Accounts

Corollary B.14 [Semantic Content Cannot Be Reduced to Shannon Fidelity]**.** *Systems whose continuity depends on structural invariance under non-trivial transformations of representation — including mythic transmission, cognitive schemas, cultural traditions, and biological identity — cannot be characterised by Shannon fidelity.*

Proof. Each of these systems instantiates a constraint structure $(\mathcal{C}, \mathcal{T})$ in which the admissible transformations $\mathcal{T}$ include non-trivial changes of symbolic representation (retelling, translation, metabolic renewal, contextual adaptation). By Theorem B.12, the structural invariants of these systems are not representable as Shannon-invariant quantities. Therefore Shannon fidelity is the wrong criterion for their continuity. □

Corollary B.15 [Maximising Fidelity May Destroy Structural Meaning]**.** *A system that maximises Shannon fidelity at the level of symbolic representation may destroy structural invariants that depend on admissible transformation.*

Proof. Shannon fidelity enforces $\hat{s} = s$ with high probability, treating all transformations T with $T(s) \neq s$ as noise. If some such T is an admissible transformation in $\mathcal{T}$ that preserves structural invariants, then enforcing $\hat{s} = s$ blocks T and thereby prevents the realisation of structurally equivalent states. The structural meaning encoded in the equivalence class $[s]_{\mathcal{T}}$ is inaccessible to a fidelity-maximising system. □

These corollaries establish the formal basis for the claims in Chapters 2, 7, 24, and 22: continuity in complex adaptive systems is not a fidelity problem but a constraint-preservation problem, and the two are not only distinct but in tension.

Chapter C

Extended RSVP Equations

This appendix develops the RSVP framework of Chapter 8 at full mathematical precision. We state the field definitions, the coupled evolution system, the coupling structure, stability conditions, and the connection to known PDE classes. The goal is to establish that RSVP is a well-posed dynamical system, not merely a conceptual schema.

C.1 Field Definitions

Let $\Omega \subset \mathbb{R}^n$ be a bounded domain with smooth boundary $\partial\Omega$, and let $T > 0$. We work on the space-time cylinder $\Omega_T = \Omega \times (0, T)$.

Definition C.1 [RSVP State]**.** The *RSVP state* is a triple

$$\mathbf{X}(x,t) = \big(\Phi(x,t),\ \boldsymbol{v}(x,t),\ S(x,t)\big),$$

where:

(i) $\Phi : \Omega_T \to \mathbb{R}$ is the *scalar constraint field*, representing local constraint density;

(ii) $\boldsymbol{v} : \Omega_T \to \mathbb{R}^n$ is the *admissible flow field*, representing the direction and magnitude of admissible evolution;

(iii) $S : \Omega_T \to \mathbb{R}_{\geq 0}$ is the *entropy field*, representing the local degree of variability or admissible drift.

Definition C.2 [Function Spaces]**.** We assume:

$$\begin{aligned}
&\Phi \in H^1(\Omega_T) \cap L^\infty(\Omega_T),\\
&\boldsymbol{v} \in L^2(0,T; H^1_0(\Omega)^n),\\
&S \in H^1(\Omega_T) \cap L^\infty(\Omega_T), \quad S \geq 0.
\end{aligned}$$

Here H^1 denotes the standard Sobolev space of square-integrable functions with square-integrable first derivatives, and H^1_0 enforces zero boundary conditions for the flow field.

C.2 Evolution Equations

Definition C.3 [RSVP System]**.** The RSVP evolution equations are:

$$\partial_t\Phi + \nabla\cdot(\Phi\,\boldsymbol{v}) = D_\Phi\,\Delta\Phi - \alpha_\Phi\,\Phi\,S + \beta_\Phi\,(1-\Phi), \tag{C.1}$$

$$\partial_t\boldsymbol{v} + (\boldsymbol{v}\cdot\nabla)\boldsymbol{v} = -\nabla\Phi - \nu\,\boldsymbol{v} + D_v\,\Delta\boldsymbol{v}, \tag{C.2}$$

$$\partial_t S = D_S\,\Delta S + \gamma\,|\nabla\Phi|^2 - \delta\,\Phi\,S + \sigma, \tag{C.3}$$

on Ω_T, with boundary conditions $\boldsymbol{v} = \mathbf{0}$ and $\nabla\Phi\cdot\mathbf{n} = \nabla S\cdot\mathbf{n} = 0$ on $\partial\Omega$, and initial conditions $(\Phi_0, \boldsymbol{v}_0, S_0)$ specified at $t = 0$. The parameters $D_\Phi, D_v, D_S > 0$ are diffusion coefficients; $\alpha_\Phi, \beta_\Phi, \nu, \gamma, \delta > 0$ and $\sigma \geq 0$ are coupling constants.

C.3 Coupling Structure

The three terms warrant explicit interpretation.

Constraint guidance of flow (Eq. (C.2)). The term $-\nabla\Phi$ in the flow equation acts as a potential gradient, directing $\boldsymbol{v}$ toward regions of lower constraint density. This encodes the principle that admissible flow moves along constraint gradients rather than against them. The term $-\nu\boldsymbol{v}$ provides linear damping, ensuring the flow does not grow without bound in the absence of driving.

Entropy regulation of constraint (Eq. (C.1)). The term $-\alpha_\Phi\,\Phi\,S$ reduces constraint density in regions of high entropy, permitting greater variability where it already exists. The term $\beta_\Phi(1-\Phi)$ drives Φ toward a baseline value of one, encoding a tendency toward restoration of moderate constraint.

Constraint production of entropy (Eq. (C.3)). The term $\gamma|\nabla\Phi|^2$ generates entropy at constraint boundaries — regions of rapid spatial variation in Φ — reflecting the physical principle that constraint gradients produce local disorder. The term $-\delta\,\Phi\,S$ reduces entropy in regions of high constraint, encoding the coherence-preserving role of strong constraint.

