Foundations of the UNNS Substrate: From Universal Admissibility to Structural Regime Theory

A Cross-Domain Synthesis of Persistence, Boundary Behavior, Operator-Selective Response, and the Geometry of Structural Admissibility

Instrument: STRUC-I v1.0.4
Domains: Atomic / Molecular / Nuclear / Hadronic / Geoid / Seismology / CMB / Cosmic Web / Condensed Matter / Biological
Constants: α, μ, α_s, α_G
Corpus: $>$ 1,500 ladders; $>$ 150,000 assessments
Protocol: Falsification-first; preregistered criteria
Status: Foundation Document — Intermediate Stage
Date: 2026

Abstract. The UNNS Substrate Research Program has assembled, across a systematic sequence of domain-specific investigations, an empirical corpus broad enough to support a synthesis that goes materially beyond the original single-inequality statement. This document presents the first complete statement of that synthesis as a foundational framework. The accumulated corpus establishes three interlocking structural claims: *universal admissibility* (the inequality $\mathrm{inv}(P_\varepsilon; L) \leq \nu(V_\varepsilon(L))$ holds at all physical constant values across every tested domain, with no hard violations at physical parameters), *stratified occupation* (physical systems do not uniformly populate the admissible region but distribute across a structured landscape of characteristic structural pressure and boundary proximity), and *operator-selective activation* (only specific constant-deformation directions in specific domains generate genuine structural reorganisation, while others act metrically). Together these establish what we call *structural regime theory*: a framework in which the central object is not merely a universal inequality but a *regime geometry* of persistence, susceptibility, and boundary approach, equipped with a formal admissibility manifold, a notion of structural direction in operator space, and a selection principle that makes the inequality constitutive rather than merely descriptive. The document further develops the theory of the phase interface, formalises the metric/structural distinction as a transformation-theoretic pillar, presents the constant-anchoring hypothesis with its full evidence chain, and derives cross-corpus deductions—including substrate-independence and the structural sensitivity tensor—that are only visible when all results are read together. It closes with open mathematical questions and experimental predictions that follow necessarily from the framework.

Contents

Part I

Core Law and Empirical Basis

1. Conceptual Grounding: The Problem of Structural Persistence

1.1. Why Ordered Sequences?

Physics has many ways of representing the properties of systems. Dynamical trajectories describe how states evolve; probability distributions describe statistical ensembles; symmetry groups describe transformation invariances; correlation functions describe statistical dependencies. Each representation has its natural domain and its natural questions.

The UNNS programme begins from a different representation: *ordered sequences of values* produced by physical systems when their energy levels, spectral gaps, harmonic coefficients, activity values, or other ranked quantities are arranged in natural order. These are called *ladders*. The natural question for a ladder is not "how does it evolve?" or "what symmetry does it have?" but: *how much structure does it contain, and can that structure persist across the variation capacity of the sequence?*

This is not a merely technical choice of representation. It reflects a substantive claim: that ordered relational structure is a primary property of physical systems, one that is not already captured by dynamical equations, symmetry arguments, or statistical descriptions. Ordered sequences encode the relational geometry of a system's observable quantities—how its energy levels relate to each other, how its spectral gaps vary, how its harmonic content is organised—in a way that is independent of the dynamical mechanism generating those quantities. A sequence of energy levels is the same structural object whether it arises from a harmonic oscillator, a Coulomb potential, or a strong-force confinement potential. The UNNS programme asks whether this structural object satisfies a universal admissibility condition—and if so, what the geometry of that condition tells us about physical reality.

1.2. What is Structural Persistence?

The concept of *structural persistence* is central to the entire framework, and it requires careful statement. Persistence is not conservation. A conserved quantity has the same value at all times; persistence in the structural sense means something different: a structural feature of an ordered sequence that is stable across a family of threshold scales.

More precisely: given a sequence $L = (L_1, L_2, \ldots, L_n)$ and a threshold $\varepsilon > 0$, a *persistent* *structural element* is a feature of L that survives the coarse-graining induced by ε. The ε-persistence set $P_\varepsilon(L)$ counts how many ordered pairs (L_i, L_j) maintain their relational order (the gap $L_j - L_i > \varepsilon$) under threshold ε. The *invariant count* $\text{inv}(P_\varepsilon; L)$ measures how many of these persistent relations are simultaneously invariant under the natural partial order of the sequence.

The variation set $V_\varepsilon(L)$ measures the total variation capacity: how many threshold-respecting relations could in principle exist, given the length and spread of the sequence. The admissibility ratio $A_\kappa = \text{inv}(P_\varepsilon; L)/\nu(V_\varepsilon(L))$ measures the fraction of the variation capacity that is actually occupied by persistent structure.

> **Structural persistence**, in the UNNS framework, is the property of an ordered sequence whereby its relational structure—the pattern of invariant gap relations among its elements—is stable under threshold coarse-graining across a range of scales $\varepsilon \in [\varepsilon_{\text{min}}, \varepsilon_{\text{max}}]$. It is not a single-threshold property but a multi-scale profile.

Why is persistence non-trivial? Because the variation capacity $\nu(V_\varepsilon(L))$ grows with the length and complexity of the sequence. A sequence can have many elements without those elements forming persistent relational structure; persistence requires that the relational pattern is organised in a way that withstands coarse-graining. The USL asserts that physical sequences cannot have more persistent structure than their variation capacity allows.

1.3. What Existing Physics Does Not Provide

The structural persistence framework fills a gap that existing physical theories do not address. This is not a criticism of those theories; it is a statement of their scope.

Dynamical theories (classical mechanics, quantum mechanics, quantum field theory) specify the equations of motion that govern how a system evolves. They predict energy eigenvalues, transition rates, correlation functions, and scattering amplitudes. They do not, however, provide any general constraint on the *ordered structure* of those eigenvalues. The Schrödinger equation does not forbid any particular arrangement of energy levels; it only specifies which arrangements are consistent with a given Hamiltonian. There is no dynamical theorem that bounds the ratio of persistent to total relational structure in an energy spectrum.

Symmetry-based approaches (group theory, representation theory) classify which energy level structures are consistent with a given symmetry. They predict degeneracies, selection rules, and spectral multiplets. They do not constrain the internal relational organisation of the levels within each symmetry sector. Two systems with the same symmetry group can have very different structural pressures $\bar{\rho}$.

Statistical mechanics describes ensembles of systems and predicts average properties. It provides no constraint on the structural organisation of individual ordered sequences; it averages over them.

Renormalisation group and effective field theory provide scaling laws for how observables change with energy. They do not address the admissibility of the ordered structure of those observables.

In summary: existing physics tells us what values physical observables can take, and how those values change under dynamics or symmetry transformations. It does not constrain the structural admissibility of ordered sequences of those values. The UNNS programme addresses this gap.

1.4. The UNNS Perspective: What Is the Substrate?

The term "substrate" in the UNNS programme refers to something more specific than "everything." The substrate is *the collection of relational structures that can stably exist*—where "stably exist" means: persist under the natural coarse-graining of threshold variation, satisfy the admissibility inequality, and maintain those properties across the domain-appropriate deformations of fundamental parameters.

The substrate is not a material. It is not defined by what physical systems are made of. It is defined by what kind of relational organisation they can sustain. This is why the programme can include, within a single framework, atomic energy levels, nuclear spectra, planetary gravity field harmonics, cosmic web orientation patterns, and biological fitness landscapes. All of these are ordered relational structures. The substrate is the set of such structures that satisfy the admissibility condition.

> **The Substrate** is the set of all relational configurations—ladder representations of physical, biological, or other ordered systems—that satisfy the admissibility inequality at the relevant operator values. It is defined by structural organisation, not by physical composition or governing equations.

This framing explains why the programme's two key findings—universal boundedness and substrate-independence—are related. Universal boundedness says the substrate contains all physically observed relational structures. Substrate-independence says the boundary of the substrate is drawn by relational organisation alone, not by the ontological category (quantum system, gravitational field, biological polymer) of the underlying medium.

2. The Empirical Corpus: Scope and Organisation

2.1. Domains, Instruments, and Data Sources

The corpus was assembled using a family of falsification chambers implementing the core inequality with $\kappa \in [0.01, 1.0]$ and $\varepsilon = \kappa \cdot \text{median}(\Delta L)$. The principal instrument, STRUC-I v1.0.4, computes $A_\kappa = \text{inv}(P_\varepsilon;\, L)\,/\,\nu(V_\varepsilon(L))$ for each (κ, L) pair and aggregates over a 17-point κ-grid to produce $\bar{\rho}$ and $A_\kappa^{\min}$ for each ladder. Specialised instruments (QM-I/QM-II for atomic spectra, GRAV-I for geoid harmonics, CW-I for cosmic web structure, STRUC-BIO-I/II for biological ladders) implement domain-appropriate ladder constructions before passing to STRUC-I for the admissibility evaluation.

All protocols are preregistered with explicit falsification criteria before runs are executed. Results are reported without post-hoc reclassification.

The tested domains span ten families:

- **Atomic spectra** (QM-I / QM-II): H, He, Li, Na; α-sweep; 1,304+ ladders.
- **Molecular rovibrational** (μ-column, HITRAN): CO, H_2, N_2, HCl, HD; Tier-A decomposition.
- **Nuclear spectroscopy** (ENSDF): ^{48}Ca, ^{150}Nd, ^{208}Pb and 12 further isotopes; α and α_s columns.
- **Hadronic spectroscopy** (PDG): charmonium J/ψ family; α_s-column.
- **Planetary geoid** (GRAV-I): Earth (EIGEN-6C4), Mars (JGM85F01), Moon (AIUB-GRL350A); Tier-A $C_{20}^{\rm rot}$ decomposition; α_G-column.
- **Seismology** (LXV): global arrival-time sequences.
- **Cosmic microwave background** (CMB-I through CMB-SPECTRA-Σ): Planck 2018 TT/TE/EE power spectra.
- **Cosmic web** (CW-I v2.1.0): DESI, SDSS, 2MRS orientation ladders.
- **Condensed matter**: SiO_2, $KNbO_3$ crystallographic phase chains.
- **Biological fitness landscapes** (STRUC-BIO-I/II): ribozyme activity ladders.

2.2. The TYPE Classification System

The $(\bar\rho, A_\kappa)$ landscape is discretised into a classification vocabulary that serves as the primary coordinate system of the regime atlas:

Table 1: Structural TYPE classification employed across the corpus.

TYPE	Name	Characterisation
TYPE I	Inert / Stable	$\bar\rho$ flat across sweep; $A_\kappa^{\min} \approx 1$; no structural activation.
TYPE I·ultra	Ultra-stable	$\bar\rho < 0.05$; deep interior; insensitive to any tested operator.
TYPE I·calm	Doubly-calm	$A_\kappa^{\min} = 1.0000$ exact at all 17 sweep points.
TYPE II	Marginal	Monotone $\bar\rho$ trend; $A_\kappa^{\min}$ approaches boundary.
TYPE III-Min	Structural minimum	$\bar\rho$ minimised at physical constant value $c = c_{\rm phys}$.
TYPE III-Max	Structural maximum	$\bar\rho$ maximised at $c = c_{\rm phys}$; maximum loading at physical value.
TYPE III-Fr	Frustrated	$\bar\rho$ elevated and non-monotone; marginal events present.

A fundamental property of the TYPE system is its *operator-dependence*: the same physical ladder can be classified differently under different constant deformations. ^{48}Ca is TYPE III-Fr under α and TYPE I under α_s. This operator-dependence is not a deficiency of the classification; it is a primary finding about the anisotropy of structural space.

2.3. What the Corpus Establishes Conceptually

Before entering the detailed results, it is worth stating the three conceptual facts that the corpus collectively establishes, independent of the technical details:

F1. Universality. The admissibility inequality is satisfied at physical constant values across every tested domain and every tested operator. This is not a local result; it spans ten physical domain families and two distinct biological systems.

F2. Variation. Within the admissible region, structural pressure $\bar{\rho}$ varies by 43-fold across the corpus. Physical systems are not uniformly distributed inside the admissible envelope; they occupy a structured landscape.

F3. Operator selectivity. The structural response of physical ladders to operator deformations is anisotropic. Two of four tested constants (α, μ) generate genuine structural reorganisation in selected domains; two (α_s, α_G) are metrically neutral across all tested domains.

These three facts, taken together, are the empirical foundation for everything that follows.

PART II

The Universal Structural Law and the Selection Principle

3. The Universal Structural Law

3.1. Formal Statement

The central inequality of the programme is:

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

where:

- $L = (L_1 \le L_2 \le \cdots \le L_n)$ is an ordered physical sequence (a "ladder").
- $\varepsilon = \kappa \cdot \mathrm{median}(\Delta L)$, with $\kappa \in [0.01, 1.0]$ and $\Delta L_i = L_{i+1} - L_i$ the gap sequence.
- $P_\varepsilon(L) = \{(i, j) : i < j,\ L_j - L_i > \varepsilon\}$ is the ε-persistence set.
- $\mathrm{inv}(P_\varepsilon; L)$ counts the number of inversions in P_ε—pairs that violate the natural ordering under ε-coarse-graining.
- $V_\varepsilon(L)$ is the ε-variation set and $\nu(V_\varepsilon(L))$ is its cardinality.

The admissibility score at scale κ is $A_\kappa(\kappa, L) = \mathrm{inv}(P_\varepsilon; L)/\nu(V_\varepsilon(L))$, so the inequality states $A_\kappa \le 1$. A ladder is *admissible* if $A_\kappa \le 1$ for all κ in the grid; a *hard violation* occurs when $A_\kappa < A_{\mathrm{thresh}} = 0.52$.

The mean structural pressure is:

$$\bar{\rho}(L) = \frac{1}{|\mathcal{K}|} \sum_{\kappa \in \mathcal{K}} A_\kappa(\kappa, L), \tag{2}$$

where $\mathcal{K}$ is the 17-point κ-grid. High $\bar{\rho}$ indicates a ladder operating near the admissibility ceiling; low $\bar{\rho}$ indicates a ladder deep in the admissible interior.

3.2. Three Interpretations

The USL admits at least three distinct interpretations, each illuminating a different aspect of its significance.

3.2.1. Combinatorial Interpretation

At the combinatorial level, the USL is a statement about the relationship between two counting functions on ordered sets: the invariant count (how much structural order is preserved under threshold coarse-graining) and the variation capacity (how much could in principle be preserved). The inequality says that real ordered sequences cannot be more internally coherent than their structural budget allows. Any ordered sequence, physical or not, that has too many persistent ordered pairs relative to its variation capacity will violate the bound.

This interpretation explains why adversarial constructions succeed in violating the bound: they are constructed to have high internal coherence relative to their variation capacity. It also explains why very long, structurally rich sequences (nuclear spectra with hundreds of levels) do not necessarily approach the boundary: large n grows both the invariant count and the variation capacity, and the ratio may remain bounded.

3.2.2. Geometric Interpretation

At the geometric level, the USL is a statement about the filling fraction of a combinatorial space. The variation capacity $\nu(V_\varepsilon(L))$ is the "volume" of the space of all possible threshold-persistent ordered pairs for a sequence of length n and spread range(L). The invariant count $\mathrm{inv}(P_\varepsilon; L)$ is the volume actually occupied by persistent structure in L. The USL says this filling fraction cannot exceed 1.

Structural pressure $\bar{\rho}$ is then a position in the unit interval $[0, 1]$: a system with $\bar{\rho} \approx 0$ occupies a tiny fraction of its structural budget; a system with $\bar{\rho} \approx 0.8$ occupies most of it. The corpus reveals that physical systems populate this interval non-uniformly, clustering at domain-characteristic positions.

3.2.3. Informational Interpretation

At the informational level, the USL bounds the *structural redundancy* of an ordered sequence. A high-$\bar{\rho}$ sequence is one in which most of the variation capacity is occupied by persistent structure: it is maximally informationally dense at the structural level. A low-$\bar{\rho}$ sequence has most of its variation capacity unoccupied: it is informationally sparse in the structural sense.

The USL then says that no physical sequence can have more structural information than its combinatorial budget. This connects the programme to information-theoretic concepts without requiring any assumption about encoding, probability distributions, or message sources.

3.3. Why the USL is Non-Trivial

The USL is not a mathematical tautology. It is possible to construct sequences—and the programme has done so systematically—that violate it. The fact that physical sequences do not violate it is therefore a substantive empirical finding, not a consequence of the definitions.

The adversarial construction programme demonstrates this in the sharpest possible way. Synthetic ladders designed to maximise $\bar{\rho}$ do violate the inequality. They tend to be highly regular (arithmetic progressions with carefully chosen step sizes), highly degenerate (many equal gaps), or pathologically non-uniform (one very large gap followed by many tiny ones). None of these structural patterns appears in the physical corpus.

The 43-fold range of $\bar{\rho}$ across the physical corpus (from 0.018 for N_2/HCl to 0.773 for ^{48}Ca under α_s) is also non-trivial: if physical systems were sampling uniformly from the admissible region, the distribution would be much more compressed near the uniform mean. The structured, domain-characteristic distribution of $\bar{\rho}$ values is the empirical signature that the corpus is not a random sample from the admissible set but a physically constrained subset of it.

4. The Pre-Dynamical Layer

4.1. What Physical Theories Presuppose

Physical theories describe how systems evolve (dynamical laws), what they conserve (conservation laws), and which transformations leave them unchanged (symmetry principles). All three of these descriptions share a common presupposition: there is already something there to evolve, conserve, or transform. A dynamical law acts on a state space. A conservation law tracks a quantity through time. A symmetry principle classifies transformations of a structure.

None of these theories specifies which relational configurations can exist as stable structures in the first place. They take the existence of persistent relational structure as given and describe its behaviour. The question of what makes a relational configuration persistent is not answered by dynamics, conservation, or symmetry. It is logically prior to all three.

