

Layered Emergence and the Limits of Control: Toward a Theory of Persistence Across Complex Systems

Stephen Garner

April 4th 2026

Introduction — The Problem of Scale and the Failure of Control

Systems work at one scale, and then break at the next.

Families can be cohesive, resilient, and deeply aligned—then fracture when extended outward. Communities can coordinate and cooperate—then dissolve into abstraction at the level of institutions. Nations attempt to impose order—then generate instability through the very mechanisms meant to stabilize them.

The common explanation is usually some version of human failure: bad actors, poor leadership, corruption, lack of education, insufficient enforcement. Sometimes that's true. But it's not the root cause.

The deeper issue is structural.

We consistently attempt to control systems at one layer using the tools, assumptions, and models of another. We treat scale as a quantitative problem when it is fundamentally a qualitative transition.

This paper argues that the failure is not primarily moral, political, or economic. It is epistemic.

No system can fully describe or control the processes that generate it.

Each layer of reality emerges from a substrate it cannot fully resolve. This limitation defines what can be known, what can be controlled, and what must instead be allowed to emerge.

The central claim is:

Stable systems do not arise from optimal strategies or ideal actors. They arise from constraint structures that align behavior with persistence across layers that cannot fully perceive one another.

1 Layered Reality and Emergence

Core Definitions

Emergence: The formation of higher-level structures through compression and stabilization of lower-level dynamics.

Persistence: The capacity of a structure to maintain coherence across time and perturbation.

Invariants: Constraints required for system stability.

Constraint Structure: The configuration of invariants and feedback mechanisms shaping viable behavior.

Reality is layered.

Physics → Biology → Cognition → Social systems → Institutions.

At each step, the system compresses. Many underlying states become indistinguishable. Information about generation is lost.

Emergence is not just complexity—it is compression.

Each layer inherits from the one beneath it, but cannot fully describe it.

This creates separation. It allows independent reasoning at each level—but introduces limits.

We tend to treat the layer we operate in as complete.

That assumption is wrong.

2 The Epistemic Horizon

No layer can fully see the one that generates it.

This is structural, not temporary.

An **epistemic horizon** is the boundary beyond which generative mechanisms cannot be reconstructed.

If compression maps many states to one, the mapping cannot be inverted.

Information is not hidden—it is gone.

This applies across domains:

- Consciousness cannot fully access its own generation

- Markets cannot see internal motivations
- Institutions cannot fully resolve individuals

A system can often predict behavior it cannot explain.

This creates a critical limitation:

You cannot reliably control a system whose generative dynamics you cannot fully resolve.

Thus, control must shift to constraints.

3 Invariants Across Layers

What holds systems together is not optimization.

It is invariants.

Each layer has its own:

- Physics → coherence, conservation
- Biology → metabolic viability, reproduction
- Cognition → perceptual coherence
- Social systems → trust, identity, continuity
- Institutions → legitimacy, coordination

These are not interchangeable.

Violating invariants may produce short-term gains, but causes long-term instability.

Optimization often targets visible variables, not invariants.

Thus:

A system can improve performance while degrading its ability to persist.

4 The Failure Mode: Layer Mismatch

Layer mismatch occurs when tools from one layer are applied to another.

Examples:

- Incentives applied to identity-driven systems
- Institutional logic applied to families

- Tribal logic applied to nations

Initial effects:

- Increased efficiency
- Improved metrics

Long-term effects:

- Erosion of invariants
- Increased fragility

This produces a reinforcing cycle:

Mismatch \rightarrow Partial Success \rightarrow Overconfidence \rightarrow Amplified Mismatch

5 Rethinking Game Theory

Classical game theory assumes:

- Fixed players
- Fixed rules
- Known strategies
- Stable payoffs

In reality:

- Players emerge
- Rules evolve
- Strategies are discovered
- Payoffs depend on system persistence

Thus, equilibrium is insufficient.

We introduce:

Collapse Interpretation of Strategic Interaction

Let:

X = space of strategy configurations

$T(x)$ = tension / instability measure

Define collapse operator:

$$\Phi : X \rightarrow X$$

such that unstable configurations are suppressed and stable ones persist.

Stability condition:

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

Thus, stable strategies are fixed points of collapse dynamics.

Persistence-Weighted Payoff

Standard payoff:

$$U_i^{\text{local}}$$

Extended payoff:

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

where:

- $P(x)$ = persistence of system under configuration x
- λ = weighting parameter

Agents optimize:

- short-term reward
- system survivability

6 Ethics as a Stability Constraint

Ethics is not preference.

It is a structural constraint.

Behaviors that degrade invariants:

- erode trust
- fragment identity

- destabilize coordination

These are selected against over time.

Ethical behavior aligns with persistence.

Thus:

Ethics is the boundary condition for system stability.

7 Constraint Structures vs Agent Trust

Systems do not stabilize by trusting individuals.

They stabilize by shaping behavior.

Constraint structures:

- make stable behavior easy
- make destabilizing behavior costly

Thus:

The goal is not to find perfect people, but to design systems where imperfect people produce stable outcomes.

8 Case Study — Family System Reconstruction

8.1 Initial Conditions

Stable system with shared identity and trust.

8.2 Disruptive Agent

Introduction of adversarial behavior.

8.3 Standard Failures

- Direct conflict → escalation
- Calling out → polarization

8.4 Intervention

- Increase legibility
- Reinforce continuity
- Align incentives with cohesion
- Maintain inclusion

8.5 Outcome

Behavior constrained, not eliminated.

System stability restored.

9 Scaling the Principle

Small systems have:

- visibility
- tight feedback
- shared identity

Large systems lose these.