Proposition C.4 [Constraint-Flow Interaction]**.** *Under the RSVP dynamics, regions of high Φ channel $\boldsymbol{v}$-flow into low-entropy structured trajectories: the scalar field acts as a geometric constraint on admissible flow paths.*

Proof. The potential $-\nabla\Phi$ in Eq. (C.2) creates a conservative force directed toward low-Φ regions. The damping term $-\nu\boldsymbol{v}$ ensures that flow trajectories remain bounded. In regions where Φ is large and spatially uniform, $|\nabla\Phi|^2$ is small, entropy production is low, and S decreases due to the $-\delta\Phi S$ term. Low entropy implies reduced variability, so $\boldsymbol{v}$-trajectories are constrained to a narrow, structured family. □

C.4 Stability Analysis

Definition C.5 [Equilibrium Configuration]**.** An *equilibrium* of the RSVP system is a time-independent solution $(\Phi^*, \boldsymbol{v}^*, S^*)$ satisfying

$$\begin{aligned} D_\Phi\,\Delta\Phi^* - \alpha_\Phi\,\Phi^*\,S^* + \beta_\Phi\,(1-\Phi^*) &= 0,\\ \nabla\Phi^* + \nu\,\boldsymbol{v}^* - D_v\,\Delta\boldsymbol{v}^* &= \mathbf{0},\\ D_S\,\Delta S^* + \gamma\,|\nabla\Phi^*|^2 - \delta\,\Phi^*\,S^* + \sigma &= 0. \end{aligned}$$

Proposition C.6 [Existence of Uniform Equilibrium]. *The spatially uniform configuration* $(\Phi^*, \mathbf{0}, S^*)$ *with*

$$\Phi^* = \frac{\beta_\Phi}{\beta_\Phi + \alpha_\Phi S^*}, \qquad S^* = \frac{\sigma}{\delta \Phi^*}$$

is an equilibrium. A unique positive solution exists when $\sigma, \delta, \alpha_\Phi, \beta_\Phi > 0$.

Proof. At a spatially uniform equilibrium, all Laplacian terms vanish and $\boldsymbol{v}^* = \mathbf{0}$. The equations reduce to two scalar equations in (Φ^*, S^*), which form a closed system. Existence and uniqueness of a positive solution follow from the monotone structure of the right-hand sides. □

Proposition C.7 [Lyapunov Candidate]. *The functional*

$$\mathcal{V}[\Phi, \boldsymbol{v}, S] = \int_\Omega \left[\frac{(\Phi - \Phi^*)^2}{2} + \frac{|\boldsymbol{v}|^2}{2} + \frac{(S - S^*)^2}{2} \right] dx$$

is a Lyapunov candidate for the equilibrium $(\Phi^*, \mathbf{0}, S^*)$*:* $\mathcal{V} \geq 0$ *with equality if and only if* $(\Phi, \boldsymbol{v}, S) = (\Phi^*, \mathbf{0}, S^*)$.

Proof. Non-negativity is immediate from the squared terms. Equality $\mathcal{V} = 0$ requires $\Phi = \Phi^*$, $\boldsymbol{v} = \mathbf{0}$, and $S = S^*$ almost everywhere. □

Verifying that $\dot{\mathcal{V}} \leq 0$ along trajectories of the RSVP system requires sign conditions on the coupling parameters; this is parameter-dependent and is verified by direct computation in specific regimes.

C.5 Connection to Known PDE Classes

Proposition C.8 [Reaction-Diffusion Structure]. *The RSVP system is of reaction-diffusion type: each field evolves under linear diffusion plus nonlinear reaction terms.*

Proof. Each equation in Definition C.3 has the form $\partial_t u = D\,\Delta u + R(u, \mathbf{u})$, where $D > 0$ is a diffusion coefficient and R is a nonlinear reaction term depending on the full state $\mathbf{u} = (\Phi, \boldsymbol{v}, S)$. This is the standard reaction-diffusion structure; see [11]. □

Proposition C.9 [Gradient Flow Structure of Scalar Fields]. *In the absence of advection (*$\boldsymbol{v} = \mathbf{0}$*), the equations for* Φ *and* S *are gradient flows of an energy functional* $\mathcal{E}[\Phi, S]$.

Proof. Setting $\boldsymbol{v} = \mathbf{0}$ decouples the scalar equations. Define

$$\mathcal{E}[\Phi, S] = \int_\Omega \left[\frac{D_\Phi}{2} |\nabla \Phi|^2 + \frac{\alpha_\Phi}{2} \Phi^2 S - \beta_\Phi \Phi + \frac{D_S}{2} |\nabla S|^2 + \frac{\delta}{2} \Phi S^2 - \sigma S \right] dx.$$

The variational derivatives $\delta\mathcal{E}/\delta\Phi$ and $\delta\mathcal{E}/\delta S$ reproduce the right-hand sides of Eqs. (C.1) and (C.3) up to sign, confirming the gradient flow structure. □

Remark C.10 [Navier–Stokes Analogy]. The flow equation (C.2) is structurally analogous to the incompressible Navier–Stokes equations, with $-\nabla\Phi$ playing the role of a pressure gradient and $D_v\,\Delta\boldsymbol{v}$ providing viscous regularisation. Unlike Navier–Stokes, the RSVP flow is driven by the constraint gradient rather than an externally imposed pressure, and damping $-\nu\boldsymbol{v}$ replaces the incompressibility constraint.

C.6 Constraint Interpretation

We now connect the RSVP system to the abstract constraint framework of Chapter 5.

