Darwinian Ethics as Collapse-Stable Invariant Structure in Relational Cognitive Space

Stephen Garner

April 20, 2026

Abstract

We propose a structural formulation of ethics within the framework of Quantum Collapse Geometry (QCG), interpreting ethical systems as collapse-stable invariant structures arising from memetic dynamics. Rather than treating ethics as externally imposed rules or subjective preference, we model ethical configurations as relational structures subject to constraint-driven selection under repeated interaction.

Within this framework, ethical norms correspond to invariant families that persist across agents, contexts, and scales. Collapse dynamics eliminate unstable configurations that degrade coordination, trust, or system persistence, while reinforcing those that maintain coherence. This yields a Darwinian interpretation of ethics in which stability under constraint, rather than prescription, defines ethical structure. The result is a unified account of ethical emergence consistent with QCG's generative ontology of selection, persistence, and invariant structure.

1 Introduction

Across social systems, ethical structures exhibit a recurring pattern: certain behaviors persist, while others degrade coordination and are eliminated over time. Traditional accounts interpret ethics either as normative prescriptions or as culturally contingent constructs.

In contrast, we propose that ethical structure may be understood as an emergent consequence of selection under constraint. Within the framework of Quantum Collapse Geometry (QCG), stable structure arises through repeated application of admissibility conditions acting on relational configurations. We extend this framework to the domain of cognition and social interaction, interpreting ethical systems as invariant structures selected under collapse dynamics.

2 Relational Ethical Configuration Space

Let Σ_E denote a relational configuration space whose elements

$$E \in \Sigma_E$$

encode patterns of behavior, norms, and interaction structures across agents.

These configurations are not evaluated in isolation, but through their effect on system-level persistence under repeated interaction.

Define an admissible set

$$A_E \subseteq \Sigma_E$$

such that $E \in A_E$ if and only if it supports persistence under constraint, including:

- maintenance of coordination,
- preservation of trust,
- stability under repeated interaction,
- compatibility with multi-agent coexistence.

3 Collapse-Selection Dynamics

Let

$$\Phi_E : \Sigma_E \to \Sigma_E$$

be a collapse operator representing selection under constraint induced by interaction, feedback, and consequence.

Repeated application of Φ_E eliminates unstable configurations and reinforces those that maintain persistence.

Define the invariant set:

$$\mathrm{Inv}(\Phi_E) = \{E \in \Sigma_E \mid \Phi_E(E) = E\}.$$

These invariant configurations correspond to stable ethical structures.

4 Definition of Darwinian Ethics

Definition. A Darwinian ethical structure is an invariant family

$$\mathcal{I}_E \subseteq A_E$$

such that:

1. $\Phi_E(E) \approx E$ for all $E \in \mathcal{I}_E$,
2. E remains stable under repeated multi-agent interaction,
3. E supports persistence of the system across scales.

Ethical structure is thus identified with persistence under constraint rather than external prescription.

5 Selection of Ethical Invariants

Proposition. Let $E \in \Sigma_E$ be an initial configuration. Suppose:

1. interaction induces repeated application of Φ_E,
2. configurations that degrade persistence are suppressed,
3. configurations that maintain coherence are reinforced.

Then, for a class of systems,

$$\lim_{n\to\infty} \Phi_E^n(E) = E^*$$

exists, where E^* is a collapse-stable ethical invariant.

6 Local and Global Stability in Ethical Structure

The collapse-selection framework naturally distinguishes between locally stable and globally stable ethical configurations. This distinction is essential for explaining the persistence of both cooperative and pathological systems within the same ontology.

6.1 Local Stability

An ethical configuration $E \in \Sigma_E$ is said to be locally stable if:

- it persists under limited interaction,
- it remains stable over short time horizons,
- it satisfies admissibility conditions within a restricted domain of agents or contexts.

Formally, E is locally stable if there exists a restricted domain $D \subset \Sigma_E$ such that:

$$\Phi_E(E)|_D \approx E|_D.$$

Local stability does not guarantee persistence under extension beyond D.

6.2 Global Stability

An ethical configuration is globally stable if:

- it persists under scale extension,
- it remains stable under repeated interaction across heterogeneous agents,
- it preserves system-level coherence over long time horizons.

Formally, E is globally stable if:

$$\Phi_E(E) \approx E$$

holds under arbitrary admissible extensions of the interaction domain.

6.3 Local Attractors and Global Failure

Certain configurations satisfy local stability while failing global stability. These include:

- exploitative or coercive systems,
- rigid hierarchical structures,
- trust-degrading interaction patterns.

Such configurations function as *local attractors* of Φ_E but do not correspond to invariant families under extended collapse dynamics.