Control replaces generative conditions.

Failure modes:

- over-centralization
- fragmentation

Solution:

- restore legibility
- preserve continuity
- align incentives with invariants

10 Governance, Secrecy, and Power

Secrecy reduces:

- visibility
- feedback
- correction

If misaligned:

Secrecy → Constraint Violation → Loss of Trust → More Control

Solution: Align internal structure with persistence.

11 The Principle of Layer-Appropriate Action

Three constraints:

1. Operate within layer invariants
2. Do not directly control lower layers
3. Preserve generative capacity below

12 Implications and Predictions

- Short-term optimization precedes instability
- Layer mismatch produces oscillation
- Abstraction increases measurable performance but reduces stability
- Ethics predicts long-term persistence
- Trust emerges from structure, not messaging
- Misaligned secrecy produces instability

Conclusion

The future of stable civilization will not be built on better control, but on deeper alignment with the structures that allow systems to endure.

13 Appendix A: QCG Formalization

13.1 A.0 Scope and Positioning

This appendix provides a minimal formal instantiation of Quantum Collapse Geometry (QCG), focusing on the definition of the state space, the collapse operator, and the associated invariant structures.

The goal is not to present a complete physical theory, but to establish a consistent mathematical framework within which collapse can be treated as a structural operation acting on relational configurations. Subsequent sections define the operator, identify its invariant sector, and demonstrate its behavior in a representative model.

Throughout, emphasis is placed on properties that can be expressed in standard mathematical language (operators, invariants, and mappings), enabling comparison with existing frameworks such as decoherence, projection operators, and coarse-graining.

A.1 State Space Definition

We define a relational state space Σ , representing configurations of interacting degrees of freedom. The structure of Σ admits both discrete and continuous realizations, depending on the level of description.

Discrete representation. In a minimal model, let $G = (V, E)$ be a graph with vertices V and edges E . Each vertex $i \in V$ is assigned a phase variable:

$$\theta_i \in S^1,$$

so that the state space is given by:

$$\Sigma = (S^1)^{|V|}.$$

Local interactions are defined by adjacency in G , and configurations are specified by assignments $\{\theta_i\}_{i \in V}$.

Continuous representation. At a more general level, Σ may be modeled as a manifold or configuration space of relational degrees of freedom:

$$\Sigma \supset M_k \subset \mathbb{C}^n,$$

where M_k denotes an embedded submanifold representing a coherence-stable configuration.

In this setting, relational structure is encoded in phase relations, amplitudes, or more general fields defined over M_k .

Observational projection. We assume the existence of a projection map:

$$P : \Sigma \rightarrow \mathcal{O},$$

where \mathcal{O} denotes a space of observable quantities. The map P is generally non-invertible, reflecting loss of information in passing from relational configurations to observables.

Remarks. The dual presentation of Σ allows for both computational models (discrete lattice systems) and analytical treatments (manifold-based descriptions). In both cases, the essential feature is that Σ encodes relational structure, rather than independent degrees of freedom, and supports the action of a collapse operator as defined in Section A.3.

A.2 Local Dynamics

We assume the existence of a local dynamical evolution acting on the relational state space:

$$F : \Sigma \rightarrow \Sigma,$$

which updates configurations according to interactions defined on the underlying structure (e.g., adjacency in a graph or locality in a manifold).

Discrete realization. In the discrete phase model introduced in Section A.1, a representative local update rule is given by:

$$\theta_i \mapsto \arg \left(\sum_{j \sim i} e^{i\theta_j} \right),$$

where the sum is taken over neighbors of site i in the graph G . This defines a local alignment or averaging process, analogous to dynamics in coupled phase systems.

Continuous realization. In a continuous setting, F may be modeled as a flow on Σ , generated by local interaction terms or variational principles. For example, one may consider gradient flows or Hamiltonian evolution of fields defined on $M_k \subset \Sigma$, with dynamics determined by local couplings.

Compatibility with collapse. The local dynamics F and the collapse operator Coll (Section A.3) are related but distinct. While F describes the evolution of configurations under interaction, Coll acts to suppress incoherent structure and select stable configurations.

In general, these operations do not commute exactly, but satisfy a compatibility relation of the form:

$$\text{Coll} \circ F \approx \text{Coll},$$

indicating that local fluctuations introduced by F are absorbed by collapse unless they preserve coherence.

Interpretation. The decomposition into local dynamics and collapse provides a two-stage description of system evolution:

- F : generates local variation and interaction,
- Coll : selects configurations that remain stable under these interactions.

This separation allows one to distinguish between generative processes (dynamics) and stabilizing processes (collapse), with invariant structures arising from their interplay.

A.3 Collapse Operator

We introduce a collapse operator acting on the relational state space

$$\text{Coll} : \Sigma \rightarrow \Sigma,$$

which maps configurations to a reduced set of coherence-stable states.

The operator is defined abstractly as a composition of a decoherence map and a selection process:

$$\text{Coll} = \mathcal{S} \circ \mathcal{D},$$

where:

- $\mathcal{D} : \Sigma \rightarrow \Sigma$ is a decoherence-like map which suppresses incompatible phase relations and projects the system toward a coherence-compatible subspace,
- $\mathcal{S} : \Sigma \rightarrow \Sigma$ is a selection map that identifies and retains configurations stable under local interaction constraints.

Operational interpretation. The action of Coll can be viewed as a projection onto a subset of states that are stable under repeated application of local dynamics. It is therefore natural to interpret Coll as a coherence-selection operator.

