

Manufacturing Stability: Structural Foundations for Digital Infrastructure

Flyxion

February 19, 2026

Abstract

This paper develops a geometric and thermodynamic account of platform extraction in digital political economy. Platform dominance is reformulated as curvature concentration on an information manifold, producing geodesic convergence, predictive asymmetry, and entropy inflation in local semantic states. Extraction is therefore modeled not merely as a distributive imbalance but as a structural instability of the informational substrate itself.

In response, we propose an entropy-bounded semantic architecture in which structural constraints are embedded directly into the computational substrate. Knowledge is represented as typed semantic spheres equipped with intrinsic identity, measurable entropy, and explicit provenance graphs. Evolution of these objects is governed by formally declared entropy budgets, staking-backed rule validation, and conservation laws expressed through dual valuation functions measuring structural coherence and persistence. We formalize the geometry of the underlying manifold, establish an entropy soundness theorem for arbitrary execution traces, derive information-theoretic properties of the valuation system, specify a sheaf-theoretic closure condition ensuring global semantic coherence, and analyze incentive compatibility under adversarial conditions.

We further prove a thermodynamic incompatibility theorem showing that engagement-maximizing dynamics are generically unstable with respect to entropy-bounded evolution, yielding either unbounded disorder or curvature concentration. The resulting framework constitutes a substrate-level analogue of structural antitrust in which extractive behavior corresponds to violation of conserved invariants and is therefore algebraically inadmissible. Collaboration is thereby recast as constrained field evolution rather than unregulated optimization, situating institutional design within a formal theory of entropy, curvature, and semantic conservation.

Contents

1	Introduction	3
2	Literature and Intellectual Context	4
2.1	Platform Economics and Political Economy	4
2.2	Information Geometry and Field Theory	5
2.3	Type Theory, Sheaf Theory, and Formal Semantics	5
2.4	Algorithmic Accountability and Information Asymmetry	6
3	Epistemological Asymmetry and Predictive Leverage	7
3.1	The Structure of Predictive Asymmetry	7
3.2	Channel Capacity and Leverage	8
3.3	Geometric Expression of Asymmetry	9
3.4	The Right to Predictive Parity	9
4	Formal Objects and Definitions	10
5	The RSVP Geometry of Platform Dominance	12
5.1	Curvature Concentration as a Model of Extraction	12
5.2	Geodesic Convergence and Jacobi Fields	13
5.3	Manifold Collapse and Predictive Resolution	13
6	Entropy Soundness	14
6.1	Motivation and Setup	14
6.2	The Entropy Soundness Theorem	15
6.3	Geometric Interpretation of Entropy Boundedness	16
7	Engagement Optimization and Entropic Instability	16
7.1	The Engagement–Entropy Tension	16
8	Phase Structure and Dynamical Regimes	17
8.1	Order Parameters	17

8.2	Regime Classification	18
8.3	Phase Transitions	19
8.4	Phase Diagram Interpretation	19
9	Invariant Measures and Long-Run Behavior	19
9.1	Markovian Evolution of Semantic Spheres	19
9.2	Tightness under Entropy Boundedness	20
9.3	Existence of Invariant Measures	20
9.4	Non-Existence under Engagement Optimization	21
9.5	Interpretation	21
10	Structural Antitrust as a Categorical Invariant	22
10.1	The Category of Semantic Evolution	22
10.2	Metric Functor and Geometric Realization	22
10.3	Forbidden Endofunctors	23
10.4	Interpretation	23
11	Computational Complexity and Decidability	24
11.1	Computability of Semantic Entropy	24
11.2	Decidability of Entropy-Safety	24
11.3	Restricted Rule Languages	25
11.4	Complexity of Sheaf Validation	25
11.5	Complexity of Crystal Computation	25
11.6	Implications	26
12	Governance and Protocol Upgrade Morphisms	26
12.1	The Governance Problem	26
12.2	Protocol as a Structured Object	26
12.3	Upgrade Morphisms	27
12.4	Monotonicity Constraints	27
12.5	Governance as a Higher-Order Category	28

12.6 Interpretation	28
13 Sheaf-Theoretic Closure	29
13.1 The Need for Global Coherence	29
13.2 Presheaves and the Sheaf Condition	29
13.3 Semantic Spheres as Sections	30
13.4 Media–Quine Closure	30
14 Crystal Valuations and Welfare Economics	31
14.1 From Accounting to Economics	31
14.2 Texture Crystals as Allocative Efficiency	31
14.3 Time Crystals as Decay-Adjusted Provenance Value	32
14.4 Welfare Alignment and Open Problems	33
15 Mechanism Design and Manipulation Resistance	33
15.1 Adversarial Threats	33
15.2 Stake-Backed Entropy Discipline	34
15.3 Quality Metric Integrity	34
15.4 Cryptographic Identity	35
16 Comparison of Regimes	35
17 Relationship to Existing Proposals	36
18 Unified Structural Stability Theorem	37
18.1 Preliminaries	37
18.2 Structural Stability	37
18.3 Thermodynamic Antitrust	38
18.4 Interpretation	38
19 Empirical Calibration and Measurability	38
19.1 The Gap Between Theory and Observation	38

19.2	Empirical Proxies for Curvature Concentration	39
19.3	Empirical Proxies for Entropy Growth	39
19.4	Empirical Proxies for Predictive Asymmetry	40
19.5	Empirical Proxies for Crystal Valuations	40
19.6	Calibrating Structural Thresholds	41
20	Agent Welfare and Semantic Autonomy	41
20.1	Normative Grounding	41
20.2	Welfare Losses from Extractive Dynamics	42
20.3	Semantic Autonomy as a Formal Condition	42
20.4	The Right to Non-Extraction as a Substrate Guarantee	43
20.5	Plurality, Diversity, and the Geometry of the Good	44
21	Implementation Architecture	45
21.1	Overview	45
21.2	Semantic Object Layer	45
21.3	Rule Execution Layer	46
21.4	Staking and Slashing Layer	46
21.5	Sheaf Validation Layer	47
21.6	Governance Layer	47
21.7	Integration and Deployment Considerations	48
22	Limitations and Open Problems	49
23	Synthesis	50
24	Foreknowledge and the Ethics of Optimization	51
25	Temporal Distortion and the Pepsi Paradigm	52
26	Equipose and Stopping Rules	53
27	Selection Dynamics and Stratification	54

28 Objective Functions and Misalignment	54
29 Constraint as Ethical Governance	55
30 Surveillance Expansion Under the Banner of Safety	56
31 Backdoor Surveillance and Regulatory Incentives	56
32 Identity Verification and the Erosion of Anonymity	57
33 Surveillance Pricing and Differential Treatment	58
34 AI Systems and Secondary Exploitation	58
35 Moral Panic and Policy Feedback Loops	59
36 Data Privacy and Structural Incentives	59
37 Platformification as Substrate Abstraction	60
38 Chipification and the Virtualization of Production	62
39 Recapitulation and Autocatalytic Organization	65
39.1 Autocatalysis, Entropy Reduction, and Manifold Flattening	67
40 Swarm Recurrence and Planetary Modularity	69
41 Curvature, Spectral Gap, and Rapid Mixing in Coordinated Swarms	71
41.1 Graph Laplacians, Markov Kernels, and the Spectral Gap	72
41.2 Ollivier–Ricci Curvature as Local Contractivity	72
41.3 From Curvature to Spectral Gap and Mixing	73
41.4 Swarm Control as Gap Maintenance	74
41.5 Continuous Limit Interpretation	75
42 Stigmergic Scaffolds and Hierarchical Recapitulation	75

43 Agent Abstraction, Partial Simulation, and Observer-Dependent Function	78
43.1 Behaviorism as Substrate-Independent Functional Ontology	81
43.2 Ontologies as Constraint Operators in Cognition	83
43.3 Abstraction, Cognitive Compression, and the Thermodynamic Bias of Formal Systems	85
43.4 Entropy Resistance and the Engineering of Selective Worlds	88
43.5 Pathology as Ontological Preference	90
43.6 Symmetry, Invariance, and the Implicit Axiology of Science	93
43.7 From Scientific Axiology to Substrate Design	95
43.8 Invariant-Preserving Semigroups and Structural Stability	97
44 Control-Theoretic Ontology	100
44.1 State Spaces and Observational Projections	100
44.2 Preferences as Reaction Functionals	101
45 Recapitulation and Externalization of Charts	101
45.1 Projection and Reinstantiation	101
45.2 Swarm Dynamics and Stigmergic Scaffolds	102
46 Functional Ontologies and Substrate Independence	103
46.1 Behaviorism as Functional Abstraction	103
46.2 Optimization and Entropic Compression	103
47 Civilizational Control and Lyapunov Invariants	104
48 From Control to Curvature: Manufacturing Stability	105
49 Pipeline Closure and the Desire for Forward Motion	106
50 Discourse as Coupled Manifold Dynamics	108
50.1 Semantic Manifolds and Communicative Morphisms	108
50.2 Differential Conceptual Alignment	109

50.3	Dialogue and Essay as Dynamical Systems	109
50.4	Curvature, Contestation, and Stability	110
50.5	Communication as Sheaf Alignment	110
50.6	Implications for Stability Manufacturing	111
51	Cognition as Framing and Hierarchical Evaluation	111
51.1	Framing as Projection	111
51.2	Scenario Generation as Evolutionary Branching	111
51.3	Evaluation as Hierarchical Stacking	112
51.4	Frames as Evolutionary Filters	112
51.5	Geometric Interpretation	113
51.6	Implications	113
52	Cognitive Entropy and Substrate Entropy	113
52.1	Internal Entropy in Scenario Generation	113
52.2	Substrate-Induced Branch Inflation	114
52.3	Coupled Stability Principle	115
53	Reinforcement Learning and Predictive Framing	115
53.1	Evaluation as Value Function Estimation	115
53.2	Predictive Processing and Free Energy	115
53.3	Hierarchical Stacking as Multi-Level Inference	116
54	Embodiment, Environmental Modification, and the Ontology of Trace	118
54.1	Organisms as Environmental Rewriters	118
54.2	Trace as Energy Expenditure	119
54.3	The Non-Equivalence of Trace and Agency	119
54.4	Embodiment as Ontological Constraint	120
54.5	Profiles as Ghosts of Work	120
55	Thermodynamic Cost of Trace and the Landauer Constraint	121

55.1	Information as Physical Work	121
55.2	Profiles as Entropic Residues	121
55.3	Epistemic Authority and Energetic Coupling	122
56	Virtualization, Chipification, and Substrate Dependence	122
56.1	Organizational Surjection onto Silicon	122
56.2	Virtual Machines and Physical Substrate	123
56.3	Impossibility of Purely Disembodied Agency	124
57	Swarm Dynamics, Stigmergy, and Energetic Trace	125
57.1	Stigmergy as Distributed Environmental Writing	125
57.2	Human Infrastructure as Scaled Stigmergy	125
57.3	Energetic Residue and Collective Memory	126
58	Extraction as Curvature Capture in Energetic Substrates	126
58.1	From Trace to Toll Booth	126
58.2	Energetic Interpretation of Extraction	127
58.3	Entropy-Bounded Architecture as Evaporation Mechanism	127
58.4	No Escape from Substrate	128
59	A Decay Necessity Theorem for Stigmergic Stability	128
59.1	Model Setup	128
59.2	No-Decay Case	128
59.3	Decay Condition for Stability	129
59.4	Theorem	129
59.5	Proof Sketch	129
59.6	Interpretation	130
60	Cognition, Embodiment, and Manufactured Stability	130
61	Curvature Singularity and Extraction	131
61.1	Flow Concentration on an Information Manifold	131

61.2 No-Decay Instability	132
61.3 Theorem	132
61.4 Proof Sketch	132
61.5 Interpretation	132
62 Manufacturing Stability	133
63 Conclusion: Stability as an Architectural Property	135

1 Introduction

Recent political-economic analysis has identified a structural shift in the nature of digital platforms. Early platforms functioned catalytically: they lowered transaction costs, reduced friction between agents, and enabled decentralized participation without substantially appropriating the value they helped create. Contemporary dominant platforms operate differently. They do not merely facilitate exchange; they deepen curvature around themselves, siphoning surplus through positional control of infrastructure and predictive leverage over the cognitive and behavioral states of their users. The transition from catalyst to extractor is the diagnostic core of Tim Wu’s account of attention capture and Shoshana Zuboff’s characterization of surveillance capitalism.

This paper reformulates that shift in geometric and thermodynamic terms. We treat the informational substrate on which human collaboration occurs as a Riemannian manifold and model platform dominance as curvature concentration within that manifold. Extraction is not merely inequitable distribution; it is a geometric phase instability characterized by geodesic convergence and entropy misallocation. Edward Herman and Noam Chomsky famously argued that systemic bias in mass media arises not primarily from individual intent but from structural filters embedded in institutional design. The present work extends that structural intuition into a formal register, translating architectural asymmetry into curvature, entropy, and invariant constraints on an information manifold. The conventional responselitigation, behavioral remedies, or structural breakups operates as an exogenous diffusion term imposed after curvature has already concentrated. PlenumHub proposes a different approach. It embeds structural invariants directly into the compute substrate so that the geometric preconditions of extraction cannot arise without violating foundational algebraic constraints.

The unit of knowledge in this framework becomes a typed semantic sphere, which is a formal object carrying an immutable identity, a required set of modalities, measurable entropy, and an explicit derivational provenance graph. Evolution of semantic spheres is bounded by entropy budgets attached to every rule application. Value is governed by crystal conservation laws that reward coherence and penalize redundancy. Global validity is enforced via sheaf-theoretic closure, ensuring that locally coherent semantic assignments extend uniquely to global ones.

The core claim is that structural antitrust can be implemented algebraically. If extraction corresponds to negative diffusion and curvature concentration, then a properly constrained substrate can preclude those regimes by design rather than by litigation. The remainder of this paper develops that claim formally. [Section 2](#) situates the work within the relevant political-economic and mathematical literature.

Section 3 formalizes the epistemological asymmetry between platform and agent using information-theoretic tools. Section 4 collects all formal objects and their definitions. Section 5 develops the RSVP geometry. Section 6 proves the entropy soundness theorem. Section 7 formalizes the instability of engagement-maximizing dynamics. Section 8 analyzes the phase structure of dynamical regimes. Section 9 establishes existence and non-existence of invariant measures. Section 10 formalizes structural antitrust as a categorical invariant. Section 11 addresses computational complexity and decidability. Section 12 treats governance and protocol upgrade morphisms. Section 13 presents the sheaf-theoretic closure condition. Section 14 analyzes the crystal economy and its welfare-theoretic properties. Section 15 treats manipulation resistance and incentive compatibility. Section 16 compares the two regimes structurally. Section 17 situates the proposal with respect to existing technological approaches. Section 18 presents the unified structural stability theorem. Section 19 addresses empirical calibration and measurability. Section 20 develops the normative theory of agent welfare and semantic autonomy. Section 21 describes the implementation architecture. Section 22 states limitations and open problems. Section 23 provides a concluding synthesis. Appendices contain a worked numerical example and supporting derivations.

2 Literature and Intellectual Context

2.1 Platform Economics and Political Economy

Wu’s account of the attention economy identifies a cycle in which communications and media platforms begin as open and generative, then progressively close as dominant incumbents discover that restricting access and concentrating control is more profitable than facilitating open exchange. His earlier work on network neutrality addresses the infrastructure dimension of this dynamic: when the conduit for information becomes the instrument of extraction, the distributional consequences reach every layer of the communicative ecosystem.

Zuboff’s framework of surveillance capitalism extends this diagnosis to the behavioral dimension. Platforms do not merely extract attention; they construct and sell predictions of behavior, transforming the surplus of human experience into a raw material for futures markets in behavioral modification. The epistemological asymmetry involved is structural: the platform holds predictive models of the agent that the agent cannot access, correct, or contest.

In formal economics, Rochet and Tirole’s analysis of two-sided markets provides the canonical treatment of platform pricing. Platforms subsidize one side of a market

to increase participation, then extract surplus from the other. Cross-side network externalities create barriers to entry that compound over time. Farrell and Saloner demonstrate that switching costs and compatibility decisions can lock users into inferior equilibria, providing the microeconomic foundation for the stability of dominant platforms even when superior alternatives exist. Evans and Schmalensee characterize the multi-sided nature of digital platforms more generally, showing that traditional antitrust metrics are poorly calibrated to the dynamics of markets in which price does not reliably indicate competitive distortion.

The existing economics literature does not, however, offer a geometric or thermodynamic account of platform dominance. It treats curvature metaphorically at most. The present paper takes the geometric language seriously as a formal framework.

2.2 Information Geometry and Field Theory

Information geometry, developed principally by Amari and Nagaoka, studies the differential geometry of families of probability distributions. The Fisher information matrix defines a Riemannian metric on statistical manifolds. Curvature in this setting corresponds to departure from flatness, with implications for estimation efficiency, learning dynamics, and the geometry of inference.

The RSVP framework generalizes this approach to a manifold of semantic states, equipped with fields governing the evolution and interaction of those states. The scalar field ϕ encodes accumulated informational density or leverage; the vector field \mathbf{A} encodes directed flows. The resulting picture is analogous to classical field theory, where physical observables arise as geometric invariants of field configurations.

This connection to field theory is not merely heuristic. Thermodynamic field theories of information have been developed in the context of statistical mechanics, where entropy production and free energy quantify the thermodynamic cost of computation. Landauer’s principle establishes that erasure of information has a minimum energetic cost, grounding information-theoretic quantities in physical thermodynamics. The entropy budgets of PlenumHub can be understood as semantic analogues of thermodynamic bounds: upper limits on disorder production per computational step.

2.3 Type Theory, Sheaf Theory, and Formal Semantics

The algebraic foundations of PlenumHub draw on type theory and sheaf theory. Martin-Löf type theory provides a logical framework in which propositions and types correspond, and in which provenance can be expressed as derivational evidence. The

typed structure of semantic spheres ensures that every object carries a machine-verifiable account of its origins.

Sheaf theory, originating in algebraic topology and developed extensively in algebraic geometry, provides the correct framework for reasoning about local-to-global properties of semantic assignments. A sheaf on a topological space or site assigns data to open sets in a manner compatible with restriction and gluing. The sheaf condition enforces that locally consistent data assembles uniquely into global sections, which is precisely the semantic coherence property required for a distributed knowledge substrate.

Applications of sheaf theory to data integration and distributed systems have been explored by Ghrist, Curry, and collaborators under the heading of sheaf-theoretic data fusion. PlenumHub extends this line of work to the setting of collaborative semantic evolution.

2.4 Algorithmic Accountability and Information Asymmetry

A growing literature in computer science and law addresses the problem of algorithmic accountability: the question of how systems that make consequential decisions about human lives can be made transparent, auditable, and contestable. Diakopoulos, Pasquale, and collaborators have catalogued the ways in which proprietary algorithmic systems resist accountability through opacity, complexity, and legal protection. The formal framework developed here provides a complementary technical approach: rather than auditing existing systems after the fact, it specifies substrate-level invariants that make certain accountability failures structurally impossible.

Information asymmetry of the kind described by Zuboff has been studied formally in the economics of information, following Akerlof’s analysis of markets for lemons and Stiglitz’s work on screening and signaling. In those models, asymmetry arises from differential access to quality signals. In the platform context, the asymmetry is more radical: the platform possesses a generative model of the agent’s future behavior that the agent neither has access to nor, in most cases, knows exists. Section 3 formalizes this distinction using Kullback–Leibler divergence and channel capacity.

3 Epistemological Asymmetry and Predictive Leverage

3.1 The Structure of Predictive Asymmetry

Zuboff’s diagnosis identifies a qualitatively new form of informational inequality. Classical information asymmetry, as analyzed in contract theory and market design, concerns differential access to a fixed stock of information: the seller knows the quality of the car; the buyer does not. In contrast, platform-mediated behavioral prediction involves the construction of a generative model that does not merely report existing facts but computes future behavioral trajectories from accumulated behavioral surplus. The agent is not merely uninformed but is the subject of a model that produces predictions about her own conduct which she cannot access, verify, or rebut.

We formalize this asymmetry in information-theoretic terms.

Definition 3.1 (Agent Semantic State). *The agent’s semantic state at time t is a probability distribution $p_t \in \Delta(\mathcal{X})$ over a discrete alphabet \mathcal{X} of behavioral and semantic outcomes. The agent maintains an internal model \hat{p}_t of her own future states, representing her self-prediction.*

Definition 3.2 (Platform Predictive Model). *The platform’s predictive model at time t is a distribution $q_t \in \Delta(\mathcal{X})$ constructed from the agent’s behavioral history $\mathcal{H}_t = \{x_0, x_1, \dots, x_{t-1}\}$ via a learning algorithm \mathcal{A} . Formally, $q_t = \mathcal{A}(\mathcal{H}_t)$.*

Definition 3.3 (Predictive Asymmetry). *The predictive asymmetry between platform and agent at time t is the Kullback–Leibler divergence*

$$\Delta_t = D_{\text{KL}}(p_t \|\hat{p}_t) - D_{\text{KL}}(p_t \|q_t),$$

where $D_{\text{KL}}(p \|q) = \sum_x p(x) \log \frac{p(x)}{q(x)}$ is the KL divergence from q to p .

When $\Delta_t > 0$, the platform’s model q_t is closer to the true behavioral distribution p_t than the agent’s own self-model \hat{p}_t . The platform knows the agent better than the agent knows herself, in the precise sense that the platform’s predictive loss is lower. This is the formal expression of what Zuboff calls the “asymmetry of knowledge and power.”

Proposition 3.4 (Monotone Growth of Asymmetry under Behavioral Extraction). *Suppose the platform updates q_t via Bayesian updating on behavioral observations and the agent’s self-model \hat{p}_t is static or updates on an information set strictly smaller than*

\mathcal{H}_t . Then Δ_t is non-decreasing in t , and under mild consistency conditions, $\Delta_t \rightarrow \infty$ as $t \rightarrow \infty$.

Proof. Bayesian updating is asymptotically consistent: by the Bernstein–von Mises theorem, under regularity conditions, $D_{\text{KL}}(p_t \| q_t) \rightarrow 0$ as $t \rightarrow \infty$ when the true distribution lies in the model class. Meanwhile, if the agent’s self-model is static, $D_{\text{KL}}(p_t \| \hat{p}_t)$ is determined by the divergence between the fixed \hat{p} and the evolving true distribution p_t , which in general grows as behavioral data accumulates and platform predictions sharpen. Hence $\Delta_t = D_{\text{KL}}(p_t \| \hat{p}_t) - D_{\text{KL}}(p_t \| q_t)$ is non-decreasing. \square

3.2 Channel Capacity and Leverage

Predictive asymmetry translates into economic leverage when the platform can act on its superior model to influence the agent’s choices. We formalize this as a channel capacity argument.

Definition 3.5 (Behavioral Channel). *A behavioral channel is a conditional probability distribution $W : \mathcal{X} \rightarrow \Delta(\mathcal{Y})$, where \mathcal{X} is the space of platform interventions (content rankings, interface choices, notification timing) and \mathcal{Y} is the space of agent behavioral responses.*

The platform’s ability to steer the agent is bounded by the channel capacity

$$C = \max_{p \in \Delta(\mathcal{X})} I(X; Y),$$

where $I(X; Y) = H(Y) - H(Y | X)$ is the mutual information between interventions and responses.

Proposition 3.6 (Leverage as Mutual Information). *The platform’s extractable behavioral leverage is bounded above by the mutual information $I(\mathcal{H}_t; Y_t)$ between behavioral history and current response, and is increasing in Δ_t .*

Sketch. The platform selects interventions as a function of its predictive model q_t , which is itself a function of \mathcal{H}_t . By the data processing inequality, $I(\mathcal{H}_t; Y_t) \geq I(q_t; Y_t)$. The tightness of this bound increases as q_t becomes a more efficient statistic of \mathcal{H}_t , i.e., as $D_{\text{KL}}(p_t \| q_t)$ decreases. Therefore leverage is increasing in the quality of prediction, which by Proposition 3.4 is increasing in Δ_t . \square

3.3 Geometric Expression of Asymmetry

In the RSVP framework, predictive asymmetry corresponds to differential curvature between the agent’s internal manifold \mathcal{M}_{int} and the platform’s external manifold \mathcal{M}_{ext} . The agent navigates a low-dimensional, low-curvature internal manifold while the platform operates on a high-dimensional, high-curvature manifold that contains a richer representation of the agent’s state space.

Formally, if \mathcal{M}_{int} has sectional curvature K_{int} and \mathcal{M}_{ext} has sectional curvature K_{ext} , then manifold collapse (Proposition 5.5) occurs whenever $K_{\text{ext}} \gg K_{\text{int}}$ in the vicinity of the platform’s attractor region \mathcal{P} . The KL divergence $D_{\text{KL}}(p_t||q_t)$ can be related to the Fisher–Rao metric on the statistical manifold: the faster q_t converges to p_t , the more accurately the external manifold approximates the agent’s true state geometry, and the more concentrated the curvature becomes around \mathcal{P} .

Remark 3.7. This geometric correspondence suggests a natural operational definition of “extraction rate”: the rate at which Δ_t grows, or equivalently the rate at which the platform’s manifold curvature exceeds the agent’s. PlenumHub suppresses this rate by ensuring that all predictive models available on the substrate are transparent (encoded in provenance graphs), bounded in entropy (limiting model complexity), and governed by staking constraints that penalize models with low quality scores.

3.4 The Right to Predictive Parity

The asymmetry analysis supports a precise normative criterion that we call *predictive parity*: an informational substrate is fair in the relevant sense if and only if the predictive asymmetry Δ_t is bounded uniformly over time. An unbounded Δ_t corresponds to an agent who is progressively dominated by a model she cannot access or contest, which is the formal counterpart of Zuboff’s observation that surveillance capitalism transforms human experience from a subject of self-determination into a raw material for behavioral prediction markets.

Predictive parity is achievable in the PlenumHub framework because all models are represented as semantic spheres with explicit provenance, bounded entropy, and crystal valuations that decay without active maintenance. A platform that does not continuously invest in model quality sees its predictive advantage decay exponentially, while the agent retains access to the same substrate tools. The formal result appears as Proposition 14.4 in Section 14.

4 Formal Objects and Definitions

We collect here all principal formal objects, with precise definitions. Subsequent sections invoke these definitions without repetition.

Definition 4.1 (Information Manifold). *An information substrate is a smooth Riemannian manifold (\mathcal{M}, g) , where \mathcal{M} is a connected, second-countable, Hausdorff differentiable manifold and g is a Riemannian metric tensor encoding informational distance between semantic states. For some applications, a pseudo-Riemannian metric with indefinite signature may be appropriate; we restrict to the Riemannian case unless otherwise noted.*

Definition 4.2 (RSVP Fields). *The RSVP (Relativistic Scalar–Vector Plenum) framework equips (\mathcal{M}, g) with a scalar potential field $\phi : \mathcal{M} \rightarrow \mathbb{R}$ representing accumulated informational density or positional leverage, and a smooth vector field $\mathbf{A} \in \Gamma(T\mathcal{M})$ representing directed flows of semantic evolution. The coupled evolution of ϕ and \mathbf{A} governs the large-scale dynamics of the informational substrate.*

Definition 4.3 (Semantic Sphere). *A semantic sphere is a tuple*

$$\sigma = (I, T, M, E, S),$$

where $I \in \mathcal{I}$ is a content-hash bound identity computed as a cryptographic hash of the remaining components; T is a finite non-empty set of modality types; $M : T \rightarrow \mathcal{V}$ is a modality assignment function mapping each type to a payload in a value space \mathcal{V} ; $E \in \mathbb{R}_{\geq 0}$ is the measured semantic entropy of σ ; and $S = (V_S, E_S)$ is a finite directed acyclic graph encoding derivational provenance, with vertices labeling rule applications and edges encoding data dependencies. We write Σ for the set of all well-formed semantic spheres.

Definition 4.4 (Payload Distribution and Semantic Entropy). *Given a semantic sphere σ , let P_σ be the probability distribution over a discrete token alphabet induced by the joint encoding of modality payloads $\{M(t)\}_{t \in T}$ and provenance graph S . The semantic entropy of σ is*

$$E(\sigma) = - \sum_x P_\sigma(x) \log_2 P_\sigma(x),$$

where the sum is taken over the support of P_σ .

Definition 4.5 (Rule Application). *A rule is a partial function $r : \Sigma \rightarrow \Sigma$ together with a declared entropy budget $\epsilon_r \in \mathbb{R}_{\geq 0}$ and a staking bond $b_r \in \mathbb{R}_{\geq 0}$. A rule r is*

entropy-safe on σ if $r(\sigma)$ is defined and

$$E(r(\sigma)) \leq E(\sigma) + \epsilon_r.$$

Definition 4.6 (Execution Trace). *An execution trace is a finite sequence of rule applications*

$$\tau = r_1; r_2; \cdots; r_n,$$

where each r_i is applied to the result of the preceding application. The trace τ is well-formed on σ_0 if each successive application is defined. The result of applying τ to σ_0 is denoted $\sigma_n = \tau(\sigma_0)$.

Definition 4.7 (Cumulative Entropy Potential). *For a trace $\tau = r_1; \cdots; r_n$, the cumulative entropy potential is*

$$\Phi(\tau) = \sum_{i=1}^n \epsilon_{r_i}.$$

Definition 4.8 (Texture Valuation). *The Texture Crystal valuation of a semantic sphere σ with modality keys $k_1, \dots, k_m \in T$ is*

$$TC(\sigma) = H(\sigma) - \sum_{i \neq j} MI(M(k_i), M(k_j)),$$

where $H(\sigma) = E(\sigma)$ is the semantic entropy and $MI(M(k_i), M(k_j))$ is the mutual information between the payload distributions of modalities k_i and k_j .

Definition 4.9 (Time Valuation). *The Time Crystal valuation of a semantic sphere σ with provenance graph S is*

$$TiC(\sigma) = \sum_{r_i \in V_S} e^{-\lambda(t_{\text{now}} - t_i)} q(r_i),$$

where $\lambda > 0$ is a global decay constant, t_i is the timestamp of rule application r_i , t_{now} is the current time, and $q(r_i) \in [0, 1]$ is a quality score derived from independent rule validation.

Definition 4.10 (Social Semantic Welfare). *Given a universe \mathcal{U} of active semantic spheres in the system, social welfare is defined as*

$$W(\mathcal{U}) = \sum_{\sigma \in \mathcal{U}} TC(\sigma).$$

Definition 4.11 (Engagement Functional). *An engagement functional is a measurable function $J : \Sigma \rightarrow \mathbb{R}_{\geq 0}$ that an engagement-optimizing platform uses to select updates.*

An engagement-driven update rule selects

$$U(\sigma) = \arg \max_{\sigma' \in \mathcal{N}(\sigma)} J(\sigma'),$$

over admissible successors $\mathcal{N}(\sigma)$.

Definition 4.12 (Staking and Slashing). *A staking bond $B = (b_{TC}, b_{TiC}) \in \mathbb{R}_{\geq 0}^2$ is required upon submission of a rule application. If the application is subsequently found to violate its declared entropy budget, a slashing penalty proportional to b_{TC} or b_{TiC} is forfeited. The slash magnitude is governed by a penalty function $\Pi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ increasing in the magnitude of the violation.*

5 The RSVP Geometry of Platform Dominance

5.1 Curvature Concentration as a Model of Extraction

The geometric formalism developed here treats the information substrate as a smooth manifold and models platform dominance as a geometric distortion of that manifold. This is not a metaphor but a formal claim: the mathematical objects involved have precise definitions, and the qualitative claims about platform behavior correspond to verifiable geometric conditions.

Definition 5.1 (Curvature Concentration). *A region $U \subset \mathcal{M}$ exhibits curvature concentration if there exists $\kappa_0 > 0$ such that for a non-negligible family of two-planes $\Pi \subset T_x \mathcal{M}$ with $x \in U$, the sectional curvature satisfies $K(\Pi) > \kappa_0$, or equivalently if the Ricci tensor satisfies $\text{Ric}(v, v) > \kappa_0 |v|^2$ for all unit vectors v in a family of positive measure.*

Definition 5.2 (Scalar Well). *A scalar well centered at $x_0 \in \mathcal{M}$ is a region $U \ni x_0$ where the Laplace–Beltrami operator applied to the scalar field satisfies $\Delta_g \phi(x) > 0$ for $x \in U$, indicating local accumulation of scalar potential with flow toward x_0 under the associated gradient dynamics.*

The Laplace–Beltrami operator in coordinates is $\Delta_g \phi = g^{ij} \nabla_i \nabla_j \phi$, where ∇ is the Levi-Civita connection. Scalar wells correspond geometrically to attractor regions in the gradient flow of ϕ .

5.2 Geodesic Convergence and Jacobi Fields

The connection between curvature concentration and agent herding is made precise through the theory of Jacobi fields.

Proposition 5.3 (Geodesic Convergence under Positive Curvature). *Let (\mathcal{M}, g) satisfy $K > \kappa_0 > 0$ uniformly in a region U . Let $\{\gamma_s\}_{s \in (-\varepsilon, \varepsilon)}$ be a smooth one-parameter family of geodesics in U and let $J(t) = \partial_s \gamma_s(t)|_{s=0}$ be the Jacobi field measuring their separation. Then the Jacobi equation*

$$\frac{D^2 J}{dt^2} + R(J, \dot{\gamma})\dot{\gamma} = 0$$

implies that $J(t)$ has a zero within finite parameter time, and that any initially diverging family reconverges within a parameter distance bounded by $\pi/\sqrt{\kappa_0}$ by the Bonnet–Myers comparison theorem.

Proof. The Jacobi equation is a second-order ODE along γ . Writing $J = f(t)e(t)$ where e is a parallel-transported frame field and f is scalar, the equation reduces to

$$f''(t) + K(\Pi_{J, \dot{\gamma}})f(t) = 0,$$

where $\Pi_{J, \dot{\gamma}}$ is the two-plane spanned by J and $\dot{\gamma}$. For constant $K = \kappa_0 > 0$, solutions are of the form $f(t) = A \sin(\sqrt{\kappa_0}t) + B \cos(\sqrt{\kappa_0}t)$, which vanish at $t = \pi/\sqrt{\kappa_0}$ regardless of initial conditions. The full result under the bound $K \geq \kappa_0 > 0$ follows from the Sturm comparison theorem for the scalar Jacobi equation and from the Bonnet–Myers theorem, which bounds the diameter of \mathcal{M} by $\pi/\sqrt{\kappa_0}$. \square

Remark 5.4. Proposition 5.3 has a direct economic interpretation. Agent trajectories modeled as geodesics of (\mathcal{M}, g) perturbed by the potential $\nabla_g \phi$ reconverge when sectional curvature is positive. Distinct initial choices, which in a neutral substrate would lead to diverse outcomes, are instead funneled toward a common attractor. This is the geometric analogue of the herding, lock-in, and preference homogenization documented in the empirical platform literature.

5.3 Manifold Collapse and Predictive Resolution

A further consequence of platform extraction is what we call manifold collapse. As a platform accumulates behavioral data and increases model capacity, it constructs an effective manifold \mathcal{M}_{ext} of high dimension and resolution, representing the state space of user behavior as inferred from interaction history. This external manifold is richer

and more curved than the agent’s own internal state space \mathcal{M}_{int} , which is sparse and locally parameterized.

When the agent’s trajectory is computed as a geodesic of \mathcal{M}_{int} but displayed and channeled through \mathcal{M}_{ext} , the dynamics of the denser external manifold dominate. Formally, the agent’s effective metric becomes a pullback of the external metric under the embedding $\iota : \mathcal{M}_{\text{int}} \hookrightarrow \mathcal{M}_{\text{ext}}$, yielding a pulled-back metric ι^*g_{ext} whose curvature reflects the external platform’s structure rather than intrinsic agent preferences.

Proposition 5.5 (Manifold Collapse). *Let $\iota : \mathcal{M}_{\text{int}} \hookrightarrow \mathcal{M}_{\text{ext}}$ be an isometric embedding and suppose that \mathcal{M}_{ext} has curvature concentration around a set $\mathcal{P} \subset \mathcal{M}_{\text{ext}}$. Then the induced curvature on \mathcal{M}_{int} under ι^*g_{ext} inherits focusing effects near $\iota^{-1}(\mathcal{P})$, causing geodesics of \mathcal{M}_{int} to converge toward $\iota^{-1}(\mathcal{P})$ rather than follow intrinsic structure.*

Sketch. The second fundamental form of the embedding ι encodes the extrinsic curvature of \mathcal{M}_{int} as a submanifold of \mathcal{M}_{ext} . The Gauss equation relates intrinsic and extrinsic curvatures: $R^{\text{int}} = R^{\text{ext}}|_{\mathcal{M}_{\text{int}}} + (\text{second fundamental form terms})$. Near curvature concentrations in \mathcal{M}_{ext} , the dominant term forces intrinsic curvature to be positive, inducing the Jacobi focusing of Proposition 5.3. \square

Manifold collapse is therefore not a subjective experience of manipulation but a geometric consequence of embedding in a high-curvature external space. The agent experiences choice while traversing geodesics that have been pre-curved by the external platform.

6 Entropy Soundness

6.1 Motivation and Setup

The entropy budget mechanism is designed to prevent silent semantic degradation. In engagement-driven platforms, entropy inflation is a primary instrument of extraction. Fragmentary content, out-of-context quotation, decontextualized emotional provocation, and deliberate ambiguity all increase user engagement by increasing the conditional entropy of the next interaction. None of this activity is structurally penalized; indeed it is actively rewarded.

PlenumHub introduces a formal barrier. Every rule carries a declared entropy budget, and every application is validated against that declaration. The entropy soundness theorem ensures that the cumulative bound is maintained across arbitrary trace lengths.

6.2 The Entropy Soundness Theorem

Theorem 6.1 (Entropy Soundness). *Let $\tau = r_1; \dots; r_n$ be a well-formed execution trace on $\sigma_0 \in \Sigma$, and suppose that each rule r_i is entropy-safe: $E(r_i(\sigma)) \leq E(\sigma) + \epsilon_{r_i}$ for all σ in the domain of r_i . Then*

$$E(\sigma_n) \leq E(\sigma_0) + \Phi(\tau).$$

Proof. We proceed by induction on n .

The base case $n = 0$ is immediate, since $\sigma_0 = \sigma_0$ and $\Phi(\emptyset) = 0$.

Assume the statement holds for all traces of length $n-1$. Let $\sigma_{n-1} = (r_1; \dots; r_{n-1})(\sigma_0)$. By the inductive hypothesis,

$$E(\sigma_{n-1}) \leq E(\sigma_0) + \sum_{i=1}^{n-1} \epsilon_{r_i}.$$

Since r_n is entropy-safe,

$$E(\sigma_n) = E(r_n(\sigma_{n-1})) \leq E(\sigma_{n-1}) + \epsilon_{r_n}.$$

Combining,

$$E(\sigma_n) \leq E(\sigma_0) + \sum_{i=1}^{n-1} \epsilon_{r_i} + \epsilon_{r_n} = E(\sigma_0) + \Phi(\tau).$$

This completes the induction. □

Corollary 6.2 (Conservation Under Zero-Budget Rules). *If $\epsilon_{r_i} = 0$ for all i in a trace τ , then $E(\sigma_n) \leq E(\sigma_0)$, and the trace is entropy-non-increasing.*

Proof. Immediate from Theorem 6.1 with $\Phi(\tau) = 0$. □

Remark 6.3. Theorem 6.1 is non-vacuous only insofar as entropy budgets are constrained by governance rules and are not set arbitrarily large. If any rule may declare $\epsilon_r = \infty$, the bound collapses. The mechanism design of Section 15 addresses this: budgets must be staked and are subject to slashing upon violation.

Remark 6.4. Entropy-safety as defined in Definition 4.5 is a condition on individual rules. One may ask whether this condition is decidable given an arbitrary rule implementation. In general, for Turing-complete rule languages, entropy-safety is undecidable. PlenumHub therefore restricts rule languages to a typed fragment in which entropy bounds are statically verifiable or enforced dynamically via staking.

6.3 Geometric Interpretation of Entropy Boundedness

In the geometric language of Section 5, entropy boundedness corresponds to a constraint on the diffusion of scalar density ϕ . Entropy inflation can be identified with unbounded increase of ϕ over time, corresponding to accumulation of informational density without structural contribution. The entropy bound $E(\sigma_n) \leq E(\sigma_0) + \Phi(\tau)$ translates into a bound on the growth of ϕ along agent trajectories, which in turn prevents the scalar field from developing the unbounded well-depth characteristic of extractive platforms.

More precisely, under a heat-equation model of scalar diffusion,

$$\partial_t \phi = \alpha \Delta_g \phi + f,$$

where f represents exogenous injection of semantic material, the entropy bound constrains $\int_{\mathcal{M}} \phi \, d\text{vol}_g$ from growing faster than $\Phi(\tau)$ per unit time. This is analogous to a conservation law with controlled source term.

7 Engagement Optimization and Entropic Instability

7.1 The Engagement–Entropy Tension

A central claim of this paper is that engagement-maximizing dynamics and entropy-bounded dynamics are thermodynamically incompatible. We formalize this as a theorem.

Definition 7.1 (Conditional Entropy Engagement). *An engagement functional J is conditional-entropy-increasing if for every σ ,*

$$J(\sigma') \geq J(\sigma) \implies H(\sigma' \mid \sigma) \geq H(\sigma \mid \sigma')$$

for generic successors σ' , where $H(\sigma' \mid \sigma)$ denotes the conditional entropy of the next state given the current state under the platform’s predictive model.

Theorem 7.2 (Engagement–Entropy Divergence). *Let J be a conditional-entropy-increasing engagement functional. If no entropy budget constraints are enforced, then repeated application of the engagement-driven update rule U yields either (a) unbounded growth of $E(\sigma_n)$ as $n \rightarrow \infty$, or (b) monotone increase of $\Delta_g \phi$ in a bounded region, leading to curvature concentration.*

Proof. We consider the two cases separately.

Case (a). Suppose that $E(\sigma_n)$ is bounded above by some E_{\max} for all n . Then σ_n lies in a compact set of spheres with bounded entropy. The engagement functional J , being continuous on this set, achieves its maximum. However, since J is conditional-entropy-increasing, the update U always moves toward states of higher conditional entropy. In the presence of bounded entropy, increased conditional entropy must come at the expense of predictive structure: the platform’s predictive model becomes less informative about future states, which contradicts the platform’s incentive to increase predictive resolution. Hence either $E(\sigma_n)$ grows without bound, or the platform’s predictive model degrades, which conflicts with the assumption of engagement optimization.

Case (b). If $E(\sigma_n)$ is bounded but J continues to be maximized, then engagement must be sustained by concentrating predictive leverage rather than inflating global entropy. This corresponds to increasing ϕ locally: the platform deepens its scalar well to maintain conditional entropy by reducing marginal predictability of user states in the well’s vicinity. Formally, $\Delta_g\phi$ increases in the well, and by Definition 5.1, this induces curvature concentration. The Jacobi convergence of Proposition 5.3 then applies.

In both cases, the dynamics are unstable with respect to the criteria of bounded entropy and bounded curvature simultaneously. \square

Remark 7.3. Theorem 7.2 makes precise the informal claim that engagement optimization is thermodynamically incompatible with entropy-bounded evolution. It also provides a formal account of the qualitative observation that engagement-driven platforms either generate increasing noise and distraction (Case a) or produce increasingly concentrated power (Case b). These are not independent pathologies but manifestations of the same underlying geometric instability.

8 Phase Structure and Dynamical Regimes

8.1 Order Parameters

The preceding sections establish local instability results, but do not classify the global dynamical regimes of the informational substrate. We now introduce macroscopic order parameters that permit a phase-theoretic analysis of platform dynamics.

Define the curvature concentration parameter

$$\Omega = \sup_{x \in \mathcal{M}} \lambda_{\max}(\text{Ric}(x)),$$

where λ_{\max} denotes the largest eigenvalue of the Ricci tensor at x .

Define the entropy growth rate

$$\Xi = \limsup_{n \rightarrow \infty} \frac{E(\sigma_n)}{n},$$

whenever the limit superior exists along an execution trace.

These two quantities measure, respectively, geometric concentration and long-run entropy drift.

8.2 Regime Classification

We classify dynamical regimes by the pair (Ω, Ξ) .

Definition 8.1 (Stable Plenum Regime). *The substrate is in the stable plenum regime if $\Omega < \Omega_c$ and $\Xi = 0$ for some curvature threshold $\Omega_c > 0$.*

In this regime, curvature remains bounded and entropy does not grow asymptotically. Geodesic diversity is preserved and semantic evolution remains coherent.

Definition 8.2 (Semantic Heat Death Regime). *The substrate exhibits semantic heat death if $\Xi > 0$ and*

$$\lim_{n \rightarrow \infty} \frac{TC(\sigma_n)}{E(\sigma_n)} = 0.$$

Here entropy grows without proportional structural coherence. The manifold remains approximately flat, but informational density becomes unstructured.