Proposition C.11 [RSVP as Constraint Structure]**.** *The RSVP system defines a constraint structure* $(\mathcal{C}_{\mathrm{RSVP}}, \mathcal{T}_{\mathrm{RSVP}})$ *where:*

(i) $\mathcal{C}_{\mathrm{RSVP}}$ *consists of all field triples* $(\Phi, \boldsymbol{v}, S)$ *satisfying the equilibrium conditions and the function space regularity of Definition C.3;*

(ii) $\mathcal{T}_{\mathrm{RSVP}}$ *consists of the flows generated by the RSVP dynamics over finite time intervals.*

Admissible trajectories are solutions of the RSVP system with admissible initial data.

Proof. The RSVP dynamics define a continuous semigroup on the appropriate Sobolev spaces, and the flow maps admissible initial data to admissible trajectories by standard PDE existence theory for parabolic systems. □

The structural invariants of the RSVP constraint structure — the conserved and dissipated quantities along admissible trajectories — are those identified by the Lyapunov candidate of Proposition C.7 and the energy functional of the gradient flow structure. These correspond to the abstract structural invariants of Definition 6.1 instantiated in the field-theoretic setting.

Chapter D

Sheaf Cohomology and Obstruction Theory

This appendix develops the sheaf-theoretic and cohomological machinery underlying Chapters 10, 13, 20, and 21. The central goal is to make precise the notion of *obstruction*: the cohomological measure of the failure of locally consistent data to assemble into a globally coherent structure.

D.1 Čech Construction

Let X be a topological space and let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of X.

Definition D.1 [Presheaf]. A *presheaf* $\mathcal{F}$ of abelian groups on X assigns:

(i) to each open set $U \subseteq X$ an abelian group $\mathcal{F}(U)$, whose elements are called *sections* over U;

(ii) to each inclusion $V \subseteq U$ a group homomorphism $\rho_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$, called the *restriction map*, satisfying:

   (a) $\rho_{UU} = \mathrm{id}_{\mathcal{F}(U)}$;
   (b) $\rho_{VW} \circ \rho_{UV} = \rho_{UW}$ for all $W \subseteq V \subseteq U$.

We write $s|_V := \rho_{UV}(s)$ for $s \in \mathcal{F}(U)$.

Definition D.2 [Sheaf]. A presheaf $\mathcal{F}$ is a *sheaf* if it satisfies:

(i) *Locality*: if $U = \bigcup_i U_i$ and $s, t \in \mathcal{F}(U)$ satisfy $s|_{U_i} = t|_{U_i}$ for all i, then $s = t$;

(ii) *Gluing*: if $\{s_i \in \mathcal{F}(U_i)\}$ is a collection satisfying the compatibility condition

$$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j} \quad \text{for all } i, j,$$

then there exists $s \in \mathcal{F}(U)$ such that $s|_{U_i} = s_i$ for all i.

The locality condition ensures that a global section is uniquely determined by its local restrictions. The gluing condition ensures that compatible local sections assemble into a global section. When the gluing condition fails, the failure is measured by the first Čech cohomology group.

D.2 Čech Cochain Complex

Definition D.3 [Čech Cochains]. For a presheaf $\mathcal{F}$ and open cover $\mathcal{U}$, the *Čech cochain groups* are:

$$\check{C}^0(\mathcal{U},\mathcal{F}) = \prod_i \mathcal{F}(U_i),$$
$$\check{C}^1(\mathcal{U},\mathcal{F}) = \prod_{i<j} \mathcal{F}(U_i \cap U_j),$$
$$\check{C}^2(\mathcal{U},\mathcal{F}) = \prod_{i<j<k} \mathcal{F}(U_i \cap U_j \cap U_k).$$

Elements of $\check{C}^p$ are called *p-cochains*.

Definition D.4 [Čech Coboundary Maps]. The *coboundary maps* $\delta^p : \check{C}^p \to \check{C}^{p+1}$ are defined as follows.

For a 0-cochain $f = (f_i) \in \check{C}^0$:

$$(\delta^0 f)_{ij} = f_j|_{U_i\cap U_j} - f_i|_{U_i\cap U_j}.$$

For a 1-cochain $g = (g_{ij}) \in \check{C}^1$:

$$(\delta^1 g)_{ijk} = g_{jk}|_{U_i\cap U_j\cap U_k} - g_{ik}|_{U_i\cap U_j\cap U_k} + g_{ij}|_{U_i\cap U_j\cap U_k}.$$

Proposition D.5 [Cochain Complex]. $\delta^1 \circ \delta^0 = 0$.

Proof. Direct computation shows that $(\delta^1\delta^0 f)_{ijk} = (f_k - f_j) - (f_k - f_i) + (f_j - f_i) = 0$ on each triple intersection. □

The sequence

$$0 \longrightarrow \check{C}^0 \xrightarrow{\delta^0} \check{C}^1 \xrightarrow{\delta^1} \check{C}^2 \longrightarrow \cdots$$

is therefore a cochain complex.

D.3 Čech Cohomology Groups

Definition D.6 [Čech Cohomology]. The *Čech cohomology groups* of $\mathcal{F}$ with respect to $\mathcal{U}$ are:

$$\check{H}^0(\mathcal{U},\mathcal{F}) = \ker\delta^0 = \{(f_i) : f_i|_{U_i\cap U_j} = f_j|_{U_i\cap U_j} \text{ for all } i,j\},$$
$$\check{H}^1(\mathcal{U},\mathcal{F}) = \ker\delta^1 / \operatorname{im}\delta^0.$$

Proposition D.7 [Global Sections]. *$\check{H}^0(\mathcal{U},\mathcal{F}) \cong \mathcal{F}(X)$, the group of global sections.*

Proof. A 0-cocycle is a collection (f_i) satisfying $f_i = f_j$ on all pairwise intersections — precisely the gluing condition for a global section. Since $\mathcal{F}$ is a sheaf, global sections are uniquely determined by such compatible local data. □

D.4 The Obstruction Class

Definition D.8 [Obstruction Class]. Let $\{s_i \in \mathcal{F}(U_i)\}$ be a collection of local sections. Define the *transition functions*

$$g_{ij} = s_j|_{U_i \cap U_j} - s_i|_{U_i \cap U_j} \in \mathcal{F}(U_i \cap U_j).$$

The collection $(g_{ij}) \in \check{C}^1(\mathcal{U}, \mathcal{F})$ is a 1-cocycle (it lies in $\ker \delta^1$), and its cohomology class

$$[\omega] = [(g_{ij})] \in \check{H}^1(\mathcal{U}, \mathcal{F})$$

is the *obstruction class* of the local data $\{s_i\}$.

