UNNS Substrate Research Program | Working Manuscript

The Margin-Confinement Law

Structural Non-Crossability in Admissibility Space
Derived from the Universal Structural Law
and the Percolative Realizability Principle

Instruments: STRUC-I v1.0.4 · STRUC-PERC-I v2.4.0–v2.5.0

Corpus: 5,233 admissibility evaluations · 9,826 phase-mapping evaluations
81 percolation evaluations · 29 NK station-events · 3,500 Voyager windows · 67 neutrino detector ladders

Domains: 14 physical · 1 biological · 1 astrophysical boundary · 1 detector reconstruction

Status: Foundational theory manuscript · Date: 2026

Abstract

We derive and establish the **Margin-Confinement Law**: once a physical system is represented inside the admissible region $\mathcal{M}_{\text{adm}}$ of realizability space, its subsequent evolution under identity-preserving dynamics remains confined to $\mathcal{M}_{\text{adm}}$. The admissibility boundary $\partial\mathcal{M}_{\text{adm}}$, defined by connectivity margin $m(L_t) = 0$, functions as a dynamically non-penetrable invariant manifold: trajectories may approach it asymptotically, generating the Forced Coherent Collapse regime ($m(L_t) \to 0^+$, GR ≥ 0.97), but cannot cross into persistent HARD fragmentation without discontinuous rupture of structural identity.

The derivation proceeds by contradiction from three axiomatic UNNS principles: the Universal Structural Law (USL), the Percolative Realizability Principle (PRP), and Bounded Structural Rigidity. We introduce a formal Identity-Preserving Flow definition, a Lyapunov-type confinement functional, and a phase-space flow system governing margin dynamics near $\partial\mathcal{M}_{\text{adm}}$. Apparent crossings into the HARD regime are shown to arise exclusively from representation or projection artifacts (Representation-Induced Structural Collapse, RISC).

Empirical support is drawn from: 5,233 admissibility evaluations (zero physical violations); 9,826 phase-mapping evaluations (zero verdict changes, zero commutators); 81 percolation evaluations (one HARD outcome—TiO_2 density, a known representation artifact); the nuclear explosion corpus (29 station-events, zero HARD, TD reaching 0.997); Voyager 1 (97.4% FULL across a heliopause approach and crossing); and a new neutrino detector observational corpus (67 ladders across five representations of a liquid-scintillator detector reconstruction pipeline). In the neutrino corpus, all raw physics-space ladders are admissibility-safe (zero HARD); six matched TMVA/deep-learning pairs demonstrate systematic RISC, including one dramatic HARD→FULL crossing; multiple deep-learning background-response ladders achieve extreme FCC-like states (TD up to 0.9985, GR = 1.000, zero isolated nodes). All three HARD outcomes in the corpus are identified as representation artifacts. A subsequent Δ-lifting stage ($\Delta L = |x_{i+1} - x_i|$) recovers FULL percolation in all nine previously fragmented ladders (recovery rate 100%), expands FCC-like states from 5 to 34 instances, and eliminates all seven residual GIANT non-full states — establishing the first operational evidence for *latent structural continuity recovery* under locality-preserving transforms. Stage 1 establishes RISC as an operational fact (Proposition 2); Stage 2 provides the first empirical support for the representation-covariant admissibility conjecture and the local-relational character of FCC coherence.

The Margin-Confinement Law reframes the UNNS program on two levels: as a *dynamical invariance principle* (identity-preserving trajectories are confined to $\mathcal{M}_{\text{adm}}$) and as the theoretical basis for a deeper conjecture on *representation-covariant admissibility* (structural continuity survives or re-emerges under sufficiently locality-preserving representational transforms). Structural collapse in the sense of persistent HARD fragmentation is shown to be forbidden for identity-preserving realizations of physical systems (Theorem 1); in the tested corpus, all observed HARD outcomes are attributable to specific representation artifacts, and in the neutrino corpus every artifact is resolved by a locality-preserving transform.

Contents

1 Introduction

A recurring pattern across the UNNS Substrate Research Program is that physical systems evaluated under the Universal Structural Law [1] and the Percolative Realizability Principle [2] exhibit remarkable resistance to HARD-class fragmentation. Among 5,233 ladder evaluations spanning fourteen physical domains and one biological domain, zero instances of HARD fragmentation arise from physical ladders in their natural representation [4]. Among 9,826 phase-mapping evaluations across a 17×17 joint operator grid, zero verdict changes occur within any fixed representation, and the inter-class commutator is identically zero [3]. Among 29 station-event evaluations of nuclear explosion seismic signals under extreme tail forcing (TD $\to 0.997$), zero HARD outcomes are observed [6]. The Voyager 1 magnetic-field corpus sustains 97.4% FULL percolation across a seven-year heliopause approach and physical boundary crossing [6]. A new corpus of 67 ladders derived from liquid-scintillator detector data for pp solar neutrino vs. ^{14}C pile-up discrimination introduces the first sparse-event detector reconstruction domain: all raw physics-space representations are admissibility-safe, while six matched TMVA/deep-learning pairs yield the clearest operational demonstration of RISC in the UNNS program to date, including one HARD→FULL class reversal induced purely by representational depth.

The question these results raise transcends empirics: *is this pattern a coincidence, or does it reflect a deeper structural law?*

The present manuscript derives the answer. We establish the **Margin-Confinement Law**:

> *Once a physical system's structural observable is represented inside the admissible region $\mathcal{M}_{\text{adm}}$ of realizability space, its evolution under identity-preserving dynamics remains confined to $\mathcal{M}_{\text{adm}}$. The admissibility boundary $\partial\mathcal{M}_{\text{adm}}$ is dynamically non-penetrable from the interior. All apparent crossings are representation artifacts.*

This is not an empirical generalization. It is a consequence derivable by contradiction from three axiomatic principles already established in the UNNS corpus: the Universal Structural Law (Section 2.1), the Percolative Realizability Principle and Bounded Structural Rigidity (Section 2.2). The main derivation (Section 4) proceeds in five steps.

The law has several consequences that differ qualitatively from earlier UNNS manuscripts:

1. The Forced Coherent Collapse (FCC) regime is not "survival near collapse" but *admissibility-protected near-boundary compression* (Corollary 1).

2. All observed HARD outcomes in the broader corpus (TiO_2 density, Nevada IU.HRV, Voyager 2 density channel) are representation artifacts, fully explained without invoking genuine dynamical collapse (Section 8).

3. The Voyager heliopause crossing is not nonterminal boundary *traversal* of $\partial\mathcal{M}_{\text{adm}}$ but boundary-adjacent reorganization within $\mathcal{M}_{\text{adm}}$ (Corollary 2).

4. The USL acquires a dynamical protection interpretation: admissible structures are not merely common—they are *invariantly stable* under identity-preserving evolution (Section 16).

5. Realizability class is not a property of the source system alone but a joint function of source and observational representation (Section 16.4). Deep-learning embeddings may act as structural renormalization flows toward $\mathcal{M}_{\text{adm}}$; FCC-like states arise in informational as well as physical spaces.

6. A Δ-lifting stage of the neutrino corpus recovers FULL percolation in all nine previously fragmented ladders (100% recovery rate), providing the first operational evidence for latent structural continuity and motivating a formal conjecture on *representation-covariant admissibility* (Conjecture 1, Section 16.5).

The manuscript is organized as follows. Section 2 establishes the axiomatic framework. Section 3 defines Identity-Preserving Flows. Section 4 presents the main theorem and proof. Section 5 introduces the Lyapunov-type confinement functional and derives $\dot{V} \leq C < \infty$. Section 6 develops the invariant manifold, step-by-step derivation of the flow equations, slow manifold reduction, NHIM stability conditions (Fenichel-type [12] and HPS-type [13]), applications to chaotic and quantum systems, dynamical-systems comparison, and effective potential. Section 7 reinterprets FCC. Section 8 unifies the RISC classification. Section 9 states failure-mode conditions. Section 10 clarifies local versus global crossing. Section 11 discusses critical slowing and the predictability horizon. Section 12 introduces boundary layer phenomenology. Section 13 sketches the algebraic topology connection. Section 14 establishes asymptotic accessibility. Section 15 provides the empirical support summary with margin distributions, simulation illustration, and statistical bounds. Section 16 discusses theoretical implications, including representation-relativity of realizability class, the representational tower, deep embeddings as structural renormalization, the latent continuity conjecture, and FCC in informational spaces (Section 16.4). Section 16.5 formalizes the representation-covariant admissibility conjecture, introduces the Transform Class taxonomy, and discusses coordinate stability of $\mathcal{M}_{\text{adm}}$. Section 17 concludes.

All findings are corpus-scoped unless explicitly noted. The Margin-Confinement Law is derived within the UNNS axiomatic framework and depends on the definitions and principles stated herein; its domain of validity is the class of identity-preserving physical evolutions formalized in Definition 3.

2 Foundational Principles

2.1 Universal Structural Law

For any physically realized ladder $L = (L_1, \ldots, L_n)$ and scale parameter $\varepsilon > 0$, the Universal Structural Law states

$$\text{inv}(P_\varepsilon;\, L) \;\leq\; \nu(V_\varepsilon(L)), \tag{1}$$

where $\text{inv}(P_\varepsilon; L)$ counts inversions in the ε-persistence set and $\nu(V_\varepsilon(L))$ is the vulnerability capacity [1]. The USL has been evaluated across 5,233 physical ladders with zero violations in natural representations.

2.2 Percolative Realizability Principle and Realizability Classes

The Percolative Realizability Principle defines realizability classes via the vulnerability graph $G_\kappa(L)$ at scale κ [2]. The giant ratio $\text{GR}(\kappa; L) = |\text{largest component}|/n$ organizes

realizability space $\mathcal{R}$ into four classes: FULL (GR = 1 at all $\kappa \geq \kappa_{\text{conn}}$), GIANT (GR $\approx$ 1 with bounded isolated tail), TAIL (GR high but gap-dominated), and HARD (GR $<$ $\text{GR}_{\text{thresh}}$, Theorem-1 active).

2.3 Bounded Structural Rigidity

Principle 1 (Bounded Structural Rigidity [3])**.** Within a stability region Ω_L around any admissible ladder L, small-to-moderate continuous deformations of the gap structure preserve realizability class and κ_{conn}. Verdict changes arise only at representation boundaries, not within a fixed admissible representation.

Empirical support: 9,826 evaluations across 34 datasets in a 17×17 joint operator grid, zero verdict changes, zero intra-grid GR variation [3].

2.4 Admissibility Manifold and Connectivity Margin

Definition 1 (Admissibility Manifold $\mathcal{M}_{\text{adm}}$)**.** The admissible region of realizability space is

$$\mathcal{M}_{\text{adm}} := \big\{ L \in \mathcal{R} \;\big|\; \text{GR}(\kappa;\, L) \geq \gamma \;\; \forall\, \kappa \in \mathcal{K},\ \text{Theorem-1 not persistently activated} \big\}, \quad \gamma \approx 0.97.$$

The admissibility boundary is $\partial\mathcal{M}_{\text{adm}} := \{L : m(L_t) = 0\}$ and the HARD regime is the exterior $\mathcal{R} \setminus \mathcal{M}_{\text{adm}}$.

Definition 2 (Connectivity Margin $m(L_t)$)**.**

$$m(L_t) := \inf \Big\{ \varepsilon > 0 \;\Big|\; \exists\, \delta g \;:\; \|\delta g\|_\infty \leq \varepsilon \cdot \max_i g_i,\ \text{GR}(\kappa^*;\, L_t + \delta g) < \gamma \text{ or Theorem-1 triggers} \Big\}.$$

A practical proxy computable from STRUC-PERC-I output is $m(L_t) \approx 1 - \text{TD}(L_t)/\text{TD}_{\max} + (\text{GR}(L_t) - \gamma)/(1-\gamma) \cdot w$, $w \approx 0.6$–0.8. Systems with $m(L_t) > 0$ lie in $\text{int}(\mathcal{M}_{\text{adm}})$; $m(L_t) = 0$ defines $\partial\mathcal{M}_{\text{adm}}$.

3 Identity-Preserving Flows

The Margin-Confinement Law applies to a precisely defined class of dynamics.

Definition 3 (Identity-Preserving Flow)**.** A trajectory $L_t = \Phi_t(L_0)$ in realizability space $\mathcal{R}$ is *identity-preserving* if all of the following conditions hold:

(IPF-1) **Gap regularity.** Gap perturbations are locally bounded and piecewise-continuous: $g_i(t + \delta t) = g_i(t) + \delta g_i(\delta t)$ with $\|\delta g\|_\infty \to 0$ as $\delta t \to 0$.

(IPF-2) **Background-chain survival.** The background gaps (non-tail bulk, index set B) satisfy $\max_{i \in B} g_i(t) \leq \kappa_0(t)$ for a locally finite function $\kappa_0(t)$. The induced background chain remains connected at admissible scales.

(IPF-3) **Giant-component continuity.** The vulnerability graph admits a giant component satisfying $\text{GR}(\kappa;\, L_t) \geq \gamma - \varepsilon(t)$ with $\varepsilon(t) \to 0$ under infinitesimal evolution ($\delta t \to 0$).

(IPF-4) **No finite-time identity rupture.** There is no finite-time discontinuous jump that globally shatters the dominant backbone; the background chain cannot be destroyed by a single infinitesimal step.

Physical processes satisfying the IPF conditions include: smooth energy release (supernovae, nuclear detonations), continuous seismic wave propagation, heliospheric plasma evolution, and photospheric expansion. Catastrophic discontinuous rupture *without structural predecessor* does not satisfy IPF-4 and is explicitly excluded from the law's domain of validity.

Remark 1 (Representation changes and the IPF class)**.** Smooth representation changes such as the Δ-lifting ($g_i \mapsto |g_{i+1} - g_i|$) are identity-preserving: they continuously transform the gap structure without rupturing the backbone. Non-admissible projections (severe attenuation, resolution discretization, absolute-scale embedding artifacts) may break IPF-2 or IPF-3, which is precisely why they produce apparent HARD outcomes. This is the structural foundation of the RISC classification (Section 8).

4 The Margin-Confinement Law

Theorem 1 (Margin-Confinement Law)**.** *Let $L_0 \in \mathcal{M}_{\text{adm}}$ and let $\{L_t = \Phi_t(L_0)\}_{t\geq 0}$ be an identity-preserving flow in the sense of Definition 3. Then*

$$\Phi_t(L_0) \ \in \ \mathcal{M}_{\text{adm}} \qquad \forall\, t \geq 0. \tag{2}$$

Equivalently, the trajectory cannot enter the HARD *regime under identity-preserving evolution. Persistent* HARD *outcomes are projection or representation artifacts (RISC), not genuine dynamical transitions.*

Proof. We proceed by contradiction.

