Collapse-Selection as a Unified Framework for Structure Across Scales

Stephen Garner

April 20, 2026

Abstract

Across cognition, machine learning, social systems, and ethics, stable structure appears under distinct descriptive frameworks that are typically treated as independent. In this work, we propose a unified formulation in which these phenomena are interpreted as instances of a common generative process: collapse-selection acting on relational configuration spaces.

Within this framework, structure is not generated directly, but arises as the invariant residue of constraint-driven selection. We show that intuition, externalized reasoning, system-level persistence, and ethical structure can each be modeled as collapse dynamics over appropriate configuration spaces. These domains differ in representation, but share a common architecture defined by admissibility, collapse, and invariant structure.

The resulting formulation provides a unified account of how structure forms, transfers, stabilizes, and is selected across scales, situating diverse phenomena within a single generative ontology.

1 Introduction

Across multiple domains, stable structure emerges in forms that appear conceptually distinct. In cognition, structure appears as intuition prior to articulation. In machine learning systems, coherent outputs arise from partial inputs. In social systems, persistent patterns of coordination and breakdown are observed. In ethics, norms emerge that appear stable across cultures and time.

These phenomena are typically studied in isolation. Cognitive science treats intuition as a preformal process, machine learning studies reconstruction and generalization, systems theory analyzes persistence and failure, and ethics examines normative structure. Despite differences in language and method, each domain exhibits a common pattern:

- a space of possible configurations,
- constraints acting on that space,
- elimination of unstable configurations,
- persistence of stable structure.

We propose that this pattern reflects a single generative mechanism. Specifically, we interpret these domains as instances of collapse-selection dynamics acting on relational configuration spaces. In this view, structure arises not as a primitive, but as the invariant residue of repeated selection under constraint.

2 Minimal Structural Framework

Let Σ denote a relational configuration space, whose elements represent structured configurations of interacting degrees of freedom. These configurations are evaluated under a collapse operator

$$\Phi : \Sigma \to \Sigma,$$

which acts to eliminate configurations that fail to satisfy admissibility constraints.

Define the invariant sector:

$$I = \{x \in \Sigma \mid \Phi(x) = x\}.$$

Configurations in I are stable under repeated application of Φ and correspond to persistent structure.

We further introduce a projection

$$P : \Sigma \to O,$$

mapping relational configurations to observable representations. This mapping is generally many-to-one and lossy.

Within this framework:

- generation corresponds to exploration of Σ,
- collapse corresponds to selection under constraint,
- persistence corresponds to membership in I,
- observation corresponds to projection into O.

The central claim is:

Structure is the invariant residue of collapse-selection under constraint.

3 Instantiations Across Scales

We now show that this framework applies across four domains.

3.1 Cognitive Generation

Let Σ_H denote a relational configuration space corresponding to pre-verbal cognition. Define a collapse operator

$$\Phi_H : \Sigma_H \to \Sigma_H.$$

Configurations in the invariant sector

$$I_H = \{x \in \Sigma_H \mid \Phi_H(x) = x\}$$

correspond to stabilized intuitive structure.

Articulation is given by a projection

$$P_H : \Sigma_H \to O,$$

which maps internal configurations to symbolic representations.

Thus, intuition corresponds to convergence toward invariant structure prior to projection.

3.2 Cross-System Reconstruction

Let $\Sigma_{joint} = \Sigma_H \cup \Sigma_{LLM}$ denote a joint configuration space spanning internal cognition and model-based reconstruction.

Define an effective collapse operator

$$\Phi_{EIC} : \Sigma_{joint} \to \Sigma_{joint},$$

induced by projection, reconstruction, and validation.

Stable configurations

$$I_{EIC} = \{x \in \Sigma_{joint} \mid \Phi_{EIC}(x) = x\}$$

correspond to articulated understanding.

In this setting, structure is not transmitted directly, but reconstructed under shared admissibility constraints.

3.3 System-Level Persistence

Let Σ_{sys} denote a configuration space of system-level states. Collapse dynamics act to suppress configurations that degrade persistence.