Proposition D.9 [Cocycle Condition]. *(g_{ij}) is a 1-cocycle: $\delta^1(g_{ij}) = 0$.*

Proof. On a triple intersection $U_i \cap U_j \cap U_k$:

$$g_{jk} - g_{ik} + g_{ij} = (s_k - s_j) - (s_k - s_i) + (s_j - s_i) = 0.$$

□

Theorem D.10 [Obstruction Criterion]. *The local sections $\{s_i \in \mathcal{F}(U_i)\}$ admit a global section $s \in \mathcal{F}(X)$ with $s|_{U_i} = s_i$ for all i if and only if the obstruction class vanishes: $[\omega] = 0$ in $\check{H}^1(\mathcal{U}, \mathcal{F})$.*

Proof. $[\omega] = 0$ means $(g_{ij}) = \delta^0(h_i)$ for some $(h_i) \in \check{C}^0$, i.e., $g_{ij} = h_j - h_i$ on each $U_i \cap U_j$. Then the sections $s'_i = s_i - h_i$ satisfy $s'_i|_{U_i \cap U_j} = s'_j|_{U_i \cap U_j}$ for all i, j — they are compatible — and by the sheaf gluing condition they assemble into a global section s' with $s'|_{U_i} = s'_i$. Conversely, if a global section exists, (g_{ij}) is the coboundary of its restrictions, so $[\omega] = 0$. □

This theorem is the formal content of the Yarncrawler framework: global reconstruction is possible if and only if the cohomological obstruction vanishes.

D.5 Connection to Chapter 13: Incompatibility as Obstruction

The cross-framework incompatibilities of Chapter 13 have a direct cohomological interpretation.

Proposition D.11 [Incompatibility as Non-Vanishing Obstruction]. *A cross-framework constraint incompatibility between frameworks $(\mathcal{C}_i, \mathcal{T}_i)$ and $(\mathcal{C}_j, \mathcal{T}_j)$ on overlapping domains corresponds to a non-vanishing transition function $g_{ij} \neq 0$ in the Čech 1-cochain of the associated sheaf. When such incompatibilities cannot be resolved locally, the obstruction class $[\omega] \neq 0$ in $\check{H}^1$.*

Proof. Each framework defines a local section of the sheaf of admissible constraint structures. On the overlap where both frameworks apply, the sections disagree: the transition function g_{ij} measures this disagreement. If the incompatibility is irresolvable by local adjustment — if no local modification of the sections makes them agree on overlaps — then (g_{ij}) is not a coboundary, and the class $[\omega] \neq 0$. □

Corollary D.12 [Expansion as Obstruction Trivialisation]. *Adaptive expansion of the constraint structure in the sense of Theorem 14.7 corresponds to passage to a refined cover $\mathcal{U}'$ of a larger space X' in which the pullback of the obstruction class $[\omega]$ vanishes: $[\omega'] = 0$ in $\check{H}^1(\mathcal{U}', \mathcal{F}')$.*

Proof. An admissible medium shift (Definition 21.4) lifts the system to X' with a refined cover such that the transition functions become coboundaries in the new complex. By Theorem D.10, this is equivalent to the existence of a global section in the lifted framework. □

D.6 Connection to TARTAN: Local Consistency Is Not Global Coherence

Proposition D.13 [TARTAN Obstruction]. *A TARTAN tiling satisfying inter-scale consistency (Definition 9.3) defines a Čech* 0*-cochain, but trajectory consistency (Definition 9.2) corresponds to the vanishing of the obstruction class, which may fail even when inter-scale consistency holds.*

Proof. Inter-scale consistency ensures that local states agree within each scale — this is the data of a 0-cochain. However, agreement on pairwise overlaps across scales is an additional condition: it requires the transition functions g_{ij} to vanish, which is the cohomological condition. Inter-scale consistency alone does not imply this: the tiling may be locally valid at each scale while the g_{ij} form a non-trivial cocycle, preventing global reconstruction. □

D.7 Phase Transitions and Obstruction

Proposition D.14 [Phase Transition as Persistent Obstruction]. *A phase transition in the sense of Chapter 20 generates a non-trivial obstruction class* $[\omega] \in \check{H}^1(\mathcal{U}, \mathcal{F})$ *that cannot be trivialised by any local modification of the cover* $\mathcal{U}$*.*

Proof. By Theorem 20.6, the failure at a phase transition is global: no sequence of local repairs converges to a global section. Equivalently, the obstruction class is not in the image of any coboundary δ^0 with respect to any local modification of the 0-cochains. Therefore $[\omega] \neq 0$ and cannot be trivialised within the current framework. □

Corollary D.15 [Medium Shift as Cover Refinement]. *An admissible medium shift corresponds to passing to a refined cover* $\mathcal{U}'$ *of a strictly larger space* X' *such that the induced obstruction class satisfies* $[\omega'] = 0$ *in* $\check{H}^1(\mathcal{U}', \mathcal{F}')$*.*

Proof. By Definition 21.4(iii), the obstruction class vanishes in the lifted system. This is precisely the condition $[\omega'] = 0$, which by Theorem D.10 guarantees the existence of a global section in the lifted framework. □

The sheaf-theoretic picture therefore unifies the framework of Chapters 10–21 into a single cohomological narrative: knowledge is reconstructibility, incompatibility is obstruction, expansion is trivialisation, and medium shifts are the passage to covers in which previously non-trivial classes vanish.