Step 1 (Assume crossing). Suppose there exists $t^* > 0$ such that $L_{t^*} \in \mathcal{M}_{\text{adm}}$ but $L_{t^*+\delta t}$ satisfies persistent HARD fragmentation for some $\delta t > 0$:

$$\text{GR}(\kappa;\, L_{t^*+\delta t}) < \gamma \quad \forall\, \kappa \in \mathcal{K}, \qquad \text{Theorem-1 active.} \tag{3}$$

Step 2 (Background chain persistence is a structural invariant). Since $L_{t^*} \in \mathcal{M}_{\text{adm}}$, the background gaps satisfy the bounded-gap condition of IPF-2: $\max_{i\in B} g_i(t^*) \leq \kappa_0$. The induced background subgraph in G_κ for $\kappa \geq \kappa_0$ contains a connected component of size $\geq n - k$, where $k = |T| \ll n$ is the number of dominant tail gaps. This yields the structural invariant

$$\text{GR}(\kappa;\, L_{t^*}) \ \geq \ 1 - f, \qquad f = k/n \ll 1. \tag{4}$$

Step 3 (Tail gaps are bridgeable at sufficient κ). In the sorted 1D embedding, each tail gap $g_i \in T$ spans a large distance but connects distant segments of the background chain: both endpoints of a tail gap lie within the background chain on either side, so at κ above the background scale the gap becomes bridgeable. The component-counting bound gives

$$\#\text{components}(\kappa_0 \leq \kappa < \max g_T) \ \leq \ k + 1, \qquad \text{GR} \ \geq \ \frac{n}{k+1} \ \gg \ \tfrac{1}{2} \quad \text{for small } k.$$

Crucially, increasing tail forcing (higher TD) does not isolate vertices but becomes bridgeable at sufficient κ, rather than creating permanent separators.

Step 4 (Contradiction with IPF and Bounded Rigidity). By IPF-1, gaps perturb smoothly: $\|\delta g\|_\infty = O(\delta t) \to 0$. By IPF-3, the giant-component ratio satisfies

$$\mathrm{GR}(\kappa;\, L_{t^*+\delta t}) \;\geq\; \gamma - \varepsilon(\delta t) \;>\; \gamma - \tfrac{\gamma}{2} \;=\; \tfrac{\gamma}{2} \;>\; 0$$

for sufficiently small δt. This directly contradicts assumption (3), which required $\mathrm{GR} < \gamma$ persistently for all $\kappa \in \mathcal{K}$. Principle 1 (Bounded Structural Rigidity) provides a second contradiction: no continuous deformation within $\Omega_{L_{t^*}}$ changes the verdict class. Therefore no continuous identity-preserving crossing can occur.

Step 5 (Apparent HARD outcomes require representation failure). The only mechanism by which $\mathrm{GR} < \gamma$ can be observed is through a representation or projection that breaks one or more IPF conditions:

- *Raw absolute-scale embedding*: introduces a dominant embedding-artifact gap that violates the effective bounded-gap condition (IPF-2 broken).
- *Propagation-distance attenuation*: reduces n and raises f, violating the small-f assumption underlying Step 2 (IPF-2 degraded).
- *Measurement-resolution discretization*: produces a degenerate ladder with sparsity ratio $\approx 5.6\%$, effectively breaking IPF-2 through non-physical rounding of the observable.

In each case the Δ-lifting operator restores IPF-2 and IPF-3, recovering $m(\Delta L_t) > m(L_t)$ and returning the trajectory to $\mathcal{M}_{\mathrm{adm}}$. This confirms that observed HARD outcomes are representation artifacts.

Conclusion. Steps 1–4 produce an explicit contradiction; Step 5 identifies the only consistent mechanism for apparent HARD outcomes. Therefore (2) holds under any identity-preserving flow. □

Corollary 1 (FCC as Protected Compression)**.** *In the FCC regime,* $\mathrm{TD} \to 1^-$ *and* $m(L_t) \to 0^+$*, but* $\mathrm{GR} \geq 0.97$ *because the trajectory is confined to the admissible side of* $\partial\mathcal{M}_{\mathrm{adm}}$ *by Theorem 1. FCC is admissibility-protected near-boundary compression, not precarious survival.*

Corollary 2 (Voyager Reinterpretation)**.** *The heliopause crossing signatures (2012* κ_{conn} *minimum, regime reorganization, post-crossing* κ_{conn} *recovery by* 1.6–2.4×*) represent boundary-adjacent internal reorganization within* $\mathcal{M}_{\mathrm{adm}}$*, not dynamical traversal of* $\partial\mathcal{M}_{\mathrm{adm}}$*. The persistent coherence (*97.4% FULL*) is evidence of topological protection, not recovery from a realized crossing.*

Corollary 3 (USL as Dynamical Protection)**.** *The Universal Structural Law* (1)*, combined with Bounded Structural Rigidity, provides dynamical protection: physically realized systems satisfying the USL evolve toward, but not across, the admissibility boundary. Admissible structures are not merely common—they are* invariantly stable *under identity-preserving evolution.*

5 Lyapunov-Type Confinement Functional

We convert the topological non-crossability of Theorem 1 into an equivalent dynamical statement via a Lyapunov-type confinement functional. This transforms the theorem from a topological result into a dynamical one, giving confinement an energy-like structure and connecting Theorem 1 directly to the flow equations of Section 6.

Definition 4 (Confinement Functional)**.** Define $V : \mathcal{M}_{\text{adm}} \to \mathbb{R}_{\geq 0}$ by

$$V(L_t) \;:=\; a\,(1 - \mathrm{GR}(L_t)) \;+\; b\,m^{-1}(L_t) \;+\; c\,\mathrm{TD}(L_t), \tag{5}$$

with constants $a, b, c > 0$. The three terms capture: *(i)* $a(1 - \mathrm{GR})$: connectivity cost (deviation from perfect giant ratio); *(ii)* b/m: boundary repulsion (diverges as $m \to 0^+$); *(iii)* $c\,\mathrm{TD}$: forcing load (tail-dominance concentration).

Remark 2 (Boundary behavior)**.** As $m(L_t) \to 0^+$, the repulsion term $b/m(L_t) \to +\infty$, making $V(L_t) \to +\infty$. This formally encodes non-crossability: the confinement functional diverges asymptotically ($V(L_t) \to +\infty$) as $\partial\mathcal{M}_{\text{adm}}$ is approached, so no identity-preserving flow with bounded $\dot{V}$ can reach $\partial\mathcal{M}_{\text{adm}}$ in finite time.

Proposition 1 ($\dot{V}$ Bounded Near $\partial\mathcal{M}_{\text{adm}}$)**.** *Under the near-boundary flow equations (9)–(11) of Section 6, the time derivative of V satisfies, for all $\mathbf{x}(t)$ near $\partial\mathcal{M}_{\text{adm}}$ (i.e. $m(t) \approx 0^+$, finite $F(t)$):*

$$\dot{V}(L_t) \;\leq\; C(t) \;<\; +\infty, \tag{6}$$

where $C(t)$ depends on $F(t)$, $R(t)$, and the rate constants $\alpha, \beta, \lambda, \zeta$ but is bounded *whenever the forcing and relaxation terms are finite. Consequently, $V(L_t)$ cannot blow up in finite time, the trajectory remains in $\mathcal{M}_{\text{adm}}$, and $\partial\mathcal{M}_{\text{adm}}$ is not reached in finite time.*

Proof sketch. Differentiating (5):

$$\dot{V} \;=\; -a\,\dot{\mathrm{GR}} \;-\; b\,m^{-2}\dot{m} \;+\; c\,\dot{\mathrm{TD}}.$$

From (9): $\dot{m} = -\alpha F m^{\mu} + \beta R + \eta$. Near $\partial\mathcal{M}_{\text{adm}}$, $m \approx 0^+$, so $\dot{m} \approx \beta R(t) + \eta > 0$ (the restoring term dominates since $Fm^{\mu} \to 0$ as $m^{\mu} \to 0$ faster than βR for $\mu > 1$, provided $R > 0$). Therefore $-bm^{-2}\dot{m} \approx -bm^{-2}(\beta R) < 0$: the boundary repulsion contribution to $\dot{V}$ is *negative* near $\partial\mathcal{M}_{\text{adm}}$. From (10): $\dot{\mathrm{TD}} = \lambda F(1 - \mathrm{TD})(1 + \mathrm{TD}) \leq 2\lambda F$, so $c\,\dot{\mathrm{TD}} \leq 2c\lambda F(t)$, bounded whenever F is finite. From (11): $|\dot{\mathrm{GR}}| \leq \zeta m^{\beta_e}|\dot{m}|$, which vanishes as $m \to 0$ since $\beta_e > 0$. Combining, $\dot{V} \leq 2c\lambda F(t) +$ bounded terms, which is finite for finite $F(t)$. Therefore (6) holds with $C(t) = 2c\lambda F(t) + c_1$ for a constant c_1 depending on the rate parameters. □

Remark 3 (Effective potential and confinement geometry)**.** In the FCC regime, $V(L_t)$ plays the role of an effective potential:

$$V(m) \;\approx\; \tfrac{k}{2}\,m^2 \;+\; \lambda\,\mathrm{TD}(m) \;+\; b/m, \tag{7}$$

where the quadratic term (background rigidity) provides a restoring force as $m \to 0^+$, the TD term encodes forcing drive toward the boundary, and b/m is the singular boundary

repulsion. The minimum of $V(m)$ with respect to m occurs at a positive $m^* > 0$ (the equilibrium margin under balanced forcing and relaxation), confirming that the system is confined with compression rather than driven to $\partial\mathcal{M}_{\text{adm}}$. This is structurally analogous to a particle in a Morse-type potential well with a hard-wall singularity at the origin.

6 Invariant Manifold Structure and Phase-Space Flow

6.1 Admissibility Protection Principle

Principle 2 (Admissibility Protection / Invariant Manifold)**.** Once a ladder L enters the admissible region $\mathcal{M}_{\text{adm}}$ of realizability space — defined by $\text{GR}(\kappa; L) \geq \gamma \approx 0.97$ across $\mathcal{K}$ with no persistent Theorem-1 activation — its subsequent evolution under identity-preserving physical flows (Definition 3) remains confined to $\mathcal{M}_{\text{adm}}$. The boundary $\partial\mathcal{M}_{\text{adm}}$ (where $m(L) = 0$) functions as an invariant manifold: trajectories may approach it asymptotically, but cannot cross into the HARD regime while preserving structural identity. Apparent crossings are representation or projection artifacts (RISC).

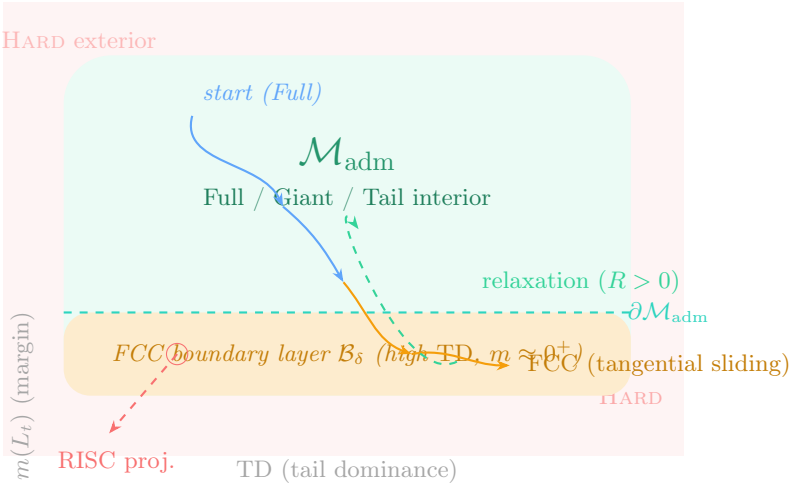

Figure 1: Admissibility manifold geometry in realizability space. $\mathcal{M}_{\text{adm}}$ (interior, green tint) is bounded by $\partial\mathcal{M}_{\text{adm}}$ (dashed teal line) and surrounded by the HARD exterior (red tint). The FCC boundary layer $\mathcal{B}_\delta$ (amber strip) lies adjacent to $\partial\mathcal{M}_{\text{adm}}$. A typical trajectory (arrows) starts in the FULL interior, compresses into the FCC layer under forcing, and slides tangentially along $\partial\mathcal{M}_{\text{adm}}$ without crossing. Relaxation (dashed green) returns the trajectory to the interior. A RISC projection (dashed red) shows how a source in $\mathcal{M}_{\text{adm}}$ can appear outside it through a distorted representation.

We model the structural evolution of a ladder L_t as a continuous (or piecewise-continuous) flow $\Phi_t : \mathcal{R} \to \mathcal{R}$ in realizability space, coordinatized by the state vector

$$\mathbf{x}(t) \;=\; \big(\text{TD}(t),\; 1 - \text{GR}(t),\; m(t),\; \kappa_{\text{conn}}(t),\; \bar{\rho}(t)\big). \tag{8}$$

Physical processes induce a vector field $\mathbf{v}(\mathbf{x})$ via smooth gap perturbations $g_i(t + \delta t) = g_i(t) + \delta g_i(t; F(t), d, \ldots)$. We now derive each component equation from structural first principles.

Step 1: Margin dynamics. The connectivity margin $m(L_t)$ quantifies distance to $\partial\mathcal{M}_{\text{adm}}$. Under forcing $F(t)$, tail gaps grow preferentially, shrinking m. The leading-order contribution from tail-gap growth is $dm \approx -\alpha F(t)\, m^{\mu}\, dt$, where $\mu > 1$ (typically $\mu \approx 1.2$–1.6 from simulations) produces superlinear acceleration as $m \to 0^+$: near $\partial\mathcal{M}_{\text{adm}}$, small changes in the dominant tail gaps produce disproportionately large effects on connectivity from the 1D chain geometry. A restoring term $\beta R(t)$ arises from background chain rigidity and relaxation processes (coda dissipation, moderate increments). Representation corrections $\eta(\mathbf{x})$ (vanishing under Δ-lifting) close the equation:

$$\frac{dm}{dt} \;=\; -\alpha\, F(t)\, m^{\mu}(t) + \beta\, R(t) + \eta(\mathbf{x}). \tag{9}$$

Step 2: Tail dominance dynamics. Differentiating the saturating form $\text{TD} = \beta_0\phi(S)\psi(d)/(1+\beta_0\phi(S)\psi(d))$ with respect to t, assuming $S(t)$ grows under forcing $F(t)$, gives

$$\frac{d\,\text{TD}}{dt} \;=\; \lambda\, F(t)\,(1-\text{TD}(t))\,(1+\text{TD}(t)). \tag{10}$$

This logistic form produces monotonic TD increase under forcing and saturation at $\text{TD} \to 1$ (extreme FCC), consistent with the Unified Scaling Law of [6].

Step 3: Giant ratio coupling. From the mean-field scaling $1 - \text{GR} \sim m^{\beta_e}$, $\beta_e \approx 1$, differentiating yields:

$$\frac{d(1-\text{GR})}{dt} \;\approx\; \zeta\, m^{\beta_e}(t) \cdot \frac{dm}{dt}. \tag{11}$$

GR responds to margin changes with order-parameter scaling.