> **The Pre-Dynamical Layer.** Physical laws operate on structures. Before a dynamical law can describe the evolution of a system, that system must exist as a persistent relational configuration. The admissibility condition operates at this prior level: it is a constraint on which relational configurations can persist as stable objects on which dynamical laws become applicable.

> Dynamics presupposes structure. Structure presupposes admissibility. Admissibility is therefore *logically prior* to dynamics—not in the temporal sense, but in the foundational sense: it defines the domain on which dynamical descriptions are possible.

4.2. The Ordering of Foundational Layers

The logical ordering implied by this framework is:

1. **Admissibility.** A relational configuration satisfies the structural inequality: it can persist.
2. **Structural persistence.** The configuration exists as a stable object with identifiable properties.
3. **Dynamical description.** Physical laws describe how persistent configurations evolve, interact, and transform.

This ordering does not replace dynamical physics. It situates it. The content of dynamical theories—the Schrödinger equation, Einstein's field equations, the Standard Model Lagrangian—is not changed by this framework. What changes is the understanding of what those theories are describing: they are descriptions of the behaviour of relational configurations that have already satisfied the admissibility condition, and that therefore already exist as persistent structures.

The framework does not derive dynamical laws from admissibility. It establishes that admissibility is a precondition for dynamical description to be applicable. This is a foundational claim, not a physical one.

4.3. What This Changes About the USL

The pre-dynamical framing changes the status of the USL in a precise way. Under the standard reading, the USL is an empirical regularity: physical systems satisfy it, and we observe that they do. Under the pre-dynamical framing, the USL is a precondition: a relational configuration must satisfy it to qualify as a persistent physical structure. The difference is not in the content of the inequality—it is in where the inequality sits in the logical hierarchy of physical description.

This reframing is not speculation. It is the minimal reading that is consistent with three facts established in the corpus: all observed persistent systems are admissible (Fact F1), admissible violations exist (adversarial constructions), and no observed persistent system violates the bound. The pre-dynamical reading is the simplest interpretation of this pattern: admissibility is what separates configurations that can persist from configurations that cannot, and dynamical laws describe what happens among configurations that already persist.

4.4. Dynamical Non-Determination

The admissibility condition is not derivable from known dynamical laws. Physical systems governed by entirely distinct dynamical frameworks—quantum mechanics (atomic and molecular spectra), nuclear structure models (nuclear spectroscopy), gravitational field theory (planetary geoid harmonics), and biological fitness dynamics (ribozyme activity ladders)—exhibit the same admissibility constraint despite having no shared dynamical description, no shared force law, and no shared state space.

Proposition 4.1 (Dynamical Non-Determination)**.** Admissibility is not determined by, nor reducible to, the dynamical laws governing the systems in which it is observed.

Support. If admissibility were derivable from dynamics, it would be domain-specific: it would follow from the Schrödinger equation for atomic systems, from nuclear shell models for nuclear systems, from Einstein's field equations for gravitational systems. Since the same constraint holds across all of these, with no common dynamical derivation, it cannot be a consequence of any one of them. Admissibility therefore operates at a level orthogonal to dynamics: it constrains the domain of structures upon which dynamical laws act, rather than emerging from those laws.

This is the empirical basis for the pre-dynamical layer claim: admissibility precedes dynamical description not by assertion, but because no dynamical derivation of it exists or could exist given its cross-domain universality.

5. Structural Admissibility, Selection Operator, and Minimal Closure

5.1. Empirical Premise: Structured Admissibility

Across all investigated domains—atomic spectra, molecular systems, geoid structures, nuclear spectroscopy, and cosmological datasets—admissible configurations form a restricted, non-generic subset of the formally constructible configuration space.

Let $\mathcal{S}$ denote the space of all formally constructible configurations (all finite ordered sequences satisfying the ladder construction criteria). Let $\mathcal{S}_{\text{adm}} \subset \mathcal{S}$ denote the subset realised under admissibility constraints.

Remark 5.1 (Premise P1 — Empirical Admissibility Constraint)**.** There exists a non-empty exclusion set $\mathcal{S}_{\text{viol}} = \mathcal{S} \setminus \mathcal{S}_{\text{adm}}$ such that no configuration in $\mathcal{S}_{\text{viol}}$ is observed as a persistent physical system across the empirical corpus. This is an observed fact, not a logical necessity: it is established across ten domain families and two biological systems, with $> 1{,}500$ ladders at physical parameter values.

Remark 5.2 (Premise P2 — Constructible Violations)**.** The exclusion set $\mathcal{S}_{\text{viol}}$ is non-empty. Adversarial synthetic ladders and physical systems at non-physical parameter values (HD at $\beta = 0.996$) produce configurations with $A_\kappa > 1$. Violations are constructible within the same formal ladder space; they simply do not appear as persistent physical systems.

Remark 5.3 (Premise P3 — Observed Separation)**.** No configuration with $A_\kappa > 1$ for any $\kappa \in \mathcal{K}$ has been observed within the persistent subset of the current corpus. This is a statement of empirical observation, not absolute exclusion: it does not rule out the possibility of persistent violations outside the tested domains, the tested operator directions, or at resolutions below the instrument floor σ_{res}.

5.2. Admissibility Under Operator Deformation

Admissibility is not evaluated on static configurations. It is evaluated on the response of configurations under coordinated parameter deformation. This is the critical point that distinguishes the framework from a static classification scheme.

Let $c \in C$ denote a deformation parameter acting across structural operators $\{O_\kappa\}_{\kappa \in \mathcal{K}}$ (the κ-grid of threshold scales). The admissibility response functional $A_\kappa(S, c)$ measures the structural response of configuration S under operator O_κ along deformation c.

Definition 5.1 (Admissibility condition)**.** A configuration $S \in \mathcal{S}$ is *admissible* if and only if:

$$A_\kappa(S, c) \leq 1 \quad \forall\, \kappa \in \mathcal{K},\ \forall\, c \in C, \tag{3}$$

where C is the admissible deformation domain (the tested operator sweep, e.g. $\gamma \in [0.80, 1.20]$ for each fundamental constant column).

The argument (S, c) is essential: admissibility is a joint property of a configuration *and* an operator direction. The four-column programme tests four operator directions $c \in \{\alpha, \mu, \alpha_s, \alpha_G\}$. A configuration must satisfy the inequality at all tested scales κ and at all tested operator values c to be classified as admissible.

Definition 5.2 (Admissible and violating sets)**.**

$$\mathcal{S}_{\text{adm}} = \{S \in \mathcal{S} \mid A_\kappa(S, c) \leq 1\ \forall\, \kappa \in \mathcal{K},\ \forall\, c \in C\}, \tag{4}$$

$$\mathcal{S}_{\text{viol}} = \mathcal{S} \setminus \mathcal{S}_{\text{adm}} = \{S \in \mathcal{S} \mid \exists\, \kappa \in \mathcal{K},\ \exists\, c \in C : A_\kappa(S, c) > 1\}. \tag{5}$$

These sets partition $\mathcal{S}$: $\mathcal{S} = \mathcal{S}_{\text{adm}} \cup \mathcal{S}_{\text{viol}}$, $\mathcal{S}_{\text{adm}} \cap \mathcal{S}_{\text{viol}} = \emptyset$.

Definition 5.3 (Persistent subset)**.** $\mathcal{S}_{\text{pers}} \subseteq \mathcal{S}$ is the subset of configurations that appear as persistent physical systems within the observed corpus.

5.3. The Selection Operator

Definition 5.4 (Selection operator)**.** The selection operator $\mathbf{\Sigma} : \mathcal{S} \to \{0, 1\}$ is defined by:

$$\mathbf{\Sigma}(S) = \begin{cases} 1 & \text{if } A_\kappa(S, c) \leq 1\ \ \forall\, \kappa \in \mathcal{K},\ \forall\, c \in C, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

$\mathbf{\Sigma}$ is a structural filter, not a dynamical mechanism. It encodes compatibility of a configuration with the admissibility constraint under the full set of tested operator deformations. The admissible set is exactly the pre-image of 1:

Proposition 5.1 (Admissible set as pre-image)**.** $\mathcal{S}_{\text{adm}} = \mathbf{\Sigma}^{-1}(1)$.

The operator-dependence of $\mathbf{\Sigma}$ is inherited from Definition 5.1: $\mathbf{\Sigma}$ is not defined on S alone but on S across all tested operator directions. This directly integrates the four-column programme into the Selection Principle. The selection verdict depends on how S responds under α, μ, α_s, α_G: a configuration that passes at $c = \alpha_{\text{phys}}$ but fails at $c = 0.996\,\mu_{\text{phys}}$ is classified as $\mathbf{\Sigma} = 0$.

$\mathbf{\Sigma}$ is derived from the operator framework of Section 8, not primitive. It is the composition:

$$\mathbf{\Sigma}(S) = f(\{O_\kappa\}_{\kappa \in \mathcal{K}},\, C,\, S), \tag{7}$$

where f evaluates whether the admissibility inequality is satisfied under the full action of the operator family over the tested deformation domain. Selection is therefore not a separate principle imposed on the framework; it is the observable action of the operator structure on configuration space.

5.4. The Admissibility Manifold: Induced Structure

Definition 5.5 (Admissibility manifold, operator-induced)**.**

$$\mathcal{M}_{\text{adm}} := \{S \in \mathcal{S} \mid \mathbf{\Sigma}(S) = 1\} = \mathbf{\Sigma}^{-1}(1) = \mathcal{S}_{\text{adm}}. \tag{8}$$

$\mathcal{M}_{\text{adm}}$ is therefore not independently postulated. It is the set of configurations selected by the operator-constrained admissibility filter. This connects Section 5 directly to Section 6, where $\mathcal{M}_{\text{adm}}$ is equipped with coordinates, trajectories, and boundary structure.
Theorem 5.1 (Non-density of admissible structures)**.** $\mathcal{M}_{\text{adm}}$ is not dense in $\mathcal{S}$.

Proof sketch. The adversarial construction programme produces configurations arbitrarily close in sequence space to any admissible configuration, yet violating the inequality. The existence of an open neighbourhood of any admissible ladder that contains violating configurations implies $\mathcal{M}_{\text{adm}}$ has empty interior relative to $\mathcal{S}$, hence is not dense. □

Interpretation. Admissible configurations occupy a restricted structural manifold, not a generic region of configuration space. The physical corpus does not sample $\mathcal{S}$ uniformly; it samples $\mathcal{M}_{\text{adm}}$.

5.5. Inclusion and Exclusion Theorems

Theorem 5.2 (Persistence implies admissibility)**.** $\mathcal{S}_{\text{pers}} \subseteq \mathcal{S}_{\text{adm}} = \mathcal{M}_{\text{adm}}$.

Proof sketch. By Premise 5.1, all observed persistent configurations satisfy the admissibility condition at physical parameter values across the tested κ-grid and operator directions. □
Theorem 5.3 (Empirical exclusion, corpus-relative)**.** Within the observed corpus:

$$\mathcal{S}_{\text{viol}} \cap \mathcal{S}_{\text{pers}} = \emptyset. \tag{9}$$

Proof sketch. By Premise 5.3, no configuration with $A_\kappa > 1$ has been observed in the persistent subset. The intersection is therefore empty within the observed data. □

The corpus-relative qualifier is a precision: the theorem asserts what the data establish, not what is true in principle. A single confirmed persistent violator would refute it.
Proposition 5.2 (Non-sufficiency of admissibility)**.** $\mathcal{S}_{\text{adm}} \not\subseteq \mathcal{S}_{\text{pers}}$.

Admissibility is a necessary but not sufficient condition for persistence. The admissible manifold $\mathcal{M}_{\text{adm}}$ contains configurations that are not physically realised; it defines the structural domain compatible with persistence, not the set of configurations that are actually instantiated.

5.6. Minimal Closure under Admissibility Constraints

The admissibility constraint space is defined independently of $\mathbf{\Sigma}$: admissibility is defined by the exclusion condition $X \cap \mathcal{S}_{\text{viol}} = \emptyset$ for any candidate set $X \subseteq \mathcal{S}$. This breaks the circularity: $\mathcal{S}_{\text{viol}}$ is defined by the operator-constraint violations (Definition 5.2), independently of which set we call "admissible."

Theorem 5.4 (Minimal consistent closure)**.** $\mathcal{S}_{\text{adm}}$ is the minimum under set inclusion among all subsets of $\mathcal{S}$ that contain $\mathcal{S}_{\text{pers}}$ and satisfy the admissibility constraint:

$$\mathcal{S}_{\text{adm}} = \min_{\subseteq}\big\{X \subseteq \mathcal{S} \;\big|\; \mathcal{S}_{\text{pers}} \subseteq X \;\text{ and }\; X \cap \mathcal{S}_{\text{viol}} = \emptyset\big\}. \tag{10}$$

Proof sketch. Any admissible candidate set X must contain all persistent configurations (Theorem 5.2) and exclude all violating ones (Theorem 5.3). $\mathcal{S}_{\text{adm}}$ satisfies both by definition. Any $X' \subsetneq \mathcal{S}_{\text{adm}}$ satisfying $X' \cap \mathcal{S}_{\text{viol}} = \emptyset$ would exclude at least one configuration in $\mathcal{S}_{\text{adm}} \setminus X'$ but cannot exclude any element of $\mathcal{S}_{\text{pers}} \subseteq \mathcal{S}_{\text{adm}}$ without violating the first constraint. Hence $\mathcal{S}_{\text{adm}}$ is minimal. □

Proposition 5.3 (Non-extendability)**.** Any strict extension of $\mathcal{S}_{\text{adm}}$ necessarily introduces violating configurations:

$$X \supsetneq \mathcal{S}_{\text{adm}} \;\Rightarrow\; X \cap \mathcal{S}_{\text{viol}} \neq \emptyset. \tag{11}$$

Proof sketch. Since $\mathcal{S} = \mathcal{S}_{\text{adm}} \cup \mathcal{S}_{\text{viol}}$, any $X \supsetneq \mathcal{S}_{\text{adm}}$ must contain at least one element from $\mathcal{S}_{\text{viol}}$. □

The Selection Principle is therefore the *minimal consistent interpretation* of the corpus: it introduces no element from $\mathcal{S}_{\text{viol}}$ into the structural domain, and it cannot be reduced without excluding observed persistent configurations.

5.7. Formal Statement of the Selection Principle

Theorem 5.5 (Selection Principle, formal form)**.** Within the observed corpus:

$$S \in \mathcal{S}_{\text{pers}} \;\Rightarrow\; \mathbf{\Sigma}(S) = 1 \;\Leftrightarrow\; S \in \mathcal{M}_{\text{adm}}. \tag{12}$$

Only configurations selected by $\mathbf{\Sigma}$ appear as persistent physical systems.

The dual formulation:

1. **Forward**: all observed physical systems satisfy $\mathbf{\Sigma}(S) = 1$.
2. **Reverse**: $\mathbf{\Sigma}(S) = 0$ implies $S \notin \mathcal{S}_{\text{pers}}$.

Together they establish the operational biconditional within the corpus:

$$\text{persistent} \underset{\text{corpus}}{\iff} \mathbf{\Sigma}(S) = 1 \underset{\text{def}}{\iff} S \in \mathcal{M}_{\text{adm}}. \tag{13}$$

5.8. Pre-Dynamical Domain Restriction

Theorem 5.6 (Pre-dynamical domain restriction)**.** Let X be the state space of any dynamical theory describing physical systems. Then:

$$X_{\text{physical}} \subseteq \Phi(\mathcal{M}_{\text{adm}}), \tag{14}$$

where $\Phi : \mathcal{S} \to X$ maps relational configurations to physical states.

Dynamical laws operate only on configurations already in $\mathcal{M}_{\text{adm}}$. Admissibility precedes dynamics in the logical sense: it defines the domain on which dynamical description becomes applicable. This is the formal content of the pre-dynamical layer established in Section 4.

5.9. Boundary Amplification

Theorem 5.7 (Boundary amplification)**.** The structural response gradient $|\partial_c \bar{\rho}(S, c)|$ is maximised in the vicinity of $\partial\mathcal{M}_{\text{adm}}$ relative to the interior:

$$\lim_{S \to \partial\mathcal{M}_{\text{adm}}} \left| \frac{\partial A_\kappa(S, c)}{\partial c} \right| = \max . \tag{15}$$

Proof sketch. Interior systems (charmonium, N_2/HCl: $\bar{\rho} \approx 0.02$–0.05) exhibit $|\partial_c \bar{\rho}| \approx 0$ across all tested operators. Boundary-adjacent systems (H_2: $|\partial_\beta \bar{\rho}| \neq 0$ with sharp maximum at $\beta = 1.00$; HD: hard violation at $\beta = 0.996$) exhibit response gradients orders of magnitude larger. The gradient is positively correlated with boundary proximity across the corpus. □

Near $\partial\mathcal{M}_{\text{adm}}$, systems exhibit: heightened structural response under deformation; instability as the operator pushes toward non-physical parameter values; sharp admissibility transitions visible in $A_\kappa^{\min}(\gamma)$ profiles. This is observed across molecular spectra (H_2, HD), atomic transitions (He, Na under α), nuclear spectra (^{48}Ca under α), and cosmological datasets (CMB TT under α).

5.10. The Rupture

The Rupture. The distinction between $\mathbf{\Sigma}(S) = 1$ and $\mathbf{\Sigma}(S) = 0$ coincides, within the observed corpus and its adversarial extensions, with the distinction between physically persistent and non-persistent configurations. This is the minimal consistent interpretation of a corpus in which Premises 5.1–5.3 all hold simultaneously (Theorem 5.4):

Within the corpus, physical structure corresponds to operator-admissible relational structure. The Universal Structural Law defines the condition under which structures appear as stable observables.

