

SELECTION WITHOUT DESIGN

Constraint, Projection, and the Geometry of the Unseen

Flyxion

Independent Researcher

April 29, 2026

ABSTRACT

What appears as optimization in biology, intrinsic function in philosophy of mind, and cumulative knowledge in epistemology are all consequences of a single deeper structure: the persistence of closed systems under constraint, combined with the irreducibly partial character of observation. This essay argues in three movements. The first establishes that natural selection is not an optimizer but a filter, and that the appearance of design in living systems is an artifact of conditioning on survival within a constraint field rather than evidence of any directed process. The second applies this conclusion to the philosophy of function: functional attribution requires a privileged coordinate over system behavior that the underlying dynamics never supply, and therefore both teleological and computational functionalism rest on an ungrounded projection rather than an intrinsic property of physical systems. The third extends the same structure to observation itself, showing that any observer is a constrained subsystem whose reachable set is a proper subset of the full state space, that invariant regions may exist beyond every reachable set, and that such inaccessible regions deform the geometry of what is accessible in ways that are in principle detectable as epistemic curvature. The essay concludes that being and knowing are not independent domains but dual projections of a single constraint-governed dynamical field, and that any claim about purpose, progress, or knowledge must first specify which constraints, which projection, and which reachable set it presupposes.

CONTENTS

Opening Frame	6
Part I: Selection Without Optimization	7
1 The Narrative of Progress and Its Presuppositions	7
2 Selection as a Sieve: The Formal Structure of Persistence	9
3 Survivor Bias and the Illusion of Design	11
4 The Origin of Life Reframed	13
Part II: Function as Projection	15
5 The Structure of Functional Explanation	15
6 The Bridge: From Persistence to the Relativity of Function	15
7 Teleological Functionalism Dissolved	17
8 Computational Functionalism and the Problem of Equivalence	18
9 Functional Descriptions as Coordinate Systems	20
Part III: Observation as Constrained Access	22
10 The Observer Within the Field	22
11 Reachability and the Limits of Access	23
12 Islands of Persistence	25
13 Epistemic Curvature	26
14 The Lifting Problem and Structural Underdetermination	28
Part IV: The Unity of Being and Knowing	30
14 The Collapse of the Distinction	30
15 Physics Without Privileged Observation	31

16 Epistemology Without Privileged Reality	32
17 Final Synthesis: Constraint as the Common Substrate	34
18 Closing Reflection	35
Part V: Consequences and Extensions	36
18 Constraint Without Teleology	36
19 Hierarchies of Closure and the Emergence of Complexity	37
20 Projection, Compression, and the Limits of Representation	38
21 Temporal Asymmetry and the Direction of Inference	39
22 Constraint Geometry and Physical Law	40
23 Universality of the Constraint-Projection Structure	41
24 Closing Extension	43
Part VI: RSVP Extension	44
24 From Abstract Constraint to Structured Field	44
25 Persistence as Field Coherence	44
26 Selection as Entropic Filtering in the Field	45
27 Function as Projection in RSVP Fields	46
28 Observation as Field-Constrained Trajectory	47
29 Islands and Field Disconnection	47
30 RSVP Unification	48
Part VII: Against the Optimization Interpretation	49
32 The Optimization Assumption and Its Requirements	49
33 Persistence Does Not Imply Ascent	50

34 The Absence of a Canonical Fitness Function	50
35 Non-Transitivity and the Impossibility of a Scalar Objective	51
36 Fitness Landscapes as Coordinate Artifacts	52
37 When Optimization Emerges as an Approximation	52
Part VIII: Selection, Obstruction, and the Sheaf of Coherent Models	53
37 Observations as Local Sections	53
38 The Obstruction Class and Its Meaning	54
39 Selection as Vanishing Obstruction	55
40 The Selection–Projection Theorem	56
41 A Variational Formulation of Selection	57
42 Noether-Type Invariants and the Structure of Persistence	58
Part IX: Photographic Contradiction and Multimodal Constraint	59
44 Photographs as Partial Projections of the Constraint Field	59
45 Universal Photographic Inconsistency	60
46 Contradiction as Entropic Obstruction	60
47 Hinge Observations and the Collapse of Apparent Coherence	61
48 Deferred Constraint Closure and the Topology of Learning	62
49 Before and After as Boundary Conditions on Trajectories	63
Part X: Multimodal Constraints and the Amplification of Obstruction	64
50 Multiple Sensor Modalities as Independent Projections	64
51 Accelerometric and Acoustic Obstruction	65
52 Temporal Alignment as an Additional Constraint	66

53 Dense Versus Sparse Observation and the Refuge of Unobserved Intervals	66
54 The Statistical Field and Entropy as Obstruction Density	68
55 Uncertainty, Incompatibility, and the Information-Theoretic Limit	69
Coda: The Geometry of What Cannot Be Seen	70
57 The Complete Picture	70
58 What the Framework Does Not Say	71
59 The Final Inversion	71

OPENING FRAME

There is a story about life that is so familiar it feels like observation rather than interpretation. Molecules assembled, reactions stabilized, replication emerged, and through the slow action of selection, organisms became increasingly well-suited to their environments. Eyes formed because seeing is useful. Enzymes became efficient because efficiency is rewarded. Cognition deepened because accurate prediction confers advantage. The arc of this narrative bends toward improvement, and improvement implies direction, and direction implies that something was being optimized.

The same story recurs at finer scales. An organ is said to be *for* its function: the heart for pumping blood, the kidney for filtering waste, the prefrontal cortex for executive control. Philosophy of mind formalizes this intuition into functionalism, the view that mental states are defined by their causal roles rather than their physical substrate. What matters is not what something is made of but what it does, and what it does is fixed by its functional description. This seems both rigorous and liberal: it accommodates multiple realizability, grounds cognition in structure rather than matter, and aligns with how both biologists and engineers talk about systems.

At the broadest scale, the story becomes epistemological. Knowledge accumulates. Science progresses. Each generation of inquiry builds on the last, refining its models, eliminating error, converging toward truth. The observer stands at one end of this process, taking in the world and constructing representations that improve over time.

This essay argues that all three narratives rest on the same suppressed assumption, and that assumption is false. The assumption is that there exists, internal to the dynamics of the relevant system, a gradient or ranking over outcomes that licenses the language of optimization, function, and progress. The argument proceeds in three movements corresponding to the three narratives, but its underlying claim is unified: what appears as direction is the shadow of constraint; what appears as purpose is the artifact of projection; what appears as knowledge is the trace of a path through a space far larger than any observer can reach.

The machinery required to make this precise is that of dynamical systems theory, elementary topology, and the theory of fiber bundles. None of it is technically deep, but all of it is necessary. The formal expressions that appear in what follows are not decoration. They are the points at which the argument

could go another way and instead becomes fixed.

PART I SELECTION WITHOUT OPTIMIZATION

1. THE NARRATIVE OF PROGRESS AND ITS PRESUPPOSITIONS

The dominant interpretation of biological evolution treats natural selection as a process of progressive refinement. Populations vary, variants differ in reproductive success, and over time the distribution of traits shifts toward those that are better adapted. In the language of optimization theory, selection acts on a fitness landscape, and populations climb toward local maxima. The metaphor is so thoroughly embedded in the discourse that even careful treatments of evolution describe organisms as being “designed” by natural selection, meaning not that any agent designed them, but that the process of selection produces outcomes that look as though they were designed for a purpose.

This framing, however, smuggles in a structure that requires examination. An optimization process is defined by three components: a space of possibilities, an objective function that assigns values to elements of that space, and a mechanism that preferentially moves the system toward higher-valued elements. All three are required. Without the objective function, there is no notion of better or worse; without the mechanism that exploits it, variation and selection are merely stochastic. The narrative of progress requires that natural selection supply all three.

What natural selection actually supplies is only a mechanism of exclusion. Organisms that fail to survive and reproduce do not contribute to the next generation. This is a boundary condition, not an objective function. It tells us which trajectories are terminated; it says nothing about the ranking of trajectories that are not terminated. To pass from “this variant was not eliminated” to “this variant is better” requires the additional claim that non-elimination tracks some independent notion of quality. That claim is not part of the mechanism. It is an interpretation imposed on outcomes.

The distinction is not merely verbal. An optimization process explores the space of possibilities in a directed way, guided by the objective function. Natural selection does not explore; it filters. Variants arise through processes—mutation, recombination, developmental noise—that are not guided by the fitness landscape. Selection then removes some of them. What remains is not

what was found by searching; it is what was not destroyed by excluding. The set of survivors is determined by the boundary condition, not by an interior ordering.

This point was recognized early in the mathematical formalization of evolutionary theory. The Price equation, developed by George Price in the 1970s, gives an exact decomposition of evolutionary change that makes no reference to progress or improvement. It states that the change in the population mean of any character z across one generation satisfies

$$\Delta\bar{z} = \frac{\text{Cov}(w, z)}{\bar{w}} + \mathbb{E} \left[\frac{\Delta z}{w} \right] \quad (1)$$

where w is fitness, \bar{w} is mean fitness, and the second term captures transmission bias. The equation is exact and general. It holds for any character, any fitness function, any population structure. What it does not contain is any reference to an external standard against which z is being optimized. The covariance term captures correlation between character and reproductive success, but that correlation is local and contingent. It does not imply convergence toward an independently specified optimum.

Fisher's fundamental theorem, often cited as evidence that selection is an optimizer, asserts that the rate of increase in mean fitness equals the additive genetic variance in fitness. This has sometimes been read as showing that fitness always increases under selection, which would make selection a genuine optimizer with fitness as its objective function. But the theorem is more limited than this reading suggests. It applies only to additive genetic variance in fitness, not to realized fitness in a changing environment. As the environment changes—including as a result of the population's own evolution—the fitness landscape shifts, and there is no guarantee that the population tracks any stable optimum. The theorem describes a local, instantaneous tendency, not a global trajectory toward improvement.

The philosophical consequence is that the appearance of optimization in biological systems is not evidence of an optimizing process. It is evidence of filtering. We observe organisms that have not been eliminated, and because they have not been eliminated, they satisfy the constraints that define non-elimination. From within those constraints, their properties appear coordinated, even purposeful. But this appearance is an artifact of conditioning on survival, not a signature of directed search.

2. SELECTION AS A SIEVE: THE FORMAL STRUCTURE OF PERSISTENCE

To make this argument precise, we need a minimal formal language for what persistence under constraint actually requires. The framework developed here is deliberately spare; it isolates the necessary structure without importing additional assumptions.

Definition 2.1 (State Space and Dynamics). Let \mathcal{X} be a set, called the *state space*, together with a map $T : \mathcal{X} \rightarrow \mathcal{X}$ called the *dynamics*. The pair (\mathcal{X}, T) is called a *dynamical system*. Elements of \mathcal{X} are called *states*, and the sequence $x, T(x), T^2(x), \dots$ is called the *trajectory* of the state x under T .

Remark 2.2. We work here with discrete-time dynamics for clarity of exposition. All results extend naturally to continuous-time systems defined by a flow $\phi : \mathcal{X} \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{X}$ satisfying the semigroup property $\phi(x, 0) = x$ and $\phi(\phi(x, s), t) = \phi(x, s + t)$.

Definition 2.3 (Constraint and Admissibility). A *constraint* on (\mathcal{X}, T) is a subset $C \subseteq \mathcal{X}$. A state $x \in \mathcal{X}$ is *admissible* with respect to C if $x \in C$. A family $\mathcal{C} = \{C_i\}_{i \in I}$ of constraints defines the *admissible region*

$$\mathcal{A}(\mathcal{C}) = \bigcap_{i \in I} C_i \subseteq \mathcal{X}. \quad (2)$$

A state is admissible with respect to \mathcal{C} if and only if it belongs to $\mathcal{A}(\mathcal{C})$.

Definition 2.4 (Persistent Structure). A subset $S \subseteq \mathcal{X}$ is *persistent* with respect to the dynamical system (\mathcal{X}, T) and constraint family \mathcal{C} if two conditions hold simultaneously. First, $S \subseteq \mathcal{A}(\mathcal{C})$, meaning every state in S is admissible. Second,

$$T(S) \subseteq S, \quad (3)$$

meaning that the dynamics map S into itself. Condition (3) is called the *closure condition* or *persistence condition*.

The closure condition is the minimal content of what it means to persist. A set of states that is mapped into itself under the dynamics will continue to contain its own images indefinitely: $T^n(S) \subseteq S$ for all $n \geq 0$. This is not a statement about the quality of the states in S , nor about their ranking relative to states in $\mathcal{X} \setminus S$. It is a statement about closure.

The following lemma makes explicit what the closure condition does and does not imply.

Lemma 2.5 (Closure Without Ranking). *Let $S \subseteq \mathcal{X}$ be persistent. Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be any function. Define $S_\alpha = \{x \in S : f(x) \geq \alpha\}$ for any $\alpha \in \mathbb{R}$. Then S_α need not be persistent, and the dynamics T need not preserve any ordering induced by f on S .*

Proof. It suffices to construct a counterexample. Let $\mathcal{X} = \{a, b, c\}$, and define T by $T(a) = b$, $T(b) = a$, $T(c) = c$. Let $S = \{a, b\}$. Then $T(S) = \{b, a\} = S$, so S is persistent. Now define $f(a) = 2$, $f(b) = 1$, $f(c) = 0$. The ordering induced by f on S is $a > b$. But $T(a) = b$, so the dynamics move states from higher to lower f -value within S . There is no sense in which the dynamics optimize f . Moreover, $S_{1.5} = \{a\}$, and $T(S_{1.5}) = \{b\} \not\subseteq S_{1.5}$, so $S_{1.5}$ is not persistent. The ranking induced by f is entirely disjoint from the closure condition. \square

Lemma 2.5 establishes that persistence is orthogonal to any external ranking. A persistent set need not be the set of high-scoring states under any function f , and the dynamics need not move states in the direction of increasing f . This rules out the possibility of reading optimization off from persistence as a consequence.

The converse direction is equally important. Not every set of high-scoring states is persistent.

Lemma 2.6 (Optimality Without Persistence). *Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be any function and let $\alpha^* = \sup_{x \in \mathcal{X}} f(x)$. Define $S^* = \{x \in \mathcal{X} : f(x) = \alpha^*\}$. Then S^* need not be persistent.*

Proof. Let $\mathcal{X} = \{a, b\}$, $T(a) = b$, $T(b) = a$. Let $f(a) = 1$, $f(b) = 0$. Then $S^* = \{a\}$ and $T(S^*) = \{b\} \not\subseteq S^*$. The globally optimal state under f is not persistent. \square

Lemmas 2.5 and 2.6 together establish that persistence and optimality are logically independent. A structure may persist without being optimal under any reasonable ranking, and may be optimal without persisting. The mechanisms of selection—whatever specific form they take in chemistry, biology, or elsewhere—enforce the closure condition, not optimality.

Definition 2.7 (Minimal Persistent Set). A persistent set S is *minimal* if no proper subset $S' \subsetneq S$ satisfies both $S' \subseteq \mathcal{A}(\mathcal{C})$ and $T(S') \subseteq S'$. Minimal persistent sets are the irreducible units of long-term dynamical behavior within the admissible region.

In the context of the origin of life, the relevant state space is the space of possible chemical configurations of a prebiotic system—molecular species, concentrations, reaction networks, energy states. The constraints are chemical stability conditions, thermodynamic admissibility conditions, and conditions imposed by the physical environment: temperature, pressure, available reactants, ultraviolet flux, pH. A chemical network persists if its states remain within the admissible region under the dynamics of chemical reaction.

The critical point is that this framework places no condition on the efficiency, complexity, or information content of the persistent network. A chemical loop that is highly inefficient, energetically wasteful, and informationally trivial can satisfy the closure condition provided it does not drive itself outside the admissible region. It persists not because it is good at anything, but because it does not fail the exclusion criterion.

3. SURVIVOR BIAS AND THE ILLUSION OF DESIGN

The illusion of design arises from a statistical artifact: we observe only what persists, and we construct our interpretations from that biased sample. This section makes the structure of the bias precise.

Let Ω be a probability space representing the set of all possible chemical trajectories in a prebiotic environment, with probability measure μ reflecting the initial distribution of chemical configurations and environmental conditions. Let $E \subseteq \Omega$ be the event that a trajectory gives rise to a persistent chemical network—that is, a network satisfying the closure condition with respect to the relevant constraints.

An observer embedded in a world produced by one of these trajectories samples from the conditional distribution $\mu(\cdot | E)$, not from μ itself. This conditioning is not an epistemic choice; it is a structural necessity. An observer can only exist in a world that has produced the conditions for observation, which at minimum requires a world that has produced some persistent structure.

Now let $Q : \Omega \rightarrow \mathbb{R}$ be any measurable function that assigns to each trajectory a quantity interpretable as “degree of optimization” or “functional efficiency.” The unconditional distribution of Q under μ may be quite spread out, reflecting the fact that most trajectories in the full distribution produce configurations with no particular functional organization. The conditional distribution of Q under $\mu(\cdot | E)$, however, may appear concentrated in a high- Q

region, because the constraints defining E may correlate with properties that Q also tracks.

The observer who samples from $\mu(\cdot | E)$ and observes that Q takes high values will naturally interpret this as evidence that the process was directed toward high- Q outcomes. But this inference conflates the conditional distribution with the generative process. The process generated trajectories according to μ ; the conditioning on E is a filter applied afterward. The high values of Q in the sample reflect the constraints defining E , not a gradient in μ pointing toward high Q .

This is survivor bias in its precise form. It is not merely the psychological tendency to notice successes; it is the mathematical consequence of conditioning on a non-trivial event. The bias is computable: it is the ratio of the conditional expectation $\mathbb{E}_\mu[Q | E]$ to the unconditional expectation $\mathbb{E}_\mu[Q]$. When E is correlated with high Q —as it will be whenever the constraints defining persistence overlap with the conditions that Q tracks—this ratio exceeds 1, and the observer will systematically overestimate the degree to which the generative process was directed toward high- Q outcomes.

Proposition 3.1 (Survivor Bias as Conditional Inflation). *Let $Q, \mathbf{1}_E$ be positively correlated random variables on (Ω, μ) , i.e., $\text{Cov}_\mu(Q, \mathbf{1}_E) > 0$. Then*

$$\mathbb{E}_\mu[Q | E] > \mathbb{E}_\mu[Q]. \quad (4)$$

An observer conditioning on E will therefore infer a higher mean value of Q in the generative distribution than actually obtains.