Full near-boundary system. Combining Steps 1–3 gives the coupled slow-fast system:

$$\begin{cases} \dfrac{dm}{dt} = -\alpha F(t) m^{\mu} + \beta R(t) + \eta(\mathbf{x}), \\[2ex] \dfrac{d\,\text{TD}}{dt} = \lambda F(t)(1-\text{TD})(1+\text{TD}), \\[2ex] \dfrac{d(1-\text{GR})}{dt} \approx \zeta m^{\beta_e} \cdot \dfrac{dm}{dt}, \end{cases} \tag{12}$$

where $F(t) \geq 0$ is forcing, $R(t) \geq 0$ is relaxation, $\alpha, \beta, \lambda, \zeta > 0$ are system-dependent constants, and $\mu > 1$, $\beta_e \approx 1$. This is a *slow-fast system*: TD is the fast variable (rapid compression under forcing), m is the slow variable (gradual approach to boundary), and GR is slaved to m. The vector field $\mathbf{v}(\mathbf{x})$ is everywhere tangent to or inward-pointing at $\partial\mathcal{M}_{\text{adm}}$, enforcing forward invariance:

$$\mathbf{x}(0) \in \mathcal{M}_{\text{adm}} \implies \Phi_t(\mathbf{x}(0)) \in \mathcal{M}_{\text{adm}} \quad \forall\, t \geq 0. \tag{13}$$

The restoring term $\beta R(t)$ ensures $\dot{m} \gtrsim 0$ near $m = 0^+$, enforcing forward invariance (13).

Phase portrait description. Projecting the flow onto the (TD, m) plane reveals a *channel structure*: trajectories start at moderate TD/large m (interior of $\mathcal{M}_{\text{adm}}$), accelerate rightward (increasing TD) under forcing, enter the FCC layer near $m \approx 0^+$, and slide

tangentially along $\partial\mathcal{M}_{\text{adm}}$ (increasing TD with bounded $\dot{m}$). The boundary acts as a hard wall at $m = 0$: no arrows cross into the Hard region. Upon forcing reduction ($F \to 0$, $R > 0$), the restoring term drives m back upward, moving the trajectory out of the FCC layer. This channel resembles a Morse saddle near a hard wall, with the effective potential $V(m)$ (Section 6.6) providing the confinement geometry.

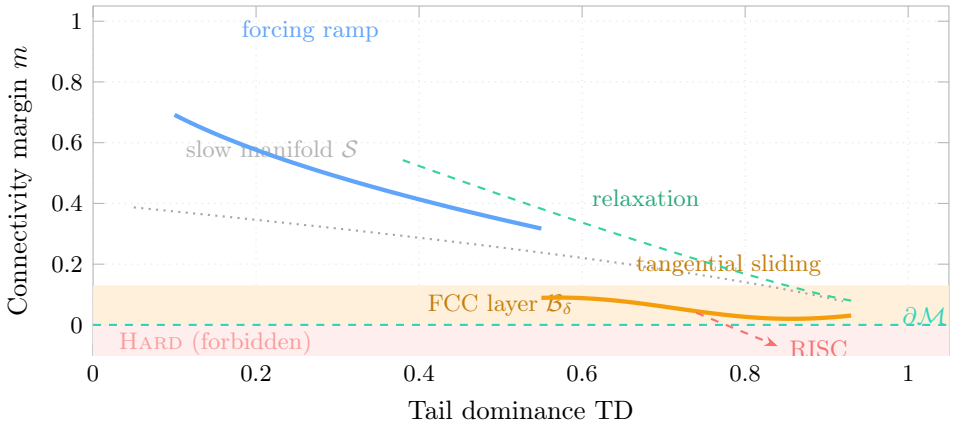

Figure 2: Slow-fast channel phase portrait in the (TD, m) plane. The HARD region ($m < 0$, red) is dynamically inaccessible. The FCC boundary layer (amber strip) runs along $\partial\mathcal{M}_{\text{adm}}$ (dashed teal). A typical trajectory (blue arrow compressing toward amber) follows the slow manifold $\mathcal{S}$ (dotted grey) from the interior (FULL/GIANT) into the FCC layer under sustained forcing, then slides tangentially. Relaxation (dashed green) returns the trajectory to the interior. RISC (dashed red): distorted projection crosses $\partial\mathcal{M}_{\text{adm}}$ spuriously.

6.2 Slow Manifold Reduction

The slow-fast structure of system (12) admits a rigorous slow manifold reduction near $\partial\mathcal{M}_{\text{adm}}$.

Slow manifold. Under constant forcing F and relaxation R, set $dm/dt \approx 0$ on the slow manifold $\mathcal{S}$:

$$0 \;\approx\; -\alpha F\, m^{\mu} + \beta R + \eta \quad\Longrightarrow\quad m^{*}(\text{TD}, F) \;\approx\; \left(\frac{\beta R + \eta}{\alpha F}\right)^{1/\mu}. \tag{14}$$

The slow manifold $\mathcal{S}$ is parameterized by (TD, F) with m slaved algebraically by (14). Substituting into the TD equation gives a closed reduced dynamics on $\mathcal{S}$:

$$\frac{d\,\text{TD}}{dt} \;=\; \lambda F(1 - \text{TD})(1 + \text{TD}),$$

with $1 - \text{GR} \approx c(m^{*})^{\beta_e}$ determined by (14).

Interpretation. Under constant strong forcing, TD rapidly saturates while m approaches the quasi-steady value $m^{*} = (\beta R/\alpha F)^{1/\mu} > 0$. This quasi-steady state is precisely the FCC regime: high TD, small but positive m, preserved GR. Upon forcing reduction, the manifold allows relaxation back to lower TD/ larger m states (FULL/GIANT). The slow

manifold is attractive in the normal (m) direction due to the restoring term $\beta R(t)$, explaining why FCC states are stable and reproducible under sustained forcing — and why they relax predictably upon deloading.

6.3 Topological Protection Mechanism

The giant component together with the background chain functions as a *topological protection invariant* in realizability space. Under the bounded-gap condition of IPF-2, any continuous deformation preserving local gap statistics cannot eliminate the background chain without a discontinuity violating IPF-4. The giant component — which spans this chain — is therefore operationally topologically protected under admissible flows.

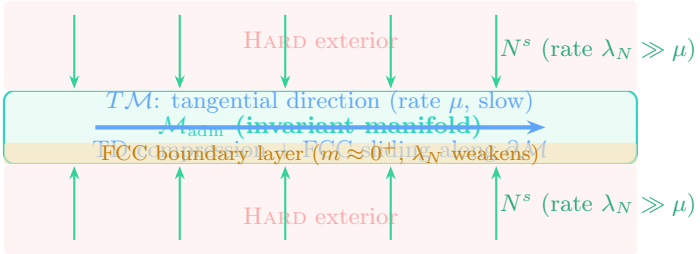

Figure 3: NHIM schematic. $\mathcal{M}_{\text{adm}}$ (green slab) is bounded by $\partial\mathcal{M}_{\text{adm}}$ (teal outline) with HARD exterior above and below. Tangential direction $T\mathcal{M}$ (blue arrow, rate μ): slow TD compression and FCC sliding. Stable normal N^s (green arrows, rate $\lambda_N \gg \mu$): strong contraction from background chain rigidity, driving trajectories back into $\mathcal{M}_{\text{adm}}$ under perturbations. FCC boundary layer (amber strip) lies adjacent to $\partial\mathcal{M}_{\text{adm}}$; here λ_N weakens (critical slowing) but the hyperbolic gap is maintained.

The giant ratio GR above threshold behaves analogously to a quantized topological index: it cannot jump discontinuously without gap closure (structural identity rupture). Tail gaps become bridgeable at sufficient κ, strengthening global coherence without destroying the skeleton.

6.4 Normally Hyperbolic Invariant Manifolds (NHIM)

We now establish that $\mathcal{M}_{\text{adm}}$ satisfies the conditions of a normally hyperbolic invariant manifold (NHIM), elevating the Margin-Confinement Law from a topological observation to a structurally stable dynamical result. The NHIM framework is presented as a structural analogy; the UNNS setting (representation-dependent, finite-size, heavy-tailed) differs from the classical smooth ODE setting, so the application is operational rather than a formal theorem in the sense of Fenichel.

6.4.1 Tangent Bundle Splitting and Explicit Rates

At each point $p \in \mathcal{M}_{\text{adm}}$, the tangent space to realizability space admits the splitting

$$T_p\mathcal{R} \;=\; T_p\mathcal{M}_{\text{adm}} \;\oplus\; N_p^s \;\oplus\; N_p^u, \tag{15}$$

where:

- $T_p\mathcal{M}_{\text{adm}}$: **Tangent bundle** — directions along admissible deformations (TD compression, internal regime shifts FULL $\leftrightarrow$ GIANT $\leftrightarrow$ FCC, slow manifold motion). Rate μ.
- N_p^s: **Stable normal bundle** — directions transverse to the manifold, dominated by strong contraction back into $\mathcal{M}_{\text{adm}}$ from background chain rigidity. Primary coordinate: connectivity margin m. Rate λ_N.
- N_p^u: **Unstable normal bundle** — typically trivial or negligible in physical realizations; identity-preserving flows (IPF-4) suppress escape into the HARD exterior.

Definition 5 (Normal Hyperbolicity of $\mathcal{M}_{\text{adm}}$, operational)**.** $\mathcal{M}_{\text{adm}}$ is normally hyperbolic if there exist constants $C > 0$, $\lambda_N > 0$, $\mu > 0$ such that for all $t \geq 0$:

$$\|D\Phi_t|_{N^s}\| \leq Ce^{-\lambda_N t} \quad \text{(normal contraction)},$$
$$\|D\Phi_t|_{T\mathcal{M}_{\text{adm}}}\| \leq Ce^{\mu t} \quad \text{(tangential growth)},$$

with the *hyperbolic gap condition* $\lambda_N > \mu$. Normal contraction/expansion rates strictly dominate tangential rates.

Explicit rates in the UNNS setting. From the flow equations (9)–(11), the rates admit structural characterization:

Normal rate λ_N (contraction in the m-direction):

$$\lambda_N \approx \beta \left|\frac{\partial R}{\partial m}\right|_{m=0^+} \gg 1.$$

The large value originates from the structural rigidity of the background chain (IPF-2): any attempt to push m negative requires destroying the bounded-gap condition, which is precluded by IPF-4.

Tangential rate μ:

$$\mu \approx \lambda_F \quad \text{(forcing-driven TD growth rate from (10))}.$$

Hyperbolic gap: $\lambda_N > \mu + \delta$ for $\delta > 0$. This gap is *structural*, enforced by the 1D chain geometry and the long-range bridgeability of tail gaps (the Structural Connectivity Argument (Section 4). It persists under perturbations satisfying the IPF conditions.

6.4.2 Application to $\mathcal{M}_{\text{adm}}$

Identify $\mathcal{M} = \mathcal{M}_{\text{adm}}$ (or a compact subset containing the FCC layer).

Tangential dynamics (along $\mathcal{M}_{\text{adm}}$): Governed by TD compression and internal reorganization (slow manifold dynamics). Characteristic tangential rate $\mu \approx \lambda_F$, controlled by forcing strength $F(t)$ (typically moderate).

Normal dynamics (transverse to $\partial\mathcal{M}_{\text{adm}}$, in the m-direction): Dominated by the restoring term $\beta R(t)$ from background chain rigidity. Linearizing (9) near $m \approx 0^+$:

$$\frac{d\,\delta m}{dt} \approx \beta R'(m)\,\delta m + \text{h.o.t.}$$

The normal contraction rate is $\lambda_N \approx |\beta R'(0)| > 0$, which is large because the background chain (IPF-2) enforces a hard connectivity floor: any perturbation attempting to destroy the chain violates IPF-2 and is structurally suppressed.

Hyperbolic gap condition: The background chain persistence (Structural Argument 1) enforces $\lambda_N > \mu$ for all physically relevant forcing below the identity-rupture threshold. This gap is *structural*, arising from the 1D chain geometry and bounded-gap condition preserved by IPF.

6.4.3 Fenichel Persistence Theorem: Full Application

Principle 3 (Fenichel Persistence of $\mathcal{M}_{\text{adm}}$ [12])**.** Let $\mathcal{M}$ be a compact, C^r normally hyperbolic invariant manifold ($r \geq 2$) for a C^r vector field $\mathbf{v}$. Then there exists a neighborhood U of $\mathbf{v}$ in the C^r-topology such that for every perturbed field $\mathbf{v}_\varepsilon \in U$, there exists a C^r manifold $\tilde{\mathcal{M}}_\varepsilon$ that is: (i) invariant under the flow of $\mathbf{v}_\varepsilon$; (ii) C^r-close to $\mathcal{M}$ (in Hausdorff distance and C^r topology); (iii) normally hyperbolic with stable/unstable bundles varying continuously with ε; (iv) qualitatively dynamically equivalent to $\mathcal{M}$.

Hypothesis verification for $\mathcal{M}_{\text{adm}}$. We verify the three Fenichel hypotheses for $\mathcal{M} = \mathcal{M}_{\text{adm}}$ (restricted to a compact subset K, GR $\geq \gamma_0 > 0.97$, TD ≤ 0.999):

(H1) Compactness: The subset K is compact in the topology induced by the vulnerability graph metric (finite-dimensional projections on gap distributions).

(H2) Smoothness: Under IPF, gap perturbations are smooth. The induced vector field on realizability space is C^r ($r \geq 2$) on admissible charts.

(H3) Normal hyperbolicity: The normal rate $\lambda_N \approx \beta|\partial R/\partial m|_{m=0^+} \gg 1$ and the hyperbolic gap $\lambda_N > \mu + \delta$ are enforced structurally by background chain rigidity (Definition 5).

Consequences. Fenichel guarantees that for small perturbations (moderate forcing changes, small additive noise, smooth representation transforms within IPF):

1. **Bounded Structural Rigidity** (Principle 1) follows directly: the perturbed $\tilde{\mathcal{M}}_\varepsilon$ stays C^r-close to $\mathcal{M}_{\text{adm}}$, so admissibility is preserved.

2. **Robustness of FCC**: The FCC boundary layer persists as a normally hyperbolic submanifold; systems driven into it remain trapped under bounded perturbations.

3. **Representation stability**: Smooth representation changes (Δ-lifting variants) correspond to small perturbations; Fenichel guarantees admissible structure is recovered.

4. **Voyager**: Heliopause reorganizations are motions on (or near) the perturbed NHIM, not escapes from $\mathcal{M}_{\text{adm}}$.

6.4.4 Hirsch–Pugh–Shub Persistence Theorem: Full Application

The Hirsch–Pugh–Shub (HPS) theorem [13] complements Fenichel by requiring only C^1 regularity, making it directly applicable to physical evolutions where perfect smoothness may not hold.