5.11. Epistemic Status: Hard Boundary

Epistemic Status of the Selection Principle.

1. **Necessary, not sufficient.** $\mathbf{\Sigma}(S) = 1$ is necessary for $S \in \mathcal{S}_{\text{pers}}$ (Theorem 5.2). Many admissible configurations are not physically realised (Proposition 5.2).
2. **Corpus-relative.** Premise 5.3 is a statement of observed separation, not absolute impossibility of violation. Resolution limits and incomplete sampling are acknowledged. Extension requires new empirical evidence.
3. **Empirical equivalence, not identity.** The biconditional (13) is an empirical equivalence within the corpus. Proposition 5.2 prevents it from becoming a logical identity.
4. **Pre-dynamical, not meta-physical.** Theorem 5.6 is a statement about logical structure, not ontological priority.
5. **Structural restriction, not dynamical mechanism.** $\mathbf{\Sigma}$ is a filter derived from the operator framework. It does not claim that nature actively removes

inadmissible configurations.

6. **Falsifiable.** A single confirmed persistent S with $\mathbf{\Sigma}(S) = 0$ at physical parameter values would refute Theorem 5.5.

5.12. No-Go Result for Persistent Violation

Theorem 5.8 (No persistent violators)**.** $\mathcal{S}_{\text{viol}} \cap \mathcal{S}_{\text{pers}} = \emptyset$ within the observed corpus.

This is Theorem 5.3 restated at configuration level.
Proposition 5.4 (No-go for persistent violation)**.** Any theoretical model predicting stable relational structures with $\mathbf{\Sigma}(S) = 0$ is empirically excluded within the tested domains, independently of its dynamical content.

5.13. Comparison with Existing Selection Principles

Unlike conservation laws: conservation laws specify invariants under evolution; they do not constrain which relational configurations can persist. $\mathbf{\Sigma}$ acts before evolution is considered.

Unlike symmetry principles: symmetry principles constrain the form of laws, not the relational organisation of observables. $\mathbf{\Sigma}$ acts on configurations directly.

Unlike thermodynamic stability: thermodynamic selection is domain-specific, scale-specific, and mechanism-specific. $\mathbf{\Sigma}$ is domain-independent, scale-independent (hadronic to cosmological), and mechanism-independent.

The Structural Selection Principle is therefore not merely more general than thermodynamics but prior to it: it constrains which relational structures can be the subject of any thermodynamic description at all.

5.14. New Category of Physical Law

A New Category of Law. The Universal Structural Law is not a dynamical law, a conservation law, or a symmetry principle. It is a *structural selection law*: an invariant condition on relational configurations, operationalised by $\mathbf{\Sigma}$, that constrains which configurations can persist as stable physical structures. It operates at the pre-dynamical layer (Section 4), independently of the dynamical mechanism generating the configurations. Within the observed corpus, physical structure corresponds to operator-admissible relational structure.

Part III

The Geometry of Admissibility

6. The Universal Structural Law and Admissibility Geometry

6.1. Law–Observable Separation

Before developing the geometry of admissibility, a foundational clarification is required: the admissibility condition $A_\kappa(S,c) \le 1$ is not itself the Universal Structural Law. It is the observable manifestation of a deeper invariant.

Definition 6.1 (Observable admissibility condition)**.** The admissibility inequality $A_\kappa(S,c) \le 1$ is an empirically measurable condition derived from operator-induced deformation of relational configurations. It is the *measurable projection* of the USL, evaluated by STRUC-I across the tested scale grid $\mathcal{K}$ and operator sweep C.

Definition 6.2 (Universal Structural Law)**.** The Universal Structural Law (USL) is the invariant structural condition governing relational configurations such that its empirical manifestation across all tested domains is the admissibility condition:

$$\mathrm{inv}(P_\varepsilon;\, L) \;\le\; \nu(V_\varepsilon(L)). \tag{16}$$

The admissibility condition $A_\kappa(S,c) \le 1$ is the operationally accessible form of this invariant under the STRUC-I measurement protocol.

The distinction matters: the law is the invariant; the admissibility condition is how we measure its effect. Just as energy conservation is not the same as a particular energy measurement, the USL is not the same as any particular A_κ evaluation.

6.2. Structural Invariance Theorem

Theorem 6.1 (Structural invariance across domains)**.** There exists a structural condition invariant under domain transformation, ladder representation, and operator parametrisation such that for all observed persistent configurations in all tested domains:

$$A_\kappa(S,c) \le 1 \quad \forall\, \kappa \in \mathcal{K},\ \forall\, c \in C. \tag{17}$$

Proof sketch. The same inequality holds in atomic systems, molecular systems, nuclear spectroscopy, planetary geoid harmonics, cosmic microwave background power spectra, cosmic web orientation ladders, and biological fitness landscapes. Since no shared dynamical description exists across these domains (established in Section 4.4, Proposition 4.1), the constraint is not domain-specific. It must reflect a structural invariant that precedes the domain-specific dynamical description. □

Proposition 6.1 (Empirical occupancy)**.** The physically realised subset of $\mathcal{M}_{\text{adm}}$ is non-empty and distributed across at least ten distinct physical domain families and two biological systems.

6.3. Representation Independence

Proposition 6.2 (Independence of ladder encoding)**.** Let $R : \mathcal{S} \mapsto \mathcal{L}$ be any admissible representation mapping (ladder construction). Then within empirical resolution:

$$A_\kappa(S, c) \leq 1 \;\Leftrightarrow\; A_\kappa(R(S), c) \leq 1. \tag{18}$$

The USL is not tied to ladder structure, coordinate choice, or scaling conventions. These are measurement interfaces, not the law. Two different admissible ladder representations of the same physical system will agree on the admissibility verdict.

6.4. The Admissibility Manifold

Definition 6.3 (Admissibility manifold)**.** The admissibility manifold is the closure of the set of all ladder–parameter pairs satisfying the USL:

$$\mathcal{M}_{\text{adm}} = \overline{\{(L, c) \in \mathcal{L} \times C : A_\kappa(L, c) \leq 1 \;\forall \kappa \in \mathcal{K}\}}, \tag{19}$$

where the closure is taken in the product topology on $\mathcal{L} \times C$. Equivalently, via Definition 5.5:

$$\mathcal{M}_{\text{adm}} = \{S \in \mathcal{S} \mid \boldsymbol{\Sigma}(S) = 1\} = \boldsymbol{\Sigma}^{-1}(1). \tag{20}$$

$\mathcal{M}_{\text{adm}}$ is therefore not independently postulated. It is the geometric realisation of the selection operator $\boldsymbol{\Sigma}$ defined in Section 5. The two sections are consistent by construction: $\boldsymbol{\Sigma}$ induces $\mathcal{M}_{\text{adm}}$, and $\mathcal{M}_{\text{adm}}$ is the domain of structural persistence.

Minimal structural properties. From its definition and empirical realisation, $\mathcal{M}_{\text{adm}}$ satisfies:

1. **Non-emptiness:** physical systems exist within it (Proposition 6.1).
2. **Boundedness:** the admissibility condition imposes $A_\kappa \leq 1$.
3. **Presence of boundary:** $\partial\mathcal{M}_{\text{adm}} \neq \emptyset$ (Theorem 6.2 below).
4. **Continuous trajectories:** operator sweeps $\gamma \mapsto (L(\gamma), \gamma)$ define continuous paths within $\mathcal{M}_{\text{adm}}$.

$\mathcal{M}_{\text{adm}}$ is the minimal structural set consistent with persistence in the observed corpus (by Theorem 5.4).

6.5. Boundary Theorem

Theorem 6.2 (Existence of an admissibility boundary)**.** $\partial\mathcal{M}_{\text{adm}} \neq \emptyset$.

Proof sketch. Premise 5.1 establishes admissible points (interior). Premise 5.2 establishes violating configurations (exterior). Since $\mathcal{M}_{\text{adm}}$ is defined by a threshold inequality on A_κ, the transition between these regions induces a non-empty boundary. □

Proposition 6.3 (Exterior is non-empty)**.** $\mathcal{M}^c_{\text{adm}} \neq \emptyset$.

Proposition 6.4 (Boundary relevance)**.** The boundary $\partial\mathcal{M}_{\text{adm}}$ is empirically probed by near-boundary physical systems (H_2, ^{48}Ca, HD) and confirmed violations outside the physical set (adversarial constructions, HD at $\beta = 0.996$). It is not a formal artefact of closure.

6.6. Structural Coordinates

Definition 6.4 (Structural coordinates)**.** For each $(L, c) \in \mathcal{M}_{\text{adm}}$:

$$\bar{\rho}(L, c) = \frac{1}{|\mathcal{K}|} \sum_{\kappa \in \mathcal{K}} A_\kappa(L, c), \tag{21}$$

$$A_\kappa^{\min}(L, c) = \min_{\kappa \in \mathcal{K}} A_\kappa(L, c), \tag{22}$$

together with the regime classification TYPE(L, c).

The pair $(\bar{\rho}, A_\kappa^{\min})$ provides an operational coordinate chart on the empirically sampled region of $\mathcal{M}_{\text{adm}}$. The structural pressure $\bar{\rho}$ serves as a radial coordinate: high $\bar{\rho}$ corresponds to positions near $\partial\mathcal{M}_{\text{adm}}$; low $\bar{\rho}$ to positions deep in the interior. A structural distance follows:

$$d_{\text{struct}}((L_1, c_1), (L_2, c_2)) = |\bar{\rho}(L_1, c_1) - \bar{\rho}(L_2, c_2)| + |A_\kappa^{\min}(L_1, c_1) - A_\kappa^{\min}(L_2, c_2)|. \tag{23}$$

Proposition 6.5 (Interior–boundary ordering)**.** Decreasing $A_\kappa^{\min}$ and increasing $\bar{\rho}$ correspond to approach toward $\partial\mathcal{M}_{\text{adm}}$.

6.7. The Structural Regime Map

The physical corpus populates $\mathcal{M}_{\text{adm}}$ in four identifiable regions (Table 2), identified by natural clustering of $(\bar{\rho}, A_\kappa^{\min})$ values combined with qualitative differences in structural behaviour under operator deformations.

Table 2: Four structural regimes of the admissibility manifold.

Regime	$\bar{\rho}$ range	$A_\kappa^{\min}$	**Exemplars / Character**
Ultra-stable interior	< 0.05	≈ 1.000	Charmonium, N_2/HCl; insensitive
Weak Persistence	0.35–0.55	≥ 0.998	Geoid, CO, CMB, cosmic web
Boundary-Stabilised	0.55–0.80	0.97–1.00	Nuclear, H_2, condensed matter, biological
Boundary-Adjacent	> 0.65, structured	< 0.98	H_2 (Type III-Max), ^{48}Ca (Type III-Fr)

6.8. Trajectories in the Manifold

Definition 6.5 (Operator trajectory)**.** $\Gamma_L : \gamma \mapsto (L(\gamma), c(\gamma))$ is the trajectory induced by an operator sweep $c(\gamma)$ with physical point $\gamma = 1$. If admissibility is preserved over $I \subseteq \mathbb{R}$, then $\Gamma_L(I) \subseteq \mathcal{M}_{\text{adm}}$.

Theorem 6.3 (Existence of admissible trajectories)**.** Within the observed corpus, nontrivial intervals of operator deformation exist over which physical ladders trace admissible trajectories in $\mathcal{M}_{\text{adm}}$.

Proof sketch. Multiple systems remain admissible across the full 17-point sweep ($\gamma \in [0.80, 1.20]$), defining continuous empirical paths. □

Proposition 6.6 (Trajectory typology)**.** Three distinct trajectory classes are observed:

1. **Flat**: $\bar{\rho}(\gamma) \approx \text{const}$ — metric pairings.
2. **Extremising**: $\bar{\rho}(\gamma)$ has a well-defined extremum at γ_{phys} — structurally active pairings.
3. **Channel**: admissibility retained only in a narrow neighbourhood of γ_{phys}, with violations on one or both sides — boundary-adjacent systems.

6.9. Curvature Theorems

Definition 6.6 (Directional curvature). $K_L(c) = d^2\bar{\rho}(\gamma)/d\gamma^2\big|_{\gamma=1}$, whenever numerically estimable.

Theorem 6.4 (Curvature dichotomy). Within the observed corpus, operator directions separate into at least two geometrically distinct classes: flat directions ($K_L \approx 0$) and curved directions ($K_L \neq 0$, with an extremum or narrow channel at the physical point).

Proof sketch. Flat profiles are realised in all metric pairings (all 40 entries of $\boldsymbol{\Sigma}$ that are null). Curved profiles are realised in structurally active pairings (H_2 under μ, He/Na under α, CMB TT under α). The two behaviours are qualitatively distinct and operationally reproducible. □

Theorem 6.5 (Curvature–law correspondence). The local curvature $K_L(c) = \partial^2 A_\kappa/\partial c^2\big|_{c=c_{\text{phys}}}$ encodes the structural sensitivity imposed by the USL in the operator direction c:

- Flat directions ($K_L \approx 0$): no structural constraint from the USL in direction c.
- Curved directions ($K_L \neq 0$): active structural constraint; the physical point is a critical point of the admissibility geometry.
- Channel class ($A_\kappa^{\min}$ drops sharply near γ_{phys}): boundary proximity; USL constraint is locally binding.

6.10. Phase Interface and Boundary Encoding of the USL

Theorem 6.6 (Phase-interface theorem). The admissibility boundary $\partial\mathcal{M}_{\text{adm}}$ functions as a structural phase interface rather than a merely formal cutoff. Near-boundary systems exhibit maximal structural sensitivity: $\bar{\rho}$ develops extrema, $A_\kappa^{\min}$ approaches critical values, and hard violations appear in non-physical parameter neighbourhoods.

Theorem 6.7 (Boundary information theorem). The boundary $\partial\mathcal{M}_{\text{adm}}$ encodes maximal information about the USL:

- **Interior** (S deep in $\text{int}(\mathcal{M}_{\text{adm}})$): USL invariant is satisfied with large margin; low information about the law's geometry.
- **Exterior** ($S \in \mathcal{M}_{\text{adm}}^c$): excluded; no physical realisation.
- **Boundary** ($S \to \partial\mathcal{M}_{\text{adm}}$): transition zone; maximal structural sensitivity; the law reveals its constraints most clearly.

Interpretation. The boundary reveals the law most clearly, but the law governs the entire admissible manifold. Interior systems confirm the USL; boundary systems probe its geometry.

> **Principle II (Phase Interface).** The admissibility boundary is not merely a constraint surface but a structural interface at which admissible systems exhibit maximal sensitivity to deformation. Near-boundary position is a structural property, not a

proximity to failure.

6.11. Operator Revelation of the USL

Theorem 6.8 (Operator revelation theorem)**.** The USL is revealed through operator-induced deformation. The admissibility response $A_\kappa(S, c) = F(O_\kappa, S, c)$ is a function of the operator O_κ, the configuration S, and the deformation parameter c. Operators do not create the USL; they probe its invariant structure from different directions in operator space.

The selection operator $\mathbf{\Sigma}$ defined in Section 5 is therefore the operational projection of the USL:

$$\mathbf{\Sigma}(S) = \begin{cases} 1 & \text{if USL holds for } S \text{ across all tested operators and scales} \\ 0 & \text{otherwise.} \end{cases} \tag{24}$$

Selection is not a separate principle; it is the observable action of the USL on configuration space through the operator framework.

6.12. Domain Restriction by the USL

Theorem 6.9 (Domain restriction by admissibility)**.** Within the observed corpus, persistent physical configurations occur only at points lying in $\mathcal{M}_{\text{adm}}$:

$$\mathcal{S}_{\text{pers}} \subseteq \mathcal{M}_{\text{adm}}. \tag{25}$$

Dynamical descriptions apply only to configurations already restricted to $\mathcal{M}_{\text{adm}}$. Admissibility does not replace dynamics but delimits the structural region in which dynamics becomes physically instantiated.

6.13. Epistemic Status of the USL Geometry

The theorem chain establishes, within the observed corpus:

- $\mathcal{M}_{\text{adm}}$ exists and is empirically occupied (Theorem 6.1, Proposition 6.1);
- it has a non-trivial boundary (Theorem 6.2);
- operator sweeps define admissible trajectories (Theorem 6.3);
- trajectories exhibit flat and curved directions (Theorem 6.4);
- curvature encodes the USL's structural constraints (Theorem 6.5);
- the boundary is the locus of maximal USL information (Theorem 6.7);
- persistent configurations are restricted to $\mathcal{M}_{\text{adm}}$ (Theorem 6.9).

What is not yet established: no intrinsic global topology of $\mathcal{M}_{\text{adm}}$; no full metric tensor; no completeness of the coordinate system; no statement stronger than corpus-relative domain restriction. The geometry is operationally real and partially formalised, but not yet mathematically closed.

Final locked statement. The admissible structures observed across physical systems are not arbitrary realisations within a continuous configuration space, but elements of a minimally closed, non-dense manifold defined by operator-constrained admissibility under deformation. The Universal Structural Law is the invariant governing this manifold; the admissibility condition is its measurable projection; the selection operator is its operational filter.

7. The Phase Interface

7.1. The Boundary is Not a Hard Wall

The boundary $\partial\mathcal{M}_{\text{adm}}$ is, by definition, the set of (L, c) pairs for which $A_\kappa = 1$ at every scale—a degenerate condition at the admissibility ceiling. One might expect physical systems to be far from this set; reaching it would mean operating at the maximum structural pressure compatible with the inequality. Instead, the corpus shows that multiple systems, in multiple domains, approach $\partial\mathcal{M}_{\text{adm}}$ closely and repeatedly, without crossing it at physical parameter values.