Proof. By definition, $\mathbb{E}_\mu[Q | E] = \mathbb{E}_\mu[Q \cdot \mathbf{1}_E] / \mu(E)$. We have

$$\text{Cov}_\mu(Q, \mathbf{1}_E) = \mathbb{E}_\mu[Q \cdot \mathbf{1}_E] - \mathbb{E}_\mu[Q] \cdot \mu(E) > 0, \quad (5)$$

so $\mathbb{E}_\mu[Q \cdot \mathbf{1}_E] > \mathbb{E}_\mu[Q] \cdot \mu(E)$, and therefore $\mathbb{E}_\mu[Q | E] > \mathbb{E}_\mu[Q]$. \square

Proposition 3.1 makes the illusion of design mathematically explicit. The apparent concentration of the surviving sample in high- Q regions is not evidence of an optimizing process; it is evidence that the persistence condition E is positively correlated with Q . The correlation exists because both E and high Q depend on the system satisfying certain constraints. But the existence of a correlation does not imply a causal mechanism driving the system toward high Q . It implies only that the constraints are shared.

The finely tuned character of biological systems—the precise fit between enzyme active sites and their substrates, the coordination of metabolic pathways, the mutual compatibility of cellular components—is real. But its interpretation must be inverted. It is not that the system was driven toward a finely tuned state; it is that only finely tuned states are compatible with the persistence constraints. The precision is a consequence of filtration, not of directed search.

4. THE ORIGIN OF LIFE REFRAMED

With this framework in place, the origin of life can be reinterpreted without appeal to optimization or progress. What is required is not the emergence of an efficient replicator, but the accidental formation of a chemical system that satisfies the closure condition.

Definition 4.1 (Chemical Closure). A reaction network N is *chemically closed* with respect to a set of environmental constraints \mathcal{C} if the set of molecular species and their concentrations under the network dynamics remains within the admissible region $\mathcal{A}(\mathcal{C})$ indefinitely. Formally, letting $S_N \subseteq \mathcal{X}$ be the set of states compatible with network N ,

$$T(S_N) \subseteq S_N \quad \text{and} \quad S_N \subseteq \mathcal{A}(\mathcal{C}). \quad (6)$$

Replication, in this framework, is a specific and particularly effective mechanism for achieving closure. A replicating molecule ensures its own continued presence by producing copies of itself, thereby maintaining the network's state within the admissible region even under degradation. But replication is not the only mechanism of closure. Autocatalytic cycles, where each component catalyzes the production of the next, can achieve closure without any individual molecule being a self-replicator. Stuart Kauffman's analysis of autocatalytic sets identifies exactly this: networks where the set of molecules catalyzes its own production collectively, without requiring a single molecular species to replicate itself.

The formal condition for an autocatalytic set (M, R) of molecular species M and reactions R to be self-sustaining is that every molecule in M is produced by at least one reaction in R whose reactants and catalysts are all in M :

$$\forall m \in M, \exists r \in R \text{ such that } m \in \text{products}(r) \text{ and } \text{reactants}(r) \cup \text{catalysts}(r) \subseteq M. \quad (7)$$

This is precisely the closure condition applied to the categorical structure of a reaction network. It requires mutual admissibility and collective self-maintenance, not individual optimization.

The origin of life, on this account, is not the discovery of an optimal self-replicating molecule but the accidental entry into the admissible region by a chemical configuration that happens to satisfy collective closure. Once such a configuration exists, it persists—not because it is good at anything in an absolute sense, but because it does not drive itself out of the constraints.

From this foundation, the appearance of increasing complexity and efficiency in later evolution requires no revision of the basic framework. More complex networks can achieve closure more robustly: they are less likely to be disrupted by environmental perturbations, because they can maintain admissibility across a wider range of conditions. This is not optimization toward an external criterion; it is the natural consequence of the fact that more robust closure is less likely to fail. Robustness persists because fragility is excluded.

The conclusion of Part I can now be stated as a formal claim.

Theorem 4.2 (Selection Without Optimization). *Let (\mathcal{X}, T) be a dynamical system with constraint family \mathcal{C} . The set of persistent structures—subsets $S \subseteq \mathcal{A}(\mathcal{C})$ satisfying $T(S) \subseteq S$ —is determined entirely by the closure condition and the constraint family. It contains no reference to any objective function, ranking, or gradient. Any ranking $f : \mathcal{X} \rightarrow \mathbb{R}$ that appears to describe the persistent structures is an external imposition, not an intrinsic feature of the dynamics. Furthermore, an observer conditioned on observing only persistent structures will systematically infer higher mean values of any f correlated with persistence than obtain in the underlying generative distribution.*

Proof. The first claim follows directly from Definitions 2.2 and 2.3 and Lemmas 2.5 and 2.6: the closure condition $T(S) \subseteq S$ contains no reference to any function f , and neither optimality under f nor membership in a high- f region is a necessary or sufficient condition for persistence. The second claim is Proposition 3.1. \square

PART II FUNCTION AS PROJECTION

5. THE STRUCTURE OF FUNCTIONAL EXPLANATION

The argument of Part I shows that the language of optimization is not licensed by the mechanism of selection. Part II asks whether a weaker vocabulary—not optimization, but function—can survive. Functionalism in the philosophy of mind and biology holds that what something *is* is determined by what it *does*: by its causal role, its relational structure, its input-output profile. This seems to escape the critique of optimization because it does not assert that systems are driven toward their functions; it asserts only that their functions characterize what they are.

To assess this claim, it is necessary to be precise about what functional explanation commits one to. A functional explanation of a structure s with respect to a function f asserts at minimum the following. First, s reliably produces or contributes to a certain type of output or relation $f(s)$ in a certain type of environment. Second, this production is not accidental; it is a stable feature of s that persists across relevant perturbations. Third, and most importantly, the description of s in terms of f is more fundamental than, or at least not reducible to, its description in terms of physical composition.

The third commitment is what gives functionalism its philosophical interest. It motivates the claim that two physically distinct structures—silicon and carbon, neurons and transistors—can be instances of the same type, because they realize the same function. It grounds the possibility of a science of mind that is autonomous from neuroscience: the relevant description is functional, and the physical implementation is in principle interchangeable.

It is this third commitment that Part II will dismantle. The argument does not deny that functional descriptions are useful or that they often track stable and important structural features. It denies that functional descriptions are ontologically primary, that they describe intrinsic properties of systems independent of any observer or framework.

6. THE BRIDGE: FROM PERSISTENCE TO THE RELATIVITY OF FUNCTION

The argument connecting Part I to the critique of functionalism proceeds in four steps, each of which follows necessarily from what precedes it.

Step 1. By Theorem 4.2, selection guarantees closure, not optimization. A persistent structure S satisfies $T(S) \subseteq S$ under the relevant constraints. This is the entire content of what selection establishes about S .

Step 2. A functional attribution of the form “ S is for f ” asserts that among all the invariant properties of S —all the regularities in its behavior that are preserved across its dynamics—one particular regularity f is privileged as the function of S . The statement is not merely that S reliably exhibits the regularity f ; it is that f is what S is for, that f explains the existence or persistence of S .

Step 3. For f to be intrinsically privileged, there must be something in the dynamics of S that distinguishes f from all other invariant regularities. There must be a sense in which f is not just one of many patterns that S exhibits, but the pattern around which S is organized. In an optimization framework, this distinction is supplied by the objective function: f is the criterion that selection was tracking, so f is what explains persistence. But by Step 1, no such criterion is part of the dynamics. The closure condition treats all constraints simultaneously and symmetrically; it does not privilege any single regularity. In particular, any claim that f explains the persistence of S requires that persistence vary as a function of f ; but by Lemma 2.5, persistence is invariant under arbitrary reorderings induced by f , and therefore no such explanatory dependence exists.

Step 4. Therefore, the privileging of f over other invariant regularities of S is not read off from the dynamics. It is imposed by an external choice of coordinate—a projection that selects f as the relevant axis of description. Different projections impose different functions. The function of S is not intrinsic to S ; it is relative to the projection.

This argument can be stated more formally. Let S be a persistent structure in (\mathcal{X}, T) , and let $\text{Inv}(S) = \{f : S \rightarrow \mathbb{R} : f \circ T|_S = f\}$ be the set of all real-valued invariants of S under the restricted dynamics $T|_S$. This is the set of all quantities that are preserved along trajectories within S .

Proposition 6.1 (Underdetermination of Function by Dynamics). *For any persistent structure S in a dynamical system (\mathcal{X}, T) with more than one trajectory, $\text{Inv}(S)$ contains infinitely many elements. The closure condition $T(S) \subseteq S$ does not determine which element of $\text{Inv}(S)$, if any, is the “function” of S . Any selection of a function from $\text{Inv}(S)$ requires an external criterion not contained in the dynamics.*

Proof. If S contains more than one trajectory, then there exist at least two distinct trajectory classes in S . For any partition of S into trajectory classes, the indicator function of any union of classes is an invariant. The set of such

partitions is at least countably infinite (it includes the trivial partition and all finer refinements), so $|\text{Inv}(S)| \geq \aleph_0$. Given any two distinct elements $f, g \in \text{Inv}(S)$, the closure condition $T(S) \subseteq S$ is satisfied by both and provides no basis for preferring one over the other as the “function.” Any such preference requires an external criterion. \square

Proposition 6.1 establishes that functional attribution is underdetermined by the dynamics. The persistent structure S has infinitely many invariant regularities, and the mechanism that produced S —the closure condition under constraint—does not privilege any of them. To select one as the function of S is to impose an external coordinate.

7. TELEOLOGICAL FUNCTIONALISM DISSOLVED

Teleological accounts of function, associated with the work of Larry Wright and refined by Karen Neander in terms of biological proper functions, hold that the function of a trait is what it was selected *for*: what it did in the ancestral lineage that caused it to be retained and spread. The function of the heart, on this account, is to pump blood because pumping blood is what caused hearts to be selected for in the ancestral lineage of vertebrates.

This account was designed to escape the charge of anthropomorphism: it grounds function in actual causal history rather than in the intentions of an observer. But the argument of Section 6 shows that the escape is illusory. The claim that pumping blood is what the heart was “selected for” presupposes exactly what has been shown to be absent from the dynamics of selection: a gradient or criterion that selection was tracking, such that pumping blood is privileged over the indefinitely many other invariant regularities of the heart—its weight, its periodic motion, its electromagnetic signature, its role in circulating immune cells.

More precisely, the teleological account requires that among all the co-varying properties of ancestral hearts, one property—pumping blood—was causally responsible for the differential reproductive success of organisms possessing those hearts. This is an empirical claim, and it may be true in a statistical sense: pumping blood may have been more consistently correlated with reproductive success across environments than the heart’s other properties. But establishing this statistical correlation does not establish that pumping blood is intrinsically privileged as the function. It establishes only

that pumping blood was, in the relevant historical environments, part of the constraint set that defined the persistence region.

By Proposition 6.1, this historical correlation is itself one invariant among many. The selection history carves out a persistent structure, but it does not designate which of that structure's invariants is the function. The teleological account adds to the dynamics a retrospective interpretation—"this is what the trait was selected for"—that the dynamics do not contain. The word "for" in "selected for" is doing normative work that the causal history alone does not supply.

Corollary 7.1 (Teleological Function as Retrospective Projection). *Let S be a persistent structure whose persistence is explained by historical selection. Let $f \in \text{Inv}(S)$ be the property designated as S 's function by the teleological account. Then f is one element of the infinitely many invariants of S , designated by an external criterion (the choice of which covariate to treat as causally relevant to selection) that is not determined by the dynamics or the selection history alone. The teleological function is a projection, not an intrinsic property.*

8. COMPUTATIONAL FUNCTIONALISM AND THE PROBLEM OF EQUIVALENCE

Computational functionalism, associated with Hilary Putnam's original formulation and subsequent development by Jerry Fodor and others, advances a more sophisticated claim. It does not invoke selection history. It asserts instead that mental states are defined by their abstract causal-functional roles: the pattern of relations between inputs, internal states, and outputs that characterizes a type of processing. What makes something a belief that it is raining is not its physical constitution, not its evolutionary history, but its functional organization: the way it interacts with perception, with other beliefs, with desires, and with behavior.

The philosophical appeal of this view is its combination of ontological liberalism with explanatory rigor. Multiple realizability follows immediately: any physical system that instantiates the relevant functional organization has the mental state in question, regardless of substrate. Computational functionalism thus promises to ground psychology as an autonomous science whose kinds are functional rather than physical.

The argument of Section 6 applies here as well, but its target must be located more precisely. Computational functionalism does not assert that

functional states were optimized toward. It asserts that functional states are defined by equivalence classes under a certain relation: two physical states are functionally equivalent if and only if they occupy the same causal role within the relevant system. The ontological commitment is not to optimization but to equivalence.

But equivalence relations are not given by the dynamics. They are imposed by a choice of description. To say that two physical systems s_1 and s_2 instantiate the same computation, one must specify a mapping $\phi : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ between their state spaces such that $\phi \circ T_1 = T_2 \circ \phi$ —that is, such that ϕ is a morphism of dynamical systems. Such a morphism is a *conjugacy* or *topological equivalence* between the systems.

Definition 8.1 (Dynamical Conjugacy). Two dynamical systems (\mathcal{X}_1, T_1) and (\mathcal{X}_2, T_2) are *conjugate* if there exists a bijection $\phi : \mathcal{X}_1 \rightarrow \mathcal{X}_2$ such that $\phi \circ T_1 = T_2 \circ \phi$. The map ϕ is called a *conjugacy*.

Proposition 8.2 (Non-Uniqueness of Functional Equivalence). *Let (\mathcal{X}, T) be a dynamical system. For any persistent subset $S \subseteq \mathcal{X}$, there exist in general multiple non-equivalent dynamical systems (\mathcal{X}', T') that are conjugate to $(\mathcal{X}, T|_S)$ under different conjugacies. The choice of which system counts as “realizing the same function” as S depends on the choice of conjugacy, which is not determined by the dynamics of S alone.*

Proof. Conjugacy is an equivalence relation on dynamical systems, and its equivalence classes are generally large. Any automorphism of \mathcal{X} that commutes with T generates a different self-conjugacy of (\mathcal{X}, T) . If $\text{Aut}(\mathcal{X}, T)$ is the group of such automorphisms, then any two systems (\mathcal{X}_1, T_1) and (\mathcal{X}_2, T_2) connected by conjugacies ϕ_1 and ϕ_2 to $(\mathcal{X}|_S, T|_S)$ may not be conjugate to each other, because $\phi_2^{-1} \circ \phi_1$ may not be a legitimate morphism if it does not preserve the relevant structure. The collection of systems conjugate to $(\mathcal{X}|_S, T|_S)$ under all possible conjugacies therefore contains systems that are not mutually conjugate. There is no dynamics-internal criterion to select one of these conjugacies as the “correct” one for determining functional equivalence. \square

This proposition makes the problem with computational functionalism precise. The claim that two systems realize the same function presupposes a choice of conjugacy—a choice of how to map the state spaces of the two systems onto each other. Different choices of conjugacy yield different judgments about functional equivalence. A silicon circuit and a neural network might be conjugate under one mapping but not under another. The claim that they

instantiate “the same computation” is not a fact about either system in isolation; it is a fact about the relation induced by a particular choice of conjugacy. The issue is not that conjugacy fails as a mathematical notion of equivalence, but that the selection of a particular conjugacy class as explanatorily relevant is not determined by the dynamics; it depends on which structural features are chosen to be preserved.

The equivalence relation that computational functionalism invokes is not determined by the dynamics. It is imposed by an observer who selects a level of description, a notion of relevant similarity, and a mapping between systems. This selection is what Proposition 6.1 called an external criterion. It is a projection.

The situation is analogous to the choice of coordinate system in physics. Two descriptions of the same physical system in different coordinate systems are related by a transformation, but neither coordinate system is intrinsically correct. The physical facts are coordinate-independent; the coordinate-dependent descriptions are useful representations, not ontological primitives. Multiple realizability is the functionalist’s coordinate freedom: it acknowledges that the same functional description can be applied to multiple physical systems. But it does not follow that the functional description captures an intrinsic, observer-independent property. It follows only that a certain coordinate system—a certain projection—can be applied to multiple systems consistently.

9. FUNCTIONAL DESCRIPTIONS AS COORDINATE SYSTEMS

Having dissolved both teleological and computational functionalism, we can now give a positive account of what functional descriptions are and why they are useful, without treating them as ontologically primary.

A functional description of a persistent structure S is a choice of projection $\pi_F : S \rightarrow F$, where F is a space of functional states. The projection maps physical states in S to functional states in F in a way that is consistent with the dynamics: if $T(x) = y$ in S , then the functional dynamics T_F on F satisfy $T_F(\pi_F(x)) = \pi_F(y)$. In other words, π_F is a morphism from the dynamical system $(S, T|_S)$ to the dynamical system (F, T_F) .

Definition 9.1 (Functional Projection). *A functional projection of a dynamical system $(S, T|_S)$ is a surjective morphism $\pi_F : (S, T|_S) \rightarrow (F, T_F)$. The dynamical system (F, T_F) is called a functional description of S .*

Functional projections exist in abundance—by Proposition 6.1, every invariant of S gives rise to a functional projection, and there are infinitely many invariants. The choice of which functional projection to use is a choice of which aspects of the system’s behavior to represent. This choice may be guided by many considerations: explanatory parsimony, predictive utility, relevance to the observer’s practical interests, compatibility with higher-level theories. But it is a choice, and it is made by an observer, not read off from the system.

Different levels of biological description induce different functional projections. At the molecular level, a protein’s function is its catalytic activity: the functional projection maps protein conformations to catalytic states. At the cellular level, the same protein’s function may be its role in a signaling pathway: a different projection maps protein activities to pathway states. At the organismal level, the signaling pathway’s function is its contribution to development or homeostasis: yet another projection. Each level is a legitimate coordinate system. None is intrinsically correct.