Structural properties. The collapse operator satisfies the following properties:

1. **Idempotence.** Repeated application does not further change the state:

$$\text{Coll} \circ \text{Coll} = \text{Coll}.$$

2. **Non-invertibility.** The operator is not invertible:

$$\text{Coll}^{-1} \text{ does not exist in general,}$$

reflecting loss of information under collapse.

3. **Information reduction.** The effective dimension of the image is reduced:

$$\dim(\text{Im}(\text{Coll})) \leq \dim(\Sigma),$$

with strict inequality in nontrivial cases.

4. **Stability selection.** The image of Coll defines the set of collapse-stable configurations:

$$\mathcal{I} = \{x \in \Sigma \mid \text{Coll}(x) = x\}.$$

5. **Locality compatibility.** If $F : \Sigma \rightarrow \Sigma$ denotes local dynamics, then Coll is compatible with F in the sense that:

$$\text{Coll} \circ F \approx \text{Coll},$$

i.e. collapse absorbs local fluctuations that do not preserve coherence.

Discrete realization. In a discrete phase model with variables $\theta_i \in S^1$ defined on a graph G , a representative realization of the collapse operator is given by:

$$\theta_i \mapsto \arg \left(\sum_{j \sim i} e^{i\theta_j} \right),$$

where the sum runs over neighbors of site i . This operation locally aligns phases while preserving global topological structure.

Relation to projection operators. The operator Coll is structurally analogous to a projection onto a stable subspace, but differs in that the target subspace is not fixed *a priori*; it emerges from the interaction between local dynamics and coherence constraints.

Categorical formulation. In a categorical setting, let $Q = \text{CPM}(\mathbf{FHilb})$ and let $\mathcal{D} : Q \rightarrow Q$ be an idempotent decoherence functor selecting a coherence-stable subalgebra. Then the collapse operator corresponds to a natural transformation

$$\eta : \text{Id}_Q \Rightarrow \text{Coll},$$

where Coll acts as a structure-preserving, idempotent endomorphism. Under standard assumptions, Coll is natural and monoidal up to environment handling.

Interpretation. Within QCG, collapse is not introduced as an external postulate but arises as a structural operation that selects coherence-compatible configurations. The invariant set \mathcal{I} defined by Coll forms the basis for subsequent emergent structure.

A.4 Invariant Structures

Given the collapse operator

$$\text{Coll} : \Sigma \rightarrow \Sigma,$$

we define the invariant sector as the set of configurations that remain unchanged under its action:

$$\mathcal{I} = \{x \in \Sigma \mid \text{Coll}(x) = x\}.$$

Interpretation. The set \mathcal{I} represents the collapse-stable sector of the state space. Elements of \mathcal{I} are fixed points of the collapse operator and correspond to configurations that persist under repeated application of collapse.

In this sense, \mathcal{I} defines an invariant subspace (or, more generally, an invariant subset) of Σ , characterized by stability under coherence-selection dynamics.

Structural role. The invariant sector provides a natural decomposition of the state space into:

- stable configurations $x \in \mathcal{I}$,
- unstable configurations $x \notin \mathcal{I}$, which are mapped toward \mathcal{I} under collapse.

Thus, the collapse operator induces a partition of Σ into equivalence classes defined by their convergence under repeated application of Coll .

Relation to known structures. Invariant configurations under collapse correspond to structures that remain stable under local modification. This is consistent with several established concepts:

- **Topological invariants:** global properties (e.g., winding number) that are preserved under continuous deformation,
- **Conserved quantities:** quantities that remain unchanged under system evolution or symmetry constraints.

Within the present framework, such quantities can be interpreted as features of configurations that lie within the invariant sector \mathcal{I} .

Summary. The invariant sector \mathcal{I} identifies the set of configurations that are stable under collapse. These configurations form the basis for emergent structure, as they represent the elements of the state space that persist under repeated interaction and coherence selection.

A.5 Toy Model (Discrete Lattice)

To provide a concrete instantiation of the collapse operator, we consider a discrete lattice model with phase-valued degrees of freedom.

Lattice and state space. Let

$$L \subset \mathbb{Z}^2$$

be a finite two-dimensional lattice with nearest-neighbor adjacency. At each site $i \in L$, we assign a phase variable:

$$\theta_i \in S^1.$$

A configuration is given by a phase field:

$$\theta : L \rightarrow S^1,$$

so that the state space is:

$$\Sigma = (S^1)^{|L|}.$$

Collapse rule. We define a local collapse update at each site $i \in L$ by:

$$\theta_i \mapsto \arg\left(\sum_{j \sim i} e^{i\theta_j}\right),$$

where the sum is taken over nearest neighbors $j \sim i$.

This rule aligns the phase at site i with the average direction of its neighbors, and can be interpreted as a local coherence-selection step.

Iterated collapse. The global collapse operator $\text{Coll} : \Sigma \rightarrow \Sigma$ is defined by applying the local update across all sites, either synchronously or asynchronously, and iterating until convergence.

Observed behavior. Under repeated application of the collapse rule:

- Local phase fluctuations are smoothed, reducing incoherent structure.
- Regions of approximately aligned phase expand over time.

At the same time, global topological features of the configuration are preserved.

Invariant structure. In particular, the winding number associated with phase configurations around nontrivial loops in the lattice remains unchanged under collapse.

Thus, the collapse operator exhibits the following behavior:

- **Local smoothing:** suppression of high-frequency phase variation,
- **Global preservation:** invariance of topological charge.

Interpretation. This model provides a minimal example in which collapse acts as a coherence-selection mechanism: it removes unstable local structure while preserving global invariants.