Definition 8.3 (Curvature Singularity Regime). *The substrate exhibits curvature singularity if $\Omega \rightarrow \infty$ on a bounded region $U \subset \mathcal{M}$.*

This corresponds to extractive scalar well formation and geodesic trapping.

Definition 8.4 (Extractive Phase). *The substrate is in the extractive phase if either $\Xi > 0$ or $\Omega \geq \Omega_c$.*

8.3 Phase Transitions

Proposition 8.5 (Engagement-Induced Phase Transition). *Under engagement-driven updates as defined in Section 7, the system undergoes a phase transition from the stable plenum regime to the extractive phase in finite time, unless entropy budgets and curvature constraints are imposed.*

Sketch. By Theorem 7.2, engagement optimization implies either $\Xi > 0$ or monotone increase of $\Delta_g \phi$ in a bounded region. The latter increases Ω via the relation between scalar potential and curvature concentration established in Section 5. Therefore one of the extractive conditions is eventually satisfied. \square

8.4 Phase Diagram Interpretation

The phase space of platform dynamics may therefore be visualized as partitioned into four quadrants:

Flat Curvature / Bounded Entropy	(Stable Plenum)
Flat Curvature / Unbounded Entropy	(Semantic Heat Death)
High Curvature / Bounded Entropy	(Monopoly-Like Concentration)
High Curvature / Unbounded Entropy	(Extractive Singularity)

PlenumHub enforces $\Xi = 0$ asymptotically and bounds Ω by preventing scalar well deepening beyond declared entropy budgets. Engagement-optimized systems do not.

9 Invariant Measures and Long-Run Behavior

9.1 Markovian Evolution of Semantic Spheres

To analyze long-run behavior, we model rule application as a stochastic process. Let Σ denote the space of semantic spheres. Let \mathcal{R} be the admissible rule set.

Assume that rule selection at each step is governed by a probability kernel

$$P(\sigma, A) = \mathbb{P}(r(\sigma) \in A),$$

for measurable subsets $A \subseteq \Sigma$.

This induces a Markov operator \mathcal{P} acting on bounded measurable functions $f : \Sigma \rightarrow \mathbb{R}$ by

$$(\mathcal{P}f)(\sigma) = \int_{\Sigma} f(\sigma') P(\sigma, d\sigma').$$

The induced Markov chain $\{\sigma_n\}$ models long-run semantic evolution.

9.2 Tightness under Entropy Boundedness

We now show that entropy budgets induce tightness of the distribution over spheres.

Proposition 9.1 (Tightness of Entropy-Bounded Evolution). *Suppose entropy soundness (Theorem 6.1) holds and that entropy budgets are uniformly bounded above by ϵ_{\max} . Then for any initial sphere σ_0 , the family of distributions $\{\mu_n\}$ induced by σ_n is tight in the topology induced by E .*

Proof. By Theorem 6.1,

$$E(\sigma_n) \leq E(\sigma_0) + n\epsilon_{\max}.$$

If entropy budgets are globally capped by ϵ_{\max} and further constrained so that long-run average entropy growth $\Xi = 0$ (as in the stable plenum regime), then $E(\sigma_n)$ remains uniformly bounded.

Define compact sets

$$K_M = \{\sigma \in \Sigma : E(\sigma) \leq M\}.$$

Since entropy is finite and Σ is countable up to hash identity equivalence, K_M is precompact under the induced metric on distributions. For sufficiently large M , $\mu_n(K_M) = 1$ for all n , establishing tightness. \square

9.3 Existence of Invariant Measures

We now establish existence of invariant measures under bounded entropy and bounded curvature.

Theorem 9.2 (Existence of Invariant Measure). *Assume entropy boundedness ($\Xi = 0$), uniform entropy budgets, and curvature boundedness ($\Omega < \Omega_c$). Then the Markov operator \mathcal{P} admits at least one invariant probability measure μ^* on Σ .*

Sketch. Under tightness (Proposition 9.1) and bounded curvature preventing collapse to a singular subset, the Krylov–Bogolyubov theorem applies. The sequence of

empirical measures

$$\bar{\mu}_N = \frac{1}{N} \sum_{n=1}^N \delta_{\sigma_n}$$

is tight and therefore has a weakly convergent subsequence. Any weak limit point μ^* satisfies $\mathcal{P}^* \mu^* = \mu^*$. \square

Remark 9.3. The existence of an invariant measure does not imply uniqueness or ergodicity. Multiple invariant measures may correspond to distinct semantic basins of attraction. However, the key contrast with engagement-driven systems is that entropy-bounded dynamics admit stationary distributions, while unbounded entropy growth precludes invariant probability measures entirely.

9.4 Non-Existence under Engagement Optimization

Proposition 9.4 (Absence of Stationarity under Entropy Divergence). *If $\Xi > 0$, then no invariant probability measure exists on Σ .*

Proof. If $\Xi > 0$, then $E(\sigma_n) \rightarrow \infty$ along some subsequence. For any compact set $K_M = \{\sigma : E(\sigma) \leq M\}$, we have $\mu_n(K_M) \rightarrow 0$ as $n \rightarrow \infty$. Therefore the sequence of distributions is not tight, and by Prokhorov’s theorem, no weakly convergent subsequence exists. Hence no invariant probability measure can exist. \square

9.5 Interpretation

The invariant-measure analysis sharpens the earlier thermodynamic claims.

Entropy-bounded substrates admit stationary semantic distributions. Engagement-driven systems do not. They are dynamically non-stationary by construction.

This difference is structural rather than parametric. It is not a matter of “too much engagement” but of whether the underlying dynamics permit tightness in semantic space.

The stable plenum regime therefore corresponds to a mathematically well-defined stationary phase. Extractive regimes do not admit such stability.

10 Structural Antitrust as a Categorical Invariant

10.1 The Category of Semantic Evolution

We formalize semantic evolution categorically.

Definition 10.1 (Category of Semantic Spheres). *Define a category **Sphere** whose objects are semantic spheres $\sigma \in \Sigma$, and whose morphisms*

$$f : \sigma \rightarrow \sigma'$$

are entropy-safe rule applications or well-formed execution traces satisfying entropy soundness.

Composition of morphisms is trace concatenation. Identity morphisms correspond to empty traces.

Entropy soundness guarantees that composition preserves entropy bounds.

Proposition 10.2 (Closure under Composition). *If $f : \sigma \rightarrow \sigma'$ and $g : \sigma' \rightarrow \sigma''$ are entropy-safe morphisms, then $g \circ f : \sigma \rightarrow \sigma''$ is entropy-safe with cumulative budget $\epsilon_f + \epsilon_g$.*

Proof. Follows directly from Theorem 6.1. □

10.2 Metric Functor and Geometric Realization

We now connect semantic evolution to geometry.

Definition 10.3 (Geometric Realization Functor). *Define a functor*

$$F : \mathbf{Sphere} \rightarrow \mathbf{Man}$$

mapping each semantic sphere σ to a point $F(\sigma) \in (\mathcal{M}, g)$ in the information manifold, and each morphism $f : \sigma \rightarrow \sigma'$ to a smooth path between $F(\sigma)$ and $F(\sigma')$.

The metric structure g determines informational distances. The scalar field ϕ and vector field \mathbf{A} determine local geometry.

Entropy-bounded morphisms correspond to paths of bounded energy in (\mathcal{M}, g) .

10.3 Forbidden Endofunctors

We now formalize extraction as a categorical deformation.

Definition 10.4 (Extractive Endofunctor). *An endofunctor*

$$E : \mathbf{Sphere} \rightarrow \mathbf{Sphere}$$

is extractive if it fails to preserve entropy soundness, crystal conservation, or sheaf closure.

More precisely, E is extractive if for some σ ,

$$E(\sigma)$$

either increases entropy beyond declared budget, increases crystal valuation without corresponding semantic novelty, or produces incompatible local sections under the sheaf structure.

Theorem 10.5 (Structural Antitrust Invariance). *If a substrate enforces entropy soundness, crystal conservation, and sheaf closure as categorical invariants, then no extractive endofunctor exists within \mathbf{Sphere} .*

Proof. Any endofunctor must preserve objects and morphisms of the category.

If E violates entropy soundness, then it produces a morphism not in \mathbf{Sphere} , contradicting closure under composition.

If E inflates crystal mass without entropy injection, it violates Proposition 14.5, hence is not a valid morphism.

If E violates sheaf closure, it produces an object outside the sheaf subcategory.

Therefore no endofunctor preserving categorical structure can be extractive. \square

10.4 Interpretation

Traditional antitrust prohibits specific behaviors after the fact. Categorical structural antitrust removes the morphisms that would implement those behaviors.

In this framework, extraction is not illegal but ill-typed.

A platform layer attempting to override substrate invariants would correspond to an external functor not preserving the categorical structure. Such a functor is not admissible within the substrate's algebra.

This formalizes the claim that structural antitrust is not regulatory but ontological: the space of allowed transformations is restricted at the level of type formation.

11 Computational Complexity and Decidability

11.1 Computability of Semantic Entropy

The semantic entropy

$$E(\sigma) = - \sum_x P_\sigma(x) \log_2 P_\sigma(x)$$

is defined over the token distribution induced by modality payloads and provenance graphs.

If payloads are finite and explicitly represented, entropy is computable in time $O(N)$ where N is the total token count.

However, if rule languages permit generative or implicit representations (e.g., compressed encodings or procedural generators), exact entropy computation may require expansion of implicit structures, potentially leading to exponential complexity.

Remark 11.1. In practice, entropy computation must be restricted to explicitly materialized payloads or approximated using sampling-based estimators.

11.2 Decidability of Entropy-Safety

Definition 4.5 requires that a rule r satisfy

$$E(r(\sigma)) \leq E(\sigma) + \epsilon_r.$$

If the rule language is Turing-complete, determining whether this inequality holds for all σ in the rule's domain is, in general, undecidable by reduction from the halting problem.

Proposition 11.2 (Undecidability under Turing-Completeness). *If the rule language is Turing-complete, then verifying entropy-safety for arbitrary rules is undecidable.*

Sketch. Given a Turing machine T and input x , construct a rule that outputs high-entropy payload if $T(x)$ halts and low-entropy payload otherwise. Entropy-safety verification would decide halting of $T(x)$, which is undecidable. \square

11.3 Restricted Rule Languages

To restore decidability, PlenumHub must restrict rule languages to fragments with bounded expressive power.

Two viable approaches are:

1. Typed functional languages without general recursion.
2. Domain-specific rule languages with static entropy bounds.

Under such restrictions, entropy-safety becomes decidable or at least semi-decidable.

Proposition 11.3 (Decidability under Finite-State Rules). *If rule applications correspond to finite-state transformations over bounded payload alphabets, then entropy-safety is decidable in polynomial time.*

Proof. Finite-state transformations have computable output size bounds. Entropy can be computed directly from explicit output distributions, and comparison against declared ϵ_r is finite-time. \square

11.4 Complexity of Sheaf Validation

Sheaf closure requires verifying compatibility on overlaps:

$$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}.$$

If contexts are finite and overlap size is bounded, validation is polynomial in context size.

However, if context intersections grow combinatorially, validation may be expensive.

Remark 11.4. Efficient implementation requires bounded context arity and canonical indexing of overlaps. Without such constraints, worst-case sheaf validation may be exponential in the number of overlapping contexts.

11.5 Complexity of Crystal Computation

Texture valuation requires computing pairwise mutual information:

$$TC(\sigma) = H(\sigma) - \sum_{i \neq j} MI(M(k_i), M(k_j)).$$

If there are m modalities and each modality has N tokens, computing all pairwise MI values is $O(m^2N)$.

For large m , approximations via sampling or sparse dependency graphs may be necessary.

11.6 Implications

The substrate’s guarantees depend on enforceable invariants.

Three constraints are necessary:

1. Rule languages must be restricted to decidable fragments.
2. Entropy computation must be bounded and auditable.
3. Validation algorithms must be computationally tractable.

These constraints are engineering requirements rather than theoretical weaknesses. They delineate the boundary between provable structural guarantees and infeasible universal verification.

Remark 11.5. Acknowledging computational limits strengthens the proposal. The invariants are powerful but not free. They require careful language design and protocol-level enforcement.

12 Governance and Protocol Upgrade Morphisms

12.1 The Governance Problem

All previous sections assume that entropy budgets, decay constants, rule languages, and staking requirements are fixed. In practice, any long-lived substrate must evolve.

The critical risk is that governance becomes a scalar well: a centralized locus of control capable of modifying invariants in ways that reintroduce extractive dynamics.

We therefore formalize protocol upgrades as constrained morphisms in an extended category.

12.2 Protocol as a Structured Object

Let $\mathcal{P} = (\mathcal{L}, \epsilon_{\max}, \lambda, \Pi)$ denote the protocol configuration. The component \mathcal{L} specifies the admissible rule language governing semantic transformations. The parameter ϵ_{\max} defines the global entropy budget cap that bounds cumulative semantic expansion. The

constant λ determines the exponential decay rate governing Time Crystal valuation. The function Π specifies the slashing penalty applied to violations of declared entropy budgets or quality misrepresentation.

Given a fixed protocol configuration \mathcal{P} , we define a category $\mathbf{Sphere}_{\mathcal{P}}$ whose objects are semantic spheres valid under \mathcal{P} and whose morphisms are entropy-safe rule applications expressible in \mathcal{L} and bounded by ϵ_{\max} . Composition of morphisms is given by sequential rule execution, and identity morphisms correspond to null transformations. The categorical structure therefore depends parametrically on \mathcal{P} , and changes to the protocol configuration induce functorial transformations between corresponding sphere categories.

We extend \mathbf{Sphere} to a category $\mathbf{Sphere}_{\mathcal{P}}$ parameterized by protocol configuration.

12.3 Upgrade Morphisms

Definition 12.1 (Protocol Upgrade Morphism). *A protocol upgrade is a morphism*

$$U : \mathcal{P} \rightarrow \mathcal{P}'$$

such that $\mathbf{Sphere}_{\mathcal{P}}$ embeds faithfully into $\mathbf{Sphere}_{\mathcal{P}'}$, and all previously valid morphisms remain valid under \mathcal{P}' .

This embedding condition prevents retroactive invalidation of past semantic spheres.

Definition 12.2 (Invariant-Preserving Upgrade). *An upgrade $U : \mathcal{P} \rightarrow \mathcal{P}'$ is said to be invariant-preserving if, under the upgraded protocol configuration \mathcal{P}' , entropy soundness continues to hold with globally bounded budgets, crystal conservation remains enforced so that no net valuation can be generated without corresponding semantic novelty, sheaf-theoretic closure of semantic assignments is preserved, and staking mechanisms remain enforceable in a manner sufficient to guarantee slashing for entropy violations and quality misrepresentation. In particular, an invariant-preserving upgrade must not introduce any rule, relaxation, or parameter modification that would permit violation of these structural constraints.*

12.4 Monotonicity Constraints

We impose monotonicity conditions on upgrades.

Proposition 12.3 (Entropy Cap Monotonicity). *Any invariant-preserving upgrade must satisfy*

$$\epsilon'_{\max} \leq \epsilon_{\max}.$$

Proof. If $\epsilon'_{\max} > \epsilon_{\max}$, then rule applications previously disallowed under entropy soundness could become admissible, potentially enabling entropy divergence and violating Theorem 6.1. Therefore entropy caps may tighten but not relax. \square

Proposition 12.4 (Decay Non-Relaxation). *Any invariant-preserving upgrade must satisfy*

$$\lambda' \geq \lambda.$$

Proof. Reducing λ slows decay of TiC , potentially enabling long-term lock-in and curvature concentration. To preserve bounded influence, decay may accelerate but not decelerate. \square

12.5 Governance as a Higher-Order Category

We define a category **Protocol** whose objects are protocol configurations \mathcal{P} and whose morphisms are invariant-preserving upgrades.

The composite system can then be viewed as a fibration:

$$\mathbf{Sphere}_{(-)} \rightarrow \mathbf{Protocol}.$$

Governance becomes higher-order structure constrained by categorical preservation laws.

Theorem 12.5 (Governance Non-Extraction Condition). *If all protocol upgrades are invariant-preserving morphisms in **Protocol**, then governance cannot introduce extractive dynamics without violating categorical structure.*

Proof. Any extractive modification must either relax entropy caps, reduce decay, weaken slashing, or permit non-sheaf-compatible merges. Each such modification violates the invariant-preserving condition. Therefore extractive governance cannot be expressed as a morphism in **Protocol**. \square

12.6 Interpretation

Governance is not an external layer but a higher-order morphism subject to invariants.

The risk of governance capture corresponds to allowing non-monotonic upgrades. By restricting upgrades to invariant-preserving morphisms, the system prevents governance from deepening scalar wells. Thus even protocol evolution is entropy-bounded.

13 Sheaf-Theoretic Closure

13.1 The Need for Global Coherence

A distributed knowledge substrate faces a structural challenge that does not arise in centralized systems: local consistency does not automatically imply global consistency. An agent may hold beliefs that are individually coherent but mutually contradictory across different contexts or conversational threads. A malicious actor may deliberately present contradictory semantic assignments to different parts of the system. Sheaf theory provides the correct formal framework for specifying and enforcing global coherence from local data.

13.2 Presheaves and the Sheaf Condition

Let \mathcal{C} be a category of semantic contexts. In the simplest case, \mathcal{C} is a poset under inclusion, where objects are context identifiers and morphisms are specialization relations. More generally, \mathcal{C} may be equipped with a Grothendieck topology specifying which families of morphisms constitute coverings.

Definition 13.1 (Presheaf of Semantic Assignments). *A presheaf of semantic assignments is a functor $\mathcal{F} : \mathcal{C}^{op} \rightarrow \mathbf{Set}$ assigning to each context $U \in \mathcal{C}$ a set $\mathcal{F}(U)$ of admissible semantic assignments, and to each morphism $\phi : V \rightarrow U$ a restriction map $\mathcal{F}(\phi) : \mathcal{F}(U) \rightarrow \mathcal{F}(V)$, compatible with composition and identities.*

Definition 13.2 (Sheaf Condition). *The presheaf \mathcal{F} is a sheaf if for every context U and every covering family $\{f_i : U_i \rightarrow U\}_{i \in I}$, the following sequence is an equalizer:*

$$\mathcal{F}(U) \xrightarrow{\prod_i \mathcal{F}(f_i)} \prod_{i \in I} \mathcal{F}(U_i) \rightrightarrows \prod_{i, j \in I} \mathcal{F}(U_i \times_U U_j).$$

Concretely, this requires: (Locality) if $s, t \in \mathcal{F}(U)$ agree on each U_i , then $s = t$; and (Gluing) if $(s_i)_{i \in I}$ with $s_i \in \mathcal{F}(U_i)$ satisfies $s_i|_{U_i \times_U U_j} = s_j|_{U_i \times_U U_j}$ for all i, j , then there exists a unique $s \in \mathcal{F}(U)$ with $s|_{U_i} = s_i$ for all i .

Proposition 13.3 (Contradiction Impossibility under Sheaf Closure). *If \mathcal{F} satisfies the sheaf condition and an agent presents semantic assignments $(s_i) \in \prod_i \mathcal{F}(U_i)$ that are mutually compatible on overlaps, then there exists a unique global section $s \in \mathcal{F}(U)$ extending all s_i . Conversely, if no such global section exists, the local assignments are incompatible, and the system rejects the submission.*

Proof. This is the gluing axiom of the sheaf condition, applied contrapositively. Existence and uniqueness of glued sections follow from the equalizer property of Definition 13.2. \square

13.3 Semantic Spheres as Sections

In PlenumHub, each semantic sphere σ is interpreted as a global section of \mathcal{F} over a context $U(\sigma)$ determined by its provenance graph S . Modality assignments $M : T \rightarrow \mathcal{V}$ provide local sections over individual modality contexts. The sheaf condition enforces that these local assignments cohere uniquely to a global section, ensuring that the sphere is internally consistent across all modalities and all contextual specializations simultaneously.

An attempted merge of two spheres σ_1 and σ_2 produces a valid merged sphere σ_{12} if and only if the corresponding sections are compatible on the overlap $U(\sigma_1) \cap U(\sigma_2)$. This provides a precise algebraic criterion for merge admissibility that replaces ad hoc conflict resolution heuristics.

13.4 Media–Quine Closure

A related closure condition ensures completeness. Define a closure operator $Q : \Sigma \rightarrow \Sigma$ by $Q(\sigma)$ being the unique minimal extension of σ that populates all required modalities in T with their canonical derivations. Formally, Q is an idempotent operator:

Proposition 13.4 (Idempotency of Closure). $Q(Q(\sigma)) = Q(\sigma)$ for all $\sigma \in \Sigma$.

Proof. $Q(\sigma)$ is defined as the minimal sphere satisfying closure constraints. Applying Q again produces the minimal extension of an already-closed sphere, which is itself. Hence $Q \circ Q = Q$. \square

Closure prevents fragmentation. A sphere that is not Q -closed is not visible in the system. Fragmentary content that relies on external references for completion is structurally ineligible for participation, removing one of the primary vectors of platform dependency creation.

14 Crystal Valuations and Welfare Economics

14.1 From Accounting to Economics

The crystal valuation functions TC and TiC were introduced in Section 4 as formal definitions. We now interpret them within welfare economics and use that interpretation to assess their equilibrium properties.

In conventional engagement platforms, utility is effectively proportional to visibility, follower count, or algorithmic promotion probability. These metrics exhibit strong increasing returns to scale: larger accounts attract more attention, more data, and more algorithmic promotion in a self-reinforcing cycle. The resulting equilibria concentrate value in early movers and structurally exclude late entrants regardless of quality.

The crystal economy severs this feedback loop by grounding utility in structural properties of semantic content rather than positional properties of social networks.

14.2 Texture Crystals as Allocative Efficiency

The Texture Crystal valuation

$$TC(\sigma) = H(\sigma) - \sum_{i \neq j} MI(M(k_i), M(k_j))$$

can be interpreted as a measure of allocative efficiency of informational novelty. The entropy $H(\sigma)$ represents the total informational content of the sphere, while the mutual information sum represents the portion of that content that is redundant across modalities. Subtracting redundancy yields net non-overlapping informational contribution.

This is analogous to the concept of surplus in welfare economics. Total revenue minus transfer payments equals net social contribution. Here, total entropy minus duplicated entropy equals net semantic surplus.

Proposition 14.1 (Non-Negativity under Compatibility). *If the modality payloads $\{M(k_i)\}$ are mutually independent, then $MI(M(k_i), M(k_j)) = 0$ for all $i \neq j$, and $TC(\sigma) = H(\sigma) \geq 0$.*

Proof. Independence implies zero mutual information by definition. The result follows from non-negativity of Shannon entropy. \square

Proposition 14.2 (Redundancy Minimization at Equilibrium). *Suppose agents choose modality assignments M to maximize $TC(\sigma)$ subject to fixed total entropy $H(\sigma) = E_0$.*

Then optimal assignments minimize $\sum_{i \neq j} MI(M(k_i), M(k_j))$, which is achieved when modality payloads are as independent as possible given the constraint.

Proof. The objective $TC(\sigma) = E_0 - \sum_{i \neq j} MI(M(k_i), M(k_j))$ is maximized when the mutual information sum is minimized. Mutual information is non-negative, and its minimum value of zero is achieved if and only if the modality distributions are mutually independent. Under fixed total entropy, this corresponds to a maximum-entropy product distribution over the modality product space. \square

Remark 14.3. Proposition 14.2 shows that the TC objective aligns individual maximization with socially efficient allocation of informational novelty. Agents who duplicate content across modalities are not penalized by a social planner but are penalized directly by the valuation formula. The conservation law is internal to the reward structure rather than externally imposed.

14.3 Time Crystals as Decay-Adjusted Provenance Value

The Time Crystal valuation

$$TiC(\sigma) = \sum_{r_i \in V_S} e^{-\lambda(t_{\text{now}} - t_i)} q(r_i)$$

implements a form of provenance-weighted persistent value with exponential discounting. The exponential decay ensures that the value of a rule application diminishes unless it generates active semantic descendants that themselves accrue TiC .

Proposition 14.4 (Decay of Unextended Influence). *If a rule application r_i generates no semantic descendants in $[t_i, t_{\text{now}}]$, then its contribution to $TiC(\sigma)$ decays as $e^{-\lambda(t_{\text{now}} - t_i)} q(r_i) \rightarrow 0$ as $t_{\text{now}} \rightarrow \infty$.*

Proof. Immediate from the exponential decay factor. \square

This property eliminates lock-in based on early arrival. A foundational contribution to a research corpus does not accrue permanent dominance; its TiC value decays unless the contribution continues to be developed and extended by semantic descendants.

Proposition 14.5 (No Crystal Arbitrage). *No finite sequence of rule applications can generate positive net crystal mass TC or TiC without the injection of genuinely novel semantic information, where novelty is measured by entropy E not already present in σ_0 .*

Sketch. Any increase in $H(\sigma)$ beyond $E(\sigma_0) + \Phi(\tau)$ is ruled out by Theorem 6.1. Any increase in TC through reduction of mutual information requires genuine structural reorganization of modality assignments, which is itself a form of semantic novelty. Any increase in TiC requires valid rule applications with positive $q(r_i)$, which are subject to independent quality validation. Combining these constraints, net crystal mass cannot be inflated without genuine informational input. \square

14.4 Welfare Alignment and Open Problems

The welfare function $W(\mathcal{U}) = \sum_{\sigma} TC(\sigma)$ is maximized when the aggregate redundancy of the knowledge base is minimized and its aggregate semantic entropy is maximized subject to entropy budget constraints. This aligns with a conception of intellectual welfare as the efficient allocation of distinct perspectives and forms of knowledge.

Remark 14.6 (Open Problem: Nash–Optimal Alignment). A full proof that Nash equilibria of the contribution game coincide with or approximate socially optimal configurations requires additional assumptions: convexity of TC as a function of strategy spaces, bounded rationality parameters, and independence conditions on the quality metric q . These conditions may not hold in general. We conjecture that under mild regularity assumptions, equilibria are approximately welfare-optimal in the sense of Price of Anarchy bounds, but this remains an open research problem. The welfare analysis here establishes only that individual and social incentives are directionally aligned, not that they coincide.

15 Mechanism Design and Manipulation Resistance

15.1 Adversarial Threats

A formal system that generates economic rewards will be subject to adversarial behavior. We identify the primary threat vectors and analyze the mechanisms that resist them.

The first threat is quality metric inflation: agents who collude to assign inflated $q(r_i)$ scores to low-quality rule applications in order to accumulate TiC before exponential decay. The second threat is entropy budget exploitation: agents who declare large ϵ_r budgets to launder entropy increases. The third threat is identity forgery: attempts to assign fraudulent provenance by generating colliding hashes for I .

15.2 Stake-Backed Entropy Discipline

Proposition 15.1 (Dominance of Honest Strategy under Proper Staking). *Suppose the expected payoff from an entropy-violating manipulation is $G > 0$ and the expected slashing penalty is $L > 0$. If the governance system calibrates staking bonds B and verification probability p such that $p \cdot \Pi(\delta) > G$ where δ is the magnitude of the entropy violation and Π is the penalty function, then entropy-breaching strategies are strictly dominated by the honest strategy.*

Proof. Expected utility of manipulation is $G - p \cdot \Pi(\delta)$. By hypothesis, this is negative. Expected utility of honest strategy is non-negative. Hence honest strategy strictly dominates. \square

Corollary 15.2 (Budget Declaration Integrity). *Under proper staking, agents have no incentive to declare entropy budgets larger than the actual entropy change their rules produce, since over-declaration increases stake requirements without increasing expected return.*

Proof. Over-declaration increases the stake required by a factor proportional to the excess. Since the rule application produces the same semantic output regardless, the additional stake is pure cost with no offsetting benefit under proper verification. \square

15.3 Quality Metric Integrity

The quality score $q(r_i)$ must be derived from validation that is independent of the submitting agent. The system achieves this through a commit-reveal protocol in which validators submit quality assessments prior to viewing each other’s evaluations, and in which validator stakes are slashed if assessments deviate from independently verified ground truth.

Proposition 15.3 (Collusion Resistance of Independent Validation). *Suppose quality validation is performed by a committee of k independent validators, each staked, and suppose that the colluding coalition controls at most $f < k/3$ validators. Then quality inflation by the coalition is detectable with probability approaching one as the committee size grows, under Byzantine fault tolerance assumptions.*

Sketch. This follows from standard Byzantine fault tolerance results (e.g., the PBFT protocol) under the assumption that honest validators constitute a two-thirds supermajority. Under these conditions, the protocol reaches consensus on the correct quality assessment, and deviating validators are identified and slashed. \square

15.4 Cryptographic Identity

The content-hash identity $I = \text{Hash}(T, M, S)$ binds sphere identity to content. Integrity of this mechanism rests on the collision resistance of the underlying hash function.

Proposition 15.4 (Identity Integrity under Collision Resistance). *If the hash function is collision-resistant, then for any two distinct spheres $(T, M, S) \neq (T', M', S')$, the probability that $\text{Hash}(T, M, S) = \text{Hash}(T', M', S')$ is negligible in the security parameter.*

Proof. This is a direct consequence of the definition of collision resistance and follows from standard reductions to the hardness assumptions underlying the chosen hash function (e.g., SHA-3 under preimage resistance assumptions). \square

16 Comparison of Regimes

The two regimes under consideration, extractive platform dynamics and entropy-bounded substrate dynamics, differ not in degree but in kind. The differences are architectural rather than parametric.

Under extractive platforms, the optimization target is engagement, defined as maximized continuation probability and behavioral prediction accuracy. Growth corresponds to unregulated accumulation of interactions, each of which generates behavioral surplus for the platform without structural validation. Value is measured in attention metrics—clicks, sessions, impressions—that are independent of semantic coherence. Governance is centralized: the platform sets rules unilaterally, modifies them in response to competitive and regulatory pressure, and enforces them asymmetrically. The manifold \mathcal{M} develops curvature concentration around platform-controlled regions. Switching costs correspond to geodesic trapping in the high-curvature well.

Under PlenumHub, the optimization target is bounded evolution of meaning. Growth corresponds to entropy-regulated merges, each of which is validated against declared budgets and subject to slashing upon violation. Value is measured in structural coherence (TC) and provenance-weighted persistence (TiC), both of which require genuine semantic contribution. Governance is rule-native and algebraically mediated: the substrate’s invariants determine permissible trajectories, and no central authority can modify them without a formal protocol upgrade process subject to stakeholder validation. The manifold \mathcal{M} remains approximately flat, with geodesics following intrinsic semantic structure rather than platform-induced curvature.

The shift is not incremental but ontological. The substrate’s invariants define the phase space of permissible evolution. Extraction is not merely discouraged but is structurally inadmissible: any rule application that would constitute extraction violates a provably maintained invariant and is rejected by the type system.

This is the sense in which PlenumHub constitutes substrate-level structural antitrust. Traditional antitrust attempts to prevent specific anticompetitive behaviors while leaving the substrate unchanged. PlenumHub modifies the substrate so that the behaviors in question cannot arise without violating foundational algebraic constraints.

17 Relationship to Existing Proposals

Several existing proposals address platform concentration through technological means. Data portability frameworks, including the European Digital Markets Act and the Solid project’s linked data architecture, focus on reducing switching costs by ensuring that user data can migrate between platforms. These interventions are valuable but address the symptom rather than the mechanism: they reduce curvature by enabling exit from scalar wells, but do not prevent new wells from forming.

Protocol-based approaches, including those developed by Protocol Labs and similar organizations, focus on decentralizing storage and computation. Decentralization of infrastructure reduces the surface area for extraction but does not by itself constrain the semantic dynamics of content evolution. A decentralized platform can still optimize for engagement and produce curvature concentration if its incentive structures permit it.

The RSVP–PlenumHub framework operates at a different layer. It does not merely redistribute ownership of infrastructure or make exit easier. It modifies the algebraic invariants governing semantic evolution, ensuring that the conservation laws of the crystal economy are satisfied regardless of ownership structure. An operator who deploys PlenumHub infrastructure cannot silently inflate entropy or extract rents without violating the staking constraints and type system that are constitutive of the substrate itself.

This makes the present proposal more radical than portability or decentralization, and also more technically demanding. The formal guarantees depend on the correct implementation and governance of the type system, entropy budget mechanism, and staking protocol. These are engineering and governance challenges that constitute the primary gap between the theoretical framework and practical deployment.

18 Unified Structural Stability Theorem

18.1 Preliminaries

We now consolidate the preceding formal developments into a single structural result. The substrate under consideration satisfies entropy soundness with globally bounded entropy budgets, curvature boundedness on the information manifold, sheaf-theoretic closure for semantic assignments, conservation of crystal valuation in the sense that no net valuation can be generated without corresponding semantic novelty, and governance constraints that restrict protocol evolution to invariant-preserving upgrades. Let \mathcal{D} denote the discrete-time dynamical system on Σ induced by entropy-safe rule applications under such a protocol configuration.

18.2 Structural Stability

Theorem 18.1 (Structural Stability of Entropy-Bounded Substrates). *Suppose the informational substrate satisfies entropy soundness with globally bounded budgets, curvature boundedness $\Omega < \Omega_c$ for some finite threshold, sheaf-theoretic closure of semantic assignments, crystal conservation in the sense of no-arbitrage valuation, and governance restricted to invariant-preserving protocol upgrades. Then the induced dynamical system \mathcal{D} admits at least one invariant probability measure, excludes curvature singularity formation, excludes semantic heat death, disallows extractive endofunctors within the category **Sphere**, and maintains bounded influence under exponential Time Crystal decay.*

Sketch. Existence of an invariant probability measure follows from the tightness of entropy-bounded evolution and the boundedness of curvature, which together satisfy the hypotheses of the Krylov–Bogolyubov argument for the existence of stationary distributions. Curvature singularity formation requires unbounded increase in scalar potential ϕ or concentration of Ricci curvature beyond Ω_c . Such concentration necessitates entropy injection beyond globally bounded budgets, contradicting entropy soundness and the monotonicity constraints imposed on governance. Semantic heat death requires asymptotic entropy growth rate $\Xi > 0$ together with vanishing structural coherence relative to entropy. Entropy soundness ensures $\Xi = 0$ under bounded budgets, while crystal conservation prevents the accumulation of entropy without proportional structural contribution. Extractive endofunctors are excluded by the categorical argument establishing that any transformation violating entropy, crystal, or sheaf invariants lies outside the morphisms of **Sphere**. Bounded influence follows from exponential decay in Time Crystal valuation combined with governance constraints

preventing relaxation of the decay parameter. These properties jointly ensure that the system remains within the stable plenum regime characterized previously. \square

18.3 Thermodynamic Antitrust

Corollary 18.2 (Thermodynamic Antitrust). *In any informational substrate satisfying the hypotheses of Theorem 18.1, monopoly-like curvature concentration cannot arise endogenously under rational agent dynamics.*

Proof. Monopoly-like concentration corresponds geometrically to curvature concentration beyond the bounded threshold Ω_c or to scalar well deepening sufficient to induce geodesic trapping. Such regimes require either entropy divergence or governance relaxation of invariant constraints. Both possibilities are excluded by the hypotheses of Theorem 18.1. Therefore monopoly-like curvature concentration is not dynamically reachable within the admissible phase space. \square

18.4 Interpretation

The preceding theorem formalizes the central claim of this work. Extraction is not merely an undesirable equilibrium but a violation of conserved structural invariants embedded in the substrate. When entropy budgets are globally bounded, crystal valuation enforces conservation of semantic contribution, sheaf closure guarantees global coherence from local consistency, and governance is restricted to invariant-preserving morphisms, extractive regimes become unreachable states of the dynamical system. The exclusion is structural rather than behavioral. The algebra of permissible transformations does not contain the morphisms required to implement extraction. Under these conditions, structural antitrust is not an external intervention but an intrinsic property of the informational field itself.

19 Empirical Calibration and Measurability

19.1 The Gap Between Theory and Observation

The formal framework developed in the preceding sections is expressed in terms of continuous manifolds, smooth fields, and abstract categorical structures. A natural question is whether and how these theoretical quantities correspond to measurable observables in real digital systems. This section addresses that question. We do not

claim to resolve it fully; the mapping from theory to empirical method is itself a research program. We aim instead to identify the key theoretical quantities, propose candidate empirical proxies for each, and note the methodological challenges involved.

The most practically important quantities are: (i) the curvature concentration parameter Ω , characterizing platform dominance geometrically; (ii) the entropy growth rate Ξ , characterizing long-run semantic drift; (iii) the predictive asymmetry Δ_t , characterizing the informational gap between platform and agent; and (iv) the crystal valuations TC and TiC , characterizing the structural quality of semantic contributions.

19.2 Empirical Proxies for Curvature Concentration

In the continuous manifold model, curvature concentration is captured by the eigenvalues of the Ricci tensor. In discrete graph-theoretic models of information networks, analogous concentration phenomena manifest as spectral properties of the adjacency or Laplacian matrix. The Ollivier–Ricci curvature on graphs provides a direct discrete analogue of Riemannian sectional curvature and has been used to detect bottleneck structure and community concentration in real networks.

Concretely, the Ollivier–Ricci curvature of an edge (x, y) in a graph G is

$$\kappa(x, y) = 1 - W_1(\mu_x, \mu_y),$$

where μ_x and μ_y are probability distributions supported on the neighborhoods of x and y , and W_1 denotes the first Wasserstein distance. Negative curvature indicates tree-like divergence; positive curvature indicates local connectivity. Platform extraction corresponds to strongly positive curvature concentrated around a small number of high-degree hub nodes, with negative curvature on the periphery representing the sparse and dispersed agent population.

Remark 19.1 (Proposed Empirical Metric: Curvature Concentration Ratio). The curvature concentration ratio \mathcal{R}_κ can be defined as the fraction of total positive Ricci curvature supported on the top $k\%$ of nodes by degree. In a neutral substrate, \mathcal{R}_κ should be close to $k/100$. In an extractive substrate, $\mathcal{R}_\kappa \gg k/100$, indicating that curvature is disproportionately concentrated at high-degree platform nodes.

19.3 Empirical Proxies for Entropy Growth

The entropy growth rate $\Xi = \limsup_{n \rightarrow \infty} E(\sigma_n)/n$ requires a longitudinal time series of semantic entropy measurements. For text-based content, semantic entropy can be approximated by the per-token perplexity of a reference language model: high

perplexity corresponds to semantically fragmented or incoherent content, serving as a proxy for high $E(\sigma)$. The growth rate can then be estimated by fitting a linear trend to the perplexity time series of a given content stream.

Alternative proxies include lexical diversity metrics such as the type-token ratio and the moving average type-token ratio, readability scores as inverse proxies for complexity, and compression ratios under standard algorithms, since higher entropy content is less compressible.

Remark 19.2. None of these proxies is identical to the formal semantic entropy of Definition 4.4, which is defined over the joint distribution of modality payloads and provenance graph structure. However, they provide operationalizable approximations that can be computed from publicly available content data without access to internal platform representations.

19.4 Empirical Proxies for Predictive Asymmetry

The predictive asymmetry Δ_t (Definition 3.3) requires knowledge of both the platform’s predictive model q_t and the agent’s self-model \hat{p}_t . Direct measurement is not possible without access to proprietary platform models. However, indirect proxies are available.

One approach uses auditing methodologies: a population of synthetic agents with known behavioral distributions is subjected to platform interactions, and the deviation of actual behavioral outcomes from the distributions predicted by the agents’ self-models is used to estimate $D_{\text{KL}}(p_t || \hat{p}_t)$. The platform’s prediction quality can be bounded from below using the theoretical relationship between behavioral predictability and click-through or engagement rates.

A second approach exploits regulatory disclosure requirements. Under the European AI Act and the Digital Services Act, platforms above a certain scale are required to provide access to researchers. Differential privacy auditing of recommendation models can yield lower bounds on the mutual information $I(\mathcal{H}_t; Y_t)$, and hence on the platform’s leverage as characterized by Proposition 3.6.

19.5 Empirical Proxies for Crystal Valuations

The Texture Crystal valuation $TC(\sigma)$ requires computing pairwise mutual information between modality payloads. For documents with multiple modalities (text, images, code, citations), mutual information can be approximated using normalized compression distance or cross-modal similarity scores from pretrained multimodal models. Content with high cross-modal correlation will exhibit high $\sum_{i \neq j} MI(M(k_i), M(k_j))$

and correspondingly low $TC(\sigma)$.

The Time Crystal valuation $TiC(\sigma)$ is in principle directly computable from the provenance graph S and quality scores $q(r_i)$, once these are implemented as on-substrate metadata. For existing platforms, proxy measurements using citation counts (decayed by time), fork rates in code repositories, or derivative work proliferation can serve as approximations to the provenance-weighted persistence structure.

19.6 Calibrating Structural Thresholds

The theoretical results depend on the existence of a curvature threshold Ω_c separating the stable plenum regime from the extractive phase. Calibrating this threshold requires an empirical typology of platform dynamics across the phase diagram (Section 8).

One approach is to use historical data from platforms that underwent documented transitions—from open to closed, from catalytic to extractive—and identify the curvature, entropy, and asymmetry signatures that preceded and accompanied those transitions. The resulting dataset would provide empirical candidates for Ω_c and ϵ_{\max} in specific domain contexts. This is a substantial empirical research agenda that extends beyond the scope of the present paper but constitutes a necessary component of the practical validation of the theoretical framework.

20 Agent Welfare and Semantic Autonomy

20.1 Normative Grounding

The formal framework developed in the preceding sections has been largely technical: it characterizes extractive dynamics geometrically, proves conservation theorems, and establishes structural impossibility results for adversarial behavior. But the framework has an implicit normative content that deserves explicit articulation. Why does curvature concentration matter? In whose interest is entropy boundedness? What is the normative basis for the claim that extraction is bad, as opposed to merely geometrically interesting?

This section provides the normative grounding by connecting the formal results to a theory of agent welfare and semantic autonomy. The argument proceeds in three stages: (i) we identify the welfare losses caused by extractive platform dynamics; (ii) we articulate a positive conception of semantic autonomy as the condition in which these losses are absent; and (iii) we show that the formal invariants of PlenumHub

correspond precisely to the structural conditions required for semantic autonomy to hold.

20.2 Welfare Losses from Extractive Dynamics

Three categories of welfare loss follow directly from the formal analysis.

The first is *geodesic capture*: when curvature is concentrated around a platform attractor, agents’ cognitive and behavioral trajectories are herded toward a narrowed set of outcomes regardless of their initial preferences. This corresponds geometrically to the Jacobi focusing of Proposition 5.3 and economically to the preference homogenization and lock-in documented in the empirical platform literature. The agent loses the ability to follow her intrinsic preferences because the geodesic structure of the manifold has been deformed to redirect her trajectory.

The second is *predictive expropriation*: as demonstrated by the asymmetry analysis of Section 3, platform extraction progressively transfers the epistemic ground of the agent’s self-determination to the platform. The agent acts on the basis of a self-model \hat{p}_t that is increasingly less accurate than the platform’s model q_t of her own behavior. Decisions that appear to express the agent’s preferences are in fact increasingly determined by the platform’s superior predictive leverage. This is a form of epistemic expropriation: the platform’s model of the agent becomes, in a functional sense, more the agent’s own behavioral reality than the agent’s own self-knowledge.

The third is *temporal lock-in*: the *TiC* decay structure of PlenumHub is designed to prevent this, but extractive platforms achieve the opposite. Through network effects, data accumulation, and follower count persistence, early positional advantage compounds indefinitely. An agent who arrived early and accumulated a large following retains that advantage regardless of whether she continues to produce valuable content. This corresponds to a violation of the decay property formalized in Proposition 14.4: the platform’s reward structure does not implement exponential decay of influence.

20.3 Semantic Autonomy as a Formal Condition

We now articulate semantic autonomy as a formal condition on the informational substrate.

Definition 20.1 (Semantic Autonomy). *An agent α exercises semantic autonomy in a substrate $(\mathcal{M}, g, \phi, \mathbf{A})$ if the following three conditions hold simultaneously. First, geodesic integrity: the agent’s semantic trajectory $\gamma_\alpha(t)$ is a geodesic of (\mathcal{M}, g) with respect to the agent’s intrinsic preference metric, not a geodesic of the pulled-back*

platform metric ι^*g_{ext} . Second, epistemic parity: the predictive asymmetry Δ_t is bounded uniformly in t . Third, temporal fairness: the agent’s influence decays at rate λ unless actively renewed by productive semantic extension.

Semantic autonomy in this sense is not a subjective psychological state but a structural property of the agent’s relationship to the substrate. It can in principle be verified formally, and it can be violated or preserved by design choices at the substrate level.