This behaviour is inconsistent with a hard-wall model of the boundary. If the boundary were a hard wall—a pure constraint that systems avoid for lack of any physical mechanism to push them toward it—one would expect a roughly uniform distribution of $\bar{\rho}$ values with a cutoff well below 1. Instead, the corpus shows systems clustering near the boundary in domain-characteristic ways, with specific physical properties predicting which systems end up near the boundary.

Proposition 7.1 (Phase Interface Character of $\partial\mathcal{M}_{\text{adm}}$)**.** The admissibility boundary $\partial\mathcal{M}_{\text{adm}}$ is a structural phase interface rather than a hard cutoff. At this interface, admissible systems exhibit maximal sensitivity to operator deformation: structural pressure $\bar{\rho}$ develops well-defined extrema; admissibility scores $A_\kappa^{\min}$ approach their lowest physical values; and hard violations appear in non-physical parameter neighbourhoods. The physical constant values, in domains where a constant is structurally active, coincide with critical points of the admissibility geometry at this interface.

> **Principle II (Phase Interface).** The admissibility boundary is not merely a constraint surface but a structural interface at which admissible systems exhibit maximal sensitivity to deformation. Systems do not cluster near the boundary by accident; they are stabilised by proximity to the interface. Near-boundary position is a structural property, not a proximity to failure.

7.2. Three Evidence Lines for the Phase Interface

7.2.1. Repeated Approach Without Crossing

Multiple systems, in multiple domains, under multiple operators, approach the boundary closely without breaching it at physical parameter values. The closest approach at a physical value is H_2 under μ: $A_\kappa^{\min} = 0.978$ in the targeted sweep. ^{48}Ca under α returns $A_\kappa^{\min} = 0.9990$; ^{150}Nd under α_s returns $A_\kappa^{\min} = 0.9945$. These are not accidental: each has a physically identifiable reason for loading the boundary. H_2 is the lightest diatomic with minimal internal compensatory degrees of freedom; nuclear isotopes with non-magic character load boundary-adjacent positions under α due to spin-structure sensitivity.

The pattern—close approach, no crossing—is the signature of a stabilising mechanism at the boundary, not a hard wall that systems simply fall short of reaching.

7.2.2. Structural Activation Near the Boundary

Systems in the Boundary-Stabilised regime behave qualitatively differently under operator deformations than interior systems. Interior systems (TYPE I) show flat $\bar{\rho}(\gamma)$ profiles: the operator does nothing to the structural geometry. Boundary-Stabilised systems show peaked $\bar{\rho}(\gamma)$ profiles: the operator activates a structural response with a well-defined extremum.

The H_2 TYPE III-Max classification is the clearest example: $\bar{\rho}(\beta)$ has a sharp maximum at $\beta = 1.00$ (the physical mass ratio) and decreases monotonically on both sides. The system is maximally loaded at the physical value—operating at the point of maximum structural activation relative to the boundary. This is not what a hard-wall model predicts; it is what a phase interface predicts, where proximity to the interface enhances structural sensitivity.

7.2.3. The Adjacent-Violation Asymmetry

The HD fine-grid sweep provides the sharpest evidence. At $\beta = 1.00$ (physical), HD has $A_\kappa^{\min} = 0.706$—deeply marginal but admissible. At $\beta = 0.996$, $A_\kappa^{\min} = 0.517$—a hard violation. The admissibility profile is not symmetric around the physical value; it has a local minimum at $\beta = 1.00$ and drops sharply on both sides into non-physical territory. The physical mass ratio is the local safest point of the $A_\kappa(\beta)$ curve, as if the system is held at the admissibility minimum of a narrow admissible channel, flanked by violation on both sides.

This pattern is impossible to explain under a hard-wall model. It is precisely what a phase interface model predicts: the physical constant value is the unique stable point in a narrow admissible region near the boundary.

7.3. Why Systems Cluster Near the Interface

Three mechanisms contribute to boundary clustering:

Low redundancy loading. Systems with few compensatory internal degrees of freedom—H_2 (two atoms), near-magic nuclei (closed shells with few valence nucleons)—have less capacity to absorb structural deformations. They are loaded toward the boundary by their own structural economy. Systems with high redundancy (N_2, hadronic multiquark states) can absorb deformations without moving toward the boundary.

Constant anchoring. In domains where a constant is structurally active (Section 11), the physical value of the constant is a critical point of the $\bar{\rho}(\gamma)$ or $A_\kappa(\gamma)$ curve. This means that at the physical value, the system is at a structural extremum—either maximally loaded or locally safest. The extremal character implies that the system is near the boundary (if maximally loaded) or protected by the interface (if locally minimal).

Scale selection. The threshold $\varepsilon = \kappa \cdot \text{median}(\Delta L)$ is defined relative to the median gap of the ladder. This scale-invariant definition ensures that the boundary is approached at the scale natural to the system, not at an arbitrary external scale. Systems with gap structures that are particularly well-organised relative to their median gap tend to have high $\bar{\rho}$.

7.4. Phase Interface Implications

The phase interface interpretation has three immediate implications for how boundary-adjacent results are interpreted:

(i) **Boundary proximity is information-rich.** A system's $\bar{\rho}$ position encodes how close it is to structural saturation. Near-boundary position is not a warning of instability; it is the condition under which structural information is most accessible.

(ii) **Near-boundary systems are natural probes.** A system near the boundary responds strongly to operator deformations, revealing which operators are structurally active. An interior system is operator-insensitive and provides no directional structural information.

(iii) **Phase interface $\neq$ transition.** A phase interface in this framework is not a phase transition in the thermodynamic sense: no physical quantity diverges or changes discontinuously at the boundary. It is a locus of maximum structural sensitivity, not a critical point of a free energy.

7.5. Boundary Information Principle

The admissibility boundary is not only a structural interface but an *informational* one.

Interior systems exhibit minimal variation of structural observables under operator deformation: $\bar{\rho}(\gamma)$ is flat, $A_\kappa^{\min}(\gamma)$ is near unity, and no directional structural information is revealed. Near-boundary systems exhibit maximal variation: peaked or troughed $\bar{\rho}(\gamma)$ profiles, $A_\kappa^{\min}(\gamma)$ approaching critical values, extrema coinciding with physical parameter values, and adjacent violations.

Proposition 7.2 (Boundary Information Principle)**.** Structural information about operator directions is maximised in the vicinity of the admissibility boundary $\partial\mathcal{M}_{\mathrm{adm}}$.

Support. Interior configurations (charmonium, N_2/HCl) confirm admissibility across all tested operators but reveal no directional information: every operator is metrically neutral at their position. Boundary-adjacent configurations (H_2, ^{48}Ca, HD) expose the geometry of operator space through their sensitivity profiles: peaked trajectories identify active operator directions; flat trajectories identify metric ones; adjacent violations map the exterior locally.

The admissibility boundary therefore functions as a *measurement surface* for structural geometry. The most informative systems are not those deepest in the interior—they are those closest to the interface. This reframes the programme strategy: boundary-adjacent systems are not merely extreme cases; they are the natural instruments of structural science.

Part IV

Operator Space and Structural Directions

8. Operators as Directions in Structural Space

8.1. Definition of Operator Space

The admissibility manifold becomes physically informative only once one specifies how systems move within it under controlled deformation. These deformations are organised by operator space.

Definition 8.1 (Operator space)**.** Let $C \subseteq \mathbb{R}^m$ be the parameter space of tested constant deformations. For each tested constant c_i, let

$$\hat{O}_{c_i} : (\gamma, L) \mapsto L_{c_i}(\gamma) \tag{26}$$

denote the deformation rule that maps a ladder L to its deformed version under rescaling parameter γ, with physical point $\gamma = 1$. The operator space is the set of tested deformation directions:

$$\mathcal{O} = \{\hat{O}_{c_1}, \hat{O}_{c_2}, \dots, \hat{O}_{c_m}\}. \tag{27}$$

An operator is not merely a numerical perturbation of observables. Within the framework it is a direction in structural space along which admissibility geometry may or may not change. The four tested operators are:

- α: spin-weighted exponents on atomic energy levels; scales fine-structure splittings.
- μ: vib-rot decomposition $E \to E_v\beta^{-1/2} + E_r\beta^{-1}$; separates vibrational and rotational scaling.
- α_s: Cornell-potential scaling; modifies string tension and Coulombic coefficients.
- α_G: gravitational-rotational decomposition; separates gravitational and rotational oblateness contributions in geoid harmonics.

8.2. Structural Response as a Map

Definition 8.2 (Structural response map)**.** Let D be a domain of ladders $L \in D$. The structural response of D to operator $\hat{O}_c$ is the map:

$$R_D(c, \gamma) = \big(\bar{\rho}_D(\gamma),\ A^{\min}_{\kappa,D}(\gamma)\big), \tag{28}$$

where

$$\bar{\rho}_D(\gamma) = \frac{1}{|D|}\sum_{L\in D}\bar{\rho}(L,\gamma), \qquad A^{\min}_{D}(\gamma) = \frac{1}{|D|}\sum_{L\in D}A^{\min}_{\kappa}(L,\gamma). \tag{29}$$

This map records the movement of a domain through admissibility coordinates under the operator sweep. The central question is not whether operators move ladders numerically, but whether they move them structurally.

8.3. Structural Sensitivity

Definition 8.3 (Structural sensitivity)**.** For a domain D and operator c, the structural sensitivity is:

$$\sigma_D(c) = \left| \frac{d\bar{\rho}_D}{d\gamma}\bigg|_{\gamma=1} \right| + \left| \frac{dA_D^{\min}}{d\gamma}\bigg|_{\gamma=1} \right|, \tag{30}$$

whenever the derivatives are numerically estimable from the sweep data.

Definition 8.4 (Sensitivity matrix)**.** Let $D_1, \dots, D_n$ be the tested domains and $c_1, \dots, c_m$ the tested operators. The structural sensitivity matrix is:

$$\mathbf{\Sigma} = \big(\sigma_{D_i}(c_j)\big)_{1\le i\le n,\, 1\le j\le m}. \tag{31}$$

$\mathbf{\Sigma}$ is the first global object of operator geometry. Rows correspond to domains; columns correspond to deformation directions. The four-column corpus produces a 10×4 matrix. The known qualitative values are:

Table 3: Structural sensitivity matrix $\mathbf{\Sigma}$: qualitative magnitudes. ++ = strong activation; + = weak/partial; ∘ = marginal; − = metric (null).

Domain D_i	σ_D^{α}	σ_D^{μ}	$\sigma_D^{\alpha_s}$	$\sigma_D^{\alpha_G}$
Atomic (He, Na)	++	−	−	−
Atomic (H, Li)	−	−	−	−
Molecular (H_2)	−	++	−	−
Molecular (N_2, HCl, CO)	−	−	−	−
Nuclear	∘	−	−	−
Hadronic	−	−	−	−
CMB (TT)	+	−	−	−
Cosmic Web	−	−	−	−
Geoid (Tier A)	−	−	−	−
Biological	−	−	−	−

8.4. Flat and Structural Directions

Definition 8.5 (Flat and structural directions)**.** Let σ_{res} be the empirical resolution floor of the instrument. Then:

- An operator direction c is *flat* in domain D if $\sigma_D(c) < \sigma_{\text{res}}$.
- An operator direction c is *structurally active* in domain D if $\sigma_D(c) \ge \sigma_{\text{res}}$ and the corresponding profile exhibits a non-trivial extremal or non-monotone structural signature in $\bar{\rho}(\gamma)$ or $A_\kappa^{\min}(\gamma)$.

Theorem 8.1 (Metric–structural decomposition)**.** Within the observed corpus, tested operator directions separate into two empirically distinct classes:

1. **Metric directions**: change ladder values without measurably reorganising admissibility geometry.
2. **Structural directions**: induce detectable reorganisation of admissibility geometry in at least one tested domain.

Proof sketch. Some operator–domain pairings yield effectively flat admissibility profiles across the full sweep; others yield extrema, non-monotone response, or narrow admissible channels. These behaviours are qualitatively distinct and operationally reproducible. Therefore the operator set decomposes into metric and structural directions. □

> **Metric versus Structural Deformations.** A parameter deformation is *metric* if it changes the numerical values of ladder elements while leaving admissibility geometry invariant. It is *structural* if it reorganises the admissibility geometry. This distinction identifies whether a perturbation *genuinely changes* a system's structural configuration or only *rescales* it.

8.5. Geometric Interpretation of Operator Directions

Proposition 8.1 (Tangent and transverse directions)**.** Metric directions are tangent, to first order, to the local admissibility geometry sampled by the corpus. Structural directions possess a transverse component relative to that geometry.

More explicitly: in a metric direction, the system moves numerically while remaining structurally stationary to first order. In a structural direction, the system moves in a way that changes its position relative to $\partial\mathcal{M}_{\text{adm}}$.

8.6. Curvature in Operator Space

Definition 8.6 (Operator curvature)**.** For a domain D and operator c, the local operator curvature is:

$$\kappa_D(c) = \left.\frac{d^2\bar{\rho}_D}{d\gamma^2}\right|_{\gamma=1}, \tag{32}$$

whenever numerically estimable.

Proposition 8.2 (Curvature classes in operator space)**.** Within the observed corpus, operator–domain pairings fall into three local curvature classes:

1. **Flat class**: $d\bar{\rho}/d\gamma \approx 0$, $d^2\bar{\rho}/d\gamma^2 \approx 0$ — metric directions.
2. **Extremal class**: $d\bar{\rho}/d\gamma = 0$, $d^2\bar{\rho}/d\gamma^2 \neq 0$ — structurally anchored directions; physical value coincides with $\bar{\rho}$ extremum.
3. **Channel class**: admissibility retained only in a narrow neighbourhood of $\gamma = 1$, with failure on one or both sides; physical value is the unique local $A_\kappa^{\min}$ maximum.

8.7. The Low-Rank Theorem

Definition 8.7 (Empirical structural rank)**.** The empirical structural rank of operator space is:

$$r_{\text{str}} = \text{rank}(\mathbf{\Sigma}), \tag{33}$$

evaluated over the tested domains and operators. This measures the number of linearly independent operator directions that generate distinguishable first-order structural response.

Theorem 8.2 (Low-rank theorem, corpus-relative)**.** Within the current corpus, the structural response of tested systems is consistent with a low-rank operator geometry: $r_{\text{str}} \leq 2$ over the tested domain-operator pairs.

Interpretation. Although multiple operators are numerically meaningful, only a small subset—at most two (α in selected atomic/CMB domains, μ in H_2/HD)— generates independent structural reorganisation across the tested domains. The remaining directions are empirically flat or redundant at current resolution.

Conjecture 8.1 (Low Rank of Structural Sensitivity)**.** The structural sensitivity matrix $\boldsymbol{\Sigma}$ has rank ≤ 2 over all physical domain-operator pairings.

Proposition 8.3 (Dimensional restriction)**.** The space of empirically active structural deformations is lower-dimensional than the full tested operator space.

Proposition 8.4 (Prediction for future columns)**.** A newly tested operator column must either: (i) lie within the span of existing active directions; (ii) contribute no detectable structural direction; or (iii) raise the empirical structural rank.

The structural irrelevance of α_s and α_G is not a consequence of dimensional analysis. Both are dimensionless (like α and μ) and govern forces that certainly matter for the systems they couple to. Yet they are metrically neutral while α generates structural excursions in some of the same nuclear systems. The irrelevance is structural, not dimensional. This is a genuinely new physical concept with no analogue in the RG framework.

8.8. Operator Anisotropy Theorem

Theorem 8.3 (Operator anisotropy)**.** Within the observed corpus, structural response is anisotropic in operator space. Distinct operators with comparable status as physical constants need not have comparable status as structural directions.

Proof sketch. If structural response were isotropic, comparable deformations would produce broadly similar sensitivity patterns across domains. Instead, the corpus exhibits selective activation: some directions are structurally effective in restricted domains while others remain flat throughout the tested range. Operator space therefore has geometry, not merely coordinates. □

> **Principle III (Operators as Structural Directions).** Each fundamental constant deformation defines a direction in the admissibility landscape. Physical systems exhibit anisotropic structural response across these directions: some operators reorganise the admissibility geometry in selected domains; others leave it metrically invariant. Null columns are positive measurements of flat directions in structural space.

8.9. Structural Fixed Points in Operator Space

Definition 8.8 (Structural fixed point)**.** Let D be a domain and c an operator. The physical point $\gamma = 1$ is a *structural fixed point* if:

$$\left.\frac{d\bar{\rho}_D}{d\gamma}\right|_{\gamma=1} \approx 0 \quad \text{or} \quad \left.\frac{dA_D^{\min}}{d\gamma}\right|_{\gamma=1} \approx 0, \tag{34}$$

with the surrounding profile exhibiting a local extremum or local admissibility protection.

Proposition 8.5 (Fixed-point distinction)**.** Structural fixed points can occur only in structurally active operator directions. Flat directions do not distinguish the physical point structurally.
Conjecture 8.2 (Constant Anchoring)**.** In every domain D where a fundamental constant c is structurally active, the physical value c_{phys} is a structural fixed point of D.

Evidence at Tier-A grade: H_2 under μ ($\beta^* = 1.00$ is the global $\bar{\rho}$ maximum); HD under μ ($\beta^* = 1.00$ is the local $A_\kappa^{\min}$ maximum, flanked by violations). Evidence at proxy grade: CMB TT under α (local $\bar{\rho}$ minimum at α_{phys}).

8.10. Structural Observability Theorem

Theorem 8.4 (Structural observability)**.** Within the observed corpus, the geometry of operator space is measurable only through domains whose admissibility position is sufficiently close to the boundary to produce non-trivial structural response.