Corollary 9.2 (Hierarchy of Functional Descriptions). *For any persistent structure S , the collection of functional projections $\{\pi_F : (S, T|_S) \rightarrow (F, T_F)\}$ forms a partial order under the refinement relation: π_{F_1} refines π_{F_2} if there exists a morphism $\psi : F_1 \rightarrow F_2$ such that $\pi_{F_2} = \psi \circ \pi_{F_1}$. Coarser projections lose information; finer projections preserve more of the structure of $(S, T|_S)$. No element of this partial order is privileged by the dynamics.*

The conclusion of Part II is therefore not eliminativist. Functional descriptions are genuine representations of real structural features of persistent systems. They are useful, often indispensable, and capable of capturing regularities that purely physical descriptions miss or obscure. But they are representations, not ontological primitives. The function of a system is what an observer with a particular coordinate system sees when they look at it. Different observers, with different coordinate systems, see different functions. This is not a defect to be corrected; it is the correct account of what functional description is.

PART III OBSERVATION AS CONSTRAINED ACCESS

10. THE OBSERVER WITHIN THE FIELD

Parts I and II have established that persistence is orthogonal to optimization, and that function is a projection rather than an intrinsic property. Both arguments have treated the observer as an external analyst, someone who formulates the state space, identifies the constraints, and chooses a projection. Part III removes this privilege. The observer is not external to the dynamical field; it is a subsystem within it, subject to the same framework of constraint and closure.

The term *observer* should be understood broadly throughout what follows as any persistent subsystem capable of inducing a projection, not necessarily a conscious agent. A measuring instrument, a sensory organ, a recording medium—any system that maps states of the world into a representation space qualifies. This generality is deliberate: it prevents the argument from being read as a claim about human cognition specifically, and connects it to the framework already developed for biological and physical systems.

This relocation has consequences that are not merely philosophical. It introduces a second layer of structure that was absent from the earlier analysis: the observer's own dynamical constraints limit which parts of the state space can be accessed, and this limitation is not an epistemic inconvenience to be overcome but a structural feature of what observation is.

Definition 10.1 (Observer as Subsystem). An *observer* in a dynamical system (\mathcal{X}, T) is a tuple $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$ where $\mathcal{O}_X \subseteq \mathcal{X}$ is the *observer's domain*—the set of states accessible to the observer—satisfying the closure condition $T(\mathcal{O}_X) \subseteq \mathcal{O}_X$, $T_{\mathcal{O}} = T|_{\mathcal{O}_X}$ is the restricted dynamics, and $\pi : \mathcal{O}_X \rightarrow \mathcal{M}$ is a non-injective *observation map* from the observer's domain into a *measurement space* \mathcal{M} .

The closure condition $T(\mathcal{O}_X) \subseteq \mathcal{O}_X$ expresses the requirement that the observer's dynamics are self-consistent: the observer, as a physical system, does not leave its own admissible region under its own dynamics. This is just the persistence condition applied to the observer itself. It is not a guarantee about what the observer can learn; it is a constraint on what the observer can be.

The observation map $\pi : \mathcal{O}_X \rightarrow \mathcal{M}$ represents the process by which the observer's physical states in \mathcal{O}_X are converted into representations in \mathcal{M} . That

π is non-injective means that distinct states of the world can produce identical observations. This is not a contingent limitation; it is a structural consequence of the fact that \mathcal{M} is a representation space, necessarily lower-dimensional than \mathcal{X} in any system of interest.

Definition 10.2 (Reachable Set). Let $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$ be an observer, and let $x_0 \in \mathcal{O}_X$ be an initial state. The *reachable set* of \mathcal{O} from x_0 is

$$\mathcal{R}(x_0) = \{T^n(x_0) : n \geq 0\} \subseteq \mathcal{O}_X. \quad (8)$$

More generally, if the observer's initial state is drawn from a distribution μ_0 on \mathcal{O}_X , the reachable set is $\mathcal{R} = \overline{\bigcup_{x_0 \in \text{supp}(\mu_0)} \mathcal{R}(x_0)}$.

The reachable set is the set of states that the observer can, in principle, visit given its initial conditions and dynamics. It is in general a proper subset of \mathcal{O}_X , which is itself a proper subset of \mathcal{X} . The observation of the world is therefore doubly filtered: first by the restriction of the observer's domain $\mathcal{O}_X \subseteq \mathcal{X}$, and second by the further restriction to the reachable set $\mathcal{R} \subseteq \mathcal{O}_X$.

11. REACHABILITY AND THE LIMITS OF ACCESS

The reachable set of an observer is determined by the observer's initial conditions and its constraint field. It is not determined by the structure of \mathcal{X} alone, nor by the observer's theoretical interests. This creates a form of structural limitation that is categorically different from ordinary information loss.

Ordinary information loss occurs when a projection π collapses distinct states to the same representation. This is the non-injectivity of π , and it implies that the observer cannot distinguish certain states even if their trajectory passes through those states. Structural limitation is different: it occurs when the observer's trajectory does not pass through a region at all, so that the question of distinguishability does not arise. The region is not compressed into indistinguishability; it is simply never sampled.

To make this distinction precise, define the set of states that are observationally accessible as the image of the reachable set under the observation map:

$$\pi(\mathcal{R}) = \{\pi(x) : x \in \mathcal{R}\} \subseteq \mathcal{M}. \quad (9)$$

The fiber of an observation $m \in \pi(\mathcal{R})$ is the set of states in \mathcal{X} that would produce observation m :

$$\pi^{-1}(m) = \{x \in \mathcal{X} : \pi(x) = m\}. \quad (10)$$

Because π is non-injective, fibers may contain many states. The observer's uncertainty about the world given observation m is represented by the fiber $\pi^{-1}(m)$: the observation is consistent with any state in the fiber.

But the observer's actual epistemic access is more restricted than this. The states in the fiber that are consistent with the observer's history are only those in $\pi^{-1}(m) \cap \mathcal{R}$. States in $\pi^{-1}(m) \setminus \mathcal{R}$ are in principle consistent with the observation m —they would produce the same measurement—but they are not reachable by the observer. They are absent not because they have been ruled out by evidence, but because the observer's trajectory has never visited them.

Definition 11.1 (Epistemically Accessible Fiber). Given an observer \mathcal{O} with reachable set \mathcal{R} and observation map π , the *epistemically accessible fiber* over an observation $m \in \mathcal{M}$ is

$$\pi^{-1}(m) \cap \mathcal{R}. \quad (11)$$

This is the set of states that are both consistent with observation m and reachable by the observer.

The lifting problem—the problem of inferring the state of the world from an observation—is correctly posed as: given observation m , identify the element(s) of $\pi^{-1}(m) \cap \mathcal{R}$ that are consistent with all prior observations. This is a more constrained problem than the naive lifting problem, which would identify elements of the full fiber $\pi^{-1}(m)$. It is also a more limited one: explanations that lie in $\pi^{-1}(m) \setminus \mathcal{R}$ are unavailable to the observer, not because they are false or inconsistent with the evidence, but because no admissible trajectory reaches them.

Theorem 11.2 (Structural Underdetermination of Explanation). *Let $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$ be an observer with reachable set \mathcal{R} . For any observation $m \in \pi(\mathcal{R})$, if $|\pi^{-1}(m) \cap \mathcal{R}| > 1$, then the observation m is consistent with multiple distinct states, and no dynamical information available to \mathcal{O} can determine which state obtains. Moreover, any state in $\pi^{-1}(m) \setminus \mathcal{R}$ represents an explanation of m that is structurally unavailable to \mathcal{O} , regardless of the quality or completeness of the evidence.*

Proof. The first claim follows from the definition of the fiber: if $|\pi^{-1}(m) \cap \mathcal{R}| > 1$, then distinct states $x_1, x_2 \in \pi^{-1}(m) \cap \mathcal{R}$ both produce observation m and are both reachable. The observation m alone does not determine which state obtains, and because both states are reachable, the observer's trajectory history does not rule out either without additional observations. The second claim follows from the definition of the reachable set: if $x \in \pi^{-1}(m) \setminus \mathcal{R}$, then no trajectory of \mathcal{O} passes through x , so x never contributes to any observation of \mathcal{O} , and no sequence of observations can lead \mathcal{O} to identify x as the actual state. \square

Theorem 11.2 establishes two distinct sources of epistemic limitation: the non-injectivity of π (which produces the accessible fiber), and the restriction of the trajectory to \mathcal{R} (which produces the structural inaccessibility of $\pi^{-1}(m) \setminus \mathcal{R}$). Standard epistemology and philosophy of science attend primarily to the first source. Theorem 11.2 insists on the second.

12. ISLANDS OF PERSISTENCE

The framework now allows a precise formulation of the possibility raised in the opening commentary: that there exist regions of \mathcal{X} that are dynamically stable—persistent under the closure condition—yet structurally inaccessible to any observer in a given class.

Definition 12.1 (Inaccessible Persistent Region). A region $I \subseteq \mathcal{X}$ is an *inaccessible persistent region* with respect to an observer class $\{\mathcal{O}_\alpha\}_{\alpha \in A}$ if two conditions hold simultaneously. First, I is persistent: $I \subseteq \mathcal{A}(\mathcal{C})$ and $T(I) \subseteq I$. Second, I is structurally inaccessible to every observer in the class:

$$\forall \alpha \in A : \quad \mathcal{R}_\alpha \cap I = \emptyset, \quad (12)$$

where \mathcal{R}_α is the reachable set of observer \mathcal{O}_α .

The existence of inaccessible persistent regions is not a speculative posit; it is a structural consequence of the framework. The following proposition establishes that such regions arise generically.

Proposition 12.2 (Generic Existence of Inaccessible Persistent Regions). *Let (\mathcal{X}, T) be a dynamical system with constraint family \mathcal{C} , and let $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$ be an observer with $\mathcal{O}_X \subsetneq \mathcal{X}$. If $\mathcal{X} \setminus \mathcal{O}_X$ contains at least one persistent subset $I \subseteq \mathcal{A}(\mathcal{C})$ satisfying $T(I) \subseteq I$, then I is an inaccessible persistent region with respect to \mathcal{O} .*

Proof. Since $I \subseteq \mathcal{X} \setminus \mathcal{O}_X$, we have $I \cap \mathcal{O}_X = \emptyset$. The reachable set $\mathcal{R} \subseteq \mathcal{O}_X$, so $\mathcal{R} \cap I = \emptyset$. The conditions $I \subseteq \mathcal{A}(\mathcal{C})$ and $T(I) \subseteq I$ are given. Therefore I satisfies both conditions of Definition 11.1. \square

Remark 12.3. The condition $\mathcal{O}_X \subsetneq \mathcal{X}$ is nearly trivially satisfied in any realistic setting. An observer is a physical system of finite complexity embedded in a universe of greater complexity. The observer's state space is necessarily a proper subspace of the full state space of the universe. The only question is whether the complement $\mathcal{X} \setminus \mathcal{O}_X$ contains persistent structure, which is generically true for any system rich enough to support observers.

The existence of inaccessible persistent regions has an immediate philosophical consequence: empirical emptiness does not imply ontological absence. A region of the world that produces no observations in any observer in a given class may nonetheless exist as a coherent, dynamically stable structure. The absence of evidence is, in this case, not merely inconclusive; it is structurally guaranteed. The observer cannot encounter the region, not because it is absent, but because the observer's constraint field does not permit trajectories that intersect it.

This conclusion parallels results in dynamical systems theory about the decomposition of state spaces into invariant components. In a system with multiple basins of attraction, trajectories in one basin never visit another, even if the other basin contains globally stable attractors. The existence of multiple basins is not a deficiency of the dynamics; it is a structural feature. The framework here generalizes this: the separation of the observer's reachable domain from inaccessible persistent regions is a consequence of the constraint structure, not an accident.

13. EPISTEMIC CURVATURE

The most important—and most original—consequence of the foregoing is that inaccessible persistent regions are not simply absent from the observer's experience. They deform it. The reachable domain \mathcal{R} of an observer is not an arbitrary subset of \mathcal{X} ; it is shaped by the same constraint field that defines what lies beyond it. The boundary of \mathcal{R} encodes information about what is on the other side.

To develop this idea formally, we introduce the notion of the boundary of the reachable set and show that it carries structural information about inaccessible regions.

Definition 13.1 (Boundary of Reachable Set). Let \mathcal{O} be an observer with reachable set $\mathcal{R} \subseteq \mathcal{O}_X \subseteq \mathcal{X}$. Assume \mathcal{X} carries a topology. The *boundary of the reachable set* is $\partial\mathcal{R} = \overline{\mathcal{R}} \cap \overline{\mathcal{X} \setminus \mathcal{R}}$, the set of points that are limit points of both \mathcal{R} and its complement.

Definition 13.2 (Epistemic Curvature). The *epistemic curvature* at a point $x \in \mathcal{R}$ is the rate at which the reachable set “bends away” from the inaccessible region near x . If \mathcal{R} is a submanifold of a Riemannian manifold (\mathcal{X}, g) , the epistemic curvature is the second fundamental form of \mathcal{R} in \mathcal{X} , which measures the curvature of \mathcal{R} as a submanifold. Informally, it measures how the constraint field that defines $\partial\mathcal{R}$ deforms the geometry of the accessible region.

The key insight is that the second fundamental form of \mathcal{R} in \mathcal{X} depends on the ambient geometry—the geometry of \mathcal{X} near $\partial\mathcal{R}$ —which is determined by the constraint field. If inaccessible persistent regions lie just beyond $\partial\mathcal{R}$, their existence shapes the ambient geometry, which in turn shapes the second fundamental form, which in turn shapes the observable geometry of \mathcal{R} .

This can be stated as a detection principle: the presence of inaccessible persistent regions may be inferred from anomalies in the curvature of the reachable domain. An observer who cannot see beyond $\partial\mathcal{R}$ may nevertheless detect the presence of structure beyond it by observing that the geometry of \mathcal{R} near $\partial\mathcal{R}$ is not what would be expected in the absence of nearby persistent structure.

Theorem 13.3 (Curvature as Signature of Inaccessible Structure). *Let (\mathcal{X}, g) be a Riemannian manifold with dynamics T derived from a smooth vector field V on \mathcal{X} . Let \mathcal{O} be an observer with reachable set \mathcal{R} , and let $I \subseteq \mathcal{X} \setminus \mathcal{R}$ be an inaccessible persistent region. If the constraint field \mathcal{C} that defines both $\partial\mathcal{R}$ and I induces a non-trivial curvature on the boundary hypersurface $\partial\mathcal{R}$ as an embedded submanifold of \mathcal{X} , then the principal curvatures of $\partial\mathcal{R}$ contain information about the proximity and structure of I .*

More precisely, if I and \mathcal{R} are separated by a constraint barrier $\Sigma \supseteq \partial\mathcal{R}$ with associated potential function $\Phi : \mathcal{X} \rightarrow \mathbb{R}$ (so that $\Sigma = \Phi^{-1}(0)$, $\mathcal{R} \subseteq \Phi^{-1}(\mathbb{R}_{<0})$, and $I \subseteq \Phi^{-1}(\mathbb{R}_{>0})$), then the Hessian $\nabla^2\Phi|_{\Sigma}$ encodes second-order information about the structure of Φ on the I side of Σ , which is detectable from measurements of the curvature of Σ from within \mathcal{R} .

Proof. The second fundamental form of Σ in (\mathcal{X}, g) is determined by the normal derivative of the unit normal field $\hat{n} = \nabla\Phi/|\nabla\Phi|$ along Σ . This is

computed as $II(u, v) = g(\nabla_u \hat{n}, v)$ for tangent vectors u, v to Σ . The Hessian of Φ decomposes as $\nabla^2 \Phi = II \cdot |\nabla \Phi| + (\nabla^2 \Phi)_{\hat{n}\hat{n}} \hat{n} \otimes \hat{n}$, so the eigenvalues of II (the principal curvatures of Σ) are determined by the Hessian of Φ restricted to the tangent space of Σ , divided by $|\nabla \Phi|$. The Hessian of Φ at a point $p \in \Sigma$ depends on the behavior of Φ in a neighborhood of p , which includes points on the I side of Σ . Therefore, the principal curvatures of Σ , which are measurable from within \mathcal{R} by observing the geometry of trajectories near $\partial \mathcal{R}$, carry information about Φ on the I side, which encodes information about the structure of the inaccessible region I . \square

Theorem 13.3 establishes the detectability principle: the observer cannot directly sample I , but the observer can measure the geometry of the accessible region near the boundary, and that geometry is shaped by I 's existence and structure. This establishes that inaccessibility is not epistemically neutral: the structure of what cannot be reached enters into the observable world as a deformation of its geometry. The unseen is not merely absent; it is present as curvature.

The analogy with dark matter is instructive. Dark matter is inferred not from direct observation but from gravitational effects on visible matter: galaxy rotation curves, gravitational lensing, large-scale structure. The inferred matter is "dark" not because it is ontologically absent but because it does not interact with the electromagnetic sector that our instruments are tuned to. It lies in an inaccessible region of the space of observable states, but its presence is encoded in the curvature of the accessible region. Theorem 13.3 is the abstract version of this inference pattern.

14. THE LIFTING PROBLEM AND STRUCTURAL UNDERDETERMINATION

We are now in a position to give a complete account of the epistemological situation. The observer is a subsystem with a reachable set $\mathcal{R} \subseteq \mathcal{O}_X \subseteq \mathcal{X}$. Observations are elements of $\pi(\mathcal{R}) \subseteq \mathcal{M}$. Explanation is the lifting of observations to states: given $m \in \mathcal{M}$, identify the element(s) of $\pi^{-1}(m) \cap \mathcal{R}$ that are consistent with the observed trajectory.

This lifting is subject to two distinct forms of underdetermination, which we have now established separately.

The first, established in Part II, is *functional underdetermination*: even if the state $x \in \mathcal{R}$ is known, the functional description of x is underdetermined by

the dynamics, because there are infinitely many invariant regularities of x and no dynamics-internal criterion for privileging one. Functional attribution requires an external choice of projection.

The second, established in this Part, is *structural underdetermination*: the state x itself is underdetermined by the observation m , because the fiber $\pi^{-1}(m) \cap \mathcal{R}$ may contain multiple elements, and additionally because $\pi^{-1}(m) \setminus \mathcal{R}$ may contain elements that are consistent with m but structurally unavailable.