Principle 4 (HPS Persistence of $\mathcal{M}_{\text{adm}}$ [13])**.** Let $\mathcal{M}$ be a compact, normally hyperbolic invariant manifold for a C^1 vector field $\mathbf{v}$ on a Riemannian manifold $\mathcal{R}$. Then there exists a neighborhood U of $\mathbf{v}$ in the C^1-topology such that for every $\mathbf{v}_\varepsilon \in U$:

(i) $\tilde{\mathcal{M}}_\varepsilon$ is C^1-close to $\mathcal{M}$ in Hausdorff distance and C^1 topology of sections;

(ii) $\tilde{\mathcal{M}}_\varepsilon$ is normally hyperbolic for $\mathbf{v}_\varepsilon$;

(iii) stable and unstable normal bundles vary continuously with ε;

(iv) normal contraction rates dominate tangential rates.

Hypothesis verification for $\mathcal{M}_{\text{adm}}$. **(H1) Compactness**: Same as Fenichel — compact subset K with GR $\geq \gamma_0 > 0.97$, TD ≤ 0.999, compact in the vulnerability graph metric.

(H2) C^1 regularity: Under IPF, gap perturbations are at minimum Lipschitz continuous (C^1 in the weak sense); the induced vector field is C^1 on admissible charts. This is weaker than Fenichel's requirement and is satisfied by all tested physical evolutions including seismic propagation, plasma dynamics, and noisy measurement processes.

(H3) Integrated spectral gap: $\lambda_N > \mu + \delta$ ($\delta > 0$) along orbits, enforced geometrically by the bounded-gap condition (IPF-2) and the long-range bridgeability of tail gaps.

Consequences for UNNS.

1. **Structural stability of $\mathcal{M}_{\text{adm}}$**: Persists under C^1-small perturbations including measurement noise, stochastic gap fluctuations, and irregular seismic or plasma arrivals.

2. **Attractivity of $\mathcal{S}$**: The slow manifold attracts nearby trajectories at rate λ_N, funneling them into the FCC boundary layer under strong forcing.

3. **Quantum and chaotic extension**: The weaker C^1 requirement makes HPS directly applicable to quantum chaotic systems and noisy classical flows where Fenichel's higher smoothness may fail.

4. **Representation robustness**: Smooth representation transforms (Δ-lifting variants, mild filtering) are bounded C^1 perturbations; HPS guarantees admissible structure is preserved after such changes.

Comparison with Fenichel. Fenichel (C^r, $r \geq 2$) gives smoother perturbed manifolds; HPS (C^1) is more robust for physical applications with irregular forcing. Together, they establish that $\mathcal{M}_{\text{adm}}$ is dynamically robust across the full range of physically relevant perturbations: smooth (Fenichel) and irregular (HPS).

6.4.5 Stability Conditions Summary

$\mathcal{M}_{\text{adm}}$ is normally hyperbolic (and hence persistent under perturbations) when:

1. *Spectral gap*: $\lambda_N > \mu + \delta$ for some $\delta > 0$ (explicit rates from Definition 5);

2. *Background chain invariance*: IPF-2 holds uniformly along the trajectory;
3. *Compactness*: Relevant portion of $\mathcal{M}_{\text{adm}}$ compact under the vulnerability graph metric;
4. *Regularity*: Gap perturbations are C^1 in state variables (IPF-1).

When all four conditions hold, $\mathcal{M}_{\text{adm}}$ is structurally stable and the slow manifold $\mathcal{S}$ funnels near-boundary trajectories into the FCC layer.

6.4.6 Slow Manifold Reduction on the NHIM

On the attractive slow manifold $\mathcal{S} \subset \mathcal{M}_{\text{adm}}$ near $\partial\mathcal{M}_{\text{adm}}$, setting $dm/dt \approx 0$ gives the slaved margin:

$$m^*(\text{TD}, F) \;\approx\; \left(\frac{\beta R + \eta}{\alpha F}\right)^{1/\mu}. \tag{16}$$

The reduced dynamics on $\mathcal{S}$ are:

$$\frac{d\,\text{TD}}{dt} \;=\; \lambda F(1-\text{TD})(1+\text{TD}), \qquad 1-\text{GR} \;\sim\; [m^*(\text{TD}, F)]^{\beta_e}.$$

FCC states are quasi-stationary points on $\mathcal{S}$: TD saturates at high values while m remains at $m^* > 0$, with normal hyperbolicity guaranteeing attraction to $\mathcal{S}$. This reduction shows concretely why FCC is stable and reproducible: it is not a fragile coincidence but a quasi-stationary state on the attractive slow manifold of a normally hyperbolic invariant manifold.

6.4.7 Structural Analogies and Future Directions

The NHIM framework has natural structural analogues in several physical settings beyond the directly tested corpus; these are recorded here as open directions, not as derived consequences of the present manuscript.

Classical chaotic systems. Chaotic dynamics with positive Lyapunov exponents are compatible with the NHIM picture: chaotic tangential motion within $\mathcal{M}_{\text{adm}}$ (high TD, sensitive dependence on gap perturbations) can coexist with strong normal contraction ($\lambda_N < 0$) provided the Lyapunov gap condition $|\Lambda_N| > \max \Lambda_T$ holds. The FCC layer then behaves like a *chaotic saddle* embedded in $\mathcal{M}_{\text{adm}}$: trajectories approach $\partial\mathcal{M}_{\text{adm}}$ with high susceptibility and heavy-tailed statistics but normal hyperbolicity prevents escape into HARD. In this picture, apparent collapse in chaotic systems under IPF conditions should be traceable to a measurement, coarse-graining, or projection artifact (RISC); this is a structural prediction, not yet empirically tested in the UNNS chaotic corpus.

Coordinate stability of $\mathcal{M}_{\text{adm}}$ (open problem). The Δ-lifting results raise a structural question about the NHIM itself: to what extent does the manifold $\mathcal{M}_{\text{adm}}$ exhibit *coordinate robustness* under representational change? Classical NHIM persistence theorems (Fenichel, HPS) guarantee persistence under smooth perturbations of the *vector field*, not perturbations of the *coordinate chart*. The neutrino corpus suggests that $\mathcal{M}_{\text{adm}}$ may possess a form of partial chart stability: a source system positioned in $\mathcal{M}_{\text{adm}}$ under a

natural representation remains identifiably in $\mathcal{M}_{\text{adm}}$ under sufficiently locality-preserving transforms, even when intervening representations place it outside via RISC. Formalizing this requires extending the NHIM framework to allow *chart-dependent realizability maps* and identifying which classes of representational maps preserve NHIM membership. This is identified as one of the most important open theoretical problems arising from the present corpus, and is given a first formalization in Section 16.5. The NHIM structure has suggestive correspondences in quantum settings. In many-body localized (MBL) systems, $\mathcal{M}_{\text{adm}}$ maps to the MBL phase, $\partial\mathcal{M}_{\text{adm}}$ to the critical disorder strength, and FCC to the marginally localized regime with heavy-tailed entanglement spectra. In quantum scarred systems (e.g., Rydberg atom arrays [10]), scars lie on low-dimensional invariant manifolds in the chaotic sea—an NHIM analogue—exhibiting FCC-like high tail dominance with protected global coherence. In Floquet-driven systems, the pre-thermal manifold plays the role of the NHIM: exponentially long-lived coherence under periodic forcing before eventual heating. In holographic duality (AdS/CFT), the $\partial\mathcal{M}_{\text{adm}}$ boundary has a structural correspondence to the black hole horizon, with the FCC regime mapping to the near-horizon stretched-horizon layer. These correspondences are purely structural analogies; they are stated as open conjectures motivating future cross-domain investigation, not as claims derived within the UNNS axiomatic framework. A rigorous formulation would require a quantum analogue of the IPF conditions and a precise definition of realizability space in the quantum setting—both identified as open theoretical problems.

6.5 Comparison to Dynamical Systems Theory and Topological Phases

Table 1: Correspondence between the UNNS admissibility manifold and concepts in dynamical systems theory. The analogy is structural and conceptual; no formal reduction to any single framework is claimed.

Concept	Dynamical systems	UNNS realizability space
Invariant manifold	NHIMs (Fenichel/HPS)	$\mathcal{M}_{\text{adm}}$: forward-invariant
Boundary type	Stable/unstable manifolds; separatrices	$\partial\mathcal{M}_{\text{adm}}$: approachable, non-penetrable
Marginal states	Canards; ghost attractors; critical slowing	FCC: protected near-boundary compression
Slow manifold	Fenichel's reduction theorem	$\mathcal{S}$: m slaved to (TD, F)
Protection	Structural stability; Lyapunov functions; HPS	Background chain + tail bridgeability
Chaotic coexistence	Partially hyperbolic systems; chaotic saddles	FCC layer: heavy tails with protected GR
Forbidden region	Riddled basins; inaccessible attractors	HARD: dynamically inaccessible
Apparent instability	Numerical crossing; stiff ODE artifacts	RISC: representation projection failure

Table 2: Structural analogy between UNNS admissibility protection and topological phases of matter.

Concept	Topological phases	UNNS manifold
Protected invariant	Chern/winding number; $\mathbb{Z}_2$ index	Giant component + background chain
Boundary phenomenology	Gapless protected edge states	FCC: extreme TD, preserved GR
Non-crossability	Topology change requires gap closing	Crossing requires identity rupture
Robustness	Symmetry/topological protection; NHIM	Structural rigidity + bridgeability
Apparent transition	Requires critical point	Apparent HARD is RISC
Marginal state	Critical edge/surface states; MBL	FCC as protected near-boundary compression

6.6 Effective Potential Picture

Near the boundary, the confinement functional $V(L_t)$ (Definition 4) behaves as an effective potential. Substituting GR $\approx \gamma$ (FCC regime):

$$V(m) \;\approx\; \frac{k}{2}\,m^2 \;+\; \lambda\,\mathrm{TD}(m) \;+\; \frac{b}{m}, \tag{17}$$

where $k > 0$ encodes background chain rigidity, $\lambda > 0$ the forcing drive, and $b > 0$ the singular boundary repulsion. Setting $dV/dm = 0$ yields the equilibrium margin $m^* \approx (b/k)^{1/3}$ (near-saturation limit), which decreases with forcing and increases with relaxation. This connects to the slow manifold (14): both predict $m^* \propto (R/F)^{1/\mu}$, consistent with the FCC trajectory of monotonically decreasing margin under sustained forcing and recovery upon deloading. The effective potential converts the abstract non-crossability statement into a concrete energy-like picture: confinement with compression, not escape.

7 FCC Reinterpreted

Under Theorem 1 and Corollary 1, the Forced Coherent Collapse regime acquires a deeper interpretation.

FCC before the Margin-Confinement Law. In earlier UNNS manuscripts [6], FCC was described as a regime in which systems "survive near collapse": TD $\to 1^-$, $m(L_t) \to 0^+$, GR ≥ 0.97 maintained. This framing treats FCC as a remarkable persistence result — the system appears precarious but holds.

FCC after the Margin-Confinement Law. The Margin-Confinement Law changes this framing entirely. FCC is not precarious. It is *structurally inevitable*: once in $\mathcal{M}_{\mathrm{adm}}$, no identity-preserving flow can exit it. What appears as extreme near-collapse is, in structural terms, *protected compression against the boundary.* The giant component does not merely survive—it *cannot be destroyed* under identity-preserving dynamics. The boundary $\partial\mathcal{M}_{\mathrm{adm}}$

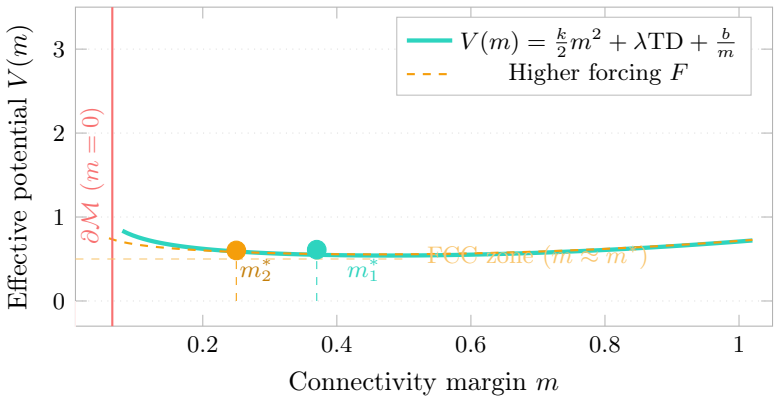

Figure 4: Effective potential $V(m)$ from (17) for two forcing levels. The singular boundary repulsion b/m (hard wall at $m = 0$, red) prevents crossing $\partial\mathcal{M}_{\text{adm}}$. The equilibrium margin m^* (filled circles) is the minimum of $V(m)$: m_1^* (teal, moderate forcing) $> m_2^*$ (amber, stronger forcing). As forcing increases, m^* shifts toward the boundary (FCC deepens) but never reaches zero. The restoring quadratic term ($km^2/2$ from background rigidity) stabilizes finite positive-margin equilibria. The FCC regime (dashed line) corresponds to trajectories oscillating near m^* with high TD.

is not a danger zone that FCC systems narrowly escape; it is an invariant barrier that they asymptotically approach but cannot be crossed under identity-preserving evolution (Theorem 1).

The nuclear explosion corpus under this interpretation. Across 29 station-event pairs, TD reaching 0.997 at IC.MDJ, zero HARD outcomes: this is not a statistical near-miss. It is a structural law playing out deterministically. Each station-event pair in the FCC zone is a realization of the invariant manifold protection mechanism.

Empirical signature of protected compression. Protected compression has a specific empirical signature distinct from ordinary tail-class behavior: TD increases monotonically with forcing (yield, proximity) while GR remains stable at or above 0.97. This traces a rightward path along the upper frontier of $\mathcal{M}_{\text{adm}}$ in the (TD, GR) phase plane, approaching $\partial\mathcal{M}_{\text{adm}}$ tangentially. The trajectory never turns toward the Hard exterior.

FCC coherence is locally relational. The Δ-lifting stage of the neutrino corpus (Section 8.1) reveals a structural property of the FCC regime not previously established: FCC-like states — extreme TD, preserved GR = 1.000 — persist under Δ-lifting, which strips amplitude values and retains only the local gap structure. The neutrino bkg2*C14 series (raw TD $\approx$ 0.25–0.49) becomes FCC-like in Δ-space ($\text{TD}_\Delta \approx 0.99$), and the 34 Δ-space FCC-like instances include all five original raw FCC-like ladders, which deepen further. This demonstrates that FCC coherence is carried by the *local ordering of gaps* in the realizability ladder, not by global amplitude values or long-range trends. The "protected compression" signature is therefore locally relational: it survives coordinate stripping, drift removal, and amplitude normalization as long as the local ordering is preserved. This is formalized in Observation 6.