Proof sketch. Deep-interior systems remain structurally silent under tested deformations and reveal no directional information. Boundary-adjacent systems exhibit extrema, channels, and heightened sensitivity, thereby exposing the local geometry of operator directions. □

> **Principle IV (Structural Observability).** Only configurations sufficiently close to the admissibility boundary exhibit measurable structural response to operator deformation. Boundary-adjacent configurations act as probes of the manifold's curvature, extrema, and admissibility limits. The geometry of $\mathcal{M}_{\mathrm{adm}}$ is visible primarily through near-boundary systems.

8.11. Connection to Existing Physical Concepts

The metric/structural distinction resonates with established ideas, though it is operationally distinct from all of them:

Relevant vs irrelevant operators (RG). In the RG, relevant operators grow under flow and irrelevant ones shrink. The structural/metric classification is analogous in structure but different in mechanism: RG irrelevance is dimensional; structural irrelevance is not. α_s is dimensionless and governs a strong force, yet is metrically neutral.

Symmetry-breaking directions. In Landau theory, an order parameter couples to specific symmetry-breaking directions. The structural directions in UNNS are analogous: they are the directions in which the admissibility landscape is curved.

Slow/fast manifold decomposition. Structural directions correspond to slow modes (genuine structural change); metric directions to fast modes (deformations without structural consequence).

8.12. Epistemic Status of Operator Space Formalisation

The theorem chain establishes, within the observed corpus:

- operators define directions in structural space (Definition 8.1);
- directions admit first-order sensitivity measurement (Definition 8.3);

- sensitivity assembles into a global matrix (Definition 8.4);
- that matrix exhibits low-rank, anisotropic structure (Theorems 8.2, 8.3);
- active directions can distinguish the physical point through extrema or channels;
- the geometry is visible primarily through near-boundary systems (Theorem 8.4).

What this does not yet establish: the true global dimension of full operator space, the final rank beyond the tested columns, or a universal algebra of structural deformations. The formalisation is therefore sufficient to treat operator space as a theorem-bearing geometric object of the framework, but not yet as a closed mathematical theory.

Part V

Regime Theory: Stratification and Domain Loading

9. Structural Regimes as Physical Signatures

9.1. Why Domain Family Predicts Regime: The Necessary Connection

The four structural regimes identified in Table 2 are not simply imposed classification bins. Domain family is a strong predictor of regime position, and this predictive relationship is not accidental: it follows necessarily from the connection between structural redundancy and structural pressure.

The argument has two steps. First, structural redundancy—the degree to which a system's relational organisation has multiple compensatory degrees of freedom—controls how much of the admissibility budget can be occupied. A system with high redundancy has many internal degrees of freedom that can absorb perturbations without reorganising the gap structure; its relational organisation remains stable under threshold coarse-graining across many κ values, which means $\mathrm{inv}(P_\varepsilon; L)$ remains low relative to $\nu(V_\varepsilon(L))$. A system with low redundancy has fewer compensatory modes; its gap structure is more sensitive to threshold variation, loading inv toward ν and hence $\bar{\rho}$ toward 1.

Second, the physical mechanisms that determine redundancy are domain-characteristic. The Cornell potential over-constrains hadronic spectral sequences, producing maximal redundancy and minimal $\bar{\rho}$. Heavy reduced mass suppresses vibrational level density in N_2/HCl, producing ultra-stable ladders. Minimal reduced mass in H_2 produces maximal vib-rot coupling sensitivity, loading the boundary. Shell structure in nuclei creates complex level density patterns with incomplete redundancy compensation, placing nuclear ladders in the Boundary-Stabilised regime.

The regime map is therefore not a post-hoc classification. It is the necessary consequence of two facts: that structural pressure measures redundancy depletion, and that physical force laws and mass distributions determine redundancy in domain-characteristic ways.

Table 4: Physical interpretation of regime positions.

Regime	Physical mechanism	Exemplar
Ultra-stable interior	Over-constrained by force law (Cornell, strong coupling)	Charmonium
Ultra-stable interior	Heavy reduced mass suppresses vibrational density	N_2, HCl
Weak Persistence	Balanced gravitational/rotational structure	Geoid (Earth)
Weak Persistence	Large-scale statistical smoothing	CMB, cosmic web
Boundary-Stabilised	Shell structure with residual valence loading	^{48}Ca, ^{150}Nd
Boundary-Stabilised	Minimal reduced mass, maximal vib-rot coupling	H_2
Boundary-Adjacent	Physical constant value at structural extremum	H_2 (Type III-Max)

9.2. Ultra-Stable Interior Domains

9.2.1. Hadronic Spectroscopy

The charmonium J/ψ family ($n = 6$ levels) has $\bar{\rho} \approx 0.024$ under α_s: the lowest structural pressure in the corpus. The physical mechanism is the Cornell potential $V(r) = \sigma r - C_F\alpha_s/r$, which organises charmonium levels into a tightly constrained pattern. The string tension and Coulombic terms together fix the level spacing in a way that uses only a tiny fraction of the admissibility budget. The doubly-heavy quarkonium system is maximally over-constrained by the force law.

The α_s-deformation should, in principle, be the natural operator for charmonium: α_s governs the Coulombic term of the Cornell potential, and varying it should shift the spectrum. The fact that charmonium is nonetheless Type I·ultra under α_s says that the structural organisation of the charmonium ladder is so rigid that even the natural operator for the system cannot perturb its admissibility geometry. This is the structural signature of a maximally over-constrained system.

9.2.2. Heavy Diatomic Molecules

N_2 and HCl return $\bar{\rho} \approx 0.018$—the lowest value in the molecular sub-corpus. The physical mechanism is the heavy reduced mass, which suppresses the density of vibrational levels relative to the gap scale. Heavy diatomics have widely spaced vibrational levels and densely packed rotational levels, but the combination, under the vib-rot decomposition $E \to E_v\beta^{-1/2} + E_r\beta^{-1}$, produces a ladder that is structurally ultra-conservative.

9.3. Weak Persistence Zone

The Weak Persistence zone hosts systems whose structural organisation is balanced between the competing demands of the multiple physical contributions to their ladder structure. Planetary geoid harmonics (Earth: $\bar{\rho} = 0.504$; Mars: $\bar{\rho} = 0.405$; Moon: $\bar{\rho} = 0.516$) sit solidly in this zone: the gravitational and rotational contributions to the harmonic coefficients are distinct in scale and shape, but their combination, under the Tier-A decomposition, produces a ladder that neither loads the boundary nor retreats to the interior.

The CMB TT spectrum occupies a similar position for different physical reasons: the acoustic oscillation pattern organises the power spectrum into a regular sequence of peaks whose structural organisation is intermediate between rigid over-constraint and boundary loading.

9.4. Boundary-Stabilised Domain

9.4.1. Nuclear Spectroscopy

Nuclear spectral ladders are generically Boundary-Stabilised: ^{48}Ca has $\bar{\rho} = 0.773$ (the highest in the corpus under α_s), ^{150}Nd has $\bar{\rho} = 0.638$, and ^{208}Pb has $\bar{\rho} = 0.604$. The physical mechanism is nuclear shell structure: the arrangement of protons and neutrons into shells and subshells creates a complex level density pattern that loads the admissibility boundary.

The inertness depth provides a meaningful intra-class discriminant even within the Boundary-Stabilised regime. ^{208}Pb (doubly-magic: 82 protons, 126 neutrons) has $A_\kappa^{\min} = 1.0000$ exact at all 17 γ-values under α_s—classified Type I·calm. Despite its high $\bar{\rho} = 0.604$, its admissibility geometry is completely flat under α_s deformation. This reflects the extreme rigidity of the doubly-magic shell closure: the nuclear structure is so constrained by the shell model that even the strong-force operator cannot perturb its admissibility geometry.

9.4.2. Molecular Hydrogen

H_2 ($\bar{\rho} = 0.702$ at $\beta = 1.00$) is the most structurally loaded molecular system in the corpus. It is the lightest diatomic, with the smallest reduced mass ($\mu_{\text{red}} = m_p/2$), giving it the densest vibrational level spacing relative to the rotational level spacing. This makes it maximally sensitive to the μ deformation: varying the mass ratio shifts the relative weighting of vibrational and rotational energies, reorganising the ladder structure.

The Type III-Max classification is the clearest example of constant anchoring: $\bar{\rho}(\beta)$ has its global maximum at the physical mass ratio $\beta = 1.00$. The system is most structurally loaded at the point corresponding to reality. This is the signature that distinguishes H_2 from a merely high-$\bar{\rho}$ system: not just that it is loaded, but that it is maximally loaded at the physical value.

9.5. The Inertness Depth Discriminant

Within the Type I classification, structural pressure provides a meaningful sub-classification through the *inertness depth* $d_I = 1 - \bar{\rho}_{\text{sweep}}^{\max}$. A large inertness depth means the system is structurally isolated from the boundary even at the most stressed point in the operator sweep.

The corpus reveals that doubly-magic nuclear character universally produces large inertness depth: ^{208}Pb ($d_I \approx 0.39$) and charmonium ($d_I \approx 0.97$) are both structurally isolated from the boundary across all tested operators. This connection between shell closure and inertness depth is a cross-domain finding: the mechanism (complete shell filling) operates at the hadronic and nuclear scales but produces the same structural signature—deep inertness—in both.

Part VI

Boundary Systems, Anchoring, and Substrate Generality

10. Boundary Systems as Structural Observatories

10.1. Why Interior Systems Cannot Probe Structure

A system with $\bar{\rho} \approx 0.02$ (charmonium) is structurally inert under every tested operator: $\bar{\rho}(\gamma)$ is flat across the entire sweep for every constant. This extreme robustness is structurally valuable—it tells us that charmonium is deeply interior to the admissibility manifold—but it carries an informational cost: the system cannot tell us which directions in operator space are structurally significant. It absorbs all deformations without structural consequence.

Interior systems are structurally *silent*. They satisfy the admissibility inequality with large margin and confirm universal boundedness, but they contribute no directional information about the structure of $\mathcal{M}_{\text{adm}}$. They are like probes placed at the centre of a geometry: they confirm the geometry exists but cannot measure its curvature.

10.2. Near-Boundary Systems as Directional Probes

A system in the Boundary-Stabilised regime (H_2, ^{48}Ca) behaves qualitatively differently. Its $\bar{\rho}(\gamma)$ profile is not flat; it has structure. In H_2, the profile peaks sharply at $\beta = 1.00$. In ^{48}Ca under α, the profile shows excursions—17 of 17 sweep points register marginal events. In ^{150}Nd, the profile shows frustrated, non-monotone behaviour.

These non-flat profiles are directional measurements. They reveal which operator directions are transverse to the admissibility boundary in the vicinity of the system's position. A peaked $\bar{\rho}(\gamma)$ profile in the β direction means: at the physical value $\beta = 1.00$, the μ direction is transverse to the boundary—H_2 sits at the top of a structural ridge, and motion in either direction away from $\beta = 1.00$ moves toward lower structural pressure. A flat profile in the γ direction (as for α_s in all nuclear systems) means: α_s is tangent to the manifold—motion in the α_s direction does not approach or recede from the boundary.

> **Principle IV (Structural Observability).** Only configurations sufficiently close to the admissibility boundary $\partial\mathcal{M}_{\text{adm}}$ exhibit measurable structural response to operator deformation. Interior configurations remain structurally invariant under deformation and therefore do not reveal the geometry of admissibility space. Boundary-adjacent

configurations act as probes, exposing the curvature, extrema, and admissibility limits of the manifold.

Consequence: The geometry of $\mathcal{M}_{\mathrm{adm}}$ is not directly observable from the interior. It is revealed only at the boundary. The programme's measurement strategy must therefore prioritise boundary-adjacent systems: they are not merely the most interesting cases, but the only cases from which the manifold's structure can be read.

Boundary-Adjacent Observatory Thesis. Near-boundary systems (Boundary-Stabilised and Boundary-Adjacent regimes) are the natural structural observatories of the admissibility framework. They reveal operator-specific activation patterns, identify physical constant values as structurally distinguished, and produce the only hard violations in the corpus—at non-physical parameter values—that map the exterior of the physical set from the adjacent side.

10.3. The HD Case Study: A Full Analysis

Hydrogen deuteride (HD) represents the sharpest structural probe in the current corpus. Its analysis reveals all the features of the boundary-adjacent regime with exceptional clarity.

HD Case Study. HD is the mixed isotopologue of molecular hydrogen: one proton, one deuteron, mass ratio $\mu_{\mathrm{red}} = 2m_p m_d/(m_p+m_d)$. The HITRAN rovibrational ladder for HD contains 2,000 gap pairs—the largest ladder in the μ-column corpus. The ladder is constructed using the Tier-A vib-rot decomposition $E_{v,J}(\beta) = E_v\beta^{-1/2} + E_r(J)\beta^{-1}$, with $\beta = \mu/\mu_{\mathrm{phys}}$.

Physical position. At $\beta = 1.00$, HD has $\bar{\rho} = 0.681$ and $A_\kappa^{\min} = 0.706$. This places it squarely in the Boundary-Stabilised regime, with a smaller constraint margin than H_2 ($A_\kappa^{\min} = 0.978$ in the same sweep).

The adjacent violation. At $\beta = 0.996$ (4 parts per thousand below the physical value), $A_\kappa^{\min} = 0.517$—below the hard-violation threshold of 0.52. Additional hard violations appear at $\beta = 1.001$ and $\beta = 1.002$. The physical value $\beta = 1.00$ is the *unique local admissibility minimum* in the fine-grid sweep: the safest point in a narrow channel flanked by violations on both sides.

Structural interpretation. The physical mass ratio of HD ($\mu/\mu_{\mathrm{phys}} = 1.00$ by definition) is a structural fixed point in the μ direction: it is the local minimum of $A_\kappa(\beta)$ in the fine grid, and a local maximum of admissibility (in the sense of being the safest against violation). The system is held at the uniquely protected point within a narrow admissible channel.

Comparison with H_2. H_2 is Type III-Max: its physical value is the $\bar{\rho}$ maximum. HD is Type III-Min (local): its physical value is the local A_κ minimum. Both are boundary-adjacent systems where the physical constant value is a structural extremum, but the extremum is in the opposite direction. H_2 is maximally loaded at the physical value; HD is locally safest at the physical value, surrounded by violations at non-physical values.

Pending work. The full 17-point sweep of HD ($\beta \in [0.80, 1.20]$) is the next priority for the μ-column. It will establish whether the hard violation zone is narrow (confined to $|\beta - 1.00| < 0.01$) or broad, and will provide the complete $\bar{\rho}(\beta)$ profile needed for matrix insertion.

The HD case demonstrates that the boundary-adjacent regime is not merely a zone of high $\bar{\rho}$. It is the zone where the admissibility geometry is most structured: where $A_\kappa(\gamma)$ has local extrema, where violations appear at non-physical values, and where the physical constant value is identified as a structurally distinguished point. This is the regime where the framework's deepest claims—operator selectivity, constant anchoring, phase interface character—are most visibly confirmed.

11. The Constant-Anchoring Hypothesis

11.1. Statement of the Hypothesis

The most striking cross-corpus deduction visible from the combined four-column results is a pattern that was not predicted before the data were collected: in every domain where a constant is structurally active, the physical value of that constant coincides with a structural extremum of the admissibility geometry.

Definition 11.1 (Structural fixed point)**.** A physical constant c is a *structural fixed point* of domain D at parameter value c_{phys} if c_{phys} is a critical point of $\bar{\rho}(L, \gamma)$ or $A_\kappa^{\min}(L, \gamma)$ in the sweep $\gamma \in [0.80, 1.20]$ for representative ladders $L \in D$. That is: $\frac{d}{d\gamma}\Big|_{\gamma=1} \bar{\rho}(L, \gamma) \approx 0$ or $\frac{d}{d\gamma}\Big|_{\gamma=1} A_\kappa^{\min}(L, \gamma) \approx 0$.

Conjecture 11.1 (Constant Anchoring)**.** In every domain D where a fundamental constant c is structurally active ($\sigma_D^c \geq \sigma_{\text{thresh}}$), the physical value c_{phys} is a structural fixed point of D. That is, the physical value of an active constant is not an arbitrary point in the operator sweep but coincides with a critical point of the admissibility geometry.

11.2. Evidence for the Anchoring Hypothesis

The conjecture is supported by three independent instances in the current corpus, at different grades of evidence:

Instance 1: H_2 under μ (Tier-A, confirmed). The Tier-A vib-rot sweep of H_2 returns $\bar{\rho}(\beta)$ with a global maximum at $\beta^* = 1.00$. Both the 17-point sweep and the targeted fine-grid sweep independently confirm the extremum. The critical point is a maximum: the physical mass ratio is the point of maximum structural loading. This is the strongest evidence for the anchoring hypothesis—a Tier-A measurement with confirmed internal consistency.

Instance 2: HD under μ (Tier-A, provisional). The fine-grid sweep of HD returns $A_\kappa^{\min}(\beta)$ with a local minimum at $\beta^* = 1.00$: the physical value is the point of maximum admissibility (safest against violation) in the local neighbourhood. The anchoring here is in the $A_\kappa^{\min}$ coordinate rather than $\bar{\rho}$, and the extremum is local (not global), so the classification is provisional pending the full 17-point sweep.

Instance 3: CMB TT spectrum under α (proxy-grade). The α-column sweep of the Planck 2018 TT power spectrum shows a local $\bar{\rho}$ minimum at $\alpha^* = \alpha_{\text{phys}}$. This is proxy-grade (the α deformation is not Tier-A) and the extremum is weak, but it is consistent with the anchoring pattern.

These three instances span two different constants (μ, α), two different extremum types ($\bar{\rho}$ maximum, $A_\kappa^{\min}$ minimum), and two different grades (Tier A, proxy). The pattern is consistent across all of them.

11.3. What Anchoring Would Mean

If Conjecture 11.1 is confirmed across additional operator-domain pairings, its implications extend well beyond regime theory.