Theorem 14.1 (Layered Underdetermination). *Let $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$ be an observer and let $m \in \pi(\mathcal{R})$ be an observation. Then:*

1. (State Underdetermination) *The observation m underdetermines the state: any element of $\pi^{-1}(m) \cap \mathcal{R}$ is consistent with m , and the accessible fiber $\pi^{-1}(m) \cap \mathcal{R}$ generically contains more than one element.*
2. (Structural Inaccessibility) *Any element of $\pi^{-1}(m) \setminus \mathcal{R}$ represents an explanation of m that is structurally unavailable: consistent with m , but unreachable by \mathcal{O} .*
3. (Functional Underdetermination) *Given a state $x \in \pi^{-1}(m) \cap \mathcal{R}$, the functional description of x is further underdetermined by $|\text{Inv}(S_x)| = \infty$, where S_x is the persistent set containing x .*

The total underdetermination is therefore the composition of all three layers.

Proof. Part (1) follows from Theorem 11.2. Part (2) follows from Definition 10.2 and Proposition 12.2. Part (3) follows from Proposition 6.1. The composition follows from the sequential dependence: to identify a functional description, one must first identify a state; to identify a state, one must lift an observation; and each step introduces irreducible underdetermination. \square

The philosophical consequences of Theorem 14.1 are substantial. Knowledge, in this framework, is not a convergent process of accumulating representations that approach the truth. It is a trajectory through a space of observations, constrained by the observer's reachable set, partially lifting to a space of states via the fiber structure of π , and further constrained in interpretation by the underdetermination of functional descriptions. At each step, there are choices that are not forced by the evidence, and at each step, there are possibilities that are excluded not by evidence but by the geometry of the observer's constraint field.

This is not skepticism. It is not the claim that knowledge is impossible or that truth is inaccessible. It is the more precise claim that what can be known

is structurally constrained by what is reachable, and that the reachable set is a historical artifact—the result of the observer’s own persistence under its own constraints across time.

PART IV THE UNITY OF BEING AND KNOWING

14. THE COLLAPSE OF THE DISTINCTION

The argument has proceeded in three movements, each apparently about a different domain: biology, philosophy of mind, epistemology. In fact, all three have been about the same structure, applied at increasing levels of generality.

In Part I, the state space was the space of chemical configurations, the dynamics were chemical reactions, the constraints were thermodynamic and environmental, and the persistent structures were self-maintaining chemical networks. The observer was implicit: us, looking back at the history of life and inferring optimization.

In Part II, the same structure was applied to functional explanation. The state space was the space of physical states of a system, the dynamics were its causal evolution, and the persistent structures were those satisfying closure under the relevant constraints. The observer was explicit: the entity who chooses a projection and applies a functional description.

In Part III, the observer was itself relocated within the same framework. The state space is now the full state space of the world, the observer’s domain is a proper subset, and the observer’s trajectory is a constrained path through that subset. The observer is not exempt from the framework; it is an instance of it.

This convergence is not an accident. The framework is defined by three elements: a state space \mathcal{X} , a dynamics T , and a constraint family \mathcal{C} . Everything else—persistence, function, observation, knowledge—is derived from these three elements by applying the closure condition and the projection apparatus at different levels of description. The diversity of the essay’s topics reflects not the application of a framework to three different domains, but the discovery that three apparently different domains are projections of the same framework.

Theorem 14.1 (Unification by Constraint). *Let $(\mathcal{X}, T, \mathcal{C})$ be a constrained dynamical system. Then:*

1. Persistence is defined by the closure condition $T(S) \subseteq S$ within $\mathcal{A}(\mathcal{C})$. It is intrinsic to $(\mathcal{X}, T, \mathcal{C})$.
2. Function is defined by a choice of projection $\pi_F : (S, T|_S) \rightarrow (F, T_F)$, where S is a persistent set. It is relative to the projection, not intrinsic to $(\mathcal{X}, T, \mathcal{C})$.
3. Observation is defined by the reachable set \mathcal{R} of an observer subsystem $\mathcal{O} \subseteq (\mathcal{X}, T)$ and the observation map $\pi : \mathcal{R} \rightarrow \mathcal{M}$. It is relative to the observer's constraint field and initial conditions.
4. Knowledge is the progressive refinement of the lifting map from $\pi(\mathcal{R})$ to elements of $\pi^{-1}(\cdot) \cap \mathcal{R}$, subject to the layered underdetermination of Theorem 14.1.

Of these four, only persistence is intrinsic. Function, observation, and knowledge are projections: coordinate choices that depend on the observer and are not determined by the dynamics alone.

The theorem is not a proof of a non-trivial mathematical fact; it is a synthesis of the results established in Parts I through III, collecting them under a unified statement. Its purpose is to make the unification explicit: the essay has been developing a single framework, not applying three different frameworks to three different problems.

15. PHYSICS WITHOUT PRIVILEGED OBSERVATION

The physical reading of this framework is a claim about ontology: the world contains persistent structure that is independent of any observer. This claim is not trivial. A tradition in physics, particularly in interpretations of quantum mechanics, has been drawn to the idea that physical reality is constituted by or dependent on observation. The measurement problem, the role of the observer in collapse interpretations, the participatory universe of Wheeler—all of these position observation as a foundational element of physical ontology.

The framework developed here takes a different position. Observation is a local, constrained, trajectory-dependent process within a field of constraint. The field itself—the state space \mathcal{X} , the dynamics T , the constraint family \mathcal{C} —does not depend on the existence of observers for its structure. Persistent regions satisfy the closure condition independently of whether any trajectory passes through them.

This position is closest to the structural realism of Eddington and later Ladyman and Ross: what exists are structural relations, not intrinsic properties of substances. But it adds a dynamical dimension: the structural relations are

those defined by constraint and closure, and their persistence is a dynamical fact, not a logical or mathematical one.

Proposition 15.1 (Observer-Independent Persistence). *Let $I \subseteq \mathcal{X}$ be an inaccessible persistent region with respect to an observer class $\{\mathcal{O}_\alpha\}$. Then the persistence of I —its continued satisfaction of $T(I) \subseteq I$ within $\mathcal{A}(\mathcal{C})$ —is independent of the existence, structure, or trajectories of any \mathcal{O}_α . Removing all observers from the system does not alter the dynamics of I .*

Proof. The closure condition $T(I) \subseteq I$ and the admissibility condition $I \subseteq \mathcal{A}(\mathcal{C})$ are defined entirely in terms of the dynamics T and the constraint family \mathcal{C} , neither of which contains any reference to the observer class $\{\mathcal{O}_\alpha\}$. The dynamics of I under T are therefore unaffected by any property of the observer class, including its existence. Formally, if $T_{\text{with obs}}$ and $T_{\text{without obs}}$ denote the global dynamics with and without observers respectively, and if observers do not alter the dynamics of I (i.e., $T_{\text{with obs}}|_I = T_{\text{without obs}}|_I$), then the closure condition is preserved in both cases. \square

The physical consequence is a form of realism about unobserved structure: the universe may contain dynamically stable, internally coherent regions that no observer ever samples, and these regions are part of the ontology in the same sense that observed regions are. Their inaccessibility is not a sign of their unreality; it is a sign of the constraint field's geometry.

The analogy with the decomposition of a dynamical system into invariant components is exact. A system with two invariant basins of attraction—say, the phase portrait of a bistable system—contains trajectories in each basin that never visit the other. The existence of each basin is independent of whether any trajectory visits it. Both basins are real components of the dynamical structure, even if a particular trajectory is confined to one.

16. EPISTEMOLOGY WITHOUT PRIVILEGED REALITY

The epistemic reading of the framework is the dual claim: what can be known is constrained by the observer's position within the field, not by a privileged access to reality in itself. This is not a form of idealism. It does not deny that there is a fact of the matter about the world independent of observers. It denies that any observer has access to that fact of the matter in a form unconditioned by its own constraint field.

The traditional epistemological picture, from Descartes through Kant to contemporary analytic philosophy, posits a relation between mind and world in which the mind receives inputs from the world and constructs representations. The representations may be more or less accurate, but the world is the standard against which accuracy is measured. There is, in principle, a view from nowhere: an idealized representation that captures the world as it is independent of any perspective.

The framework developed here denies the view from nowhere, but it denies it on structural rather than metaphysical grounds. The view from nowhere would require an observer with $\mathcal{O}_X = \mathcal{X}$ —an observer whose domain is the entire state space—and an injective observation map $\pi : \mathcal{X} \rightarrow \mathcal{M}$ that preserves all distinctions. The first condition would require the observer to exist at every state simultaneously, which violates the closure condition: an observer is a persistent subsystem, and persistent subsystems are proper subsets of \mathcal{X} in any non-trivial system. The second condition would require \mathcal{M} to have at least the cardinality of \mathcal{X} , which violates the representational character of observation: a representation that preserves all distinctions is not a representation but a copy.

Proposition 16.1 (Non-Existence of the View from Nowhere). *In any dynamical system (\mathcal{X}, T) with $|\mathcal{X}| > 1$, there is no observer $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$ such that (a) $\mathcal{O}_X = \mathcal{X}$, (b) π is injective, and (c) \mathcal{O} is persistent in the sense that $T(\mathcal{O}_X) \subseteq \mathcal{O}_X$. Condition (c) is trivially satisfied by (a) since $T(\mathcal{X}) \subseteq \mathcal{X}$ always holds, but condition (b) requires $|\mathcal{M}| \geq |\mathcal{X}|$, which means the measurement space is at least as large as the state space. An observer satisfying all three conditions would not be an observer in any meaningful sense—it would be a copy of the system itself.*

The result is not paradoxical; it simply formalizes the intuition that any genuine observer is finite and partial. What the formalization adds is precision about the source of the partiality: it is not a contingent limitation of current technology or cognitive capacity but a structural consequence of what it means to be a persistent subsystem with a representation space.

Knowledge, then, is path-dependent in a strong sense. What the observer knows is determined not only by the current state of its representations, but by the entire trajectory of its constraint-field through the state space. Two observers with different initial conditions, even in the same state space under the same dynamics, will in general have different reachable sets, different epistemically accessible fibers, and therefore different available explanations. Their knowledge is not merely quantitatively different (one has more evidence

than the other); it is structurally different (some explanations available to one are structurally unavailable to the other).

17. FINAL SYNTHESIS: CONSTRAINT AS THE COMMON SUBSTRATE

The thesis of this essay can now be stated in its final form. Being and knowing are not independent domains that happen to be related through a fortunate correspondence between mind and world. They are dual projections of a single underlying structure: the constrained dynamical field $(\mathcal{X}, T, \mathcal{C})$.

Persistence is the intrinsic property of this field: structures that satisfy $T(S) \subseteq S$ within $\mathcal{A}(\mathcal{C})$ exist as dynamical facts, independent of any observer. Function and observation are projections of this field: coordinate choices made by embedded subsystems that select certain invariants for representation and certain equivalence classes for identification. Knowledge is the progressive refinement of these projections, subject to the layered underdetermination established in Theorem 14.1 and the epistemic curvature established in Theorem 13.3.

The symmetry between being and knowing that emerges from this synthesis is not the symmetry of identity—being is not the same as knowing—but the symmetry of dual projections. Just as a three-dimensional object has multiple two-dimensional projections, none of which is the object itself, the constrained dynamical field has both an ontological projection (what persists) and an epistemic projection (what is known). Neither projection exhausts the field. Neither is prior to the other. Both are constrained by the same underlying structure.

Definition 17.1 (Ontological and Epistemic Projections). Given a constrained dynamical system $(\mathcal{X}, T, \mathcal{C})$:

The *ontological projection* is the map that assigns to each region $S \subseteq \mathcal{X}$ its persistence status: $\Phi_{\text{ont}}(S) = 1$ if $T(S) \subseteq S$ and $S \subseteq \mathcal{A}(\mathcal{C})$, and $\Phi_{\text{ont}}(S) = 0$ otherwise.

The *epistemic projection* for an observer \mathcal{O} is the map $\Phi_{\text{ep}} : \pi(\mathcal{R}) \rightarrow 2^{\pi^{-1}(\cdot) \cap \mathcal{R}}$ that assigns to each observation the set of accessible states consistent with it.

Both projections are defined on the same underlying field $(\mathcal{X}, T, \mathcal{C})$. The ontological projection is observer-independent; the epistemic projection is observer-relative.

The relationship between the two projections is mediated by the epistemic curvature: the geometry of the epistemic projection (the shape of the accessible region $\pi(\mathcal{R})$ and the structure of its fibers) is influenced by the ontological projection (which regions of \mathcal{X} are persistent and where they lie relative to the observer's boundary). The two projections are not independent; they are constrained by the same dynamics and the same constraint field.

This is the deepest result of the essay: not merely that being and knowing are related, but that they are related by the same mechanism that relates all persistent structures to each other—the mechanism of constraint and closure. *Constraint is prior; persistence is intrinsic; function and knowledge are projections.*

18. CLOSING REFLECTION

We began with three intuitions: that life improves, that organs are for things, that knowledge accumulates toward truth. Each felt like an observation rather than an interpretation. Each seemed to describe something that any careful examination of the world would confirm.

What the foregoing has shown is that all three are artifacts of position within a constraint field. The appearance of improvement in life is a consequence of conditioning on persistence, which inflates the apparent quality of survivors above the mean of the generative distribution. The appearance of intrinsic function in organisms is a consequence of applying an observer's coordinate system to a dynamical structure that is indifferent to any particular coordinate. The appearance of convergent knowledge is a consequence of the progressive refinement of a lifting map within a reachable set, which can appear to converge from within while remaining profoundly incomplete from without.

The inversion is now fully visible. What appears as direction is the shadow of a boundary: the constraint that excludes, which produces the appearance of selection toward, when observed from within the surviving set. What appears as purpose is the trace of a projection: the observer's coordinate system, which reads function into a system that contains only constraint-satisfying dynamics. What appears as progress in knowledge is the accumulation of trajectory within a reachable set, which looks like convergence toward truth from inside but is properly understood as the exploration of an accessible region whose boundary encodes the presence of the inaccessible.

The world, on this account, is not organized around observers or their interests. It is a field of constraint within which certain structures persist, among them the observer. The observer's experience of the world is shaped by its position within that field: by the constraints that define its domain, by the trajectory that defines its reachable set, and by the projection that defines its representations. Knowledge is the trace of that trajectory, refracted through that projection.

To ask what exists is to ask what persists under constraint. To ask what is known is to ask what is reachable from a given position under a given constraint field. These are not the same question, and their answers are not the same. But they are answers given within the same framework, about the same underlying structure. The separation between being and knowing, long treated as a foundational distinction in both physics and philosophy, is here revealed as a consequence of the geometry of constraint: a distinction between an intrinsic property of the field and a relational property of embedded subsystems, both of which are equally features of the single constrained dynamical system that contains both world and observer.

What the observer sees is not the world. It is the world's projection onto the observer's reachable set. The world is larger than any projection. It is larger than any observer. And it persists, whether or not it is seen, because persistence requires only closure—not witness.

PART V CONSEQUENCES AND EXTENSIONS

18. CONSTRAINT WITHOUT TELEOLOGY

The framework developed in this essay removes teleology not by denying structure, but by relocating its source. What appears as direction in a system is not the result of a goal toward which the system is moving, but the result of constraints that eliminate incompatible trajectories. Direction is not intrinsic to the dynamics; it is the shape of what remains after exclusion.

This inversion has a general form. Let P be any property of trajectories in (\mathcal{X}, T) that appears to exhibit directionality—*increase, improvement, convergence*. If P is correlated with persistence under \mathcal{C} , then conditioning on persistence will produce an apparent gradient in P even if no such gradient exists in the underlying dynamics. The direction is therefore not a feature of T , but of the conditional distribution induced by \mathcal{C} .

Proposition 18.1 (Apparent Direction as Conditional Artifact). *Let $P : \mathcal{X} \rightarrow \mathbb{R}$ be any measurable property positively correlated with persistence under \mathcal{C} in the sense of Proposition 3.1. Then the expected value of P along persistent trajectories exceeds its expectation over all trajectories:*

$$\mathbb{E}_\mu[P \mid E] > \mathbb{E}_\mu[P]. \quad (13)$$

Any observed monotonicity in P along surviving trajectories may therefore arise from conditioning on E rather than from an intrinsic gradient in T .

Proof. This is an immediate application of Proposition 3.1 with $Q = P$: positive correlation between P and $\mathbf{1}_E$ yields $\mathbb{E}_\mu[P \mid E] > \mathbb{E}_\mu[P]$. \square

Proposition 18.1 generalizes the survivor-bias result from Part I to arbitrary measurable properties. It establishes that any property correlated with persistence will appear to trend upward in the sample of survivors, including properties that have nothing to do with the generating mechanism. The elimination of teleology is thus not eliminative but explanatory: it reveals that the source of apparent direction lies in the conditional distribution, not in the dynamics.

19. HIERARCHIES OF CLOSURE AND THE EMERGENCE OF COMPLEXITY

The framework also clarifies the emergence of hierarchical organization in living systems. Higher levels of organization do not arise because they are more optimal in an abstract sense, but because they permit more robust forms of closure.

Let S_1, S_2, \dots, S_n be persistent structures at different levels of description, with S_{i+1} defined over configurations of elements in S_i . A hierarchy emerges when closure at level $i + 1$ depends on the coordinated behavior of elements at level i , and when this coordination itself satisfies a closure condition.

Definition 19.1 (Hierarchical Closure). A sequence of persistent sets $\{S_i\}_{i=1}^n$ exhibits *hierarchical closure* if for each $i < n$ there exists a map $\Phi_i : S_i \rightarrow S_{i+1}$ such that

$$T_{i+1}(\Phi_i(S_i)) \subseteq \Phi_i(S_i) \quad (14)$$

and Φ_i preserves admissibility: $\Phi_i(S_i) \subseteq \mathcal{A}(\mathcal{C}_{i+1})$ where \mathcal{C}_{i+1} is the constraint family at level $i + 1$.

Hierarchical complexity arises when closure at a higher level stabilizes the dynamics at lower levels. A cell membrane stabilizes the chemical reactions within it; an organism stabilizes its constituent cells; an ecosystem stabilizes its constituent organisms. In each case, the higher-level closure expands the set of perturbations under which the lower-level dynamics remain admissible. Complexity persists not because it is intrinsically valuable, but because it enlarges the basin of closure—the set of initial conditions from which trajectories remain within the admissible region.