Under iteration, these systems exhibit:

- degradation of coordination,
- erosion of trust,
- accumulation of instability across layers,
- eventual collapse or fragmentation.

Thus, their persistence is conditional and scale-limited.

6.4 Mechanism of Failure Under Extension

The failure of locally stable configurations under global extension arises from:

1. **Constraint mismatch:** local admissibility conditions do not generalize across agents or contexts,
2. **Information loss:** compressed structure fails to encode necessary invariants for larger-scale coordination,
3. **Feedback amplification:** small instabilities accumulate under repeated interaction,
4. **Invariant violation:** trust, identity, or continuity constraints are degraded.

These mechanisms reflect the general QCG principle that invariance must be preserved under admissible transformations to sustain persistence.

6.5 Interpretation

This distinction explains the coexistence of:

- short-term success of unethical or exploitative strategies,
- long-term persistence of cooperative and trust-aligned systems.

Ethical systems that appear advantageous locally may be eliminated under broader collapse dynamics, while globally stable structures emerge as invariant families across scales.

6.6 Relation to Darwinian Ethics

Within Darwinian Ethics, this yields a refinement:

- ethical validity corresponds to global invariance under collapse,
- unethical behavior corresponds to locally stable but globally unstable configurations.

Thus, ethics is not defined by immediate outcome or preference, but by persistence under extended constraint.

6.7 Conclusion

The distinction between local and global invariance is necessary to account for the observed diversity of ethical systems. It allows the framework to accommodate both stability and failure without contradiction.

In this view, ethical structure is identified not with what temporarily persists, but with what remains invariant under repeated collapse across scales.

7 Global Ethical Invariance Across Collapse Classes

The distinction between local and global stability may be formalized using the notion of collapse classes.

Let $\mathcal{C}$ denote a family of admissible collapse classes, where each

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

represents a distinct regime of interaction, context, or scale.

For each collapse class $\mathcal{C}_i$, define the corresponding invariant set:

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

Definition (Global Ethical Invariance). An ethical configuration E is globally invariant if:

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

That is, E remains stable under all admissible collapse dynamics across contexts and scales.

7.1 Interpretation

This condition refines the definition of Darwinian Ethics:

- **Local invariants** satisfy stability under a restricted subset of collapse classes,
- **Global invariants** satisfy stability across the full admissible family $\mathcal{C}$.

Thus:

- exploitative or coercive systems may satisfy

  $$E \in \mathrm{Inv}(\Phi_E^{\mathcal{C}_{\mathrm{local}}})$$

  for some restricted class $\mathcal{C}_{\mathrm{local}}$,
- but fail to satisfy

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

  leading to instability under broader interaction.

7.2 Proposition (Selection of Globally Ethical Structure)

Let $E \in \Sigma_E$ be an initial configuration. Suppose:

1. interaction induces exposure to a family of collapse classes $\mathcal{C}$,
2. configurations unstable under any $\Phi_E^{\mathcal{C}_i}$ are eliminated,
3. configurations stable across all $\mathcal{C}_i$ persist.

Then the asymptotic structure satisfies:

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

and defines a globally stable ethical invariant.

7.3 Relation to Invariant Families

This formulation aligns ethical structure with invariant families across collapse classes.

Ethical systems are not identified with invariants of a single collapse operator, but with relational structures preserved across admissible classes of collapse dynamics.

Thus, ethical validity corresponds to cross-class invariance rather than operator-specific stability.

7.4 Conclusion

The requirement of invariance across collapse classes provides a formal criterion distinguishing globally stable ethical structure from locally stable configurations.

It explains the persistence of cooperative and trust-aligned systems, and the eventual failure of systems that are stable only within restricted domains.

In this view, ethics is identified with the intersection of invariant structure across admissible collapse regimes.

8 Distance from Global Invariance and Ethical Instability

The criterion of global ethical invariance identifies configurations that remain stable across all admissible collapse classes. To refine this distinction, we introduce a QCG-native measure of distance from global invariance based on persistence defect.

Let $\mathcal{C} = \{\mathcal{C}_i\}$ denote a family of admissible collapse classes, with associated collapse operators

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

For each collapse class $\mathcal{C}_i$, define a *persistence defect functional*

$$T_{\mathcal{C}_i}(E) \geq 0,$$

which measures the degree to which E fails to remain invariant under $\Phi_E^{\mathcal{C}_i}$, i.e. deviation from collapse-stability under admissible transformations.

Interpretation:

- $T_{\mathcal{C}_i}(E) = 0$ if and only if E is invariant under $\Phi_E^{\mathcal{C}_i}$,
- $T_{\mathcal{C}_i}(E) > 0$ quantifies the structural tension induced by collapse under $\mathcal{C}_i$,
- larger values indicate greater instability and required reconfiguration.