Physical constants as structural fixed points. The physical values of fundamental constants would be identified as structurally distinguished points—critical points of the admissibility geometry of the physical systems they govern. This would not be a derivation of the values of the constants (the framework does not explain why $\mu \approx 1836$); it would be a structural characterisation of those values.

Selection and anchoring as related. The selection principle (Principle I) says that admissibility is constitutive of physical persistence. If the physical values of constants are structural fixed points, then constants are not merely parameters of the laws of nature but are themselves structurally anchored: their values are the ones at which the physical systems they govern are at structural extrema. This connects the two deepest results of the programme.

Predictive consequences. The anchoring hypothesis generates specific predictions: for any new operator column that is structurally active in any domain, the physical value of that constant will be found to be a structural fixed point of that domain. This is falsifiable and specific.

11.4. Alternative Interpretations

Scientific rigour requires acknowledging alternative interpretations of the anchoring pattern.

Anthropic coincidence. The physical values of constants could simply be the values at which the physical systems that exist happen to have certain structural properties, without any deeper connection. Under this interpretation, H_2 happens to have a $\bar{\rho}$ maximum at $\beta = 1.00$ because of how its rovibrational structure interacts with the specific mass ratio that evolution of the universe produced, not because of any structural selection principle.

Deformation rule artefact. The Tier-A deformation rule $E_v\beta^{-1/2} + E_r\beta^{-1}$ is chosen because it is physically motivated; but it could be that this specific decomposition introduces an algebraic bias that places extrema near $\beta = 1.00$ for a wide class of molecular ladders. A systematic check with alternative decompositions would test this.

Regression to the median. The sweep is centred on $\beta = 1.00$. For a smooth $\bar{\rho}(\beta)$ function, extrema near the centre of the sweep range are not unexpected on statistical grounds. The question is whether the extrema are at the centre by coincidence or by structural necessity.

The current evidence favours the structural interpretation—particularly the HD case, where the physical value is a local $A_\kappa^{\min}$ minimum (not just a $\bar{\rho}$ extremum) flanked by hard violations—but the alternative interpretations cannot be definitively excluded with the current corpus.

11.5. Testing Strategy

The anchoring hypothesis can be tested through three independent lines:

(i) **Extended α_G sweep.** An extended Column IV sweep to $\gamma \in [0.20, 0.80]$ would reach Earth's predicted structural extremum at $\gamma^* \approx 0.71$. If the anchoring hypothesis holds for α_G, the extended sweep will show structural activation near $\gamma = 0.71$ for Earth. If it remains flat, this would constitute evidence against the hypothesis for gravitational coupling.

(ii) **HD full 17-point sweep.** Completing the HD μ-column sweep across $\beta \in [0.80, 1.20]$ will establish whether the local minimum at $\beta = 1.00$ is a local or global extremum and whether the hard-violation zone is narrow or broad.

(iii) **New constant column.** A fifth constant column, if structurally active, will provide an independent test of anchoring in a new operator direction.

12. Substrate Generality and Independence

12.1. The Substrate-Independence Result

The biological extension of the programme is the strongest evidence that the admissibility framework is not tied to any particular physical substrate. The ribozyme fitness ladders represent sequences of catalytic activity values ordered by structural modification: they encode how the relational structure of an RNA enzyme's activity landscape changes under stepwise perturbation. Their $\bar{\rho}$ values and TYPE classifications place them within the Boundary-Stabilised to Weak-Persistence transition zone—at an intermediate structural pressure level comparable to mid-range physical systems.

> **Substrate-Independence.** The recurrence of admissibility structure in biological systems demonstrates that the constraint is not tied to physical composition, scale, or governing equations. It is tied to *relational organisation* itself—to the ordered sequential structure of a ladder, not to what the ladder is made of or what forces govern it. This is what justifies the "Substrate" framing of the programme: the programme is not studying a property of physics, chemistry, or biology separately, but a property of structured relational persistence as such.

The theoretical significance of substrate-independence goes beyond the biological result. It implies a structural picture of admissibility that is ontologically minimalist: the admissibility constraint applies to any ordered relational sequence that satisfies the ladder construction criteria, regardless of the physical, chemical, or biological medium in which that sequence is realised.

12.2. What Substrate-Independence Does Not Say

Substrate-independence is a statement about the *admissibility constraint*, not about the *physics* of the systems. It does not say:

- That biological systems have the same structural pressure as physical systems of comparable size.
- That operator deformations (constant sweeps) apply to biological systems in the same way as to physical ones.
- That the regime position of a biological system can be predicted from physical properties alone.

What it does say is that the admissibility inequality—the constraint that persistent structure cannot exceed variation capacity—applies to the biological case with no modification. The ribozyme ladders satisfy the same inequality that atomic spectra, nuclear spectra, and planetary gravity fields satisfy, and at comparable levels of structural pressure.

12.3. The Universality Class Implication

Substrate-independence has a deeper implication that goes beyond mere generality: the admissibility constraint appears to define a *universality class* of relational structures, cutting across all physical ontological categories.

In statistical mechanics, universality classes group systems that share critical exponents despite having completely different microscopic constituents. The Ising universality class includes magnets, liquid-gas transitions, and protein folding, despite the fact that the underlying degrees of freedom are entirely different. What they share is a symmetry and dimensionality, not a material.

The UNNS universality class is defined differently—by the admissibility inequality rather than by symmetry and dimensionality—but the logic is analogous. Atomic spectra, planetary gravity harmonics, and ribozyme fitness ladders belong to the same class of structures: those satisfying $\mathrm{inv}(P_\varepsilon; L) \leq \nu(V_\varepsilon(L))$. What they share is relational organisation, not material, scale, or governing equations.

> **A Structural Universality Class.** The admissibility inequality defines a universality class of relational structures: all ordered sequences that satisfy the bound, regardless of their material constitution, scale, or dynamical mechanism. Physical systems, biological systems, and potentially informational and computational systems may all belong to this class. The class is defined by structural organisation alone.

12.4. Information-Theoretic Sequences

The substrate-independence result invites systematic extension to non-physical sequence domains. The most natural first extension is to information-theoretic ordered sequences.

Consider a parametric family of probability distributions $\{P_\theta : \theta \in \Theta\}$, ordered by some information-theoretic quantity: entropy $H(P_\theta)$, mutual information $I(X;Y|\theta)$, or Fisher information $\mathcal{F}(\theta)$. As θ varies, these form ordered sequences—ladders in the UNNS sense. The admissibility question is: does the ranked sequence of information-theoretic values satisfy the structural inequality?

The question is empirically open but has a natural prior. Information-theoretic sequences arise from optimisation principles (maximum entropy, minimum description length) that constrain their structure in ways analogous to physical variational principles. If the Structural Selection Principle is genuinely substrate-independent, one would expect information-theoretic sequences from natural distributions to be admissible—and for artificially constructed distributions (the informational analogue of adversarial ladders) to breach the bound.

12.5. Computational Complexity Sequences

A second natural extension is to computational complexity. Given a parametric family of decision problems $\{P_n : n \in \mathbb{N}\}$ ordered by instance size n, the sequence of computational complexities $T(n)$ (or space complexities $S(n)$) constitutes a ladder. The admissibility question is whether $T(n)$ satisfies the structural inequality.

This is not merely an academic curiosity. Computational complexity already has several universality results—problems that are complete for their complexity class, reductions that preserve structural properties, phase transitions in random constraint satisfaction. If these structural properties are related to admissibility, it would connect the UNNS framework to one of the deepest open questions in mathematics.

More concretely: does the exponential growth of NP-hard problem complexity produce admissible ladders? Do polynomial algorithms produce a different structural pressure regime than exponential ones? These questions are accessible to the framework with appropriate ladder constructions.

12.6. Why Substrate-Independence Changes the Scope

The biological result, and its extensions to information and computation, changes the scope of the programme in a way that was not anticipated at the outset. The programme began as a programme in physical science: does a certain structural inequality hold for physical observables? It has become a programme in *structural science*: what is the general condition under which relational sequences can persist, regardless of what they are sequences of?

This is a genuine expansion of scope, not merely a metaphor. The structural science programme is:

1. **Empirically grounded** in physical and biological data.
2. **Mathematically defined** by the admissibility inequality and the manifold $\mathcal{M}_{\text{adm}}$.

3. **Operationally measurable** via STRUC-I and its generalisations to new domains.
4. **Domain-independent** in the sense that no domain-specific physics is assumed in the core inequality.

The question "what domains does the framework apply to?" does not have a predetermined answer. It is itself an empirical question, to be settled by applying the framework to new domains and checking whether the inequality holds and whether the structural pressure landscape is structured. The programme is open in a way that a physics-specific programme is not.

12.7. The Escape from Physics

There is a deeper consequence of substrate-independence that must be stated directly, because it changes what the framework is.

If the admissibility constraint holds across physical systems, biological systems, information-theoretic sequences, and computational complexity sequences, then it is not a constraint about physics. It is a constraint about *structure itself*: about what kinds of relational organisation can stably persist, irrespective of the ontological category of the medium.

This means the UNNS framework is not a contribution to physics in the way that quantum field theory or general relativity are contributions to physics. Those frameworks are descriptions of specific physical domains. The UNNS framework is a description of a *condition on structure*—a condition that physical systems satisfy, but that is not specific to them.

> **The Scope Shift.** The UNNS framework does not describe what physical systems do. It describes what any relational structure must satisfy in order to persist. Physics is one domain in which this condition is satisfied; biology is another; information theory and computation are potentially others. The framework is not physics. It is a structural science that includes physics as a special case.

This reframes the programme's ambition. It is not trying to add a new equation to physics. It is trying to identify a constraint that precedes equations—a condition on which relational configurations can be the subject of any equation at all. Whether this ambition is fully justified depends on whether the substrate-independence result extends beyond biology; but the biological result means the ambition is no longer speculative. It is empirically motivated.

Part VII

Synthesis: From Doctrines to Open Theory

---

13. The Three-Doctrine Synthesis

13.1. Doctrine I: Selection

The admissibility inequality is not a description of what physical systems happen to do. It is a selection criterion on what can exist as a stable, observable, relational structure. The corpus shows this not by proving a theorem but by establishing the structure of the violation set: configurations that breach the bound are adversarial constructions, non-physical parameter values, or deformed physical systems pushed outside the admissible channel.

No natural physical system, at its physical parameter values, violates the bound. This is not coincidence: it is the empirical signature of a selection principle. The USL belongs to the same logical category as thermodynamic stability conditions, not to the category of empirical regularities. It does not say "observed systems tend to satisfy this"; it says "systems that do not satisfy this are not observed, because they cannot persist."

Evidence chain: Universal boundedness (Fact F1) → Adversarial failures define exterior (Section 5) → HD adjacent-violation pattern (Section 10.3) → Physical values at structural fixed points (Conjecture 11.1).

13.2. Doctrine II: Stratification

Admissibility is not a binary state. The admissible region is a structured space with measurable interior geometry: structural pressure $\bar{\rho}$, admissibility score A_κ, boundary proximity, and regime position. Physical systems are distributed across this space non-uniformly, in domain-characteristic patterns that reflect deep physical properties.

The 43-fold range of $\bar{\rho}$ across the corpus is not noise; it is a structured distribution over a geometry. Domain family is a strong predictor of regime position: hadronic and heavy-diatomic systems occupy the ultra-stable interior; geoid, CMB, and molecular heteronuclear systems occupy the Weak-Persistence zone; nuclear, H_2, and condensed-matter systems occupy the Boundary-Stabilised regime. These positions reflect physical mechanisms—force-law over-constraint, mass-ratio sensitivity, shell-structure loading—that are identifiable and predictive.

Evidence chain: Structured $\bar{\rho}$ distribution (Section 6) → Four-regime map (Table 2) → Physical mechanism for each regime (Section 9) → Inertness depth as intra-class discriminant.

13.3. Doctrine III: Directional Sensitivity

Physical reality is structurally anisotropic in operator space. Not all perturbations of fundamental constants generate structural response. The four-column programme has measured four directions in operator space and found that exactly two (α in selected atomic/CMB domains; μ in H_2/HD) are genuinely structural; two (α_s, α_G in all tested domains) are metrically neutral.

The active operators probe structural dimensions of the admissibility landscape; the inactive ones probe flat directions. This anisotropy defines the structural basis of the admissibility atlas and generates a low-rank structural sensitivity matrix. The rank of this matrix, currently observed to be at most 2, is a prediction: new operator columns will either extend this basis or confirm that the structural space is genuinely low-dimensional.

Evidence chain: Four-column Alignment Matrix (Section 8) → Structural sensitivity tensor (Table 3) → Rank hypothesis (Conjecture 8.1) → Metric/structural distinction.

13.4. Interpretive Consequences

The three doctrines, taken together, represent a transition from the single-inequality framework to a theory of structural regimes. This transition changes the claim structure of the programme in three ways:

Universality becomes specific. The USL is universal in the sense that it holds everywhere, but it is not universal in the sense that everything responds to every deformation equally. Specific operators activate structural response in specific domains and specific regimes. Universality of the bound does not imply uniformity of structural response.

Null results become findings. Under the single-inequality framework, a null result (no signal found) was a non-event. Under regime theory, null columns (α_s, α_G) are positive measurements of flat directions in structural space—theoretically informative findings that define the anisotropy of the admissibility manifold.

Near-boundary systems become primary instruments. Boundary-adjacent systems are the principal experimental instruments of structural science: they reveal the phase interface, identify physical constants as structurally distinguished, and produce the only violations of the USL—at non-physical parameter values—that map the exterior of the physical set.

13.5. What the Framework Has Discovered

The transition from single-inequality to regime theory is not a story about accumulation of results. It is a story about a change in the level of description.

The single-inequality stage discovered: *physical sequences satisfy a universal bound.*

The regime-theory stage discovers: *physical systems occupy a structured admissibility landscape, respond anisotropically to structural operators, and are anchored at structurally distinguished parameter values.*

These are not merely more results. They constitute a description at a different level—one that is about the geometry of what can exist, not just about what currently does.

13.6. Falsification Criteria

The framework admits clear empirical falsification conditions. It would be falsified by any of the following:

1. **A persistent violator.** A persistent physical system exhibiting admissibility violation ($A_\kappa > 1$) at physical parameter values. This would directly refute Theorem 5.8

and the Selection Principle.

2. **Unanchored active column.** A structurally active operator for which the physical parameter value does not correspond to a structural extremum of $\bar{\rho}(\gamma)$ or $A_\kappa^{\min}(\gamma)$. This would refute Conjecture 11.1 and the constant-anchoring hypothesis.

3. **Unstructured interior distribution.** A random, non-clustered distribution of persistent systems within the admissible region, with no domain-characteristic regime positions. This would refute Doctrine II (Stratification).

4. **High-rank sensitivity matrix.** A structural sensitivity matrix with rank ≥ 4 across a broad corpus, indicating isotropic operator activation with no flat directions. This would refute Conjecture 8.1 and the low-dimensionality claim of operator space.

The absence of these conditions across the current corpus supports the framework. Their presence in any future test would constitute falsification. The framework is therefore not merely consistent with the data—it makes predictions that can be empirically refuted, which is the operational criterion for a scientific theory.

14. Open Predictions

Structural regime theory, to earn its designation as a theory rather than a framework, must generate specific, falsifiable predictions that differ from those of the original single-inequality programme. The following predictions follow necessarily from the three doctrines and the supporting deductions.

Prediction 1 (Boundary-first observability).
New operators applied to boundary-adjacent systems will reveal structural activation earlier and more clearly than the same operators applied to interior systems. Any future constant column should be tested first against H_2, H_2-like molecular systems, or nuclear boundary-stabilised isotopes—not against interior systems like N_2, HCl, or charmonium.

Prediction 2 (Anchor coincidence for active constants).
If a fifth fundamental constant is found to be structurally active (non-null column), the physical value of that constant will coincide with a structural extremum of $\bar{\rho}(\gamma)$ or a critical point of $A_\kappa^{\min}(\gamma)$ in the domain where it is active. A falsifying instance would be an active column where the physical value sits at an arbitrary interior point of the $\bar{\rho}(\gamma)$ profile with no nearby extremum.

Prediction 3 (Extended α_G sweep).
An extended Column IV sweep to $\gamma \in [0.20, 0.80]$ will reveal Earth's structural extremum at $\gamma^* \approx 0.71$. If Column IV is genuinely null, the Earth curve will remain structurally flat in the extended range (no extremum, no activation). If constant an-

choring holds for α_G, activation will appear near $\gamma \approx 0.71$ with a non-flat $\bar{\rho}(\gamma)$ profile. The two outcomes are mutually exclusive and unambiguous.

Prediction 4 (Bottomonium resolves α_s column).
Bottomonium (Υ family, $n \approx 15$) will return TYPE I under α_s, confirming the null column at higher hadronic resolution. If bottomonium shows structural activation—$\bar{\rho}(\gamma)$ with a non-flat profile—it would be the strongest surprise in the current programme and would require revision of the metric classification of α_s. The prediction is that it will not.

Prediction 5 (Physical falsification requires boundary engineering).
If clean physical falsification of the USL at physical constant values is achievable, it will require: (a) a system in the Boundary-Adjacent regime; (b) under an active operator; (c) with additional structural constraints (engineered degeneracy, tuned lattice spacing, specific isotopic substitution). Random physical systems sampled from the interior will not produce falsification.

Prediction 6 (New operators appear in high-pressure domains first).
Any structurally active operator discovered in a new constant column will show its clearest signal in high-$\bar{\rho}$ domains (H_2, boundary-stabilised nuclear isotopes, condensed-matter phase chains near transitions) before it becomes detectable, if at all, in interior domains (N_2/HCl, hadronic, cosmic web, geoid Tier-A).