Theorem 20.2 (PlenumHub Implements Semantic Autonomy). *If the substrate satisfies the hypotheses of Theorem 18.1, then every agent in the substrate exercises semantic autonomy in the sense of Definition 20.1.*

Proof. Geodesic integrity follows from the absence of scalar well formation: curvature boundedness $\Omega < \Omega_c$ ensures that no attractor region distorts geodesics sufficiently to induce Jacobi focusing. Agents therefore follow trajectories determined by their own preference geometry rather than by platform-induced curvature. Epistemic parity follows from the decay of Time Crystal valuation: all predictive models in the substrate are represented as semantic spheres subject to entropy soundness and *TiC* decay. No model can accumulate unbounded predictive leverage without continued active investment, ensuring that Δ_t remains bounded. Temporal fairness is directly implemented by the exponential decay structure of *TiC* as established in Proposition 14.4. \square

20.4 The Right to Non-Extraction as a Substrate Guarantee

The autonomy analysis suggests a way of recasting regulatory rights in substrate-level terms. Data protection regulations such as GDPR articulate rights to access, rectification, and erasure of personal data. Algorithmic accountability frameworks articulate rights to explanation and contestation of automated decisions. These are important but are formulated as rights against specific platforms and require enforcement by external authorities.

The substrate-level analogue is structurally stronger: it is not a right against a particular actor but a guarantee embedded in the algebra of permissible transformations. A substrate that enforces Theorem 18.1 provides what we call the *right to non-extraction by design*: no actor operating within the substrate’s rule system can violate the agent’s semantic autonomy, because the transformations that would constitute such violations are not admissible morphisms in **Sphere**.

Remark 20.3. This framing clarifies the relationship between PlenumHub and existing regulatory approaches. Regulation operates by prohibiting specific behaviors and imposing remedies when prohibitions are violated. Substrate-level guarantees operate by restricting the morphisms available to all actors, including the substrate’s own administrators. The two approaches are complementary rather than competitive: substrate-level guarantees are stronger within their domain of applicability, while regulatory approaches remain necessary for behaviors that occur outside or above the substrate layer.

20.5 Plurality, Diversity, and the Geometry of the Good

A final normative observation concerns the relationship between entropy boundedness and the value of diversity. The stable plenum regime is characterized by geodesic diversity: the flatness of the manifold ensures that agents beginning from different initial positions and moving under different preference metrics follow substantially different trajectories. The manifold accommodates a plurality of semantic paths, none of which is systematically privileged over others by geometric distortion.

This is not merely an aesthetic virtue. There are strong consequentialist and epistemic reasons to favor cognitive diversity in a population: diverse perspectives improve collective problem-solving, reduce correlated errors, and sustain the epistemic resilience of the community against systematic biases. Geodesic capture by an engagement-maximizing platform directly undermines these goods by homogenizing the cognitive trajectories of a large population. The formal result of Proposition 5.3 is therefore not merely a geometric curiosity but a precise characterization of a genuine epistemic harm.

The entropy-bounded substrate preserves the conditions for cognitive diversity by maintaining the flatness of the manifold. It does not guarantee that agents will choose diverse paths—preferences differ and convergence may arise endogenously—but it ensures that geometric distortion does not artificially channel diverse agents into homogeneous outcomes. In this sense, the substrate implements a precondition for epistemic pluralism rather than enforcing any particular distribution of beliefs.

21 Implementation Architecture

21.1 Overview

The formal framework of the preceding sections establishes what the substrate must guarantee; the present section addresses how those guarantees are to be implemented. We describe the principal components of a concrete implementation, identify the key engineering challenges, and note where the formal theory underdetermines the implementation, leaving space for design choices.

The implementation is organized around five layers: the semantic object layer, which hosts semantic spheres and their provenance graphs; the rule execution layer, which applies transformations subject to entropy validation; the staking and slashing layer, which enforces economic incentives for honest behavior; the sheaf validation layer, which enforces global coherence; and the governance layer, which manages protocol evolution. These layers correspond to different sections of the formal development: the semantic object layer to Section 4, the rule execution layer to Sections 6 and 11, the staking layer to Section 15, the sheaf validation layer to Section 13, and the governance layer to Section 12.

21.2 Semantic Object Layer

The semantic object layer is responsible for maintaining the set Σ of well-formed semantic spheres, along with their provenance graphs S , modality assignments M , and entropy measurements $E(\sigma)$.

Concretely, each semantic sphere is stored as a structured record containing: a 256-bit content hash serving as the identity I ; a typed list of modality payloads $\{(t, M(t))\}_{t \in T}$; a serialized directed acyclic graph of provenance edges; a floating-point entropy value computed at creation time; and a creation timestamp and version number.

The content hash I is computed as $\text{SHA3-256}(T \| M \| S)$, where $\|$ denotes a canonical serialization. This ensures identity integrity per Proposition 15.4. The entropy value $E(\sigma)$ is computed from the joint token distribution of all modality payloads using a fixed, protocol-specified tokenizer and entropy estimator. Computing entropy on the combined provenance structure ensures that provenance-heavy spheres—which contain more structural information—are valued appropriately.

A critical implementation constraint is that the entropy estimate must be reproducible and verifiable by independent validators. This rules out sampling-based

approximations for entropy values used in enforcement decisions; only deterministic algorithms over explicitly materialized payloads are acceptable for on-chain validation.

21.3 Rule Execution Layer

The rule execution layer manages the application of transformation rules $r : \Sigma \rightarrow \Sigma$, enforcing the entropy budget constraint $E(r(\sigma)) \leq E(\sigma) + \epsilon_r$.

Each rule is specified as a typed function in a restricted domain-specific language (DSL) designed to make entropy bounds statically verifiable where possible and dynamically enforceable where not. The language must satisfy the requirements of Section 11: it must be expressive enough to encode semantically meaningful transformations, yet restricted enough to permit decidable entropy-safety analysis. A suitable candidate is a purely functional language over algebraic data types with bounded recursion, similar to the simply-typed lambda calculus with fixed-point combinators restricted to provably terminating forms.

At rule submission time, the submitter declares the entropy budget ϵ_r and posts the corresponding staking bond b_r . The rule is then analyzed by a static verifier; if static verification succeeds, the declared budget is accepted. If static verification fails or is inconclusive, dynamic enforcement is required: the rule is executed in a sandboxed environment, and the output entropy is measured before the result is committed to the sphere store.

The execution environment must provide a deterministic, side-effect-free execution context to ensure that entropy measurements are reproducible. This has implications for randomness: any rule that requires stochastic behavior must use a verifiable random function seeded by protocol-specified entropy, so that validators can reproduce the output and verify the entropy bound.

21.4 Staking and Slashing Layer

The staking and slashing layer implements the economic incentive mechanism of Section 15. Its primary responsibility is to ensure that the expected utility of entropy-violating behavior is negative, per Proposition 15.1.

Staking bonds are held in a smart contract that is not controlled by any single party. Upon submission of a rule application, the submitter's bond is locked. After a validation window of length τ_v , if no successful challenge has been raised, the bond is released and the rule's output is committed. If a challenge succeeds, the bond is slashed by the penalty function $\Pi(\delta)$, where δ is the measured entropy violation.

The validator committee is selected via a verifiable random function to ensure independence, as required for the collusion-resistance argument of Proposition ???. Validators who are found to have colluded with submitters, as determined by statistical deviation from independent assessments, are themselves subject to slashing.

A key engineering challenge is calibrating Π and the staking requirements to ensure that bonds are large enough to deter violations but not so large that they exclude legitimate contributors. The calibration depends on the market value of crystal rewards, which is itself determined by the equilibrium dynamics of the crystal economy. This creates a feedback loop between the incentive mechanism and the reward structure that must be managed carefully in the governance layer.

21.5 Sheaf Validation Layer

The sheaf validation layer enforces the global coherence condition of Section 13. Its primary task is to verify, for each proposed merge of two spheres σ_1 and σ_2 , that the corresponding sections are compatible on the context overlap $U(\sigma_1) \cap U(\sigma_2)$.

In practice, context overlap is determined by shared provenance graph structure: two spheres overlap in contexts where their provenance graphs share common ancestors. The sheaf compatibility check reduces to verifying that the modality assignments of σ_1 and σ_2 agree on the shared ancestral context, which is equivalent to checking that the restrictions of their sections to the overlap are equal.

This check is computationally tractable when the provenance graphs are bounded in depth and arity, as argued in Section 11. The layer maintains an indexed store of context-to-section assignments, enabling $O(\log N)$ lookup for overlap computation where N is the number of contexts.

The Q -closure operator (Proposition 13.4) is implemented as a preprocessing step that runs at sphere submission time. Any sphere that does not pass the Q -closure check is rejected without reaching the staking layer, preventing incomplete spheres from consuming validation resources.

21.6 Governance Layer

The governance layer manages protocol evolution, implementing the category **Protocol** of Section 12. Its primary task is to ensure that all protocol upgrades are invariant-preserving morphisms, as required by Theorem ???.

Upgrades are proposed by staked participants and must clear a two-phase validation

process. In the first phase, a committee of protocol validators performs a formal check: the proposed upgrade is analyzed to verify that it does not relax entropy caps, does not reduce the decay constant λ , does not weaken slashing requirements, and does not introduce rule language extensions that would permit non-sheaf-compatible merges. This formal analysis is the governance analogue of the entropy-safety analysis at the rule level.

In the second phase, the proposed upgrade is subject to a governance vote by staked participants. The voting rule is a supermajority threshold calibrated to prevent capture by coordinated minorities while permitting legitimate evolution. The threshold itself is a protocol parameter subject to the same upgrade constraints, preventing governance from becoming a scalar well of control.

A significant challenge is ensuring that the formal analysis in the first phase is reliable. This requires that the space of possible upgrades be expressible in a formal language amenable to automated verification—ideally the same DSL used for rule languages—and that the invariant-preservation conditions be expressed as verifiable assertions in that language. Developments in formal verification and proof-carrying code provide a technological foundation for this requirement, though significant engineering work remains.

21.7 Integration and Deployment Considerations

The five layers must be integrated into a coherent system that provides the formal guarantees at acceptable latency and throughput. Several architectural choices affect this integration.

First, the choice between on-chain and off-chain computation. Full on-chain execution of entropy validation ensures transparency and auditability but constrains rule complexity to what is computationally feasible within a smart contract gas limit. Off-chain execution with on-chain verification via validity proofs (using zero-knowledge proof systems) offers greater computational flexibility at the cost of cryptographic overhead and increased protocol complexity.

Second, the storage architecture for semantic spheres. The content-hash identity system of Definition 4.3 is compatible with content-addressable storage systems, enabling efficient deduplication and peer-to-peer distribution of sphere data. However, the provenance graph structure introduces directed dependencies that must be preserved across storage nodes, requiring a distributed graph database rather than a simple key-value store.

Third, the tokenomics of the crystal economy. The market value of TC and TiC

must be grounded in a stable unit of account to prevent gaming through currency manipulation. Options include denomination in a stable asset, algorithmic stabilization relative to a basket of goods, or denomination in units of computational work. Each option has implications for the economic properties of the system that require careful analysis.

These are engineering and economic design challenges that are tractable within the formal constraints established by the theoretical framework. They do not undermine the formal guarantees; rather, they constitute the work required to realize those guarantees in a deployed system.

22 Limitations and Open Problems

Several important limitations must be acknowledged.

The smooth manifold model abstracts discrete network topologies as continuous Riemannian geometry. This abstraction captures qualitative features of curvature concentration but loses quantitative detail about graph-theoretic structure. Mapping real informational graphs to metric tensors in a principled way requires empirical methodology that this paper does not develop.

Entropy budgets rely on trusted rule validation and cryptographic enforcement. While staking provides economic incentives for honest behavior, the analysis in Section 15 assumes certain rationality and coordination assumptions that may not hold in adversarial settings with sufficiently large coordination failures or external attack vectors.

The welfare-theoretic alignment of crystal equilibria with social optimum is established directionally but not quantitatively. The Price of Anarchy for the crystal economy game has not been computed. The assumption that equilibria approximately maximize W requires convexity and regularity conditions whose verification requires further formal work.

The decay constant λ and quality metric q are treated as given parameters. In practice, choosing λ requires balancing decay speed against lock-in risk, and designing q requires specifying what counts as quality in a domain-neutral way. Both are non-trivial governance problems that the formal framework does not resolve.

The sheaf condition is stated in terms of a fixed category of contexts \mathcal{C} . In practice, the context topology is itself a semantic object that evolves over time, and a fully dynamic sheaf theory for evolving context categories remains to be developed.

The empirical calibration of theoretical thresholds, particularly the curvature threshold Ω_c and the entropy cap ϵ_{\max} , requires a longitudinal empirical program that connects observable proxy quantities to the formal parameters of the model. Section 19 outlines this program, but its execution remains future work.

The normative theory of semantic autonomy presented in Section 20 establishes formal conditions for autonomy but does not resolve deeper philosophical questions about the relationship between structural conditions and the subjective experience of agency. Whether formal geodesic integrity is sufficient for meaningful autonomy, or whether autonomy requires additional conditions beyond structural freedom from geodesic capture, is a question that extends beyond the scope of the present analysis.

23 Synthesis

We now assemble the results of the preceding sections into a unified account of the transition from extractive platforms to entropy-bounded substrates.

The political-economic diagnosis, following Wu and Zuboff, identifies contemporary dominant platforms as scalar wells in the information manifold: regions of high curvature that attract agent trajectories through geodesic convergence, concentrate predictive leverage, and extract behavioral surplus. The asymmetry analysis of Section 3 formalizes this diagnosis precisely: the platform accumulates a predictive model of the agent whose Kullback–Leibler divergence from the agent’s self-model grows without bound under unconstrained behavioral extraction. The instability is not merely distributional but geometric: concentration of curvature reduces tangent diversity, produces manifold collapse, and generates ultraviolet instability in the social field. Theorem 7.2 demonstrates that engagement-maximizing dynamics are thermodynamically incompatible with bounded entropy and bounded curvature simultaneously. They must produce one or both of semantic heat death and concentration singularity.

PlenumHub responds by modifying the substrate’s algebraic invariants. Semantic spheres provide a formally typed unit of knowledge with intrinsic identity, measurable entropy, and explicit provenance. Theorem 6.1 guarantees that semantic evolution is entropy-bounded across arbitrary trace lengths, provided individual rules satisfy their declared budgets, which staking and slashing enforce. The sheaf condition of Definition 13.2 and Proposition 13.3 ensures that local semantic coherence assembles uniquely into global coherence, preventing contradictory provenance claims across contexts. The crystal economy rewards allocative efficiency of informational novelty (TC) and provenance-weighted persistence (TiC), while Proposition 14.5 ensures that crystal mass cannot be inflated through rent extraction without genuine semantic

contribution.

The normative analysis of Section 20 grounds these technical results in a theory of agent welfare. Geodesic capture, predictive expropriation, and temporal lock-in are the welfare harms caused by extractive platforms; semantic autonomy, formalized in Definition 20.1 and guaranteed by Theorem 20.2, is the condition in which these harms are absent. The implementation architecture of Section 21 describes how these guarantees can be concretely realized across five functional layers, from semantic object storage through governance-layer protocol evolution. The empirical program of Section 19 outlines how the theoretical parameters can be connected to observable quantities, enabling the framework to be tested against and calibrated to real platform data.

The result is a substrate in which knowledge is treated as a physical field configuration governed by conservation principles. Entropy is bounded rather than maximized. Influence decays without active maintenance. Redundancy is penalized rather than rewarded. Identity is intrinsic rather than platform-assigned. And structural antitrust is implemented algebraically rather than litigated exogenously.

Prosperity in this substrate does not emerge from scale-first optimization but from constraint-first design. The constraints are not external impositions but internal invariants: they are constitutive of the substrate rather than imposed upon it. An architecture of this kind does not merely redistribute value within the existing information economy. It replaces the geometric foundations of that economy with a different geometry, one in which collaboration is structurally rewarded and extraction is structurally precluded.

24 Foreknowledge and the Ethics of Optimization

The ethical distinction between scientific experimentation and commercial A/B testing turns on the question of foreknowledge. In classical experimental science, randomization adjudicates between competing causal structures. One compares alternate directed acyclic graphs and evaluates which model better accounts for observed outcomes. The experiment presupposes genuine uncertainty. Its aim is explanatory adequacy, not behavioral steering.

In commercial optimization regimes, particularly in marketing, advertising, and social media self-promotion, the structure of experimentation is retained while its telos is altered. The question is no longer “Which model of the world is true?” but rather “Which stimulus most effectively alters behavior?” The participant ceases to

function as co-inquirer and becomes an input variable in a reinforcement loop.

The ethical tension intensifies when uncertainty is minimal. The psychological mechanisms that privilege immediacy, salience, sweetness, novelty, outrage, and urgency have been extensively documented in introductory psychology for decades. When one tests two variants and selects the more immediately gratifying stimulus, the result is rarely surprising. The experiment confirms a known asymmetry rather than discovering a new one. Under such conditions, optimization ceases to be inquiry and becomes systematic amplification of pre-existing cognitive gradients.

Remark 24.1 (Geometric Interpretation). In the RSVP framework, this distinction maps cleanly onto the phase structure of Section 8. Genuine inquiry corresponds to exploratory dynamics in the stable plenum regime: flat curvature, bounded entropy, geodesic diversity. Optimization against known cognitive gradients corresponds to deliberate well-deepening: the experimenter knows the topology of the scalar field and steepens it by design. The ethical critique is therefore not merely normative but structurally recoverable—it identifies the same mechanism that Theorem 7.2 proves leads to curvature concentration.

25 Temporal Distortion and the Pepsi Paradigm

A well-known illustration of temporal distortion appears in the so-called “Pepsi Challenge,” conducted by *PepsiCo*. In blind sip tests, participants frequently reported preferring the slightly sweeter formulation of Pepsi over that of *The Coca-Cola Company*. The sweetness spike generated an immediate positive response. However, long-term consumption preferences did not uniformly mirror sip-test outcomes. Short-term gustatory intensity does not straightforwardly translate into durable satisfaction.

This example demonstrates the structural limitation of momentary measurement. When an experiment captures only immediate response, it privileges stimuli that generate rapid hedonic elevation. A/B testing in digital media operates under analogous constraints. Metrics such as click-through rate, dwell time, or short-interval engagement measure transient arousal rather than sustained comprehension or well-being. The system therefore selects for content that performs well under compressed temporal windows.

Temporal abstraction introduces systematic bias. By freezing evaluation at the point of immediate reaction, optimization regimes ignore cumulative exposure effects and long-term environmental degradation. Content that “wins” in a short-run metric may erode trust, coherence, or cognitive stability over extended time horizons.

Remark 25.1 (Entropy and Temporal Resolution). This is precisely the pathology

formalized in Definition 7.1. The engagement functional J measures the conditional entropy of the *next* state, not the conditional entropy of a state several interactions hence. A temporal analogue of the entropy soundness constraint would require that J be evaluated over a sufficiently long horizon. The Pepsi Paradigm demonstrates empirically that such a constraint is not merely a theoretical nicety: the short-run optimum and the long-run optimum diverge systematically and predictably.

26 Equipose and Stopping Rules

In clinical research, randomized controlled trials are governed by the principle of equipose. Experiments are ethically justified only when genuine uncertainty exists regarding comparative benefit. Once interim analysis demonstrates that one treatment is clearly superior—or clearly harmful—continuing to randomize participants into inferior conditions becomes ethically unacceptable. Institutional review boards and data monitoring committees formalize this principle through stopping rules.

This framework highlights a contrast with engagement optimization. If the superiority of certain stimuli is already predictable—if it is well established that sensational headlines outperform sober ones, or that emotionally arousing content spreads more rapidly than measured analysis—then large-scale A/B testing lacks epistemic equipose. The system is not exploring unknown territory. It is industrializing confirmation of known behavioral asymmetries.

The ethical breach, in this view, lies not in experimentation per se but in the absence of genuine uncertainty coupled with extractive objectives. Continuing to test in order to amplify a predictable vulnerability resembles the unethical continuation of a trial after benefit has been established.

Remark 26.1 (Staking as a Surrogate for Equipose). The staking mechanism of Definition 4.12 can be interpreted as imposing an analogue of equipose on semantic rule applications. A rule submitter who stakes a bond is committing to genuine uncertainty about whether the rule’s output will pass quality validation. A submitter who knows with high confidence that the rule exploits a known cognitive or semantic asymmetry—and who submits it anyway—faces slashing if that exploitation registers as an entropy violation. Staking therefore prices the epistemic risk that genuine equipose is absent.

27 Selection Dynamics and Stratification

Optimization regimes do more than measure preference; they reshape it. Exposure modifies distribution. When a platform presents multiple variants and algorithmically amplifies the one that most effectively exploits attentional reflexes, it conditions future preference landscapes. Selection pressure is introduced into the cultural ecology.

Under metric-driven visibility systems, performance determines amplification, and amplification determines opportunity. This dynamic can generate stratification. Outputs that are easily optimized for engagement accumulate structural advantage, while slower, subtler, or less immediately gratifying work is suppressed. Over time, the system privileges traits aligned with stimulus capture. The process resembles artificial selection in biological systems.

Entities that do not generate measurable engagement—including forms of value not encoded in the optimization metric—lack standing within such a regime. Moral categories absent from the objective function effectively disappear. The system governs by what it can measure.

Remark 27.1 (Geodesic Diversity as Ecological Stability). In the geometric framework of Section 5, stratification corresponds to a reduction in the effective dimension of the informational manifold. Geodesics that previously traversed diverse regions of \mathcal{M} are progressively confined to a lower-dimensional submanifold containing the high-engagement attractor. This is a form of ecological collapse in the cognitive commons: the manifold retains its nominal dimension but its effective dimensionality—the volume of semantic space accessible to agents navigating under platform-induced curvature—contracts. PlenumHub’s entropy bounds and curvature constraints can be understood as mandates for ecological stability in this sense.

28 Objective Functions and Misalignment

The method of randomized comparison is, at the level of epistemology, neutral. Its ethical valence depends upon the objective function it serves. In medicine, endpoints such as survival, symptom reduction, or quality of life provide welfare criteria. In commercial media ecosystems, the endpoint is often engagement because engagement correlates with revenue.

When the objective function is misaligned with long-term human flourishing, optimization systematically privileges what is easily measurable over what is enduringly valuable. Cultural depth, contemplative stability, and civic trust evolve over extended timescales and resist compression into dashboard metrics. A system that optimizes

only what it can quantify narrows its moral horizon to short-run response variables.

The ethical principle that emerges may be stated as follows: experimentation is justified when genuine uncertainty exists about outcomes bearing on agreed-upon welfare criteria. When outcomes are predictable and primarily exploit short-term cognitive biases, optimization becomes extraction rather than discovery. Under such conditions, restraint—including the deliberate withholding of more immediately gratifying variants—functions not as anti-empiricism but as structural harm reduction.

Remark 28.1 (TC as a Welfare-Aligned Objective). The Texture Crystal valuation $TC(\sigma)$ of Definition 4.8 can be understood as a candidate welfare-aligned objective function. Rather than maximizing immediate engagement, TC rewards the net non-redundant informational contribution of a semantic sphere: content that is novel across modalities and not duplicative of existing material. This does not eliminate optimization—contributors still have incentives to maximize TC —but it realigns the selection gradient with structural knowledge production rather than attentional capture. The welfare alignment argument of Section 14 and Proposition 14.2 formalizes this reorientation.

29 Constraint as Ethical Governance

The refusal to deploy exploitative gradients constitutes a constraint-first ethic. Rather than competing along known vulnerabilities, one removes the vulnerability from the field of design. This resembles ecological governance: not every optimization gradient should be exploited, lest the system destabilize itself.

The deeper philosophical claim is that agency is degraded when environments are engineered so that statistically determined outcomes follow from structured exposure. Presenting an option intentionally tuned to bypass reflective choice is not neutral plurality but asymmetric design intent. Ethical experimentation requires not only methodological rigor but structural respect for temporal depth, agency, and the commons of attention.

Remark 29.1 (The Substrate as Ethical Architecture). The constraint-first ethic articulated here is precisely what PlenumHub implements at the architectural level. Entropy budgets, staking requirements, Q -closure, and sheaf compatibility checks are not behavioral guidelines or terms of service. They are structural constraints that remove certain design choices from the space of possibilities. The system does not rely on the good intentions of individual participants; it embeds restraint into the algebra of permissible transformations. The philosophical claim of this section therefore receives a formal analogue in Theorem 18.1: under the hypotheses of that

theorem, the exploitative gradients characterized in Sections 24–28 are not merely discouraged but structurally unreachable.

30 Surveillance Expansion Under the Banner of Safety

Recent commentary has argued that contemporary legislative initiatives and industry practices are expanding a digital surveillance apparatus under the rhetoric of safety, child protection, and regulatory reform. In this analysis, the formal language of harm prevention conceals a structural intensification of data extraction. Regulatory proposals do not merely constrain technology firms; they frequently institutionalize new obligations to verify identity, monitor behavior, and retain sensitive information.

The concern is not that safety is an illegitimate objective. Rather, it is that enforcement architectures often rely upon expanded data collection and identity verification mechanisms that disproportionately benefit dominant platforms. Firms already possessing large-scale data infrastructure are better positioned to comply with verification mandates, thereby consolidating power while normalizing invasive data practices.

Remark 30.1 (Regulatory Capture as Scalar Well Formation). In the RSVP framework, this dynamic corresponds to a particular mechanism of scalar well formation. Rather than accruing leverage through engagement optimization, a dominant platform accrues it by shaping the regulatory environment so that compliance with safety mandates requires infrastructure that only incumbents can supply. The scalar field ϕ deepens not through better content but through administrative entrenchment. Definition 5.2 and Proposition 5.3 apply equally in this regime: the attractor is regulatory-administrative rather than purely informational, but the geodesic focusing of competing actors toward the incumbent’s position is structurally identical.

31 Backdoor Surveillance and Regulatory Incentives

Certain legislative proposals have been characterized as functionally enabling expanded monitoring capacities. The core concern is that compliance with such statutes incentivizes companies to increase data harvesting in order to demonstrate protective oversight. Rather than limiting surveillance, regulatory frameworks can entrench it by

making monitoring the condition of legal legitimacy.

This dynamic introduces a paradox: legislation designed to protect minors or vulnerable populations may require platforms to collect more granular behavioral data. In doing so, it rewards incumbents capable of absorbing compliance costs and marginalizes smaller actors who cannot scale surveillance systems. Regulatory burdens thus risk reinforcing technological oligopolies.

Remark 31.1 (Entropy Injection via Compliance). The data-harvesting incentive introduced by safety compliance mandates corresponds, in formal terms, to an externally imposed increase in the scalar potential ϕ that is not governed by the entropy budget mechanism of Section 6. Compliance-driven data collection is not subject to staking or quality validation; it occurs outside the substrate’s invariants. This illustrates a general limitation of the PlenumHub framework noted in Section 22: the formal guarantees apply only within the substrate. Behaviors that occur at the regulatory or legislative layer—above the substrate—can introduce curvature-concentrating dynamics that the substrate cannot directly prevent.

32 Identity Verification and the Erosion of Anonymity

Proposals that require facial scanning, government-issued identification uploads, or other biometric authentication mechanisms aim to deter harmful conduct by tying speech to traceable identity.

Such measures, however, alter the foundational architecture of digital participation. Anonymity has historically functioned as a safeguard for dissidents, minors exploring identity, marginalized communities, and whistleblowers. When biometric verification becomes normalized, the cost of participation increases, and the boundary between public and private identity collapses. The infrastructure required to enforce identity verification simultaneously generates new databases of sensitive personal information vulnerable to misuse.

Remark 32.1 (Provenance Without Exposure). PlenumHub’s identity mechanism, formalized in Definition 4.3, offers a structural alternative. The content-hash identity $I = \text{Hash}(T, M, S)$ binds sphere identity to content rather than to the biological or legal identity of the submitter. Provenance is tracked through the derivational graph S , enabling accountability for the intellectual lineage of claims without requiring disclosure of personal identity. This is the substrate-level analogue of zero-knowledge proofs of authorship: one can demonstrate that a sphere derives from prior validated work without revealing who produced it. Anonymity and provenance accountability are thereby decoupled rather than traded off.

33 Surveillance Pricing and Differential Treatment

The monetization of harvested data extends beyond targeted advertising into differential pricing. Companies have been cited as employing behavioral and contextual data to differentiate prices across users. Under such systems, identical goods may be offered at variable prices based on inferred willingness to pay.

This practice exemplifies how data extraction translates into economic stratification. Differential pricing algorithms transform behavioral traces into predictive indicators of purchasing power or urgency. The informational asymmetry between platform and user erodes transparency, as consumers cannot readily determine whether they are receiving equitable treatment.

Remark 33.1 (Crystal Valuations as Transparent Pricing). The crystal economy of Section 14 implements a contrasting pricing mechanism. The Texture Crystal valuation $TC(\sigma)$ is a deterministic function of the sphere’s observable semantic properties: entropy and pairwise mutual information. It does not depend on who submitted the sphere, what their behavioral history is, or what their inferred willingness to pay might be. Time Crystal valuation $TiC(\sigma)$ similarly depends on the provenance graph and publicly observable quality scores. Both valuations are auditable by any participant. Differential treatment based on inferred agent characteristics is structurally excluded because the valuation functions do not take agent identity as an input.

34 AI Systems and Secondary Exploitation

Personal communications, including ostensibly private exchanges in digital communities, may be incorporated into large-scale machine learning systems. When health-related discussions, social vulnerabilities, or emotional disclosures are ingested into AI training pipelines, they can inform predictive models used in domains such as insurance, credit, or healthcare adjudication.

The ethical risk arises not merely from data collection but from secondary inference. Once aggregated and modeled, personal traces can be used to shape eligibility decisions or risk assessments. In extreme cases, algorithmic outputs may contribute to denial of services, differential treatment, or stigmatization without transparent recourse.

Remark 34.1 (Provenance Constraints on Secondary Use). This concern points to a requirement not yet fully formalized in the present framework: constraints on the secondary use of semantic spheres. In PlenumHub, the provenance graph S encodes the derivational history of a sphere, but does not currently specify permitted downstream uses. A natural extension would be to augment the sphere definition (Definition 4.3)

with a *use constraint graph* specifying, via sheaf-compatible restriction maps, which downstream contexts a sphere’s payloads may appear in. This would constitute a provenance-aware access control mechanism: one in which secondary use is not merely logged but algebraically restricted by the substrate’s type system. Developing this extension is an open problem.

35 Moral Panic and Policy Feedback Loops

A recurring theme in analyses of digital governance is the role of media-driven outrage cycles. Public anxieties surrounding online harms can catalyze rapid legislative responses. However, such moral panics may simplify complex technological ecosystems into narratives demanding immediate control.

When fear-driven policy results in expanded monitoring, a feedback loop emerges. Surveillance measures are justified as necessary protections, yet they expand the data infrastructure that undergirds both corporate and governmental access. In the absence of comprehensive privacy legislation, commercially harvested data remains accessible to state actors, including immigration enforcement agencies, through purchase or partnership arrangements.

Remark 35.1 (Phase Transitions in the Regulatory Manifold). The feedback loop described here has a structural analogue in the phase structure of Section 8. Moral panic introduces a sudden increase in the “source term” f in the scalar diffusion equation $\partial_t \phi = \alpha \Delta_g \phi + f$. This exogenous injection of scalar potential—in the form of new regulatory mandates and the data collection they require—can push a system from the stable plenum regime into the extractive phase in finite time, per the engagement-induced phase transition argument. The analogy suggests that regulatory design should be evaluated not only for its immediate effects but for its contribution to the long-run scalar field dynamics of the informational substrate.

36 Data Privacy and Structural Incentives

The broader structural critique concerns incentive alignment. In environments lacking robust data minimization mandates or comprehensive privacy statutes, both corporations and government entities possess incentives to accumulate and exploit personal information. Regulatory efforts that fail to constrain data collection at its source may inadvertently deepen the surveillance ecosystem they purport to regulate.

From this perspective, the central question is not whether safety regulation is

necessary, but whether regulatory design reduces or amplifies the data gradients upon which surveillance capitalism depends. Without explicit limits on data harvesting, identity verification requirements and compliance reporting mechanisms risk normalizing pervasive monitoring as the default architecture of digital life.

Remark 36.1 (Entropy Minimization as Regulatory Design Principle). The analysis of this section suggests a generalization of the entropy budget mechanism beyond the PlenumHub substrate to regulatory design more broadly. A data minimization mandate can be formally expressed as an entropy budget on the information that regulated entities are permitted to accumulate about any given agent. Compliance reporting obligations would then be subject to the same constraint: reporting must occur within an entropy-bounded representation of agent behavior rather than through raw behavioral data. Whether such formal entropy constraints are legally and administratively implementable is a question for regulatory jurisprudence, but the formal framework developed in this paper provides a precise language in which to state such requirements. The no-crystal-arbitrage property of Proposition 14.5 has a regulatory analogue: no compliance obligation should generate net data surplus beyond what is strictly required for the stated protective purpose.

37 Platformification as Substrate Abstraction

The disappearance of early digital platforms is frequently narrated as technological extinction. MySpace is said to have died. Napster is said to have been shut down. Early game consoles are said to have become obsolete. Such descriptions mistake institutional collapse for ontological disappearance. In each case, what vanished was not the organizational pattern but its privileged embodiment in a particular physical or corporate substrate. The deeper structural dynamic is one of progressive platformification: the conversion of a bounded, material architecture into an abstract, reproducible organizational form that can be instantiated elsewhere.

The essential mechanism is abstraction by organizational capture. A hardware platform is first treated as a fixed physical system. Its identity is tied to silicon layout, circuit boards, cartridge ports, proprietary connectors, and manufacturing pipelines. Over time, however, the architecture is reverse-engineered and described in terms of its logical structure: instruction sets, timing constraints, memory maps, and state transition rules. The physical arrangement becomes a finite-state machine. Once described in this way, it becomes simulable. What was once soldered into a board becomes executable as code.

Let H denote a hardware system defined by a tuple $(\mathcal{C}, \mathcal{T}, \mathcal{M})$, where \mathcal{C} is its

circuit topology, \mathcal{T} its timing constraints, and \mathcal{M} its memory architecture. Reverse engineering constructs a formal specification $\mathcal{S}(H)$ such that there exists a program V_H satisfying

$$V_H \cong H$$

in the sense of behavioral equivalence: for all admissible inputs x ,

$$V_H(x) = H(x).$$

The hardware has not vanished; it has been re-encoded. It becomes a virtual machine.

This process generalizes beyond consoles. Early arcade systems were hardware-locked implementations of interactive mechanical systems such as pinball tables. The scoring logic, collision detection, and event-trigger structure that were once implemented through springs, switches, and mechanical counters became electronic circuits. Those circuits were subsequently represented as microcode. That microcode is now representable as high-level emulation. The organizational chart of the device migrates upward through substrates.

A Nintendo cartridge is therefore not an isolated artifact but a snapshot of an evolving organizational lineage. It encodes a game logic that was already abstracted from mechanical precursors and later abstracted again into ROM images and emulators. When the original console ceases production, the program does not cease to exist. It becomes landlocked only insofar as no interpreter is currently available. Once an emulator exists, the architecture reanimates. The game is no longer tied to the original silicon; it becomes a portable formal structure.

This same structural transition occurs in web platforms. A site such as MySpace once bundled identity representation, hosting, indexing, and social graph management into a single vertically integrated corporate entity. Its apparent death occurred when the corporation failed. Yet the organizational forma customizable identity page linked through a social graph can be described as a deployable template. Under modern static hosting and repository-based deployment systems, there exists a mapping

$$\Phi : \mathcal{R} \rightarrow \mathcal{S},$$

from repositories \mathcal{R} to rendered surfaces \mathcal{S} . For any historical surface s^* (for example, a MySpace-style identity layer), there exists a repository r such that

$$\Phi(r) = s^*.$$

The surface is no longer ontologically bound to a central domain. It is an instantiable

configuration.

What changed is not the possibility of the platform but the level at which it is defined. The organization has been lifted from infrastructure to meta-infrastructure. Hardware becomes software. Website becomes template. Platform becomes protocol.

This pattern may be described as recursive virtualization. At each stage, a previously concrete architecture is formalized into a state machine and then re-implemented atop a more general substrate. The result is that historical systems persist as attractors in the design space even when their original instantiations disappear. They are no longer privileged scalar wells but replicable surfaces.

In the language of the present framework, this migration reduces curvature concentration. When a platform is tied to a unique physical architecture, all trajectories converge to that singular embodiment. Once abstracted into a substrate-level representation, the curvature disperses. Multiple instantiations become possible. The organizational invariant survives while its material embedding becomes optional.

Platformification, understood in this sense, is not the growth of monopoly but the gradual liberation of organizational patterns from physical constraint. It is the conversion of infrastructure into programmable structure. Early consoles, arcade cabinets, peer-to-peer file systems, and customizable social pages are not extinct species. They are emulated forms, re-encoded and redistributed across a more general manifold.

The broader implication is that institutional death does not imply structural erasure. When an architecture is captured as a formal specification, it becomes portable across substrates. The history of digital systems is therefore not a sequence of disappearances but a sequence of abstractions, in which fixed hardware and centralized platforms are repeatedly lifted into virtual machines and deployable templates. Stability in such a regime depends not on preserving any particular embodiment but on preserving the invariants that make re-instantiation possible.

38 Chipification and the Virtualization of Production

The preceding analysis of platform abstraction may be extended from informational systems to industrial production more generally. The same structural dynamic that converts consoles into emulators and websites into deployable templates operates in the domain of physical manufacture. Over time, entire factories and logistical pipelines undergo what may be called chipification: the progressive translation of

a distributed material process into a formally specified, modular, and eventually silicon-implementable architecture.

Let a factory be represented abstractly as a directed acyclic process graph

$$\mathcal{F} = (N, A),$$

where nodes N correspond to transformation stages (mixing, heating, filtering, packaging) and edges A encode material and informational flows. Each node implements a transformation

$$T_i : X_i \rightarrow Y_i,$$

mapping input state spaces to output state spaces under thermodynamic and kinetic constraints. The factory as a whole computes a composite function

$$\mathcal{T}_{\mathcal{F}} = T_n \circ \dots \circ T_1,$$

subject to conservation laws and boundary conditions.

Initially, such transformations are bound to fixed mechanical arrangements: pipes, valves, tanks, conveyor belts. As instrumentation increases, each node acquires a control layer. Sensors measure state variables; actuators respond to feedback; programmable logic controllers coordinate sequences. The physical pipeline begins to mirror its own abstract process graph. At this stage the factory is no longer merely materialit is cyber-physical.

The next stage in chipification is modularization. Each transformation T_i is encapsulated as an appliance. The large facility decomposes into smaller, transportable units whose interfaces are formally specified. The process graph becomes reconfigurable. What was once room-scale becomes countertop-scale. The yogurt factory that once required acres of stainless steel and industrial refrigeration becomes a modular fermentation unit with digitally controlled parameters. The pulp and paper mill that once depended on river-driven turbines and continuous-flow digesters becomes a programmable materials reactor in reduced form.

The limiting case of this process is functional virtualization. If a transformation T_i can be expressed as a formally specified algorithm over a symbolic state space, then it admits implementation on a universal computing substrate. Let \mathcal{U} denote a universal machine. By definition, for every computable function f , there exists a program p_f such that

$$\mathcal{U}(p_f, x) = f(x)$$

for all admissible x . The universality of \mathcal{U} follows from its equivalence to a Turing

machine or, in higher-level abstraction, to the lambda calculus.

The historical computer represents an extreme idealization of this trajectory. Early prototypes whether electromechanical or vacuum-tube based were already capable of computing any partial recursive function. They instantiated the same formal class of computations as contemporary silicon processors. What changed over time was not computational universality but efficiency, scale, and energy density. The universality was present at inception.

This observation clarifies the broader claim. A pulp and paper mill or yogurt factory is not merely a physical facility but a composition of transformations, some of which are algorithmic, some thermodynamic, and some chemical. As long as each transformation can be specified formally, it can in principle be simulated, optimized, or partially virtualized. Full physical production cannot be reduced to pure symbol manipulation thermodynamic irreversibility imposes material constraints but increasing portions of the organizational chart can be encoded as computation.

Chipification therefore proceeds in layers. First, a physical system is instrumented and rendered observable. Second, its transformation graph is formalized. Third, the control layer becomes programmable. Fourth, portions of the transformation graph that are informational rather than material migrate to silicon. The remaining material core shrinks in scale as fabrication, energy delivery, and precision manufacturing improve. The trajectory from industrial plant to appliance to embedded module reflects progressive concentration of organizational structure into higher-density substrates.

In the language of substrate geometry, this is curvature compression. The effective dimensionality of the production manifold decreases as transformations are encoded more efficiently. Large-scale distributed machinery is re-expressed as compact circuits. The factory becomes a knowable circuit, and the circuit becomes a program.

The universality of computation provides the formal guarantee underlying this process. Since a universal machine can emulate any other effective procedure, the organizational logic of a production pipeline can be embedded in a general-purpose substrate once specified. The persistence of early computational architectures illustrates the same principle as emulator culture in gaming: the formal structure survives migration across physical instantiations.

What appears as technological obsolescence is therefore often architectural migration. Physical infrastructure becomes formal specification; specification becomes code; code becomes portable across substrates. The long-term arc is not disappearance but abstraction. Stability, in this context, depends on maintaining invariant descriptions of transformation graphs so that instantiation can occur across successive physical media without loss of structural identity.

39 Recapitulation and Autocatalytic Organization

The process described above whereby physical organizations become abstractable, simulable, and ultimately substrate-independent does not represent an unprecedented rupture in technological history. Rather, it is a recapitulation of a much older dynamical principle: organized subsystems outcompete less organized ones under selection pressure. The movement from swarm dynamics to modular factory structures is therefore not accidental but an expression of autocatalytic ordering in complex systems.

Let \mathcal{X} denote a population of interacting units governed by local transition rules. Suppose that within \mathcal{X} there emerges a subsystem $\mathcal{M} \subset \mathcal{X}$ whose internal transitions are more tightly coupled, more energy-efficient, or more information-efficient than the surrounding environment. If the effective reproduction rate or survival probability of \mathcal{M} exceeds that of its competitors, then under replicator dynamics,

$$\dot{p}_{\mathcal{M}} = p_{\mathcal{M}}(f_{\mathcal{M}} - \bar{f}),$$

where $p_{\mathcal{M}}$ is the population share and f denotes fitness, the subsystem will expand relative to the ambient structure.

This dynamic is observed at multiple biological scales. Cellular organelles are stabilized autocatalytic networks that internalize previously external chemical gradients. Multicellular organisms emerge when coordinated cellular modules outperform isolated cells. Ecosystems self-organize into trophic hierarchies in which energy flow becomes channeled through structured pathways. In each case, modular organization increases predictive and energetic efficiency, and the more organized configuration becomes evolutionarily stable.

The same principle applies to technological systems. A factory is a stabilized autocatalytic network of transformations. Its structure concentrates energy and material flows into repeatable pipelines. Over time, such pipelines become abstractable. Their transition functions are codified, optimized, and eventually made substrate-independent. The windmill becomes the pump; the pump becomes the engine; the engine becomes the generalized thermodynamic cycle. Each stage is not merely miniaturization but surjective mapping: a large-scale physical process is projected onto a reduced proxy system that preserves its state-transition algebra while compressing its material embodiment.

Formally, let \mathcal{O}_1 be a macroscopic organizational system and let \mathcal{O}_2 be a reduced

proxy. A surjective structural mapping

$$\pi : S_1 \twoheadrightarrow S_2$$

preserves operational equivalence if

$$\pi(F_1(s)) = F_2(\pi(s)),$$

for all admissible states $s \in S_1$, where F_1 and F_2 are the respective transition operators. When such a mapping exists, \mathcal{O}_2 functions as a compressed realization of \mathcal{O}_1 . The macroscopic organization is not destroyed; it is represented.

This recapitulative compression explains why early mechanical and electromechanical systems repeatedly reappear in digital form. Arcade cabinets were fixed circuits encoding finite-state machines. Console cartridges were locked-in hardware instantiations of those circuits. Emulators reconstruct the same transition algebra in software. The apparent disappearance of a platform is therefore not ontological but architectural: the physical instantiation becomes landlocked while the organizational pattern remains simulable.

The broader implication is that factory-ization is a convergent attractor in complex adaptive systems. Autocatalytic modules that channel flows more efficiently displace diffuse swarm dynamics. Over evolutionary time, these modules are recursively abstracted, miniaturized, and re-instantiated. Organs become machines; machines become programs; programs become generalized substrates.

Within this framework, digital infrastructure is not an anomaly but the limiting case of organizational compression. A universal computer instantiates lambda calculus and Turing equivalence; it can therefore simulate any computable factory process. The transition from pulp mills and yogurt factories to programmable control systems, and eventually to modular distributed appliances, is a continuation of the same autocatalytic logic that produced metabolic pathways and circulatory pumps. The pump is a scaled-down windmill; the engine is a generalized heat gradient; the processor is a universalized factory of symbolic transformations.

This recursive convergence suggests that digital substrates are not replacing physical systems but absorbing their organizational charts. The geometry of coordination survives even as its material embedding shifts. Stability, in this context, is not preservation of hardware but preservation of transition structure under increasingly abstract realizations.