Chapter E

Implementation Sketches

This appendix provides pseudocode and computational sketches for the core algorithmic components of the monograph. The goal is not a production implementation but a precise operational description: each sketch should be sufficient to specify what the algorithm does, what it requires, and what failure looks like. All components are tied explicitly to the formal structures of the main text.

E.1 Spherepop Log Validation

The Spherepop event calculus of Chapter 11 requires that every event in a log be admissible with respect to the constraint relation $\mathcal{R} \subseteq S \times S$.

Algorithm (Spherepop Log Validation)

Input: Event log $\mathcal{L} = (e_1, \ldots, e_n)$, constraint relation $\mathcal{R}$, initial state s_1.

Output: VALID if $\mathcal{L}$ is admissible; INVALID(k) at first violation.

Set $s \leftarrow s_1$. For $k = 1$ to n: extract $e_k = (s \to s')$. If $(s, s') \notin \mathcal{R}$, return INVALID($k$). Set $s \leftarrow s'$. Return VALID.

Failure mode. An INVALID(k) return corresponds to the failure of coherence described in Proposition 11.15: the log contains an illegal event that breaks admissibility and prevents reliable reconstruction.

Complexity. $O(n \cdot |\mathcal{R}|)$ for a naive membership check; $O(n \log |\mathcal{R}|)$ if $\mathcal{R}$ is indexed.

Relationship to CLIO. When the validator returns INVALID(k), the CLIO operator (Section E.2) should be invoked to attempt repair of the constraint relation, not to modify the log. The log is the immutable record; the constraint structure is what CLIO adjusts.

E.2 CLIO Gradient Flow Discretisation

The CLIO operator of Chapter 12 is defined as a gradient flow on the space of constraint structures: $\dot{\mathcal{C}} = -\nabla_{\mathcal{C}} \mathcal{E}(\mathcal{C})$. In practice, the constraint structure is represented by a finite set of parameters $\theta \in \mathbb{R}^d$, and the coherence energy $\mathcal{E}(\theta)$ is a differentiable function of these parameters.

Algorithm (CLIO Gradient Descent)

Input: Initial parameters θ_0; energy $\mathcal{E}(\theta)$ with gradient $\nabla_\theta \mathcal{E}$; step sizes $\{\eta_t\}$; tolerance $\varepsilon > 0$; maximum iterations $T_{\max}$.

Output: Fixed point θ^* or NonConvergent.

Set $\theta \leftarrow \theta_0$, $t \leftarrow 0$. Repeat: $g \leftarrow \nabla_\theta \mathcal{E}(\theta)$; $\theta \leftarrow \theta - \eta_t g$; $t \leftarrow t + 1$. Until $\|g\| < \varepsilon$ or $t \geq T_{\max}$. If $\|g\| < \varepsilon$ return θ^*, else return NonConvergent.

Coherence energy. The energy $\mathcal{E}(\theta)$ should aggregate:

(i) within-framework incoherence: violations of admissibility conditions internal to each framework;

(ii) cross-framework incompatibility: states admissible in one framework but inadmissible in another, weighted by overlap.

A simple instantiation: $\mathcal{E}(\theta) = \sum_i w_i \cdot \mathrm{viol}_i(\theta)$ where viol_i counts constraint violations in framework i.

Convergence criteria. Proposition 12.4 guarantees convergence when $\mathcal{E}$ is proper and bounded below and the flow is complete. In practice, monitor:

(i) gradient norm $\|g\|$ (approaches zero at fixed point);

(ii) energy decrease per step (slowing signals approach to framework limit);

(iii) step count to convergence (increasing cost is a diagnostic indicator per Proposition 20.8).

Non-convergence. A NonConvergent return after $T_{\max}$ steps indicates that CLIO cannot find a fixed point within the current framework — the phase transition regime of Chapter 20. The appropriate response is not to increase $T_{\max}$ but to trigger a medium shift (Section E.4).

E.3 TARTAN Recursive Tiling

The TARTAN framework of Chapter 9 decomposes a domain into a hierarchy of tiles and enforces trajectory consistency across scales.

Algorithm (TARTAN Tiling with Trajectory Consistency)

Input: Domain Ω; depth K; local state $s : \Omega \to Y$; trajectory data $\{x_t\}$; tolerance $\varepsilon_{\mathrm{traj}}$.

Output: Recursive tiling $\{U_i^{(k)}\}$ with consistency flags.

Set $\mathcal{T}^{(0)} \leftarrow \{\Omega\}$. For $k = 1$ to K: subdivide each $U \in \mathcal{T}^{(k-1)}$ into subtiles $\{V_j\}$; check pairwise overlap consistency within $\varepsilon_{\mathrm{traj}}$; flag inconsistent pairs; check inter-scale consistency. Compute Yarncrawler obstruction class from inconsistencies. Return tiling with consistency report.

Local-global gap. Overlap consistency within each level does not guarantee a vanishing obstruction class (Proposition D.13). The Yarncrawler check in the final step is essential for detecting global incoherence that local checks miss.

E.4 Diagnostics and Phase Transition Detection

The three diagnostic indicators of Proposition 20.8 can be monitored concurrently with the CLIO run.

Algorithm (Phase Transition Diagnostics)

Input: CLIO run history $\{(\theta_t, g_t, \mathcal{E}(\theta_t))\}_{t=1}^{T}$; obstruction computation oracle.