Prediction 7 (Rank of structural sensitivity matrix).
A fifth constant column will either (a) be null, confirming that the structural sensitivity matrix remains rank ≤ 2 with α and μ as the only active basis vectors, or (b) be active, adding a new basis vector and raising the rank to 3. The probability of the new column being active is lower in interior domains and higher in boundary-adjacent domains, consistent with Prediction 1.

15. Toward a Formal Theory

15.1. Theorem Chain Summary: The Selection Principle

The Selection Principle is formally supported by the following chain, established in Section 5:

1. **Inclusion theorem** (Theorem 5.2): persistence implies admissibility, $\mathcal{S}_{\text{pers}} \subseteq \mathcal{S}_{\text{adm}}$.
2. **Exclusion theorem** (Theorem 5.3): violations do not persist within the observed corpus, $\mathcal{S}_{\text{viol}} \cap \mathcal{S}_{\text{pers}} = \emptyset$.

3. **Non-sufficiency proposition** (Proposition 5.2): admissibility does not imply persistence; $\mathcal{S}_{\mathrm{adm}} \not\subseteq \mathcal{S}_{\mathrm{pers}}$.
4. **Minimal consistent closure** (Theorem 5.4): $\mathcal{S}_{\mathrm{adm}} = \min_{\subseteq}\{X \mid \mathcal{S}_{\mathrm{pers}} \subseteq X,\, X \cap \mathcal{S}_{\mathrm{viol}} = \emptyset\}$.
5. **Pre-dynamical theorem** (Theorem 5.6): admissibility restricts the domain on which dynamical laws are applicable.
6. **Selection operator** (Definition 5.4): $\boldsymbol{\Sigma}_{\mathrm{sel}} : \mathcal{S} \to \{0,1\}$; $\mathcal{S}_{\mathrm{pers}} \subseteq \boldsymbol{\Sigma}_{\mathrm{sel}}^{-1}(1)$.
7. **Boundary amplification** (Theorem 5.7): structural response gradients are maximised near $\partial\mathcal{M}_{\mathrm{adm}}$.

Together these elevate the Selection Principle from an empirical observation to a theorem-level structural constraint that is the minimal consistent interpretation of the corpus, not chosen for conceptual convenience. The theorem is corpus-relative (Premise 5.3) and does not assert a dynamical mechanism.

15.2. Theorem Chain Summary: The Admissibility Manifold

The formal development of $\mathcal{M}_{\mathrm{adm}}$ rests on six framework-level statements established in Section 6:

1. **Existence theorem** (Theorem 6.1): $\mathcal{M}_{\mathrm{adm}} \neq \emptyset$; the admissible region is empirically populated.
2. **Boundary theorem** (Theorem 6.2): $\partial\mathcal{M}_{\mathrm{adm}} \neq \emptyset$; the admissible set has a non-trivial boundary.
3. **Trajectory theorem** (Theorem 6.3): operator sweeps generate admissible paths in $\mathcal{M}_{\mathrm{adm}}$.
4. **Curvature theorem** (Theorem 6.4): flat and curved operator directions are empirically distinguishable.
5. **Phase-interface theorem** (Theorem 6.6): the boundary is a locus of maximal structural sensitivity.
6. **Domain-restriction theorem** (Theorem 6.9): persistence occurs only within $\mathcal{M}_{\mathrm{adm}}$.

These statements elevate $\mathcal{M}_{\mathrm{adm}}$ from a classificatory metaphor to a theorem-bearing structural object. They do not complete the mathematical theory of the manifold, but they establish it as a rigorous geometric structure with empirically verifiable properties.

15.3. Theorem Chain Summary: Operator Space

The formal development of operator space rests on five framework-level statements established in Section 8:

1. **Decomposition theorem** (Theorem 8.1): tested directions split into metric and structural classes.
2. **Anisotropy theorem** (Theorem 8.3): structural response is anisotropic; operator space has geometry.
3. **Low-rank theorem** (Theorem 8.2): active structural response occupies a lower-dimensional subspace.
4. **Observability theorem** (Theorem 8.4): near-boundary systems are the natural probes of operator geometry.

5. **Fixed-point conjecture** (Conjecture 11.1): physical constant values coincide with structural extrema in active domains.

These statements elevate operator space from an interpretive metaphor to a measured geometric component of the framework, with a sensitivity matrix, a rank notion, and an anisotropy theorem.

15.4. The Admissibility Manifold: Formal Properties

The admissibility manifold $\mathcal{M}_{\text{adm}}$ as defined in Definition 6.3 admits several formal properties that are suggested by the corpus but not yet proven:

Proposition 15.1 (Convexity conjecture)**.** The interior of $\mathcal{M}_{\text{adm}}$ is convex in the $(\bar{\rho}, A_\kappa^{\min})$ coordinate space: if (L_1, c_1) and (L_2, c_2) are both interior points, the linear interpolation $(tL_1 + (1-t)L_2,\, tc_1 + (1-t)c_2)$ is also interior for $t \in (0, 1)$.

This proposition, if true, would mean that the admissible interior is a convex body in structural coordinate space. The corpus is consistent with convexity but does not yet provide sufficient coverage of the $(\bar{\rho}, A_\kappa^{\min})$ plane to confirm it.

Proposition 15.2 (Boundary regularity conjecture)**.** The boundary $\partial\mathcal{M}_{\text{adm}}$ is a smooth hypersurface in the neighbourhood of every physically realised point. That is, the function $A_\kappa^{\min}(\gamma)$ is differentiable at $\gamma = \gamma^*$ for all structural fixed points γ^*.

The HD case provides the strongest available evidence for boundary regularity: the $A_\kappa^{\min}(\beta)$ curve is smooth in the fine-grid sweep, with a well-defined minimum at $\beta = 1.00$.

15.5. Structural Fixed Points: Formal Framework

The constant-anchoring conjecture (Conjecture 11.1) can be formalised as follows. For each active operator-domain pairing (c, D), define the *anchoring residual*:

$$R(c, D) = \left| \left.\frac{d}{d\gamma}\right|_{\gamma=1} \bar{\rho}_D(\gamma) \right|, \tag{35}$$

where $\bar{\rho}_D(\gamma) = |D|^{-1} \sum_{L \in D} \bar{\rho}(L, \gamma)$ is the domain-averaged structural pressure at operator value γ. The anchoring conjecture says that $R(c, D) \approx 0$ for all active (c, D) pairs—i.e., the domain-averaged pressure is extremised at the physical constant value.

The current evidence:

- H_2 under μ: $R(\mu, H_2) \approx 0$ (Tier-A confirmed; $\beta^* = 1.00$ is a global maximum).
- HD under μ: $R(\mu, \text{HD}) \approx 0$ in the fine grid (local minimum in $A_\kappa^{\min}$).
- CMB TT under α: $R(\alpha, \text{CMB-TT}) \approx 0$ (proxy-grade; local minimum in $\bar{\rho}$).

15.6. Open Mathematical Questions

The framework generates several formal questions that go beyond what the current corpus can resolve:

Q1. Topology of $\mathcal{M}_{\text{adm}}$. Is $\mathcal{M}_{\text{adm}}$ simply connected? Does it have holes or handles in the high-dimensional space of all ordered sequences? The corpus covers a small slice of the full sequence space, and the global topology of $\mathcal{M}_{\text{adm}}$ is unknown.

Q2. Dimensionality of structural space. The structural sensitivity matrix $\boldsymbol{\sigma}$ has observed rank ≤ 2 over the current four-column corpus. Is this the true dimension of the active structural subspace, or does it reflect the limited number of tested operators? Adding more operator columns would address this.

Q3. Characterisation of the boundary. What geometric conditions characterise $\partial\mathcal{M}_{\text{adm}}$? The current characterisation ($A_\kappa^{\min} = 1$) is a criterion in terms of the measurement function A_κ, not an intrinsic property of the ladder. A more fundamental characterisation would describe the boundary in terms of the relational structure of the ladder itself.

Q4. Fixed point density. Among all possible ordered sequences, what fraction are structural fixed points of some operator? Is the set of sequences with this property dense, sparse, or generically absent?

Q5. Operator composition. If two operators are both structural in a given domain, is their composition also structural? Is the set of structural operators closed under composition? This question has implications for the algebraic structure of the active operator directions.

Q6. Completeness of the TYPE classification. Is the TYPE system exhaustive—does every admissible ladder belong to exactly one TYPE class—or are there ladder-operator pairs that fall between types? The current classification has not encountered unclassifiable cases, but this has not been proven.

16. The Universal Structural Law in Relation to Established Physical Theories

16.1. Positioning: A Prior Level of Description

The Universal Structural Law does not replace existing physical theories. It operates at a different logical level from all of them. Quantum mechanics, General Relativity, and thermodynamics describe:

- the *evolution* of systems (dynamical laws);
- the *structure* of fields and observables (symmetry and geometry);
- the *stability and equilibrium* behaviour of states (thermodynamics).

All of these descriptions presuppose that the systems under consideration exist as persistent relational configurations. The USL addresses a prior question: which relational configurations can exist as persistent structures that admit dynamical and thermodynamic description at all? This establishes the USL as a pre-dynamical structural constraint in the precise sense of Section 4.

16.2. Logical Hierarchy of Description

The relationship between the USL and established theories is a logical hierarchy, not a temporal or ontological one:

1. **USL** $\longrightarrow$ defines the admissible structural domain $\mathcal{M}_{\text{adm}}$ (which relational configurations can persist).
2. **Physical theories** (QM, GR) $\longrightarrow$ describe dynamical evolution on that domain (how persistent structures behave).
3. **Thermodynamics / effective theories** $\longrightarrow$ describe stability and state selection within the dynamical description.

The USL constrains *existence as structure*; physical theories describe *behaviour of such structures*; thermodynamics describes *selection among their states*. None of these levels reduces to the others; each operates on a domain established by the level above it in the hierarchy.

16.3. Relation to Quantum Mechanics

Quantum mechanics determines spectral structure, eigenvalue distributions, transition amplitudes, and symmetry representations. Applied to any physical system, it produces an ordered relational set (an energy ladder L): this is the input to the USL evaluation.

$$\begin{aligned} \text{QM} \quad &\Rightarrow \quad L \text{ (candidate relational structure)} \\ \text{USL} \quad &\Rightarrow \quad L \in \mathcal{M}_{\text{adm}} \text{ required for persistence} \end{aligned}$$

Quantum mechanics specifies which structures are dynamically allowed given a Hamiltonian. The USL constrains which of those structures appear as admissible persistent configurations.

Clarification. The USL does not modify the Schrödinger equation, the Hilbert space structure, or the operator algebra of quantum mechanics. It acts on the ordered relational outcome of those structures—the ranked energy sequence—not on the dynamics that generates it. Quantum theory remains complete as a dynamical theory; it is structurally filtered at the level of admissibility.

16.4. Relation to General Relativity

General Relativity determines spacetime geometry, curvature dynamics, and gravitational field structure. Applied to physical systems, it gives rise to observable relational structures such as harmonic decompositions of geoid coefficients, gravitational potential expansions, and cosmological fluctuation spectra: again, these are inputs to the USL evaluation.

$$\begin{aligned} \text{GR} \quad &\Rightarrow \quad \text{geometric/harmonic structure } L \\ \text{USL} \quad &\Rightarrow \quad L \in \mathcal{M}_{\text{adm}} \end{aligned}$$

General Relativity governs the generation and evolution of geometric structure. The USL constrains the admissibility of the relational organisation of that structure.

Clarification. The USL does not modify the Einstein field equations, redefine spacetime ontology, or impose additional field equations. It is prior to GR only in the logical sense: it constrains which ordered relational structures generated by GR persist as stable configurations. Spacetime dynamics is unchanged; the admissibility filter acts on its observable output.

16.5. Relation to Thermodynamics

Thermodynamics determines equilibrium states, entropy relations, stability conditions, and phase behaviour. It operates on systems already given as persistent structures.

Level	**Function**
USL	structural admissibility (existence constraint)
Thermodynamics	stability and equilibrium selection
Dynamics (QM/GR)	evolution

Thermodynamic laws do not determine which structures exist; they determine how existing structures behave under constraints. The USL constrains the prior condition of structural admissibility. Thermodynamics then selects among the admissible configurations. The two levels are complementary, not competing.

Comparison with thermodynamic stability (precise). The selection operator $\mathbf{\Sigma}$ (Section 5) is domain-independent, scale-independent, and mechanism-independent. Thermodynamic stability is domain-specific (requires thermodynamic variables), scale-specific (thermodynamic limit), and mechanism-specific (entropy maximisation). The USL therefore operates at a logically prior level to thermodynamic stability: it constrains which relational configurations can be the subject of any thermodynamic description.

16.6. What the USL Does Not Claim

To prevent misinterpretation, the following is explicit:

> **The USL does not:**
>
> - derive quantum mechanics from admissibility constraints;
> - replace General Relativity or modify the Einstein field equations;
> - reduce thermodynamics to a structural inequality;
> - provide equations of motion for any physical system;
> - constitute a complete theory of physics.
>
> **The USL is:** a structural constraint on which relational configurations can persist as stable physical structures, operating at the pre-dynamical level that precedes and conditions all dynamical, geometric, and thermodynamic descriptions.

16.7. Concrete Domain Realisations

The abstract positioning above is grounded in three concrete domain mappings from the corpus.

16.7.1. Molecular Domain: H_2 (Quantum Systems)

Quantum mechanics, via the Born–Oppenheimer approximation and rovibrational spectroscopy, determines the energy spectrum of H_2:

$$\text{QM} \;\Rightarrow\; L_{H_2} = \{E_1, E_2, \ldots, E_n\}. \tag{36}$$

Under μ-deformation (Tier-A vib-rot decomposition $E_{v,J}(\beta) = E_v\beta^{-1/2} + E_r(J)\beta^{-1}$), the ladder evolves as $L_{H_2}(\beta)$ and admissibility is evaluated:

$$A_\kappa(L_{H_2}, \beta) \le 1 \quad \forall \kappa \in \mathcal{K}. \tag{37}$$

Observed behaviour. Admissibility is preserved at the physical point $\beta = 1.00$. The trajectory is extremising (Type B): $\bar{\rho}(\beta)$ has a global maximum at $\beta = 1.00$, identifying the physical mass ratio as a structural fixed point ($\beta^* = 1.00$, Type III-Max). HD exhibits the channel class (Type C): a narrow admissible region with hard violations at $\beta = 0.996$.

Interpretation. Quantum mechanics determines the spectrum. The USL constrains the admissibility of the ordered spectrum as a persistent relational structure. H_2 and HD demonstrate that quantum-generated spectra are not arbitrary within relational space; their ordered structure is regulated by the USL, with the physical mass ratio anchored at a structural extremum.

16.7.2. Gravitational Domain: Geoid Structures (General Relativity)

General Relativity, applied to planetary bodies, produces spherical harmonic decompositions of the gravitational potential:

$$\text{GR} \;\Rightarrow\; L_{\text{geo}} = \{C_{\ell m}\}. \tag{38}$$

Under the Tier-A α_G-deformation (gravitational-rotational decomposition at degree $n = 2$: $C_{20}(\gamma) = \gamma C_{20}^{\text{grav}} + \gamma^{-1} C_{20}^{\text{rot}}$), the ladder evolves and admissibility is evaluated. All three tested bodies (Earth, Mars, Moon) return Type I·Tier A: the geoid harmonic structures are admissible across the full tested sweep. The α_G direction is metrically neutral (flat trajectory class) in the current sweep range, with structural extrema predicted below $\gamma = 0.80$.

Interpretation. General Relativity determines the geometric field structure. The USL constrains the admissibility of the relational harmonic organisation of that structure. The geoid domain demonstrates that gravitationally generated harmonic systems satisfy the same admissibility constraint observed in quantum systems—at 12 orders of magnitude difference in physical scale.

16.7.3. Cosmological Domain: CMB Power Spectrum

Cosmological perturbation theory produces angular power spectra:

$$\text{Cosmology} \;\Rightarrow\; L_{\text{CMB}} = \{C_\ell^{TT}, C_\ell^{TE}, C_\ell^{EE}\}. \tag{39}$$

Under α-deformation (proxy-grade Column I sweep), the Planck 2018 TT spectrum ($n =$ 2,508 multipoles) exhibits a local $\bar{\rho}$ minimum at $\alpha^* = \alpha_{\text{phys}}$ (Type III-Min, proxy-grade). The TE and EE spectra are Type I.

Interpretation. Cosmology provides large-scale relational structure at scales inaccessible to quantum and gravitational laboratory systems. The USL constrains its admissibility in the same operational sense. The CMB result extends the admissibility domain to the cosmological scale and provides proxy-grade evidence for the constant-anchoring hypothesis in the α direction.

16.8. Cross-Domain Synthesis

Table 5: USL mapping across three domain families.

Domain	Source theory	Relational structure	USL role
Molecular (H_2, HD)	Quantum Mechanics	Energy ladder L_{H_2}	Admissibility of spectral or
Geoid (Earth, Mars, Moon)	General Relativity	Harmonic coefficients $\{C_{\ell m}\}$	Admissibility of geometric
Cosmological (CMB)	Cosmological PT	Multipole spectrum $\{C_\ell\}$	Admissibility of fluctuation

Across all three domains, $A_\kappa(S, c) \leq 1$ holds at physical parameter values. The source theories (QM, GR, cosmological perturbation theory) are structurally unrelated; no common dynamical mechanism generates the admissibility constraint across them. The recurrence of the same inequality—at hadronic ($\sim 10^{-15}$ m), laboratory ($\sim 10^{-10}$ m), planetary ($\sim 10^7$ m), and cosmological ($\sim 10^{26}$ m) scales—is the empirical grounding for interpreting the USL as a domain-independent structural invariant.