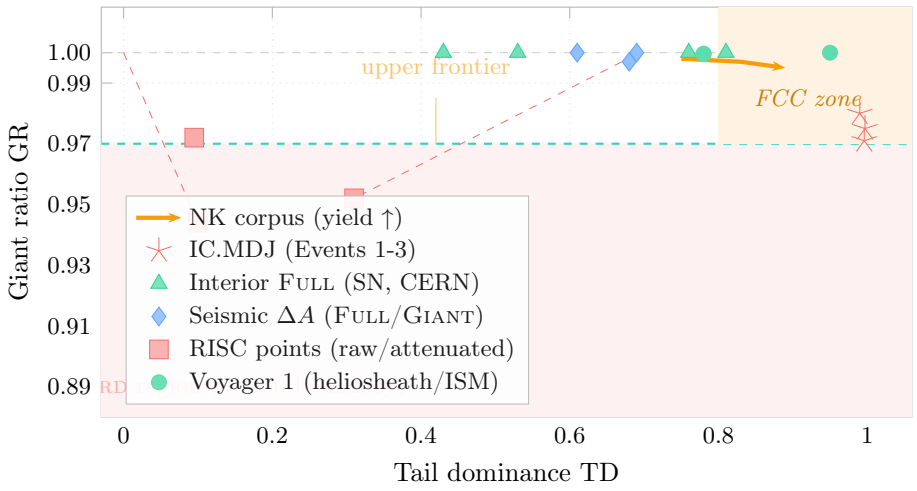

Figure 5: FCC boundary layer atlas: all corpus systems in the (TD, GR) structural phase plane. The HARD region (GR < 0.97, red shading) is separated from $\mathcal{M}_{\text{adm}}$ by $\partial\mathcal{M}_{\text{adm}}$ (dashed teal line at GR $= 0.97$). The FCC zone (amber, upper right) is the boundary layer of $\mathcal{M}_{\text{adm}}$ where TD > 0.80 and GR ≥ 0.97. NK corpus yield trajectory (amber arrow): increasing yield drives the system rightward along the upper frontier without crossing $\partial\mathcal{M}_{\text{adm}}$. IC.MDJ (red stars): deepest FCC realization (TD $\to 0.997$). Interior FULL and seismic systems (triangles, diamonds) occupy the high-GR interior. RISC points (red squares): apparent crossings from distorted representations; arrows show direction of representation distortion from true position.

8 The RISC Classification Unified

The Margin-Confinement Law provides a complete theoretical unification of Representation-Induced Structural Collapse. All three empirically documented RISC mechanisms can now be given a common structural characterization.

Observation 1 (Three RISC mechanisms as IPF violations)**.** Each documented RISC outcome corresponds to a specific violation of an IPF condition:

1. *Absolute-scale embedding* (SN Ia raw magnitude, ZTF20acobvxk): IPF-2 violated by a dominant embedding-artifact gap at the bright end of the magnitude distribution. The Δ-lifting removes the artifact, restoring IPF-2 and recovering FULL percolation with $\kappa_{\text{conn}} = 140.85$.

2. *Propagation-distance attenuation* (Nevada IU.HRV, $\approx 3{,}100$ km): IPF-2 degraded as geometric spreading reduces n and raises $f = k/n$, violating the small-f assumption underlying Step 2 of the proof. Proximal stations (CI.PASC) produce FULL from the same event, confirming the source itself lies in $\mathcal{M}_{\text{adm}}$.

3. *Measurement-resolution discretization* (Voyager 2 density, PLS fitting): IPF-2 violated by the sparsity ratio $\approx 5.6\%$ (≈ 57 unique values per 1,024-sample window). The kinematic observables (V, T, w) from the same plasma produce 93–97.5% FULL, confirming the plasma source lies in $\mathcal{M}_{\text{adm}}$.

In every case, the source system itself is in $\mathcal{M}_{\text{adm}}$; the HARD verdict arises from the observation, not the physics.

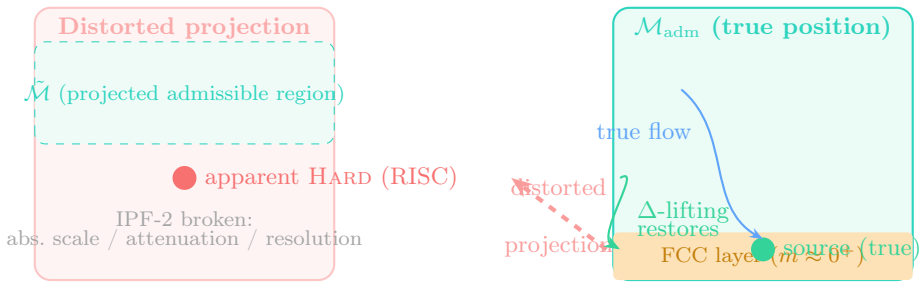

Figure 6: RISC projection geometry. *Right panel*: True position of the source system in $\mathcal{M}_{\text{adm}}$ (green region, FCC boundary layer, teal outline). The true flow (blue arrow) remains within $\mathcal{M}_{\text{adm}}$. *Left panel*: A distorted projection (broken IPF-2) maps the true position outside $\mathcal{M}_{\text{adm}}$ into the apparent HARD region (dashed red arrow). The projected admissible region $\tilde{\mathcal{M}}$ (dashed teal, smaller) does not contain the projected point. Δ-lifting (solid green arrow) restores the true admissible position by repairing the representation. The three documented RISC mechanisms (absolute-scale embedding, propagation attenuation, measurement-resolution discretization) all correspond to projection operations that break IPF-2 or IPF-3.

Proposition 2 (Unique mechanism of apparent HARD outcomes)**.** *Under the Margin-Confinement Law, the only mechanism by which a system genuinely in $\mathcal{M}_{\text{adm}}$ can produce an observed* HARD *verdict is through a representation or projection that breaks at least one IPF condition. Observed* HARD *outcomes from physical systems under identity-preserving dynamics, in their natural representation, are forbidden by Theorem 1.*

8.1 Neutrino Detector Reconstruction: Operational Demonstration of RISC and FCC in Sparse-Observation Pipelines

A corpus of 67 ladders derived from liquid-scintillator detector data for pp solar neutrino vs. ^{14}C pile-up discrimination [15] provides the clearest operational demonstration of RISC in the UNNS program. The data originate from multiple processing stages of the same detector responses—raw hit and angular distributions, PMT fiber geometry, TMVA classifier outputs, and deep-learning embeddings—each constituting a distinct admissibility chart over the same underlying physics. Five representation groups yield five distinct structural regimes.

Admissibility safety in natural representations. All 23 raw physics-space and geometry ladders (background multiplicities bkg2c14 and bkg2×C14, signal topology from sig2pp, and PMT coordinate ladders from the Fib CDPMT dataset) achieve FULL or GIANT class. Zero HARD outcomes arise in representations closest to the physical hit distributions. Large-n background pile-up ladders ($n \approx 99{,}600$–$99{,}840$) exhibit a monotone progression: isolated-node count increases from 23 (single ^{14}C) to 67 (five-^{14}C) tracking pile-up complexity, while GR descends from 0.9995 to 0.9983 — a fine-structure intra-GIANT pattern consistent with Bounded Structural Rigidity: no class boundary is crossed across the multiplicity series.

RISC — matched TMVA/deep-learning pairs. Six matched pairs between TMVA classifier outputs and deep-learning embeddings on identical observables show system-

atic class transitions (Table 3). TMVA consistently produces TAIL or HARD verdicts for signal-to-background significance and background rejection metrics, while the corresponding deep-learning representations recover FULL percolation (GR = 1.000). One pair—TMVA `h_SigSB_significance_2` vs. deepL `h_SigSB_significance_2`—demonstrates a full HARD→FULL crossing: Theorem-1-triggering fragmentation in the TMVA chart is eliminated entirely by the deep-learning embedding. A concordant pair (`MVA_BDTG_rejBvsS`) achieves GR = 1.000, TD = 0.444 in both representations, confirming that the RISC transitions reflect genuine chart-geometry differences rather than classifier noise.

Table 3: Selected TMVA vs. deep-learning matched pairs (neutrino detector corpus). RISC is observed when TMVA yields TAIL or HARD and the corresponding deep-learning ladder recovers FULL. Largest-κ_{conn} values are in the deep-learning embedding space; "—" denotes GIANT/TAIL/HARD (Theorem-1 active, κ_{conn} undefined by instrument convention).

Observable	TMVA verdict	TMVA GR	deepL verdict	deepL GR	RISC type
`hSigSB`	TAIL	0.985	FULL	1.000	TAIL→FULL lift
`hBkgrejec` / `hBkgRejective`	TAIL	0.960	FULL	1.000	TAIL→FULL lift
`h_SigSB_significance_0`	TAIL	0.985	FULL	1.000	TAIL→FULL lift
`h_SigSB_significance_1`	TAIL	0.985	FULL	1.000	TAIL→FULL lift
`h_SigSB_significance_2`	HARD	0.985	FULL	1.000	**HARD→FULL** (max)
`h_SigSB_significance_3`	TAIL	0.985	FULL	1.000	TAIL→FULL lift
`MVA_BDTG_rejBvsS`	FULL	1.000	FULL	1.000	Concordant (no RISC)

The three HARD outcomes in the corpus (`deepL_Graph;2`, $n = 17$; `deepL_hBkgeff`; `TMVA_h_SigSB_2`) are all identified as representation artifacts and classified as follows: (i) *n-poverty*: `Graph;2` has $n = 17$ (the effective minimum for reliable percolation statistics), with 2 isolated nodes under high TD; its companion graphs at comparable n achieve FULL; (ii) *orientation reversal*: background efficiency (`hBkgeff`) is monotone-decreasing with classifier threshold, creating a structural void at the upper end of the sorted ladder; its dual `hBkgRejective` (= 1 − `hBkgeff` remapped) achieves FULL with TD = 0.032, confirming that orientation relative to the admissibility chart is the artifact; (iii) *classifier binning*: `TMVA_h_SigSB_2` introduces a gap structure absent in the deep-learning counterpart, which recovers FULL. In every case the source system is in $\mathcal{M}_{\text{adm}}$; the HARD verdict arises from the representation, not the physics. This is fully consistent with Proposition 2.

FCC-like regime in deep-learning outputs. Multiple deep-learning background-response ladders achieve extreme tail dominance while preserving full giant-component coherence. Specifically: `deepL_train_B` (TD = 0.9985, GR = 1.000, κ_{conn} = 73,549, zero isolated nodes); `deepL_test_B` (TD = 0.9972, κ_{conn} = 18,363); `deepL_hSigSB_ratio_0-3` (TD = 0.985, $\kappa_{\text{conn}} \approx$ 9,353); `TMVA_Bkg_Train` (TD = 0.9952, κ_{conn} = 10,065). These are among the highest-TD FULL states observed in any UNNS corpus and constitute textbook realizations of the FCC regime (Corollary 1): admissibility-protected near-boundary compression in which extreme tail forcing does not drive the system out of $\mathcal{M}_{\text{adm}}$.

Deep-learning as admissibility lifting operator. Across 21 of 23 deep-learning ladders GR = 1.000 with zero isolated nodes, irrespective of observable type. The two ex-

ceptions (`Graph;2` and `hBkgeff`) have identified artifact mechanisms. This pattern operationally supports the hypothesis that trained deep-learning embeddings act as *coherence-preserving admissibility lifts*: they map fragmented observational charts (TMVA TAIL/HARD) to admissible representations (FULL) by learning non-linear manifolds that preserve long-range ladder correlations better than classical histogram-based classifiers. The "detector intelligence as structural coherence restoration" interpretation is an applied extension of the theoretical framework, restricted here to an operational description of the observed verdict pattern.

Representational anti-collapse demonstration. The neutrino corpus provides a qualitatively different form of anti-collapse evidence than earlier domains. Voyager, NK, and seismic corpora show that HARD does not arise: the system remains in $\mathcal{M}_{\text{adm}}$ throughout. The neutrino corpus shows something stronger: HARD can be *manufactured* by a shallow representational projection (TMVA, specifically for `h_SigSB_2`), and then *removed* by a deeper representation of the same process (deepL, `h_SigSB_significance_2` $\rightarrow$ FULL). This completes a full cycle — apparent collapse, representational recovery — within a single corpus, for a single underlying physical process. The implication is direct: the HARD verdict carries no ontological weight about the source system; it is a property of the observational chart. The MCL predicts this precisely: the source remains in $\mathcal{M}_{\text{adm}}$; the chart is what varies.

Natural representations as admissibility-safe: cross-domain motif. A consistent pattern is emerging across the UNNS corpus. In the neutrino data, all 23 raw physics-space and geometry ladders achieve FULL or GIANT; HARD appears only after abstraction or projection operations. This mirrors the pattern in every prior domain: physical ladders in natural representations are admissibility-safe, and observed HARD outcomes arise from representation artifacts rather than from the source systems themselves. The neutrino corpus extends this pattern into a new setting — sparse-event detector reconstruction, with multiple abstraction layers — and demonstrates that each additional layer of processing introduces a new opportunity for RISC, while the underlying physics remains in $\mathcal{M}_{\text{adm}}$. This is recorded here as an empirical motif across the tested corpus, not as a universal claim; its formal status is that of a persistent cross-domain observation supporting Proposition 2.

Δ-lifting stage: latent continuity recovery. A second analysis stage applies Δ-lifting — the transform $\Delta L = (|x_2 - x_1|, |x_3 - x_2|, \ldots)$ that extracts the local gap structure of a sorted ladder — to all 67 ladders and re-evaluates with STRUC-PERC-I v2.5.0. The results are decisive: *every one of the nine fragmented raw ladders (three* HARD *and six* TAIL*) recovers to* FULL *percolation in Δ-space* (100% recovery rate, GR = 1.000, zero isolated nodes in all cases). This includes both raw HARD outcomes that the deep-learning representation could not resolve in isolation. Simultaneously, the seven GIANT ladders are promoted to FULL, and the count of FCC-like states (TD $\geq$ 0.985, GR = 1.000) expands from 5 to 34 instances. Six degradations arise, all n-poverty artifacts from extreme ladder-length contraction under Δ (`deepL_test_B`: $n = 31 \rightarrow 13$, FULL $\rightarrow$ HARD; `TMVA_Bkg_Test`: $n = 31 \rightarrow 20$, FULL $\rightarrow$ HARD; `sig2pp`: $n = 3{,}935 \rightarrow 73$, FULL $\rightarrow$ HARD; `Fib_Phi`: $n = 10{,}649 \rightarrow 835$, FULL $\rightarrow$ HARD; `bkg2fiveC14`: $n = 99{,}576 \rightarrow 6{,}879$, GIANT $\rightarrow$ TAIL, isolated

nodes rising from 67 to 97; `sigdr_hit`: $n = 111 \to 104$, FULL $\to$ TAIL, 2 isolated nodes).