Proposition 19.2 (Complexity as Robustness of Closure). *Let S_{simple} and S_{complex} be two persistent sets with $S_{\text{simple}} \subseteq S_{\text{complex}}$. If the constraint family \mathcal{C} is subject to perturbations $\mathcal{C} \mapsto \mathcal{C}_\epsilon$ for $\epsilon > 0$, then the probability that S_{complex} remains persistent under \mathcal{C}_ϵ is at least as large as the probability that S_{simple} remains persistent, provided that the hierarchical structure of S_{complex} admits more redundant pathways to closure.*

This proposition is informal in the sense that it depends on a precise specification of what counts as “redundant pathways,” but its qualitative content is clear and matches the empirical pattern: more complex biological systems tend to be more robust to environmental perturbation, not because robustness is selected for directly, but because complex hierarchical closure provides more routes to satisfying the persistence condition when individual routes are disrupted.

20. PROJECTION, COMPRESSION, AND THE LIMITS OF REPRESENTATION

Functional description and observation both involve projection, which is necessarily compressive. The projection $\pi : \mathcal{O}_X \rightarrow \mathcal{M}$ reduces the dimensionality of the state space, collapsing distinctions. This compression is not merely a loss; it is a structural necessity for representation.

Let $d_{\mathcal{X}} = \dim(\mathcal{X})$ and $d_{\mathcal{M}} = \dim(\mathcal{M})$ denote the effective dimensions of the state and measurement spaces in the case where both carry smooth manifold structure. In any non-trivial observer, $d_{\mathcal{M}} < d_{\mathcal{X}}$. The projection therefore induces a many-to-one mapping, and the fibers $\pi^{-1}(m)$ are submanifolds of dimension $d_{\mathcal{X}} - d_{\mathcal{M}}$.

Proposition 20.1 (Compression-Induced Ambiguity). *If $\pi : \mathcal{X} \rightarrow \mathcal{M}$ is a smooth surjective map between manifolds of dimensions $d_{\mathcal{X}} > d_{\mathcal{M}}$, then for almost*

all $m \in \mathcal{M}$ (in the measure-theoretic sense), the fiber $\pi^{-1}(m)$ is a submanifold of \mathcal{X} of positive dimension $d_{\mathcal{X}} - d_{\mathcal{M}}$, and therefore has positive measure in \mathcal{X} . The ambiguity in lifting observations to states is therefore generic, not exceptional.

Proof. By the regular value theorem (Sard's theorem), almost every $m \in \mathcal{M}$ is a regular value of π , meaning $\pi^{-1}(m)$ is a smooth submanifold of \mathcal{X} of dimension $d_{\mathcal{X}} - d_{\mathcal{M}} > 0$. A submanifold of positive dimension has positive measure with respect to the ambient measure on \mathcal{X} . Therefore, for almost every observation m , the fiber $\pi^{-1}(m)$ has positive measure, and the lifting problem has uncountably many solutions. \square

Proposition 20.1 formalizes the inevitability of underdetermination. The observer cannot avoid it by improving precision or acquiring more data, because the ambiguity is a structural consequence of representing a higher-dimensional space in a lower-dimensional one. Additional observations reduce the effective fiber (by restricting to the epistemically accessible portion $\pi^{-1}(m) \cap \mathcal{R}$), but they cannot eliminate it. The residual underdetermination is the dimensional gap $d_{\mathcal{X}} - d_{\mathcal{M}}$, which is not closed by any finite sequence of observations.

21. TEMPORAL ASYMMETRY AND THE DIRECTION OF INFERENCE

The framework also clarifies the origin of temporal asymmetry in inference. While the dynamics T may in principle be time-reversible—as is the case for Hamiltonian dynamics in classical mechanics—the process of observation is not. The reachable set \mathcal{R} is constructed forward in time from initial conditions, and the lifting of observations depends on this trajectory.

Definition 21.1 (Forward and Backward Reachability). Given initial state $x_0 \in \mathcal{O}_{\mathcal{X}}$, the *forward reachable set* is

$$\mathcal{R}^+(x_0) = \{T^n(x_0) : n \geq 0\}. \quad (15)$$

If T is invertible, the *backward reachable set* is

$$\mathcal{R}^-(x_0) = \{T^{-n}(x_0) : n \geq 0\} = \{x : \exists n \geq 0, T^n(x) = x_0\}. \quad (16)$$

In general, $\mathcal{R}^+(x_0) \neq \mathcal{R}^-(x_0)$ even when T is invertible, because the constraint family \mathcal{C} may not be symmetric under time reversal. The observer's

accumulation of representations is a forward process: at time n , the observer has sampled $\{x_0, T(x_0), \dots, T^n(x_0)\}$ and constructed representations in \mathcal{M} of each. The epistemically accessible fiber at time n is

$$\pi^{-1}(m_n) \cap \mathcal{R}^+(x_0)|_{[0,n]} = \pi^{-1}(m_n) \cap \{T^k(x_0) : 0 \leq k \leq n\}. \quad (17)$$

This fiber shrinks as n increases, because additional observations provide additional constraints on which states in $\pi^{-1}(m_n)$ are consistent with the full observed trajectory. The apparent direction of knowledge—the sense in which understanding accumulates over time—is a consequence of this forward construction. It does not imply that there is an intrinsic temporal direction in the underlying dynamics; it implies that the observer is a forward-running process.

Proposition 21.2 (Asymmetry of Epistemic Access). *The epistemically accessible fiber $\pi^{-1}(m_n) \cap \mathcal{R}^+(x_0)|_{[0,n]}$ is in general a proper subset of $\pi^{-1}(m_n) \cap \mathcal{R}^-(x_0)$. The observer's forward trajectory excludes states that are backward-reachable but not forward-reachable, generating an asymmetry in available explanations that is not a feature of the dynamics but of the directionality of observation.*

Temporal direction in knowledge is therefore not imposed by physical law but by the structure of constrained observation. The arrow of time in epistemology is the arrow of forward reachability under constraint, not a fundamental feature of the dynamical equations.

22. CONSTRAINT GEOMETRY AND PHYSICAL LAW

The final extension concerns the interpretation of physical law itself. In the framework developed here, laws are not external prescriptions imposed on a system, but descriptions of the constraint geometry that defines admissible regions and governs which dynamics are compatible with persistence.

Let \mathcal{C} be specified by a family of constraint functions $\{\Phi_i : \mathcal{X} \rightarrow \mathbb{R}\}_{i \in I}$ such that the admissible region is

$$\mathcal{A}(\mathcal{C}) = \{x \in \mathcal{X} : \Phi_i(x) \leq 0 \text{ for all } i \in I\}. \quad (18)$$

The geometry of $\mathcal{A}(\mathcal{C})$ is determined by the level sets of the Φ_i . For a continuous-time dynamical system generated by a vector field $V : \mathcal{X} \rightarrow T\mathcal{X}$, the condition that trajectories remain within $\mathcal{A}(\mathcal{C})$ is the requirement that V point inward or tangentially at the boundary $\partial\mathcal{A}(\mathcal{C})$.

Proposition 22.1 (Physical Law as Tangency Condition). *Let \mathcal{C} be defined by smooth constraint functions $\{\Phi_i\}$. A smooth vector field V on \mathcal{X} satisfies the constraints—meaning $x(t) \in \mathcal{A}(\mathcal{C})$ for all $t \geq 0$ whenever $x(0) \in \mathcal{A}(\mathcal{C})$ —if and only if for all $x \in \partial\mathcal{A}(\mathcal{C})$ and all active constraints i (those with $\Phi_i(x) = 0$),*

$$\langle \nabla \Phi_i(x), V(x) \rangle \leq 0. \quad (19)$$

Proof. The condition $\Phi_i(x(t)) \leq 0$ for all $t \geq 0$ requires that $\frac{d}{dt}\Phi_i(x(t)) \leq 0$ whenever $\Phi_i(x(t)) = 0$. By the chain rule, $\frac{d}{dt}\Phi_i(x(t)) = \langle \nabla \Phi_i(x(t)), \dot{x}(t) \rangle = \langle \nabla \Phi_i(x(t)), V(x(t)) \rangle$. The condition therefore becomes $\langle \nabla \Phi_i(x), V(x) \rangle \leq 0$ at all boundary points where $\Phi_i = 0$. \square

Proposition 22.1 expresses physical law as a geometric condition: the vector field generating the dynamics must not point outward from the admissible region at its boundary. Conservation laws, symmetry principles, and equations of motion can all be understood as specifying the constraint geometry within which dynamics must remain. The laws of physics, on this reading, are not imposed from outside; they are the description of which vector fields are compatible with the admissible region defined by the constraint family.

This reinterpretation connects the framework to variational principles. The Lagrangian and Hamiltonian formulations of classical mechanics can be understood as selecting, among all vector fields satisfying the tangency condition, those that extremize a certain functional. The variational principle is a selection criterion within the space of admissible dynamics—itsself a space defined by constraint geometry. In this sense, even the choice of dynamics within the admissible region is governed by a higher-level constraint: the constraint that the dynamics be variational.

23. UNIVERSALITY OF THE CONSTRAINT-PROJECTION STRUCTURE

The framework developed in this essay is minimal but universal. It applies to any system that can be described as a constrained dynamical field. This includes not only biological and cognitive systems but physical systems at all scales: from chemical reaction networks to cellular automata, from neural architectures to cosmological models.

The universality arises from the fact that the three operations used throughout—closure, projection, and reachability—are not domain-specific. They are struc-

tural operations that can be defined on any set equipped with a dynamics and a constraint family.

Theorem 23.1 (Universality of the Constraint-Projection Structure). *Let $(\mathcal{X}, T, \mathcal{C})$ be any constrained dynamical system. Then the following hold without further assumptions:*

1. (Persistence) *The collection of persistent subsets $\{S \subseteq \mathcal{A}(\mathcal{C}) : T(S) \subseteq S\}$ is closed under arbitrary intersection, and therefore the intersection of all persistent subsets containing a given point x is itself persistent.*
2. (Functional Description) *For any persistent subset S , the collection of functional projections $\pi_F : (S, T|_S) \rightarrow (F, T_F)$ is non-empty (the identity is always a functional projection) and forms a category under composition of morphisms.*
3. (Observation) *For any observer subsystem $\mathcal{O} = (\mathcal{O}_X, T_{\mathcal{O}}, \pi)$, the reachable set $\mathcal{R} \subseteq \mathcal{O}_X$ is forward-invariant under $T_{\mathcal{O}}$, and the image $\pi(\mathcal{R}) \subseteq \mathcal{M}$ is the set of observations accessible to \mathcal{O} .*

These three structures—persistent sets, functional projections, and observable images—are well-defined for any $(\mathcal{X}, T, \mathcal{C})$ and require no additional primitives.

Proof. Part (1): If $\{S_\alpha\}$ is a collection of persistent subsets, then for any $x \in \bigcap_\alpha S_\alpha$, we have $T(x) \in S_\alpha$ for each α (since $T(S_\alpha) \subseteq S_\alpha$), so $T(x) \in \bigcap_\alpha S_\alpha$. Thus $T(\bigcap_\alpha S_\alpha) \subseteq \bigcap_\alpha S_\alpha$, and the intersection is persistent provided it lies in $\mathcal{A}(\mathcal{C})$, which follows since each $S_\alpha \subseteq \mathcal{A}(\mathcal{C})$ and $\mathcal{A}(\mathcal{C})$ is closed under intersection. Part (2): The identity map $\text{id} : S \rightarrow S$ is always a surjective endomorphism, hence a functional projection. Composition of morphisms of dynamical systems is itself a morphism, so the collection is closed under composition and forms a category. Part (3): \mathcal{R} is forward-invariant by definition ($T^n(x_0) \in \mathcal{R}$ for all $n \geq 0$), and $\pi(\mathcal{R})$ is the direct image of \mathcal{R} under π , which is well-defined for any map. \square

The consequence of Theorem 23.1 is that the distinctions between biological, cognitive, and physical domains are not fundamental. They arise from different projections of the same underlying structure. The essay has moved between these domains not because the argument changes, but because the same argument appears differently when different projections are applied. Biology, philosophy of mind, and epistemology are coordinate systems over a single constraint-governed field.

24. CLOSING EXTENSION

The arc of the argument can now be seen in full. Selection does not optimize; it filters. Function is not intrinsic; it is projected. Observation is not neutral; it is constrained. Knowledge is not complete; it is path-dependent. Complexity is not progressive; it is a consequence of hierarchical closure. Physical law is not imposed; it is the tangency condition of the constraint geometry. These are not separate claims about separate domains. They are different faces of a single structure, revealed by applying the same three operations—closure, projection, reachability—at different levels of description.

The framework has a further consequence that should be stated explicitly, though it cannot be fully developed here. If the constraint geometry of physical law is itself subject to variation—if the Φ_i are not fixed but evolve—then the admissible region $\mathcal{A}(\mathcal{C})$ changes over time, and with it the set of persistent structures. This opens the possibility that the constraint geometry itself has a dynamics, and that the emergence of new forms of persistence corresponds to the evolution of new constraint structures. In this direction, the framework points toward a dynamical theory of laws, in which what counts as physically possible is itself subject to constraint-governed change. This remains a direction for future development.

What can be said here, without speculation, is this. The world is not organized around what is seen, what is useful, or what is known. It is organized around what does not fail under constraint. Everything else—purpose, design, explanation, progress—is the structure that appears when a constrained subsystem moves through that field and attempts to describe it.

Constraint is prior. Persistence is intrinsic. Function and knowledge are projections.

What remains unseen is not absent. It is present as boundary, as curvature, as the shape of what can be reached and what cannot. The geometry of the accessible is the shadow of the inaccessible. And it is within that shadow that all claims about progress, function, and knowledge must be understood: not as discoveries about the world, but as features of a path taken through it.

PART VI RSVP EXTENSION: CONSTRAINT FIELDS AND STRUCTURED PERSISTENCE

24. FROM ABSTRACT CONSTRAINT TO STRUCTURED FIELD

The framework developed in Parts I through V has treated constraint abstractly, as a family \mathcal{C} of admissible subsets of the state space \mathcal{X} . This abstraction is sufficient to establish the logical structure of persistence, projection, and reachability without importing assumptions specific to any physical domain. It does not, however, specify how constraints are organized, how they interact with one another, or how they evolve as the system moves through state space.

The Relativistic Scalar–Vector Plenum (RSVP) framework provides a structured realization of \mathcal{C} as a field over \mathcal{X} , giving the abstract constraint family a concrete geometric and dynamical interpretation. In RSVP, the state of the system is not given solely by a point $x \in \mathcal{X}$ but by a triple

$$X(t) = (\Phi(t), \mathbf{v}(t), S(t)), \quad (20)$$

where Φ is a scalar field representing constraint density or admissibility pressure, \mathbf{v} is a vector field representing directed organization or flow through the state space, and S is an entropy field representing unresolved incompatibility or degeneracy within the current configuration. In this formulation, the admissible region $\mathcal{A}(\mathcal{C})$ is not specified by an arbitrary family of sets but emerges from the joint structure of (Φ, \mathbf{v}, S) : constraint becomes a field-theoretic object rather than a set-theoretic one, and the geometry of persistence inherits the geometry of the field.

This refinement does not change the conclusions of the preceding parts. It sharpens them by locating the abstract operations of closure, projection, and reachability within a physically interpretable substrate.

25. PERSISTENCE AS FIELD COHERENCE

The closure condition $T(S) \subseteq S$ can now be refined in RSVP terms. A persistent structure is one whose field configuration remains self-consistent under the coupled evolution of (Φ, \mathbf{v}, S) .

Definition 25.1 (RSVP Persistence). A region $R \subseteq \mathcal{X}$ is *RSVP-persistent* if its

associated fields satisfy the evolution equations

$$\partial_t \Phi = F_\Phi(\Phi, \mathbf{v}, S), \quad \partial_t \mathbf{v} = F_v(\Phi, \mathbf{v}, S), \quad \partial_t S = F_S(\Phi, \mathbf{v}, S), \quad (21)$$

with the additional requirement that the evolution preserves admissibility:

$$(\Phi, \mathbf{v}, S)(t) \in \mathcal{A}(\mathcal{C}) \quad \text{for all } t \geq 0. \quad (22)$$

In this form, persistence is no longer merely set closure but field coherence: the system must maintain a mutually compatible configuration of constraint density, directional flow, and entropy throughout its evolution. The closure condition $\nabla \Phi \cdot \mathbf{v} \leq 0$ on persistent structures, established in the tangency condition of the preceding part, now admits the reading that the flow field \mathbf{v} must not drive the system toward increasing constraint violation. The scalar field Φ encodes the shape of the admissible region, and the vector field \mathbf{v} must remain compatible with that shape at all times.

This refinement clarifies the nature of robustness. A structure is robust not because it occupies a region of high fitness but because its field configuration admits a wide basin of closure under perturbations in (Φ, \mathbf{v}, S) . Robustness is therefore a geometric property of the constraint field, not a property of any scalar objective.

26. SELECTION AS ENTROPIC FILTERING IN THE FIELD

Within RSVP, the interpretation of selection becomes more precise. The entropy field S governs admissibility by registering incompatibility between the current configuration and the constraints. Regions of high S correspond to configurations in which local field values cannot be globally integrated without contradiction; regions of low S correspond to coherent, mutually compatible configurations.

Selection is expressed as an entropic filtering process in which configurations with diverging entropy are excluded from persistence:

$$\frac{dS}{dt} \leq 0 \quad \text{on persistent structures,} \quad (23)$$

while configurations whose entropy grows without bound are driven outside the admissible region and eliminated. This makes explicit what was implicit in Part I: selection does not maximize Φ or optimize \mathbf{v} . It enforces admissibility through the entropy field. Apparent optimization arises when

low-entropy field configurations correlate with persistence, but the governing mechanism is entropic exclusion, not ascent toward a goal.

The survivor-bias result of Proposition 3.1 now acquires a field interpretation. The event E of persistence corresponds to the condition that entropy remains bounded, and any property Q correlated with low entropy will appear elevated in the surviving sample. The direction of evolution is the direction of decreasing entropy within the constraint field, not the direction of any external criterion.

27. FUNCTION AS PROJECTION IN RSVP FIELDS

Functional description can be recast as a projection of the RSVP field configuration onto a reduced representation space. A functional description is a map

$$\pi_F : (\Phi, \mathbf{v}, S) \mapsto F, \quad (24)$$

which selects a subset of field invariants and represents them in a lower-dimensional space F . Different choices of π_F correspond to different functional interpretations of the same underlying field. A projection that retains local structure in Φ yields a description in terms of constraint density, appropriate to molecular-level function. A projection that retains flow patterns in \mathbf{v} yields a description in terms of dynamical roles, appropriate to cellular-level function. A projection that retains entropy gradients in S yields a description in terms of information processing or coherence, appropriate to cognitive-level function. No projection is privileged by the field itself.