We then define the *global persistence defect* by

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

or, when appropriate,

$$\overline{T}_{\text{glob}}(E) = \frac{1}{|\mathcal{C}|} \sum_{\mathcal{C}_i \in \mathcal{C}} T_{\mathcal{C}_i}(E).$$

Definition (Ethical Instability). The ethical instability of a configuration $E \in \Sigma_E$ is measured by $T_{\text{glob}}(E)$ (or $\overline{T}_{\text{glob}}(E)$), representing its distance from global invariance.

Interpretation. This quantity measures the extent to which E fails to satisfy:

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

Thus:

- $T_{\mathrm{glob}}(E) = 0$ implies global ethical invariance,
- small $T_{\mathrm{glob}}(E)$ indicates near-invariant but fragile structure,
- large $T_{\mathrm{glob}}(E)$ indicates strong dependence on restricted collapse classes.

Local versus Global Stability. A configuration may satisfy

$$T_{\mathcal{C}_j}(E) = 0$$

for some restricted class $\mathcal{C}_j$, while still having

$$T_{\mathrm{glob}}(E) > 0.$$

Such configurations are locally stable but globally unstable.

This captures systems that persist under narrow conditions (e.g., coercive or exploitative structures) but fail under broader interaction across agents or scales.

Proposition (Characterization of Global Ethical Structure). A configuration E is globally ethically invariant if and only if

$$T_{\mathrm{glob}}(E) = 0.$$

Proof Sketch. If $T_{\mathrm{glob}}(E) = 0$, then $T_{\mathcal{C}_i}(E) = 0$ for all $\mathcal{C}_i$, implying

$$\Phi_E^{\mathcal{C}_i}(E) = E$$

for all collapse classes, hence global invariance.

Conversely, if E lies in the intersection of all invariant sets, then all persistence defects vanish, and $T_{\mathrm{glob}}(E) = 0$.

Interpretive Consequence. The functional $T_{\mathrm{glob}}(E)$ provides a graded notion of ethical instability, allowing one to distinguish:

- globally stable ethical invariants,
- near-stable but fragile configurations,
- locally persistent but globally unstable systems,
- configurations requiring extensive collapse to restore admissibility.

Conceptual Interpretation. The quantity $T_{\mathrm{glob}}(E)$ measures the amount of corrective collapse required for E to remain viable across contexts.

In this sense, ethical instability is identified with unresolved relational tension under admissible collapse dynamics, and measures how much corrective collapse a configuration would require in order to remain viable across contexts.

9 Collective Moral Memory

Let Σ_M denote the space of memetic representations of ethical structure.

Define a memetic collapse operator

$$\Phi_M : \Sigma_M \to \Sigma_M$$

acting through communication, interpretation, and cultural transmission.

Collective moral memory is given by the collapse-stable invariant:

$$M^* = \lim_{n\to\infty} \Phi_M^n(M),$$

representing a coarse-grained ethical structure emerging across distributed cognitive systems.

10 Interpretation

Within this framework:

- cooperation emerges as a persistence-stable configuration,
- trust functions as an invariant required for coordination,
- destabilizing behaviors are eliminated under repeated collapse,
- ethical norms correspond to constraints preserving system viability.

Ethics is therefore not external to system dynamics, but a structural consequence of persistence under constraint.

Global invariance is a necessary condition for ethical structure, though not sufficient to fully characterize its experiential qualities.

11 Relation to Memetic and Moral Landscape Frameworks

The present formulation may be situated in relation to two influential approaches to ethics and cultural evolution: memetic theory and the moral landscape framework.

11.1 Relation to Memetic Theory

Memetic theory, as proposed in the replicator-based framework of *The Selfish Gene*[1], treats cultural units ("memes") as replicators subject to variation, selection, and transmission. In this view, cultural evolution proceeds analogously to biological evolution, with successful memes propagating through populations.

The present framework is compatible with this perspective but differs in emphasis. Rather than treating memes as primitive replicators, we interpret them as observable projections of underlying relational configurations. Memetic persistence is not attributed solely to replication, but to stability under collapse-selection dynamics.

In particular:

- Memetic theory emphasizes replication and transmission,
- the present framework emphasizes admissibility, collapse, and invariant structure.

Thus, memes are reinterpreted as coarse-grained representatives of invariant families in relational cognitive space, rather than as fundamental units.

11.2 Relation to the Moral Landscape

The moral landscape framework, as developed in *The Moral Landscape*[2], proposes that ethical systems may be evaluated in terms of the well-being of conscious agents. In this view, moral states correspond to regions in a high-dimensional space of possible experiences, with peaks representing maximal well-being.

The present framework does not reject this interpretation, but situates it at the level of descriptive projection.