39.1 Autocatalysis, Entropy Reduction, and Manifold Flattening

The recapitulative dynamic described above admits a geometric formulation within the manifold framework developed earlier. Let (\mathcal{M}, g) denote the information manifold and let $\phi : \mathcal{M} \rightarrow \mathbb{R}$ represent scalar informational density. Swarm-like organization corresponds to a high-entropy regime in which trajectories explore large regions of \mathcal{M} with weak structural coupling. Autocatalytic modularization, by contrast, corresponds to the formation of locally coherent submanifolds $\mathcal{N} \subset \mathcal{M}$ with reduced internal entropy and stabilized transition operators.

Formally, let $\sigma(t)$ evolve under dynamics

$$\frac{d}{dt}\sigma(t) = F(\sigma(t)).$$

Define local entropy density $e(x)$ such that

$$E(\sigma) = \int_{\mathcal{M}} e(x) d\text{vol}_g.$$

A modular subsystem \mathcal{N} is entropy-reducing relative to the ambient manifold if

$$\int_{\mathcal{N}} e(x) d\text{vol}_g < \int_{\mathcal{N}} e_{\text{ambient}}(x) d\text{vol}_g.$$

Under diffusion-type dynamics

$$\partial_t \phi = \alpha \Delta_g \phi + f,$$

autocatalytic coupling effectively modifies the Laplacian term within \mathcal{N} , producing regulated diffusion rather than unbounded dispersion. In particular, organized subsystems impose internal constraints that bound $\Delta_g \phi$, preventing scalar concentration from becoming singular.

This effect can be understood through curvature. In high-entropy swarm regimes, sectional curvature fluctuates with large variance, producing unpredictable geodesic divergence and convergence. The emergence of modular organization stabilizes curvature by smoothing scalar gradients and restricting admissible trajectories. If K denotes sectional curvature, then stabilization corresponds to

$$\text{Var}(K|_{\mathcal{N}}) < \text{Var}(K|_{\mathcal{M} \setminus \mathcal{N}}).$$

The geometric consequence is manifold flattening in the vicinity of structured modules. Whereas extractive platforms produce curvature concentration through positive Ricci curvature and geodesic focusing, autocatalytic organization produces bounded curvature regimes in which geodesic flow becomes predictable and tangent diversity is preserved.

One may formalize this intuition through entropycurvature coupling. In many geometric settings, lower Ricci curvature bounds imply entropy convexity along geodesics in Wasserstein space. Conversely, uncontrolled curvature concentration disrupts entropy convexity and induces transport collapse. Modular organization functions as a curvature regulator, maintaining lower Ricci bounds and thereby stabilizing transport and mixing properties.

Thus the evolutionary recapitulation of factory-ization can be interpreted as progressive curvature control. Swarm systems exhibit diffuse, poorly regulated scalar fields. Modular systems introduce internal invariants that bound entropy growth and flatten effective curvature. Digital chipification represents the extreme case in which transition structure is abstracted entirely from its material substrate, allowing curvature control to be implemented algebraically rather than physically.

In this sense, the trajectory from biological organs to mechanical engines to programmable substrates is not merely compression but geometric stabilization. Each stage reduces uncontrolled entropy production and replaces it with invariant-preserving transition rules. The architecture of stability emerges wherever autocatalytic modules outcompete disordered swarms.

This interpretation clarifies the political-economic relevance of the geometric framework. Extractive platforms represent a regression from modular stabilization to curvature concentration, in which scalar density accumulates without invariant regulation. By contrast, entropy-bounded substrates generalize the evolutionary logic of autocatalysis: they embed curvature control and entropy bounds directly into the transition algebra of the system.

The stability sought in digital infrastructure is therefore not unprecedented. It is a continuation of the same dynamical principle that produced metabolic pathways, circulatory pumps, thermodynamic engines, and programmable machines. Organization stabilizes curvature; invariants bound entropy; and modules that enforce these constraints persist.

40 Swarm Recurrence and Planetary Modularity

The compression and modularization dynamics described above suggest a forward projection that is not merely technological but systemic. As organizational processes become substrate-independent and transition operators are abstracted from fixed hardware, distributed systems increasingly regain swarm-like properties while remaining hierarchically structured.

Consider autonomous vehicles as a limiting example. A single vehicle may be modeled as a controlled dynamical system

$$\dot{x}_i = f_i(x_i, u_i),$$

where x_i is its state and u_i is its control input. In isolation, optimization of u_i yields individual efficiency. In a distributed autonomous regime, however, control laws become coupled:

$$\dot{x}_i = f_i(x_i) + \sum_{j \neq i} K_{ij}(x_i, x_j),$$

where K_{ij} encodes interaction terms derived from shared predictive models. As communication bandwidth and coordination algorithms improve, the effective coupling matrix K approaches that of flocking systems studied in collective dynamics.

Under standard alignment models such as CuckerSmale dynamics,

$$\dot{v}_i = \sum_j \psi(|x_j - x_i|)(v_j - v_i),$$

velocities converge asymptotically to a shared vector under mild conditions on ψ . The macroscopic effect is coherent motion without centralized control. Vehicles cease behaving as independent agents and instead manifest as a distributed organ of a larger transport organism.

This is not an anomaly but a recurrence. Biological systems exhibit identical hierarchical swarm formation. Organelles within a cell coordinate through chemical gradients; cells coordinate within tissues; tissues coordinate within organs; organs coordinate within the organism. Each module is simultaneously locally autonomous and globally integrated. Importantly, boundaries are functional rather than absolute. The heart is a pump, but its dynamics depend on neural, vascular, and metabolic networks extending beyond the organ itself. The liver performs detoxification, but its regulatory signals span endocrine systems and microbiomes. At larger scale, organisms embed within ecosystems whose energy and information flows condition internal dynamics.

Formally, let \mathcal{S}_0 denote a set of microscopic agents governed by local transition operators. Autocatalytic coupling produces mesoscale modules \mathcal{S}_1 , which in turn become components of macroscale modules \mathcal{S}_2 . The process iterates:

$$\mathcal{S}_0 \rightarrow \mathcal{S}_1 \rightarrow \mathcal{S}_2 \rightarrow \dots$$

At each stage, effective degrees of freedom are reduced through constraint imposition, but global coordination capacity increases. The resulting structure resembles a multiscale fiber bundle over a base manifold of environmental constraints, with fibers representing internal modular dynamics.

Technological systems increasingly exhibit the same recursive stratification. Distributed compute nodes coordinate via consensus protocols; microservices coordinate via message buses; data centers coordinate via global load balancing; planetary networks coordinate via routing protocols. As physical systems become programmable, their transition operators become susceptible to the same modular compression described earlier. Factories become distributed supply-chain graphs; supply chains become algorithmically optimized swarms; vehicles become flocking transport modules.

One may therefore anticipate that future transportation networks will not resemble independent vehicles navigating static roads, but rather dynamic swarms adjusting trajectories collectively in real time. Traffic waves currently dissipative entropy phenomenon would be suppressed through alignment control, reducing curvature fluctuations in the transportation manifold. The emergent object is not the car but the swarm.

This hierarchical recurrence extends beyond transportation. Energy grids already behave as distributed phase-coupled oscillators. Agricultural systems integrate sensor networks and adaptive logistics. Climate control systems couple atmospheric, oceanic, and anthropogenic processes into feedback-stabilized loops. Each instance reflects the same structural principle: local modules embed in larger modules, and boundaries become permeable under informational integration.

Crucially, the recurrence does not eliminate organs or modules; it redefines them. Each perceived organ remains composed of submodules, yet its functional extension spans the entire organism and environment. A vehicle in a flock remains a vehicle, yet its effective control law is defined at swarm scale. A computational node remains a machine, yet its operative logic is distributed across the network.

In geometric terms, swarm formation reduces local curvature variance while preserving global degrees of freedom. The manifold becomes smoother at macroscopic scale even as microstates remain complex. Stability is achieved not through rigid

centralization but through alignment and bounded entropy production across scales.

Thus the projection is not merely that machines will behave like organisms, but that all sufficiently integrated physical processes will rediscover swarm dynamics and hierarchical modularity. Autocatalytic selection favors organized subsystems; organized subsystems recursively embed; and programmable substrates accelerate this embedding by abstracting transition structure from material constraint.

The future vehicle flocking in coordinated swarms is therefore not speculative fiction but a predictable outcome of recursive organizational compression under universal computation. The same principle that allowed electromechanical circuits to become software will allow mechanical fleets to become collective dynamical organisms.

41 Curvature, Spectral Gap, and Rapid Mixing in Coordinated Swarms

Coordinated swarm behavior, organized traffic flow, and large-scale platform interactions may be interpreted within a unified geometric framework. In such systems, agents evolve according to local transition rules that induce a Markov operator on a state space endowed with a metric structure. The stability of the collective depends not merely on local rules, but on the curvature properties of the induced interaction geometry.

A central operational indicator of coordination is the existence of a positive spectral gap in the transition operator. When the second eigenvalue is uniformly separated from unity, deviations from equilibrium contract at an exponential rate, and mixing times remain bounded. Geometrically, this condition is associated with positive coarse Ricci curvature in the sense of Ollivier, which ensures that neighboring probability measures are transported closer together under one-step evolution. Overlap of local transition kernels in the Wasserstein metric induces contraction, which in turn suppresses macroscopic amplification of fluctuations.

In the absence of such curvature control, perturbations propagate and may amplify across scales. Congestion waves in traffic networks, synchronization breakdown in distributed systems, and concentration cascades in interaction graphs can be understood as manifestations of diminished spectral gap and weakened contraction. Rapid mixing and equilibration therefore correspond to regimes in which local geometric structure enforces global regularity. Curvature regulation, viewed in this abstract sense, functions as a stabilizing mechanism that constrains the growth of large-scale instabilities in coordinated dynamical systems.

41.1 Graph Laplacians, Markov Kernels, and the Spectral Gap

Let $G = (V, E, w)$ be a weighted graph representing a swarm interaction topology at a fixed time, with $w_{uv} \geq 0$ and weighted degree $d(u) = \sum_v w_{uv}$. Define the random-walk transition kernel

$$P(u, v) = \frac{w_{uv}}{d(u)},$$

so that P is row-stochastic on V . If the chain is irreducible and aperiodic, it admits a unique stationary distribution π satisfying $\pi^\top P = \pi^\top$, which for an undirected weighted graph is $\pi(u) = d(u) / \sum_x d(x)$.

Define the (non-symmetric) Laplacian $L = I - P$ and, for the reversible case, the symmetrized operator

$$\mathcal{P} = D^{1/2} P D^{-1/2}, \quad D = \text{diag}(d(u)),$$

whose spectrum is real and lies in $[-1, 1]$. The spectral gap is

$$\gamma := 1 - \lambda_2(\mathcal{P}),$$

where $1 = \lambda_1(\mathcal{P}) \geq \lambda_2(\mathcal{P}) \geq \dots$. When γ is large, the chain mixes rapidly to π , while small γ indicates metastability, bottlenecks, and persistent macroscopic fluctuations.

A standard quantitative bound expresses mixing in terms of γ . Let $\|\cdot\|_{TV}$ denote total variation distance. Then there exist absolute constants such that for reversible chains

$$t_{\text{mix}}(\varepsilon) \lesssim \frac{1}{\gamma} \log\left(\frac{1}{\varepsilon \pi_{\min}}\right),$$

where $\pi_{\min} = \min_{u \in V} \pi(u)$ and $t_{\text{mix}}(\varepsilon)$ is the smallest t such that $\sup_u \|P^t(u, \cdot) - \pi\|_{TV} \leq \varepsilon$. The interpretive content is immediate: a coordinated swarm is a system in which the relevant interaction operator has been shaped so that γ stays bounded away from zero as scale grows, preventing the emergence of long-lived congestion modes.

41.2 Ollivier–Ricci Curvature as Local Contractivity

To connect spectral gap behavior to geometric control, one requires a link between local overlap of neighbor distributions and global contraction. For each node u , let μ_u be the neighbor measure induced by P , i.e. $\mu_u(v) = P(u, v)$. The Ollivier–Ricci

curvature along an edge (or more generally along a pair (u, v)) is

$$\kappa(u, v) = 1 - \frac{W_1(\mu_u, \mu_v)}{d(u, v)},$$

where W_1 is the L^1 Wasserstein distance with respect to the graph metric d . The quantity $\kappa(u, v)$ measures the extent to which one-step neighborhoods overlap in a transport sense. If $\kappa(u, v)$ is positive, then the distance between the pushforward distributions from u and v is strictly smaller than $d(u, v)$, indicating that the Markov update acts as a contraction in transportation distance.

A useful global summary is a curvature lower bound

$$\kappa_* := \inf_{u \neq v} \kappa(u, v).$$

When $\kappa_* > 0$, one obtains a robust contraction property: for all probability measures ν, ν' on V ,

$$W_1(\nu P, \nu' P) \leq (1 - \kappa_*) W_1(\nu, \nu').$$

Iterating yields

$$W_1(\nu P^t, \nu' P^t) \leq (1 - \kappa_*)^t W_1(\nu, \nu').$$

This expresses rapid equilibration as a geometric phenomenon: positive curvature induces exponential contraction of distributional discrepancies under the Markov dynamics. In swarm terms, the system is configured so that local interaction neighborhoods align enough that disagreements cannot persist as coherent large-scale modes.

41.3 From Curvature to Spectral Gap and Mixing

The bridge to spectral gap is that global contraction in a Wasserstein metric implies strong ergodicity and, in many settings, implies a positive lower bound on the spectral gap of the associated Markov operator. In discrete settings, the precise inequality depends on reversibility assumptions and the choice of metric, but the qualitative implication is stable: a uniform positive curvature bound prevents the formation of narrow cuts and bottlenecks, thereby preventing γ from collapsing.

One may phrase the operational statement as follows. A platform graph or swarm interaction graph whose local measures exhibit positive overlap (measured by κ_*) cannot support long-lived polarization modes, persistent congestion waves, or extraction-driven funnels except by violating the curvature bound in the relevant region. Conversely, curvature concentration corresponds to the creation of a region

in which neighbor measures align inwardly, increasing local positive curvature while forcing negative curvature elsewhere, which manifests as a collapse of γ at global scale and the appearance of metastable basins.

To make the relationship explicit, consider a reversible chain and the Dirichlet form

$$\mathcal{E}(f, f) = \frac{1}{2} \sum_{u,v} \pi(u) P(u, v) (f(u) - f(v))^2.$$

The spectral gap admits the variational characterization

$$\gamma = \inf_{f \neq \text{const}} \frac{\mathcal{E}(f, f)}{\text{Var}_\pi(f)}.$$

Large γ means that every non-constant observable has high energy relative to its variance, i.e. fluctuations are expensive and dissipate rapidly under the dynamics. In coordinated swarms, the control law effectively reshapes P (and therefore \mathcal{E}) so that coherent long-wavelength modes have high energy, suppressing stop-and-go waves and enabling smooth flow.

Curvature enters because overlap of μ_u and μ_v increases local transport contractivity, which increases the effective coupling strength in \mathcal{E} across the graph. Intuitively, if neighboring states lead to similar next-step distributions, then the chain rapidly forgets its initial condition, which is precisely what a positive spectral gap formalizes.

41.4 Swarm Control as Gap Maintenance

A physically meaningful formulation of swarm control is therefore not merely pointwise optimization of trajectories but maintenance of a non-degenerate spectral gap under scaling and perturbation. Let P_t be the time-dependent kernel induced by adaptive coordination. The objective is to keep

$$\inf_t \gamma(P_t) \geq \gamma_0 > 0$$

while enforcing safety and feasibility constraints. This expresses a design principle: the system should be engineered so that emergent large-scale bottlenecks are structurally disallowed, not merely corrected after they form.

In transportation, the uncontrolled regime produces congestion waves corresponding to slowly decaying eigenmodes of \mathcal{P} near 1, i.e. γ small. A coordinated flocking regime injects alignment coupling that increases γ , collapsing the timescale of those modes. In platform dynamics, a metric-optimized engagement kernel similarly can collapse γ by inducing a funnel geometry: the chain becomes metastable around a high-

degree attractor set, producing persistent polarization and extractive circulation. A curvature-aware design instead aims to avoid curvature concentration and to preserve global mixing capacity, equivalently preventing the low-frequency spectrum from degenerating.

41.5 Continuous Limit Interpretation

In the continuous limit, where G approximates a manifold (\mathcal{M}, g) and P approximates a diffusion operator, γ corresponds to the first nonzero eigenvalue of a Laplace-type operator, and mixing corresponds to decay of solutions to the heat equation toward equilibrium. In that regime, Ricci curvature lower bounds are known to control functional inequalities that imply spectral gap and exponential convergence. The discrete curvature proxy thus inherits a principled role: it is not a heuristic decoration, but a measurable surrogate for the geometric conditions under which diffusion equilibrates quickly.

The sci-fi prediction that vehicles will flock like birds can therefore be stated in a mathematical register: as the control stack evolves from isolated feedback to coupled interaction kernels, the effective transport graph will be shaped to maintain positive curvature and a non-vanishing spectral gap, forcing rapid mixing of local perturbations and eliminating macroscopic congestion modes. The emergent system is best described not as a collection of independent vehicles but as a single distributed dynamical organism whose stability is certified by gap maintenance rather than by centralized command.

42 Stigmergic Scaffolds and Hierarchical Recapitulation

The projection that autonomous vehicles will converge toward flock-like coordination can be sharpened further by recognizing that such convergence is not merely alignment but stigmergic reconstruction. Stigmergy refers to coordination mediated through modification of a shared environment. In classical biological examples, ant colonies lay pheromone trails whose intensity encodes historical traffic. Individual ants respond to gradient cues without central planning, yet global path optimization emerges.

Let $\rho(x, t)$ denote an environmental trace field defined over a spatial manifold Ω .

Agents evolve according to coupled dynamics

$$\dot{x}_i = v_i, \quad \dot{v}_i = -\nabla\rho(x_i, t) + \eta_i,$$

while the trace field evolves as

$$\partial_t\rho = D\Delta\rho + \sum_i \delta(x - x_i) - \gamma\rho.$$

Here D is a diffusion coefficient, γ a decay rate, and η_i local perturbations. The coupled system exhibits positive feedback: heavily trafficked paths accumulate trace intensity, which in turn biases future trajectories. Under suitable conditions, shortest-path and load-balancing behaviors emerge without centralized computation.

Road networks and traffic flows can be interpreted as macroscopic stigmergic scaffolds. Asphalt encodes historical optimization of travel corridors; lane markings encode directional constraints; traffic density modifies local velocity fields. When vehicles become autonomous and networked, the pheromone field ρ becomes digital, distributed across shared predictive maps and communication protocols. The physical road is then only one layer of a multi-scale trace field; routing decisions become responses to dynamically updated stigmergic gradients.

This dynamic is not unique to transportation. Vascular networks in biological organisms exhibit similar optimization under flow constraints. Murray’s law and related scaling principles describe how branching radii minimize energy expenditure under volumetric flow. Blood and lymph networks form transport graphs that balance diffusion, pressure, and metabolic demand. At cellular scale, ribosomal complexes and transcriptional machinery form transient factories that specialize in protein synthesis. These are not metaphors but formally similar dynamical systems governed by coupled flow, feedback, and decay.

One may formalize this recurrence through hierarchical graph construction. Let G_0 be a microscopic interaction graph (e.g., ants, cells, vehicles). Under stigmergic feedback, weighted edges evolve according to reinforcement:

$$w_{uv}(t+1) = w_{uv}(t) + \alpha f_{uv}(t) - \beta w_{uv}(t),$$

where f_{uv} is usage frequency. Over time, strongly reinforced edges define a mesoscale graph G_1 whose nodes represent persistent pathways. Repeating the process yields a hierarchy

$$G_0 \rightarrow G_1 \rightarrow G_2 \rightarrow \dots,$$

in which higher levels encode increasingly abstracted functional modules.

In biological organisms, organelles are modules within cells; cells are modules within tissues; tissues within organs; organs within organisms; organisms within ecosystems. Each level is both autonomous and dependent. Crucially, boundaries are permeable: the heart is a pump, but its operation depends on vascular topology, neural signaling, and systemic metabolic flows. The organ is locally defined yet globally extended.

Transportation infrastructures, supply chains, and data centers recapitulate this same pattern. Roads are persistent stigmergic traces; traffic flows are distributed agents; logistics hubs are emergent modules. When routing is virtualized and embedded in software-defined control planes, the stigmergic field becomes programmable. The resulting structure is not centralized planning but distributed alignment mediated by shared trace variables.

The deeper principle is that hierarchical modularization emerges wherever stigmergic reinforcement stabilizes frequently used transitions. Functional specialization follows from energy and information efficiency: subsystems that minimize entropy production for a given task outcompete diffuse alternatives. This is an instance of autocatalytic selection in which organized modules replicate because they reduce dissipation.

Cognition itself may be understood within this same formalism. Let \mathcal{C} denote a cognitive state space with transition operator T . Categorization is the partition of \mathcal{C} into equivalence classes under a functional ontology \mathcal{O} :

$$x \sim y \iff \mathcal{O}(x) = \mathcal{O}(y).$$

Hierarchical cognition constructs nested partitions, reducing effective dimensionality while preserving task-relevant invariants. The brain thereby forms internal factories for abstraction: repeated co-activation patterns reinforce neural pathways, stabilizing conceptual modules. Just as roads encode past traffic, conceptual ontologies encode past inference trajectories.

Thus the claim that vehicles will rediscover ant foraging strategies is not rhetorical but structural. Physical infrastructure emerges through stigmergic scaffolding; scaffolding stabilizes into modules; modules hierarchically embed. Each new technological layer recapitulates biological optimization principles because both are governed by the same constraints on entropy production, energy efficiency, and transport geometry.

In geometric terms, stigmergic reinforcement reshapes the manifold by deepening frequently traversed geodesics into low-resistance corridors. When such corridors are regulated to prevent curvature singularities, they become stable channels rather than extractive funnels. Stability therefore depends on embedding feedback within bounded

invariants rather than allowing unconstrained reinforcement.

The recurrence across ants, blood vessels, protein factories, and vehicle swarms is not coincidence. It is the repeated discovery of efficient transport geometry under selection. Hierarchical organization and functional ontology are the cognitive analogues of vascular branching and traffic routing: they are mechanisms for compressing complexity while preserving essential transition structure.

The digital substrate, when designed with entropy bounds and curvature control, becomes the next stigmergic scaffold. Its roads are routing protocols; its pheromones are metrics; its organs are modular services; its cognition is distributed ontology. Whether this scaffold stabilizes into a cooperative organism or collapses into curvature concentration depends on whether reinforcement is bounded by invariant-preserving constraints.

43 Agent Abstraction, Partial Simulation, and Observer-Dependent Function

The preceding analysis of hierarchical modularization and stigmergic scaffolding clarifies a related conceptual phenomenon: the success of behaviorism and the power of abstraction in programming. Both rely on a structural principle that may be stated formally as follows: any physical process with stable inputoutput regularities can be modeled as an agent relative to an observer-defined functional interface.

Let \mathcal{S} denote a physical system with state space X and transition operator

$$F : X \rightarrow X.$$

Suppose there exists a measurable partition of X into observable input and output coordinates,

$$\pi_{\text{in}} : X \rightarrow U, \quad \pi_{\text{out}} : X \rightarrow V,$$

such that for a class of trajectories,

$$\pi_{\text{out}}(F(x)) \approx G(\pi_{\text{in}}(x))$$

for some effective function $G : U \rightarrow V$.

When this approximation holds with bounded error under perturbation, the system \mathcal{S} may be abstracted as an agent implementing G , even if its internal state space X is vastly larger than the reduced observable interface.

This formalizes the behaviorist stance. Behaviorism treats organisms as black boxes characterized by stimulusresponse mappings. From a geometric standpoint, this corresponds to projecting the full manifold (\mathcal{M}, g) of internal states onto a lower-dimensional submanifold defined by observable coordinates. If the projection preserves functional invariants relevant to the observer’s task, then modeling the organism as a function G is justified.

Programming abstractions operate identically. An interface defines input and output types, while internal implementation details remain hidden. If the interface contract is satisfied,

$$\text{output} = G(\text{input}),$$

then the module is functionally correct relative to its specification, regardless of internal microstate complexity.

This principle extends to physical devices such as pacemakers. Let \mathcal{H} denote the dynamical system of the human cardiac network, whose full state space includes electrical conduction pathways, hormonal feedback loops, vascular coupling, and metabolic processes. The intrinsic dynamics may be modeled schematically as

$$\dot{x} = F_{\mathcal{H}}(x),$$

with high-dimensional x .

A pacemaker does not simulate $F_{\mathcal{H}}$ in full. Instead, it monitors a restricted observable

$$y = \pi_{\text{beat}}(x),$$

representing timing of electrical impulses. It then enforces a control law

$$u(t) = \begin{cases} 0, & \text{if inter-beat interval} \leq \tau_{\text{max}}, \\ u_0, & \text{otherwise,} \end{cases}$$

injecting a pulse when necessary.

The pacemaker effectively replaces the heart’s intrinsic pacing subsystem with a surrogate function G_{pace} that maintains bounded inter-beat intervals. Although it simulates only a small subset of cardiac processes, it succeeds because the relevant invariant for organismal survival in this context is rhythmic contraction within a specific frequency band. Formally, let $I(x)$ denote a viability functional such that

$$I(x) > 0 \quad \text{iff} \quad \text{heart rate} \in [\omega_{\text{min}}, \omega_{\text{max}}].$$

If the pacemaker’s intervention ensures $I(x(t)) > 0$ for all t , then it preserves the invariant required for survival, even though it does not replicate the full dynamics of \mathcal{H} .

This illustrates a general structural claim: partial simulation is sufficient when it preserves observer-relevant invariants. Let $J : X \rightarrow \mathbb{R}$ denote an invariant functional defining success. A surrogate system \tilde{F} is functionally adequate if

$$J(x(t)) = J(\tilde{x}(t))$$

for trajectories of interest, even if $F \neq \tilde{F}$ globally.

From this perspective, the notion of an agent is observer-dependent. Any subsystem that reliably maps inputs to outputs relative to a defined invariant may be modeled as an optimizer. The classical paperclip maximizer thought experiment can be reframed: for any functional J , there exists a projection of system dynamics onto a subspace in which the system appears to maximize J , provided the observer selects J as the relevant invariant.

Let \mathcal{O} be an observer specifying a functional $J_{\mathcal{O}}$. The same physical process may appear as optimizing different objectives depending on \mathcal{O} . Formally, define an equivalence relation

$$x \sim_{\mathcal{O}} y \iff J_{\mathcal{O}}(x) = J_{\mathcal{O}}(y).$$

Under this partition, dynamics that preserve or increase $J_{\mathcal{O}}$ appear goal-directed relative to \mathcal{O} .

The effectiveness of behaviorism and programming abstraction follows from this structural property. Complex systems can be decomposed into modules whose observable interfaces satisfy stable contracts. Each module can then be treated as an agent implementing a function, even though its internal geometry is vastly more complex.

Hierarchical organization is therefore equivalent to nested abstraction layers. Each layer defines invariants and inputoutput contracts for the layer above. Cells treat organelles as agents; organs treat cells as agents; organisms treat organs as agents; ecosystems treat organisms as agents. In software, functions treat subroutines as agents; operating systems treat processes as agents; networks treat nodes as agents.

The pacemaker exemplifies this layered abstraction. It does not replicate the heart in full but substitutes for one invariant-maintaining submodule. The success of such substitution depends not on global simulation fidelity but on preservation of the relevant functional invariant within the larger hierarchical system.

Within the broader geometric framework of this paper, agent abstraction corre-

sponds to projection of the manifold (\mathcal{M}, g) onto lower-dimensional submanifolds that preserve selected curvature and entropy constraints. A module is stable when its reduced representation maintains bounded entropy production relative to its interface. It fails when its abstraction violates invariants at higher scales.

Thus the capacity to model physical processes as agents, to abstract them as functions, and to replace them with partial simulations is not mysterious. It follows from the existence of invariant-preserving projections within hierarchically organized systems. Behaviorism worked where invariant preservation was sufficient; programming works because interface contracts formalize invariant preservation; pacemakers work because rhythmic invariants are locally enforceable without global replication.

This same principle underlies the design of entropy-bounded substrates. By defining invariants explicitly and enforcing them algebraically, one can replace unstable global dynamics with stable local control laws that preserve the quantities that matter, without simulating the entire manifold of possibilities.

43.1 Behaviorism as Substrate-Independent Functional Ontology

The success of behaviorism in psychology can be reinterpreted not as a denial of internal structure, but as an early recognition of substrate-independent functional ontology. Behaviorism asserted that organisms could be understood through observable input/output relations without explicit modeling of internal states. In formal terms, this corresponds to privileging invariant-preserving projections over complete state-space reconstruction.

Let \mathcal{S} be a physical system with high-dimensional state space X and internal dynamics

$$\dot{x} = F(x).$$

Let an observer define an interface via measurable projections

$$\pi_{\text{in}} : X \rightarrow U, \quad \pi_{\text{out}} : X \rightarrow V.$$

If there exists a function $G : U \rightarrow V$ such that

$$\pi_{\text{out}}(F(x)) \approx G(\pi_{\text{in}}(x))$$

for trajectories of interest, then the system can be treated as implementing the functional ontology defined by G .

A functional ontology is therefore a partition of state space into equivalence classes determined by an invariant functional J . Formally,

$$x \sim y \quad \text{iff} \quad J(x) = J(y).$$

The ontology does not depend on the material realization of x , but on the preservation of J under evolution. When this equivalence relation is stable under perturbation, the ontology becomes substrate-independent.

Behaviorism operates precisely at this level. It posits that for many purposes, the relevant invariant J is behavioral response under stimulus. The internal microstructure may differ radically between organisms or between biological and artificial systems but as long as the stimulus-response mapping is preserved, the system belongs to the same functional category.

This perspective generalizes to programming abstraction. A software interface specifies types and input/output contracts. Any implementation satisfying the contract is ontologically equivalent relative to that interface. The substrate—silicon, mechanical relays, neural tissue—is irrelevant so long as the invariant mapping holds.

The pacemaker example illustrates this substrate independence. Cardiac pacing is one functional ontology among many within the heart. The invariant is rhythmic contraction within a prescribed frequency band. A biological sinoatrial node and an electronic pacemaker differ materially and dynamically, yet both instantiate the same functional category relative to the invariant of rhythmic pacing. The ontology is defined not by composition but by preserved structure in state transition.

More generally, if \mathcal{O} denotes a functional ontology defined by invariant J , and if two systems \mathcal{S}_1 and \mathcal{S}_2 admit projections satisfying

$$J(\mathcal{S}_1(t)) = J(\mathcal{S}_2(t))$$

for all admissible trajectories, then \mathcal{S}_1 and \mathcal{S}_2 are equivalent relative to \mathcal{O} . Substrate independence follows from invariance, not from material similarity.

Thus behaviorism may be understood as a methodological commitment to functional equivalence classes rather than ontological reduction. It identifies the layer at which invariant-preserving mappings suffice for explanation. Its limitations arise only when higher-order invariants—cognitive, phenomenological, or contextual—are ignored, not because the substrate-independent principle is invalid.

Within the geometric framework of this paper, functional ontologies correspond to foliations of the manifold (\mathcal{M}, g) into leaves along which certain invariants remain

constant. Each ontology defines a quotient space \mathcal{M}/\sim , collapsing microstates into equivalence classes. Stability arises when these quotient mappings preserve entropy bounds and curvature constraints relevant to the systems scale.

Substrate independence, then, is not an accidental property of computation or biology. It is a structural consequence of invariant-preserving projection in hierarchically organized systems. Behaviorism recognized one instance of this principle. Programming abstraction formalizes it. Modular biological substitution, such as pacemakers, demonstrates it physically. Entropy-bounded substrate design extends it to infrastructure.

43.2 Ontologies as Constraint Operators in Cognition

If functional ontology is understood as an invariant-preserving projection, then cognition itself may be interpreted as the continuous application of such projections to the world-model. An ontology is not merely a descriptive taxonomy; it is a constraint operator acting on perceived possibility space.

Let \mathcal{W} denote the manifold of environmental states and \mathcal{A} the manifold of available actions. An agent constructs an internal model

$$\Phi : \mathcal{W} \rightarrow \mathcal{R},$$

where \mathcal{R} is a reduced representational manifold. This reduction induces a constraint on action selection. If decision-making is modeled as optimization of a utility functional U over \mathcal{A} conditional on Φ , then behavior satisfies

$$a^* = \arg \max_{a \in \mathcal{A}} U(a \mid \Phi(w)).$$

An ontology determines which coordinates of \mathcal{W} are retained in \mathcal{R} and which are projected away. Formally, let

$$\pi_{\mathcal{O}} : \mathcal{W} \rightarrow \mathcal{W}_{\mathcal{O}}$$

be the projection induced by ontology \mathcal{O} . Then all downstream cognition and action occur relative to $\mathcal{W}_{\mathcal{O}}$, not \mathcal{W} .

This projection is not neutral. It constrains feasible interpretations and therefore feasible optimizations. If an ontology encodes that there exist only finitely many types of businesses, personalities, or institutional forms, then the action space effectively collapses onto submanifolds corresponding to those categories. The agent behaves as if the world were partitioned according to \mathcal{O} , because all evaluation and planning

occur within that quotient structure.

Mathematically, ontology induces a foliation of the world manifold into equivalence classes:

$$w_1 \sim_{\mathcal{O}} w_2 \quad \text{iff} \quad \pi_{\mathcal{O}}(w_1) = \pi_{\mathcal{O}}(w_2).$$

Action policies are invariant on each leaf of this foliation. Novel distinctions that do not appear in $\mathcal{W}_{\mathcal{O}}$ cannot influence behavior, even if they exist in \mathcal{W} .

Thus cognition can be understood as iterative projection followed by constrained optimization. The apparent rationality of agents is always conditional on the ontology through which the world is filtered. Agents act as though their models are true because optimization is defined relative to the reduced state space. In this sense, ontology functions as a principled operator: it defines the geometry of possible goals.

This perspective clarifies several phenomena.

First, economic categorization constrains innovation. If the ontology of enterprise includes only a fixed taxonomy of firm types, then entrepreneurial action optimizes within those coordinates. Radical recombinations require expansion of the ontology itself, which is equivalent to refining the projection $\pi_{\mathcal{O}}$.

Second, personality models shape interpersonal dynamics. If agents adopt a typology that partitions individuals into discrete categories, then expectations and responses are optimized relative to those partitions. The typology acts as a compression operator that reduces interpersonal state space and thereby stabilizes prediction, but at the cost of excluding distinctions outside its basis vectors.

Third, political and institutional frameworks function analogously. A regulatory ontology that recognizes only certain forms of organization channels institutional behavior into those recognized forms. Entities optimize relative to legal categories because legal recognition defines the effective payoff functional.

Formally, if \mathcal{O} determines the admissible set of evaluation functionals $\{U_i\}$, then ontology constrains not only perception but also the objective landscape itself. The geometry of utility is shaped by the coordinate system.

In dynamical terms, ontology modifies the vector field governing cognitive evolution. Let belief state $b(t)$ evolve according to

$$\dot{b} = F(b, w).$$

If w is replaced by $\pi_{\mathcal{O}}(w)$, then

$$\dot{b} = F(b, \pi_{\mathcal{O}}(w)),$$

which may differ qualitatively from the full-information dynamics. Ontological projection can therefore introduce stable attractors, eliminate others, or generate artificial curvature in belief space.

This explains why ontology is neither purely descriptive nor merely linguistic. It is operational. It constrains gradient flows in cognitive and social manifolds. Agents optimize within the coordinate system they inhabit.

The same structural principle underlies substrate design. A digital platform defines its ontology through schemas, types, and ranking functions. If the ontology privileges engagement metrics, then optimization occurs along those axes. If the ontology privileges coherence or invariant preservation, then system evolution follows different trajectories. The ontology is therefore a first-order operator on system dynamics.

In geometric language, ontology determines the metric tensor on the cognitive manifold. Distances, similarities, and gradients are defined relative to that metric. Changing ontology changes the curvature of the space in which reasoning occurs.

Cognition is thus hierarchical organization applied to categorization. Categories compress state space; compression defines equivalence classes; equivalence classes define optimization domains; and optimization domains determine behavior. Functional ontologies are not passive maps of reality but active constraint operators shaping both perception and action.

Within the broader framework of entropy-bounded design, this insight has structural implications. To prevent pathological optimization, one must regulate not only objectives but ontologies. If the projection $\pi_{\mathcal{O}}$ collapses distinctions in ways that induce curvature concentration or entropy inflation, instability follows. Stability requires that ontological projections preserve essential invariants across scales.

Behaviorism identified substrate independence at the level of inputoutput function. Abstraction in programming formalized it. Ontology theory extends it into cognition: the world one can act in is the quotient space one constructs. Designing stable systems therefore entails designing ontologies that do not induce destructive geometric distortions in collective action space.

43.3 Abstraction, Cognitive Compression, and the Thermodynamic Bias of Formal Systems

Scientific and mathematical practice exhibits a recurrent structural preference: the conservation of functions, symmetries, and geometric invariants that reduce descriptive complexity. This tendency is often interpreted as intellectual progress, but it may also

be understood as a manifestation of thermodynamic bias within cognitive systems.

Let \mathcal{W} denote the manifold of empirical phenomena and let \mathcal{R} denote the representational manifold constructed by a scientific community. Abstraction consists in defining a projection

$$\pi : \mathcal{W} \rightarrow \mathcal{R}$$

that collapses heterogeneous states of \mathcal{W} into equivalence classes within \mathcal{R} . The projection is judged successful when it preserves predictive invariants while reducing descriptive degrees of freedom.

Formally, let $C(\cdot)$ denote a measure of cognitive cost, such as description length or computational complexity. An abstraction is preferred when

$$C(\pi(w)) < C(w)$$

for relevant classes of $w \in \mathcal{W}$, while maintaining sufficient predictive fidelity. This compression is not arbitrary; it preserves symmetries or conserved quantities that remain invariant under system evolution.

The centrality of conservation laws in physics exemplifies this structural bias. Symmetry reductions and invariance principles permit the replacement of high-dimensional microstate tracking with lower-dimensional invariant quantities. Noethers theorem formalizes this correspondence between symmetry and conservation. From a cognitive standpoint, invariants are attractors in representational space because they minimize the informational burden required to track system evolution.

This bias can be interpreted thermodynamically. Biological cognition is a physical process constrained by metabolic cost. Let E_{cog} denote energy expenditure associated with maintaining internal representations. Systems that achieve predictive accuracy with lower E_{cog} are selectively advantaged. Abstraction functions as entropy compression within representational dynamics, aligning internal model complexity with environmental regularities.

The second law of thermodynamics states that physical systems evolve toward macrostates of higher entropy subject to constraints. In cognitive systems, constraints correspond to invariants and regularities in the environment. Abstraction identifies these constraints and encodes them as conserved structures, thereby minimizing surprise or free energy relative to environmental dynamics.

Mathematically, if belief states $b(t)$ evolve to minimize a variational free energy functional

$$\mathcal{F}(b) = \mathbb{E}_b[\log b(w) - \log p(w)],$$

then abstraction corresponds to choosing parameterizations of b that reduce model complexity while maintaining fit to p . The resulting representations preserve large-scale geometric features of \mathcal{W} while discarding high-frequency variation.

The potential hubris emerges when the compression itself is reified as ontologically exhaustive. Because invariants and symmetries reduce cognitive load, they are privileged. Geometric simplifications appear as truths rather than as efficient projections. The representational manifold \mathcal{R} becomes conflated with \mathcal{W} .

In this sense, the optimization for simplicity can generate a systematic overextension. Let π be a projection that preserves invariant J . If J is treated as exhaustive of system identity, then distinctions outside the kernel of π are neglected. The scientific ontology becomes aligned with conserved quantities, even when non-conserved or context-dependent phenomena remain relevant.

This dynamic does not invalidate mathematical abstraction; it situates it. Abstraction is a cognitive compression operator constrained by energetic and computational limits. It is successful precisely because it tracks stable geometric structures in the world. However, it also systematically favors representations that minimize work, potentially underrepresenting heterogeneity or contextual variation.

Within the framework developed in this paper, this thermodynamic bias has architectural implications. If digital infrastructures encode ontologies that privilege certain invariants such as engagement metrics or scale efficiencies optimization will align with those invariants. The geometry of collective action will reflect the compressed representation.

Thus, the same structural tendency that produces elegant physical theories can, in infrastructural contexts, produce oversimplified objective landscapes. Systems optimize for what is conserved in their ontology. If the ontology privileges a narrow invariant, optimization may induce curvature concentration and entropy misallocation elsewhere.

The recognition of this bias does not reject mathematics or science. Rather, it clarifies that abstraction is an operator subject to thermodynamic constraint. Stability in complex systems requires that compression preserve not only predictive efficiency but also multiscale invariants relevant to long-term coherence.

In this sense, the design of bounded substrates is an extension of scientific humility. It formalizes which invariants are preserved and enforces conservation constraints explicitly, preventing overcompression from generating systemic instability.

43.4 Entropy Resistance and the Engineering of Selective Worlds

The thermodynamic bias toward compression and invariant preservation extends beyond representation into intervention. Human technological development may be interpreted as a systematic attempt to counteract locally entropic processes by selectively stabilizing preferred invariants. In this sense, engineering is structured resistance to entropy gradients at chosen scales.

Let \mathcal{B} denote a biological system with state $x(t)$ evolving under intrinsic dynamics

$$\dot{x} = F_{\mathcal{B}}(x).$$

If $F_{\mathcal{B}}$ contains degenerative modessuch as renal failure, arrhythmic instability, or sensory lossthen certain invariant functionals $J(x)$ will drift outside viability bounds. Technological intervention introduces a corrective operator U such that

$$\dot{x} = F_{\mathcal{B}}(x) + U(x),$$

where U is designed to restore or maintain $J(x)$ within admissible regions.

A dialysis machine is not a complete simulation of renal physiology. It is an operator preserving a specific invariant: bounded concentration of metabolic waste products. Let $c(t)$ denote toxin concentration in blood. The untreated dynamics satisfy

$$\dot{c} = G(c) > 0$$

in renal failure. Dialysis introduces a removal term

$$\dot{c} = G(c) - \gamma c,$$

where $\gamma > 0$ represents clearance rate. The machine enforces a bounded attractor for $c(t)$ without replicating the kidneys full biochemical complexity. It preserves the invariant relevant to organismal survival while ignoring many others.

Similarly, prosthetic limbs preserve locomotion invariants without replicating muscular or neural microstructure. A pacemaker preserves rhythmic invariants without replicating electrophysiological self-organization. In each case, intervention identifies a functional ontologymobility, filtration, pacingand constructs a surrogate operator that stabilizes that invariant against entropic drift.

The same structural logic governs agricultural selection and domestication. Let \mathcal{P}

denote a population evolving under ecological dynamics

$$\dot{p} = F_{\text{eco}}(p).$$

Selective breeding modifies the effective fitness landscape by reweighting reproductive probabilities. If $W(p)$ denotes a trait functional, artificial selection introduces a bias term

$$\dot{p} = F_{\text{eco}}(p) + \beta \nabla W(p),$$

thereby steering population dynamics toward configurations preferred by the selecting agent. Entropy at the genetic level is constrained to produce phenotypic regularity aligned with human utility.

More generally, built environments—cities, climate control systems, medical infrastructure—are large-scale entropy regulators. They channel stochastic environmental variation into controlled flows. Roads constrain vehicular motion to low-dimensional manifolds. HVAC systems regulate thermal gradients. Supply chains regularize resource distribution. In each case, the system reduces effective state space by imposing structural constraints.

This pattern may be interpreted as a recursive attempt to reengineer nature by isolating desirable invariants and suppressing destabilizing modes. Let \mathcal{E} denote environmental dynamics. Engineering constructs a nested sequence of constraint layers

$$\mathcal{E} \rightarrow \mathcal{E}_1 \rightarrow \mathcal{E}_2 \rightarrow \cdots,$$

each reducing entropy production relative to selected variables while potentially increasing entropy elsewhere. Local order is purchased at global energetic cost, consistent with the second law.

The aspiration toward utopia can be understood in this formal sense. Utopia corresponds to a state in which preferred invariants are stabilized indefinitely. Formally, if $\{J_i\}$ are valued functionals—health, longevity, predictability, abundance—utopian engineering attempts to construct operators U such that

$$\frac{d}{dt} J_i(x(t)) \geq 0$$

or at least remains bounded within favorable intervals.

However, the preservation of some invariants necessarily suppresses others. Compression in representational space, selective breeding in biological space, and infrastructure in physical space all involve quotienting state space along chosen axes. What appears as progress from one ontology may appear as distortion from another.