Output: One of: CONVERGING, SLOWINGCONVERGENCE, PHASETRANSITIONINDICATED.

Compute repair cost $\bar{t}$ and convergence rate $\rho = \|g_T\|/\|g_{T/2}\|$. Compute $[\omega]$ from TARTAN output. If $\|g_T\| < \varepsilon$: return CONVERGING. If $\rho > \rho_{\max}$ or $[\omega] \neq 0$: return PHASETRANSITIONINDICATED. Else: return SLOWINGCONVERGENCE.

Degenerative drift detection. Separately, monitor discriminative power $\mathrm{Disc}(\mathcal{C}_t)$ over time. A monotonically decreasing sequence with no corresponding decrease in $\mathcal{E}$ signals degenerative drift (Chapter 19) rather than healthy convergence: the constraint structure is weakening, not improving.

E.5 Simulation Notes

RSVP numerical stability. The RSVP system (Appendix C) is stiff when diffusion coefficients D_Φ, D_S are small relative to coupling terms. Use implicit time-stepping (Crank–Nicolson or backward Euler) for the diffusion terms and explicit treatment of the nonlinear reaction terms. Monitor the Courant–Friedrichs–Lewy condition for the advection term in Eq. (C.2).

CLIO step size. Use adaptive step sizes: reduce η_t when $\mathcal{E}(\theta_t)$ increases (line search) and increase when convergence is fast and stable. A simple rule: $\eta_{t+1} = \eta_t \cdot 0.9$ after any step that increases $\mathcal{E}$.

Obstruction computation. For small covers, the Čech obstruction class $[\omega]$ can be computed directly by solving the linear system arising from the coboundary conditions of Appendix D. For large covers, approximate methods based on persistent homology are available; see [10] for the algebraic foundation.

Visualisation. For RSVP fields, plot Φ, $|\boldsymbol{v}|$, and S as heatmaps over Ω at successive time steps. The constraint surface Σ^* can be visualised as the level set $\{V = 0\}$ of the Lyapunov function. For TARTAN tilings, use a quad-tree or oct-tree rendering with colour coding for consistency flags. For CLIO convergence, plot $\mathcal{E}(\theta_t)$ and $\|g_t\|$ against iteration count on a logarithmic scale; linear decay of $\log \|g_t\|$ indicates geometric convergence.

Chapter F

Algebraic-Geometric Grounding: K-Trivial Fibrations and Boundedness

This appendix develops the algebraic-geometric foundations underlying the fibration theory of Chapter 15, §§15.3–15.6. It states the main definitions and results from the geometry of K-trivial varieties in a form that makes the correspondence with the constraint-theoretic framework explicit. References are given to the primary literature throughout.

F.1 K-Trivial Varieties and Their Fibrations

Definition F.1 [K-Trivial Variety]**.** A smooth projective variety $\mathcal{X}$ over $\mathbb{C}$ is *K-trivial* if its canonical bundle $K_{\mathcal{X}}$ is numerically trivial:

$$K_{\mathcal{X}} \equiv 0.$$

The three principal classes are:

(i) *Calabi–Yau varieties*: $K_{\mathcal{X}} \sim_{\mathbb{Q}} 0$ with $h^{i,0}(\mathcal{X}) = 0$ for $0 < i < \dim \mathcal{X}$;

(ii) *Abelian varieties*: complex tori $\mathcal{X} \cong \mathbb{C}^n/\Lambda$ with $K_{\mathcal{X}}$ trivial;

(iii) *Irreducible holomorphic symplectic varieties* (hyperkähler): $\mathcal{X}$ admits a holomorphic symplectic form $\sigma \in H^0(\mathcal{X}, \Omega^2_{\mathcal{X}})$ and $h^{2,0} = 1$.

K-trivial varieties are by themselves wildly unbounded in general: there is no bound on their topology or geometry that holds across all examples in each class. The central observation of [15] is that imposing a fibration structure introduces hidden constraints that collapse this freedom.

Definition F.2 [Admissible Fibration]**.** A *fibration* on a K-trivial variety $\mathcal{X}$ is a surjective morphism with connected fibers,

$$f : \mathcal{X} \to Y,$$

where Y is a normal projective variety with $0 < \dim Y < \dim \mathcal{X}$, and the general fiber $\mathcal{X}_y = f^{-1}(y)$ is a smooth K-trivial variety of the same type (abelian or symplectic).

F.2 The Canonical Bundle Formula

The fundamental tool for analysing fibrations of K-trivial varieties is the canonical bundle formula of Kawamata, Kodaira, and their successors.

Theorem F.3 [Canonical Bundle Formula [18, 19]]. *Let $f : \mathcal{X} \to Y$ be an admissible fibration of K-trivial varieties. Then there exists an effective $\mathbb{Q}$-divisor B_Y and a nef $\mathbb{Q}$-divisor M_Y on Y such that*

$$K_{\mathcal{X}} \sim_{\mathbb{Q}} f^*(K_Y + B_Y + M_Y)\,.$$

The divisors B_Y and M_Y have precise geometric meanings.

Definition F.4 [Boundary Divisor]. The *boundary divisor* $B_Y = \sum_P b_P \cdot P$ is an effective $\mathbb{Q}$-divisor on Y supported on the discriminant locus of f: the set of points $y \in Y$ over which the fiber $\mathcal{X}_y$ is singular or reducible. The coefficient b_P encodes the severity of degeneration over P.

Definition F.5 [Moduli Divisor]. The *moduli divisor* M_Y is a nef $\mathbb{Q}$-divisor on Y encoding the variation of fiber type across Y. It vanishes identically if and only if all smooth fibers of f are isomorphic (isotrivial fibration). In general, $M_Y = \mathcal{P}^*\lambda$ where $\mathcal{P} : Y \dashrightarrow \mathcal{M}$ is the period map into a moduli space of fiber types and λ is the Hodge bundle on $\mathcal{M}$.