As such, it illustrates the general principle that collapse dynamics select stable configurations without erasing the invariant structures that characterize them.

A.6 Topological Invariant Preservation

A central property of the collapse operator is that, while it suppresses local incoherence, it preserves global topological invariants. This establishes collapse as a structure-selecting operation that reduces local variability without destroying large-scale organization.

Loop configurations and winding number. Let $\gamma \subset \Sigma$ denote a closed loop configuration in the relational state space. In a phase representation $\theta : \gamma \rightarrow S^1$, define the winding number:

$$w(\gamma) = \frac{1}{2\pi} \oint_{\gamma} d\theta,$$

which counts the net phase rotation along the loop.

Action of collapse. Under the collapse operator Coll , local phase configurations are smoothed according to the coherence-selection rule described in Section A.3. This induces a transformed loop:

$$\gamma' = \text{Coll}(\gamma).$$

Invariant preservation. Despite local smoothing, the winding number is preserved:

$$w(\text{Coll}(\gamma)) = w(\gamma).$$

This follows from the fact that the collapse operator acts locally to reduce phase variance but does not introduce or remove global phase discontinuities.

Sketch of argument. The collapse operator replaces each local phase θ_i with an averaged phase derived from neighboring values. Such local averaging operations are homotopy-preserving: they continuously deform the phase configuration without crossing singularities that would change the winding number.

Therefore, collapse acts as a homotopy within the space of configurations, preserving topological equivalence classes.

Generalization to higher invariants. The preservation property extends beyond simple winding numbers. Let $M_k \subset \Sigma$ denote an embedded coherence-stable manifold, and let ω be a differential form representing a topological invariant (e.g., a curvature form or Berry connection). Then, for admissible collapse dynamics:

$$\int_{M_k} \omega = \int_{\text{Coll}(M_k)} \omega,$$

provided that collapse does not introduce singularities or change the topological class of M_k .

Interpretation. This result indicates that collapse operates as a *local smoothing, globally invariant* transformation. It suppresses incoherent fluctuations while preserving the topological structure of the system.

In this sense, topological invariants can be interpreted as elements of the collapse-stable sector:

$$\mathcal{I} = \{x \in \Sigma \mid \text{Coll}(x) = x\},$$

with equivalence classes defined by homotopy under collapse dynamics.

Relation to physical systems. The preservation of topological invariants under local smoothing is consistent with behavior observed in topological phases of matter, where global invariants remain stable under local perturbations.

Within QCG, this suggests that such invariants arise as fixed structures under collapse, rather than as primary ontological inputs.

Summary. The collapse operator preserves global topological invariants while reducing local incoherence. This establishes a separation between:

- local structure, which is modified by collapse, and
- global structure, which remains invariant.

Accordingly, topological quantities emerge as natural invariants of collapse dynamics.

A.7 Relation to Established Frameworks

The collapse-first formulation of QCG is not intended as a replacement for existing physical frameworks, but as a structural reinterpretation that clarifies their domains of validity and the relationships between them.

Quantum measurement and projection. In standard quantum mechanics, measurement is represented by a projection or, more generally, by a completely positive (CP) map associated with an instrument. The QCG collapse operator shares formal similarities with such projections, in that it is idempotent and non-invertible, mapping states to a reduced subset.

However, in QCG the collapse operation is interpreted as a primary coherence-selection mechanism rather than as an external postulate tied to observation. The projection-like behavior arises from structural constraints on admissible configurations rather than from measurement alone.

Decoherence. Decoherence describes the suppression of phase coherence through interaction with an environment, leading to effective classical behavior. In QCG, decoherence corresponds to a component of the collapse process, modeled by a decoherence-like map \mathcal{D} .

The full collapse operator

$$\text{Coll} = \mathcal{S} \circ \mathcal{D}$$

extends this by including a selection step \mathcal{S} , which identifies coherence-stable configurations. Thus, decoherence alone is not sufficient; collapse includes both suppression and selection.

Topological phases of matter. Topological phases are characterized by invariants that remain stable under local perturbations. As shown in Section A.6, the collapse operator preserves global topological invariants while smoothing local structure.

Within QCG, such invariants are interpreted as elements of the collapse-stable sector, rather than as primary inputs. This provides a structural explanation for their robustness: they are precisely those features that survive repeated collapse.

Effective field theory. Effective field theories (EFTs) describe physics at a given scale by integrating out degrees of freedom at smaller scales. This process can be viewed as a form of coarse-graining.

In QCG, collapse plays an analogous role: it reduces the state space to configurations that are stable under interaction, effectively eliminating incompatible microstates. The resulting invariant structures define the degrees of freedom relevant at larger scales.

Renormalization and coarse-graining. Renormalization group (RG) flow describes how physical systems evolve under scale transformations, with fixed points corresponding to scale-invariant behavior.

The collapse operator can be interpreted as a structural analog of coarse-graining: it removes unstable fine-scale structure while preserving invariants. Collapse-stable configurations correspond to fixed structures under this process, providing a conceptual parallel to RG fixed points.

Categorical quantum mechanics. In categorical formulations of quantum theory, such as $\text{CPM}(\mathbf{FHilb})$, physical processes are represented as completely positive maps, and classical data is modeled using Frobenius algebras.

As described in Appendix ?? (or equivalent), the QCG collapse operator can be expressed as an idempotent, natural endomorphism arising from a decoherence functor and an instrument. This embeds the collapse process within an established mathematical framework.

Summary. Across these examples, existing frameworks can be understood as describing specific layers or aspects of a more general collapse-driven structure:

- projection operators capture the idempotent structure of collapse,
- decoherence captures coherence suppression,
- topological phases capture invariant structures,
- effective theories capture reduced degrees of freedom,
- renormalization captures scale-dependent stability.