The underdetermination result of Proposition 6.1 is therefore strengthened in the RSVP setting. Function is not merely projection-dependent in the abstract sense that different invariants of a persistent set can be chosen as functions. It is field-component dependent: different observers, attending to different components of (Φ, \mathbf{v}, S) , will identify different functions even when examining the same persistent structure. The relativization of function is therefore not a formal artifact but a physical consequence of the multi-component nature of the constraint field.

28. OBSERVATION AS FIELD-CONSTRAINED TRAJECTORY

In RSVP, the observer is a localized, persistent configuration of the field, satisfying RSVP closure within its own domain. Its reachable set \mathcal{R} is determined by the evolution of its internal field configuration $(\Phi_{\mathcal{O}}, \mathbf{v}_{\mathcal{O}}, S_{\mathcal{O}})$, and observation is a projection

$$\pi : \mathcal{R} \rightarrow \mathcal{M} \quad (25)$$

defined over the reachable trajectory of that configuration. Observation is therefore a process of sampling the global field through a local field configuration, and the limitations of observation decompose into three field-theoretic components. Restrictions on $\Phi_{\mathcal{O}}$ determine which regions of constraint density are accessible to the observer. Restrictions on $\mathbf{v}_{\mathcal{O}}$ determine which trajectories through state space the observer can follow. Restrictions on $S_{\mathcal{O}}$ determine which distinctions the observer can resolve without violating its own coherence. The layered underdetermination of Theorem 14.1 thus reflects the joint limitations imposed by all three field components simultaneously.

29. ISLANDS AND FIELD DISCONNECTION

In the RSVP framework, inaccessible persistent regions correspond to field configurations separated from the observer's domain by constraint barriers in Φ and entropy gradients in S . Such regions are not merely spatially distant; they are field-disconnected, meaning that no admissible trajectory of the observer's field configuration $(\Phi_{\mathcal{O}}, \mathbf{v}_{\mathcal{O}}, S_{\mathcal{O}})$ can reach them without first violating either the constraint field Φ or the entropy bound S .

Definition 29.1 (RSVP Island). A region $I \subseteq \mathcal{X}$ is an *RSVP island* with respect to an observer \mathcal{O} if I is RSVP-persistent and if no trajectory generated by $\mathbf{v}_{\mathcal{O}}$ within the observer's domain intersects I , where the barrier to intersection is constituted by the constraint geometry of Φ or by entropy incompatibility in S .

The epistemic curvature theorem of Part III acquires a direct physical interpretation here. The constraint potential Φ induces curvature in the geometry of accessible trajectories, and the Hessian $\nabla^2\Phi$ governs the deformation of those trajectories near constraint boundaries. Inaccessible persistent regions appear in the structure of Φ on the far side of these boundaries, and their

presence deforms the geometry of Φ on the near side in ways that are in principle detectable by the observer through anomalies in the curvature of its reachable domain. The unseen is encoded in the geometry of the seen because both are governed by the same field.

30. RSVP UNIFICATION

The abstract framework of this essay and the RSVP framework align precisely, with each abstract concept finding a concrete field-theoretic realization. The state space \mathcal{X} is realized as the space of field configurations (Φ, \mathbf{v}, S) . The constraint family \mathcal{C} is realized as the joint structure of constraint density Φ and entropy S . The dynamics T are realized as the coupled field evolution equations. Persistence becomes field coherence under those equations. Functional description becomes projection of field components. Observation becomes the local field trajectory of a persistent subfield. Reachability becomes the support of trajectories generated by \mathbf{v} within the domain defined by Φ and S .

Theorem 30.1 (RSVP Unification). *Let $\mathcal{F} = (\Phi, \mathbf{v}, S)$ be a constraint field over \mathcal{X} . Persistent structures are regions where \mathbf{v} is compatible with Φ and does not increase S . Selection is the restriction of trajectories to such regions, with no reference to optimization. Function is any projection of \mathcal{F} onto a lower-dimensional representation space. Observation is a projection defined by a persistent subfield with restricted flow. Knowledge is trajectory-dependent inference within that restricted flow, subject to entropy accumulation and constraint curvature. None of these require the introduction of a scalar objective, a privileged observer, or an externally specified goal.*

Proof. Persistence: the condition \mathbf{v} compatible with Φ is the tangency condition $\nabla\Phi \cdot \mathbf{v} \leq 0$, which is the RSVP form of the closure condition $T(S) \subseteq S$ established in Proposition 22.1. Non-increase of S on persistent structures follows from the entropic filtering condition. Function: by Proposition 6.1, any invariant of the persistent structure can serve as a functional description, and different projections of (Φ, \mathbf{v}, S) realize different such invariants. Observation: the observer is a persistent subfield satisfying RSVP closure, and its observation map π is the restriction of the global projection to the subfield's reachable set. Knowledge: by Theorem 14.1, knowledge is bounded by state underdetermination, structural inaccessibility, and functional underdetermination, all of which have field-theoretic realizations in the RSVP structure. \square

The consequence of this unification is that physics, biology, and epistemology are not separate domains requiring separate frameworks. They are regimes of the same constraint field viewed under different projections. The apparent differences arise from the position of the observer within the field, not from differences in the underlying mechanism.

PART VII AGAINST THE OPTIMIZATION INTERPRETATION: A DIRECT ARGUMENT

32. THE OPTIMIZATION ASSUMPTION AND ITS REQUIREMENTS

The dominant interpretation of natural selection treats it as an optimization process operating over a fitness function. In this view, a population evolves on a landscape $f : \mathcal{X} \rightarrow \mathbb{R}$, and selection is understood as a mechanism that drives the system toward regions of higher f . This interpretation pervades not only popular accounts of evolution but also formal treatments in evolutionary biology and theoretical ecology, where the language of fitness maximization is standard.

For this interpretation to be well-founded, three conditions must hold simultaneously. First, there must exist a well-defined objective function f whose values rank configurations against one another. Second, there must be a dynamical tendency in the system toward increasing f , so that the trajectory of evolution under selection moves toward high- f regions. Third, there must be a correspondence between persistence and high values of f , so that what is observed as surviving is what ranks highly under the objective. The optimization interpretation requires all three, and each is independently problematic.

The fundamental difficulty is that the mechanism of natural selection satisfies none of these conditions as a consequence of its definition. Selection enforces a boundary condition: configurations that violate constraints are eliminated. This is a restriction on which trajectories continue, not a gradient that guides trajectories toward any particular destination. The argument of Part I established this formally through the closure condition, Lemmas 2.5 and 2.6, and Theorem 4.2. The present part develops the critique further by targeting specific features of the optimization interpretation that survive those lemmas.

33. PERSISTENCE DOES NOT IMPLY ASCENT

The closure condition $T(S) \subseteq S$ is invariant under arbitrary reparameterizations of the state space. For any function $f : \mathcal{X} \rightarrow \mathbb{R}$, there is no requirement derived from the closure condition that $f(T(x)) \geq f(x)$ along persistent trajectories. This was the content of Lemma 2.5. But the point deserves amplification in the biological context.

Suppose a lineage persists for one thousand generations. The closure condition guarantees only that the lineage's states remained within the admissible region throughout. It does not imply that any measurable property of the lineage—morphological complexity, metabolic efficiency, genome size, behavioral repertoire—increased monotonically. In fact, the fossil record contains abundant examples of evolutionary change that reduce measurable indicators of complexity, efficiency, or structural elaboration, while the lineage persists robustly. Regressive evolution, simplification of life cycles, and reduction of metabolic pathways under stable environmental conditions are all consistent with the closure condition and inconsistent with any simple optimization narrative.

The persistence of a lineage over geological time is therefore evidence only of the following: the lineage's states satisfied the constraints defining the admissible region throughout the relevant period. No stronger inference about optimization or directed change is licensed by the mechanism.

34. THE ABSENCE OF A CANONICAL FITNESS FUNCTION

For selection to constitute optimization, there must exist a privileged function f whose maximization explains persistence. The difficulty is that given any persistent set S , infinitely many functions can be constructed that are maximized on S . The indicator function $f(x) = \mathbf{1}_S(x)$ is one such construction; any smooth function that takes its maximum on S and decays away from S is another; any weighted combination of the infinitely many invariants established in Proposition 6.1 is another. The existence of a function consistent with persistence is therefore entirely uninformative about the generative process.

Any attempt to identify a fitness function after the fact reduces to encoding persistence itself. One observes which lineages persist, constructs a function that ranks them highly, and then interprets the function as the objective that selection was tracking. The reasoning is circular: the function is derived

from the outcome it is supposed to explain. This is precisely the retrospective projection identified in Corollary 7.1, now applied to the fitness concept itself. The fitness function is not a causal driver of persistence; it is a representation of persistence in a coordinate system chosen by the observer.

35. NON-TRANSITIVITY AND THE IMPOSSIBILITY OF A SCALAR OBJECTIVE

A more decisive objection is that natural selection operates in regimes where fitness comparisons are non-transitive, and no scalar function can represent non-transitive orderings. Frequency-dependent selection provides the clearest examples. In many systems, the success of a trait depends on its prevalence relative to other traits in the population. When this dependence is strong, it generates cyclic dominance relations: configuration x outcompetes y when y is common, y outcompetes z when z is common, and z outcompetes x when x is common. No scalar function f can represent such cycles, because a scalar ordering is by definition acyclic.

The mathematical underpinning is precise. Let \succ_P denote the dominance relation under population state P . If there exist configurations x, y, z and population states P_1, P_2, P_3 such that $x \succ_{P_1} y$, $y \succ_{P_2} z$, and $z \succ_{P_3} x$, then there is no function $f : \mathcal{X} \rightarrow \mathbb{R}$ satisfying $f(x) > f(y)$, $f(y) > f(z)$, and $f(z) > f(x)$ simultaneously. The dominance structure cannot be embedded in a total order. Frequency-dependent selection produces exactly this situation in rock-paper-scissors dynamics, and analogous structures appear in predator-prey systems, host-parasite coevolution, and social evolution with strategic interaction.

Path dependence introduces a further obstruction. The dynamics of selection are in general not Markovian in the state of the population alone; the future evolution of a lineage depends on its history of environmental exposures, developmental trajectories, and prior selection events. This dependence violates the assumption required for a well-defined fitness landscape, namely that fitness is a function of the current state. When path dependence is present, the state space must be extended to include history, and the resulting high-dimensional system cannot in general be represented as ascent on a scalar surface.

36. FITNESS LANDSCAPES AS COORDINATE ARTIFACTS

The fitness landscape, as a representation, requires that evolutionary dynamics take place on a fixed, low-dimensional surface whose topology determines the accessible optima. This representation requires several conditions that are generically violated. Environmental stability is required because a changing environment reshapes the landscape while the population is traversing it; under change, there is no guarantee that uphill movement in the old landscape corresponds to persistence in the new one. Additivity of fitness contributions is required because without it, the landscape becomes an astronomically high-dimensional object with no useful low-dimensional projection. Transitivity is required as argued above. And limited path dependence is required so that the current state of the population is sufficient to determine its future.

When any of these conditions fails, the landscape metaphor misleads rather than guides. The admissible region $\mathcal{A}(\mathcal{C})$ and the closure condition $T(S) \subseteq S$ contain no hidden assumptions of this kind. They describe the actual constraint structure without presupposing the additional regularity required for a scalar objective. The landscape is therefore not the fundamental object; it is a projection of the constraint geometry onto a one-dimensional coordinate, valid only when the constraint structure has the special regularity needed to support that projection.

37. WHEN OPTIMIZATION EMERGES AS AN APPROXIMATION

If natural selection is not fundamentally an optimization process, the question arises of why optimization appears so often as an adequate description. The answer is that optimization is an emergent approximation that arises when the constraint structure possesses sufficient regularity to admit reduction to a scalar coordinate. This emergence can be stated precisely.

Suppose there exists a function $f : \mathcal{X} \rightarrow \mathbb{R}$ satisfying two conditions: first, persistence implies high f , in the sense that $x \in S$ implies $f(x) \geq \alpha$ for some threshold α ; and second, the dynamics approximately preserve ordering, in the sense that $f(T(x)) \gtrsim f(x)$ along trajectories. Under these conditions, the system behaves as though it were maximizing f , and the optimization description is approximately valid.

Proposition 37.1 (Emergence of Optimization). *Let $(\mathcal{X}, T, \mathcal{C})$ be a constrained dynamical system. Suppose there exists $f : \mathcal{X} \rightarrow \mathbb{R}$ such that $S \subseteq f^{-1}([\alpha, \infty))$ for*

the persistent set S , and such that $f \circ T \geq f - \epsilon$ on S for some small $\epsilon > 0$. Then the dynamics on S are approximately described as ascent on f with error bounded by ϵ . The function f is not intrinsic to the dynamics; it is a coordinate on the constraint manifold that preserves approximate ordering.

Proof. The condition $f \circ T \geq f - \epsilon$ means that f decreases by at most ϵ per step along trajectories in S . Over n steps, f decreases by at most $n\epsilon$. For small ϵ and moderate n , the trajectory remains near the level set $f^{-1}(\alpha)$ and above it, producing behavior indistinguishable from approximate maximization of f . \square

Optimization is therefore a coordinate choice that becomes approximately valid when the constraint geometry is sufficiently regular. It requires environmental stability so that the constraint structure does not shift during the observation window, weak interaction so that fitness contributions are approximately additive, transitivity so that pairwise comparisons extend to a global order, and limited path dependence so that the current state is approximately sufficient. These are real conditions that are sometimes satisfied, particularly over short timescales and in stable environments. When they are satisfied, the optimization description provides a useful compression of the constraint dynamics. When they fail, the description misleads.

The deeper point is that even when optimization provides a valid approximate description, it remains a representational choice, not an intrinsic feature of the system. Different choices of f that preserve approximate ordering are all equally valid, and no dynamics-internal criterion distinguishes among them. Optimization is the shadow of constraint closure under a particular projection, not the mechanism of selection itself.

PART VIII SELECTION, OBSTRUCTION, AND THE SHEAF OF COHERENT MODELS

37. OBSERVATIONS AS LOCAL SECTIONS

The preceding parts have treated observation as a projection $\pi : \mathcal{R} \rightarrow \mathcal{M}$ from the observer's reachable set into a measurement space. A single observation $m \in \mathcal{M}$ lifts to a fiber $\pi^{-1}(m) \cap \mathcal{R}$ of accessible states. But observations are not isolated events; they form sequences whose joint structure imposes

constraints that no individual observation captures. To formalize this, we introduce the sheaf-theoretic language of local sections and global coherence.

Let $\{U_i\}_{i \in I}$ be an open cover of the measurement space \mathcal{M} , and for each U_i let $\mathcal{F}(U_i)$ denote the set of locally admissible field configurations consistent with observations in U_i . The assignment $U \mapsto \mathcal{F}(U)$ defines a presheaf over \mathcal{M} : if $V \subseteq U$, restriction maps $\rho_{UV} : \mathcal{F}(U) \rightarrow \mathcal{F}(V)$ send global sections to their restrictions on smaller domains. A local section $s_i \in \mathcal{F}(U_i)$ is a field configuration consistent with the observations available in U_i .

A global model of the world is a global section $s \in \mathcal{F}(\mathcal{M})$ such that $s|_{U_i} = s_i$ for all i : a single field configuration that is simultaneously consistent with all locally observed constraints. The question of whether a global model exists is the question of whether the local sections can be coherently glued into a global one. When gluing succeeds, the local observations are jointly realizable. When it fails, there is irreducible contradiction in the observation set, measured by a cohomological invariant.

38. THE OBSTRUCTION CLASS AND ITS MEANING

The failure of local sections to glue into a global section is measured by the first Čech cohomology group $H^1(\mathcal{F})$ of the sheaf \mathcal{F} . A non-vanishing element of $H^1(\mathcal{F})$ is an obstruction class that encodes the specific pattern of inconsistency preventing global coherence.

Definition 38.1 (Obstruction Energy). The obstruction energy of a candidate model \mathcal{M} with respect to a cover $\{U_i\}$ and local sections $\{s_i\}$ is

$$\mathcal{E}_{\text{obs}}(\mathcal{M}) = \|\text{Obs}(\mathcal{M})\|^2 + \sum_{i,j} d\left(s_i|_{U_i \cap U_j}, s_j|_{U_i \cap U_j}\right)^2, \quad (26)$$

where the first term measures the norm of the obstruction class and the second measures explicit pairwise disagreement on overlapping regions. The obstruction energy vanishes if and only if a global section exists.

The CLIO repair flow is then the constrained gradient descent on \mathcal{E}_{obs} :

$$\frac{d\mathcal{M}}{dt} = -\nabla_{\mathcal{M}} \mathcal{E}_{\text{obs}}(\mathcal{M}), \quad \mathcal{M}(t) \in \mathcal{A}(\mathcal{C}). \quad (27)$$

This flow attempts to move the model toward configurations in which local sections agree on overlaps. Where the flow converges to zero obstruction,

a global model is achieved. Where it stalls at a nonzero minimum, the observation set contains irreducible contradiction relative to the current model class, and no refinement within that class can resolve it.

39. SELECTION AS VANISHING OBSTRUCTION

The sheaf-theoretic language makes possible a precise unification of the selection principle with the coherence requirement.

Theorem 39.1 (Selection as Vanishing Obstruction). *A model \mathcal{M} is persistent under repeated observation, write-back, and constraint repair if and only if it approaches a region of field space in which obstruction energy is minimized. In the ideal case, selection converges to configurations satisfying $H^1(\mathcal{F}_{\mathcal{M}}) = 0$. Equivalently, what persists is what admits global gluing under the constraints imposed by the full observation history.*

Proof. Each observation induces a local section of the sheaf of admissible field configurations. A model persists only if successive observations can be incorporated into the field without producing contradictions that violate the constraint manifold $\mathcal{A}(\mathcal{C})$. If the obstruction class vanishes, local sections glue into a global section, so the field can absorb the observations while remaining admissible. If the obstruction class does not vanish, no globally coherent section exists within the current model space, and CLIO repair either stalls at a nonzero obstruction minimum or forces expansion of the field. Selection therefore favors configurations that minimize, and in the limiting case eliminate, obstruction. \square

The biological reading of this theorem is direct. A lineage corresponds to a sequence of local sections, one per generation, indexed by the environmental conditions encountered. Persistence requires that these local sections can be glued into a trajectory through the admissible region. Lineages whose local adaptations are mutually inconsistent—where what works in one environment produces configurations incompatible with survival in another—exhibit positive obstruction and are eventually eliminated. Lineages whose adaptations cohere globally persist. Selection is therefore the filtration of lineages by the vanishing of their obstruction class.