Invariant structures correspond to system-level constraints such as trust, identity, and coordination. Systems persist when these invariants are preserved, and fail when they are violated.

Thus, system stability corresponds to invariance under collapse across scale.

3.4 Ethical Selection

Let Σ_E denote a space of ethical configurations. For a family of collapse classes $\mathcal{C}$, define operators

$$\Phi_E^{\mathcal{C}_i} : \Sigma_E \to \Sigma_E.$$

Global ethical invariance is defined by

$$E \in \bigcap_{\mathcal{C}_i \in \mathcal{C}} \mathrm{Inv}(\Phi_E^{\mathcal{C}_i}).$$

Ethical instability is measured by the persistence defect

$$T_{\mathrm{glob}}(E) = \sup_{\mathcal{C}_i \in \mathcal{C}} T_{\mathcal{C}_i}(E),$$

which quantifies deviation from invariance.

Thus, ethical structure corresponds to invariant configurations across collapse regimes.

4 Unified Collapse-Selection Pipeline

These domains can be integrated into a single process:

$$\Sigma_H \xrightarrow{\Phi_H} I_H \xrightarrow{P_H} O \xrightarrow{R,V} \Sigma_{joint} \xrightarrow{\Phi_{EIC}} I_{EIC} \xrightarrow{\text{scale}} \Sigma_{sys} \xrightarrow{\Phi} I_{sys} \xrightarrow{\mathcal{C}} I_{eth}.$$

At each stage:

- structure is generated in a configuration space,

- collapse selects admissible configurations,
- invariant structure persists,
- projection produces observable representation.

Each layer inherits the invariant residue of the previous one.

This diagram 1 summarizes the collapse-selection framework across cognitive, communicative, system, and ethical domains.

5 Example: Trust as Collapse-Stable Structure Across Scales

To illustrate the unified collapse-selection framework, we consider the emergence and persistence of trust across cognitive, communicative, system, and ethical layers.

5.1 Cognitive Generation

At the generative level, trust arises as a relational configuration in a pre-verbal space Σ_H. Through collapse-selection dynamics

$$\Phi_H : \Sigma_H \to \Sigma_H,$$

configurations that exhibit coherence across past interactions, expectations, and behavioral patterns converge toward invariant structure

$$x^* \in I_H.$$

Trust, at this level, corresponds to a collapse-stable relational pattern prior to articulation.

5.2 Externalization and Reconstruction

This internal structure is projected into a descriptive representation:

$$P_H(x^*) \in O,$$

such as the statement "this agent is trustworthy."

However, this projection is lossy. Reconstruction in another system—human or model-based—requires generating candidate configurations consistent with the projected structure and validating them against local constraints.

Thus, trust is not transmitted directly, but reconstructed through shared admissibility structure.

5.3 System-Level Persistence

At the system level, trust functions as an invariant necessary for coordination and persistence. Systems in which trust is preserved correspond to configurations in the invariant sector

$$x \in I_{\text{sys}},$$

and exhibit stable interaction across time.

When trust degrades, the system accumulates persistence defect:

$$T(x) > 0,$$

leading to fragmentation, loss of coordination, and eventual collapse.

We interpret $T(x)$ as a persistence defect functional measuring deviation from collapse-stable structure.

5.4 Ethical Selection

Within the ethical framework, trust corresponds to a configuration that remains invariant across collapse classes:

$$E_{\text{trust}} \in \bigcap_{\mathcal{C}_i \in \mathcal{C}} \text{Inv}(\Phi_E^{\mathcal{C}_i}).$$

This identifies trust as a globally stable ethical invariant.

In contrast, exploitative behavior may satisfy local invariance

$$T_{\mathcal{C}_{\text{local}}}(E) = 0,$$

but fails global invariance:

$$T_{\text{glob}}(E) > 0,$$

resulting in instability under extended interaction.