Within the collapse-selection ontology:

- relational configurations $E \in \Sigma_E$ define the generative structure,
- projection maps $P : \Sigma_E \to O$ encode observable or experiential states,
- well-being corresponds to features of the projected structure.

Thus, well-being may be interpreted as an observable property associated with collapse-stable configurations, rather than as the primitive criterion of ethical validity.

11.3 Synthesis

These relationships may be summarized as follows:

Framework	Structural Role
Memetic theory	Describes transmission and selection of cultural representations
Moral landscape	Describes experiential or phenomenological evaluation
Present framework	Models generative selection of stable relational structure

From this perspective:

- memetic theory describes how structures propagate,
- the moral landscape describes how they are experienced,
- the present framework describes how they are generated and stabilized.

11.4 Interpretation

The collapse-selection formulation does not replace these frameworks, but provides a generative layer beneath them.

Memetic replication may be understood as the transmission of coarse-grained invariant structure, while moral evaluation may be understood as a projection of that structure into experiential space.

In this way, both frameworks are recovered as descriptive layers within a broader generative architecture defined by constraint, selection, and persistence.

12 Bridge to Layered Emergence and Persistence Across Scales

The formulation of Darwinian Ethics developed in this work admits a direct correspondence with the framework presented in *Layered Emergence and the Limits of Control.* While the two treatments differ in emphasis—one formal, the other interpretive—they describe the same underlying structure at different levels of resolution.

12.1 Invariants and Collapse-Stable Structure

In the layered emergence framework, stable systems are characterized by invariants: constraint structures that preserve coherence across time and scale. These include trust, identity, continuity, and coordination, which function as the structural conditions for persistence.

Within the present formulation, these invariants correspond to collapse-stable configurations:

$$E \in \mathrm{Inv}(\Phi_E),$$

and, more strongly, to structures invariant across collapse classes:

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

Thus, the notion of invariants across layers is equivalent to cross-class invariance under admissible collapse dynamics.

12.2 Layer Mismatch and Collapse Instability

The layered emergence framework identifies a primary failure mode: the application of models or control structures from one layer to another. This produces a degradation of invariants, leading to instability.

In the collapse-selection formulation, this corresponds to an increase in persistence defect:

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

indicating that a configuration fails to remain stable under one or more collapse classes.

Layer mismatch may therefore be interpreted as a structural incompatibility between collapse regimes, in which a configuration that is locally stable becomes globally unstable under extended interaction.

12.3 Persistence and Ethical Structure

In both frameworks, persistence is the defining criterion of stability:

- In layered emergence, systems persist when invariants are preserved across scale.
- In Darwinian Ethics, ethical structures persist when they remain invariant under collapse across contexts.

These statements are structurally equivalent.

Thus, ethical systems are not imposed constructs, but the invariant structures that remain when configurations are subjected to constraint-driven selection across layers.

12.4 System Failure as Collapse Dynamics

The degradation of systems described in the layered emergence framework—loss of trust, fragmentation of identity, breakdown of coordination—corresponds directly to the failure of invariance under collapse dynamics.

In this view:

- erosion of trust corresponds to violation of relational invariants,

- fragmentation corresponds to failure of coherence under projection,
- instability corresponds to accumulation of persistence defect.

System collapse is therefore not an external event, but the natural consequence of unresolved relational tension:

$$T_{\text{glob}}(E) \uparrow \;\Rightarrow\; \text{collapse likelihood increases.}$$

12.5 Unified Interpretation

Taken together, the two frameworks describe a single generative structure:

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

The layered emergence formulation provides an intuitive account of this process across domains, while the Darwinian Ethics formulation provides a formal characterization in terms of relational configuration space, collapse operators, and invariant families.

12.6 Conclusion

The correspondence established here shows that the behavior of real-world systems across scale—coordination, stability, failure, and recovery—is not separate from ethical structure, but is its direct manifestation.

Ethics, in this unified view, is the invariant structure that remains when systems are subjected to constraint-driven selection across layers of reality.

13 Conclusion

We have proposed a formulation of Darwinian Ethics within the framework of Quantum Collapse Geometry, identifying ethical systems as collapse-stable invariant structures over relational cognitive space.

This perspective reframes ethics as an emergent property of selection under constraint, unifying cognitive, social, and physical domains under a shared generative architecture:

$$\text{Constraint} \to \text{Selection} \to \text{Persistence} \to \text{Structure.}$$

In this view, ethical systems are not imposed, but discovered as the invariant structures that remain when unstable modes of interaction are eliminated.

References

[1] Richard Dawkins. *The selfish gene*. Oxford university press, 2016.

[2] Sam Harris. *The moral landscape: How science can determine human values*. Simon and Schuster, 2010.