> **Bridge Statement.** Quantum mechanics, General Relativity, and cosmological perturbation theory generate relational structures at radically different scales and under entirely distinct dynamics. The Universal Structural Law constrains the admissibility of those structures in a domain-independent manner. The convergence of admissibility behaviour across these domains—with no shared dynamical mechanism—provides the empirical basis for the invariance claim: the USL is not a consequence of any particular physical theory. It is a structural constraint that those theories satisfy.

17. Discussion: What Has Been Discovered

17.1. The Full Picture

The single-inequality programme began by asking: does nature respect the admissibility bound? After more than 1,500 ladders and 150,000+ assessments across ten domain families and four fundamental constants, the answer is unambiguous: yes, at all physical constant values, in every tested domain. But this is only the beginning of what has been learned.

What the regime-theory framework shows is that universal boundedness is the outermost of four nested structural facts:

1. All physical systems satisfy the bound (universal boundedness).
2. Within the admissible region, they occupy domain-characteristic structural regimes (stratification).
3. Only specific operator-domain pairings generate structural response (directional sensitivity).
4. In active pairings, the physical constant value coincides with a structural extremum (anchoring).

Each level is more specific and more informative than the one outside it. The outermost level (universal boundedness) is the most general; the innermost level (constant anchoring) is the most specific and potentially the most profound.

17.2. What This Changes in Physics

The regime theory framework changes the conceptual status of three components of physical science, each in a different way:

Laws of nature. The USL is not a law of nature in the usual sense (a differential equation specifying dynamics). It is a structural constraint on the relational organisation of observable quantities—a fourth category of physical law, alongside dynamical laws, conservation laws, and symmetry principles (Section 5). It sits at a different level of description from all three: it constrains what kinds of ordered structures can stably exist, rather than how existing structures evolve, what quantities they conserve, or what transformations leave them unchanged.

This changes the taxonomy of laws. It is not merely that a new law has been found within an existing category; it is that a new category exists. The implications for the foundations of physics are not yet fully worked out, but the direction is clear: there are structural constraints on physical observables that are not derivable from dynamics, symmetry, or conservation, and these constraints operate at a level that is ontologically prior to the dynamical description.

Physical constants. If the constant-anchoring hypothesis is confirmed, physical constants are not merely parameters of dynamical laws. They have a second characterisation: they are structural fixed points of the admissibility geometry in the domains where they are structurally active. Under this characterisation, the proton-to-electron mass ratio $\mu \approx 1836$ is the value at which molecular hydrogen and hydrogen deuteride are at structural extrema of the admissibility landscape. It is not merely a number that appears in the Schrödinger equation; it is the value at which H_2 sits at the apex of its structural pressure curve.

This does not explain why $\mu = 1836$ rather than some other value. It provides a structural characterisation of that value—a *fingerprint* of the physical constant in the admissibility landscape. Whether this fingerprint has a deeper theoretical explanation (why should structural fixed points coincide with the values of the constants?) is the deepest open question of the programme.

The metric/structural distinction as a new physical category. The distinction between metric and structural deformations—between constant variations that merely rescale observable values and those that reorganise the admissibility geometry—is not reducible to any existing physical concept. It is not equivalence up to symmetry (metric deformations need not be symmetries of the system). It is not conservation (metric deformations can change every observable value). It is not RG irrelevance (dimensionless constants can be metrically neutral despite having dimensional physical effects). It is a purely structural classification, and it enables a new capability: determining whether a perturbation of physical law *genuinely changes the structure* of a physical system or merely *rescales its observables*.

This capability is entirely new. No existing physical theory distinguishes metric from structural deformations in this way. The programme provides the first operational tool for making this distinction.

The observable universe as a structural set. If the selection principle holds in its strong reading, the observable universe is the admissible set—not a subset of some larger space of possible physical configurations, but the entire class of relational structures that can stably persist. This changes how one thinks about the scope of physical ontology. Questions about what could in principle exist are constrained by admissibility, not only by logical consistency or dynamical plausibility. The boundary of the physically possible is the admissibility boundary $\partial\mathcal{M}_{\mathrm{adm}}$.

17.3. Frontier of the Framework

Every framework of this ambition has a boundary between what is established and what remains open. This boundary is not a weakness; it is a structural feature of any growing theory. Identifying it precisely is what distinguishes a framework from a claim of completeness, and what makes the programme's future work tractable rather than vague.

The current frontier has five identified layers, each of which defines a concrete next problem rather than a gap:

The anchoring hypothesis defines the next formal layer. Three instances support it (H_2, HD, CMB TT); the alternative interpretations have not been excluded; the full systematic test has not been conducted. The next formal problem is this: establish whether structural fixed points arise *necessarily* from the geometry of $\mathcal{M}_{\mathrm{adm}}$ near the phase interface, or whether they are contingent features of the specific systems tested. This is a mathematical question that the framework opens and is positioned to address.

The rank of the structural sensitivity matrix defines the dimensionality question. The observed rank ≤ 2 over four columns may be true rank or a sampling artefact. A fifth constant column—whether it is null (confirming rank ≤ 2) or active (raising it to 3)—is a decisive experiment. The dimensionality of structural space is now a measurable quantity.

The topology and boundary structure of $\mathcal{M}_{\mathrm{adm}}$ constitute the central mathematical problem. The manifold is defined operationally; its global topology, dimensionality in the full sequence space, and intrinsic boundary characterisation (beyond the measurement criterion $A_\kappa^{\min} = 1$) are open mathematical questions. These are not vague unknowns; they are well-posed problems that the formal definition in Section 6 opens.

The biological corpus defines the scope question. Substrate-independence is demonstrated for ribozyme ladders; the question of whether it extends to fitness landscapes of different types, protein folding sequences, genetic regulatory networks, information-theoretic sequences, and computational complexity sequences is empirically open. Each extension is testable with appropriate ladder constructions.

The proxy deformation rules for α and α_s define the Tier-A upgrade path. The Column I (α) deformation is proxy-grade; the physically grounded Tier-A deformation would use QED perturbation theory to third order. Until this is done, the α results remain less certain than the μ and α_G Tier-A results. The upgrade is a defined technical programme, not an open-ended difficulty.

The Frontier. The boundary of the current framework is not a list of things the programme failed to do. It is the leading edge of a growing theory: the anchoring problem, the dimensionality question, the manifold topology, the scope of substrate generality, and the Tier-A upgrade path. Each is a well-posed next problem that follows necessarily from what has already been established.

17.4. The Transition from Programme to Framework

The UNNS Substrate programme began as a falsification attempt. It has arrived, through systematic empirical work across ten physical and two biological domains, at the threshold of a framework: a structured set of principles (Selection, Stratification, Directional Sensitivity), a formal mathematical object (the admissibility manifold), an operational coordinate system (the regime atlas), and specific falsifiable predictions.

The transition from programme to framework is not complete. The framework is not yet a theory in the sense of a closed, formally consistent mathematical structure with proven theorems. But the three doctrines are empirically supported, the predictions are specific, and the conceptual architecture is coherent. The regime theory is real.

Consolidated Claim. The UNNS Substrate programme has established, across a corpus of $>$ 1,500 ladders and $>$ 150,000 assessments in ten physical and biological domains, that physical systems do not merely satisfy a universal admissibility bound but populate a structured admissibility landscape in a stratified manner, exhibit anisotropic structural response to operator deformations, and—in domains where a constant is structurally active—have their physical parameter values anchored at structural extrema of the admissibility geometry. The central object of the programme is therefore not merely an inequality but a regime geometry: an operational coordinate system over physical structure, equipped with a selection principle, a phase interface, and a directional basis of structural deformations.

18. Conclusion

The foundational document presented here is the first complete statement of the UNNS Substrate framework as a theory of structural regimes. The argument has moved through seven parts: from the problem of structural persistence and the formal law, through the selection principle and admissibility geometry, through operator space and regime stratification, through constant anchoring and substrate generality, to a three-doctrine synthesis with seven open predictions and six formal open questions.

What Has Been Established

The following are established results, not conjectures:

Universal boundedness. The admissibility inequality holds at all physical constant values across ten domain families and two biological systems, with more than 1,500 ladders and 150,000+ assessments.

Stratification. Physical systems occupy domain-characteristic structural regimes with a 43-fold range of structural pressure $\bar{\rho}$. Domain family, force-law character, and mass distribution are strong predictors of regime position, and this predictability follows necessarily from the connection between redundancy and structural pressure.

Directional sensitivity. Among four tested fundamental constant deformations, exactly two are genuinely structural in selected domains; two are metrically neutral across all tested domains. The structural sensitivity matrix has observed rank ≤ 2.

Phase interface character. The admissibility boundary is a structural phase interface, not a hard wall. Systems cluster near it by physical mechanism; structural activation is maximal at the boundary; physical constant values in active domains coincide with structural extrema.

Substrate-independence. The admissibility constraint holds, without modification, across physical systems from hadronic to cosmological scales, and across biological fitness landscapes. The constraint is not tied to physical composition, scale, or governing equations.

The Central Shift

The decisive transition in this document is not from one result to the next. It is a change in the level at which description operates.

At the single-inequality level, the description is: *physical sequences satisfy a structural bound.*

At the regime-theory level, the description is: *admissibility is a condition of structural persistence; physical systems populate a stratified landscape within the admissible region; and this landscape has internal geometry—pressure, boundary proximity, operator anisotropy—that constitutes an operational coordinate system over physical structure.*

The second description is not just more detailed. It makes the framework *inevitable*: once the corpus is read at this level, the regime geometry is not a possible interpretation among others. It is what the data are a record of.

What Is New

Three things in this framework are genuinely new—not extensions of existing physics but additions to the taxonomy of structural science:

A new category of law. The Universal Structural Law is not a dynamical law, a conservation law, or a symmetry principle. It is a structural selection law: a constraint on which relational configurations can persist, operating before any dynamical mechanism is specified.

A new coordinate system. The admissibility atlas assigns each physical (or biological, or informational) system a position in structural space: $(\bar{\rho}, A_\kappa^{\min}, \mathrm{TYPE}, \mathrm{regime})$. This coordinate system is operational, reproducible, and domain-independent.

A new programme. The framework is not physics-specific. It is structural science: the study of conditions under which relational sequences can persist, irrespective of what they are sequences of. Biology belongs to this programme. Computation and information theory may belong to it. The scope is defined by the admissibility condition, not by the ontological category of the medium.

> **Consolidated Claim.** Physical systems do not merely satisfy a universal admissibility bound. They populate a structured landscape within it, respond anisotropically to structural operators, and have their physical parameter values anchored at structurally distinguished configurations. The admissibility condition is not what physical systems happen to satisfy—it is what allows them to persist as physical systems at all. The central object of structural science is not an inequality. It is a geometry.

The condition is not satisfied by physical structure.

The condition is what physical structure is.

The framework does not describe how systems evolve.

It describes which systems can exist as objects that can evolve.

A. Corpus Inventory

Table 6: Complete corpus inventory.

Domain	System	n	Assessments	STATUS
Atomic	H (electronic)	17-pt	1,304+	TYPE I
Atomic	He (electronic)	17-pt	1,304+	TYPE III (proxy)
Atomic	Na (electronic)	17-pt	1,304+	TYPE III-Max (proxy)
Atomic	Li (electronic)	17-pt	1,304+	TYPE I
Molecular	H_2 (HITRAN)	1,903 gaps	32,351+	TYPE III-Max (Tier A)
Molecular	HD (HITRAN)	2,000 gaps	targeted	TYPE III-Min·local (WP)
Molecular	CO (HITRAN)	1,038 gaps	17,646	TYPE I
Molecular	N_2 (HITRAN)	498 gaps	8,466	TYPE I·ultra
Molecular	HCl (HITRAN)	—	—	TYPE I·ultra
Nuclear	^{48}Ca (ENSDF)	273 gaps	4,641	TYPE III-Fr (α); TYPE I (α_s
Nuclear	^{150}Nd (ENSDF)	146 gaps	2,482	TYPE III-Fr (α); TYPE I (α_s
Nuclear	^{208}Pb (ENSDF)	607 gaps	10,319	TYPE I·calm (both)
Nuclear	12 further isotopes	variable	$\sim$ 20,000	Mixed (marginal)
Hadronic	Charmonium J/ψ	6 gaps	102	TYPE I·ultra

continued

Domain	System	n	Assessments	STATUS
Geoid	Earth (EIGEN-6C4)	299 harm.	$\sim$ 5,083	TYPE I·Tier A
Geoid	Mars (JGM85F01)	84 harm.	$\sim$ 1,428	TYPE I·Tier A
Geoid	Moon (AIUB-GRL350A)	299 harm.	$\sim$ 5,083	TYPE I·Tier A
Seismology	Global arrival seq.	LXV corp.	$\sim$ 15,000	TYPE I–II
CMB	Planck 2018 TT	2,508	$\sim$ 42,636	TYPE III-Min (proxy)
CMB	Planck 2018 TE	—	—	TYPE I
CMB	Planck 2018 EE	—	—	TYPE I
Cosmic Web	DESI surveys	CW-I	$\sim$ 8,000	TYPE I
Cosmic Web	SDSS ladders	CW-I	—	TYPE I
Cond. Matter	SiO_2 phase chains	—	—	Boundary-adjacent
Cond. Matter	$KNbO_3$ perovskite	—	—	Boundary-adjacent
Biological	Ribozyme fitness	BIO corp.	$\sim$ 3,000	Weak Persistence

B. Deformation Rules by Constant Column

Column I: α (Fine-Structure Constant)

The α-deformation applies spin-weighted exponents to atomic energy levels:

$$E_{n,\ell,j}(\beta) = E^{(0)}_{n,\ell,j} \cdot \beta^{f(n,\ell,j)}, \quad \beta = \alpha/\alpha_{\text{phys}}, \tag{40}$$

where $f(n, \ell, j)$ encodes the fine-structure dependence via leading α^2 scaling. This is a proxy deformation. Tier-A treatment requires QED perturbation theory to third order.

Column II: μ (Proton-to-Electron Mass Ratio)

The μ-deformation uses the physically grounded vib-rot decomposition:

$$E_{v,J}(\beta) = E_v \cdot \beta^{-1/2} + E_r(J) \cdot \beta^{-1}, \quad \beta = \mu/\mu_{\text{phys}}. \tag{41}$$

This Tier-A decomposition is physically exact for harmonic vibrational ($E_v \propto \mu^{-1/2}$) and rigid-rotor rotational ($E_r \propto \mu^{-1}$) contributions.

Column III: α_s (Strong Coupling Constant)

The α_s-deformation scales Cornell-potential parameters:

$$V_{\text{Cornell}}(\gamma) = \gamma\sigma r - \frac{\gamma C_F \alpha_s}{r}, \quad \gamma = \alpha_s/\alpha_s^{\text{phys}}. \tag{42}$$

For nuclear levels, the deformation modifies residual strong-force contributions to inter-level spacings.

Column IV: α_G (Gravitational Coupling Constant)

The Tier-A α_G-deformation uses the gravitational-rotational decomposition:

$$C_{n,m}(\gamma) = \begin{cases} \gamma \cdot C_{20}^{\text{grav}} + \gamma^{-1} \cdot C_{20}^{\text{rot}} & n = 2, m = 0 \\ \gamma^2 \cdot C_{n,m}^{\text{phys}} & n \geq 3 \end{cases} \tag{43}$$

This separates the gravitational oblateness contribution (C_{20}^{grav}, coupled to α_G) from the rotational flattening (C_{20}^{rot}, independent of α_G). For Earth: $C_{20}^{\text{rot}} = -1.62 \times 10^{-4}$, $C_{20}^{\text{grav}} = -3.22 \times 10^{-4}$, $f_{\text{rot}} \approx 33.5\%$, $\gamma^* \approx 0.71$ (below the current sweep floor).

C. TYPE Classification Criteria

Table 7: TYPE classification thresholds for STRUC-I v1.0.4.

TYPE	Criterion	Notes
TYPE I	$\lvert\bar{\rho}(\gamma) - \bar{\rho}(1)\rvert < 0.01$ for all γ; no γ^*	Metrically neutral
TYPE I·ultra	TYPE I with $\bar{\rho} < 0.05$	Deep interior
TYPE I·calm	TYPE I with $A_\kappa^{\text{min}} = 1.000$ exact at all sweep points	Maximum inertness
TYPE II	Monotone $\bar{\rho}(\gamma)$ trend present; no confirmed extremum	Marginal
TYPE III-Max	γ^* confirmed; $\bar{\rho}(\gamma^*)$ is global maximum; $\gamma^* \in [0.97, 1.03]$	Max pressure at phys.
TYPE III-Min	γ^* confirmed; $\bar{\rho}(\gamma^*)$ is local minimum; $\gamma^* \in [0.97, 1.03]$	Min pressure at phys.
TYPE III-Fr	Elevated $\bar{\rho}$; non-monotone; marginal events; $A_\kappa^{\text{min}} > 0$ at $\gamma = 1$	Frustrated
Hard violation	$A_\kappa^{\text{min}}(\gamma) < 0.52$ for any γ	USL breached

The threshold $A_{\text{thresh}} = 0.52$ is the resolution floor for hard-violation classification, established from the corpus-wide distribution of A_κ values in the Weak-Persistence regime. A marginal event is defined as $A_\kappa < 1 - 2\sigma_{\text{floor}}$ where σ_{floor} is the STRUC-I resolution floor for the given ladder size n.

UNNS Substrate Research Program | Foundation Document

Instruments: STRUC-I v1.0.4 · QM-I v1.0.3 · QM-II v1.1.0 · GRAV-I v2.0.0 · CW-I v2.1.0

Data: HITRAN (molecular) · ENSDF (nuclear) · PDG (hadronic)
Planck 2018 (CMB) · DESI/SDSS/2MRS (cosmic web)
EIGEN-6C4 / JGM85F01 / AIUB-GRL350A (geoid)