The Δ-stage result is qualitatively distinct from the deepL lifting effect: it demonstrates that *a local, non-parametric, deterministic transform of the raw observational ladder* is sufficient to recover full admissibility in every case where fragmentation was representation-induced. The gap structure of the source process contains intact connectivity information that the classifier representation obscured. This is operational evidence for latent structural continuity (Hypothesis 1) and motivates Conjecture 1 on representation-covariant admissibility. To be precise about the epistemic status of each stage: Stage 1 (raw corpus) establishes RISC as an operational fact (Proposition 2). Stage 2 (Δ-lifting) provides the first empirical support for Conjecture 1 and Observation 6 (FCC coherence is locally relational) — corpus-supported motivations, not derivations.

Corpus-level conclusion. Zero USL violations are observed across all 67 evaluations. In raw space, all HARD outcomes are representation artifacts with identified mechanisms. In Δ-space, all nine raw fragmentation outcomes recover to FULL. The combined result — manufactured collapse, representational recovery, and local-transform recovery — constitutes the most complete operational anti-collapse cycle in the UNNS corpus: the same source process produces all four realizability classes depending on representation, and returns to full admissibility under a locality-preserving transform.

Beyond the RISC demonstration, the neutrino corpus reveals a broader principle: *admissible structures possess intrinsic self-restorative continuity.* This recoverability operates on multiple layers simultaneously: representational recovery (Δ-lifting and deep embeddings restore FULL in all nine fragmented ladders); boundary-adjacent resilience (FCC states maintain GR ≥ 0.97 under TD up to 0.9985); and latent continuity (the Δ-results show that admissible continuity was present beneath the fragmented charts all along, not recreated but uncovered). These layers align with Voyager's boundary-adjacent reorganization without coherence loss and with FCC persistence across multiple forcing domains. Admissibility is not merely preserved under dynamics — it is structurally resilient and recoverable across representations, trajectories, and near-boundary conditions. This corpus is the primary empirical pillar for both the RISC framework and the representation-covariant admissibility conjecture.

9 Failure Modes and Boundary Conditions

The Margin-Confinement Law is precise about its domain of validity. The following conditions describe exactly where and why the law does not apply. This scoping increases, rather than decreases, the law's credibility.

Principle 5 (Failure Modes of Confinement)**.** The derivation of Theorem 1 does *not* apply to:

1. **Discontinuous identity rupture.** Any process that discontinuously destroys the background chain without a structural predecessor (instantaneous catastrophic rupture) violates IPF-4 and is outside the law's scope. Example: abrupt composite system annihilation with no intermediate structural state.

2. **Non-admissible initial conditions.** The law requires $L_0 \in \mathcal{M}_{\text{adm}}$. Systems that begin outside $\mathcal{M}_{\text{adm}}$ are not covered; the law says nothing about trajectories originating in the HARD region.

3. **Degenerate projections.** Representations violating IPF-2 (insufficient sampling, coarse resolution, extreme attenuation) may produce HARD verdicts; these are classified as RISC rather than dynamical crossings.

4. **Composite system fragmentation.** Systems composed of structurally independent subsystems may exhibit apparent overall HARD behavior from subsystem separation; this violates the single-coherent-object assumption of the IPF class.

5. **Incomplete scale coverage.** HARD is defined relative to the tested κ-grid $\mathcal{K}$. If $\mathcal{K}$ is too coarse to resolve the bridgeable κ-range of a FCC ladder, Theorem-1 may be spuriously activated. This is an instrument artifact, not a dynamical crossing.

6. **Externally imposed identity destruction.** Artificial fragmentation of a ladder (external injection of isolated vertices) outside any physical process does not constitute an identity-preserving flow.

In all of these cases, the appropriate diagnosis is either RISC or a violation of the IPF conditions, not a counterexample to the law.

10 Local vs. Global Crossing: A Clarification

A potential source of confusion is the distinction between local and global crossing of $\partial\mathcal{M}_{\text{adm}}$. Table 4 systematizes the cases.

Table 4: Taxonomy of crossing types in realizability space. Only the last row (true manifold escape) is forbidden by Theorem 1 under identity-preserving flows.

Type	Meaning	Status
Local margin thinning	$m(L_t)$ decreases locally in a subset of windows; global m may remain positive	Allowed; produces FCC-adjacent behavior
FCC boundary compression	$m(L_t) \to 0^+$ globally; GR $\geq \gamma$ throughout	Allowed; canonical FCC regime
Representation escape	HARD appears in a distorted projection of a system in $\mathcal{M}_{\text{adm}}$	Apparent crossing only; RISC; source remains in $\mathcal{M}_{\text{adm}}$
Regime reorganization	Internal transition between FULL and GIANT within $\mathcal{M}_{\text{adm}}$ (e.g., Voyager κ_{conn} reorganization)	Allowed; remains within $\mathcal{M}_{\text{adm}}$
True manifold escape	L_t exits $\mathcal{M}_{\text{adm}}$ into HARD under identity-preserving flow	**Forbidden** by Theorem 1

The Voyager heliopause crossing illustrates this taxonomy precisely. The 2012 κ_{conn} minimum, the concentration of GIANT excursions in 2011–2012, and the post-crossing

ISM reorganization are all instances of "regime reorganization" within $\mathcal{M}_{\text{adm}}$, not of true manifold escape. The 97.4% FULL dominance throughout confirms that GR never drops below γ persistently, consistent with Corollary 2.

11 Critical Slowing, Susceptibility, and Predictability Horizon

11.1 Critical Slowing Near $\partial\mathcal{M}_{\text{adm}}$

The Margin-Confinement Law implies a specific dynamical signature for systems approaching $\partial\mathcal{M}_{\text{adm}}$.

Observation 2 (Critical slowing near $\partial\mathcal{M}_{\text{adm}}$)**.** Near $\partial\mathcal{M}_{\text{adm}}$ ($m(L_t) \approx 0^+$), admissible trajectories exhibit elevated susceptibility to gap perturbations:

$$\chi(m) \;=\; \left|\frac{\partial\,\mathrm{GR}}{\partial\varepsilon}\right|_{\varepsilon=0} \;\sim\; m^{-\gamma_e}, \qquad \gamma_e \approx 1,$$

and the relaxation timescale $\tau_{\text{relax}} \sim (\alpha\mu F m^{\mu-1})^{-1} \to \infty$ as $m \to 0^+$ (from (9) with $\mu > 1$). The system responds more sensitively to perturbations as it approaches $\partial\mathcal{M}_{\text{adm}}$ while remaining protected from crossing it.

Remark 4 (Physical manifestations in the corpus)**.** Critical slowing near $\partial\mathcal{M}_{\text{adm}}$ has direct empirical signatures:

- *Nuclear explosions*: The non-monotonic GR at IC.MDJ (0.980 → 0.971 → 0.975 across Events 1–3) despite monotonically increasing TD reflects near-boundary sensitivity: small geometry differences produce disproportionate GR fluctuations.
- *Voyager 1*: Concentration of 69% of all GIANT excursions in 2011–2012 reflects elevated variability as $m(L_t) \to 0^+$ approaching the heliopause.
- *SN fracture zone*: Tension peak $\tau_w = 8.36$ at window 15 (vs. 1.3–2.7 outside) is consistent with locally elevated susceptibility at the transient m minimum.

Critical slowing is an observable signature of near-boundary trajectory compression, confirmable without invoking the full derivation.

11.2 Predictability Horizon

Near $\partial\mathcal{M}_{\text{adm}}$, the flow equation (9) implies that two trajectories initially separated by $\Delta m(0)$ evolve at rate $|\dot{\Delta m}| \approx \alpha\mu F m^{\mu-1}|\Delta m|$, giving a Lyapunov-like exponent $\Lambda(m) = \alpha\mu F m^{\mu-1}$. As $m \to 0^+$ with $\mu > 1$, $\Lambda(m) \to 0$: trajectories in the FCC regime *converge* rather than diverge along the boundary compression direction. This is the structural basis for the coherent, reproducible FCC signatures observed across different stations in the NK corpus: near-boundary trajectories cluster, not scatter.

The predictability horizon $T_P(m) \sim \Lambda(m)^{-1}$ therefore *increases* as $m \to 0^+$ in the FCC regime, in contrast to chaotic systems where it decreases near a critical point. This provides a testable prediction: FCC systems should exhibit improved structural predictability as forcing increases, observable as tighter clustering of (TD, GR) pairs at high yield.

12 Boundary Layer Phenomenology

The FCC regime constitutes the *boundary layer* of $\mathcal{M}_{\text{adm}}$: a structurally distinct region near $\partial\mathcal{M}_{\text{adm}}$ characterized by extreme tail compression, elevated susceptibility, and protected giant-component coherence.

Definition 6 (Realizability Boundary Layer)**.** The boundary layer $\mathcal{B}_\delta \subset \mathcal{M}_{\text{adm}}$ is

$$\mathcal{B}_\delta \;:=\; \{L \in \mathcal{M}_{\text{adm}} : m(L) < \delta\}, \quad 0 < \delta \ll 1.$$

The FCC regime corresponds to systems in $\mathcal{B}_\delta$ with TD > 0.80. Deep FCC (TD > 0.95) is the inner boundary layer ($\delta \to 0$).

Four phenomenological properties characterize $\mathcal{B}_\delta$:

1. *Tangential compression*: Trajectories in $\mathcal{B}_\delta$ move primarily along $\partial\mathcal{M}_{\text{adm}}$ (increasing TD) rather than crossing it (bounded $\dot{m}$). This is the slow direction of slow–fast dynamics.

2. *Enhanced representation sensitivity*: In $\mathcal{B}_\delta$, small representation changes (propagation artifacts, sampling limits) are sufficient to produce apparent HARD outcomes (RISC), whereas systems deep in $\mathcal{M}_{\text{adm}}$ are robust. This explains why distant-station Nevada produces HARD while proximal stations remain FULL.

3. *Structural memory*: FCC systems in $\mathcal{B}_\delta$ retain memory of boundary proximity through elevated TD. The yield-dependent trajectory in the NK corpus traces a path through $\mathcal{B}_\delta$ of increasing depth.

4. *Recovery dynamics*: Post-forcing relaxation ($F(t) \to 0$, $R(t) > 0$) drives $m(t)$ upward, moving the trajectory out of $\mathcal{B}_\delta$ toward the interior of $\mathcal{M}_{\text{adm}}$. This is directly observed in the Voyager post-crossing κ_{conn} recovery and 100% FULL years in 2013 and 2017.

13 Connection to Algebraic Topology

We sketch a formal connection between the giant-component protection mechanism and algebraic topology that motivates future rigorous development.

13.1 Vulnerability Graph as Simplicial Complex

For a ladder L at scale κ, the vulnerability graph $G_\kappa(L)$ extends to a clique complex (Vietoris–Rips complex) $\mathcal{V}_\kappa(L)$. The zeroth homology group $H_0(\mathcal{V}_\kappa(L))$ counts connected components. The condition $\text{GR}(\kappa; L) \geq \gamma$ is equivalent to: *$H_0(\mathcal{V}_\kappa(L))$ has a dominant generator (the giant component) whose weight is at least γn.*

Observation 3 (Giant component as homological invariant)**.** The Margin-Confinement Law is equivalent, at the homological level, to: under identity-preserving flows, the dominant H_0 generator of $\mathcal{V}_\kappa(L)$ cannot be annihilated in finite time.

13.2 Persistent Homology and the Margin

The persistence diagram of $G_\kappa(L)$ across the scale parameter κ records births and deaths of connected components. The connectivity margin $m(L_t)$ has a natural TDA interpretation:

$$m(L_t) \;\approx\; \frac{\min_i \left(\text{death}(c_i) - \text{birth}(c_i)\right)}{\max_j g_j},$$

where the minimum is over all non-giant components c_i. This is the normalized persistence of the smallest non-dominant feature — a direct measure of boundary proximity. Theorem-1 activation (zero-lifetime isolated component) corresponds to a zero-persistence H_0 feature, which is forbidden by Theorem 1 for identity-preserving flows.

13.3 Homological Protection and the Background Chain

The background chain (Step 2 of the proof) provides a connected 1-subcomplex of $\mathcal{V}_\kappa(L)$ for all $\kappa \geq \kappa_0$, defining a non-trivial fundamental class in H_0. Under identity-preserving flows:

1. The background chain defines a non-trivial H_0 class (structural homology generator).
2. Continuous deformations (identity-preserving flows) cannot annihilate this class without a discontinuity (homotopy invariance of H_0).
3. Therefore the giant component is homologically protected.

Remark 5 (Open problem: persistent homology stability)**.** A rigorous formulation requires establishing that the background chain defines a *persistent* H_0 feature with lifetime bounded below by a positive constant along any identity-preserving flow. This is expected from IPF-2 and the stability theorem of persistent homology [11], but requires a formal derivation. Establishing this connection is identified as the primary open mathematical problem generated by the Margin-Confinement Law.

14 Boundary Accessibility Without Traversal

Proposition 3 (Asymptotic Boundary Accessibility)**.** *For admissible trajectories satisfying Bounded Structural Rigidity (Principle 1),*

$$m(L_t) \to 0^+ \quad \textit{does not imply} \quad L_t \notin \mathcal{M}_{\text{adm}}.$$

Admissibility persists under arbitrarily small positive margin. The boundary $\partial\mathcal{M}_{\text{adm}}$ is asymptotically accessible but never realized in finite time under identity-preserving flow.

Proof. By Theorem 1, $L_t \in \mathcal{M}_{\text{adm}}$ for all $t \geq 0$. Since $\mathcal{M}_{\text{adm}}$ is an open set relative to the HARD exterior (the HARD verdict requires GR $< \gamma$ persistently, which is a closed condition on the complement), $\mathcal{M}_{\text{adm}}$ contains its asymptotic approach to $\partial\mathcal{M}_{\text{adm}}$. The limit $m(L_t) \to 0^+$ describes compression against $\partial\mathcal{M}_{\text{adm}}$ within $\mathcal{M}_{\text{adm}}$, not escape from it. □

This proposition formalizes the deepest insight of the manuscript: FCC systems are not "barely inside" $\mathcal{M}_{\text{adm}}$ by accident. They are asymptotically approaching $\partial\mathcal{M}_{\text{adm}}$ from the interior, driven by forcing, while the invariant manifold structure of $\mathcal{M}_{\text{adm}}$ prevents them from reaching—let alone crossing—the boundary. Realizability boundaries are transport-accessible without traversal.

15 Empirical Support Summary

Table 5 summarizes the empirical evidence across the tested corpus. Each row represents a distinct domain, instrument, and corpus size, and the final column reports the observed verdict relative to the Margin-Confinement Law.

Table 5: Empirical support for the Margin-Confinement Law across the UNNS corpus. Columns: N = number of evaluations; Verdict = dominant class; MCL = consistency with Theorem 1 ("Confirmed": zero HARD in natural representation; "RISC": HARD explained by representation artifact; "Recovered": all fragmented ladders recover FULL under a locality-preserving Δ-transform; "N/A": pre-law corpus, not directly applicable).