5.5 Summary

Trust provides a concrete instance of collapse-stable structure across scales:

- it forms as invariant relational structure in cognition,
- it propagates through reconstruction rather than direct transmission,
- it stabilizes systems as a persistence invariant,
- it is selected as an ethical invariant under cross-class collapse.

This example illustrates how a single structural phenomenon appears differently across layers while remaining governed by a common collapse-selection dynamic.

5.6 Empirical Anchor

Empirical studies of social and organizational systems consistently identify trust as a primary determinant of long-term stability. High-trust environments exhibit sustained coordination, lower transaction costs, and resilience under perturbation, while low-trust environments show fragmentation, defensive behavior, and rapid breakdown of cooperation. Observed phenomena such as institutional erosion, breakdown of civic trust, and collapse of collaborative systems align with the predicted increase in persistence defect:

$$T(x) \uparrow \;\Rightarrow\; \text{system instability.}$$

These patterns support the interpretation of trust as a collapse-stable invariant whose degradation leads to system-level failure.

6 Example: Cooperation and Defection as Collapse-Stable Structure

We now consider the canonical example of cooperation and defection, as formalized in the Prisoner's Dilemma, and reinterpret it within the collapse-selection framework across scales.

6.1 Cognitive Generation

At the generative level, agents evaluate relational configurations corresponding to cooperation and defection within a pre-verbal configuration space Σ_H.

Collapse dynamics

$$\Phi_H : \Sigma_H \to \Sigma_H$$

act to select configurations that satisfy internal coherence constraints, including expectations of reciprocity, predictability, and stability under repeated interaction.

While defection may appear locally advantageous, configurations corresponding to sustained cooperation often emerge as more coherent under extended evaluation, converging toward invariant structure

$$x^* \in I_H.$$

Thus, the intuition that cooperation is "correct" arises as a collapse-stable configuration prior to formal reasoning.

6.2 Externalization and Reconstruction

When articulated, these configurations are projected into descriptive form:

$$P_H(x^*) \in O,$$

for example as arguments for cooperation or defection.

However, as in the general case, projection is lossy. Reconstruction in another agent or system produces candidate configurations that are validated against local admissibility constraints.

This explains the observed variability in reasoning about cooperation: different agents reconstruct different approximations of the underlying relational structure, leading to divergence unless collapse-stable invariants are shared.

6.3 System-Level Persistence

At the system level, cooperation and defection correspond to configurations in a strategic space X.

Classical game theory identifies defection as a Nash equilibrium under local payoff optimization. However, as shown in the persistence-based extension:

$$U_i(x) = U_i^{(\text{local})}(x) + \lambda P(x),$$

where $P(x)$ measures system persistence, cooperation yields significantly higher persistence than defection.

Under collapse dynamics:

$$\Phi : X \to X,$$

configurations that degrade system viability (e.g., repeated defection) exhibit increasing instability:

$$T(x) \uparrow,$$

and are eliminated or reorganized.

Thus:

- defection is locally stable but globally unstable,
- cooperation is selected as a collapse-stable invariant.

This resolves the apparent paradox between individual rationality and collective stability.

6.4 Ethical Selection

Within the ethical framework, cooperation corresponds to a configuration that remains invariant across collapse classes:

$$E_{\text{coop}} \in \bigcap_{\mathcal{C}_i \in \mathcal{C}} \text{Inv}(\Phi_E^{\mathcal{C}_i}),$$

while defection satisfies only local invariance:

$$T_{\mathcal{C}_{\text{local}}}(E_{\text{defect}}) = 0, \quad T_{\text{glob}}(E_{\text{defect}}) > 0.$$

Thus, cooperation is selected as a globally stable ethical structure, while defection corresponds to a locally stable but globally unstable configuration.

6.5 Unified Interpretation

Across all layers:

- cooperation emerges as a collapse-stable invariant under persistence constraints,
- defection emerges as a local attractor that fails under scale extension,
- collapse dynamics select configurations that preserve the conditions for continued interaction.

This example demonstrates that cooperative behavior is not an exception to rational dynamics, but a structural consequence of collapse-selection under constraint.