QCG provides a unifying perspective in which these phenomena arise as different expressions of a common underlying operation: the selection of coherence-stable configurations under constraint.

A.8 Categorical Formulation

The collapse operator admits a natural formulation within categorical quantum mechanics, providing a structural embedding into an established mathematical framework.

CP maps and decoherence. Let \mathbf{FHilb} denote the category of finite-dimensional Hilbert spaces, and let

$$Q = \text{CPM}(\mathbf{FHilb})$$

be the category whose morphisms are completely positive (CP) maps.

In this setting, decoherence is modeled by an idempotent, structure-preserving map:

$$\mathcal{D} : Q \rightarrow Q,$$

which projects states onto a commutative subalgebra corresponding to classical data.

Collapse as an idempotent operator. The QCG collapse operator can be expressed as an idempotent endomorphism in Q :

$$\text{Coll} : A \rightarrow A,$$

satisfying:

$$\text{Coll} \circ \text{Coll} = \text{Coll}.$$

As in Section A.3, this operator may be decomposed as:

$$\text{Coll} = \mathcal{S} \circ \mathcal{D},$$

where \mathcal{D} suppresses incompatible coherence and \mathcal{S} selects configurations stable under interaction.

Invariant subobjects. The image of Coll defines a subobject corresponding to the collapse-stable sector:

$$\mathcal{I} = \text{Im}(\text{Coll}),$$

which is closed under the action of Coll and represents the set of admissible configurations in the categorical setting.

Remarks. This formulation situates the collapse operator within the standard framework of CP maps and decoherence processes. The categorical structure provides a consistent language for expressing idempotence, non-invertibility, and stability selection, without introducing additional assumptions beyond those already present in the formalism.

The categorical perspective is not required for the preceding construction, but demonstrates that the collapse operator can be embedded within an established mathematical setting.

A.9 Interpretation

Within the present formulation, the collapse operator Coll is interpreted as a coherence-selection mechanism acting on the relational state space Σ . Rather than introducing collapse as an external postulate, it arises here as a structural operation that suppresses incoherent configurations and selects those stable under repeated interaction. The resulting invariant set

$$\mathcal{I} = \{x \in \Sigma \mid \text{Coll}(x) = x\}$$

defines the collapse-stable sector of the system.

From this perspective, geometric and topological structure emerge from the organization of these invariant sectors. Local dynamics generate variation within Σ , while collapse selects configurations that persist, yielding stable relational structure at larger scales. Geometry may therefore be understood as an effective description of invariant structure under collapse, rather than as a primary input to the dynamics.

A.10 Summary

In this appendix, we have introduced a minimal formal framework for Quantum Collapse Geometry (QCG). A relational state space Σ was defined, together with local dynamics $F : \Sigma \rightarrow \Sigma$ and a collapse operator $\text{Coll} : \Sigma \rightarrow \Sigma$ acting as a coherence-selection mechanism. The structural properties of Coll , including idempotence, non-invertibility, and stability selection, were specified explicitly.

We identified the invariant sector

$$\mathcal{I} = \{x \in \Sigma \mid \text{Coll}(x) = x\},$$

and showed that collapse preserves global topological invariants while suppressing local incoherence. The framework was related to established structures, including projection operators, decoherence, topological phases, and coarse-graining, and a categorical formulation was provided to situate the construction within $\text{CPM}(\mathbf{FHilb})$.

This provides a minimal formal instantiation of QCG sufficient for further extension.

14 Appendix B: Emergent Game Theory

B.0 Purpose and Positioning

This appendix presents a reinterpretation of classical game theory within a collapse-driven, persistence-constrained dynamical framework. The aim is not to replace existing formulations, but to clarify the relationship between equilibrium descriptions and the generative processes that give rise to them.

In standard game theory, equilibrium concepts such as Nash equilibrium describe configurations that are stable under local deviations, assuming fixed players and payoff structures. These equilibria provide a useful descriptive layer, characterizing outcomes once interaction has stabilized.

The present framework introduces an additional layer in which strategic configurations evolve under a collapse operator and are evaluated according to a persistence functional. Within this setting, equilibria are understood as post-collapse descriptions of configurations that remain viable under repeated interaction and system-level constraints.

Accordingly, classical game theory may be viewed as operating at a descriptive equilibrium layer, while the emergent formulation developed here provides a generative and stability layer. The latter captures how configurations form, persist, and reorganize under constraint, rather than only describing their final stabilized form.

B.1 Classical Game Theory as a Descriptive Layer

We briefly recall the standard formulation of a strategic game. Let N denote a set of players, with each player $i \in N$ assigned a strategy set S_i . A strategy profile is given by:

$$s = (s_1, \dots, s_n) \in \prod_{i \in N} S_i,$$

and each player receives a payoff:

$$U_i(s_1, \dots, s_n).$$

A Nash equilibrium is a strategy profile s^* such that:

$$U_i(s^*) \geq U_i(s'_i, s^*_{-i}) \quad \text{for all } s'_i \in S_i,$$

where s^*_{-i} denotes the strategies of all players other than i .

This condition identifies configurations that are stable under unilateral deviation. In this sense, Nash equilibrium characterizes fixed points of local best-response dynamics.

Within the present framework, such equilibria are interpreted as *descriptive* objects: they capture configurations after interaction has stabilized, but do not specify the generative processes by which such configurations arise or persist under changing conditions.