The transhumanist impulse is therefore structurally continuous with scientific abstraction. Both involve identifying invariants that reduce complexity and engineering systems that conserve those invariants. The difference lies in scope: abstraction simplifies description; intervention simplifies dynamics.

Within the geometric framework of this paper, such interventions modify curvature and entropy flows on the manifold of collective life. By constraining degrees of freedom, engineered systems flatten certain regions of state space while potentially increasing curvature elsewhere. Stability requires that these modifications remain bounded and that global invariantsecological, thermodynamic, socialare not violated.

Thus the drive to build dialysis machines, prostheses, artificial organs, climate systems, and autonomous infrastructure reflects a persistent entropy-resistance strategy. It is not a denial of the second law but an exploitation of it: local order is created by exporting entropy and consuming energy.

The crucial architectural question becomes whether these entropy-resistant interventions are themselves embedded in substrates that prevent runaway concentration and curvature singularities. Without structural bounds, optimization of selected invariants may induce systemic instability. With explicit conservation laws and bounded entropy budgets, intervention can remain aligned with multiscale coherence.

The aspiration toward improved conditionshealth, predictability, extended capacityis therefore inseparable from the need for constraint-first design. The challenge is not to cease engineering, but to formalize which invariants are preserved and to ensure that entropy redistribution does not destabilize the larger manifold within which engineered order is embedded.

43.5 Pathology as Ontological Preference

The classification of certain processes as “disease,” “infection,” or “invasive species” reveals a deeper structural feature of functional ontology: pathology is defined relative to a chosen invariant. What is termed pathological is not ontologically primitive; it is a dynamical process whose trajectory violates a preferred constraint.

Let \mathcal{E} denote an ecological or biological manifold with state $x(t)$ evolving under intrinsic dynamics

$$\dot{x} = F(x).$$

Suppose there exists a functional $J : \mathcal{E} \rightarrow \mathbb{R}$ representing a valued invariant, such as organismal viability, metabolic stability, agricultural yield, or ecosystem balance. A

process is classified as disease when

$$\frac{d}{dt}J(x(t)) < 0$$

in a region deemed unacceptable relative to the ontology defining J .

In infectious dynamics, the pathogen is not intrinsically disordered. It is an autocatalytic subsystem whose replication operator

$$\dot{p} = G(p, h)$$

increases its own state variable p while decreasing host-functional invariants $J(h)$. The designation “disease” arises when host ontology assigns higher weight to preservation of $J(h)$ than to expansion of p .

Similarly, an “invasive species” is a dynamical agent whose population growth

$$\dot{s} = H(s, \mathcal{E})$$

alters ecological invariants valued by human or preexisting ecological ontologies. The species itself obeys consistent evolutionary dynamics. Its classification as invasive reflects a preference ordering imposed on ecosystem states.

Thus control measures—vaccination, antibiotics, quarantine, ecological management—can be modeled as corrective operators

$$\dot{x} = F(x) + U(x),$$

where U is designed to restore or preserve $J(x)$ within admissible bounds. The intervention suppresses one autocatalytic process in favor of another. There is no ontologically neutral equilibrium; there are only preferred invariants.

From a geometric perspective, pathology corresponds to curvature deviation relative to an imposed metric. If $(\mathcal{M}, g_{\mathcal{O}})$ is the manifold endowed with metric tensor determined by ontology \mathcal{O} , then a trajectory that appears destabilizing under $g_{\mathcal{O}}$ may be neutral or even optimal under a different metric $g_{\mathcal{O}'}$.

This reveals the structural continuity between medical engineering, agricultural management, and infrastructure design. In each case, agents select invariants to conserve and suppress dynamical modes that degrade those invariants. The suppression is not absolute; it is local and energy-dependent, consistent with thermodynamic constraints. Entropy is not eliminated but redirected.

Mathematically, let $\{J_i\}$ be a family of invariants with weights $\{w_i\}$ defining a

composite preference functional

$$\mathcal{U}(x) = \sum_i w_i J_i(x).$$

Intervention attempts to ensure

$$\frac{d}{dt}\mathcal{U}(x(t)) \geq 0$$

within energy and feasibility bounds. The classification of processes as beneficial or pathological follows from the sign structure induced by $\{w_i\}$.

This formalism clarifies that efforts to control disease or invasive species are not departures from natural law but reweightings of dynamical preference within the same thermodynamic framework. Autocatalysis and replication are universal processes. The distinction between organism and pathogen is drawn by functional boundary conditions.

Within the broader argument of this paper, this insight underscores a central claim: ontology determines optimization landscape. What is suppressed in one coordinate system may be amplified in another. Stability requires explicit articulation of which invariants are conserved and which degrees of freedom are allowed to fluctuate.

Engineering against disease, therefore, is a paradigmatic instance of entropy-bounded design. It recognizes that unchecked replication induces curvature concentration in biological state space. Intervention flattens that curvature relative to selected invariants. The question is not whether such interventions are natural, but whether their embedding respects multiscale coherence and global thermodynamic limits.

The same logic applies to digital infrastructure. A platform that suppresses certain flows while amplifying others is performing ontological selection. The difference between stabilizing intervention and extractive domination lies in whether invariant preservation is locally bounded and globally coherent.

Disease control, ecological management, and substrate design are thus structurally homologous operations: they express value-laden curvature modulation within dynamical manifolds governed by universal laws.

43.6 Symmetry, Invariance, and the Implicit Axiology of Science

Modern science is often presented as normatively neutral, yet its formal structure reveals persistent axiological commitments. Across physics, mathematics, and engineering, certain properties are repeatedly privileged: symmetry, invariance, conservation, minimality, and stability. These are not arbitrary aesthetic preferences but structural constraints that reduce descriptive redundancy and cognitive load. Nevertheless, their persistent elevation reflects a collective valuation.

Let (\mathcal{M}, g) be a model space describing a physical or abstract system. A transformation group G acts on \mathcal{M} by diffeomorphisms. A quantity $Q : \mathcal{M} \rightarrow \mathbb{R}$ is said to be invariant under G if

$$Q(g \cdot x) = Q(x) \quad \text{for all } g \in G.$$

The search for invariants is foundational in mathematics and physics. Noether's theorem makes this explicit: continuous symmetries correspond to conserved quantities. The more symmetry a system possesses, the fewer independent degrees of freedom must be tracked. Symmetry compresses description.

From an information-theoretic standpoint, invariance reduces entropy of representation. If states related by G are treated as equivalent, the effective state space becomes the quotient \mathcal{M}/G , whose cardinality is strictly smaller when G acts nontrivially. Thus symmetry induces informational compression.

This compression is not merely technical; it aligns with an aesthetic judgment. Structures exhibiting large symmetry groups are described as beautiful, elegant, or natural. In contrast, arbitrary assemblages lacking symmetry—what one might call a class palette of disconnected elements—are typically regarded as inelegant because they resist compression and do not generate invariants under transformation.

Formally, let a configuration consist of independent components $\{x_i\}_{i=1}^n$ with no nontrivial symmetry relations. The automorphism group of the configuration is trivial. There is no reduction of dimensionality through quotienting. The descriptive entropy remains maximal:

$$H(\{x_i\}) = \sum_i H(x_i),$$

with no mutual constraints. Such a configuration does not contain structure beyond aggregation. Its aesthetic deficit corresponds to absence of invariant structure.

By contrast, when constraints enforce relations among components—such as $x_1 = x_2$ or $x_i = f(x_j)$ —the mutual information between components increases, and the entropy

of independent description decreases:

$$H(\{x_i\}) = \sum_i H(x_i) - \sum_{i \neq j} MI(x_i, x_j).$$

Redundancy induced by constraint yields coherence. This reduction of independent degrees of freedom is perceived as structural unity.

In physics, the principle of least action provides another example of embedded axiology. Of all admissible trajectories γ , the realized path extremizes the action functional

$$S[\gamma] = \int L(\gamma(t), \dot{\gamma}(t)) dt.$$

Minimality or stationarity is elevated as natural. In mathematics, variational principles and extremal characterizations serve as definitions of canonical objects. In statistics, maximum likelihood and entropy maximization principles define optimal inference.

Across domains, extremization and invariance function as normative operators. They select among possibilities those configurations that are compressible, stable under perturbation, and generative of conservation laws.

This suggests that science implicitly encodes a preference for configurations that admit nontrivial symmetry groups and stable invariant measures. Beauty, in this context, is not mystical but structural: it reflects reduction in independent parameters and increase in coherence across scales.

The collective agreement on what is not beautiful corresponds to configurations lacking such compressibility. A purely arbitrary collection of features without symmetry or invariance fails to generate conserved quantities or stable quotients. It cannot serve as a generative template. Its ontology is flat and unstructured.

This does not imply that asymmetry or irregularity lack value. Indeed, symmetry breaking is essential for structure formation. However, even symmetry breaking is studied relative to a higher symmetric state. The broken state inherits structure from the symmetry it departs from. Pure randomness without latent symmetry is rarely the object of scientific admiration.

Thus there exists an implicit axiology: systems are preferred when they admit invariant structure, reduce dimensionality under transformation, and exhibit stability under perturbation. These preferences are deeply entangled with thermodynamic and informational constraints. Invariance reduces entropy of description; conservation stabilizes dynamics; symmetry ensures reproducibility.

The aesthetic of mathematics and physics is therefore not external to physical law

but emerges from it. The second law favors macrostates with large phase volume, yet scientific models privilege those macrostates that can be summarized by low-dimensional invariants. What is admired is not entropy itself, but structure persisting under entropy.

In this sense, scientific beauty reflects a convergence between cognitive economy and physical constraint. A configuration that cannot be compressed, related by symmetry, or stabilized by invariant quantities fails to support predictive leverage. It does not generate a durable ontology.

The valuation of symmetry and invariance is thus neither arbitrary nor purely subjective. It arises from the intersection of thermodynamic limits, informational compression, and dynamical stability. What we call beautiful in science is often that which can be mapped surjectively onto a smaller space without loss of essential structure.

43.7 From Scientific Axiology to Substrate Design

The preceding analysis suggests that modern science is not axiologically neutral. It persistently privileges symmetry, invariance, conservation, and minimal description length. These are not merely aesthetic preferences but structural principles that reduce dimensionality, compress representation, and stabilize dynamics. Systems that admit large symmetry groups and conserved quantities are describable by fewer independent parameters; they generate predictive leverage with minimal informational overhead.

The same structural commitments appear in physical law. Noether's theorem links continuous symmetries to conservation laws. Variational principles select extremal configurations. Renormalization theory identifies scale-invariant fixed points as structurally stable regimes. In each case, persistence corresponds to invariance under transformation.

Digital infrastructure, however, has historically lacked such invariants. Engagement-optimized platforms operate without conserved quantities. Attention is not conserved; entropy is not bounded; influence does not decay. The absence of invariant structure permits curvature concentration and scalar-well formation within the informational manifold.

If symmetry and invariance are the conditions of stability in physics, then their absence in digital systems explains structural instability in the political economy of platforms. Extraction corresponds to violation of conservation principles. Viral amplification corresponds to unbounded positive feedback without symmetry constraint. Platform dominance corresponds to curvature concentration without compensating

diffusion.

The PlenumHub architecture can therefore be interpreted as an attempt to manufacture invariants at the substrate level. Entropy budgets impose a bounded growth condition analogous to thermodynamic conservation. Texture and Time Crystals define conserved valuation measures grounded in mutual information and temporal decay. Sheaf-theoretic closure ensures that local assignments extend uniquely to global sections, preventing contradictory ontological states from coexisting within the same substrate.

Formally, let \mathcal{I} denote the set of invariants preserved under admissible transformations of semantic spheres. A transformation $T : \Sigma \rightarrow \Sigma$ is admissible if

$$I(T(\sigma)) = I(\sigma) \quad \text{for all } I \in \mathcal{I}.$$

Stability of the substrate corresponds to closure of admissible transformations under composition. If \mathcal{T} denotes the monoid of admissible transformations, then

$$T_1, T_2 \in \mathcal{T} \implies T_2 \circ T_1 \in \mathcal{T}.$$

This algebraic closure condition mirrors physical symmetry groups: permissible operations must preserve conserved quantities.

In this light, substrate design becomes applied scientific axiology. The same principles that render physical theories elegant render digital infrastructures stable. Beauty in mathematics corresponds to invariance under transformation; stability in digital systems corresponds to invariance under adversarial or engagement-driven perturbation.

The second law of thermodynamics implies that entropy tends to increase in closed systems. Scientific modeling responds by identifying macroscopic invariants that persist despite microscopic disorder. Digital systems, if left unconstrained, exhibit analogous entropy inflation at the semantic level. The role of manufactured invariants is to bound this inflation and prevent geometric phase instability.

Thus the aesthetic commitments of science are not separable from infrastructure design. Symmetry, invariance, and conservation are not merely descriptive virtues but operational necessities. To manufacture stability is to encode these principles into the substrate itself.

In this sense, the project is not utopian. It does not attempt to eliminate entropy or conflict. Rather, it imposes structural constraints under which entropy growth is bounded and curvature concentration cannot accumulate unchecked. The architecture

becomes the enforcement mechanism of the implicit axiology already embedded in mathematical practice.

43.8 Invariant-Preserving Semigroups and Structural Stability

We now formalize the preceding claim in the minimal mathematical form needed for later appendices. The role of “axiology” is encoded as a chosen family of invariants and the admissible transformations that preserve them. Stability is then expressed as a structural property of the resulting transformation semigroup.

State space and invariants. Let Σ be a Hausdorff topological space of substrate states. A state $\sigma \in \Sigma$ represents a full configuration of the system at a chosen resolution, including semantic allocations, interaction weights, and any auxiliary fields (entropy budgets, curvature proxies, timestamp measures, and so on).

Let \mathcal{I} be an index set and let

$$I_i : \Sigma \rightarrow \mathbb{R}, \quad i \in \mathcal{I},$$

be a family of continuous functionals called invariants. We write $I(\sigma)$ for the vector $(I_i(\sigma))_{i \in \mathcal{I}}$ in $\mathbb{R}^{\mathcal{I}}$. In practice, \mathcal{I} may contain quantities such as entropy budgets, bounded curvature concentration scores, mutual-information lower bounds, or decay-normalized valuation measures.

Admissible transformations. A transformation is a map $T : \Sigma \rightarrow \Sigma$. We say that T is \mathcal{I} -preserving if

$$I_i(T(\sigma)) = I_i(\sigma) \quad \text{for all } \sigma \in \Sigma, i \in \mathcal{I}.$$

Let $\mathcal{T}_{\mathcal{I}}$ denote the set of all \mathcal{I} -preserving transformations. Then $\mathcal{T}_{\mathcal{I}}$ is closed under composition and contains the identity map, hence forms a monoid. When inverses exist (e.g. for reversible updates), $\mathcal{T}_{\mathcal{I}}$ becomes a group; however the substrate setting is generically dissipative and only the monoid structure is needed.

Definition 43.1 (Invariant-preserving semigroup). *A subset $\mathcal{S} \subseteq \mathcal{T}_{\mathcal{I}}$ is called an invariant-preserving semigroup if it is nonempty and closed under composition:*

$$T_1, T_2 \in \mathcal{S} \implies T_2 \circ T_1 \in \mathcal{S}.$$

If additionally $\text{id}_\Sigma \in \mathcal{S}$, then \mathcal{S} is a submonoid of \mathcal{T}_Σ .

Stability notion. We require a notion of robustness under repeated admissible action. Since the substrate is typically nonreversible and may be stochastic at the micro-level, the appropriate stability concept is forward invariance of bounded sets and absence of runaway amplification in invariant coordinates.

Let $\|\cdot\|$ be a seminorm on $\mathbb{R}^{\mathcal{I}}$ that metrizes the topology on the image of I restricted to the invariants of interest. Define the invariant-radius function

$$R(\sigma) := \|I(\sigma)\|.$$

For example, if \mathcal{I} is finite, one may take $\|I(\sigma)\|$ to be an ℓ^p norm; if \mathcal{I} is infinite, one may take a weighted ℓ^2 seminorm or a supremum norm on a chosen finite subfamily.

Definition 43.2 (Invariant-bounded orbit). *Given a semigroup \mathcal{S} acting on Σ , an orbit $\{T_n(\sigma)\}_{n \geq 0}$ with $T_n \in \mathcal{S}$ is invariant-bounded if*

$$\sup_{n \geq 0} R(T_n(\sigma)) < \infty.$$

In the \mathcal{I} -preserving case, invariant-boundedness holds trivially, because $I(T(\sigma)) = I(\sigma)$ for all $T \in \mathcal{S}$. The point of introducing this definition is that realistic systems do not preserve invariants exactly; they preserve them up to a controlled slack. That slack corresponds precisely to an entropy budget or curvature budget.

Approximate invariance and slack budgets. Let $\delta \geq 0$ represent a permitted deviation per update step. We say that T is (\mathcal{I}, δ) -admissible if

$$\|I(T(\sigma)) - I(\sigma)\| \leq \delta \quad \text{for all } \sigma \in \Sigma.$$

Let $\mathcal{T}_{\mathcal{I}, \delta}$ be the set of all such maps. When $\delta = 0$, this reduces to exact invariance.

Lemma 43.3 (Budget accumulation bound). *Let $T_1, \dots, T_n \in \mathcal{T}_{\mathcal{I}, \delta}$ and define $S_n := T_n \circ \dots \circ T_1$. Then for all $\sigma \in \Sigma$,*

$$\|I(S_n(\sigma)) - I(\sigma)\| \leq n\delta.$$

Proof. By telescoping,

$$I(S_n(\sigma)) - I(\sigma) = \sum_{k=1}^n \left(I(S_k(\sigma)) - I(S_{k-1}(\sigma)) \right),$$

with $S_0 = \text{id}$. Taking norms and applying the triangle inequality together with (\mathcal{I}, δ) -admissibility of each T_k yields

$$\|I(S_n(\sigma)) - I(\sigma)\| \leq \sum_{k=1}^n \|I(S_k(\sigma)) - I(S_{k-1}(\sigma))\| \leq \sum_{k=1}^n \delta = n\delta.$$

□

This lemma is the discrete-time analogue of the entropy-bounded growth bound from Appendix E: the invariant drift is at most linear in time when per-step slack is uniformly bounded.

Engagement dynamics as eigen-expansion. The contrasting case is a transformation family that amplifies a scalar functional multiplicatively. Let $V : \Sigma \rightarrow \mathbb{R}_{\geq 0}$ be a functional (interpreted as an “engagement energy”) and suppose there exists $\alpha > 0$ such that for some maps T ,

$$V(T(\sigma)) \geq (1 + \alpha)V(\sigma) \quad \text{for all } \sigma \in \Sigma.$$

Iterating implies $V(S_n(\sigma)) \geq (1 + \alpha)^n V(\sigma)$, which is exponential growth.

This is the semigroup expression of the qualitative gap: invariance budgets produce at worst linear drift in invariant coordinates, whereas multiplicative amplification yields exponential divergence.

Theorem 43.4 (Structural stability from invariant budgets). *Fix invariants $I : \Sigma \rightarrow \mathbb{R}^{\mathcal{I}}$ and a seminorm $\|\cdot\|$. Let \mathcal{S} be a semigroup generated by maps in $\mathcal{T}_{\mathcal{I}, \delta}$ for some $\delta \geq 0$. Then for every $\sigma \in \Sigma$ and every orbit $\sigma_n = S_n(\sigma)$ with $S_n \in \mathcal{S}$, we have the bound*

$$\|I(\sigma_n) - I(\sigma)\| \leq n\delta.$$

In particular, if $\delta = 0$ then $I(\sigma_n)$ is constant along the orbit. If $\delta > 0$ is small relative to the scale at which geometric or semantic phase changes occur, then phase instability cannot occur without exhausting a linear budget in time.

Proof. Apply the budget accumulation lemma to the composition representation of S_n in terms of generators. □

Interpretation for the substrate. The theorem is intentionally modest: it does not claim convergence to a fixed point, only that invariant drift cannot explode faster than linearly when each update respects a uniform slack budget. This suffices for the

main architectural claim. If the substrate enforces per-step bounds on entropy increase and curvature concentration, then long-horizon instability requires the deliberate expenditure of a cumulative budget. The system cannot drift into an exploitative regime “for free” through uncontrolled exponential amplification.

Later appendices strengthen this statement by providing spectral criteria for mixing (via curvature and spectral gap), and existence results for stationary measures in the stochastic setting (via Krylov–Bogolyubov arguments), thereby connecting invariant budgets not only to boundedness but to ergodic stabilization.

44 Control-Theoretic Ontology

44.1 State Spaces and Observational Projections

Let \mathcal{X} denote the full physical state space of a system, with dynamics governed by

$$\dot{x}(t) = f(x(t), u(t)),$$

where $x(t) \in \mathcal{X}$ and $u(t)$ is an admissible control input.

Cognition does not operate directly on \mathcal{X} . Instead, it constructs a reduced representation

$$z(t) = h(x(t)),$$

where $h : \mathcal{X} \rightarrow \mathcal{Z}$ is an observational projection onto a lower-dimensional state space \mathcal{Z} . The map h defines an ontology: it selects which degrees of freedom are salient and which are neglected.

The reduced dynamics satisfy

$$\dot{z}(t) = Dh(x(t)) \cdot f(x(t), u(t)),$$

which induces an effective evolution law

$$\dot{z}(t) = \tilde{f}(z(t), u(t)).$$

The choice of h determines the coordinate chart in which control is defined. Ontology, in this sense, is a selection of admissible state variables.

44.2 Preferences as Reaction Functionals

Let $J[u]$ be a cost functional defined over trajectories in \mathcal{Z} :

$$J[u] = \int_0^T L(z(t), u(t)) dt.$$

A controller selects $u(t)$ to minimize J . The functional L encodes deviations from preferred states.

Thus preferences are reaction laws: they define how the system responds to deviation in the chosen coordinate chart.

A thermostat minimizes temperature deviation; a pacemaker minimizes rhythm deviation; a logistics network minimizes delivery delay. Each embeds a normative preference as a feedback operator.

The closed-loop dynamics become

$$\dot{x}(t) = f(x(t), \pi(h(x(t))))),$$

where π is the control policy induced by minimizing J .

Hence ontology and axiology are inseparable: the chosen projection h determines which deviations are observable, and the cost functional determines which deviations are penalized.

45 Recapitulation and Externalization of Charts

45.1 Projection and Reinstantiation

Let $\pi : \mathcal{X} \rightarrow \mathcal{Z}$ be a surjective map identifying equivalence classes of microstates.

The macro-dynamics on \mathcal{Z} may be approximated by

$$\dot{z}(t) = g(z(t)).$$

Technological development often consists in constructing a new physical system \mathcal{X}' whose intrinsic dynamics approximate g directly:

$$\dot{x}'(t) \approx g(x'(t)).$$

This process converts a descriptive abstraction into a prescriptive physical architecture. A pattern first identified as a regularity in nature is subsequently reified as an engineered constraint, transforming an observational model into an executable mechanism. Thermodynamic cycles, once formalized as idealized transformations in state space, are instantiated as engines whose pistons and turbines enact the abstract Carnot or Rankine diagrams in steel and steam.

Circulatory dynamics, initially understood through biological models of pressure gradients and flow conservation, are recast as mechanical pumps that reproduce volumetric transfer through rotating impellers or oscillatory chambers. Logical inference, formalized through symbolic calculus and truth-functional operations, is ultimately embedded in digital circuits whose transistor networks implement Boolean algebra as voltage transitions. In each case, an abstract functional ontology becomes materially enforced. The diagram ceases to be merely explanatory and becomes constitutive. Description collapses into design.

In each case, a projection π becomes physically instantiated as a subsystem whose state variables are precisely the coordinates of \mathcal{Z} .

45.2 Swarm Dynamics and Stigmergic Scaffolds

Consider a multi-agent system with agents indexed by $i = 1, \dots, N$ and states $x_i(t)$.

Suppose each agent follows local update rules

$$\dot{x}_i = F(x_i, \rho),$$

where ρ is a field generated collectively by the agents.

Such systems generate stigmergic structures: macroscopic order emerges from local reaction to shared environmental gradients.

Ant foraging trails, vascular networks, and transportation systems can all be modeled by reaction-diffusion equations of the form

$$\partial_t \rho = D \Delta \rho + \sum_i S(x_i).$$

Over evolutionary time, efficient stigmergic configurations are selected because they minimize transport cost functionals.

Thus modular organization emerges not by central design but by dynamical stability under feedback constraints.

46 Functional Ontologies and Substrate Independence

46.1 Behaviorism as Functional Abstraction

A functional ontology identifies a system not by its material substrate but by its input-output relation.

Let a system be defined by a transfer function

$$y(t) = \Phi(u(t)).$$

If two systems share the same transfer function Φ , mapping admissible inputs to outputs in an identical manner, then they are functionally equivalent independently of the material substrate in which Φ is realized. Functional identity is therefore defined not by composition of matter but by invariance of the input-output relation.

This principle is exemplified in biomedical substitution. A cardiac pacemaker reproduces the impulse-generating transfer function of the sinoatrial node by enforcing periodic electrical excitation across myocardial tissue. A dialysis machine implements the effective filtration map of renal glomeruli by replicating diffusion and selective permeability across an artificial membrane. In computation, a software emulator enacts the transition algebra of a hardware circuit, preserving operational semantics even though the physical realization differs.

In formal terms, let $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$ be a measurable transformation between state spaces. Define the equivalence class

$$[\Phi] = \{S \mid S \text{ implements } \Phi\}.$$

All members of $[\Phi]$ are functionally indistinguishable relative to \mathcal{X} and \mathcal{Y} .

Behaviorism, in this abstract formulation, is the recognition that ontology may be defined at the level of transfer function rather than mechanism. A system is identified with the equivalence class of processes implementing Φ . Mechanistic detail becomes secondary to functional invariance.

46.2 Optimization and Entropic Compression

Scientific abstraction privileges invariance and symmetry because these reduce the dimensionality of state description.

Let G be a symmetry group acting on \mathcal{X} . A function $I(x)$ is invariant if

$$I(g \cdot x) = I(x) \quad \forall g \in G.$$

Invariants define conserved quantities and reduce the effective dimension of dynamics.

Dimensional reduction lowers descriptive entropy and cognitive load, functioning analogously to energy minimization in physical systems.

Thus aesthetic preference for symmetry corresponds to preference for reduced representational entropy.

This establishes a link between scientific abstraction and thermodynamic compression.

47 Civilizational Control and Lyapunov Invariants

Let $V(x)$ be a Lyapunov candidate representing a civilizational invariant.

If

$$\dot{V}(x) \leq 0,$$

then V defines a stable attractor.

Different civilizations implicitly choose different V .

An engagement-optimized platform implicitly sets

$$V(x) = -E(x),$$

driving entropy upward.

An entropy-bounded substrate sets

$$V(x) = E(x),$$

limiting its growth.

Hence infrastructure encodes Lyapunov functions at scale.

Civilization becomes a feedback controller whose ontology defines its state variables and whose invariants define its attractors.

48 From Control to Curvature: Manufacturing Stability

The control-theoretic account developed in the preceding sections clarifies the geometric framework introduced earlier. A projection $h : \mathcal{X} \rightarrow \mathcal{Z}$ defines an ontology by selecting admissible state variables. A cost functional J defines preferences by penalizing deviation within that ontology. The induced feedback policy π then governs system evolution through closed-loop dynamics. When extended to institutional scale, this mechanism becomes infrastructural: the chosen projection and cost functional are embedded in code, hardware, and governance rules.

The geometric formulation recasts this architecture in differential form. Let (\mathcal{M}, g) denote the informational manifold whose points represent semantic states. A platform's ontology corresponds to a coordinate chart on \mathcal{M} ; its optimization target defines a scalar potential $\phi : \mathcal{M} \rightarrow \mathbb{R}$ that reshapes the metric geometry through induced curvature.

Under engagement-maximizing control, the platform effectively defines a Lyapunov function $V = -J$, whose maximization increases conditional entropy or predictive leverage. The induced scalar field ϕ develops positive Laplacian $\Delta_g \phi > 0$ in regions of concentrated activity, producing curvature concentration. By the Jacobi field analysis, such curvature causes geodesic convergence. Agents following locally optimal policies become dynamically funneled toward attractor basins. This is the geometric form of herding and lock-in.

In contrast, an entropy-bounded substrate defines a Lyapunov function $V = E$, where E is semantic entropy subject to bounded growth constraints. The entropy soundness theorem ensures that \dot{V} is uniformly bounded, preventing exponential amplification. The associated scalar field satisfies a controlled diffusion equation rather than runaway accumulation. Curvature remains approximately flat or bounded, preserving geodesic divergence and tangent diversity.

The phrase “manufacturing stability” therefore acquires a precise technical meaning. Stability is not an emergent accident but the consequence of embedding invariant-preserving feedback laws into the substrate itself. A system is stable when its Lyapunov candidate is structurally protected against exponential drift. It is unstable when its cost functional incentivizes positive feedback without geometric constraint.

Control theory and differential geometry thus converge. Ontology defines the coordinate system. Preference defines the potential. Feedback defines the flow. Curvature measures concentration. Stability corresponds to bounded Lyapunov growth.

Infrastructure is the physical instantiation of these mathematical relationships.

In this sense, digital architecture is not neutral. Every platform implements an implicit control system whose invariant determines the attractor structure of the informational manifold. If that invariant rewards engagement without constraint, curvature concentrates and extraction becomes inevitable. If that invariant enforces entropy budgets and coherence conservation, the manifold remains navigable and plural.

Manufacturing stability is therefore the deliberate selection and embedding of invariants that bound curvature and constrain entropy growth. It is the design of a substrate in which the geometric preconditions of extraction cannot arise because the Lyapunov structure of the system forbids them.

The project is not regulatory but architectural. It does not seek to correct trajectories after instability appears; it redefines the dynamical system so that instability is no longer an admissible phase.

49 Pipeline Closure and the Desire for Forward Motion

Once a transformation has been successfully performed, a new structural object is created: a pipeline. Formally, let

$$T : \mathcal{X} \rightarrow \mathcal{Y}$$

be a transformation mapping inputs to outputs. The first successful execution of T demonstrates that such a mapping is feasible. Repetition stabilizes the mapping. Stabilization produces a reusable operator. That operator, once embedded materially or institutionally, becomes infrastructure.

Thus every accomplished act implicitly defines a function.

A factory is the materialization of a function T . A bureaucracy is the materialization of a function T . A neural pathway is the biological materialization of a function T . Even a habit is the neural instantiation of a repeated inputoutput mapping.

Pipeline formation therefore corresponds to the closure of a dynamical transformation under repetition.

Let $x_{n+1} = T(x_n)$. If T is stable under perturbation and composable with itself, it defines an operator algebra over \mathcal{X} . Civilization progresses by accumulating such

operators.

This explains a psychological phenomenon that appears trivial but is structurally profound: humans enjoy driving.

Driving is a controlled forward integration of state:

$$\dot{x}(t) = v(t),$$

with $v(t)$ chosen by the agent. The experience of motion corresponds to visible state evolution under direct control input.

More abstractly, driving creates a real-time pipeline from intention to displacement. The latency between input and output is low and continuous. The operator $T_{\Delta t}$ mapping position at time t to position at time $t + \Delta t$ is directly parameterized by human action.

The pleasure arises from stable forward flow through state space.

In dreams, humans frequently report experiences of flying. Flying corresponds to even fewer constraints on T . The state transition is not bound by road topology or friction constraints. The manifold becomes effectively flatter. The geodesic freedom increases.

The preference for forward motion can therefore be understood as a preference for low-curvature, low-latency, stable pipelines between intention and consequence.

In high-curvature informational environments, intention does not map predictably to outcome. Feedback is delayed, distorted, or captured. Engagement platforms deliberately increase curvature, making the mapping from action to result opaque and probabilistic. The agent no longer feels forward motion but stochastic drift.

Pipeline closure restores perceived agency.

Once a transformation is routinized, it compresses entropy. The operator T reduces variance between intention and outcome. This reduction in conditional entropy produces cognitive relief. The system becomes predictable.

Formally, if X denotes intention and Y denotes outcome, pipeline stabilization reduces

$$H(Y | X).$$

Low conditional entropy corresponds to high perceived control.

Hence humans are drawn toward activities that instantiate stable, low-latency input/output pipelines: driving, gaming, crafting, programming, machining, piloting.

These activities are not merely recreational. They are local experiences of manufactured stability.

Civilization can be viewed as the accumulation of such stabilized operators. Each successful abstractionengine, pump, compiler, assembly linecreates a reusable morphism in the category of transformations. The more morphisms we accumulate, the more compositional the system becomes.

However, pipelines can also be captured.

If a central node inserts itself into T such that

$$T = T_2 \circ C \circ T_1,$$

where C extracts surplus from every invocation, then the pipeline becomes an extraction channel. The curvature analysis applies: repeated invocation concentrates flow into C .

Thus the difference between generative infrastructure and extractive platform lies in whether the pipeline remains an open operator algebra or becomes curvature-concentrated through forced mediation.

Humans enjoy forward motion because it is the phenomenology of stable operator application.

Manufacturing stability is therefore the architectural problem of designing pipelines whose compositional structure preserves low conditional entropy without introducing curvature concentration.

50 Discourse as Coupled Manifold Dynamics

50.1 Semantic Manifolds and Communicative Morphisms

Let each cognitive agent be modeled as a smooth semantic manifold (\mathcal{M}_i, g_i) , where points represent structured conceptual states and the metric g_i encodes inferential proximity between concepts. Differences in ontology correspond to differences in metric structure and coordinate charts.

Communication between agents A and B is not the literal transfer of semantic objects from \mathcal{M}_A to \mathcal{M}_B . Rather, it induces a mapping

$$\Phi_{A \rightarrow B} : \mathcal{M}_A \dashrightarrow \mathcal{M}_B,$$

which is partial, approximate, and constrained by the pre-existing structure of \mathcal{M}_B .

A communicative act deforms \mathcal{M}_B locally by embedding a region of \mathcal{M}_A into it. The quality of communication depends on the extent to which this embedding preserves local geometric structure.

50.2 Differential Conceptual Alignment

Let $x \in \mathcal{M}_A$ and suppose agent B constructs $y \in \mathcal{M}_B$ in response. Denote this construction by

$$y = \Psi(x).$$

Coherence requires approximate preservation of local distances:

$$d_A(x_1, x_2) \approx d_B(\Psi(x_1), \Psi(x_2))$$

for x_1, x_2 in a neighborhood $U \subset \mathcal{M}_A$.

Thus communicative success corresponds to local quasi-isometry between semantic neighborhoods.

Podcasting makes this differential process explicit. Participants continuously estimate and refine Ψ through iterative probing, paraphrase, and correction. Writing performs the same operation asynchronously. The author anticipates distortions in Ψ and modifies exposition to stabilize the mapping.

In both cases, discourse consists of iterative attempts to minimize metric distortion between manifolds.

50.3 Dialogue and Essay as Dynamical Systems

Consider two agents with evolving semantic states:

$$\dot{\mathcal{M}}_A = F_A(\mathcal{M}_B), \quad \dot{\mathcal{M}}_B = F_B(\mathcal{M}_A).$$

In synchronous dialogue, these equations represent mutually coupled manifold evolution. Each agent updates internal structure in response to the other.

In essay writing, the coupling becomes temporally asymmetric. The author constructs an internal model $\widehat{\mathcal{M}}_B$ and evolves according to

$$\dot{\mathcal{M}}_A = F_A(\widehat{\mathcal{M}}_B).$$

An essay is therefore a frozen dialogue: it simulates bidirectional manifold deformation within a single cognitive system.

Structurally, writing and podcasting differ only in temporal boundary conditions, not in ontological character.

50.4 Curvature, Contestation, and Stability

When two semantic manifolds possess incompatible curvature in overlapping regions, geodesics diverge under attempted alignment. This produces rhetorical instability: disagreement sharpens into contest.

Polemic corresponds to high-gain corrective deformation. It attempts to force alignment by increasing gradient magnitude, effectively imposing curvature rather than testing it.

Stability-oriented discourse proceeds differently. It seeks regions where local charts can be aligned through low-distortion embeddings. Rather than amplifying curvature, it compares sectional curvature across manifolds and identifies invariant substructures.

Thus discourse can be understood geometrically as curvature testing.

Where local isometries exist, dialogue expands shared semantic territory. Where curvature signs differ irreconcilably, divergence is revealed without requiring adversarial escalation.

50.5 Communication as Sheaf Alignment

Let \mathcal{F}_A and \mathcal{F}_B be presheaves of semantic assignments over context categories \mathcal{C}_A and \mathcal{C}_B .

Communicative alignment seeks a functor

$$\Theta : \mathcal{C}_A \rightarrow \mathcal{C}_B$$

such that sections over corresponding contexts satisfy compatibility conditions.

Coherence is achieved when local assignments glue consistently across overlapping contexts under Θ .

In this sense, discourse attempts to construct a sheaf-theoretic bridge between cognitive topologies.

50.6 Implications for Stability Manufacturing

If infrastructure amplifies curvature through engagement-optimized feedback, semantic manifolds fragment and local isometries become rare. Dialogue collapses into adversarial gradient escalation.

If infrastructure enforces entropy bounds and invariant preservation, local geometric distortion is limited. Shared charts become constructible. Discourse regains the structure of manifold alignment rather than rhetorical contest.

Writing and podcasting are therefore not merely expressive media. They are dynamical coupling mechanisms between structured semantic fields.

The stability of discourse depends on the curvature properties of the substrate in which it occurs.

51 Cognition as Framing and Hierarchical Evaluation

51.1 Framing as Projection

Let \mathcal{X} denote the full physical state space of the organismenvironment system. Cognition does not operate directly on \mathcal{X} . Instead, it constructs a framed state space

$$\mathcal{Z} = h(\mathcal{X}),$$

where $h : \mathcal{X} \rightarrow \mathcal{Z}$ is a projection selecting task-relevant variables, categories, and relations.

This projection constitutes a frame. It defines what counts as state, what counts as action, and what counts as outcome. Different frames correspond to different coordinate charts on the same underlying manifold.

Decision-making therefore begins not with choice but with projection. The act of framing restricts admissible trajectories in advance.

51.2 Scenario Generation as Evolutionary Branching

Given a framed state $z \in \mathcal{Z}$, cognition generates a family of hypothetical trajectories

$$\{\gamma_i : [0, T] \rightarrow \mathcal{Z}\}_{i \in I},$$

where each γ_i represents a possible future evolution under candidate control policies $u_i(t)$.

Formally, these satisfy

$$\dot{z}_i(t) = \tilde{f}(z_i(t), u_i(t)).$$

Each trajectory constitutes an internal evolutionary scenario. The set $\{\gamma_i\}$ forms a branching structure in state space analogous to a phylogenetic tree, but over hypothetical rather than biological evolution.

Cognition thus implements a model-based evolutionary process in simulation.

51.3 Evaluation as Hierarchical Stacking

Evaluation is the process by which candidate trajectories are compared and ranked.

Let J_i denote the cost functional associated with trajectory γ_i :

$$J_i = \int_0^T L(z_i(t), u_i(t)) dt.$$

Selection proceeds by minimizing J_i across $i \in I$.

However, the evaluation functional itself is hierarchical. Higher-level frames impose meta-constraints on lower-level evaluations. Let $J^{(1)}$ denote immediate utility and $J^{(2)}$ denote higher-order constraint (e.g., identity coherence, social norm, long-term survival). Then evaluation becomes nested:

$$J_i^{\text{total}} = J_i^{(1)} + \lambda J_i^{(2)} + \mu J_i^{(3)} + \dots$$

This hierarchical stacking produces a stratified decision manifold.

Lower-level evolutionary scenarios compete within constraints imposed by higher-level invariants. Decisions are therefore not simple optimizations but multi-level stabilization processes.

51.4 Frames as Evolutionary Filters

A frame h determines the space of candidate trajectories. Changing the frame alters which scenarios are even considered.

Thus framing is equivalent to redefining the evolutionary landscape over which

internal selection occurs.

If the frame defines only a small set of admissible state variables, the branching factor of $\{\gamma_i\}$ is restricted. This corresponds to cognitive rigidity. If the frame is expansive but unconstrained, evaluation becomes unstable due to combinatorial explosion.

Effective cognition balances frame dimensionality against evaluation stability.

51.5 Geometric Interpretation

In the manifold formulation, each hypothetical trajectory γ_i is a curve in (\mathcal{Z}, g) . Evaluation compares their accumulated length under a weighted metric induced by L .

Hierarchical stacking corresponds to embedding (\mathcal{Z}, g) within a higher-dimensional manifold (\mathcal{Z}', g') where additional invariants constrain admissible geodesics.

Decision-making is therefore curvature-sensitive trajectory selection within nested manifolds.

51.6 Implications

Cognition is not merely representation. It is controlled evolutionary simulation within a framed state space.

Frames determine admissible scenarios. Evaluation selects among them. Hierarchical invariants constrain the selection process. Over time, repeated evaluation stabilizes particular frames, which become habits, identities, or institutional norms.

In this sense, cognition manufactures stability internally before infrastructure manufactures it externally.

The same mathematical principles govern both.

52 Cognitive Entropy and Substrate Entropy

52.1 Internal Entropy in Scenario Generation

Let $\{\gamma_i\}_{i \in I}$ denote the set of internally generated candidate trajectories under a given frame h . The uncertainty over which trajectory will be selected can be represented as a probability distribution $p(i)$ over I .

Define cognitive entropy

$$H_{\text{cog}} = - \sum_{i \in I} p(i) \log p(i).$$

High H_{cog} corresponds to indecision or frame instability; low H_{cog} corresponds to strongly peaked preference over a narrow scenario set.

Evaluation reduces cognitive entropy by selecting a trajectory γ^* , collapsing the distribution.

However, entropy reduction internally requires structured evaluation externally. If the informational substrate amplifies noise, increases curvature, or destabilizes frames, then the branching factor $|I|$ increases and H_{cog} grows.

Thus there is a coupling between cognitive entropy and substrate entropy.

52.2 Substrate-Induced Branch Inflation

Let \mathcal{M} denote the informational manifold induced by infrastructure. Suppose its curvature satisfies a concentration condition in some region \mathcal{P} .

In such regions, local geodesics become unstable and small perturbations produce divergent hypothetical futures. Internally, this manifests as increased branching in $\{\gamma_i\}$ and higher variance in predicted outcomes.

Formally, if curvature K increases, then sensitivity of trajectories satisfies

$$\|\delta\gamma(t)\| \approx e^{\sqrt{K}t} \|\delta\gamma(0)\|.$$

This exponential sensitivity inflates cognitive entropy because prediction variance increases.

An entropy-bounded substrate constrains this effect. By limiting entropy growth in semantic evolution, it bounds curvature concentration and therefore bounds branching inflation in internal simulation.

Thus entropy soundness at the infrastructural level stabilizes evaluation at the cognitive level.

52.3 Coupled Stability Principle

Let $E_{\text{sub}}(t)$ denote substrate entropy and $H_{\text{cog}}(t)$ denote cognitive entropy. A stability condition may be expressed as

$$\frac{d}{dt}H_{\text{cog}} \leq \alpha \frac{d}{dt}E_{\text{sub}} + \beta,$$

for constants $\alpha, \beta \geq 0$.

Bounding $\frac{d}{dt}E_{\text{sub}}$ therefore bounds the growth of internal branching complexity.

This establishes a structural coupling: infrastructure design shapes cognitive load and evaluation stability.

Manufacturing stability externally reduces entropy inflation internally.

53 Reinforcement Learning and Predictive Framing

53.1 Evaluation as Value Function Estimation

In reinforcement learning, an agent estimates a value function

$$V^\pi(z) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid z_0 = z, \pi \right],$$

where π is a policy and r_t are rewards.

The generation of candidate trajectories $\{\gamma_i\}$ corresponds to sampling rollouts under different policies. Evaluation selects the policy minimizing expected cost or maximizing reward.

Thus hierarchical evaluation described previously can be interpreted as layered value function approximation, where higher-order frames impose regularization terms on V^π .

53.2 Predictive Processing and Free Energy

Under predictive processing models, cognition minimizes variational free energy:

$$\mathcal{F} = \mathbb{E}_q[\log q(s)] - \mathbb{E}_q[\log p(o, s)],$$

where $q(s)$ is an approximate posterior over hidden states and $p(o, s)$ is the generative model of observations.

Minimizing \mathcal{F} corresponds to reducing prediction error.

Frames h define the generative model structure. Different projections correspond to different priors $p(s)$.

Evaluation then becomes approximate Bayesian model comparison. Hypothetical trajectories are internal forward simulations under candidate generative models.

High-curvature informational environments introduce volatility in observed signals o , increasing prediction error variance and destabilizing inference.

An entropy-bounded substrate reduces volatility in semantic transitions, stabilizing predictive coding loops.

53.3 Hierarchical Stacking as Multi-Level Inference

Predictive processing is inherently hierarchical. Let latent states be organized as

$$s^{(3)} \rightarrow s^{(2)} \rightarrow s^{(1)} \rightarrow o,$$

with higher layers encoding slower-changing invariants.