The epistemic reading is equally direct. A body of knowledge corresponds to a set of local models, one per domain of inquiry. Coherent knowledge requires that these models can be glued into a global picture. Where they

cannot, the body of knowledge exhibits obstruction, and the theory is under pressure to expand its model class to one in which gluing becomes possible.

40. THE SELECTION–PROJECTION THEOREM

The results of the preceding parts can now be assembled into a single formal statement that unifies the main claims of this essay.

Theorem 40.1 (Selection–Projection Theorem). *Let $(\mathcal{X}, T, \mathcal{C})$ be a constrained dynamical system. The following descriptions of the system’s dynamics are equivalent, in the sense that each is derivable from the others.*

The selection formulation states that persistent structures are those subsets $S \subseteq \mathcal{A}(\mathcal{C})$ satisfying $T(S) \subseteq S$. The obstruction formulation states that persistent configurations correspond to field states in which $\mathcal{E}_{\text{obs}} \rightarrow \min$, with ideal persistence achieved when $H^1(\mathcal{F}_{\mathcal{M}}) = 0$. The projection formulation states that any scalar objective function $f : \mathcal{X} \rightarrow \mathbb{R}$ that appears to be optimized is a projection of the constraint structure satisfying $x \in S \Rightarrow f(x) \geq \alpha$ for some threshold α , with f being non-unique and observer-relative. The optimization formulation states that when there exists a projection f such that ordering is approximately preserved along trajectories, the dynamics admit an approximate description as ascent on f .

Selection, obstruction minimization, and optimization are not distinct mechanisms but different representations of the same underlying constraint–closure structure.

Proof. The implication from the selection formulation to the obstruction formulation: persistence requires that local constraints remain mutually compatible under iteration, which is equivalent to the existence of a global section of \mathcal{F} . Failure of compatibility produces non-vanishing obstruction, so persistence corresponds to minimizing \mathcal{E}_{obs} . The implication from the obstruction formulation to the projection formulation: given a persistent configuration, define f as any function that separates admissible from non-admissible regions, such as the signed distance to $\partial\mathcal{A}(\mathcal{C})$. This function is maximized on persistent states but is non-unique and determined by the observer’s choice of metric, not by the dynamics. The implication from the projection formulation to the optimization formulation: if the constraint geometry is sufficiently regular, the projected function f approximately preserves ordering along trajectories, yielding the approximate ascent dynamics established in Proposition 37.1. The non-implication from optimization to selection: optimization

over f does not imply persistence because f may fail to capture all active constraints, as established by the independence of persistence and optimality in Lemmas 2.5 and 2.6. \square

Two corollaries follow immediately.

Corollary 40.2 (Non-Fundamentality of Fitness). *There exists no canonical fitness function intrinsic to the dynamics of natural selection. Any such function is a projection of the constraint structure and is therefore observer-relative and non-unique.*

Corollary 40.3 (Selection Without Objective). *A system may exhibit persistent adaptive structure without admitting any global scalar objective. In such systems, the optimization description is unavailable as a matter of principle, not merely of computational convenience.*

41. A VARIATIONAL FORMULATION OF SELECTION

The Selection–Projection Theorem admits a variational reformulation that connects it to classical field theory. Define the action functional over a trajectory $\mathcal{M}(t)$ in field space:

$$\mathcal{S}[\mathcal{M}] = \int_0^T \left(\mathcal{E}_{\text{obs}}(\mathcal{M}(t)) + \lambda \mathcal{E}_{\text{cons}}(\mathcal{M}(t)) \right) dt, \quad (28)$$

where \mathcal{E}_{obs} is the obstruction energy measuring failure of global coherence, $\mathcal{E}_{\text{cons}}$ enforces admissibility under \mathcal{C} , and λ is a Lagrange multiplier coupling the two terms. This action measures cumulative contradiction and constraint violation along the trajectory.

Theorem 41.1 (Variational Selection Principle). *Persistent configurations arise as stationary points of \mathcal{S} subject to admissibility constraints: $\delta\mathcal{S} = 0$ with $\mathcal{M}(t) \in \mathcal{A}(\mathcal{C})$.*

Proof. In the absence of inertial terms, the Euler–Lagrange equations derived from \mathcal{S} reduce to the gradient condition $\dot{\mathcal{M}} = -\nabla \mathcal{E}_{\text{obs}}$ projected onto $\mathcal{A}(\mathcal{C})$, which is precisely the CLIO repair flow under constraint projection. Stationary points are therefore exactly the configurations where the repair flow vanishes within the admissible region, which are the configurations minimizing obstruction energy subject to the constraint manifold. These are the persistent configurations. \square

When \mathcal{E}_{obs} admits a scalar reduction of the form $\mathcal{E}_{\text{obs}}(\mathcal{M}) \approx -f(x)$ for some function f , the action reduces to $\mathcal{S} \approx \int -f(x(t)) dt$, and the stationary condition becomes the condition for maximum expected f along the trajectory. Optimization emerges in this limit as a degenerate case of the variational principle, confirming that it is a special case rather than the fundamental description.

42. NOETHER-TYPE INVARIANTS AND THE STRUCTURE OF PERSISTENCE

The variational formulation makes it possible to identify conserved quantities via an analogue of Noether’s theorem. A transformation $\mathcal{M} \mapsto \mathcal{M}^\epsilon$ is a symmetry of the action if $\mathcal{S}[\mathcal{M}^\epsilon] = \mathcal{S}[\mathcal{M}]$ for all ϵ . Such transformations correspond to changes in the field configuration that preserve both obstruction energy and admissibility.

Theorem 42.1 (Constraint Noether Principle). *For every continuous symmetry of the action \mathcal{S} , there exists a quantity conserved along admissible trajectories of the repair flow within $\mathcal{A}(\mathcal{C})$.*

Proof. Standard variational arguments show that if \mathcal{S} is invariant under a one-parameter family of transformations \mathcal{M}^ϵ , then the associated current $J = (\partial\mathcal{L}/\partial\dot{\mathcal{M}}) \cdot (d\mathcal{M}^\epsilon/d\epsilon)|_{\epsilon=0}$ is conserved along solutions of the Euler–Lagrange equations. Under the projected gradient flow, conservation holds within the admissible manifold because the projection $\Pi_{\mathcal{C}}$ preserves the constraint structure that defines the symmetry. \square

Several classes of symmetry arise naturally in the constraint-field setting. Gauge symmetry, corresponding to the freedom to reparameterize the representation of a field configuration without changing its obstruction energy, yields conserved quantities that reflect the redundancy in functional description—this is the field-theoretic version of the coordinate freedom identified in Section 8. Translational symmetry in regions where the constraint structure is homogeneous yields conserved structural quantities analogous to momentum. Scale symmetry, arising when tiling or coarse-graining operations preserve constraint structure across resolutions, yields conserved quantities associated with renormalization invariance in the RSVP field.

Symmetry breaking occurs when the constraint family \mathcal{C} reduces the symmetry group G to a proper subgroup $H \subset G$. The broken symmetry

generators correspond to new degrees of freedom that manifest as persistent structures encoding the history of symmetry breaking. This is the mechanism by which the constraint field acquires differentiated structure: not through optimization toward a predetermined form, but through the accumulation of symmetry-breaking events whose consequences are preserved by the Noether invariants of the remaining symmetry.

PART IX PHOTOGRAPHIC CONTRADICTION AND MULTIMODAL CONSTRAINT

44. PHOTOGRAPHS AS PARTIAL PROJECTIONS OF THE CONSTRAINT FIELD

The abstract framework of constraint and projection can be given a concrete instantiation through the analysis of photographic observation. A photograph is not a neutral record of a state but a projection of the underlying field $\mathcal{F} = (\Phi, \mathbf{v}, S)$ into a measurement space through a highly constrained observation map. Formally, a photograph p is given by $p = \pi(x)$, where $\pi : \mathcal{X} \rightarrow \mathcal{M}$ is a non-injective map that discards the majority of the underlying state variables: the full configuration of the constraint field is collapsed into a two-dimensional array of intensity values indexed by spatial coordinates, with all information about temporal evolution, depth structure, and field components \mathbf{v} and S compressed into the fiber over p .

Two photographs $p_1 = \pi(x_1)$ and $p_2 = \pi(x_2)$ of the same object, taken from distinct viewpoints or at different times, correspond to distinct points $x_1, x_2 \in \mathcal{X}$ lying in different regions of the field. The apparent identity of the object across the two photographs is itself a projection: it is not given by the field directly but by the choice of equivalence relation under which both x_1 and x_2 are identified as representations of the same underlying structure. Different such choices yield different notions of identity. There is no observation-independent fact about which projection correctly captures the object, for the same reasons established in the discussion of functional description.

45. UNIVERSAL PHOTOGRAPHIC INCONSISTENCY

The key structural fact is that any two photographs are generically inconsistent in observation space. There is no canonical embedding of both into the same state that respects the non-injectivity of the observation map.

Proposition 45.1 (Universal Photographic Inconsistency). *Let $\pi_1, \pi_2 : \mathcal{X} \rightarrow \mathcal{M}$ be two non-injective observation maps, and let $p_1 = \pi_1(x)$ and $p_2 = \pi_2(x)$ be observations of a state $x \in \mathcal{X}$. For generic observation maps and generic states, the preimage sets $\pi_1^{-1}(p_1)$ and $\pi_2^{-1}(p_2)$ either have empty intersection within the reachable set \mathcal{R} , or their intersection is a non-singleton set whose elements correspond to mutually distinct physical configurations. In neither case does the pair (p_1, p_2) uniquely determine an underlying state or a unique model of the state.*

Proof. By Proposition 20.1, the fiber $\pi_i^{-1}(p_i)$ has positive dimension equal to $\dim(\mathcal{X}) - \dim(\mathcal{M})$. Two such fibers in general position within \mathcal{X} have intersection of dimension $\max(0, 2 \dim(\mathcal{X}) - 2 \dim(\mathcal{M}) - \dim(\mathcal{X})) = \max(0, \dim(\mathcal{X}) - 2 \dim(\mathcal{M}))$, which is positive whenever $\dim(\mathcal{X}) > 2 \dim(\mathcal{M})$. When $\dim(\mathcal{X}) \leq 2 \dim(\mathcal{M})$, the intersection is generically empty by the transversality theorem. In either case, the intersection is not a singleton. \square

The universal inconsistency of photographs is therefore not a practical limitation to be overcome by better cameras or more sophisticated matching algorithms. It is a structural consequence of the dimensional gap between the state space and the measurement space. Any two photographs of the same object define different constraints on the underlying state, and these constraints are generically incompatible within the measurement space. Their reconciliation requires a model that lives in a higher-dimensional space than either photograph inhabits.

46. CONTRADICTION AS ENTROPIC OBSTRUCTION

Given photographs p_1 and p_2 , a model M is evaluated by its ability to jointly explain both. Define the photographic contradiction functional as the minimal joint mismatch achievable within a model class:

$$\mathcal{S}_M(p_1, p_2) = \inf_{x \in M} \left(d(\pi_1(x), p_1)^2 + d(\pi_2(x), p_2)^2 \right). \quad (29)$$

When $\mathcal{S}_M > 0$ for all M in the current model class, the pair (p_1, p_2) exhibits photographic obstruction: no admissible state within any available model can

simultaneously account for both observations. This obstruction is not noise or measurement error but a signal that the model class is insufficient to unify the two projections. In RSVP terms, the obstruction corresponds to an increase in the entropy field S under simultaneous constraint enforcement: the two photographs impose mutually incompatible constraints on (Φ, \mathbf{v}, S) , and no field configuration within the current model class can satisfy both without increasing entropy beyond the admissible bound.

The continuous extension to multiple photographs defines the cumulative contradiction across an observation set $\{p_i\}$:

$$\mathcal{E}_M = \sum_{i,j} \mathcal{S}_M(p_i, p_j). \quad (30)$$

Model selection is the search for a field configuration minimizing \mathcal{E}_M , subject to the constraint manifold $\mathcal{A}(\mathcal{C})$. This is exactly the CLIO repair flow applied to the photographic case, and the Selection–Projection Theorem guarantees that what persists is what minimizes obstruction across the full observation set.

47. HINGE OBSERVATIONS AND THE COLLAPSE OF APPARENT COHERENCE

A particularly instructive phenomenon arises when a third observation p_3 is introduced to a system in which two models M_1 and M_2 appear approximately equally consistent with (p_1, p_2) . The addition of p_3 can produce a discontinuous reordering:

$$\mathcal{S}_{M_1}(p_1, p_2) \approx \mathcal{S}_{M_2}(p_1, p_2) \quad \text{but} \quad \mathcal{S}_{M_1}(p_1, p_2, p_3) \gg \mathcal{S}_{M_2}(p_1, p_2, p_3). \quad (31)$$

Definition 47.1 (Hinge Observation). An observation p_3 is a *hinge* relative to a model pair (M_1, M_2) if it produces a large reordering in the contradiction functional: models that previously appeared approximately coherent become obstructed, and the apparent plausibility landscape is restructured.

A hinge observation does not merely refine a model by adding incremental evidence. It restructures the entire admissibility landscape by activating a constraint that was previously latent. Models that survived the earlier constraint set may be incompatible with the extended set, and the reordering

can be arbitrarily large: a model that was the best available under (p_1, p_2) may have infinite obstruction under (p_1, p_2, p_3) .

The mechanism is sheaf-theoretic. The pair (p_1, p_2) admits a local gluing over the domain covered by those two observations; the triple (p_1, p_2, p_3) requires a global section over the larger domain. A model that provides local coherence over two patches may fail to extend to global coherence over three, exhibiting a non-vanishing obstruction class in $H^1(\mathcal{F})$. The hinge observation is the datum that makes the obstruction class visible.

48. DEFERRED CONSTRAINT CLOSURE AND THE TOPOLOGY OF LEARNING

The photographic framework reveals a structural feature of learning that the standard optimization picture obscures. If all available observations are enforced simultaneously, the system collapses immediately to a minimal-obstruction configuration within the current model class. Premature collapse eliminates the possibility of expanding the model class in response to genuine obstruction, because the collapse produces a locally coherent state that appears stable under further perturbation within the current class.

Deferring enforcement of some observations corresponds to maintaining a subset of constraints as latent: $\mathcal{C} = \mathcal{C}_{\text{active}} \cup \mathcal{C}_{\text{withheld}}$. The entropy associated with the withheld constraints is not resolved but preserved as latent tension in the field. This latent entropy encodes the space of possible reconciliations that would be foreclosed by premature resolution. When the withheld constraints are eventually enforced, the system faces a larger obstruction, which may force expansion of the model class.

Proposition 48.1 (Topology of Learning Under Constraint Revelation). *Let \mathcal{M}_k denote the set of models consistent with k observations. Then $\mathcal{M}_{k+1} \subseteq \mathcal{M}_k$, with strict inclusion whenever the $(k+1)$ -th observation introduces a constraint independent of the first k . The limiting model set $\mathcal{M}_\infty = \bigcap_k \mathcal{M}_k$ consists of those models admissible under all constraints simultaneously.*

Proof. Consistency with $k+1$ observations requires consistency with the first k and additionally with the $(k+1)$ -th, so $\mathcal{M}_{k+1} \subseteq \mathcal{M}_k$. Independence of the new constraint means it eliminates at least one model in \mathcal{M}_k , producing strict inclusion. The intersection is the set of models surviving all constraints by a standard intersection argument. \square

The order in which constraints are revealed therefore shapes the topology of the learned model. Two observers who encounter the same observations in different orders will traverse different sequences of model sets, and may reach different locally stable configurations that are both consistent with the full observation set but represent different global sections of the sheaf. Path dependence in learning is not an artifact of finite sample sizes or suboptimal algorithms; it is a structural consequence of the sheaf geometry and the order of constraint enforcement.

49. BEFORE AND AFTER AS BOUNDARY CONDITIONS ON TRAJECTORIES

A particularly revealing instance of photographic constraint arises in the analysis of transformation processes. Two photographs p_0 and p_1 of the same location at different times—a site before and after some process of change—function not as static observations of states but as boundary conditions on trajectories. A model M must simultaneously account for $\pi(\gamma_M(t_0)) \approx p_0$, $\pi(\gamma_M(t_1)) \approx p_1$, and the requirement that γ_M be a trajectory of the model dynamics over $[t_0, t_1]$.

This transforms model evaluation from a static fitness question into a dynamical consistency question. The model must not only explain the initial and final states but also connect them through a physically admissible trajectory. Many models can match isolated states but fail to connect them through any plausible intermediate evolution. The photographic pair therefore acts as a hinge constraint on the model class, eliminating all models whose dynamics cannot generate a trajectory between the two observations.

The introduction of a third photograph $p_{1/2}$, taken at an intermediate time, extends the constraint to require that the trajectory pass through a third boundary condition. This rapidly eliminates models that interpolated freely through the unobserved interval: a model may generate arbitrarily many trajectories connecting p_0 to p_1 , but only a restricted subset of those trajectories also pass through $p_{1/2}$. As the density of intermediate photographs increases, the admissible trajectory set collapses toward those that accurately encode the underlying physical dynamics, and most models—those that relied on the freedom of the unobserved intervals to satisfy the endpoint constraints—are obstructed.

The epistemological consequence is precise: the unseen intervals between

sparse observations are exactly where incorrect models survive. Dense observation eliminates this refuge by transforming interpolation freedom into constraint. The correct model is not the one that best fits individual observations but the one that admits a globally consistent trajectory under all observations simultaneously—that is, the one whose obstruction energy vanishes.

PART X MULTIMODAL CONSTRAINTS AND THE AMPLIFICATION OF OBSTRUCTION

50. MULTIPLE SENSOR MODALITIES AS INDEPENDENT PROJECTIONS

The photographic analysis extends naturally to the case of multiple sensor modalities. A camera, an accelerometer, and a microphone are all, from the perspective of the constraint framework, observation maps from the same underlying state space \mathcal{X} into different measurement spaces. Each is non-injective; each retains different invariants of the underlying state; and each imposes independent constraints on the set of admissible trajectories.