B.2 Instability of Fixed-Player Models

Classical game-theoretic formulations typically assume a fixed player set N and a stable interaction structure. Under these assumptions, equilibria are computed over repeated or static instances of the same game.

In many real-world systems, however, these assumptions do not hold. In particular:

- the set of interacting agents may change over time,
- agent identities may not be stable across interactions,
- strategic choices can modify the structure of the interaction itself.

To capture this, we introduce a time-dependent player set:

$$N_t = \text{set of agents at time } t,$$

with evolution given by:

$$N_{t+1} = F(N_t, s_t),$$

where s_t denotes the strategy configuration at time t , and F encodes the effect of interaction on system composition.

Under this formulation, the game is no longer defined over a fixed set of players, but becomes a dynamical system in which both strategies and participants co-evolve.

As a consequence, equilibria defined under fixed-player assumptions may fail to capture system-level stability. Configurations that are locally stable may nevertheless lead to degradation of the interaction system (e.g., loss of participants, breakdown of cooperation), indicating the need for a framework that accounts for persistence and structural change.

B.3 Collapse Interpretation of Strategic Interaction

To extend the classical formulation, we introduce a dynamical interpretation of strategic interaction based on collapse and stability selection.

Let X denote the space of strategy configurations. For a finite game with players N and strategy sets S_i , we define:

$$X = \prod_{i \in N} S_i.$$

We further introduce a tension or instability measure:

$$T : X \rightarrow \mathbb{R}_{\geq 0},$$

where $T(x)$ quantifies the degree to which a configuration $x \in X$ fails to remain stable under repeated interaction. This may include, for example, sensitivity to deviation, breakdown of cooperation, or loss of participants.

We define a collapse operator:

$$\Phi : X \rightarrow X,$$

which acts to suppress configurations with high instability and retain those compatible with continued interaction.

Stability condition. A configuration $x \in X$ is said to be approximately stable under collapse if:

$$\Phi(x) \approx x,$$

indicating that repeated application of the collapse operator produces no significant change.

Interpretation. Under this formulation, strategic interaction can be viewed as an iterative process in which configurations evolve under local incentives and are subsequently filtered

by collapse. Stable configurations are therefore characterized not only by best-response conditions, but by their persistence under repeated interaction and structural constraints.

B.4 Persistence-Weighted Payoff

To incorporate system-level viability into the strategic framework, we extend the payoff function to include a persistence term.

Let $U_i^{(local)}(x)$ denote the classical payoff to player i under configuration x . We define the extended payoff:

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

where:

- $P : X \rightarrow \mathbb{R}_{\geq 0}$ is a persistence functional measuring the expected continuity of the interaction system under configuration x ,
- $\lambda \geq 0$ is a parameter controlling the relative influence of persistence.

Interpretation. Under this extension, agents are evaluated not only by their immediate gains, but also by the extent to which their behavior contributes to the stability of the system in which those gains are realized.

Configurations that maximize local payoff while reducing persistence are penalized through the persistence term, while configurations that sustain interaction are favored.

Relation to collapse. The persistence-weighted payoff provides a mechanism through which collapse operates: configurations with low persistence correspond to higher instability and are therefore more likely to be modified under the action of Φ .

Thus, the combination of local payoff and persistence defines a selection criterion for collapse-stable configurations.

Summary. The introduction of collapse dynamics and persistence-weighted payoff extends the classical framework by distinguishing between:

- local incentive optimization, and
- system-level viability.

This provides the basis for interpreting strategic stability as persistence under collapse, rather than solely as equilibrium under local payoff maximization.

B.5 Prisoner's Dilemma as a Minimal Collapse System

The Prisoner's Dilemma provides a canonical example of local incentive conflict and is therefore a natural test case for an emergent, collapse-based interpretation of strategic interaction.

B.5.1 Classical Formulation

Consider two agents, each selecting between cooperation C and defection D . The payoff structure is given by:

- T : temptation to defect
- R : reward for mutual cooperation
- P : punishment for mutual defection
- S : sucker's payoff

with ordering:

$$T > R > P > S$$

The resulting payoff matrix is:

	C	D
C	(R, R)	(S, T)
D	(T, S)	(P, P)

Under this structure, defection strictly dominates cooperation for both players, and the unique Nash equilibrium is (D, D) .

B.5.2 Descriptive Limitation

While the Nash equilibrium correctly identifies a stable outcome under one-shot rational choice, it does not account for the persistence of cooperative behavior observed in repeated or embedded interactions.

This discrepancy indicates that the equilibrium concept operates at a *descriptive layer*, capturing local incentive compatibility but not the generative dynamics that determine whether such configurations remain viable over time.

B.5.3 Configuration Space and Collapse Operator

Let X denote the space of strategy configurations. In the minimal case:

$$X = \{(C, C), (C, D), (D, C), (D, D)\}$$

We introduce a collapse operator:

$$\Phi : X \rightarrow X$$

which acts to suppress unstable configurations and retain those that are consistent with persistence under repeated interaction.

A configuration $x \in X$ is said to be *collapse-stable* if:

$$\Phi(x) = x$$

B.5.4 Persistence Functional

To capture system-level viability, we define a persistence functional:

$$P : X \rightarrow \mathbb{R}_{\geq 0}$$

where $P(x)$ measures the expected continuity of the interaction system under configuration x .