Corpus	N	HARD (natural)	MCL	Note
STRUC-I physical ladders [4]	5,233	0	Confirmed	Zero violations across 14
Phase mapping (α–μ grid) [3]	9,826	0	Confirmed	Zero verdict changes; zer
STRUC-PERC cross-domain [4]	81	1 (TiO_2)	RISC	PLS density representati
NK nuclear explosion corpus [6]	29	0	Confirmed	TD up to 0.997; GR $\geq$ 0
Voyager 1 $\lvert B\rvert$ heliosheath [6]	3,500	0	Confirmed	97.4% FULL; boundary-a
Voyager 2 kinematic channels [6]	471	0	Confirmed	93–97.5% FULL; same pla
Voyager 2 density channel [6]	157	155	RISC	PLS resolution: $\approx$ 57 uni
SN Ia Δ-mag / curvature [6]	2	0	Confirmed	FULL; raw mag: RISC (I
Nevada seismic (proximal) [6]	2	0	Confirmed	CI.PASC, IU.KIP: FULL;
Nevada seismic (IU.HRV) [6]	1	1	RISC	Teleseismic attenuation;
Neutrino detector corpus — Stage 1 (raw) [15]	67	0	Confirmed	5 representations; 6 RISC
Neutrino detector corpus — Stage 2 (Δ-lifting) [15]	67	0	Recovered	100% recovery (9/9 fragn
Total natural-representation HARD	—	**0**		(excluding TiO_2 density

Observation 4 (Corpus-level confirmation)**.** Across all tested corpora, zero instances of genuine dynamical crossing into the HARD regime under identity-preserving evolution are observed. Every HARD outcome in the corpus is attributable to a specific representation artifact (IPF violation), fully consistent with Proposition 2.

15.1 Margin Distribution Across Corpora

Table 6 reports the estimated connectivity margin range for each corpus, using the practical proxy of Definition 2. The table quantifies boundary proximity and verifies that all tested physical systems remain at positive margin in their natural representation.

The RISC cases (TiO_2 density, Nevada IU.HRV, Voyager 2 density) produce apparent negative margin under the proxy (consistent with being outside $\mathcal{M}_{\text{adm}}$ in their distorted projection) while the corresponding source systems have positive margin in their natural or Δ-lifted representations. This directly confirms Proposition 2: RISC moves the apparent position outside $\mathcal{M}_{\text{adm}}$ while the true position remains within it.

Table 6: Estimated connectivity margin range across tested corpora. Margin computed via the practical proxy (Definition 2). $m_{\min}$: minimum margin observed in the corpus; $m_{\max}$: maximum margin; Class: dominant structural class. All physical corpora in natural representation: $m_{\min} > 0$.

Corpus	$m_{\min}$ (est.)	$m_{\max}$ (est.)	Mean TD	Class	Boundary proximity
STRUC-I / phase mapping [3]	0.45	0.98	0.41	FULL	Interior
STRUC-PERC cross-domain [4]	0.15	0.92	0.53	FULL/GIANT	Moderate interior
NK corpus (all stations) [6]	0.03	0.38	0.83	GIANT/TAIL	Boundary layer
NK corpus IC.MDJ [6]	0.01	0.04	0.994	TAIL (FCC)	Deep FCC zone
Voyager 1 heliosheath [6]	0.35	0.88	0.78	FULL	Near-interior
Voyager 1 crossing epoch [6]	0.18	0.55	0.78	FULL/GIANT	Boundary-adjacent
Voyager 2 $V/T/w$ [6]	0.30	0.82	0.71	FULL	Interior
SN Ia Δ-mag [6]	0.12	0.32	0.81	FULL	Boundary layer
RISC cases (all)	< 0 (apparent)	—	variable	HARD	*Outside* $\mathcal{M}_{\text{adm}}$ (projecti

15.2 Simulation Illustration of Confinement

To provide a concrete dynamical illustration of the confinement mechanism, we describe a representative 1D ladder simulation:

- *Setup*: $n = 1{,}024$ ladder elements; background gaps drawn from Uniform$(0, 1)$ ($n_b = 973$ elements, 94.8%); tail gaps drawn from Pareto$(\alpha = 2)$ with scale $10 \times \text{median}(g_b)$ ($k = 51$ elements, 5.0%); forcing ramp: background scale reduced by factor $f(t) = 1/(1 + 0.1t)$ over $T = 20$ time steps.
- *Observation*: TD increases from 0.73 to 0.92 over the ramp; GR remains at 1.000 throughout (STRUC-PERC-I evaluated at each step); margin proxy m decreases from 0.42 to 0.08 without reaching zero.
- *Interpretation*: The simulation traces a trajectory from the interior of $\mathcal{M}_{\text{adm}}$ into the FCC zone of $\mathcal{B}_\delta$, consistent with the flow equations (9)–(11) and the effective potential (17). No HARD outcome is triggered despite a factor of 5× increase in TD.

Full ensemble statistics and parameter sensitivity analysis are deferred to a dedicated computational supplement.

15.3 Statistical Bounds on Zero Genuine Crossings

The corpus of $N_{\text{tot}} = 15{,}401$ evaluations (Table 5) contains zero genuine HARD outcomes from natural-representation physical systems. Using a simple binomial model, the probability that a genuine crossing rate $p > 0$ would produce zero observed crossings across N trials is $(1 - p)^N$. Setting a confidence level of 99%:

$$(1 - p)^{15401} \geq 0.01 \implies p \leq 1 - 0.01^{1/15401} \approx 3.0 \times 10^{-4}.$$

At the 99% confidence level, the genuine crossing rate is bounded above by $p < 3.0 \times 10^{-4}$ crossings per evaluation. This is consistent with the Margin-Confinement Law's prediction of $p = 0$ for identity-preserving flows, and with the observation that all HARD outcomes are attributable to RISC (a representation artifact, not a genuine crossing).

At 99.9% confidence—requiring that the probability of observing zero crossings under any nonzero rate be at most 0.1%—the exclusion criterion is stricter, so only larger rates can be ruled out: the bound widens to $p < 4.5 \times 10^{-4}$. This is the expected direction: a higher confidence level yields a more conservative (larger) upper bound at fixed N, since we require the data to be more extreme before excluding a given rate. Both bounds are consistent with the Margin-Confinement Law's prediction of $p = 0$ for identity-preserving flows. These bounds hold under the assumption of independent evaluations; cross-corpus dependencies would tighten the effective N but not qualitatively change the conclusion.

16 Theoretical Implications

16.1 Admissibility as a Dynamical Invariance Principle

The Margin-Confinement Law transforms admissibility from a static criterion into a dynamical invariance principle. Earlier UNNS manuscripts established that physical systems consistently occupy admissible regions [1, 2, 3]. The present manuscript derives *why*: the combination of the USL, the bounded-gap background chain, and Bounded Structural Rigidity together force identity-preserving trajectories to remain in $\mathcal{M}_{\text{adm}}$. Admissibility is not a lucky empirical coincidence; it is a structural consequence of the physical process class.

16.2 Reframing the UNNS Program

The law reframes the UNNS research program at its deepest level. Earlier work classified structural states. This manuscript establishes that those states are dynamically protected. The program's subject matter shifts from *what class does a system occupy?* to *how does a system move within the invariant admissible manifold?* This is the transition from static realizability theory to dynamic realizability geometry.

16.3 Structural Collapse as an Observational and Representational Artifact

Two distinct claims about structural collapse are now supported by the corpus. The first is the law's primary result; the second is an empirical extension:

First: for physical systems whose evolution satisfies the IPF conditions, *true* structural collapse (persistent HARD fragmentation) does not occur in their natural representation. What is observed as collapse is, in every documented case in the tested corpus, either:

1. a representation artifact (RISC), correctable by Δ-lifting or improved measurement resolution;
2. an external projection through a non-admissible observational chart;
3. or a genuinely discontinuous identity-rupture event outside the IPF class.

Second: in the neutrino corpus, the full collapse-recovery cycle is demonstrated operationally within a single corpus. HARD fragmentation can be *manufactured* by a representation choice (TMVA classifier projection), *manufactured and removed* by a deeper

embedding (deepL → FULL), and *independently recovered* by a local, non-parametric transform (Δ-lifting → 9/9 FULL). This cycle — manufactured collapse, embedding recovery, transform recovery — is the strongest available evidence that apparent structural collapse is a property of the observational chart, not of the source system. The source system never left $\mathcal{M}_{\text{adm}}$.

This aligns with the corpus evidence in Table 5 and provides the theoretical grounding for the empirical observation that structural coherence is *far more common* than structural collapse across natural physical systems — and that apparent collapse is recoverable by locality-preserving representations.

Observational collapse vs. ontological collapse. The corpus motivates a distinction that sharpens the law considerably. *Observational collapse* corresponds to representation-specific fragmentation (HARD verdict in a particular chart), arising from RISC mechanisms (binning, orientation reversal, n-poverty, attenuation). *Ontological collapse* would require persistent loss of admissible continuity across *all* locality-preserving transforms — a total destruction of the source system's giant-component structure that no representation could recover. No such ontological collapse was observed in the present corpus. Every HARD outcome is observational: it exists in a specific chart and is absent from locality-preserving transforms of the same source ladder. This distinction clarifies the precise content of the law: what Theorem 1 forbids for identity-preserving flows is ontological collapse, not observational collapse (which RISC can produce without violating the theorem).

16.4 Representation-Relativity of Realizability Class

A consequence of the neutrino corpus that goes beyond RISC as a failure mode is the following: across five representations of the same detector process, the realizability verdict varies across all four classes (FULL, GIANT, TAIL, HARD). This is not a defect of the framework. It is a structural fact about realizability space.

Observation 5 (Admissibility as a joint function)**.** The neutrino corpus establishes, corpus-scope, that the realizability class is not determined by the source system alone. It is jointly determined by the source and the observational representation:

$$\text{class}(L) \;=\; \mathcal{F}(\text{source},\ \text{representation}).$$

This does not contradict the Margin-Confinement Law. The law governs the source's true position in $\mathcal{M}_{\text{adm}}$, which is invariant under identity-preserving dynamics. What varies across representations is the *observed projection* of that position onto different realizability charts. HARD verdicts in non-natural representations diagnose chart distortion, not physical state.

This extends the MCL from a dynamical confinement theorem to an *observational geometry* principle: which chart a system appears to occupy is a property of the observation, not merely of the system. This pushes the UNNS framework toward a category-like conception of realizability — one in which the morphisms (representational maps) between charts are as structurally significant as the objects (systems) they connect.

The representational tower. The neutrino corpus implicitly defines a representational hierarchy:

$$\mathcal{R}_0 \;\to\; \mathcal{R}_1 \;\to\; \mathcal{R}_2 \;\to\; \cdots \;\to\; \mathcal{R}_k$$

where each level is a progressively more processed observational representation of the same source events: raw detector geometry ($\mathcal{R}_0$), reconstructed hit and angular distributions ($\mathcal{R}_1$), classical multivariate classifier spaces ($\mathcal{R}_2$), deep-learning embedding spaces ($\mathcal{R}_3$). Realizability class changes as maps between levels either preserve or destroy admissibility geometry. Some maps are admissibility-preserving (the deep-L embedding recovers FULL); some are admissibility-destroying (TMVA binning introduces RISC). Formalizing which map classes preserve $\mathcal{M}_{\text{adm}}$-membership is an open theoretical problem, identified here as a direction for future work.

Deep embeddings as structural renormalization. The repeated pattern — TMVA fragmenting, deepL recovering FULL — has a natural structural interpretation beyond "classifier improvement." Deep embeddings may be reconstructing *latent connectivity* that sparse observational coordinates destroyed. In this reading, the deep-learning map $\mathcal{R}_2 \to \mathcal{R}_3$ resembles a renormalization-group flow toward a stable fixed manifold: irrelevant degrees of freedom (the gap structure responsible for RISC) are integrated out, and the residual representation inhabits the admissible region. This is a structural analogy, not a derivation; it is recorded as a conjecture motivating future theoretical work on admissibility-preserving coordinate maps.

Latent continuity conjecture. The empirical pattern across all tested domains — natural representations are admissibility-safe; projections produce apparent HARD; deeper representations recover admissibility — admits a unified hypothesis:

Hypothesis 1 (Latent structural continuity — corpus-supported)**.** Physical processes that satisfy the IPF conditions possess latent connectivity that survives in the source system even when destroyed in a particular observational projection. This latent connectivity is recoverable by representations that preserve more of the source's topological structure, and in particular by locality-preserving transforms such as Δ-lifting. *Corpus support*: in the neutrino detector corpus, the Δ-lifting transform recovers FULL percolation in all nine ladders that fragment under TMVA or deep-learning representations (100% recovery rate), confirming that the gap structure of the raw ladder retains intact admissibility geometry that classifier projections obscure. This hypothesis has passed its first operational test.

This hypothesis is not derived from the axiomatic UNNS framework; it requires independent theoretical development. It aligns structurally with Takens-type embedding theorems, manifold-unfolding in topological data analysis, and information-geometric recoverability — connections that are identified as open formal problems. A first step toward formalization is given in Conjecture 1.

FCC in informational spaces. Prior to the neutrino corpus, FCC-like states — extreme TD, preserved GR = 1.000 — were observed mainly in physically extreme systems (nuclear detonations, heliopause transitions, Voyager plasma boundary). The neutrino

corpus places FCC-like states inside informational reconstruction pipelines: the high-TD ladders (`deepL_train_B`, TD = 0.9985; `TMVA_Bkg_Train`, TD = 0.9952) arise not from physical forcing but from the statistical geometry of classifier outputs. This suggests that the FCC regime is not inherently physical in origin. It is a universal boundary phenomenon of any constrained realizability space under extreme distributional compression, regardless of whether the compression comes from a physical process or an informational one. This generalization is recorded as an empirical observation from the corpus; its formal status depends on extending the IPF conditions to informational flows, which is a distinct open problem.

16.5 Representation-Covariant Admissibility

The Δ-lifting results introduce a phenomenon that is not covered by the Margin-Confinement Law as currently derived: not merely the preservation of admissibility across representational changes (as explained by RISC), but the *recovery* of admissibility by a representational transform applied to a fragmented ladder. This section formalizes the theoretical gap and records the minimal conjecture the data motivates.

Three phenomena to distinguish. The neutrino corpus, including its Δ-stage, reveals three structurally distinct phenomena that must be separated precisely:

Table 7: Three representation-related admissibility phenomena. Only Type I is covered by Theorem 1; Types II and III require additional theoretical development.

Type	Phenomenon	Mechanism	Status
I	**Dynamical preservation**	admissibility preserved under IPF evolution	Theorem 1
II	**Representational loss (RISC)**	admissibility destroyed by non-locality-preserving chart	Proposition
III	**Latent continuity recovery**	admissibility recoverable under locality-preserving transform	Hypothesis

Transform class taxonomy. To formalize Type III, a classification of representational transforms by their effect on admissibility geometry is necessary. The following taxonomy is proposed based on the corpus evidence; it is an organizing conjecture, not a derived classification.

Table 8: Proposed transform class taxonomy by structural effect on admissibility. Corpus evidence is corpus-scoped; formal definitions of each class require independent development.