Let $\pi_v : \mathcal{X} \rightarrow \mathcal{M}_v$, $\pi_a : \mathcal{X} \rightarrow \mathcal{M}_a$, and $\pi_s : \mathcal{X} \rightarrow \mathcal{M}_s$ denote the visual, accelerometric, and acoustic observation maps respectively. A simultaneous multimodal observation at time t is the tuple $o_t = (\pi_v(x_t), \pi_a(x_t), \pi_s(x_t))$. A model M must generate trajectories $\gamma_M(t)$ such that all three modality projections match the observed data across the entire observation window. The joint admissibility condition is the intersection of three independent constraint sets:

$$\mathcal{C}_{\text{multi}} = \mathcal{C}_v \cap \mathcal{C}_a \cap \mathcal{C}_s, \quad (32)$$

and the admissible model set under multimodal constraints is generically much smaller than under any single modality.

Theorem 50.1 (Constraint Amplification by Modal Intersection). *Let $\mathcal{C}_v, \mathcal{C}_a, \mathcal{C}_s$ be the constraint sets induced by visual, accelerometric, and acoustic observations respectively. If any two of these constraint sets impose independent restrictions—meaning neither is contained in the other—then $|\mathcal{C}_v \cap \mathcal{C}_a \cap \mathcal{C}_s| < \min\{|\mathcal{C}_v|, |\mathcal{C}_a|, |\mathcal{C}_s|\}$, and the obstruction energy of an incorrect model under the joint constraint exceeds the sum of its obstruction energies under any two individual constraints.*

Proof. Independence of constraint sets means each modality eliminates models not eliminated by the others. The intersection therefore strictly reduces the admissible set below any individual constraint set. The obstruction energy is additive across independent constraints because the photographic contradiction functionals for independent modalities contribute independently to the total mismatch $\mathcal{E}_M = \mathcal{S}_M^v + \mathcal{S}_M^a + \mathcal{S}_M^s$, and independence means each term can be non-zero even when the others vanish. \square

51. ACCELEROMETRIC AND ACOUSTIC OBSTRUCTION

Accelerometric data imposes constraints of a qualitatively different character from photographic data. Where photographic observation constrains the endpoints or waypoints of a trajectory—the values of $\gamma_M(t_i)$ at discrete times—accelerometric observation constrains the trajectory’s curvature: it requires that $\ddot{\gamma}_M(t) \approx a(t)$ throughout the observation window, where $a(t)$ is the measured acceleration sequence. This is a differential constraint on the dynamics of M rather than a pointwise constraint on its states.

Definition 51.1 (Kinematic Obstruction). A model M is kinematically obstructed by an accelerometer sequence $\{a_t\}$ if no trajectory generated by T_M has second derivative matching a_t at all times in the observation window. A model may satisfy photographic endpoint constraints while being kinematically obstructed if its dynamics cannot generate the required acceleration profile between those endpoints.

Kinematic obstruction is strictly stronger than photographic obstruction in the following sense: a model can satisfy all photographic constraints on a trajectory and yet fail to produce any trajectory between those endpoints whose acceleration matches the observed data. The accelerometer thus eliminates models that are photographically consistent but dynamically implausible.

Acoustic data introduces a third constraint domain: the frequency and temporal structure of the sound field generated by the process under observation. A model must produce trajectories whose induced acoustic emissions match the spectral decomposition of the observed signal. Two trajectories that are visually indistinguishable and kinematically consistent may produce radically different sound profiles, so acoustic constraints can separate models that are invisible to the other two modalities.

Definition 51.2 (Spectral Obstruction). A model M is spectrally obstructed by an acoustic signal $s(t)$ if no trajectory under T_M produces a signal whose

Fourier transform $\hat{s}_M(\omega)$ matches the observed $\hat{s}(\omega)$ up to the admissible tolerance.

The combination of kinematic and spectral obstruction with photographic obstruction produces a constraint network that is far more restrictive than any single modality, and the obstruction class of an incorrect model under the joint constraint is generically much larger than under any individual constraint.

52. TEMPORAL ALIGNMENT AS AN ADDITIONAL CONSTRAINT

A critical and often overlooked feature of multimodal data is that temporal alignment between modalities is itself a constraint rather than a given. Different models may align the same data streams with different temporal parameterizations, effectively reconstructing different internal clocks. Given simultaneous observations $(p(t), a(t), s(t))$, a model M must find a time parameterization $\tau(t)$ such that the visual, kinematic, and acoustic projections of $\gamma_M(\tau(t))$ match the observed data. Different models induce different $\tau(t)$.

Definition 52.1 (Temporal Obstruction). A model M is temporally obstructed if no reparameterization $\tau(t)$ exists that simultaneously aligns its trajectories with all observed modalities. Temporal obstruction is distinct from kinematic obstruction: a model may produce the correct acceleration magnitudes but in the wrong temporal order, or may require an internal clock that runs at a different rate than any consistent assignment of $\tau(t)$.

Temporal obstruction connects multimodal constraint analysis to the asymmetry of epistemic access established in Proposition 21.2. The observer's forward reachable set $\mathcal{R}^+(x_0)$ is constructed by a specific temporal parameterization, and models that require incompatible parameterizations are excluded not by direct measurement of a state but by the observer's inability to align its trajectory with the model's implied time. The arrow of inference is therefore constrained not only by the content of observations but by the temporal structure of the observation process.

53. DENSE VERSUS SPARSE OBSERVATION AND THE REFUGE OF UNOBSERVED INTERVALS

The distinction between sparse and dense observation generalizes the photographic analysis. Let $\{p_i\}_{i=0}^n$ be a finite sequence of observations at known

times, and let $p(t)$ be a continuous observation over the same interval. The dense observation imposes a path constraint—requiring that $\pi(\gamma_M(t)) \approx p(t)$ for all t —while the sparse observations impose only endpoint constraints at the discrete times t_i . The set of admissible trajectories under dense observation is always a subset of the set admissible under sparse observation:

$$\mathcal{C}_{\text{dense}} \subseteq \mathcal{C}_{\text{sparse}}, \quad (33)$$

with strict inclusion in all non-degenerate cases.

The intervals between sparse observations constitute a refuge for incorrect models. A model may satisfy endpoint constraints by interpolating through physically implausible intermediate states, visible only when intermediate observations are made. Refining the observation from sparse to dense eliminates this refuge by transforming the free interpolation problem—which admits many solutions—into a constrained path problem—which admits far fewer. Models that previously appeared consistent with the data because their violations were hidden in unobserved intervals are exposed and obstructed as observation density increases.

Proposition 53.1 (Collapse of Model Space Under Dense Observation). *For a physical process governed by dynamics T and a continuous observation $p(t)$ generated by T , the set of models consistent with $p(t)$ collapses, as observation resolution increases, to exactly those models whose dynamics generate a trajectory consistent with $p(t)$ at every instant. Models whose inconsistencies are localized to unobserved intervals are obstructed as those intervals become observed.*

Proof. A model M is consistent with a discrete set of observations $\{p(t_i)\}$ if and only if there exists a trajectory of T_M passing through $(\gamma_M(t_i) : \pi(\gamma_M(t_i)) \approx p(t_i))$ for all i . As the set $\{t_i\}$ becomes dense in the observation interval, the consistency condition approaches the requirement that $\pi(\gamma_M(t)) \approx p(t)$ for all t . Models that satisfied the endpoint constraints by violating physical plausibility in unobserved intervals fail the path constraint as soon as those intervals are sampled. \square

The epistemological consequence stated earlier in the photographic context now generalizes: the correct model is the one whose obstruction energy vanishes under the full path constraint. Any model that relies on unobserved intervals for its apparent coherence is a model of the observation gaps rather than of the underlying process. Dense observation eliminates the distinction

between observing a process and modeling it, because the path constraint leaves no room for the interpolation freedom that incorrect models exploit.

54. THE STATISTICAL FIELD AND ENTROPY AS OBSTRUCTION DENSITY

The constraint framework admits a statistical extension in which the system is described not by a single field configuration but by a distribution over configurations. Let $\mathbb{P}_t[\mathcal{M}]$ be a probability distribution over admissible field configurations at time t . The partition function

$$Z = \int \exp(-\beta \mathcal{E}_{\text{obs}}(\mathcal{M})) \mathbf{1}_{\mathcal{A}(c)}(\mathcal{M}) \mathcal{D}\mathcal{M} \quad (34)$$

defines a Boltzmann-type measure on the space of admissible configurations, with the inverse temperature β controlling the sharpness of concentration. The induced distribution is

$$\mathbb{P}[\mathcal{M}] = \frac{1}{Z} \exp(-\beta \mathcal{E}_{\text{obs}}(\mathcal{M})) \mathbf{1}_{\mathcal{A}(c)}(\mathcal{M}). \quad (35)$$

Selection corresponds to concentration of this distribution in the limit $\beta \rightarrow \infty$, where the measure collapses onto configurations minimizing obstruction energy. At finite β , multiple near-coherent configurations coexist, capturing the variation that makes evolutionary exploration possible and the stochasticity that makes learning non-deterministic.

The free energy $\mathcal{F} = \mathbb{E}[\mathcal{E}_{\text{obs}}] - \beta^{-1}H(\mathbb{P})$ balances two competing pressures: the reduction of contradiction and the maintenance of distributional breadth. A system with very low β distributes mass broadly over the admissible region, exploring many configurations but converging slowly to coherent ones. A system with very high β concentrates mass on low-obstruction configurations, converging rapidly but vulnerable to local obstruction minima.

Within this statistical picture, the entropy field $S(x)$ introduced in the RSVP formulation acquires the interpretation of local obstruction density: the expected obstruction energy in a neighborhood of x , averaged over the local distribution. Regions of high S correspond to locations where multiple partially coherent but mutually inconsistent field configurations coexist, none able to resolve the local contradiction within the current model class. Regions of low S correspond to coherent configurations where local constraints are mutually compatible. The dynamics of S under the repair flow,

$$\partial_t S = -\|\nabla_{\mathcal{M}} \mathcal{E}_{\text{obs}}\|^2 + \sigma, \quad (36)$$

express the competition between contradiction resolution (the first term, which decreases S as the repair flow reduces obstruction) and contradiction injection (the second term σ , which increases S as new observations impose new constraints). A system persists when it can resolve contradiction faster than it accumulates, maintaining S within the admissible bound. Selection eliminates configurations that cannot maintain this balance.

55. UNCERTAINTY, INCOMPATIBILITY, AND THE INFORMATION-THEORETIC LIMIT

The framework admits a final generalization to the information-theoretic setting, where the obstruction concept connects to fundamental limits on simultaneous constraint satisfaction. Two observables A and B correspond to incompatible constraint sets \mathcal{C}_A and \mathcal{C}_B when the intersection $\mathcal{C}_A \cap \mathcal{C}_B$ is empty or has measure zero. In this case, no configuration can simultaneously minimize obstruction with respect to both observables, and the irreducible obstruction under joint constraint is positive.

This provides a constraint-geometric interpretation of the uncertainty principle: uncertainty does not arise from measurement disturbance but from the structural incompatibility of certain constraint sets. A system constrained to be in a low-obstruction configuration with respect to A is necessarily in a high-obstruction configuration with respect to B when $\mathcal{C}_A \cap \mathcal{C}_B = \emptyset$. The bound on simultaneous resolution is set by the geometry of the constraint manifold, not by the precision of the measuring apparatus.

The rate at which obstruction can be resolved is bounded by the rate of information integration: if new constraints are imposed faster than the repair flow can absorb them, entropy accumulates without bound and the configuration exits the admissible region. This provides a fundamental limit on learning, adaptation, and coherence. A system that receives more observations than it can integrate—more photographs, more sensor data, more environmental signals—will exhibit increasing entropy until it either expands its model class or is eliminated by the constraint.

Proposition 55.1 (Information-Rate Bound on Persistence). *Let \dot{S}_{max} be the maximal rate at which the repair flow can reduce entropy within $\mathcal{A}(\mathcal{C})$. Persistence*

requires that the rate of contradiction injection σ satisfy $\sigma \leq \dot{S}_{\max}$. When $\sigma > \dot{S}_{\max}$, entropy diverges and the configuration exits the admissible region in finite time.

Proposition 55.1 connects the abstract persistence condition to the concrete dynamics of information processing. A biological organism persists only if its metabolic and repair processes can absorb environmental perturbations at least as fast as they arrive. A learning system persists only if its update mechanism can integrate new constraints at least as fast as they are imposed. A coherent model persists only if its reconciliation of contradictions keeps pace with the arrival of new observations that generate them. In all three cases, the condition for persistence is the same: the rate of contradiction resolution must dominate the rate of contradiction injection.

CODA: THE GEOMETRY OF WHAT CANNOT BE SEEN

57. THE COMPLETE PICTURE

The arc of argument across ten parts can now be seen as a single movement, not a series of applications of a fixed framework to different domains, but the progressive disclosure of a single structure through successive removals of obscuring assumptions. Part I removed the assumption that selection is optimization. Part II removed the assumption that function is intrinsic. Part III removed the assumption that observation is neutral. Parts IV through V unified the three removals under the constraint–projection framework. Parts VI through X gave this framework field-theoretic, sheaf-theoretic, variational, and multimodal realizations, each revealing the same structure at greater resolution.

What emerges is not a theory about any particular domain but a grammar for describing what it means to persist, to describe, and to know, in any domain. The grammar has four words: constraint, closure, projection, and reachability. Everything else—purpose, function, optimization, knowledge, observation, reality—is a sentence formed from those four words in different arrangements, under different conditions, from different positions within the field.

The sentences are not equivalent. Persistence is intrinsic; function is relational; knowledge is path-dependent. These asymmetries are genuine and important. But they are asymmetries within the same grammar, not marks of different fundamental kinds.

58. WHAT THE FRAMEWORK DOES NOT SAY

It is worth being explicit about what the framework established in this essay does not claim. It does not claim that functional descriptions are false or useless. They are coordinate systems over persistent structures, and coordinate systems are indispensable for navigation, prediction, and communication. The critique of functionalism in Part II was not a claim that the word “function” should be abandoned but a claim that it should not be read as designating an intrinsic, observer-independent property.

The framework does not claim that science cannot converge on stable descriptions of the world. The repair flow established in Part VIII can converge to a globally coherent model when the constraint structure admits one. The progressive tightening of the model set under increasing observations, established in Proposition 48.1, is the formal correlate of the convergence of scientific inquiry. What the framework adds is precision about the conditions under which convergence is possible, the conditions under which it is not, and the sense in which the limit of convergence is always a section of the sheaf of constraints rather than a transparent copy of the world.

The framework does not claim that unobserved structure is in any sense more real than observed structure. Persistence is the same condition for observed and unobserved regions of the state space. What differs is accessibility, not ontological status. The inaccessible persistent regions of Part III exist in exactly the same sense as the accessible ones; they simply do not appear in any observer’s reachable set. The philosophical consequence is not that the unobserved is more fundamental but that existence and accessibility are independent.

59. THE FINAL INVERSION

The opening of this essay described three intuitions: that life improves, that organs are for things, that knowledge accumulates toward truth. Each of these intuitions feels like a description of the world. The foregoing has argued that all three are descriptions of the observer’s position within the constraint field, mistaken for descriptions of the field itself.

The inversion is now complete and can be stated in its sharpest form. Direction in biological evolution is not a property of the dynamics but the shape of the conditioning event: the constraint that excludes, producing the

appearance of selection toward when observed from within the surviving set. Purpose in biological structures is not a property of the structures but the choice of coordinate: the observer's projection, which reads function into systems indifferent to any particular invariant. Progress in knowledge is not convergence toward truth but the progressive refinement of the epistemically accessible fiber within the reachable set, which appears to converge from inside while remaining a proper subset of the full fiber from outside.

The world does not optimize. It filters. What survives is not what is best but what is not eliminated. What is described as having function is not what intrinsically has it but what appears to have it under a chosen projection. What is known is not the world but the world's trace on a constrained trajectory through a space far larger than any observer can reach.

Constraint is prior. Persistence is intrinsic. Function and knowledge are projections.

What cannot be seen is not absent. It is present as the curvature of the boundary of what can be seen, as the shape of what can be reached, as the geometry of the possible paths through the field. The world is larger than any observation of it, and it persists—whether or not it is observed—because persistence requires only closure, not witness.

REFERENCES

- [1] G. R. Price, Selection and Covariance, *Nature* **227** (1970), 520–521.
- [2] R. A. Fisher, *The Genetical Theory of Natural Selection*, Clarendon Press, Oxford, 1930.
- [3] S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, 1993.
- [4] J. Maynard Smith and E. Szathmáry, *The Major Transitions in Evolution*, Oxford University Press, 1995.
- [5] L. Wright, Functions, *The Philosophical Review* **82** (1973), 139–168.
- [6] K. Neander, Functions as Selected Effects, *Philosophy of Science* **58** (1991), 168–184.
- [7] H. Putnam, Psychological Predicates, in *Art, Mind, and Religion*, eds. W. H. Capitan and D. D. Merrill, University of Pittsburgh Press, 1967.
- [8] J. A. Fodor, *The Language of Thought*, Harvard University Press, 1975.
- [9] D. J. Chalmers, *The Conscious Mind*, Oxford University Press, 1996.
- [10] S. Smale, Differentiable Dynamical Systems, *Bulletin of the American Mathematical Society* **73** (1967), 747–817.
- [11] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*, Cambridge University Press, 1995.
- [12] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, Springer, 1989.
- [13] M. W. Hirsch and S. Smale, *Differential Equations, Dynamical Systems, and Linear Algebra*, Academic Press, 1974.
- [14] J. Nash, The Imbedding problem for Riemannian manifolds, *Annals of Mathematics* **63** (1956), 20–63.
- [15] A. S. Eddington, *The Nature of the Physical World*, Cambridge University Press, 1928.
- [16] J. Ladyman and D. Ross, *Every Thing Must Go: Metaphysics Naturalized*, Oxford University Press, 2007.

- [17] J. A. Wheeler, *Law Without Law*, in *Quantum Theory and Measurement*, eds. J. A. Wheeler and W. H. Zurek, Princeton University Press, 1983.
- [18] P. Engel, S. Filipazzi, F. Greer, M. Mauri, and R. Svaldi, Boundedness of some fibered K-trivial varieties, *arXiv preprint* arXiv:2507.00973 (2025).