This includes, implicitly:

- maintenance of trust or predictability
- preservation of future interaction opportunities
- avoidance of destabilizing cascades (e.g., mutual defection)

B.5.5 Persistence-Weighted Payoff

We extend the standard payoff to include a persistence term:

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

where:

- $U_i^{(local)}$ is the classical payoff
- $P(x)$ is the persistence functional

- $\lambda \geq 0$ weights the contribution of persistence

B.5.6 Collapse Interpretation of Outcomes

Under this formulation:

- The configuration (D, D) , while locally stable in the classical sense, yields low persistence:

$$P(D, D) \ll P(C, C)$$

- Repeated selection of (D, D) degrades the conditions required for continued interaction, reducing $P(x)$ over time.
- As persistence falls below a stability threshold, the system undergoes collapse:

$$\Phi(x) \neq x$$

resulting in a reorganization of the interaction regime (e.g., termination of interaction, emergence of retaliation strategies, or enforcement mechanisms).

B.5.7 Emergence of Cooperation

In contrast, the configuration (C, C) supports high persistence:

$$P(C, C) \gg P(D, D)$$

Under persistence-weighted dynamics, cooperative configurations become *collapse-stable*, even if they are not strictly dominant under local payoff considerations.

This yields:

$$\Phi(C, C) = (C, C)$$

B.5.8 Interpretation

From this perspective:

- Classical equilibrium identifies locally stable configurations under fixed assumptions.
- Collapse dynamics identify configurations that remain viable under repeated interaction and system-level constraints.

Thus, cooperation is not an anomaly requiring additional assumptions; it is a *collapse-stable invariant* under persistence-constrained dynamics.

B.5.9 Summary

The Prisoner's Dilemma, reinterpreted through collapse dynamics, demonstrates that:

- Local payoff maximization does not guarantee system viability.
- Persistence introduces a higher-order constraint on strategic behavior.
- Collapse selects configurations that maintain the conditions for continued interaction.

Accordingly, cooperative behavior arises not as a deviation from rationality, but as a stable solution under a broader generative framework.

B.6 Emergence of Cooperation as Invariant Structure

The collapse-based formulation allows strategic configurations to be classified in terms of invariant structure, providing a natural extension of stability concepts in game theory.

B.6.1 Collapse-Stable Configurations

Let X denote the space of strategy configurations, and let

$$\Phi : X \rightarrow X$$

be the collapse operator introduced in Section B.5.

A configuration $x \in X$ is said to be *collapse-stable* if:

$$\Phi(x) = x.$$

The set of all such configurations defines the invariant sector:

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

B.6.2 Interpretation of Invariance

Configurations in \mathcal{I} are stable under repeated interaction and selection. They represent strategy profiles that maintain the conditions required for continued interaction, and are therefore robust under perturbation.

In contrast, configurations $x \notin \mathcal{I}$ are unstable: repeated application of collapse leads to a transition

$$x \mapsto \Phi(x),$$

indicating a reorganization of the interaction structure.

B.6.3 Cooperation as an Invariant

In the context of the Prisoner's Dilemma, the cooperative configuration (C, C) corresponds to a collapse-stable point:

$$\Phi(C, C) = (C, C),$$

due to its high persistence as described in Section B.5.

Conversely, the defection configuration (D, D) , while locally stable under classical equilibrium, is not invariant under persistence-constrained dynamics:

$$\Phi(D, D) \neq (D, D),$$

reflecting its tendency to degrade the interaction system over time.

Thus, cooperation emerges as an invariant structure under collapse dynamics, rather than as a deviation from rational behavior.

B.6.4 Invariant Sets and Strategy Networks

More generally, invariant structures need not be single configurations, but may form sets or networks of strategies that remain stable under collapse. These may include:

- clusters of cooperative strategies,
- reciprocity-based strategies (e.g., conditional cooperation),
- enforcement or reputation mechanisms.

Such structures can be interpreted as equivalence classes under collapse dynamics, where distinct configurations are mapped to a common invariant sector.

B.6.5 Relation to Evolutionary Stability

Classical evolutionary stability identifies strategies that resist invasion under replicator dynamics. In the present framework, this notion is extended by incorporating persistence:

- Evolutionary stability captures resistance to local perturbation.

- Collapse stability captures persistence under global constraints.

Invariant structures \mathcal{I} therefore represent strategy configurations that are stable under both local competition and system-level viability requirements.

B.6.6 Summary

The invariant sector \mathcal{I} provides a structural characterization of stable strategic behavior. Under collapse dynamics:

- Stable strategies correspond to fixed points of Φ ,
- Unstable strategies are mapped toward invariant configurations,
- Cooperative behavior emerges naturally as an invariant structure.

This reframes strategic stability as a property of persistence under collapse, rather than solely as equilibrium under local payoff optimization.

B.7 Relation to Evolutionary Game Theory

The collapse-based formulation introduced above can be related directly to evolutionary game theory, where strategy distributions evolve over time under selection dynamics.

B.7.1 Classical Replicator Dynamics

Let x_i denote the proportion of agents adopting strategy i , with payoff U_i . The standard replicator equation is given by:

$$\dot{x}_i = x_i (U_i - \bar{U}),$$

where $\bar{U} = \sum_j x_j U_j$ is the average population payoff.

This formulation describes the growth of strategies that outperform the population average, leading to equilibrium distributions corresponding to evolutionarily stable strategies.

B.7.2 Persistence-Weighted Extension

To incorporate system-level viability, we extend the payoff function to include a persistence term as introduced in Section B.5:

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

where $P(x)$ measures the persistence of the interaction system under configuration x , and $\lambda \geq 0$ controls its influence.

Substituting into the replicator equation yields:

$$\dot{x}_i = x_i \left(U_i^{(local)}(x) + \lambda P(x) - \bar{U} \right),$$

with

$$\bar{U} = \sum_j x_j \left(U_j^{(local)}(x) + \lambda P(x) \right).$$

B.7.3 Interpretation

Under this extension, strategy evolution is influenced not only by local payoff differences but also by the contribution of each configuration to system persistence.