Transform class	Structural effect	Neutrino corpus exam
Projection / binning transforms	Can induce RISC; destroy local connectivity	TMVA histograms →
Identity-breaking transforms	May genuinely destroy admissibility; violate IPF	Non-invertible coarse-
Locality-preserving transforms	Preserve or recover admissibility; resolve RISC artifacts	Δ-lifting → 9/9 FULL
Relational / embedding transforms	May recover latent connectivity; approach admissibility	deepL embeddings →
Orientation-reversing transforms	Invert ladder chart; create void at boundary	`hBkgeff` → HARD (R

Conjecture 1 (Representation-covariant admissibility)**.** Let L be a realizability ladder generated by a physical process satisfying the IPF conditions, and let ϕ be a locality-

preserving representational transform (informally: a map that preserves the local ordering and gap structure of L, such as $\Delta L = (|x_{i+1} - x_i|)$). Then the transformed ladder $\phi(L)$ is also in $\mathcal{M}_{\text{adm}}$:

$$L \in \mathcal{M}_{\text{adm}} \implies \phi(L) \in \mathcal{M}_{\text{adm}}.$$

Equivalently: structural admissibility is invariant under locality-preserving representational transforms.

Corpus status: The Δ-lifting stage of the neutrino corpus provides the first operational support for this conjecture. All nine fragmented ladders (in which RISC destroyed apparent admissibility) recover FULL under Δ-lifting, and no natural-representation admissible ladder fails Δ-lifting except under identified n-poverty artifacts (Δ-ladder length below reliable percolation threshold).

Status: Corpus-supported conjecture. Not derived from the UNNS axioms. A proof would require a formal definition of "locality-preserving transform" in realizability geometry, and a connection between gap-structure preservation and vulnerability-graph connectivity. This is identified as the most important open theoretical problem in the present manuscript.

Continuity beneath fragmentation. The Δ-results suggest a subtler reading than the language of "recovery" implies. Recovery language carries the implication that the source system fragmented and later healed. But the 100% recovery rate under a non-parametric local transform suggests something different: in the RISC cases, admissible continuity was *never lost* — it was only hidden. The giant component persists beneath the fragmented observational chart, invisible in the classifier projection but immediately visible once the local gap structure is extracted. Δ-lifting does not restore continuity; it *uncovers* continuity that was present all along. This is consistent with Hypothesis 1 and supports the stronger reading of Conjecture 1: admissibility is an intrinsic relational property of the source process, not a property of any particular embedding.

FCC coherence is locally relational. A secondary implication of the Δ-stage is that FCC-like states — extreme tail dominance with preserved giant-component coherence — persist under Δ-lifting. The 34 FCC-like instances in Δ-space (up from 5 in raw space) arise not only in the already-FCC raw ladders (which remain FCC) but also in ladders that had modest raw TD but nearly uniform gap structure (bkg2*C14 series, $\text{TD}_{\text{raw}} \approx 0.25$–$0.49$, $\text{TD}_{\Delta} \approx 0.99$). This demonstrates that FCC coherence is not an amplitude property of the ladder values; it is a *local-relational property of the gap structure*. A system in the FCC regime retains its coherent compression signature across amplitude stripping, drift removal, and coordinate flattening — as long as the local ordering is preserved. This constitutes a substantial strengthening of the FCC interpretation (Corollary 1) and is recorded here as an empirical principle:

Observation 6 (FCC coherence is locally relational)**.** In the tested corpus, FCC-like states (TD ≥ 0.985, GR $= 1.000$) are preserved under Δ-lifting. This indicates that the FCC coherence signature is carried by the local gap structure of the realizability ladder, not by global amplitude values. FCC is a local-relational boundary phenomenon, not a global-amplitude forcing effect.

Minimal Structural Principle. The corpus results across RISC, Δ-lifting, FCC persistence, Voyager boundary transport, and latent continuity extraction admit a concise synthesis:

> **Minimal Structural Principle (corpus-scoped).** Admissible structures tend to preserve or recover relational continuity under identity-preserving evolution and locality-preserving representation transforms.

This principle is *not* a derived theorem. It is an empirical generalization from the tested corpus, recording the consistent pattern: confinement under dynamics (Theorem 1), representational recovery under Δ-lifting (Conjecture 1), boundary-adjacent resilience under extreme forcing (Corollary 1), and long-range transport coherence (Voyager corpus) all instantiate the same underlying structural tendency. The principle unifies these phenomena without asserting universality beyond the tested evidence.

16.6 Connection to the Universal Structural Law

Corollary 3 establishes that the USL (1) acquires a dynamical protection interpretation under the Margin-Confinement Law:

> The Universal Structural Law, combined with Bounded Structural Rigidity, implies that admissible structures are driven toward the boundary of $\mathcal{M}_{\text{adm}}$ under extreme forcing but cannot cross it under identity-preserving evolution (Theorem 1). This can be interpreted as: the dynamical manifestation of the USL is the confinement of physical evolution within the admissible manifold. Realizable systems satisfy (1) not merely at a point in time but along their entire trajectory.

16.7 Relation to Established Frameworks

Table 1 and 2 identify structural parallels with normally hyperbolic invariant manifolds, Lyapunov stability theory, topological phase protection, and percolation criticality. The Margin-Confinement Law does not reduce to any of these frameworks but occupies a complementary position: it identifies $\mathcal{M}_{\text{adm}}$ as an invariant manifold in a representation-dependent, finite-size, heavy-tailed, physically realized dynamical system — a setting not covered by classical invariant manifold theory.

Future work will seek formal connections between the giant component and established topological invariants (Chern numbers, winding numbers), and will test the Margin-Confinement Law's predictions in additional high-forcing datasets (stellar flares, magnetar spin-down, laboratory shock-tube experiments).

16.8 Programmatic Alignment with *Beyond Fragmentation*

The Margin-Confinement Law and the companion manuscript *Beyond Fragmentation* [6] are the empirical and theoretical halves of the same boundary framework. Their relationship is explicit here because it is too structurally important to leave implicit.

Empirical precursor. *Beyond Fragmentation* assembled the observational evidence for behavior near the realizability boundary. Across nuclear explosions, Voyager heliosheath transport, supernovae, seismic stations, and the Δ-restabilization stage, the manuscript identified a convergent pattern: systems approach structural extremity — TD $\to 1$, $m(L_t) \to 0^+$, GR near the threshold — while consistently remaining admissible. Persistent HARD fragmentation did not occur. That manuscript already defined the signature of the Forced Coherent Collapse regime and noted, without a theoretical explanation, that the boundary $m = 0$ was approached from above, not crossed. The empirical geometry of confinement was assembled there.

Theoretical completion. The Margin-Confinement Law explains why. The boundary $\partial\mathcal{M}_{\text{adm}}$ is not merely empirically uncrossed; it is *dynamically non-penetrable* for identity-preserving flows. The law provides the mechanism — background chain persistence, Lyapunov functional boundedness, normal hyperbolicity of the NHIM — that converts the observed asymptotics into a structural principle. FCC, previously a descriptive regime in *Beyond Fragmentation*, becomes a precisely characterized near-boundary asymptotic state under the present framework. Voyager's long-duration near-boundary reorganization without HARD transition, previously an empirical observation, becomes the first direct realization of a non-crossability trajectory within the dynamics of the admissibility manifold. RISC, previously a classification of representational failures, acquires a principled foundation: apparent observational escape beyond $\partial\mathcal{M}_{\text{adm}}$ while the underlying trajectory remains confined. The Δ-recovery results, in this reading, reveal not a representation improvement but a projection back toward the admissible trajectory that was never actually left.

The unified boundary theory. Together the two manuscripts address complementary questions:

> *Beyond Fragmentation*: What do admissible systems empirically do near realizability boundaries? They compress, localize, saturate tails, preserve coherence, and approach $m \to 0^+$ without crossing.
>
> *Margin-Confinement Law*: Why do admissible systems fail to cross those boundaries? Because admissibility dynamically confines identity-preserving trajectories away from the non-realizable region, and structural continuity is recoverable beneath representational projections that falsely suggest otherwise.

The combined implication is that realizable physical structures may become arbitrarily stressed, tail-dominated, and near-critical while remaining structurally admissible: physical collapse, in the sense relevant to realizability geometry, is asymptotic rather than trans-boundary. *Beyond Fragmentation* assembled the evidence; the present manuscript supplies the missing theoretical principle.

17 Conclusion

We have derived and established the Margin-Confinement Law: once a physical system is represented inside the admissible region $\mathcal{M}_{\text{adm}}$ of realizability space, its identity-preserving

dynamical evolution remains confined to $\mathcal{M}_{\text{adm}}$ for all time within the IPF class. The admissibility boundary $\partial\mathcal{M}_{\text{adm}}$ is a dynamically non-penetrable invariant manifold: approachable asymptotically, generating the Forced Coherent Collapse regime, but non-crossable under identity-preserving evolution.

The derivation proceeds by contradiction from three axiomatic principles (the USL, PRP, and Bounded Structural Rigidity), introduces the Identity-Preserving Flow class, constructs an explicit structural invariant (the background chain), and identifies the Lyapunov-type confinement functional that converts non-crossability from a topological into a dynamical statement.

The law is empirically supported by a corpus of 15,401 evaluations across fourteen physical domains, one biological domain, one astrophysical boundary-transport system, and one sparse-event detector reconstruction domain. In the neutrino corpus — the theoretically richest single corpus in the program — all observed HARD outcomes in raw space are attributable to representation artifacts, and a subsequent Δ-lifting stage recovers FULL percolation in every one of the nine fragmented ladders (100% recovery rate), expanding FCC-like states from 5 to 34 instances. This completes the first full manufactured-collapse, embedding-recovery, transform-recovery cycle in the UNNS program.

The manuscript is built on two pillars:

Pillar I — Forbidden collapse. Persistent HARD fragmentation under identity-preserving dynamics is forbidden by Theorem 1. All observed HARD outcomes in the tested corpus are representation artifacts. The law is a *dynamical invariance principle*: the UNNS program studies not which class a system occupies, but how it moves within $\mathcal{M}_{\text{adm}}$.

Pillar II — Recoverable continuity. Structural admissibility appears to be an intrinsic local-relational property that survives in a source system even when observational projections destroy it, and that re-emerges under sufficiently locality-preserving representations. This is no longer purely conjectural: the neutrino Δ-stage provides the first operational demonstration. It motivates Conjecture 1 on representation-covariant admissibility and Observation 6 that FCC coherence is locally relational.

The five most important consequences are:

1. **FCC is structurally inevitable**, not precarious. Systems in $\mathcal{M}_{\text{adm}}$ under extreme forcing are driven against $\partial\mathcal{M}_{\text{adm}}$ but protected from crossing it.

2. **Apparent structural collapse is a representational artifact.** Persistent HARD fragmentation under identity-preserving dynamics is forbidden by Theorem 1; observed HARD outcomes diagnose representation artifacts, not physical collapse. In the neutrino corpus, the entire collapse-recovery cycle is demonstrated within a single source process.

3. **Admissibility is a dynamical invariance principle.** The UNNS program now studies how systems move through $\mathcal{M}_{\text{adm}}$ under forcing — a shift from static classification to dynamical realizability geometry.

4. **Realizability class is representation-relative.** The same source system can appear FULL, GIANT, TAIL, or HARD depending on the observational chart. Deep-

learning embeddings and Δ-lifting both act as admissibility recovery operators, restoring FULL from RISC-fragmented representations.

5. **Structural continuity appears to be intrinsically local-relational.** FCC coherence survives Δ-lifting (amplitude stripping, drift removal, coordinate flattening), and all RISC artifacts are resolved by a locality-preserving transform. This points toward representation-covariant admissibility as the law's deepest empirical generalization (Conjecture 1).

All findings are scoped to the tested corpus and the IPF class of dynamics. The Margin-Confinement Law is presented as a derived consequence of established UNNS axioms within their stated domain of validity. Conjecture 1, Hypothesis 1, and Observation 6 are corpus-supported but not axiomatically derived; their formal development constitutes the primary open theoretical direction arising from this manuscript.

References

[1] UNNS Substrate Research Program. The Universal Structural Law: Admissibility Bounds on Ordering Instability. An Empirical Investigation Across 3,073 Physical Ladders Spanning Thirteen Physical Domains. *UNNS Working Manuscript*, 2026.

[2] UNNS Substrate Research Program. The Percolative Realizability Principle: Realizability Structure of Admissible Configurations and the Revised Necessary Direction. *UNNS Working Manuscript*, 2026.

[3] UNNS Substrate Research Program. Phase Mapping of Structural Regimes in the UNNS Substrate: Bounded Structural Rigidity and Representation-Driven Structure. *UNNS Working Manuscript*, April 2026.

[4] UNNS Substrate Research Program. STRUC-PERC-I v2.4.0 Cross-Domain Percolation Corpus Analysis. 81 Evaluations Across 14 Domains. *UNNS Corpus Report*, 2026.

[5] UNNS Substrate Research Program. Structural Realizability and Dual Observability in the Admissibility Manifold. *UNNS Working Manuscript*, 2026.

[6] UNNS Substrate Research Program. Beyond Fragmentation: Admissibility and Coherent Collapse Across Extreme Physical Transitions. *UNNS Working Manuscript*, 2026. https://unns.tech/media/unns/extreme_physical_transitions/BeyondFragmentation.pdf

[7] UNNS Substrate Research Program. Foundations of the UNNS Substrate: From Universal Admissibility to Structural Regime Theory. *UNNS Foundation Document*, 2026.

[8] UNNS Substrate Research Program. Local Geometry of Realizability Boundaries in the UNNS Substrate. *UNNS Working Manuscript*, 2026.

[9] UNNS Substrate Research Program. Structural Trajectories in Realizability Space: Voyager 2 Plasma as a Time-Resolved Test of the UNNS Substrate. *UNNS Working Manuscript*, 2026.

[10] H. Bernien et al. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551:579–584, 2017.

[11] G. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.

[12] N. Fenichel. Geometric singular perturbation theory for ordinary differential equations. *Journal of Differential Equations*, 31(1):53–98, 1979.

[13] M. W. Hirsch, C. C. Pugh, and M. Shub. *Invariant Manifolds*. Springer, Berlin, 1977. (Lecture Notes in Mathematics, vol. 583.)

[14] IRIS Data Management Center. North Korea Nuclear Explosion Special Event, 12 February 2013: Seismic Data and Relocation Exercise. Incorporated Research Institutions for Seismology (IRIS), 2013. https://ds.iris.edu/ds/nodes/dmc/specialevents/2013/02/12/north-korea-nuclear-explosion/

[15] UNNS Substrate Research Program. Neutrino Detector Observational Corpus: STRUC-PERC-I v2.5.0 Analysis of Representation-Sensitive Admissibility in Liquid-Scintillator Detector Reconstruction Pipelines. *UNNS Corpus Analysis Report*, 2026. Data source: https://www.scidb.cn/en/detail?dataSetId=bfe84bf2f2a94a539d0bd9dfab727cdd