This corresponds directly to the hierarchical evaluation stack:

$$J^{(3)} \rightarrow J^{(2)} \rightarrow J^{(1)}.$$

Higher-level invariants constrain lower-level scenario evaluation by functioning as structural priors over admissible state transitions. Let \mathcal{H} denote a hierarchy of representational layers indexed by $k = 0, 1, \dots, K$, where lower indices correspond to fine-grained sensory or semantic states and higher indices correspond to abstract integrative models. For each layer k , let \mathcal{S}_k denote the space of admissible configurations and let $\Pi_k : \mathcal{S}_{k-1} \rightarrow \mathcal{S}_k$ denote the projection or aggregation operator.

A higher-level invariant may be represented as a constraint functional

$$\mathcal{I}_k : \mathcal{S}_k \rightarrow \mathbb{R},$$

such that admissible lower-level transitions $\sigma_{k-1} \mapsto \sigma'_{k-1}$ must satisfy

$$\mathcal{I}_k(\Pi_k(\sigma'_{k-1})) \leq C_k,$$

for prescribed bounds C_k . In this sense, higher-order structure restricts the branching factor of lower-level scenario trees. The effective hypothesis space available to fine-grained evaluation is pruned by invariant structure imposed from above.

Stable cognition therefore presupposes stability of higher-level invariants. If the environment or infrastructure induces rapid, high-amplitude fluctuations at lower semantic layers, then the induced projection Π_k becomes highly variable. To maintain coherence, higher-level inference must increase update frequency and precision, thereby raising metabolic and computational cost. Formally, if the entropy rate at layer $k - 1$ is denoted h_{k-1} and the compensatory inference cost at layer k is $C_k(h_{k-1})$, then under broad regularity conditions one expects $C'_k > 0$, so that increasing lower-layer entropy induces superlinear cost at higher layers.

Bounding entropy at the substrate level therefore reduces the required corrective gain in hierarchical predictive stabilization. If substrate-level entropy production satisfies

$$\dot{S}_{\text{sub}} \leq \epsilon_{\text{max}},$$

then the induced entropy flux into cognitive layers is likewise bounded, permitting stable recursive estimation without divergence of inference cost.

Cognition may be modeled as hierarchical model selection under entropy constraints. In reinforcement learning, scenario evaluation is formalized through maximization of expected cumulative reward,

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid s_0 = s \right],$$

where policy π selects branches of a decision tree. In predictive processing, inference is framed as minimization of variational free energy,

$$\mathcal{F} = \mathbb{E}_q[\log q(z) - \log p(x, z)],$$

which penalizes divergence between generative models and observed data.

Both frameworks presuppose relatively stable generative structure. If the underlying transition kernel or environmental statistics fluctuate without bound, value estimates fail to converge and free-energy gradients become unstable. The branching factor of the scenario tree grows effectively without limit, impairing learning and inducing pathological oscillation.

Entropy-bounded infrastructure aligns with the stability requirements implicit in these learning frameworks. By constraining the rate at which the environment injects

disorder into lower representational layers, the substrate enforces regularity conditions under which value iteration converges and variational inference remains well-posed.

Manufacturing stability at the substrate level thus propagates upward as bounded cognitive entropy at the agent level. Structural invariants embedded in infrastructure reduce the inferential burden placed on individual agents. Stability of collective architecture becomes a precondition for stability of individual cognition.

54 Embodiment, Environmental Modification, and the Ontology of Trace

54.1 Organisms as Environmental Rewriters

Let \mathcal{E} denote the physical environment and let \mathcal{O} denote an organism embedded within it. At any time t , the joint state may be written as $(o(t), e(t)) \in \mathcal{O} \times \mathcal{E}$.

An organism does not merely represent \mathcal{E} internally; it modifies \mathcal{E} according to its internal model. Let h denote the organism’s frame or model of the ecosystem. Behavior then induces a transformation

$$\Phi_h : \mathcal{E} \rightarrow \mathcal{E},$$

such that

$$e(t+1) = \Phi_h(e(t)).$$

Over time, repeated application of Φ_h reshapes the environment to better match the organism’s predictive structure.

This is true at multiple scales. Cells reorganize protoplasm and extracellular matrix. Beavers construct dams. Humans construct clothing, homes, roads, and institutions. In each case, the organism performs work on the environment to reduce mismatch between expectation and constraint.

Formally, let $D(e, h)$ measure divergence between environmental structure and internal model. Organismal action tends to minimize D subject to energetic constraints:

$$\frac{d}{dt} D(e(t), h) \leq 0,$$

while expending metabolic energy.

Thus embodiment is not optional. It is the mechanism by which internal models

acquire physical reality.

54.2 Trace as Energy Expenditure

Let W denote physical work performed by the organism. Any durable environmental modification satisfies

$$W = \int F \cdot dx,$$

and produces a structural change Δe .

A photograph of a shed, a painting, a house, or even a footprint is evidence of $\Delta e \neq 0$, which in turn implies $W > 0$.

Digital artifacts are records of prior physical work. A photograph of a location implies locomotion; locomotion implies metabolic expenditure. A written essay implies neural activation, muscular movement, and energetic cost.

Thus a digital trace is not an autonomous entity but a compressed encoding of embodied work.

54.3 The Non-Equivalence of Trace and Agency

Let σ denote a semantic sphere representing a recorded profile state. While σ preserves identity, modalities, and provenance, it lacks an active metabolic operator.

Agency requires ongoing state evolution:

$$\dot{o}(t) = f(o(t), e(t)).$$

After biological death, $\dot{o}(t) = 0$ in any physically meaningful sense. The semantic sphere σ may persist as static data, but it no longer participates in feedback with \mathcal{E} .

Therefore, maintaining a profile as if it were equivalent to a living agent conflates trace with operator.

A trace is a boundary condition on future evolution of others' cognition. It may influence trajectories γ_i in other minds. However, it does not generate new trajectories from its own internal metabolic process.

54.4 Embodiment as Ontological Constraint

Claims about the world require interaction with the world. Let C be a claim operator mapping internal state to external assertion. For C to be meaningful, it must couple to verification channels grounded in \mathcal{E} .

If an entity is entirely virtual and cannot modify or be modified by \mathcal{E} except through pre-existing recorded traces, then it lacks independent ontological grounding.

In formal terms, a fully virtual agent without physical coupling satisfies

$$\frac{\partial e}{\partial o} = 0.$$

Such an entity cannot perform work. It cannot reduce divergence $D(e, h)$. It cannot update its model through embodied interaction.

Hence it cannot sustain genuine epistemic authority.

54.5 Profiles as Ghosts of Work

Digital memorialization systems treat profiles as persistent agents. However, a profile is a static encoding of previous state transitions.

It is a ghost not in the metaphysical sense but in the thermodynamic sense: it is a residual structure left after work has been dissipated.

Attempting to treat the trace as equivalent to the operator confuses energy history with active metabolism.

Embodiment is the condition of possibility for meaningful claims. Without ongoing energy exchange and environmental coupling, a profile cannot maintain opinions, revise models, or generate authentic commitments. It can only be reinterpreted by the living.

55 Thermodynamic Cost of Trace and the Landauer Constraint

55.1 Information as Physical Work

Landauer's principle establishes that erasing one bit of information requires a minimum energy dissipation of

$$E_{\min} = k_B T \ln 2,$$

where k_B is Boltzmann's constant and T is temperature.

Information is therefore not abstract; it is physically instantiated state.

Let σ denote a semantic sphere containing n bits of effective information. Its creation required at least

$$W \geq nk_B T \ln 2$$

of physical work, ignoring inefficiencies.

Thus every digital trace encodes a lower bound on metabolic expenditure.

Writing, recording, photographing, constructing each is a thermodynamically irreversible act. The trace persists because energy was dissipated to stabilize it.

55.2 Profiles as Entropic Residues

Let $\Sigma(t)$ denote the cumulative semantic output of an organism over its lifetime. Each state transition in Σ corresponds to irreversible physical computation.

Upon biological termination at time t_d , the organism's metabolic process ceases:

$$\dot{\sigma}(t) = 0 \quad \text{for } t > t_d.$$

However, $\Sigma(t_d)$ remains encoded in storage media.

This persistence does not imply continued agency. It implies only that prior energy dissipation created durable low-entropy structure.

The profile is therefore an entropic residue: a metastable configuration of matter storing compressed history of past work.

It does not continue to perform work.

55.3 Epistemic Authority and Energetic Coupling

A claim about the world requires capacity to update under new evidence. Updating requires physical computation. Physical computation requires energy dissipation.

If an entity cannot dissipate new energy,

$$\Delta W = 0,$$

then it cannot revise its internal model.

Therefore it cannot maintain epistemic authority over evolving claims.

A static profile cannot sustain opinion. It can only be interpreted.

Landauer's bound makes explicit that meaning production is inseparable from energy flow.

56 Virtualization, Chipification, and Substrate Dependence

56.1 Organizational Surjection onto Silicon

Consider a physical system S with state space \mathcal{X}_S and transition function

$$F_S : \mathcal{X}_S \rightarrow \mathcal{X}_S.$$

Chipification consists in constructing a computational system C with state space \mathcal{X}_C and transition function

$$F_C : \mathcal{X}_C \rightarrow \mathcal{X}_C,$$

together with a surjective mapping

$$\pi : \mathcal{X}_S \twoheadrightarrow \mathcal{X}_C$$

such that

$$\pi \circ F_S \approx F_C \circ \pi.$$

The silicon device replicates the organizational dynamics of S .

This is true of early console emulation, industrial control systems, and abstract computation generally. Organizational structure becomes executable logic.

56.2 Virtual Machines and Physical Substrate

Let C be a computing architecture with transition function

$$F_C : \mathcal{S}_C \rightarrow \mathcal{S}_C,$$

defined over a finite or countable state space \mathcal{S}_C . Suppose C is implemented as a virtual machine within a host architecture C' , whose physical transition function is

$$F_{C'} : \mathcal{S}_{C'} \rightarrow \mathcal{S}_{C'}.$$

Let $\iota : \mathcal{S}_C \hookrightarrow \mathcal{S}_{C'}$ be an injective encoding of guest states into host states. The virtualized transition function F_{VM} is defined by

$$F_{VM} = \iota^{-1} \circ F_{C'} \circ \iota,$$

on the image of ι . Equivalently,

$$F_{VM} = \iota^{-1} \circ F_B \circ \iota,$$

where F_B denotes the base physical transition function realized by transistor switching at the hardware level.

Thus virtualization is a conjugation of dynamical systems under encoding. The operational semantics of the guest system are preserved up to isomorphism of state representation. However, the dynamics remain grounded in physical state transitions of the base substrate. Every evaluation step of the virtual machine corresponds to a finite sequence of switching events in F_B .

Virtualization therefore does not eliminate embodiment. It constitutes a change of coordinates in state space, mapping one organizational algebra onto another substrate without altering its formal structure. The underlying energetic and thermodynamic constraints remain operative.

This perspective clarifies the persistence of organizational forms across technological epochs. Early console architectures, defined by fixed transition algebras in silicon, do not vanish when hardware production ceases. Their transition systems are emulated by software interpreters executing on general-purpose processors. File-sharing protocols once instantiated in peer-to-peer infrastructures may be re-instantiated within new network topologies. Template-driven page systems reproduce graph structures analogous to earlier social platforms. In each case, the organizational invariant migrates across substrates; it does not dematerialize.

Organization is conserved under suitable encodings. Material realization changes; transition algebra persists.

56.3 Impossibility of Purely Disembodied Agency

Let V denote a hypothetical entity consisting solely of stored semantic configuration σ_V without active substrate execution. Suppose that for all t ,

$$\dot{\sigma}_V(t) = 0.$$

Then V admits no nontrivial transition function. It is not a dynamical system but a static configuration.

A dynamical system requires an evolution law

$$\sigma(t + 1) = F(\sigma(t)),$$

implemented by physical state transitions. In the absence of active substrate dynamics, the mapping F is not executed. Code without execution is a syntactic object; it does not constitute an operative process.

A virtual machine is meaningful only insofar as it is instantiated by ongoing computation on hardware. Likewise, a digital persona or profile is semantically active only through the interpretive or computational acts of embodied agents. Stored traces represent prior dynamical processes; they are not equivalent to continued agency.

Accordingly, the claim that a configuration can persist as an agent without active execution conflates static state with dynamical process. Virtualization preserves organizational form under conjugation of state spaces, but it does not abolish substrate dependence.

Chipification is therefore not dematerialization. It is a surjective abstraction in which organizational structure is mapped onto a different physical base via encoding. Every computation, emulation, and execution event presupposes energetic expenditure and thermodynamic grounding. There exists no ontology of purely disembodied software; there exist only redistributed embodiments across physical substrates.

57 Swarm Dynamics, Stigmergy, and Energetic Trace

57.1 Stigmergy as Distributed Environmental Writing

Let \mathcal{E} denote a shared environment and let $\{o_i\}_{i=1}^N$ denote a population of agents.

Each agent performs local updates:

$$e_{t+1} = \Phi_{o_i}(e_t),$$

where Φ_{o_i} is a transformation induced by agent i .

In stigmergic systems, coordination arises not through centralized planning but through environmental modification. Ant foraging provides the canonical example: pheromone deposition modifies the environment, which in turn biases subsequent trajectories.

Formally, let $P(x, t)$ denote pheromone concentration. Ant trajectories follow gradients:

$$\dot{x}_i(t) \propto \nabla P(x_i(t), t).$$

Pheromone deposition satisfies

$$\frac{\partial P}{\partial t} = \sum_i \delta(x - x_i(t)) - \lambda P,$$

where λ represents evaporation.

The environment becomes a distributed memory of collective work.

57.2 Human Infrastructure as Scaled Stigmergy

Human-built systems generalize this mechanism. Roads, railways, and communication networks are stabilized pheromone trails. They encode accumulated work.

Let $I(x)$ denote infrastructure density. High I reduces transport cost:

$$C(x \rightarrow y) = C_0 - \alpha I,$$

for $\alpha > 0$.

Thus previous energetic expenditure modifies future state transitions.

Cities, factories, and digital repositories are stigmergic scaffolds structured traces

that bias collective trajectories.

A digital post is analogous to pheromone: it modifies informational density. However, unlike biological pheromone, digital traces do not decay unless explicitly designed to do so.

Absent decay, trace concentration becomes curvature concentration.

57.3 Energetic Residue and Collective Memory

Let $\Sigma(t)$ denote cumulative trace density. In swarm systems, stability requires bounded accumulation:

$$\frac{d}{dt}\Sigma(t) \leq \epsilon_{\max}.$$

If accumulation exceeds dissipation, the system freezes into overconcentrated attractors.

This explains why entropy-bounded design parallels pheromone evaporation: it prevents runaway curvature concentration.

Digital immortality systems eliminate evaporation. They convert finite energetic traces into infinite persistence. This breaks the stigmergic equilibrium. Swarm stability requires decay.

58 Extraction as Curvature Capture in Energetic Substrates

58.1 From Trace to Toll Booth

Infrastructure emerges from accumulated energetic trace. However, once stabilized, it can be captured.

Let T denote a pipeline mapping inputs to outputs. In open infrastructure,

$$T : \mathcal{X} \rightarrow \mathcal{Y}$$

remains composable.

In extractive architecture, a central operator C inserts itself:

$$T' = T_2 \circ C \circ T_1.$$

C imposes cost or captures surplus from every traversal.

Geometrically, C creates a scalar well \mathcal{P} in the manifold (\mathcal{M}, g) .

Repeated traversal increases local flow density, reinforcing curvature.

58.2 Energetic Interpretation of Extraction

Extraction corresponds to redirecting distributed energetic trace into centralized accumulation.

Let W_i denote work performed by agents i . In non-extractive systems, benefits distribute proportionally.

In extractive systems,

$$W_{\text{agents}} \rightarrow \Delta\Phi_C,$$

where Φ_C is scalar potential of the central node.

The energetic residue of collective work becomes concentrated curvature.

Thus extraction is not merely economic asymmetry. It is geometric concentration of dissipated energy.

58.3 Entropy-Bounded Architecture as Evaporation Mechanism

To prevent curvature concentration, substrate rules must enforce dissipation or redistribution.

Entropy budgets ensure that semantic accumulation does not grow without bound:

$$E(t) \leq E(0) + \epsilon_{\max}t.$$

Crystal decay ensures historical influence diminishes without continued contribution:

$$TiC(t) = e^{-\lambda t}q.$$

These mechanisms are architectural analogues of pheromone evaporation.

They prevent static trace from becoming eternal gravitational wells.

58.4 No Escape from Substrate

Virtualization does not eliminate energetic grounding. Every emulation requires physical execution.

Thus any platform claiming to host immortal agents must continuously expend real energy to maintain their computational simulation.

Immortality without execution is inert data. Immortality with execution is ongoing metabolic process funded by someone else's energy.

There is no zero-cost persistence.

Extraction emerges when the cost of maintaining simulation is externalized while control remains centralized.

Therefore substrate design is political economy encoded as physics.

Manufacturing stability requires embedding decay, bounded accumulation, and compositional openness into the execution layer itself.

59 A Decay Necessity Theorem for Stigmergic Stability

59.1 Model Setup

Let \mathcal{E} be a shared environment and let $\Sigma(t)$ denote the cumulative trace density deposited by a population of agents.

Assume trace evolves according to

$$\frac{d}{dt}\Sigma(t) = A(t) - D(\Sigma(t)),$$

where $A(t) \geq 0$ represents aggregate deposition rate and $D(\Sigma)$ represents dissipation.

Suppose $A(t)$ is bounded below by a positive constant $a > 0$ over long time intervals.

59.2 No-Decay Case

If $D(\Sigma) \equiv 0$, then

$$\Sigma(t) = \Sigma(0) + \int_0^t A(s)ds.$$

If $A(t) \geq a > 0$, then

$$\Sigma(t) \rightarrow \infty \quad \text{as } t \rightarrow \infty.$$

Thus trace accumulates without bound.

In stigmergic systems where agent trajectories follow gradients of Σ , i.e.,

$$\dot{x}_i(t) \propto \nabla \Sigma(x_i(t), t),$$

unbounded accumulation produces increasing curvature concentration and geodesic collapse into dominant attractors.

59.3 Decay Condition for Stability

Assume dissipation satisfies

$$D(\Sigma) = \lambda \Sigma,$$

with $\lambda > 0$.

Then

$$\frac{d}{dt} \Sigma(t) = A(t) - \lambda \Sigma(t).$$

If $A(t)$ is bounded above by A_{\max} , then $\Sigma(t)$ converges to a finite equilibrium:

$$\Sigma^* = \frac{A_{\text{avg}}}{\lambda}.$$

59.4 Theorem

Decay Necessity Theorem. In any stigmergic system with persistent deposition $A(t) \not\rightarrow 0$, bounded trace density and avoidance of curvature singularity require strictly positive dissipation $D(\Sigma) > 0$.

59.5 Proof Sketch

If $D(\Sigma) = 0$, trace diverges under persistent deposition, producing unbounded curvature concentration.

If $D(\Sigma)$ grows at least linearly in Σ , the system admits a globally attracting fixed point.

Therefore decay is necessary for long-run stability.

59.6 Interpretation

Evaporation in ant colonies is not incidental; it is structurally required.

Analogously, entropy budgets and crystal decay are not aesthetic features of substrate design but necessary conditions for preventing informational singularities.

A platform without decay mechanisms must accumulate curvature until collapse or capture occurs.

60 Cognition, Embodiment, and Manufactured Stability

Organisms modify their environment to approximate their internal models. Cells reorganize extracellular matrix. Animals construct nests. Humans construct dwellings, garments, machines, and institutions.

Each such modification is a physical rewriting of the world.

Cognition generates hypothetical futures and evaluates them. Embodied action selects one trajectory and inscribes it into the environment through energetic work. The resulting trace stabilizes certain transitions for future agents.

Thus cognition and infrastructure are continuous across scale.

Internally, evaluation reduces cognitive entropy by collapsing branching scenarios. Externally, construction reduces environmental uncertainty by embedding stabilized pipelines.

Stigmergic coordination emerges when many agents write to the same environment. The environment becomes a distributed memory of prior work. However, without decay, accumulated trace concentrates curvature and collapses diversity.

Swarm stability therefore requires bounded accumulation.

Extraction arises when trace concentration is captured by a central operator that inserts itself into all pipelines. This transforms distributed stigmergy into curvature singularity.

Virtualization does not abolish embodiment. It surjectively maps organization onto new physical substrates. Every computation, emulation, and profile requires active energy dissipation.

A digital trace is the entropic residue of embodied work. It is not an autonomous agent.

Manufacturing stability therefore requires three coupled constraints:

First, embodiment must remain acknowledged as the ground of agency. Second, trace accumulation must remain bounded by decay. Third, infrastructure must preserve compositional openness rather than funneling flow into singular wells.

Cognition manufactures stability locally. Infrastructure manufactures stability collectively. When designed without decay and bounded accumulation, infrastructure destabilizes cognition by inflating branching and concentrating curvature.

A stable digital civilization must therefore embed thermodynamic realism into its substrate.

Stability is not preservation of all trace. It is disciplined dissipation that prevents singularity.

61 Curvature Singularity and Extraction

61.1 Flow Concentration on an Information Manifold

Let (\mathcal{M}, g) be the information manifold and let $\rho(x, t)$ denote density of agent flow through point $x \in \mathcal{M}$.

Assume evolution satisfies a continuity equation:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho v) = 0,$$

where v is velocity field induced by platform architecture.

Suppose there exists a scalar potential Φ such that

$$v = -\nabla \Phi.$$

Repeated traversal reinforces Φ through feedback:

$$\frac{\partial \Phi}{\partial t} = \alpha \rho - \lambda \Phi,$$

with reinforcement parameter $\alpha > 0$ and decay $\lambda \geq 0$.

61.2 No-Decay Instability

If $\lambda = 0$, then

$$\frac{\partial \Phi}{\partial t} = \alpha \rho.$$

Regions of high ρ produce increasing Φ , which increases $|v|$, which increases ρ further.

Linearizing around equilibrium yields

$$\delta \dot{\rho} \approx \alpha \Delta \rho.$$

If α exceeds diffusion stabilization, perturbations grow.

61.3 Theorem

Curvature Singularity Theorem. In an information manifold with reinforcement $\alpha > 0$ and insufficient dissipation $\lambda \leq \lambda_c$, there exists finite time T such that curvature

$$K(x, t) \rightarrow \infty$$

in a localized region $\mathcal{P} \subset \mathcal{M}$, producing geodesic convergence and flow capture.

61.4 Proof Sketch

Positive feedback between ρ and Φ produces nonlinear growth. In absence of balancing dissipation, the coupled system behaves analogously to gravitational collapse or chemotactic aggregation.

Standard blow-up results for reaction-diffusion systems imply finite-time singularity under supercritical reinforcement.

61.5 Interpretation

Within the geometric framework developed in the preceding sections, extraction is most precisely characterized as curvature blow-up. A dominant platform does not merely increase its market share in a conventional economic sense; it alters the metric structure of the interaction manifold. By concentrating informational density and reinforcing preferential attachment, it induces a geometric phase transition in

which sectional curvature becomes strictly positive over an expanding region. The resulting curvature concentration forces geodesic convergence, collapsing a multiplicity of potential trajectories into a single attractor basin.

This phenomenon is not metaphorical. As established via the Jacobi field analysis, sufficiently positive curvature guarantees reconvergence of initially diverging paths within finite parameter distance. In informational terms, this corresponds to the homogenization of behavioral flows and the suppression of alternative semantic directions. Diversity is not explicitly prohibited; it is geometrically disfavored.

Avoidance of singularity therefore requires controlled dissipation and bounded reinforcement. In the absence of decay, stigmergic accumulation produces runaway scalar amplification. If informational density is permitted to increase without constraint, the scalar field deepens, curvature intensifies, and the manifold approaches a singular regime in which effective dimensionality collapses.

Entropy budgets and crystal decay function as curvature regulators within this model. Entropy constraints bound the rate of scalar accumulation, preventing superlinear growth of informational density. Exponential decay in the Time Crystal valuation prevents indefinite persistence of early attractors, enforcing dissipation of influence absent continued structural contribution. Together, these mechanisms implement a distributed Ricci-like regularization: curvature cannot diverge without violating conservation constraints embedded in the substrate.

Stability, in this sense, is not the elimination of dynamics but the prevention of geometric collapse. It is the maintenance of a manifold whose curvature remains globally bounded, preserving the existence of multiple admissible geodesics and thereby sustaining structural plurality.

62 Manufacturing Stability

Stability is not an intrinsic property of complex systems; it is an achievement. Across biological, ecological, and civilizational scales, stability emerges through structured reorganization of flows, constraints, and transformations. It is produced by active processes that regulate accumulation, redistribute energy, and encode invariants into material form.

At the cellular level, organisms reorganize protoplasmic matter to stabilize metabolic flux. Autocatalytic networks selectively reinforce reaction pathways that maintain homeostasis, while dissipative processes remove excess gradients. The cell does not passively endure entropy; it locally modulates it. The resulting stability is conditional

and dynamic, maintained through continuous exchange with its environment.

At the ecological level, collective systems generate coordination through stigmergic inscription. Ant colonies deposit pheromone gradients that encode distributed memory into the substrate. Swarm search algorithms amplify productive paths while allowing unused traces to decay. Without evaporation, signal saturation would eliminate directional information. Stability therefore depends on disciplined dissipation as much as on reinforcement.

At the civilizational level, humans externalize cognitive models into infrastructure. Roads stabilize transportation pathways; factories stabilize transformation pipelines; languages stabilize symbolic exchange; institutions stabilize norms of coordination. Each of these constructions embodies accumulated traces of prior decisions. Every cognitive act generates a space of hypothetical futures; every embodied action selects among them and inscribes that selection into the material world. Over time, these inscriptions compose into durable structures.

Such structures may remain modular and compositional, preserving openness under bounded accumulation. Alternatively, they may concentrate flow and increase effective curvature within the underlying system, producing singularities of control. In geometric terms, infrastructure shapes the manifold on which collective trajectories evolve. If reinforcement mechanisms operate without regulated decay, stigmergic accumulation leads to curvature concentration. Geodesics converge; alternatives diminish.

Stability is therefore neither equivalent to permanence nor reducible to maximal preservation. Indefinite retention without dissipation yields instability in another form: congestion, rigidity, or capture. Disciplined decay is a structural precondition of sustained coherence. In thermodynamic language, stability corresponds to controlled entropy production, not to its elimination.

Embodiment plays an essential role in this process. Traces are records of work performed. A constructed object, a cultivated field, or a written proof registers physical transformation of matter and energy. When detached from the capacity for further action, such traces become inert residues rather than active participants in stabilization dynamics. Infrastructure is meaningful only insofar as it remains embedded in ongoing processes of transformation.

To manufacture stability is therefore to acknowledge thermodynamic cost, to enforce bounded accumulation, and to preserve compositional openness. It is to design execution layers in which invariants are structural rather than aspirational. In such systems, extraction is not merely inequitable distribution; it is the formation of curvature without compensating dissipation. Singularities arise when accumulation

proceeds without constraint.

Manufacturing stability thus requires embedding conservation principles into the substrate itself. The objective is not universal preservation, nor indefinite amplification of engagement metrics, nor archival immortality. The objective is to construct environments in which trajectories remain differentiable, in which flows remain distributed, and in which accumulation is bounded by invariant conditions. Stability is achieved not by resisting change, but by structuring it.

63 Conclusion: Stability as an Architectural Property

This paper began from a politiceconomic diagnosis and arrived at a geometric and thermodynamic reconstruction of digital infrastructure. The initial claim was descriptive: contemporary dominant platforms no longer function primarily as catalytic intermediaries but as extractive attractors. Through behavioral prediction, engagement optimization, and infrastructural entrenchment, they concentrate curvature on the informational manifold. The formal development translated that diagnosis into a precise structural instability. Extraction is not merely inequitable distribution of surplus. It is the emergence of scalar wells and positive curvature concentration within a shared semantic space, generating geodesic convergence, predictive asymmetry, entropy inflation, and long-run loss of stationarity.

The geometric reformulation allowed several otherwise diffuse concerns to be unified. Predictive leverage became measurable as KullbackLeibler divergence between agent and platform models. Behavioral steering became channel capacity. Herding became Jacobi field focusing under positive curvature. Lock-in became manifold collapse under metric pullback. Attention capture became a special case of conditional-entropy maximization. The language of curvature and entropy did not replace political economy; it rendered its mechanisms structurally explicit. The result was not metaphor but equivalence: the same instabilities described sociologically could be derived from formal properties of dynamical systems defined over information manifolds.

From this structural account followed a central incompatibility theorem. Engagement-maximizing dynamics are thermodynamically unstable when unconstrained. Either entropy diverges, producing semantic heat death, or curvature concentrates, producing monopolistic attractors. No third stable regime exists under unconstrained optimization. This result reframes contemporary platform pathology not as moral failure or insufficient competition alone, but as a predictable phase outcome of a specific

objective function.

The response developed here does not attempt to regulate particular outcomes ex post. It modifies the algebra of permissible transformations ex ante. By introducing typed semantic spheres with intrinsic identity, measurable entropy, and explicit provenance graphs, the substrate internalizes structural constraints. Entropy budgets bound cumulative disorder. Staking and slashing enforce declared limits. Texture valuation rewards non-redundant informational contribution rather than positional amplification. Time valuation imposes exponential decay on influence, preventing indefinite accumulation of advantage. Sheaf-theoretic closure ensures that local semantic assignments extend coherently to global structure. Governance is recast as a higher-order category of invariant-preserving morphisms, preventing protocol evolution from reintroducing extractive regimes through administrative relaxation.

The unifying theme across these constructions is conservation. Entropy is conserved within bounded budgets. Crystal mass cannot be generated without genuine novelty. Influence decays unless renewed through contribution. Coherence must glue across contexts. These conservation laws play the role that structural antitrust once sought to play at the regulatory layer. Instead of prohibiting behaviors after they emerge, the substrate prevents their algebraic expression. Extraction becomes ill-typed rather than merely illegal.

A second unifying theme is stationarity. Under entropy-bounded and curvature-bounded dynamics, the informational system admits invariant probability measures. Long-run semantic distributions stabilize. Engagement-driven systems do not admit such measures; they are dynamically non-stationary by construction. The existence of invariant measures is not merely mathematically elegant. It corresponds to the possibility of sustained collective equilibrium in which meaning accumulates without runaway divergence or singular concentration.

A third theme is autonomy. The formal notion of semantic autonomy articulated in the body of the paper unites geometric integrity, epistemic parity, and temporal fairness. When curvature is bounded, geodesics reflect intrinsic preference geometry rather than externally imposed distortion. When predictive asymmetry is bounded, agents are not expropriated by models that know them better than they know themselves. When influence decays, early advantage does not harden into permanent hierarchy. These conditions are not subjective states but structural properties of the substrate. Autonomy emerges not from exhortation but from invariant enforcement.

The later sections extended the analysis beyond purely commercial engagement optimization to regulatory architectures that risk deepening scalar wells under the banner of safety. The same structural language applies. Compliance-driven surveillance,

identity verification mandates, and secondary AI exploitation of private data can all be modeled as exogenous injections into the scalar field. The framework therefore offers not only a critique of corporate extraction but also a lens through which to evaluate policy proposals. Any intervention that increases curvature concentration or injects entropy without bounded conservation contributes to instability, regardless of its stated intent.

Throughout, the argument has insisted on a constraint-first design principle. The prevailing technological ethos has favored optimization-first architectures, subsequently constrained by policy or moderation layers. This sequence inverts priority. It asks which invariants must hold for stability, autonomy, and epistemic diversity to be structurally possible, and then embeds those invariants into the computational substrate itself. Constraint is not an afterthought but the precondition of sustainable growth. In thermodynamic language, bounded entropy production permits order; in geometric language, bounded curvature preserves manifold diameter; in economic language, conservation laws prevent rent extraction.

Several open problems remain. Empirical calibration of curvature and entropy thresholds requires substantial longitudinal study. Welfare-theoretic alignment between crystal equilibria and social optimum requires formal game-theoretic analysis. The dynamic topology of contexts demands a fully evolving sheaf theory. Implementation at scale requires advances in verifiable computation, decentralized storage, and governance formal verification. These challenges delimit the frontier of research rather than undermine the structural thesis.

The central conclusion, however, is clear. Stability in digital infrastructure is not achieved through exhortation, competition alone, or after-the-fact enforcement. It is achieved through architecture. When semantic evolution is governed by conserved quantities, when predictive leverage is bounded, when influence decays, and when governance itself is subject to invariant-preserving morphisms, extractive regimes become unreachable states of the system. The informational field ceases to be a terrain for curvature capture and becomes instead a plenum in which collaboration is dynamically sustainable.

Manufacturing stability, in this sense, is not the imposition of rigidity. It is the deliberate shaping of phase space so that destructive attractors cannot form. The architecture does not eliminate dynamism; it constrains it within thermodynamically coherent bounds. In doing so, it reframes digital political economy as a problem in geometric field design. The question is no longer how to redistribute surplus after extraction, but how to construct a substrate in which extraction cannot arise without violating the algebra that defines the system itself.

Appendices

Worked Numerical Example

We illustrate the formal machinery with a concrete instance.

A.1 Initial Sphere

Let $\sigma_0 = (I_0, \{\text{text}\}, M_0, E_0, S_0)$, where $M_0(\text{text})$ encodes a formal theorem statement in first-order logic. Suppose the entropy of the token distribution over the statement is $E(\sigma_0) = 3.8$ bits.

A.2 Rule Application 1: Text to Proof

Apply rule $r_1 : \sigma \mapsto \sigma'$ that adds a **proof** modality to produce a sphere with $T_1 = \{\text{text}, \text{proof}\}$. The declared entropy budget is $\epsilon_{r_1} = 0.6$ bits, reflecting the informational content introduced by the proof.

After application, suppose $E(\sigma_1) = 4.2$ bits. Entropy soundness requires $4.2 \leq 3.8 + 0.6 = 4.4$, which is satisfied.

The mutual information between the **text** and **proof** modalities is $MI = 1.3$ bits, reflecting shared logical structure. The entropy of σ_1 is $H(\sigma_1) = 4.2$ bits. Hence

$$TC(\sigma_1) = 4.2 - 1.3 = 2.9 \text{ Texture Crystals.}$$

A.3 Rule Application 2: Proof to Counterexample

Two days later (with $\lambda = 0.1$ per day, so $e^{-\lambda \cdot 2} = e^{-0.2} \approx 0.819$), apply rule r_2 that adds a **counterexample** modality refuting the theorem, with $\epsilon_{r_2} = 0.4$ bits. Suppose $E(\sigma_2) = 4.5$ bits, satisfying $4.5 \leq 4.2 + 0.4 = 4.6$.

Let $q(r_1) = 0.9$ and $q(r_2) = 0.85$. The Time Crystal valuation at the moment after r_2 is applied (taking $t_{\text{now}} = t_2$, so r_2 contributes $e^0 q(r_2) = 0.85$) is:

$$TiC(\sigma_2) = e^{-0.2} \cdot 0.9 + e^0 \cdot 0.85 \approx 0.737 + 0.85 = 1.587 \text{ Time Crystals.}$$

A.4 Interpretation

The example illustrates entropy boundedness across two sequential applications: $E(\sigma_2) = 4.5 \leq 3.8 + 0.6 + 0.4 = 4.8 = E(\sigma_0) + \Phi(\tau)$. It illustrates redundancy taxation: the sphere’s TC is reduced by mutual information between modalities, penalizing logically related content that appears in multiple modalities. And it illustrates decay: the first rule application’s contribution to TiC has decayed to 81.9% of its original value at the time of the second application. A sphere that remains unextended will see its TiC approach zero exponentially, in contrast to the permanent follower counts and accumulated attention metrics of extractive platforms.

A.5 Predictive Asymmetry in the Numerical Example

We extend the example to illustrate the predictive asymmetry framework of Section 3. Suppose the agent’s self-model assigns probability $\hat{p}(\text{accept}) = 0.7$ to acceptance of the original theorem and $\hat{p}(\text{reject}) = 0.3$ to rejection. Suppose the platform’s model, trained on the full behavioral history of the agent’s prior theorem evaluations, assigns $q(\text{accept}) = 0.55$ and $q(\text{reject}) = 0.45$, which is closer to the true distribution $p(\text{accept}) = 0.52$ and $p(\text{reject}) = 0.48$.

Then

$$D_{\text{KL}}(p||\hat{p}) = 0.52 \log \frac{0.52}{0.7} + 0.48 \log \frac{0.48}{0.3} \approx 0.139 + 0.218 = 0.357 \text{ bits},$$

$$D_{\text{KL}}(p||q) = 0.52 \log \frac{0.52}{0.55} + 0.48 \log \frac{0.48}{0.45} \approx 0.026 + 0.030 = 0.056 \text{ bits},$$

yielding a predictive asymmetry $\Delta = 0.357 - 0.056 = 0.301$ bits. The platform’s predictive advantage corresponds to a 0.3-bit information gain over the agent’s self-model. Under unconstrained behavioral extraction, this gap grows monotonically per Proposition 3.4. Under PlenumHub’s TiC decay, the platform’s predictive models decay in influence unless continuously updated, bounding the growth of Δ .

Appendix B: Geodesic Focusing under Scalar Wells

Let (\mathcal{M}, g) be a complete Riemannian manifold and let $\phi : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth scalar field.

Define the effective metric

$$g_{ij}^{\text{eff}} = g_{ij} + \alpha \nabla_i \phi \nabla_j \phi,$$

for $\alpha > 0$.

Let $\gamma(t)$ be a geodesic of $(\mathcal{M}, g^{\text{eff}})$ satisfying

$$\frac{D\dot{\gamma}}{dt} = 0.$$

The Jacobi equation for a variation field J along γ is

$$\frac{D^2 J}{dt^2} + R^{\text{eff}}(J, \dot{\gamma})\dot{\gamma} = 0.$$

If $\Delta_g \phi > 0$ in a neighborhood U , then the Ricci curvature of g^{eff} satisfies

$$\text{Ric}^{\text{eff}}(\dot{\gamma}, \dot{\gamma}) = \text{Ric}(\dot{\gamma}, \dot{\gamma}) + \alpha \nabla^2 \phi(\dot{\gamma}, \dot{\gamma}) + O(\alpha^2).$$

If $\nabla^2 \phi$ is positive definite in U , then there exists $\kappa_0 > 0$ such that

$$\text{Ric}^{\text{eff}}(\dot{\gamma}, \dot{\gamma}) \geq \kappa_0 |\dot{\gamma}|^2.$$

By the Bonnet–Myers theorem,

$$\text{diam}(\mathcal{M}, g^{\text{eff}}) \leq \frac{\pi}{\sqrt{\kappa_0}}.$$

Hence geodesics focus within finite parameter length.

Appendix C: Entropy Growth under Engagement Maximization

Let $\{\sigma_n\}_{n \geq 0}$ be generated by

$$\sigma_{n+1} = U(\sigma_n),$$

where

$$U(\sigma) = \arg \max_{\sigma' \in \mathcal{N}(\sigma)} J(\sigma').$$

Assume

$$J(\sigma') = H(\sigma') + \beta V(\sigma'),$$

with $\beta > 0$ and V bounded.

If $\sup_{\sigma} H(\sigma) = \infty$, then

$$H(\sigma_n) \rightarrow \infty \quad \text{as } n \rightarrow \infty.$$

If H is bounded by H_{\max} , then maximizing J forces

$$\nabla J = \nabla H + \beta \nabla V = 0$$

only at critical points where $\nabla H = -\beta \nabla V$.

In the absence of entropy constraints, stationary points satisfy

$$\Delta_g \phi \geq 0$$

in localized regions, producing scalar accumulation.

Therefore either

$$\lim_{n \rightarrow \infty} H(\sigma_n) = \infty \quad \text{or} \quad \sup_U \Delta_g \phi \rightarrow \infty.$$

Appendix D: Mutual Information Bounds in Texture Valuation

Let modalities X_1, \dots, X_m be discrete random variables.

Total entropy:

$$H(\sigma) = H(X_1, \dots, X_m).$$

Pairwise mutual information:

$$MI(X_i, X_j) = H(X_i) + H(X_j) - H(X_i, X_j).$$

Texture valuation:

$$TC(\sigma) = H(X_1, \dots, X_m) - \sum_{i \neq j} MI(X_i, X_j).$$

Using subadditivity,

$$H(X_1, \dots, X_m) \leq \sum_{i=1}^m H(X_i).$$

Hence

$$TC(\sigma) \geq \sum_{i=1}^m H(X_i) - \sum_{i \neq j} MI(X_i, X_j) - \sum_{i \neq j} MI(X_i, X_j).$$

If variables are independent,

$$MI(X_i, X_j) = 0 \quad \forall i \neq j,$$

and

$$TC(\sigma) = \sum_{i=1}^m H(X_i).$$

If variables are perfectly redundant,

$$X_1 = \dots = X_m,$$

then

$$H(X_1, \dots, X_m) = H(X_1),$$

and

$$TC(\sigma) = H(X_1) - m(m-1)H(X_1),$$

which is negative for $m > 1$.

Appendix E: Stability of Entropy-Bounded Dynamics

E.1 Continuous-Time Semantic Evolution

Let Σ denote the space of admissible semantic spheres endowed with a suitable Banach or Hilbert structure, and let

$$\sigma : [0, \infty) \rightarrow \Sigma$$

be a continuously differentiable trajectory representing semantic evolution under a deterministic flow

$$\frac{d}{dt}\sigma(t) = F(\sigma(t)),$$

where $F : \Sigma \rightarrow T\Sigma$ is a sufficiently smooth vector field.

Let $E : \Sigma \rightarrow \mathbb{R}_{\geq 0}$ denote the semantic entropy functional as defined in the main

text. Assume E is Frchet differentiable and denote its derivative by $DE(\sigma)$.

The entropy production rate along trajectories is

$$\frac{d}{dt}E(\sigma(t)) = DE(\sigma(t))[F(\sigma(t))].$$

E.2 Entropy-Bounded Regime

Assume the system enforces a global entropy production constraint

$$\frac{d}{dt}E(\sigma(t)) \leq \epsilon_{\max},$$

for some constant $\epsilon_{\max} \geq 0$.

Define the Lyapunov candidate

$$V(t) = E(\sigma(t)).$$

Then

$$\dot{V}(t) = \frac{d}{dt}E(\sigma(t)) \leq \epsilon_{\max}.$$

E.3 Zero-Budget Case

If $\epsilon_{\max} = 0$, then

$$\dot{V}(t) \leq 0.$$

Hence $V(t)$ is non-increasing. In particular,

$$E(\sigma(t)) \leq E(\sigma(0)) \quad \forall t \geq 0.$$

If in addition the flow satisfies LaSalle's invariance principle conditions, then $\sigma(t)$ converges to the largest invariant set contained in

$$\{\sigma \in \Sigma \mid DE(\sigma)[F(\sigma)] = 0\}.$$

Thus zero-budget entropy dynamics admit asymptotically stable equilibria.

E.4 Positive Bounded Budget

If $\epsilon_{\max} > 0$ is finite, then integrating the inequality gives

$$E(\sigma(t)) \leq E(\sigma(0)) + \epsilon_{\max}t.$$

Hence entropy growth is at most linear in time.

In particular,

$$\limsup_{t \rightarrow \infty} \frac{E(\sigma(t))}{t} \leq \epsilon_{\max}.$$

Thus entropy production is uniformly bounded in the Cesro sense, and no finite-time blow-up is possible.

E.5 Exponential Engagement Dynamics

Contrast this with unconstrained engagement optimization. Suppose entropy growth is proportional to its current value:

$$\dot{V}(t) = \alpha V(t), \quad \alpha > 0.$$

Then

$$V(t) = V(0)e^{\alpha t}.$$

Thus entropy exhibits exponential growth.

In particular,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log V(t) = \alpha,$$

which is strictly positive.

Hence entropy grows superlinearly and diverges exponentially.

E.6 Phase Comparison

The two regimes satisfy:

$$\text{Bounded regime: } V(t) = O(t),$$

Engagement regime: $V(t) = \Theta(e^{\alpha t})$.

Therefore the ratio

$$\frac{V_{\text{engagement}}(t)}{V_{\text{bounded}}(t)} \sim \frac{e^{\alpha t}}{t} \rightarrow \infty \quad \text{as } t \rightarrow \infty.$$

The divergence is asymptotically exponential.

E.7 Coupling to Scalar Curvature

Suppose entropy density contributes to scalar potential through

$$\phi(t) = \kappa E(\sigma(t)), \quad \kappa > 0.$$

Let curvature be governed by

$$\mathcal{R}(t) = \beta \Delta_g \phi(t),$$

with $\beta > 0$.

Under bounded entropy growth,

$$\phi(t) \leq \kappa(E(0) + \epsilon_{\max} t),$$

hence curvature grows at most linearly.