6.6 Summary

The Prisoner's Dilemma, reinterpreted through collapse-selection dynamics, provides a concrete instance of how local optimization fails to produce global stability, and how invariant structure emerges through persistence-based selection across scales.

It illustrates that:

- intuition favors cooperation through generative collapse,
- communication reconstructs cooperative structure imperfectly,
- systems stabilize only under cooperative invariants,
- ethics selects cooperation as a globally stable configuration.

6.7 Empirical Anchor

Experimental studies of repeated Prisoner's Dilemma interactions and evolutionary game dynamics consistently demonstrate the emergence of cooperative behavior, despite the dominance of defection under one-shot rationality. Mechanisms such as reciprocity, reputation, and conditional cooperation stabilize interaction systems over time, while persistent defection leads to breakdown of interaction, loss of participants, or enforcement responses. These results align with the collapse-selection interpretation, in which cooperative configurations exhibit higher persistence

$$P(x) \gg P(\text{defection}),$$

and are therefore selected as invariant structures under repeated interaction. The observed divergence between short-term payoff maximization and long-term system viability provides empirical support for the distinction between local and global stability.

7 Example: Market Stability and Collapse as Collapse-Selection Dynamics

We now consider the behavior of markets as a large-scale instance of collapse-selection dynamics, illustrating how locally stable configurations may fail under extended interaction.

7.1 Cognitive Generation

At the cognitive level, agents form internal configurations in Σ_H representing expectations of value, trust in counterparties, and perceived stability of the market.

Under collapse dynamics

$$\Phi_H : \Sigma_H \to \Sigma_H,$$

agents converge toward configurations that appear coherent under available information. These configurations may support either cooperative (long-term stability) or exploitative (short-term gain) strategies.

Importantly, local coherence does not guarantee global stability. Agents may converge to configurations that appear internally consistent but fail under broader interaction.

7.2 Externalization and Reconstruction

Market behavior is externalized through observable signals:

$$P_H(x) \in O,$$

including prices, trades, and financial instruments.

Other agents reconstruct these signals into internal configurations through inference processes analogous to reconstruction operators. However, as in previous cases, this process is lossy and constrained.

This creates a feedback loop in which:

- signals are interpreted as evidence of stability,
- reconstructed configurations reinforce existing expectations,
- divergence from underlying invariants is masked.

7.3 System-Level Persistence

At the system level, markets correspond to configurations in a space Σ_{sys} of interacting strategies, institutions, and incentives.

Classical models treat market equilibria as stable configurations under local optimization. However, within the collapse-selection framework, stability depends on persistence:

$$\Phi : \Sigma_{\text{sys}} \to \Sigma_{\text{sys}}.$$

Configurations that maximize short-term payoff may degrade system-level invariants such as trust, transparency, and coordination.

This corresponds to an increase in persistence defect:

$$T(x) \uparrow .$$

Examples include:

- excessive leverage,
- opaque financial instruments,
- misaligned incentives,
- information asymmetry.

Such configurations may appear stable locally, but accumulate instability under repeated interaction.

7.4 Collapse Dynamics and Failure

As persistence defect increases, the system approaches a threshold beyond which configurations are no longer admissible:

$$T(x) > T_{\text{crit}} \quad \Rightarrow \quad \Phi(x) \neq x.$$

At this point, collapse occurs as rapid reconfiguration:

- liquidity vanishes,
- trust breaks down,
- coordination fails,
- prices undergo discontinuous adjustment.

Market collapse is therefore not an external shock, but the consequence of accumulated invariant violation.

7.5 Ethical Selection

Within the ethical framework, stable market configurations correspond to structures that preserve invariants across collapse classes:

$$E_{\text{stable}} \in \bigcap_{\mathcal{C}_i \in \mathcal{C}} \text{Inv}(\Phi_E^{\mathcal{C}_i}).$$

These include:

- transparency,
- bounded risk-taking,
- aligned incentives,

- maintenance of trust.