Remark F.6. The correspondence with the canonical decomposition of Definition 15.6 is exact: $K_Y \leftrightarrow \Phi_Y$ (base constraint field), $B_Y \leftrightarrow \Phi_{\text{sing}}$ (defect field), $M_Y \leftrightarrow \Phi_{\text{var}}$ (variation memory field).

F.3 Variation of Hodge Structure and Period Maps

The moduli divisor M_Y is controlled by the variation of Hodge structure on the cohomology of the fibers.

Definition F.7 [Variation of Hodge Structure]. Let $f : \mathcal{X} \to Y$ be an admissible fibration. The *variation of Hodge structure* (VHS) associated to f is the local system $\mathbb{H} = R^k f_*\mathbb{Z}$ on the smooth locus Y° of Y, equipped with:

(i) the Hodge filtration $F^\bullet\mathcal{H}$ on the corresponding flat bundle $\mathcal{H} = \mathbb{H} \otimes_{\mathbb{Z}} \mathcal{O}_{Y^\circ}$;

(ii) the flat Gauss–Manin connection $\nabla : \mathcal{H} \to \mathcal{H} \otimes \Omega^1_{Y^\circ}$;

(iii) Griffiths transversality: $\nabla(F^p\mathcal{H}) \subseteq F^{p-1}\mathcal{H} \otimes \Omega^1$.

Definition F.8 [Period Map]. The *period map* associated to the VHS is the holomorphic map (defined on Y° and extending meromorphically to Y):

$$\mathcal{P} : Y^\circ \to \Gamma\backslash D,$$

where D is the period domain classifying Hodge structures of the given type, and Γ is the monodromy group.

The period map encodes all variation data: two fibers $\mathcal{X}_{y_1}$ and $\mathcal{X}_{y_2}$ have isomorphic Hodge structures if and only if $\mathcal{P}(y_1) = \mathcal{P}(y_2)$. The moduli divisor M_Y is the pullback under $\mathcal{P}$ of the natural ample bundle on the Baily–Borel compactification $\overline{\Gamma\backslash D}$ [26, 27].

Proposition F.9 [Boundedness of Period Map Image]. *If the base Y lies in a bounded family, then the image $\mathcal{P}(Y^\circ)$ in $\Gamma\backslash D$ is bounded. Consequently, M_Y lies in a bounded linear system.*

Proof. The period map is a holomorphic map from a bounded family of varieties into the arithmetic quotient $\Gamma\backslash D$, which has a canonical compactification with controlled geometry [27, 28]. Boundedness of Y implies the image is bounded. □

F.4 The Albanese Reduction

For abelian fibrations, the Albanese variety provides a canonical linearisation.

Definition F.10 [Albanese Fibration]. Let $f : \mathcal{X} \to Y$ be an admissible fibration with abelian fibers. The *Albanese fibration*

$$f_{\mathrm{Alb}} : \mathcal{X}_{\mathrm{Alb}} \to Y$$

is the universal abelian scheme over Y equipped with a map $\mathcal{X} \to \mathcal{X}_{\mathrm{Alb}}$ over Y. It is the *linearised* version of f: the minimal transport structure consistent with the variation data.

Remark F.11. In the RSVP language: f_{Alb} is the linear transport backbone $\boldsymbol{v}_{\mathrm{lin}}$, while the original f is the full nonlinear dynamics. The difference $\mathcal{X}/\mathcal{X}_{\mathrm{Alb}}$ measures higher-order structure not captured by linear transport.

F.5 Tate–Shafarevich Groups

Definition F.12 [Tate–Shafarevich Group]. Let $f_{\mathrm{Alb}} : A \to Y$ be an abelian scheme over Y. The *Tate–Shafarevich group* is

$$\mathrm{Sha}(A/Y) \ := \ H^1(Y_{\text{ét}},\, A),$$

the first étale cohomology group with coefficients in the group scheme A. It classifies A-torsors over Y that are locally (in the étale topology) isomorphic to A but not globally trivial.

Proposition F.13 [Finiteness of the Tate–Shafarevich Group]. *In the Calabi–Yau case (abelian fibration with trivial canonical bundle),* $\mathrm{Sha}(A/Y)$ *is a finite abelian group. In the symplectic case (Lagrangian fibration),* $\mathrm{Sha}(A/Y)$ *fits into an exact sequence*

$$0 \to \mathcal{D}_Y \to \mathrm{Sha}(A/Y) \to \mathcal{T}_Y \to 0$$

with $\mathcal{T}_Y$ finite and $\mathcal{D}_Y \cong \mathbb{C}^n$ for some $n \geq 0$.

Proof. The Calabi–Yau finiteness follows from the Zarhin trick [25] and finiteness results for abelian varieties over function fields [23, 24]. The symplectic decomposition follows from the structure of the Tate–Shafarevich group for Lagrangian fibrations on irreducible holomorphic symplectic manifolds, where the continuous part $\mathcal{D}_Y$ arises from the Beauville–Bogomolov form [20, 21]. □

The finiteness (or controlled decomposition) of $\mathrm{Sha}(A/Y)$ is the third and final boundedness input in the proof of Theorem 15.20.

F.6 The Main Boundedness Theorem

We state the result of [15] in full generality, then record the constraint-theoretic interpretation.

Theorem F.14 [Boundedness of Fibered K-Trivial Varieties [15]]**.** *Let $\mathcal{S}$ be the set of all K-trivial varieties $\mathcal{X}$ that admit a fibration $f : \mathcal{X} \to Y$ with:*

(i) Y a variety with klt singularities and $-(K_Y + B_Y + M_Y)$ ample;

(ii) general fibers abelian varieties or irreducible holomorphic symplectic varieties of fixed Hodge numbers.