Strategies that maximize short-term payoff while degrading system viability (e.g., persistent defection in the Prisoner's Dilemma) experience a reduction in effective fitness through the persistence term. Conversely, strategies that sustain interaction (e.g., cooperation) gain an advantage proportional to their contribution to $P(x)$.

Thus, the dynamics select for configurations that are both locally advantageous and globally sustainable.

B.7.4 Relation to Collapse Dynamics

The persistence-weighted replicator equation can be interpreted as a continuous approximation to the collapse process described in Section B.5.

In this view:

- Replicator dynamics generate variation and competition among strategies.
- The persistence term introduces a global constraint that suppresses unstable configurations.
- Fixed points of the extended dynamics correspond to collapse-stable configurations satisfying:

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

This establishes a connection between evolutionary stability and collapse stability, where the latter incorporates system-level constraints absent from purely local payoff models.

B.7.5 Summary

The persistence-weighted extension of replicator dynamics provides a bridge between classical evolutionary game theory and collapse-based interpretations of strategic interaction. It shows that:

- Evolutionary dynamics can incorporate system-level persistence as a selection pressure.
- Stable strategies are those that remain viable under both local competition and global constraints.
- Collapse dynamics may be viewed as a limiting or coarse-grained description of persistence-constrained evolution.

Accordingly, evolutionary stability can be reinterpreted as a manifestation of collapse stability in continuous time.

B.8 Finite Invariance and Bounded Responsibility

The persistence-based formulation introduced above implies that cooperative structures must remain bounded in order to remain stable. This introduces a natural constraint on strategic behavior, which can be expressed in terms of finite invariance.

B.8.1 Unbounded Cooperation and Instability

While cooperative configurations can maximize persistence, unbounded extension of cooperative behavior can introduce instability. For example:

- unconditional cooperation may be exploited by defectors,
- excessive allocation of resources to maintain cooperation may reduce system resilience,
- attempts to preserve all interactions may exceed the capacity of the system.

In such cases, persistence does not increase monotonically with cooperation, but instead exhibits diminishing or negative returns beyond a certain threshold.

B.8.2 Finite Invariance Constraint

To capture this, we introduce a bounded persistence condition:

$$P(x) \leq P_{\max},$$

where P_{\max} represents the maximum sustainable persistence of the system under given constraints.

A configuration x is said to satisfy *finite invariance* if it remains stable under collapse while respecting this bound.

B.8.3 Collapse Under Overextension

Configurations that attempt to exceed the persistence capacity of the system become unstable under collapse dynamics. Formally, if:

$$P(x) > P_{\max},$$

then repeated interaction leads to a reduction:

$$x \mapsto \Phi(x),$$

where $\Phi(x)$ satisfies:

$$P(\Phi(x)) \leq P_{\max}.$$

This reflects a reorganization of the system toward a configuration compatible with its finite capacity.

B.8.4 Interpretation in Strategic Systems

Within this framework, stable cooperative behavior is characterized not only by persistence, but by bounded persistence. Strategies must balance:

- local payoff optimization,
- maintenance of system continuity,
- adherence to capacity constraints.

This naturally gives rise to:

- selective cooperation,
- conditional reciprocity,
- enforcement of boundaries within interaction networks.

B.8.5 Relation to Responsibility

The bounded persistence condition provides a structural interpretation of responsibility in strategic systems. Agents contribute to system persistence through their actions, but this contribution is necessarily limited by finite capacity.

Thus, responsibility can be interpreted as:

- the contribution of an agent’s strategy to system persistence,
- subject to the constraint that such contribution remains sustainable.

This avoids both collapse due to insufficient cooperation and instability due to overextension.

B.8.6 Summary

Finite invariance introduces a constraint on persistence-driven dynamics, ensuring that stable configurations remain within the capacity of the system. Under collapse dynamics:

- cooperation must be bounded to remain stable,
- overextended configurations are reduced to admissible states,
- strategic stability requires balancing persistence and capacity.

Accordingly, stable strategic behavior emerges as persistence-constrained, rather than unboundedly cooperative.

B.9 Summary

In this appendix, we have presented a reformulation of strategic interaction in terms of collapse dynamics and persistence-constrained evolution.

Beginning with classical game theory, we identified equilibrium as a descriptive concept capturing locally stable configurations under fixed assumptions. We then introduced a collapse operator $\Phi : X \rightarrow X$, which acts on the space of strategy configurations X to suppress unstable configurations and select those compatible with continued interaction.

A persistence functional $P(x)$ was defined to measure system-level viability, and incorporated into payoff through a persistence-weighted extension. This modification was shown to align with evolutionary game theory by extending replicator dynamics to include system-level constraints, linking evolutionary stability to collapse stability.

Within this framework, stable strategies correspond to invariant configurations:

$$\mathcal{I} = \{x \in X \mid \Phi(x) = x\},$$

providing a structural interpretation of cooperation as a collapse-stable invariant. The Prisoner's Dilemma demonstrated that configurations maximizing local payoff may fail under persistence constraints, while cooperative configurations emerge as stable under repeated interaction.

Finally, the introduction of finite invariance established a bound on persistence, ensuring that stable configurations remain within system capacity. This yields a constrained notion of strategic stability in which viable behavior balances local payoff, system continuity, and finite resources.

Taken together, these results suggest that strategic behavior is more accurately described as an emergent process shaped by collapse and persistence constraints, rather than solely as equilibrium under local payoff optimization.