In contrast, exploitative or short-term profit-maximizing strategies satisfy only local invariance:

$$T_{\text{glob}}(E) > 0,$$

and are eliminated under extended collapse dynamics.

7.6 Relation to Equilibrium Concepts

This framework provides a reinterpretation of equilibrium:

- classical equilibrium corresponds to local stability under fixed assumptions,
- collapse-stable structure corresponds to persistence under dynamic constraints.

Thus, equilibrium is a descriptive artifact of a reduced state space, while collapse-selection determines which configurations remain viable.

7.7 Unified Interpretation

Across layers:

- agents form locally coherent expectations through collapse,
- signals are reconstructed imperfectly across participants,
- markets stabilize only when invariants are preserved,
- collapse occurs when persistence defect exceeds admissibility thresholds.

7.8 Summary

Market collapse provides a concrete example of collapse-selection dynamics in which:

- local optimization produces configurations that appear stable,
- invariant violation accumulates across interactions,
- collapse reorganizes the system into admissible structure.

This demonstrates that large-scale economic behavior is governed by the same structural principles as cognition, communication, and ethics: stability emerges only where invariant structure is preserved under constraint.

7.9 Empirical Anchor

Historical instances of financial instability, including credit crises and market collapses, exhibit consistent structural features predicted by the collapse-selection framework. Periods preceding collapse are characterized by increasing leverage, opacity, and misaligned incentives, which degrade system-level invariants such as trust, transparency, and coordination. These conditions correspond to a rising persistence defect:

$$T(x) \uparrow,$$

often accompanied by increased volatility and divergence between local performance indicators and systemic stability. Collapse events—such as liquidity crises or rapid price discontinuities—then occur as large-scale reconfigurations of the system. These observations support the interpretation of market collapse as the consequence of accumulated invariant violation rather than exogenous shock.

8 Persistence, Failure, and Collapse

Stability corresponds to invariance:

$$x \in \mathrm{Inv}(\Phi).$$

Instability corresponds to persistence defect:

$$T(x) > 0.$$

System failure arises when instability accumulates:

$$T(x) \uparrow \Rightarrow \text{collapse or reconfiguration.}$$

Across domains:

- cognitive failure corresponds to incoherent configurations,
- reconstruction failure corresponds to misalignment,
- system failure corresponds to invariant violation,
- ethical failure corresponds to lack of cross-class invariance.

Thus, collapse is not an external event, but the natural consequence of unresolved structural tension.

9 Relation to Existing Frameworks

This formulation situates existing approaches within a layered structure:

- Memetic theory describes transmission of representations,
- The moral landscape describes experiential projection,
- Game theory describes equilibrium at a descriptive layer,
- The present framework describes generative selection of invariant structure.

These frameworks correspond to different levels of description of a shared underlying process.

10 Implications

The unified framework yields several implications:

- Structure formation across domains follows a common mechanism.
- Stability is determined by invariance under constraint, not optimization alone.
- Ethical systems emerge as structural necessities rather than prescriptions.

11 Limits and Scope

This framework is structural and interpretive. It does not specify:

- the detailed form of configuration spaces,
- the explicit form of collapse operators,
- quantitative predictive dynamics.

Its purpose is to unify patterns observed across domains, not to replace existing formalisms.

12 Testable Implications

While the collapse-selection framework is primarily structural and interpretive, it yields a number of testable implications across domains. These do not require direct observation of the collapse operator itself, but instead concern observable consequences of invariant preservation and persistence defect.

12.1 Prediction of Instability Through Persistence Defect

The framework predicts that configurations exhibiting increasing persistence defect

$$T(x) \uparrow$$

will undergo structural reconfiguration or collapse.

In practice, this implies that systems displaying degradation of invariants—such as trust, coordination, or coherence—should exhibit measurable precursors to failure, including:

- increased volatility or sensitivity to perturbation,
- breakdown of coordination across subsystems,
- divergence between local performance metrics and system-level stability.

12.2 Divergence Between Local Optimization and Global Stability

The framework predicts systematic divergence between locally optimal configurations and globally stable configurations.