Then $\mathcal{S}$ is a bounded family*: there exists a scheme of finite type $\mathcal{T}$ and a morphism $\mathcal{U} \to \mathcal{T}$ such that every element of $\mathcal{S}$ appears as a fiber of $\mathcal{U} \to \mathcal{T}$.*

Remark F.15 [Constraint-Theoretic Reading]**.** Theorem F.14 is the precise algebraic-geometric content of Theorem 15.20. The three conditions on Y, the fibers, and the moduli data correspond to the three boundedness conditions (i)–(iii) of Theorem 15.20. The conclusion that $\mathcal{S}$ is a bounded family is the algebraic-geometric formulation of the statement that constraint closure implies bounded complexity.

F.7 Correspondence Table

We summarise the correspondence between the algebraic-geometric and constraint-theoretic languages.

Algebraic geometry	**Constraint theory**
Total space $\mathcal{X}$	Full field configuration
Base Y	Base trajectory manifold
Fiber $\mathcal{X}_y$	Local state space
Canonical bundle $K_{\mathcal{X}}$	Global constraint density $\Phi_{\mathcal{X}}$
Base term K_Y	Base constraint field Φ_Y
Boundary divisor B_Y	Defect field Φ_{sing}
Moduli divisor M_Y	Variation memory field Φ_{var}
Period map $\mathcal{P}$	Projection to latent variation space
Albanese fibration	Linear transport backbone
Tate–Shafarevich group Sha_Y	Gluing obstruction group $H^1(Y, \mathcal{P})$
Torsion part of Sha_Y	Discrete obstruction $\mathcal{O}_{\text{torsion}}$
Continuous part of Sha_Y	Continuous obstruction $\mathcal{O}_{\text{continuous}}$
Bounded family	Constraint closure

The correspondence is not merely terminological: each algebraic-geometric result translates to a constraint-theoretic theorem, and the proofs in both languages follow the same logical structure. The canonical bundle formula is a field decomposition law; the period map is a projection operator; the Tate–Shafarevich exact sequence is a constraint pipeline; and the main boundedness theorem is constraint closure.

Bibliography

[1] E. Thompson, A. Frank, and M. Gleiser, *The Blind Spot: Why Science Cannot Ignore Human Experience*, MIT Press, 2023.

[2] A. N. Whitehead, *Science and the Modern World*, Macmillan, 1925.

[3] A. N. Whitehead, *Process and Reality*, Macmillan, 1929.

[4] C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[5] I. Prigogine and I. Stengers, *Order Out of Chaos*, Bantam Books, 1984.

[6] F. J. Varela, E. Thompson, and E. Rosch, *The Embodied Mind*, MIT Press, 1991.

[7] M. Merleau-Ponty, *Phenomenology of Perception*, Gallimard, 1945.

[8] H. Bergson, *Time and Free Will*, 1889.

[9] S. Mac Lane, *Categories for the Working Mathematician*, Springer, 1998.

[10] G. E. Bredon, *Sheaf Theory*, Springer, 1997.

[11] L. C. Evans, *Partial Differential Equations*, AMS, 2010.

[12] S. H. Strogatz, *Nonlinear Dynamics and Chaos*, Westview Press, 2015.

[13] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, 2006.

[14] D. J. Chalmers, *The Conscious Mind*, Oxford University Press, 1996.

[15] P. Engel, S. Filipazzi, F. Greer, M. Mauri, and R. Svaldi, "Boundedness of some fibered K-trivial varieties," arXiv:2507.00973, 2025.

[16] C. Birkar, "Log Calabi–Yau fibrations," arXiv:1811.10709, 2018.

[17] C. Hacon, J. McKernan, and C. Xu, "ACC for log canonical thresholds," *Ann. of Math.* 180 (2014), 523–571.

[18] O. Fujino, *Foundations of the Minimal Model Program*, MSJ Memoirs, 35, 2018.

[19] Y. Prokhorov and V. Shokurov, "Towards the second main theorem on complements," *J. Algebraic Geom.* 18 (2009), 151–199.

[20] A. Beauville, "Variétés Kählériennes dont la première classe de Chern est nulle," *J. Differential Geom.* 18 (1983), 755–782.

[21] D. Huybrechts, "Compact hyperkähler manifolds: basic results," *Invent. Math.* 135 (1999), 63–113.

[22] M. Gross, "A finiteness theorem for elliptic Calabi–Yau threefolds," *Duke Math. J.* 74 (1994), 271–299.

[23] I. Dolgachev and M. Gross, "Elliptic threefolds I: Ogg–Shafarevich theory," *J. Algebraic Geom.* 3 (1994), 39–80.

[24] G. Faltings, "Endlichkeitssätze für abelsche Varietäten über Zahlkörpern," *Invent. Math.* 73 (1983), 349–366.

[25] Y. Zarhin, "Endomorphisms of abelian varieties and points of finite order in characteristic p," *Math. USSR Izv.* 11 (1977), 31–45.

[26] P. Deligne, "Théorie de Hodge II," *Publ. Math. IHÉS* 40 (1971), 5–57.

[27] P. Griffiths, "Periods of integrals on algebraic manifolds I," *Amer. J. Math.* 90 (1968), 568–626.

[28] W. Schmid, "Variation of Hodge structure: the singularities of the period mapping," *Invent. Math.* 22 (1973), 211–319.

[29] D. Matsushita, "On fibre space structures of a projective irreducible symplectic manifold," *Topology* 38 (1999), 79–83.

[30] F. Campana, "Orbifolds, special varieties and classification theory," *Ann. Inst. Fourier* 54 (2004), 499–630.