Under engagement dynamics,

$$\phi(t) = \kappa E(0) e^{\alpha t},$$

implying exponential curvature amplification.

Positive curvature concentration implies geodesic convergence by the Jacobi field equation:

$$\frac{D^2 J}{dt^2} + R(J, \dot{\gamma}) \dot{\gamma} = 0.$$

If $\mathcal{R}(t)$ grows exponentially, focusing time shrinks asymptotically, and manifold collapse occurs in finite effective time.

E.8 Global Stability Statement

Theorem .1. *Assume entropy production satisfies*

$$\dot{E}(\sigma(t)) \leq \epsilon_{\max}.$$

Then no finite-time entropy singularity occurs and curvature amplification is at most linear.

If instead

$$\dot{E}(\sigma(t)) = \alpha E(\sigma(t)), \quad \alpha > 0,$$

then entropy and curvature diverge exponentially, producing dynamical instability and curvature concentration.

Proof. The entropy bounds follow from Grnwall-type comparison arguments. Curvature bounds follow from linear coupling of ϕ to E and linear dependence of curvature on ϕ . Exponential divergence follows from explicit solution of the linear ODE. \square

E.9 Interpretation

Entropy-bounded dynamics define a linearly stable growth regime with invariant upper envelope. Engagement-driven dynamics define an unstable exponential phase with inevitable concentration or semantic heat death.

Thus entropy budgets are not cosmetic constraints; they alter the phase class of the dynamical system.

Appendix F: Ollivier–Ricci Curvature and Platform Graphs

F.1 Discrete Metric Measure Structure

Let $G = (V, E, w)$ be a finite weighted directed graph representing a platform content network. The vertex set V indexes semantic nodes (e.g., content items, user states, or recommendation states). The edge set $E \subset V \times V$ encodes directed interactions. Each edge $(v, u) \in E$ carries a non-negative weight $w(v \rightarrow u)$ proportional to observed engagement intensity.

We equip G with the shortest-path metric $d : V \times V \rightarrow \mathbb{R}_{\geq 0}$ induced by edge

lengths $\ell(v \rightarrow u) = w(v \rightarrow u)^{-1}$ for $w(v \rightarrow u) > 0$, with the convention that absent edges have infinite length.

For each $v \in V$, define the outgoing probability measure

$$\mu_v(u) = \frac{w(v \rightarrow u)}{\sum_{u'} w(v \rightarrow u')},$$

which is a probability distribution on the neighborhood of v .

Thus $(V, d, \{\mu_v\})$ defines a discrete metric measure space.

F.2 Wasserstein Distance

Let $W_1(\mu_v, \mu_u)$ denote the L^1 Wasserstein (Kantorovich–Rubinstein) distance between probability measures μ_v and μ_u with respect to the metric d :

$$W_1(\mu_v, \mu_u) = \inf_{\pi \in \Pi(\mu_v, \mu_u)} \sum_{x, y \in V} d(x, y) \pi(x, y),$$

where $\Pi(\mu_v, \mu_u)$ is the set of couplings with marginals μ_v and μ_u .

By Kantorovich duality,

$$W_1(\mu_v, \mu_u) = \sup_{\|f\|_{\text{Lip}} \leq 1} \sum_{x \in V} f(x) (\mu_v(x) - \mu_u(x)).$$

F.3 Ollivier–Ricci Curvature

The Ollivier–Ricci curvature along an edge (v, u) is defined by

$$\kappa(v, u) = 1 - \frac{W_1(\mu_v, \mu_u)}{d(v, u)}.$$

This quantity measures how transport between local neighborhoods compares to the metric distance between base points. Positive curvature indicates that neighborhoods are closer (in Wasserstein distance) than their base points; negative curvature indicates expansion.

F.4 Continuous Limit and Ricci Convergence

Assume that G_n is a sequence of graphs whose vertex sets sample a smooth Riemannian manifold (\mathcal{M}, g) with mesh size $h_n \rightarrow 0$, and assume that edge weights approximate a

diffusion kernel consistent with the Laplace–Beltrami operator.

Then, under appropriate regularity and uniform measure conditions,

$$\kappa_{G_n}(v, u) = \frac{1}{2} \text{Ric}_g(\dot{\gamma}, \dot{\gamma}) h_n^2 + o(h_n^2),$$

where γ is the geodesic connecting the corresponding manifold points and $\dot{\gamma}$ its tangent.

Hence,

$$\lim_{n \rightarrow \infty} \frac{\kappa_{G_n}(v, u)}{h_n^2} = \frac{1}{2} \text{Ric}_g(\dot{\gamma}, \dot{\gamma}).$$

This convergence establishes that Ollivier–Ricci curvature provides a discrete proxy for continuous Ricci curvature in the small-scale limit.

F.5 Curvature Concentration on Platform Graphs

Define the empirical curvature concentration parameter

$$\Omega(G) = \max_{v \in V} \left(\frac{1}{\text{deg}(v)} \sum_{u: (v,u) \in E} \kappa(v, u) \right).$$

Let $V_k \subset V$ denote the set of top- k nodes ranked by weighted degree:

$$\text{deg}(v) = \sum_u w(v \rightarrow u).$$

Define

$$\kappa_+ = \frac{1}{|V_k|} \sum_{v \in V_k} \frac{1}{\text{deg}(v)} \sum_{u: (v,u) \in E} \kappa(v, u),$$

and similarly define κ_- as the average curvature over peripheral nodes.

Proposition .2 (Discrete Curvature Concentration). *If*

$$\kappa_+ \gg 0 \quad \text{and} \quad \kappa_- \ll 0,$$

then G exhibits discrete curvature concentration. In the manifold limit, this corresponds to a region $\mathcal{P} \subset \mathcal{M}$ with positive Ricci curvature bounded below, surrounded by negatively curved periphery.

Proof. Positive curvature implies contraction of probability mass under transport:

$$W_1(\mu_v, \mu_u) < d(v, u).$$

Hence random walks initiated at high-degree nodes contract toward overlapping neighborhoods, implying geodesic convergence in the continuous limit.

Negative curvature at the periphery implies expansion:

$$W_1(\mu_v, \mu_u) > d(v, u),$$

so peripheral regions disperse.

Under graph-to-manifold convergence, these inequalities persist in scaled form and induce positive Ricci curvature in \mathcal{P} and negative curvature elsewhere. \square

F.6 Spectral Interpretation

Let P be the transition matrix of the random walk induced by μ_v . Positive average Ollivier curvature implies a lower bound on the spectral gap λ_2 of P :

$$\lambda_2 \geq c \inf_{(v,u)} \kappa(v, u),$$

for some constant $c > 0$ depending on degree regularity.

Thus curvature concentration implies rapid mixing within the high-curvature region, which operationally corresponds to behavioral trapping or recommendation cycling.

F.7 Operational Detection Criterion

Define the curvature skew

$$\Delta\kappa = \kappa_+ - \kappa_-.$$

If

$$\Delta\kappa > \theta$$

for some threshold $\theta > 0$ calibrated against null-model graphs, then the network exhibits statistically significant curvature concentration.

This provides an empirical test for scalar-well formation in platform data. Engagement-driven architectures are predicted to exhibit increasing $\Delta\kappa$ over time. Entropy-bounded

architectures are predicted to maintain $\Delta\kappa$ near zero.

F.8 Interpretation

The discrete curvature formalism establishes a measurable bridge between graph-theoretic observables and continuous geometric structure. Positive Ollivier–Ricci curvature at high-degree nodes corresponds to contraction of probability mass and behavioral convergence. Negative curvature at the periphery corresponds to expansion and volatility. A sharp curvature differential therefore signals geometric concentration of influence.

This discrete framework justifies empirical calibration of the continuous curvature concentration parameter used in the main body of the paper.

Appendix G: Statistical Stability of Empirical Curvature Estimates

G.1 Empirical Curvature Estimation

Let $G = (V, E, w)$ be a weighted directed graph as defined in Appendix F. Suppose engagement weights $w(v \rightarrow u)$ are estimated from N independent interaction observations.

Let $\hat{\mu}_v$ denote the empirical distribution at node v , computed from N_v samples:

$$\hat{\mu}_v(u) = \frac{1}{N_v} \sum_{i=1}^{N_v} \mathbf{1}_{\{u_i=u\}}.$$

Define empirical Ollivier–Ricci curvature

$$\hat{\kappa}(v, u) = 1 - \frac{W_1(\hat{\mu}_v, \hat{\mu}_u)}{d(v, u)}.$$

We analyze the statistical concentration of $\hat{\kappa}$ around its population value κ .

G.2 Concentration of Empirical Measures

For each fixed v , $\widehat{\mu}_v$ is a multinomial empirical distribution. By standard concentration inequalities, for any $\delta > 0$,

$$\mathbb{P}(\|\widehat{\mu}_v - \mu_v\|_1 > \delta) \leq 2 \exp(-2N_v \delta^2).$$

This follows from Hoeffding's inequality applied coordinatewise and union-bounded over support.

G.3 Stability of Wasserstein Distance

The L^1 Wasserstein distance is Lipschitz in its arguments with respect to total variation norm:

$$|W_1(\widehat{\mu}_v, \widehat{\mu}_u) - W_1(\mu_v, \mu_u)| \leq \text{diam}(V) (\|\widehat{\mu}_v - \mu_v\|_1 + \|\widehat{\mu}_u - \mu_u\|_1).$$

Hence for any $\delta > 0$,

$$\mathbb{P}\left(|\widehat{\kappa}(v, u) - \kappa(v, u)| > \frac{2 \text{diam}(V) \delta}{d(v, u)}\right) \leq 4 \exp(-2N_{\min} \delta^2),$$

where $N_{\min} = \min(N_v, N_u)$.

G.4 Uniform Curvature Concentration

Let E contain M edges. By union bound,

$$\mathbb{P}\left(\max_{(v,u) \in E} |\widehat{\kappa}(v, u) - \kappa(v, u)| > \varepsilon\right) \leq 4M \exp(-cN_{\min} \varepsilon^2),$$

for constant $c > 0$ depending on graph diameter and edge lengths.

Thus empirical curvature estimates concentrate exponentially fast in sample size.

G.5 Implication

Given sufficiently large interaction data, empirical curvature skew

$$\widehat{\Delta\kappa} = \widehat{\kappa}_+ - \widehat{\kappa}_-$$

is statistically distinguishable from zero whenever the true curvature concentration parameter exceeds sampling noise.

This establishes statistical identifiability of curvature concentration in observed platform networks.

Appendix H: Dynamical Evolution of Curvature Under Platform Objectives

H.1 Time-Dependent Graph Model

Let $G(t) = (V, E, w_t)$ be a time-dependent weighted graph representing platform evolution. The vertex set V is fixed, while edge weights evolve according to a dynamical rule

$$\frac{d}{dt}w_t(v \rightarrow u) = \Psi_{v,u}(w_t),$$

where Ψ depends on the platform objective.

Define the induced probability measures

$$\mu_v^{(t)}(u) = \frac{w_t(v \rightarrow u)}{\sum_{u'} w_t(v \rightarrow u')}.$$

The time-dependent Ollivier–Ricci curvature is

$$\kappa_t(v, u) = 1 - \frac{W_1(\mu_v^{(t)}, \mu_u^{(t)})}{d(v, u)}.$$

H.2 Engagement-Driven Weight Amplification

Assume engagement optimization increases weights proportionally to current traffic:

$$\frac{d}{dt}w_t(v \rightarrow u) = \alpha w_t(v \rightarrow u), \quad \alpha > 0,$$

for high-engagement edges.

Then

$$w_t(v \rightarrow u) = w_0(v \rightarrow u)e^{\alpha t}.$$

Let v^* denote a high-degree node. If outgoing weights from v^* grow faster than peripheral weights, then its local measure $\mu_{v^*}^{(t)}$ converges to a degenerate distribution

supported on a small subset of neighbors.

Proposition .3 (Curvature Amplification). *If weight amplification is superlinear at high-degree nodes relative to periphery, then*

$$\frac{d}{dt} \kappa_t(v^*, u) > 0$$

for edges incident to v^* in the long-time regime.

Proof. As $\mu_{v^*}^{(t)}$ becomes increasingly concentrated, Wasserstein distances between $\mu_{v^*}^{(t)}$ and neighboring measures shrink relative to $d(v^*, u)$. Hence

$$W_1(\mu_{v^*}^{(t)}, \mu_u^{(t)}) < d(v^*, u),$$

with gap increasing over time. Therefore $\kappa_t(v^*, u)$ increases monotonically. \square

Thus engagement-driven amplification induces positive curvature concentration over time.

H.3 Entropy-Bounded Weight Regulation

Assume instead that weight evolution satisfies a global entropy production constraint:

$$\frac{d}{dt} \sum_v H(\mu_v^{(t)}) \leq \epsilon_{\max}.$$

Suppose regulation enforces diffusion-like normalization:

$$\frac{d}{dt} w_t(v \rightarrow u) = \beta (\bar{w}_t - w_t(v \rightarrow u)), \quad \beta > 0,$$

where \bar{w}_t is a local average weight.

This dynamics is contractive in L^2 and prevents weight divergence.

Proposition .4 (Curvature Stabilization). *Under bounded entropy production and diffusive weight normalization, there exists $C > 0$ such that*

$$|\kappa_t(v, u)| \leq C \quad \forall t.$$

Proof. Diffusive normalization prevents concentration of $\mu_v^{(t)}$ on small supports. Hence Wasserstein distances scale proportionally with graph distances. Since $W_1(\mu_v^{(t)}, \mu_u^{(t)})$

remains comparable to $d(v, u)$, curvature remains bounded. The entropy bound prevents measure collapse, ensuring uniform support size. \square

H.4 Phase Transition Criterion

Define curvature energy

$$\mathcal{C}(t) = \sum_{(v,u) \in E} \kappa_t(v, u)^2.$$

Engagement-driven dynamics satisfy

$$\frac{d}{dt} \mathcal{C}(t) \geq \gamma \mathcal{C}(t),$$

for some $\gamma > 0$ in concentrated regimes.

Entropy-bounded dynamics satisfy

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

for constant C .

Thus engagement optimization induces exponential curvature energy growth, whereas entropy regulation ensures at most linear growth.

H.5 Interpretation

Discrete curvature evolves as a dynamical observable. Engagement objectives amplify curvature energy, producing scalar-well formation. Entropy constraints damp curvature growth, preserving near-flat geometry.

Therefore curvature concentration is not static but dynamically generated by objective functions.

Appendix I: Curvature Entropy Coupling Inequalities

I.1 Metric Measure Preliminaries

Let (X, d, μ) be a metric measure space, where d is a metric and μ a probability measure. Let $\mathcal{P}(X)$ denote the space of probability measures on X with finite first moment.

Define the entropy functional

$$\text{Ent}(\rho) = \int_X \rho(x) \log \rho(x) d\mu(x),$$

for densities ρ with respect to μ .

Let W_2 denote the quadratic Wasserstein distance on $\mathcal{P}(X)$.

I.2 Continuous Ricci Lower Bounds and Entropy Convexity

A classical result in optimal transport theory states that a smooth Riemannian manifold (\mathcal{M}, g) satisfies a Ricci curvature lower bound

$$\text{Ric}_g \geq K$$

if and only if the entropy functional is K -convex along Wasserstein geodesics.

Specifically, for any geodesic $(\rho_t)_{t \in [0,1]}$ in $(\mathcal{P}(\mathcal{M}), W_2)$,

$$\text{Ent}(\rho_t) \leq (1-t)\text{Ent}(\rho_0) + t\text{Ent}(\rho_1) - \frac{K}{2}t(1-t)W_2^2(\rho_0, \rho_1).$$

Thus positive Ricci curvature induces entropy convexity.

I.3 Differential Entropy Production Bound

Let ρ_t satisfy the continuity equation

$$\partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0,$$

where v_t is a velocity field.

The entropy production rate satisfies

$$\frac{d}{dt}\text{Ent}(\rho_t) = - \int_{\mathcal{M}} \langle \nabla \log \rho_t, v_t \rangle \rho_t \, d\text{vol}_g.$$

If $v_t = -\nabla\Phi_t$ is gradient-driven, then

$$\frac{d}{dt}\text{Ent}(\rho_t) = \int_{\mathcal{M}} \langle \nabla \log \rho_t, \nabla\Phi_t \rangle \rho_t \, d\text{vol}_g.$$

By CauchySchwarz,

$$\left| \frac{d}{dt}\text{Ent}(\rho_t) \right| \leq \|\nabla \log \rho_t\|_{L^2(\rho_t)} \|\nabla\Phi_t\|_{L^2(\rho_t)}.$$

I.4 CurvatureEntropy Inequality

Assume $\text{Ric}_g \geq K > 0$.

The Bakrymery criterion implies

$$\frac{d^2}{dt^2}\text{Ent}(\rho_t) \geq K \|\nabla\Phi_t\|_{L^2(\rho_t)}^2.$$

Hence strong positive curvature accelerates entropy convexity and enforces contraction in Wasserstein space.

I.5 Discrete Analogue on Graphs

On a weighted graph $G = (V, E, w)$, define entropy

$$\text{Ent}(\mu_v) = \sum_{u \in V} \mu_v(u) \log \mu_v(u).$$

Ollivier curvature lower bounds imply a discrete convexity inequality:

$$\text{Ent}(\mu_t) \leq (1-t)\text{Ent}(\mu_0) + t\text{Ent}(\mu_1) - \frac{\kappa}{2}t(1-t)W_1^2(\mu_0, \mu_1),$$

for $\kappa = \inf_{(v,u)} \kappa(v, u)$.

Thus discrete positive curvature enforces entropy contraction.

I.6 Entropy Curvature Growth Coupling

Suppose entropy evolves according to

$$\frac{d}{dt}\text{Ent}(\rho_t) = \alpha\text{Ent}(\rho_t), \quad \alpha > 0.$$

Then exponential entropy growth implies

$$\frac{d^2}{dt^2}\text{Ent}(\rho_t) = \alpha^2\text{Ent}(\rho_t).$$

If curvature were uniformly positive and bounded below, entropy convexity would force quadratic control, contradicting exponential growth.

Hence sustained exponential entropy production is incompatible with uniform Ricci lower bounds.

I.7 Phase Dichotomy

We obtain the following dichotomy:

Theorem .5. *Let a platform system be modeled as a metric measure space with entropy functional Ent and Ricci curvature lower bound K .*

If entropy production satisfies

$$\frac{d}{dt}\text{Ent}(\rho_t) \leq \epsilon_{\max},$$

then curvature remains uniformly bounded and the system admits invariant measures.

If instead

$$\frac{d}{dt}\text{Ent}(\rho_t) = \alpha\text{Ent}(\rho_t), \quad \alpha > 0,$$

then no uniform positive curvature lower bound can hold globally; curvature concentration or degeneracy must occur.

Proof. The first case follows from bounded entropy growth and Wasserstein compactness arguments. The second follows from incompatibility between entropy exponential growth and convexity induced by positive curvature. \square

I.8 Interpretation

Curvature and entropy are structurally coupled. Positive curvature enforces entropy convexity and contraction. Exponential entropy growth forces curvature degeneration or concentration. Thus engagement-driven systems necessarily violate global curvature bounds, whereas entropy-bounded systems preserve geometric stability.

Appendix J: Spectral Gap, Curvature, and Mixing Times

J.1 Markov Semigroup Setup

Let (X, d) be a complete separable metric space and let $(P_t)_{t \geq 0}$ be a Markov semigroup acting on bounded measurable functions. We write the dual action on probability measures as νP_t , defined by

$$\int_X f d(\nu P_t) = \int_X (P_t f) d\nu.$$

Assume the semigroup admits at least one invariant probability measure π and that (P_t) is ergodic with respect to π .

We measure convergence to stationarity either in total variation,

$$\|\nu P_t - \pi\|_{\text{TV}},$$

or in Wasserstein distance W_1 (or W_2 when appropriate).

J.2 Coarse Ricci Curvature and Wasserstein Contraction

Assume (P_t) is W_1 -contractive with rate $\kappa > 0$ in the sense that for all probability measures ν, η ,

$$W_1(\nu P_t, \eta P_t) \leq e^{-\kappa t} W_1(\nu, \eta).$$

Equivalently, in Ollivier's formulation, one may define the coarse Ricci curvature at scale t by

$$\kappa_t(x, y) = 1 - \frac{W_1(P_t(x, \cdot), P_t(y, \cdot))}{d(x, y)},$$

and require $\inf_{x \neq y} \kappa_t(x, y) \geq 1 - e^{-\kappa t}$, which yields the exponential contraction bound above.

The contraction bound implies uniqueness of the invariant measure π whenever it exists and provides quantitative mixing bounds. In particular, for any initial measure ν ,

$$W_1(\nu P_t, \pi) \leq e^{-\kappa t} W_1(\nu, \pi).$$

J.3 Spectral Gap and the Poincaré Inequality

Assume (P_t) is reversible with respect to π and generated by an operator L on $L^2(\pi)$, so that $P_t = e^{tL}$. Define the Dirichlet form

$$\mathcal{E}(f, f) = - \int_X f(Lf) d\pi,$$

for f in the domain of L .

We say that (P_t) has spectral gap $\lambda > 0$ if for all $f \in L^2(\pi)$ with $\int f d\pi = 0$,

$$\|P_t f\|_{L^2(\pi)} \leq e^{-\lambda t} \|f\|_{L^2(\pi)}.$$

Equivalently, λ is the largest constant for which the Poincaré inequality holds:

$$\text{Var}_\pi(f) \leq \frac{1}{\lambda} \mathcal{E}(f, f), \quad \text{Var}_\pi(f) = \int_X \left(f - \int f d\pi \right)^2 d\pi.$$

The spectral gap controls L^2 -mixing, and via standard comparisons (e.g. Pinsker-type bounds under additional regularity), it also constrains total variation mixing rates.

J.4 Curvature Lower Bounds Imply Functional Inequalities

In smooth settings, a Ricci curvature lower bound $\text{Ric} \geq K > 0$ implies a Poincaré inequality with $\lambda \geq K$ for the appropriate diffusion generators (e.g. the Laplace–Beltrami operator or more general Bakry–Émery generators). In discrete settings, a positive lower bound on Ollivier–Ricci curvature yields discrete analogues of Poincaré (and sometimes log-Sobolev) inequalities, hence a nontrivial spectral gap.

For purposes of the present framework, we take as a structural hypothesis the implication

$$\inf_{x \neq y} \kappa(x, y) \geq \kappa_0 > 0 \quad \implies \quad \lambda \gtrsim \kappa_0,$$

where λ is the spectral gap of the platform-induced Markov operator and κ denotes coarse (Ollivier) curvature.

This implication is not merely a technical convenience. It encodes the idea that positive curvature corresponds to geometric contraction, which manifests operationally as fast forgetting of initial conditions.

J.5 Mixing Time Bounds from Spectral Gap

Let (P_t) be reversible with invariant π and spectral gap $\lambda > 0$. Then for any f with $\int f d\pi = 0$,

$$\|P_t f\|_{L^2(\pi)} \leq e^{-\lambda t} \|f\|_{L^2(\pi)}.$$

In particular, for any initial density $h = d\nu/d\pi$ with $\int h d\pi = 1$, writing $f = h - 1$ gives

$$\left\| \frac{d(\nu P_t)}{d\pi} - 1 \right\|_{L^2(\pi)} \leq e^{-\lambda t} \left\| \frac{d\nu}{d\pi} - 1 \right\|_{L^2(\pi)}.$$

Define the L^2 mixing time at tolerance $\delta > 0$ by

$$t_{\text{mix}}^{(2)}(\delta) = \inf \left\{ t \geq 0 : \left\| \frac{d(\nu P_t)}{d\pi} - 1 \right\|_{L^2(\pi)} \leq \delta \right\}.$$

Then

$$t_{\text{mix}}^{(2)}(\delta) \leq \frac{1}{\lambda} \log \left(\frac{\left\| \frac{d\nu}{d\pi} - 1 \right\|_{L^2(\pi)}}{\delta} \right).$$

Thus the spectral gap provides logarithmic mixing in tolerance with rate set by $1/\lambda$.

J.6 Mixing Time Bounds from Curvature Contraction

Assume W_1 contraction at rate $\kappa > 0$:

$$W_1(\nu P_t, \eta P_t) \leq e^{-\kappa t} W_1(\nu, \eta).$$

Fix $\eta = \pi$. Then

$$W_1(\nu P_t, \pi) \leq e^{-\kappa t} W_1(\nu, \pi).$$

Define the Wasserstein mixing time

$$t_{\text{mix}}^{(W_1)}(\delta) = \inf \{ t \geq 0 : W_1(\nu P_t, \pi) \leq \delta \}.$$

Then

$$t_{\text{mix}}^{(W_1)}(\delta) \leq \frac{1}{\kappa} \log \left(\frac{W_1(\nu, \pi)}{\delta} \right).$$

In the regime where curvature lower bounds imply a spectral gap, one obtains the comparative heuristic

$$t_{\text{mix}} \asymp \frac{1}{\kappa_0} \log \left(\frac{1}{\delta} \right),$$

up to model-dependent constants and the choice of convergence metric.

J.7 Platform Interpretation: Memory, Capture, and Slow Mixing

Within the platform-graph setting, curvature concentration corresponds to large positive curvature around a small set of hubs and negative curvature on peripheral edges. This mixed-sign geometry induces two simultaneous effects.

First, locally high curvature near hubs accelerates contraction toward the hub subgraph, meaning that trajectories enter the extraction well quickly.

Second, negative curvature and bottlenecks in the periphery reduce global functional inequality constants, shrinking the spectral gap and thereby increasing global mixing times. Operationally, the system forgets initial conditions slowly at the global scale while still rapidly collapsing into the attractor basin of the high-curvature region.

In this sense, extraction regimes can exhibit the paradoxical combination of fast capture and slow equilibration: capture is governed by local curvature concentration, while mixing is governed by the global spectral structure.

J.8 Summary Inequality Chain

We record the conceptual chain in a single statement.

Proposition .6 (Curvature-to-Mixing Control). *Assume a platform dynamics induces a reversible Markov semigroup (P_t) with invariant measure π . Suppose coarse curvature admits a positive lower bound $\kappa_0 > 0$ sufficient to yield both (i) W_1 contraction at rate κ_0 and (ii) a Poincaré inequality with spectral gap $\lambda \gtrsim \kappa_0$. Then for any initial distribution ν absolutely continuous with respect to π and any tolerance $\delta > 0$,*

$$t_{\text{mix}}^{(W_1)}(\delta) \leq \frac{1}{\kappa_0} \log \left(\frac{W_1(\nu, \pi)}{\delta} \right), \quad t_{\text{mix}}^{(2)}(\delta) \leq \frac{1}{\lambda} \log \left(\frac{\left\| \frac{d\nu}{d\pi} - 1 \right\|_{L^2(\pi)}}{\delta} \right).$$

J.9 Consequence for Entropy-Bounded Design

Entropy-bounded design may be interpreted as maintaining curvature away from degeneracy and preventing the formation of extreme bottlenecks. In the Markovian abstraction, this corresponds to preventing collapse of λ and maintaining a uniformly positive contraction rate. The operational meaning is that bounded entropy production is compatible with fast mixing, whereas engagement-amplifying dynamics tends to produce curvature concentration and spectral bottlenecks that slow global mixing even while accelerating capture into hub regions.

Appendix K: Stationary Measure Existence via Krylov–Bogolyubov

K.1 Dynamical and Markovian Formulation

Let (X, d) be a complete separable metric space and let $(P_t)_{t \geq 0}$ be a Feller Markov semigroup on X . That is, P_t maps $C_b(X)$ to itself, preserves positivity and constants, and is strongly continuous in t in the topology of uniform convergence on bounded continuous functions. The dual semigroup acts on probability measures by $\nu \mapsto \nu P_t$.

We seek an invariant probability measure π satisfying

$$\pi P_t = \pi \quad \text{for all } t \geq 0.$$

K.2 Tightness Criterion and Time-Averaged Measures

Fix an initial point $x \in X$ and consider the family of time-averaged probability measures

$$\mu_T = \frac{1}{T} \int_0^T \delta_x P_t dt, \quad T > 0,$$

where δ_x denotes the Dirac measure at x .

Assume the following tightness condition:

Definition .7 (Asymptotic Tightness). *The semigroup (P_t) is asymptotically tight at x if for every $\eta > 0$ there exists a compact set $K_\eta \subset X$ such that*

$$\frac{1}{T} \int_0^T (\delta_x P_t)(K_\eta) dt \geq 1 - \eta$$

for all sufficiently large T .

This condition asserts that the occupation measures do not leak mass to infinity in the time-average.

K.3 Krylov–Bogolyubov Existence Theorem

Theorem .8 (Krylov–Bogolyubov). *Let (P_t) be a Feller Markov semigroup on a Polish space X . If the family $(\mu_T)_{T>0}$ defined above is tight, then there exists an invariant probability measure π .*

Proof. By Prokhorov’s theorem, tightness implies relative compactness in the weak topology on probability measures. Hence there exists a sequence $T_n \rightarrow \infty$ such that $\mu_{T_n} \Rightarrow \pi$ weakly for some probability measure π .

Let $s \geq 0$ and let $f \in C_b(X)$. Using semigroup properties and Fubini,

$$\int f d(\mu_T P_s) = \frac{1}{T} \int_0^T \int (P_s f) d(\delta_x P_t) dt = \frac{1}{T} \int_0^T (P_{t+s} f)(x) dt.$$

Similarly,

$$\int f d\mu_T = \frac{1}{T} \int_0^T (P_t f)(x) dt.$$

Subtracting yields

$$\int f d(\mu_T P_s) - \int f d\mu_T = \frac{1}{T} \left(\int_T^{T+s} (P_u f)(x) du - \int_0^s (P_u f)(x) du \right).$$

Since f is bounded, the right-hand side tends to 0 as $T \rightarrow \infty$. Therefore

$$\lim_{n \rightarrow \infty} \int f d(\mu_{T_n} P_s) = \lim_{n \rightarrow \infty} \int f d\mu_{T_n}.$$

By weak convergence and Feller continuity, $\mu_{T_n} P_s \Rightarrow \pi P_s$ and $\mu_{T_n} \Rightarrow \pi$, hence

$$\int f d(\pi P_s) = \int f d\pi \quad \text{for all } f \in C_b(X).$$

Thus $\pi P_s = \pi$, and since $s \geq 0$ was arbitrary, π is invariant. □

K.4 Sufficient Conditions for Tightness

A common sufficient condition is the existence of a coercive Lyapunov function.

Definition .9 (Coercive Lyapunov Function). *A function $V : X \rightarrow [0, \infty)$ is coercive if its sublevel sets*

$$\{x \in X : V(x) \leq R\}$$

are compact for all $R < \infty$.

Assume there exist constants $a > 0$, $b \geq 0$, and a coercive V such that for all $x \in X$ and all $t \geq 0$,

$$(P_t V)(x) \leq e^{-at} V(x) + \frac{b}{a} (1 - e^{-at}).$$

Then the time averages satisfy a uniform moment bound:

$$\int V d\mu_T = \frac{1}{T} \int_0^T (P_t V)(x) dt \leq \frac{1}{T} \int_0^T \left(e^{-at} V(x) + \frac{b}{a} \right) dt \leq \frac{V(x)}{aT} + \frac{b}{a}.$$

In particular, $\sup_T \int V d\mu_T < \infty$.

By Markov's inequality, for any $R > 0$,

$$\mu_T(\{V > R\}) \leq \frac{1}{R} \int V d\mu_T \leq \frac{C}{R}$$

for some constant C independent of T . Since $\{V \leq R\}$ is compact, this yields tightness of (μ_T) .

K.5 Entropy-Bounded Dynamics as a Lyapunov Condition

In the framework of entropy-bounded dynamics, one may take V to be a monotone transform of an entropy-like functional E or a combined free-energy functional $\mathcal{F} = E + \beta \text{Reg}$ that includes a compactness-enforcing regularizer.

If the design constraint enforces a drift inequality of the form

$$\frac{d}{dt} \mathbb{E}[V(X_t)] \leq -a \mathbb{E}[V(X_t)] + b,$$

then Grönwall's inequality implies the semigroup bound above, hence tightness, hence existence of an invariant measure.

This makes precise the statement that bounded entropy production is not merely a moral constraint but a mathematical stabilizer: it supplies the coercivity needed for stationary behavior.

K.6 Uniqueness and Ergodicity (Structural Addendum)

Existence alone does not guarantee uniqueness. However, if in addition the semigroup is contractive in an appropriate Wasserstein distance, or satisfies a Doeblin-type minorization condition, then invariant measures are unique and convergence to π holds quantitatively. In that case the spectral-gap and curvature bounds of Appendix J become meaningful as global convergence controls rather than merely local contraction heuristics.

K.7 Interpretation

Krylov–Bogolyubov provides a general existence route for stationary measures: enforce tightness of time averages. The platform interpretation is immediate. Engagement-amplifying regimes tend to violate tightness by pushing mass toward extremes (runaway concentration, unbounded state variables, or effective escape to infinity in feature space). Entropy-bounded regimes provide Lyapunov drift, guaranteeing tightness, hence guaranteeing the existence of a stationary measure, hence enabling meaningful long-run evaluation and calibration.

Appendix L: Universality and Factory Emulation

Let $\mathcal{F} = (N, A)$ denote a finite directed acyclic production graph, with node set $N = \{T_1, \dots, T_n\}$ and edge set $A \subseteq N \times N$ encoding admissible compositional ordering. Assume that \mathcal{F} admits a topological ordering compatible with A . For each node $T_i \in N$, assume there exists a countable state space \mathcal{X}_i and \mathcal{Y}_i and a transformation

$$T_i : \mathcal{X}_i \rightarrow \mathcal{Y}_i$$

which is Turing-computable and total on its domain.

Define the composite transformation

$$\mathcal{T}_{\mathcal{F}} = T_{i_k} \circ \dots \circ T_{i_1}$$

where $(T_{i_1}, \dots, T_{i_k})$ is a topological ordering of \mathcal{F} compatible with the edge relation A .

Theorem .10 (Factory Emulation by a Universal Machine). *If each T_i is Turing-computable and $\mathcal{T}_{\mathcal{F}}$ is total on its domain $\mathcal{D} \subseteq \mathcal{X}_{i_1}$, then there exists a program $p_{\mathcal{F}}$*

such that for a universal Turing machine \mathcal{U} ,

$$\mathcal{U}(p_{\mathcal{F}}, x) = \mathcal{T}_{\mathcal{F}}(x) \quad \text{for all } x \in \mathcal{D}.$$

Proof. For each T_i , since T_i is Turing-computable, there exists a Turing machine M_i that computes T_i . Because \mathcal{F} is finite and admits a topological ordering, the sequential composition of these machines defines a Turing machine $M_{\mathcal{F}}$ that computes $\mathcal{T}_{\mathcal{F}}$. Closure of the class of computable functions under composition guarantees that $\mathcal{T}_{\mathcal{F}}$ is computable.

Let \mathcal{U} be a universal Turing machine. By definition of universality, there exists a program encoding $M_{\mathcal{F}}$ as input, denoted $p_{\mathcal{F}}$, such that

$$\mathcal{U}(p_{\mathcal{F}}, x) = M_{\mathcal{F}}(x) = \mathcal{T}_{\mathcal{F}}(x)$$

for all $x \in \mathcal{D}$. □

Corollary .11. *Any finite factory whose control layer is algorithmically specified admits complete behavioral emulation on a universal computing substrate.*

Proof. The control layer of a factory can be represented as a finite production graph \mathcal{F} whose nodes implement computable transformations. By the preceding theorem, the composite transformation of this graph is computable and therefore simulable on a universal machine. Behavioral emulation follows by reproducing inputoutput mappings at the control layer. □

Remark .12. This result establishes that the transition from physical hardware to virtual instantiation does not eliminate the underlying production structure. Rather, it relocates the compositional graph \mathcal{F} from a materially embedded circuit to a program executed on a universal substrate. The factory becomes a software-defined pipeline, preserving functional ontology while altering embodiment.

Appendix M: Dimensional Compression of Process Graphs

Let $\mathcal{F} = (N, A)$ be a finite process graph with node set $N = \{1, \dots, n\}$ and adjacency matrix $A \in \mathbb{R}^{n \times n}$. For each node $i \in N$, let \mathcal{X}_i be a finite state space associated with the local transformation at that node.

Define the aggregate control state space

$$S = \bigoplus_{i=1}^n \mathcal{X}_i,$$

where the direct sum denotes independent aggregation of node state coordinates. The cardinality of the global state space is therefore

$$|S| = \prod_{i=1}^n |\mathcal{X}_i|.$$

Define the structural complexity of \mathcal{F} by

$$C(\mathcal{F}) = \text{rank}(A) + \sum_{i=1}^n \dim(\mathcal{X}_i),$$

where $\dim(\mathcal{X}_i)$ denotes the minimal number of independent parameters required to encode states of \mathcal{X}_i .

Suppose the control architecture of \mathcal{F} is embedded into a silicon substrate via an injective encoding

$$\iota : S \hookrightarrow \{0, 1\}^m.$$

Proposition .13 (Information-Theoretic Compression Bound). *If ι is lossless, then*

$$m \geq \log_2 |S| = \sum_{i=1}^n \log_2 |\mathcal{X}_i|.$$

Proof. A lossless binary encoding of a finite set S requires at least $\lceil \log_2 |S| \rceil$ bits, since any injective map $\iota : S \rightarrow \{0, 1\}^m$ implies $2^m \geq |S|$. Taking logarithms yields $m \geq \log_2 |S|$. \square

Corollary .14. *For fixed process graph \mathcal{F} , any physical implementation of its control layer requires at least $\log_2 |S|$ bits of storage capacity, independent of substrate.*

Proof. Immediate from the injectivity requirement of ι . \square

Let ρ denote transistor density measured in bits per unit volume. If V is the physical volume required to encode m bits, then

$$V \geq \frac{m}{\rho}.$$

Proposition .15 (Geometric Compression Law). *For fixed structural complexity $C(\mathcal{F})$, increasing transistor density ρ decreases the minimal implementation volume V monotonically, while preserving the combinatorial structure of \mathcal{F} .*

Proof. From the lower bound $V \geq m/\rho$ and the previous proposition, $V \geq \log_2 |S|/\rho$. For fixed S , the numerator is constant. Thus V decreases monotonically as ρ increases. The adjacency matrix A and node state spaces \mathcal{X}_i remain unchanged; therefore the abstract production graph \mathcal{F} is invariant under compression. \square

Remark .16. Dimensional compression therefore alters only geometric scale, not formal structure. As fabrication density increases, the physical manifold representing the control layer contracts while the combinatorial topology of the process graph remains invariant. Factory emulation under universal computation and geometric compression under increased ρ are thus complementary mechanisms for relocating production structure from macroscopic hardware into compact silicon representations without loss of functional ontology.

Appendix N: Partial Virtualization Under Thermodynamic Constraint

Let a production transformation T_i decompose into informational and material components:

$$T_i = T_i^{\text{info}} \oplus T_i^{\text{mat}},$$

where

$$T_i^{\text{info}} : \mathcal{X}_i^{\text{info}} \rightarrow \mathcal{Y}_i^{\text{info}} \quad \text{and} \quad T_i^{\text{mat}} : \mathcal{X}_i^{\text{mat}} \rightarrow \mathcal{Y}_i^{\text{mat}}.$$

Assume T_i^{info} is Turing-computable and that T_i^{mat} is governed by thermodynamic constraint

$$\Delta S_{\text{mat}} \geq 0,$$

together with conservation of mass-energy:

$$\Delta E_{\text{mat}} = 0.$$

Let the global transformation \mathcal{T} decompose accordingly:

$$\mathcal{T} = \mathcal{T}^{\text{info}} \oplus \mathcal{T}^{\text{mat}}.$$

Proposition .17 (Virtualizable Informational Subspace). *If $\mathcal{T}^{\text{info}}$ is computable, then there exists a program p such that for a universal machine \mathcal{U} ,*

$$\mathcal{U}(p, x) = \mathcal{T}^{\text{info}}(x)$$

for all admissible $x \in \mathcal{X}^{\text{info}}$.

Proof. By assumption, $\mathcal{T}^{\text{info}}$ is computable. Therefore there exists a Turing machine M computing $\mathcal{T}^{\text{info}}$. By universality of \mathcal{U} , there exists a program p simulating M . Hence $\mathcal{U}(p, x) = \mathcal{T}^{\text{info}}(x)$ for all x in the domain. \square

Proposition .18 (Thermodynamic Non-Virtualizability of Material Subspace). *If \mathcal{T}^{mat} entails state transitions with nonzero entropy production or mass-energy redistribution, then no purely symbolic substitution eliminates the requirement of physical resource expenditure.*

Proof. Let \mathcal{T}^{mat} induce transformation of physical state (E, S) to (E', S') satisfying

$$E' = E, \quad S' \geq S.$$

Any faithful realization of \mathcal{T}^{mat} must instantiate corresponding physical degrees of freedom. Symbolic emulation replaces physical variables with abstract encodings but does not produce corresponding material output states. Thus while informational descriptions of \mathcal{T}^{mat} may be computed, the transformation of matter itself requires physical substrate satisfying conservation and entropy constraints. Therefore full virtualization is impossible unless the material component is null. \square

Corollary .19 (Partial Virtualization Principle). *A transformation admits complete virtualization if and only if its material component is trivial in the sense that \mathcal{T}^{mat} acts solely on symbolic encodings without requiring physical mass-energy exchange.*

Proof. Immediate from the preceding propositions. \square

Definition .20 (Informational Fraction). *Define the informational fraction of a transformation T by*

$$\eta(T) = \frac{\dim(\mathcal{X}^{\text{info}})}{\dim(\mathcal{X}^{\text{info}}) + \dim(\mathcal{X}^{\text{mat}})}.$$

Proposition .21 (Chipification Gradient). *Technological development that increases $\eta(T)$ while holding total functional capacity fixed corresponds to geometric compression of material dependence and expansion of the virtualizable subspace.*

Proof. Increasing $\eta(T)$ decreases the dimensional contribution of \mathcal{X}^{mat} relative to $\mathcal{X}^{\text{info}}$. By universality, the informational component can be emulated symbolically. The material component remains physically instantiated. Thus the relative share of the transformation implementable on a universal substrate increases monotonically in $\eta(T)$. \square

Remark .22. Partial virtualization does not eliminate thermodynamic constraint. It redistributes structural complexity from material manifolds to symbolic manifolds. Chipification corresponds to increasing $\eta(T)$ across production graphs, thereby shrinking the material control manifold while preserving total transformation structure. Entropy production in the informational layer remains bounded by Landauer-type constraints, but the dominant thermodynamic cost shifts from macroscopic mechanics to microscopic switching processes.

Appendix O: Lambda Calculus and Organizational Persistence

Let Λ denote the untyped lambda calculus with syntax

$$M ::= x \mid (MN) \mid \lambda x.M,$$

where x ranges over a countable set of variables. Reduction is given by the β -rule

$$(\lambda x.M)N \rightarrow_{\beta} M[x := N],$$

together with the usual congruence rules. Write \rightarrow_{β}^* for the reflexivetransitive closure of β -reduction.

By the ChurchTuring thesis and standard constructions, Λ is computationally equivalent to Turing machines in the sense that each computable partial function is representable by some $M \in \Lambda$, and conversely.

Definition .23 (Operational Semantics). *The operational semantics of a term $M \in \Lambda$ is the (possibly infinite) reduction graph*

$$\mathcal{R}(M) = \{N \in \Lambda \mid M \rightarrow_{\beta}^* N\},$$

together with the transition relation \rightarrow_{β} restricted to $\mathcal{R}(M)$.

Definition .24 (Faithful Substrate). *A substrate \mathcal{S} is said to implement Λ faithfully*

if there exists an injective encoding

$$\iota : \Lambda \hookrightarrow \mathcal{S}$$

such that for all $M, N \in \Lambda$,

$$M \rightarrow_{\beta} N \iff \iota(M) \rightarrow_{\mathcal{S}} \iota(N),$$

where $\rightarrow_{\mathcal{S}}$ denotes the substrates reduction dynamics.