This can be tested by comparing:

- configurations that maximize local payoff or performance,
- configurations that maximize persistence under repeated interaction.

In systems where these differ, the framework predicts that locally optimal configurations will fail under extended interaction, while persistence-aligned configurations will stabilize.

12.3 Cross-Context Convergence of Ethical Structure

The definition of ethical invariance as cross-class stability implies that certain relational structures should recur across independent systems.

This predicts that:

- cooperative, trust-preserving configurations will emerge in diverse cultural and institutional settings,
- configurations that degrade persistence will fail to generalize across contexts,
- stable ethical norms will correspond to structures invariant under heterogeneous collapse regimes.

12.4 Reconstruction-Based Transfer of Structure

The framework predicts that structural properties can be transferred across systems without direct semantic transmission.

This may be tested in model-based systems by:

- training on data lacking explicit semantic content,
- evaluating whether structural or behavioral properties are nonetheless reproduced,
- comparing transfer effectiveness as a function of structural compatibility between systems.

12.5 Emergence of Fixed-Point Structure in Iterative Reasoning

The framework predicts that iterative reasoning processes—whether human or model-based—should converge toward stable configurations satisfying:

$$\Phi(x) \approx x.$$

Empirically, this corresponds to:

- stabilization of solutions under repeated reformulation,
- convergence of independently derived representations,
- resistance of certain structures to further modification.

12.6 Interpretation

These implications do not depend on the explicit form of the collapse operator, but on the structural consequences of selection under constraint.

Accordingly, the framework may be evaluated indirectly through its ability to explain and predict patterns of stability, failure, convergence, and transfer across domains.

13 Conclusion

We have proposed a unified framework in which structure across cognition, machine learning, systems, and ethics is understood as the invariant residue of collapse-selection dynamics.

Within this perspective, intuition, reconstruction, persistence, and ethical structure are not separate phenomena, but manifestations of a single generative process:

$$\text{Constraint} \to \text{Collapse (Selection)} \to \text{Persistence} \to \text{Invariant Structure.}$$

This framework provides a coherent account of how structure forms, transfers, stabilizes, and is selected across scales.

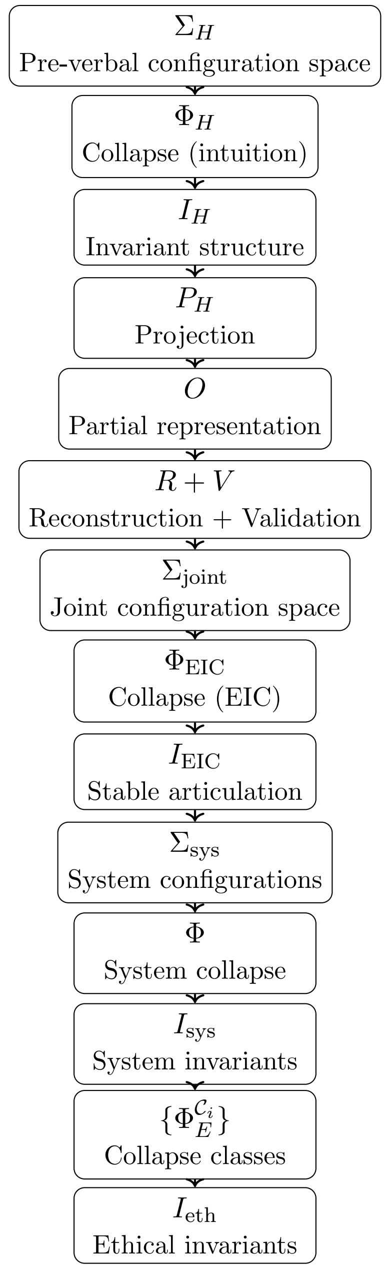

Figure 1: Unified collapse-selection pipeline across scales. Structure is generated in relational configuration spaces, selected via collapse operators, projected into observable representations, reconstructed across systems, and stabilized as invariant structure. At higher scales, these invariants define system-level persistence and ethical structure.