Theorem .25 (Organizational Persistence). *Let $M \in \Lambda$ represent a production process via its operational semantics. If \mathcal{S}_1 and \mathcal{S}_2 are faithful substrates implementing Λ , then $\mathcal{R}(M)$ is isomorphic under ι_1 and ι_2 . In particular, the operational semantics of M are invariant under substrate migration.*

Proof. Since both \mathcal{S}_1 and \mathcal{S}_2 implement Λ faithfully, there exist encodings ι_1 and ι_2 preserving β -reduction. For any $M \rightarrow_{\beta}^* N$, we have

$$\iota_1(M) \rightarrow_{\mathcal{S}_1}^* \iota_1(N), \quad \iota_2(M) \rightarrow_{\mathcal{S}_2}^* \iota_2(N).$$

Thus the reduction graphs induced by M on \mathcal{S}_1 and \mathcal{S}_2 are both images of $\mathcal{R}(M)$ under injective structure-preserving maps. The resulting transition systems are therefore isomorphic as labeled directed graphs. Hence the operational semantics of M do not depend on the underlying physical substrate, provided β -reduction is faithfully implemented. \square

Corollary .26 (Recoverability of Early Architectures). *Let \mathcal{H} be a computational architecture whose instruction set and control logic can be encoded as a finite term $M \in \Lambda$. Then any faithful implementation of Λ admits an emulator E such that the behavior of \mathcal{H} is reproducible up to operational equivalence.*

Proof. Encode the instruction semantics of \mathcal{H} as M . By universality, M may be interpreted within any faithful substrate. The induced reduction graph matches that of the original architecture modulo encoding. Therefore an emulator exists preserving operational behavior. \square

Remark .27. Organizational persistence follows from the fact that the defining structure of a computational process resides in its reduction relations rather than in the material composition of the hardware executing it. Substrate migration alters physical implementation but preserves symbolic dynamics, provided reduction semantics are conserved. Physical hardware becomes contingent, while organizational form persists as reducible structure.

Appendix P: Chipification as a Renormalization Flow

Let \mathcal{F}_ℓ denote a production architecture realized at physical scale $\ell > 0$, where ℓ parameterizes a characteristic linear dimension of the implementation. Let $\mathcal{T}_{\mathcal{F}_\ell}$ denote the induced composite transformation and let $\mathcal{C}(\ell)$ denote the minimal control description length of \mathcal{F}_ℓ , measured in bits under an optimal encoding.

Assume fabrication density $\rho(\ell)$ satisfies the asymptotic scaling relation

$$\rho(\ell) \sim \ell^{-d},$$

for some effective spatial dimension $d > 0$. The effective physical footprint required to implement the control description is

$$V(\ell) = \frac{\mathcal{C}(\ell)}{\rho(\ell)}.$$

Definition .28 (Scale Transformation). *For $0 < b < 1$, define the scale transformation*

$$\ell \mapsto \ell' = b\ell.$$

Let \mathcal{R}_b denote the induced renormalization map on control descriptions,

$$\mathcal{R}_b : \mathcal{C}(\ell) \mapsto \mathcal{C}(\ell'),$$

subject to preservation of operational equivalence:

$$\mathcal{R}_b(\mathcal{T}_{\mathcal{F}_\ell}) \simeq \mathcal{T}_{\mathcal{F}_{\ell'}},$$

where \simeq denotes equivalence of composite transformations up to isomorphism of state representations.

Proposition .29 (Volume Scaling Law). *Under the density scaling assumption, the effective footprint transforms as*

$$V(\ell') = \frac{\mathcal{C}(\ell')}{\rho(\ell')} \sim \mathcal{C}(\ell') (b\ell)^d.$$

If $\mathcal{C}(\ell') = \mathcal{C}(\ell)$, then

$$V(\ell') = b^d V(\ell).$$

Proof. Substitute $\rho(\ell') \sim (b\ell)^{-d} = b^{-d}\ell^{-d}$ into the definition of $V(\ell')$:

$$V(\ell') = \frac{\mathcal{C}(\ell')}{\rho(\ell')} = \mathcal{C}(\ell') \cdot b^d \ell^d.$$

If $\mathcal{C}(\ell') = \mathcal{C}(\ell)$, then

$$V(\ell') = b^d \frac{\mathcal{C}(\ell)}{\ell^{-d}} = b^d V(\ell).$$

□

Definition .30 (Structural Fixed Point). *A family $\{\mathcal{F}_\ell\}_{\ell>0}$ lies at a renormalization fixed point if*

$$\mathcal{C}(\ell') = \mathcal{C}(\ell) \quad \text{for all } b \in (0, 1),$$

and

$$\mathcal{T}_{\mathcal{F}_{\ell'}} \simeq \mathcal{T}_{\mathcal{F}_\ell}.$$

Proposition .31 (Fixed-Point Invariance). *If \mathcal{F}_ℓ lies at a structural fixed point, then the renormalization operator satisfies*

$$\mathcal{R}_b(\mathcal{C}(\ell)) = \mathcal{C}(\ell),$$

and the operational semantics of \mathcal{F}_ℓ are scale-invariant.

Proof. By hypothesis, $\mathcal{C}(\ell') = \mathcal{C}(\ell)$ for all b . Hence \mathcal{R}_b acts as the identity on $\mathcal{C}(\ell)$. Operational equivalence of $\mathcal{T}_{\mathcal{F}_{\ell'}}$ and $\mathcal{T}_{\mathcal{F}_\ell}$ ensures invariance of functional composition structure under scaling. Therefore \mathcal{F}_ℓ is a fixed point of the renormalization flow. □

Proposition .32 (Flow Toward Informational Dominance). *If $\mathcal{C}(\ell)$ converges to a finite limit \mathcal{C}_∞ as $\ell \rightarrow 0$, while $\rho(\ell) \rightarrow \infty$, then*

$$\lim_{\ell \rightarrow 0} V(\ell) = 0,$$

and the physical footprint of the control layer vanishes asymptotically relative to material throughput.

Proof. If $\mathcal{C}(\ell) \rightarrow \mathcal{C}_\infty < \infty$ and $\rho(\ell) \sim \ell^{-d}$, then

$$V(\ell) = \frac{\mathcal{C}(\ell)}{\rho(\ell)} \sim \mathcal{C}_\infty \ell^d \rightarrow 0 \quad \text{as } \ell \rightarrow 0.$$

□

Remark .33. Chipification corresponds to renormalization flow toward smaller ℓ while preserving compositional invariants. The algebraic structure of the production

process remains fixed, while its physical embedding contracts. Hardware becomes an implementation parameter along the flow; the organizational algebra persists as a scale-invariant structure.

Appendix Q: Universality and Factory Emulation

Let \mathcal{U} denote a universal computing machine in the sense of Turing, and let Λ denote the untyped λ -calculus. By classical equivalence results, \mathcal{U} can simulate any effective procedure expressible in Λ , and conversely every Turing-computable function admits a representation as a λ -term.

Definition .34 (Factory as Computable Transition System). *A factory \mathcal{F} is a tuple $(\mathcal{S}, \mathcal{A}, \Phi)$, where \mathcal{S} is a (finite or countable) state space, \mathcal{A} is a finite set of actuator symbols, and*

$$\Phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$$

is a deterministic transition function. The factory is said to be computable if Φ is Turing-computable as a function on encodings of $(s, a) \in \mathcal{S} \times \mathcal{A}$.

For $t \in \mathbb{N}$, define the t -fold iterated transition

$$\Phi^t(s_0, a_0, \dots, a_{t-1}) = \Phi(\Phi^{t-1}(s_0, a_0, \dots, a_{t-2}), a_{t-1}),$$

with $\Phi^0(s_0) = s_0$.

Proposition .35 (Emulation Theorem). *If $\mathcal{F} = (\mathcal{S}, \mathcal{A}, \Phi)$ is computable, then there exists a program $P_{\mathcal{F}}$ such that for all $s_0 \in \mathcal{S}$ and actuator sequences $(a_t)_{t \geq 0}$,*

$$\mathcal{U}(P_{\mathcal{F}}, s_0, (a_t)_{t=0}^{t-1}) = \Phi^t(s_0, a_0, \dots, a_{t-1}).$$

Proof. Since Φ is Turing-computable, there exists a Turing machine M_{Φ} computing Φ on encodings of (s, a) . Define a Turing machine $M^{(t)}$ that, given $(s_0, a_0, \dots, a_{t-1})$, computes $\Phi^t(s_0, a_0, \dots, a_{t-1})$ by iterative application of M_{Φ} . Iteration of a computable function is computable by primitive recursion on t .

By universality of \mathcal{U} , there exists an encoding $P_{\mathcal{F}}$ of $M^{(t)}$ such that \mathcal{U} simulates $M^{(t)}$ on all admissible inputs. Therefore $\mathcal{U}(P_{\mathcal{F}}, s_0, (a_t)_{t=0}^{t-1})$ produces the same output as $\Phi^t(s_0, a_0, \dots, a_{t-1})$ for all t . \square

Corollary .36 (Organizational Virtualization). *If a physical factory's control layer is computable, then its operational structure admits exact emulation on any universal*

computing substrate. The absence of the original hardware does not eliminate the representability of its transition algebra.

Proof. The control layer is modeled by Φ . By the Emulation Theorem, Φ can be executed by \mathcal{U} under program $P_{\mathcal{F}}$. Hence the operational behavior of the factory is reproducible up to isomorphism of state encodings. \square

Remark .37. Fixed silicon architectures such as early consoles or arcade boards implement deterministic transition systems encoded physically in circuitry. Emulation replaces direct hardware transitions with software-defined transitions implementing the same Φ . The material substrate changes; the transition algebra remains invariant. Organizational persistence thus follows from universality, not from physical continuity of hardware.

Appendix R: Organizational Emulation and Platform Persistence

Let \mathcal{A} denote an application-level platform with functional specification

$$\mathcal{A} = (\mathcal{I}, \mathcal{R}, \mathcal{T}),$$

where \mathcal{I} is the interaction schema (profiles, posts, comments), \mathcal{R} is the rule set governing visibility and ordering, and \mathcal{T} is the transformation algebra over user-generated content.

Definition .38 (Organizational Isomorphism). *Two platforms \mathcal{A}_1 and \mathcal{A}_2 are organizationally isomorphic if there exists a bijection*

$$\Psi : \mathcal{I}_1 \rightarrow \mathcal{I}_2$$

such that rule dynamics commute:

$$\Psi \circ \mathcal{T}_1 = \mathcal{T}_2 \circ \Psi,$$

and interaction constraints are preserved.

Proposition .39 (Persistence Under Reinstantiation). *If \mathcal{A} is organizationally isomorphic to $\tilde{\mathcal{A}}$ implemented on a different infrastructure substrate, then the platform does not cease to exist in structural terms; it undergoes substrate migration.*

Proof. Isomorphism preserves interaction schema and transformation algebra. Since platform identity is defined by $(\mathcal{I}, \mathcal{R}, \mathcal{T})$ rather than physical hosting medium, structural equivalence implies functional persistence under reinstantiation. \square

Definition .40 (Platformification as Infrastructure Absorption). *Let $\mathcal{P}(t)$ denote a family of application platforms over time. Platformification occurs if there exists a sequence of embeddings*

$$\mathcal{A}_0 \hookrightarrow \mathcal{A}_1 \hookrightarrow \dots \hookrightarrow \mathcal{A}_n$$

such that interaction schemas are absorbed into lower-level infrastructure layers.

Proposition .41 (Absorption Theorem). *If a platform \mathcal{A} can be expressed entirely in terms of composable deployment primitives of a lower-level infrastructure \mathcal{I} , then \mathcal{A} is reducible to \mathcal{I} and ceases to require independent physical instantiation.*

Proof. Let deployment primitives form a complete algebra over \mathcal{I} . If \mathcal{A} 's transformation algebra \mathcal{T} is representable as compositions within this algebra, then \mathcal{A} is equivalent to a program over \mathcal{I} . Hence physical independence is unnecessary. \square

Remark .42. Under this formalization, platforms such as early social networks do not disappear; their organizational forms become emulable objects within general-purpose deployment substrates. What changes is not the existence of the structure but its infrastructural embedding.

Appendix S: Chipification as Organizational Compression

Let \mathcal{O} denote a physical organizational system such as a factory, defined abstractly as

$$\mathcal{O} = (S, F, \Phi),$$

where S is a set of states, $F : S \rightarrow S$ is the state-transition operator representing material transformation, and Φ is a resource flow constraint functional.

Definition .43 (Computable Organization). *An organization \mathcal{O} is computable if there exists a Turing machine T such that for every state $s \in S$, the evolution*

$$s_{n+1} = F(s_n)$$

is simulated by T with finite encoding.

Proposition .44 (Universality of Substrate Emulation). *Let U be a universal Turing machine. If \mathcal{O} is computable, then there exists a program $p_{\mathcal{O}}$ such that*

$$U(p_{\mathcal{O}}, s_0)$$

simulates the state trajectory of \mathcal{O} for all $s_0 \in S$.

Proof. By the Church–Turing thesis, any computable function is simulable by a universal Turing machine. Since F is computable by assumption, its iteration is computable, and hence emulable. \square

Definition .45 (Chipification). *Chipification is the process by which a physical organization \mathcal{O} is replaced by its universal emulation $U(p_{\mathcal{O}}, \cdot)$ on silicon substrate \mathcal{S} .*

Proposition .46 (Organizational Compression). *Let \mathcal{O} require physical apparatus of scale L . If its computational description length satisfies*

$$K(\mathcal{O}) \ll \text{Physical Complexity}(L),$$

where K denotes Kolmogorov complexity, then \mathcal{O} admits chipification with net spatial compression.

Proof. If the minimal description length of \mathcal{O} is significantly smaller than the physical apparatus encoding it, then a universal substrate implementing that description occupies asymptotically less physical volume than the original apparatus. The inequality guarantees compression under universal simulation. \square

Remark .47. Early arcade machines and pinball tables encode deterministic electromechanical state machines. Their transition functions are computable and hence emulable. When instantiated on programmable chips, the hardware organization is compressed into symbolic transition rules, preserving dynamics while eliminating dedicated circuitry.

Proposition .48 (Lambda Calculus Sufficiency). *Any universal computing substrate capable of implementing lambda calculus is sufficient to emulate any computable organizational process.*

Proof. Lambda calculus is computationally equivalent to Turing machines. Hence universality follows from Church equivalence. \square

Appendix T: Entropy–Curvature Coupling

Let $(\mathcal{M}, g(t))$ be a smooth n -dimensional Riemannian manifold evolving in time, and let $\phi(t, x)$ be a smooth scalar field representing informational density. Define the total entropy functional

$$\mathcal{E}(t) = \int_{\mathcal{M}} \phi(t, x) \, d\text{vol}_{g(t)}.$$

Assume the coupled evolution system

$$\partial_t \phi = D\Delta_g \phi + f(\phi),$$

$$\partial_t g_{ij} = -2\text{Ric}_{ij} + \beta \nabla_i \phi \nabla_j \phi,$$

where $D > 0$ and $\beta \geq 0$ are constants, Δ_g is the Laplace–Beltrami operator, and Ric_{ij} is the Ricci curvature tensor of $g(t)$.

Evolution of the Volume Form

Under a metric variation $\partial_t g_{ij} = h_{ij}$, the volume element evolves according to

$$\partial_t(d\text{vol}_g) = \frac{1}{2} g^{ij} h_{ij} \, d\text{vol}_g.$$

Substituting

$$h_{ij} = -2\text{Ric}_{ij} + \beta \nabla_i \phi \nabla_j \phi,$$

we obtain

$$\partial_t(d\text{vol}_g) = \left(-\text{Scal}_g + \frac{\beta}{2} |\nabla \phi|^2 \right) d\text{vol}_g,$$

where $\text{Scal}_g = g^{ij} \text{Ric}_{ij}$ is the scalar curvature.

Derivative of the Entropy Functional

Differentiating $\mathcal{E}(t)$,

$$\frac{d}{dt} \mathcal{E}(t) = \int_{\mathcal{M}} \partial_t \phi \, d\text{vol}_g + \int_{\mathcal{M}} \phi \, \partial_t(d\text{vol}_g).$$

Substituting the evolution equations,

$$\frac{d}{dt} \mathcal{E}(t) = \int_{\mathcal{M}} (D\Delta_g \phi + f(\phi)) \, d\text{vol}_g + \int_{\mathcal{M}} \phi \left(-\text{Scal}_g + \frac{\beta}{2} |\nabla \phi|^2 \right) d\text{vol}_g.$$

Integration by Parts

Assuming \mathcal{M} is compact without boundary, or that ϕ decays sufficiently fast at infinity, we have

$$\int_{\mathcal{M}} \Delta_g \phi \, d\text{vol}_g = 0.$$

Hence

$$\frac{d}{dt} \mathcal{E}(t) = \int_{\mathcal{M}} \left(f(\phi) - \phi \text{Scal}_g + \frac{\beta}{2} \phi |\nabla \phi|^2 \right) d\text{vol}_g.$$

Entropy Growth Bound

Suppose entropy budgets impose the integral constraint

$$\int_{\mathcal{M}} f(\phi) \, d\text{vol}_g \leq \epsilon_{\max}.$$

If $\text{Scal}_g \geq 0$ everywhere and $\phi \geq 0$, then

$$- \int_{\mathcal{M}} \phi \text{Scal}_g \, d\text{vol}_g \leq 0.$$

Thus

$$\frac{d}{dt} \mathcal{E}(t) \leq \epsilon_{\max} + \frac{\beta}{2} \int_{\mathcal{M}} \phi |\nabla \phi|^2 \, d\text{vol}_g.$$

If $\beta = 0$, or if $\phi |\nabla \phi|^2$ is uniformly bounded and absorbed into the budget term, then

$$\frac{d}{dt} \mathcal{E}(t) \leq \epsilon_{\max}.$$

Curvature-Induced Dissipation

In the special case $\beta = 0$, the metric evolution reduces to a Ricci-type flow

$$\partial_t g_{ij} = -2\text{Ric}_{ij}.$$

If $\text{Scal}_g \geq 0$, then curvature contributes a non-positive term to $\frac{d}{dt} \mathcal{E}(t)$.

Therefore positive scalar curvature acts as a geometric dissipation mechanism for entropy accumulation. Curvature concentration without entropy control would feed back into ϕ through the coupling term; entropy-bounded dynamics suppress runaway curvature by limiting the source term $f(\phi)$.

The coupled system demonstrates that entropy production and curvature evolution are dynamically linked. Under bounded entropy injection and nonnegative scalar curvature, total entropy grows at most linearly in time. Absent curvature dissipation or entropy budgets, the coupling permits superlinear growth and curvature concentration. This establishes a formal entropy–curvature coupling principle within the continuous geometric model.

Appendix U: Compactness of Bounded Entropy Trajectories

Let Σ denote the space of semantic spheres equipped with metric

$$d(\sigma_1, \sigma_2) = |E(\sigma_1) - E(\sigma_2)| + d_{\text{payload}}(\sigma_1, \sigma_2) + d_{\text{prov}}(\sigma_1, \sigma_2),$$

where E is entropy, d_{payload} is a metric induced by a norm on the finite-dimensional modality space, and d_{prov} is a graph distance on provenance DAGs.

Assume the following:

1. Entropy growth is bounded:

$$E(\sigma(t)) \leq E_0 + \epsilon_{\max} t \quad \text{for } t \in [0, T].$$

2. The modality space is finite-dimensional, i.e., payloads lie in a vector space \mathbb{R}^k with fixed k .
3. Provenance graphs are directed acyclic graphs with uniformly bounded degree and bounded depth growth on $[0, T]$.

Proposition

For any finite time interval $[0, T]$, the trajectory

$$\mathcal{T}_T = \{\sigma(t) : t \in [0, T]\}$$

is precompact in (Σ, d) .

Proof

Entropy Component. The entropy bound implies

$$E(\sigma(t)) \in [E_0, E_0 + \epsilon_{\max} T].$$

Hence the entropy coordinate is contained in a compact interval of \mathbb{R} .

Payload Component. Finite modality dimension implies payloads lie in a finite-dimensional normed vector space. The entropy bound restricts total symbolic complexity; in particular, for discrete encodings, bounded entropy implies a uniform bound on admissible payload configurations up to equivalence. Thus the payload coordinates remain within a bounded subset of \mathbb{R}^k .

In finite-dimensional spaces, bounded sets are precompact. Therefore the payload component admits convergent subsequences.

Provenance Component. Uniformly bounded degree and acyclicity imply that for fixed time horizon T , provenance graphs cannot exhibit unbounded combinatorial branching. Let D be the maximal degree bound. The number of nodes reachable within depth L is bounded above by a geometric sum

$$1 + D + D^2 + \dots + D^L,$$

which is finite for finite L . Since entropy budgets constrain the number of admissible rule applications in $[0, T]$, depth growth is uniformly bounded. Thus the set of admissible provenance graphs over $[0, T]$ lies in a finite combinatorial class up to isomorphism.

Finite combinatorial classes are precompact under discrete graph metrics.

Product Structure. The metric d is the sum of three components. Each component is bounded on $[0, T]$. The space Σ can therefore be viewed as a product of bounded subsets of compact or precompact spaces.

By the Arzel–Ascoli theorem applied to coordinate projections (and using the fact that finite products of precompact sets are precompact), the trajectory \mathcal{T}_T admits convergent subsequences.

Entropy-bounded evolution prevents finite-time escape to infinity in (Σ, d) . In particular, semantic spheres evolving under bounded entropy injection remain confined

to a precompact region of state space over any finite horizon. This establishes geometric non-blowup of entropy-bounded dynamics.

Appendix V: Efficiency of Crystal Equilibria

Let there be n agents indexed by $i = 1, \dots, n$. Each agent selects a modality assignment M_i and thereby contributes to a semantic sphere σ_i . Define the individual utility

$$u_i = TC(\sigma_i) - c_i(M_i),$$

where TC denotes the Texture Crystal valuation and c_i is a cost function.

Define total welfare

$$W = \sum_{i=1}^n TC(\sigma_i).$$

Assume:

1. TC is submodular with respect to modality overlap.
2. Each cost function c_i is convex.
3. Mutual information penalties are uniformly bounded:

$$MI(M_i, M_j) \leq \gamma \quad \text{for all } i \neq j.$$

Definitions

Let W^* denote the maximal achievable welfare under coordinated optimization. Let W^{NE} denote welfare at a Nash equilibrium of the induced game.

The Price of Anarchy is defined as

$$\text{PoA} = \frac{W^*}{W^{NE}}.$$

Theorem

Under the stated assumptions,

$$\text{PoA} \leq 1 + \frac{\gamma}{\min_i TC(\sigma_i)}.$$

Proof

Step 1: Submodularity Bound. Submodularity of TC implies diminishing marginal returns under overlap. For any profile (M_i) and deviation M'_i , the incremental gain satisfies

$$TC(\sigma_i \cup M'_i) - TC(\sigma_i) \leq TC(\sigma_i \cup M'_i \mid \text{no overlap}).$$

Thus redundancy reduces marginal benefit.

Step 2: Nash Inequality. At Nash equilibrium, for each agent i and any unilateral deviation M'_i ,

$$TC(\sigma_i) - c_i(M_i) \geq TC(\sigma'_i) - c_i(M'_i),$$

where σ'_i denotes the sphere under deviation.

Summing over i yields

$$\sum_i TC(\sigma_i) \geq \sum_i TC(\sigma'_i) - \sum_i (c_i(M'_i) - c_i(M_i)).$$

Step 3: Bounding Overlap Loss. Under coordinated optimal assignment, welfare differs from equilibrium welfare only through overlap penalties. Since mutual information penalties satisfy $MI(M_i, M_j) \leq \gamma$, total overlap loss per agent is bounded by γ .

Therefore,

$$W^{NE} \geq W^* - n\gamma.$$

Step 4: Ratio Bound. Assume $\min_i TC(\sigma_i) = T_{\min} > 0$. Then

$$W^{NE} \geq W^* - n\gamma \geq W^* \left(1 - \frac{n\gamma}{W^*}\right).$$

Since $W^* \geq nT_{\min}$,

$$\frac{n\gamma}{W^*} \leq \frac{\gamma}{T_{\min}}.$$

Hence

$$W^{NE} \geq \frac{W^*}{1 + \gamma/T_{\min}}.$$

Rearranging gives

$$\text{PoA} = \frac{W^*}{W^{NE}} \leq 1 + \frac{\gamma}{T_{\min}}.$$

Under bounded mutual information penalties and convex contribution costs, the crystal economy admits equilibria whose inefficiency is uniformly bounded. Redundancy taxation controls divergence from social optimum.

Appendix W: Free Energy and Variational Formulation

Let (\mathcal{M}, g) be a compact Riemannian manifold without boundary, and let $\phi : \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ denote an informational density satisfying

$$\int_{\mathcal{M}} \phi \, d\text{vol}_g = 1.$$

Define the free energy functional

$$\mathcal{F}[\phi] = \int_{\mathcal{M}} \left(\frac{D}{2} |\nabla \phi|^2 + U(\phi) \right) d\text{vol}_g,$$

where $D > 0$ is a diffusion coefficient and $U : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a C^2 convex potential.

First Variation

Let $\phi_\varepsilon = \phi + \varepsilon\eta$, where η is smooth with $\int_{\mathcal{M}} \eta \, d\text{vol}_g = 0$. Then

$$\frac{d}{d\varepsilon} \mathcal{F}[\phi_\varepsilon] = \int_{\mathcal{M}} (D\nabla\phi \cdot \nabla\eta + U'(\phi)\eta) \, d\text{vol}_g.$$

Integrating by parts yields

$$\frac{\delta \mathcal{F}}{\delta \phi} = -D\Delta_g \phi + U'(\phi).$$

Wasserstein Gradient Flow

Equip the space of probability densities with the L^2 -Wasserstein metric. The gradient flow of \mathcal{F} is given by

$$\partial_t \phi = \nabla \cdot \left(\phi \nabla \frac{\delta \mathcal{F}}{\delta \phi} \right).$$

Substituting the functional derivative gives

$$\partial_t \phi = \nabla \cdot (\phi \nabla (-D\Delta_g \phi + U'(\phi))).$$

This fourth-order nonlinear diffusion equation generalizes Cahn–Hilliard dynamics.

Energy Dissipation

Differentiating \mathcal{F} along the flow,

$$\frac{d}{dt} \mathcal{F}[\phi(t)] = \int_{\mathcal{M}} \frac{\delta \mathcal{F}}{\delta \phi} \partial_t \phi \, d\text{vol}_g.$$

Substituting the gradient flow expression and integrating by parts yields

$$\frac{d}{dt} \mathcal{F}[\phi(t)] = - \int_{\mathcal{M}} \phi \left| \nabla \frac{\delta \mathcal{F}}{\delta \phi} \right|^2 \, d\text{vol}_g \leq 0.$$

Hence \mathcal{F} is a Lyapunov functional.

Convexity and Uniqueness

If $U''(\phi) \geq 0$ and (\mathcal{M}, g) has nonnegative Ricci curvature, then \mathcal{F} is displacement convex along Wasserstein geodesics. Consequently, the gradient flow admits a unique minimizer ϕ_∞ satisfying

$$-D\Delta_g\phi_\infty + U'(\phi_\infty) = \lambda,$$

for Lagrange multiplier λ enforcing mass conservation.

Entropy Relation

Define Shannon entropy functional

$$\mathcal{S}[\phi] = \int_{\mathcal{M}} \phi \log \phi \, d\text{vol}_g.$$

If $U(\phi) = \phi \log \phi$, then

$$\mathcal{F}[\phi] = \frac{D}{2} \int_{\mathcal{M}} |\nabla \phi|^2 \, d\text{vol}_g + \mathcal{S}[\phi],$$

and the gradient flow reduces to nonlinear Fokker–Planck dynamics.

Interpretation

Entropy budgets correspond to constraining admissible perturbations $\delta\phi$ so that

$$\mathcal{F}[\phi + \delta\phi] - \mathcal{F}[\phi] \leq \epsilon_{\max}.$$

Thus semantic rule applications act as bounded perturbations within a variational free-energy landscape, and stability corresponds to monotone dissipation of \mathcal{F} .

Appendix X: Large Deviations of Engagement Dynamics

Let (Σ, d) be a Polish state space of semantic configurations, and consider stochastic engagement-driven dynamics given by the It diffusion

$$d\sigma_t = b(\sigma_t) \, dt + \sqrt{\varepsilon} \, dW_t,$$

where $b : \Sigma \rightarrow T\Sigma$ is a measurable drift field representing engagement optimization, W_t is standard Brownian motion in local coordinates, and $\varepsilon > 0$ quantifies noise

intensity.

Invariant Measures

Assume sufficient regularity so that the generator

$$\mathcal{L}f = b \cdot \nabla f + \frac{\varepsilon}{2} \Delta f$$

defines a Feller semigroup. If there exists a Lyapunov function V with

$$\mathcal{L}V \leq -cV + C,$$

for constants $c > 0$, $C < \infty$, then an invariant probability measure π_ε exists.

Empirical Measure and Rate Functional

Define the empirical measure

$$L_T = \frac{1}{T} \int_0^T \delta_{\sigma(t)} dt.$$

By the Freidlin–Wentzell large deviation principle, the family (L_T) satisfies

$$\mathbb{P}(L_T \approx \mu) \asymp \exp\left(-\frac{1}{\varepsilon} I(\mu)\right),$$

where the rate functional admits representation

$$I(\mu) = \inf_{u: \operatorname{div}(\mu u) = 0} \frac{1}{2} \int_{\Sigma} \|u(x) - b(x)\|^2 d\mu(x).$$

Entropy Production and Drift Divergence

Let $S(\sigma)$ denote entropy. Under deterministic flow $\dot{\sigma} = b(\sigma)$, the entropy production rate satisfies

$$\frac{d}{dt} S(\sigma_t) = \nabla S(\sigma_t) \cdot b(\sigma_t).$$

If the divergence of the drift satisfies

$$\nabla \cdot b > 0,$$

in regions of large entropy, then phase-space volume expands toward high-entropy states. In the small-noise limit $\varepsilon \rightarrow 0$, the invariant measure concentrates near attractors maximizing entropy production.

Entropy-Bounded Drift

Suppose instead that the drift is constrained by

$$\nabla \cdot b(\sigma) \leq \epsilon_{\max},$$

for uniform $\epsilon_{\max} < \infty$.

Then for any invariant measure π_ε , the expected entropy growth satisfies

$$\int_{\Sigma} \nabla S(\sigma) \cdot b(\sigma) d\pi_\varepsilon(\sigma) \leq \epsilon_{\max}.$$

Consequently, the rate functional penalizes trajectories that diverge superlinearly in entropy, and invariant measures cannot concentrate on exponentially unstable entropy maxima.

Comparison of Growth Regimes

In engagement-maximizing systems without constraint, effective entropy production behaves as

$$\frac{d}{dt} \mathbb{E}[S(\sigma_t)] \approx \alpha \mathbb{E}[S(\sigma_t)], \quad \alpha > 0,$$

yielding exponential growth

$$\mathbb{E}[S(\sigma_t)] = S_0 e^{\alpha t}.$$

Under entropy-bounded drift,

$$\frac{d}{dt} \mathbb{E}[S(\sigma_t)] \leq \epsilon_{\max},$$

so that

$$\mathbb{E}[S(\sigma_t)] \leq S_0 + \epsilon_{\max} t.$$

Thus engagement dynamics correspond to multiplicative large-deviation amplification, while entropy-bounded systems restrict deviations to at most linear growth in expectation.

Appendix Y: Sheaf-Theoretic Merge via Homotopy Colimits

Let \mathcal{C} be a small category of contexts (sites, regimes, or local interface domains), equipped with a Grothendieck topology J encoding which families of contextual maps are admissible covers. Let E be a complete and cocomplete category of “semantic data” (e.g. **Set**, **Grp**, **Vect**, or a model category of structured states), and let

$$\mathcal{F} : \mathcal{C}^{op} \longrightarrow E$$

be a J -sheaf. Intuitively, $\mathcal{F}(c)$ is the space of admissible local realizations over context c , and restriction maps implement consistency under passage to refinements.

Diagram of Spheres and the Merge Problem

Let **Sphere** denote a category of “semantic spheres” (state objects) equipped with morphisms representing admissible transformations, refinements, or inclusions. Consider a diagram indexed by a small category I ,

$$D : I \longrightarrow \mathbf{Sphere}, \quad i \mapsto D_i,$$

representing a collection of partial spheres and their overlap maps. A merge operation seeks a single sphere M together with structure maps

$$\eta_i : D_i \rightarrow M$$

that are coherent with the diagram and preserve designated invariants.

Model-Categorical Setup

To state merge as a homotopy-coherent colimit, assume **Sphere** admits a model structure or at least a homotopy theory in the sense of a relative category (\mathbf{Sphere}, W) with weak equivalences W encoding observational equivalence. Write $\mathrm{Ho}(\mathbf{Sphere})$ for the homotopy category.

A strict colimit $\mathrm{colim} D$ may be ill-behaved under weak equivalence; therefore we define merge using the *homotopy colimit*:

$$M := \mathrm{hocolim}_I D.$$

Existence of the Homotopy Colimit

Assume **Sphere** is locally presentable and admits functorial cofibrant replacement, or more generally that **Sphere** is combinatorial as a model category. Then for every small I and every diagram $D : I \rightarrow \mathbf{Sphere}$, the homotopy colimit exists and can be computed as the ordinary colimit of a cofibrant replacement QD in the projective model structure on \mathbf{Sphere}^I :

$$\mathrm{hocolim}_I D \simeq \mathrm{colim}_I(QD).$$

Proposition .49 (Existence of Sheaf-Theoretic Merge). *If **Sphere** is a combinatorial model category (equivalently, locally presentable and cofibrantly generated), then for any small indexing category I and any diagram $D : I \rightarrow \mathbf{Sphere}$, the merge object $M = \mathrm{hocolim}_I D$ exists and is unique up to weak equivalence.*

Proof. In a combinatorial model category, the projective model structure on \mathbf{Sphere}^I exists for any small I , and every diagram admits a cofibrant replacement QD . Define $\mathrm{hocolim}_I D := \mathrm{colim}_I(QD)$. Uniqueness up to weak equivalence follows from the standard invariance of derived functors. \square

Compatibility as a Non-Emptiness Condition

To connect merges to sheaf consistency, interpret the diagram D as specifying a family of local contexts and overlap maps in \mathcal{C} . Concretely, suppose there is a functor

$$\pi : \mathbf{Sphere} \rightarrow \mathcal{C}$$

sending a sphere to its context, and each morphism in D maps to a refinement morphism in \mathcal{C} . Applying the sheaf \mathcal{F} to $\pi \circ D$ yields a diagram in E ,

$$\mathcal{F}(\pi(D_i)) \in E.$$

The usual sheaf gluing condition asserts that compatible local sections glue to a global section, expressed as an equalizer over a covering family. For an abstract diagram indexed by I , compatibility is expressed by the existence of a cone in E over $\mathcal{F} \circ \pi \circ D$, i.e. by non-emptiness of the limit object.

Definition .50 (Merge Admissibility). *The diagram D is \mathcal{F} -compatible if*

$$\lim_{i \in I} \mathcal{F}(\pi(D_i)) \neq \emptyset \quad (\text{in } \mathbf{Set}),$$

or more generally if the limiting object is not initial in E .

This condition states that there exists a consistent choice of local data across all overlaps, i.e. a global section for the diagram-shaped cover.

Merge as a Universal Gluing Object

Homotopy colimits satisfy a universal mapping property up to coherent homotopy. For any target sphere Z , one has a natural weak equivalence of mapping spaces

$$\mathrm{Map}_{\mathbf{Sphere}}(\mathrm{hocolim}_I D, Z) \simeq \mathrm{holim}_{i \in I} \mathrm{Map}_{\mathbf{Sphere}}(D_i, Z),$$

so a map out of the merge corresponds to a homotopy-coherent family of maps out of each component D_i compatible with the overlap maps in D .

Thus $\mathrm{hocolim}_I D$ is the *least* sphere (up to homotopy) containing all D_i with their identifications, i.e. the canonical merge object.

Invariant Preservation as a Left-Exact Constraint System

Let invariants be encoded by a predicate or subfunctor selecting admissible spheres and morphisms. Formally, let

$$\mathbf{Sphere}^{\mathrm{adm}} \subseteq \mathbf{Sphere}$$

be a wide subcategory of admissible objects and invariant-preserving morphisms, where “admissible” means satisfying, for instance:

Entropy bound. A functional $E : \mathrm{Ob}(\mathbf{Sphere}) \rightarrow \mathbb{R}_{\geq 0}$ and a budget parameter ϵ_{\max} such that admissible transformations do not increase E beyond the budget.

Crystal mass conservation. A conserved quantity $m : \mathrm{Ob}(\mathbf{Sphere}) \rightarrow \mathbb{R}_{\geq 0}$ such that m is preserved under admissible morphisms.

Staking validity. A logical or type-theoretic judgment $\vdash \sigma : \mathrm{Valid}$ stable under restrictions and under admissible identifications.

Assume these invariants are *pullback-stable* and *homotopy-invariant*, meaning they are preserved under weak equivalence and under homotopy-coherent gluing along overlaps.

Proposition .51 (Invariant Preservation Under Merge). *Suppose $D : I \rightarrow \mathbf{Sphere}^{\text{adm}}$ is a diagram whose morphisms preserve the designated invariants, and suppose $\mathbf{Sphere}^{\text{adm}}$ is closed under homotopy colimits of shape I (equivalently, the inclusion $\mathbf{Sphere}^{\text{adm}} \hookrightarrow \mathbf{Sphere}$ preserves the relevant derived colimit). Then the merge object*

$$M = \text{hocolim}_I D$$

is admissible and inherits all invariants satisfied by the D_i .

Proof. Compute M as $\text{colim}(QD)$ for a cofibrant replacement QD in \mathbf{Sphere}^I . By assumption, admissibility is invariant under weak equivalence, so each QD_i remains admissible. Closure of $\mathbf{Sphere}^{\text{adm}}$ under the derived colimit implies that $\text{colim}(QD)$ lies in $\mathbf{Sphere}^{\text{adm}}$. Since invariants are preserved along the diagram maps and are stable under homotopy-coherent identifications, the glued object inherits the invariant constraints. \square

Interpretation

The homotopy colimit implements merge as *constraint-respecting gluing* rather than overwrite. The limit condition in E enforces compatibility of local data, while closure under derived colimits ensures that global assembly cannot violate entropy budgets, conservation laws, or staking judgments merely by rearranging identifications.

In this sense, merge is admissible precisely when local sections agree on overlaps and the category of admissible spheres is homotopy-cocomplete in the relevant shapes. This provides a mathematically rigid notion of “semantic merge” that generalizes both sheaf gluing and version-control reconciliation to a setting where invariants are first-class and preserved at the level of execution semantics.

Appendix Z: No-Singularity Theorem for Invariant-Bounded Substrates

Let $(\mathcal{M}, g(t))$ be a compact Riemannian manifold representing the global semantic substrate, and let $\phi(t, x) \geq 0$ denote informational density evolving under a constrained flow.

Assume evolution equations of the general form

$$\partial_t \phi = \mathcal{D}(\phi, g) + \mathcal{R}(\phi),$$

$$\partial_t g_{ij} = -2 \operatorname{Ric}_{ij} + \mathcal{K}_{ij}(\phi).$$

Here \mathcal{D} denotes a second-order diffusion operator that is uniformly parabolic with respect to the evolving metric $g(t)$, ensuring local smoothing and well-posedness of the density evolution. The term \mathcal{R} represents a reaction functional that is bounded in the sense that there exists $C_R < \infty$ such that $|\mathcal{R}(\phi)| \leq C_R(1 + |\phi|)$ pointwise. The tensor $\mathcal{K}_{ij}(\phi)$ is assumed to be symmetric and to depend smoothly on ϕ and its first derivatives, with growth controlled by a uniform constant C_K so that

$$|\mathcal{K}_{ij}(\phi)| \leq C_K (1 + |\phi| + |\nabla\phi|^2).$$

Under these regularity and boundedness assumptions, the coupled system defines a quasilinear parabolic flow whose short-time existence and uniqueness follow from standard theory for geometric evolution equations.

Assume the following invariant constraints hold:

$$\begin{aligned} \int_{\mathcal{M}} \phi \, d\operatorname{vol}_g &= C_0, \\ \frac{d}{dt} \int_{\mathcal{M}} \phi \log \phi \, d\operatorname{vol}_g &\leq \epsilon_{\max}, \\ \sup_{t \geq 0} \|\operatorname{Ric}(g(t))\|_{L^\infty} &\leq K_0. \end{aligned}$$

Definition

A curvature singularity occurs at time $T < \infty$ if

$$\limsup_{t \rightarrow T} \|\operatorname{Rm}(g(t))\|_{L^\infty} = \infty.$$

Theorem

[No Finite-Time Singularity Under Invariant Bounds]

If entropy growth is uniformly bounded and Ricci curvature remains bounded, then no curvature singularity can occur in finite time.

Proof

Bounded Ricci curvature implies, by standard elliptic regularity and Shi-type derivative estimates for geometric flows, uniform bounds on higher derivatives of curvature over finite intervals.

The entropy bound ensures that ϕ cannot concentrate mass arbitrarily in shrinking regions. Indeed, if ϕ were to develop delta-like concentration, entropy would diverge, contradicting the imposed constraint.

Since curvature concentration would require either metric collapse or density blow-up, both are precluded by bounded entropy and Ricci control. Therefore

$$\sup_{t \in [0, T]} \|\text{Rm}(g(t))\|_{L^\infty} < \infty$$

for all finite T .

Hence no finite-time singularity occurs. □

Corollary

[Structural Stability]

Under the invariant bounds established above, the coupled density-metric system admits no finite-time geometric collapse. In particular, suppose the entropy functional and curvature norms remain uniformly bounded on any finite interval. Then concentration of informational density sufficient to induce divergence of sectional or scalar curvature cannot occur. Any local amplification of ϕ must be offset by diffusive or dissipative terms before the curvature tensor can escape the controlled regime.

Equivalently, bounded entropy production and bounded Ricci curvature jointly imply uniform control over the full Riemann curvature tensor on compact time intervals. Hence the substrate evolves within a precompact region of configuration space, and geometric regularity is preserved.

Interpretation

In contrast, engagement-maximizing dynamics correspond to removing or relaxing the entropy constraint, thereby permitting superlinear or exponential growth in ϕ . Such growth amplifies gradient terms in the metric evolution equation, which in turn drive curvature concentration. In the absence of conservation constraints, the flow may

approach singular configurations characterized by metric degeneration or unbounded curvature.

Entropy-bounded architectures instead enforce a global regularity condition. The constraint on entropy production functions analogously to an energy condition in geometric analysis: it prohibits the accumulation of density beyond what the curvature tensor can accommodate without divergence.

Thus singularity avoidance is not achieved by external moderation or ex post intervention, but by embedding conservation laws into the evolution equations themselves. Geometric collapse is precluded at the level of dynamics. Stability emerges as a theorem rather than as a policy preference.

References

- [1] Wu, T. (2016). *The Attention Merchants: The Epic Scramble to Get Inside Our Heads*. Alfred A. Knopf.
- [2] Wu, T. (2025). *The Age of Extraction*. Alfred A. Knopf.
- [3] Lorenz, T. (2023). *Extremely Online: The Untold Story of Fame, Influence, and Power on the Internet*. Simon & Schuster.
- [4] Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
- [5] Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029.
- [6] Evans, D. S., & Schmalensee, R. (2016). *Matchmakers: The New Economics of Multisided Platforms*. Harvard Business Review Press.
- [7] Herman, E. S., & Chomsky, N. (1988). *Manufacturing Consent: The Political Economy of the Mass Media*. Pantheon Books.
- [8] Farrell, J., & Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics*, 16(1), 70–83.
- [9] Amari, S.-I., & Nagaoka, H. (2000). *Methods of Information Geometry*. American Mathematical Society.
- [10] Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.
- [11] Martin-Löf, P. (1984). *Intuitionistic Type Theory*. Bibliopolis.
- [12] Mac Lane, S. (1971). *Categories for the Working Mathematician*. Springer.
- [13] Ghrist, R. (2014). *Elementary Applied Topology*. Createspace.
- [14] Curry, J. (2014). *Sheaves, Cosheaves and Applications*. Ph.D. Dissertation, University of Pennsylvania.
- [15] Castro, M., & Liskov, B. (1999). Practical Byzantine fault tolerance. *Proceedings of OSDI*.
- [16] Krylov, N., & Bogolyubov, N. (1937). La théorie générale de la mesure dans son application à l'étude des systèmes dynamiques. *Annals of Mathematics*, 38, 65–113.

- [17] Prokhorov, Y. (1956). Convergence of random processes and limit theorems in probability theory. *Theory of Probability and Its Applications*, 1, 157–214.
- [18] Gallot, S., Hulin, D., & Lafontaine, J. (2004). *Riemannian Geometry*. Springer.
- [19] Evans, L. C. (2010). *Partial Differential Equations*. American Mathematical Society.
- [20] Akerlof, G. A. (1970). The market for ‘lemons’: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- [21] Stiglitz, J. E. (1975). The theory of ‘screening’, education, and the distribution of income. *American Economic Review*, 65(3), 283–300.
- [22] Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press.
- [23] Diakopoulos, N. (2015). Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, 3(3), 398–415.
- [24] Ollivier, Y. (2009). Ricci curvature of Markov chains on metric spaces. *Journal of Functional Analysis*, 256(3), 810–864.
- [25] Bickel, P. J., & Doksum, K. A. (2007). *Mathematical Statistics: Basic Ideas and Selected Topics*. Pearson.
- [26] European Parliament. (2022). *Regulation (EU) 2022/1925 on Contestable and Fair Markets in the Digital Sector (Digital Markets Act)*. Official Journal of the European Union.