Structural Persistence Across Four Empirical Domains

A Unified Cross-Domain Certification of the

UNNS Admissibility Geometry Framework

UNNS Research Program

March 2026 | Version 1.1

Abstract

The UNNS Substrate Program proposes that structural persistence in diverse physical systems arises from admissibility constraints governing recursive operator transformations. We present a unified empirical certification of this hypothesis across the first four independent physical domains examined in the program: seismic displacement fields (LXV chambers, three earthquake events), the cosmic microwave background radiation (CMB chambers I–III and SPECTRA-Σ, Planck 2018), planetary gravity fields (GRAV-I, three planetary bodies), and the large-scale galaxy distribution (CW-I, DESI and 2MRS catalogs).

Despite having no common physical mechanism, all four domains exhibit the same three structural invariants under admissible operator perturbation: rank rigidity of ordered structural signatures, directional rigidity of dominant orientation axes, and topological persistence of large-scale connectivity structures. Across nineteen chamber runs, the primary admissibility falsifier—the condition $\mathrm{inv}(p) > \nu(V(p))$—is never triggered. Gravity provides the most direct demonstration: three independent planetary models all exhibit exactly zero axis drift across harmonic degrees $L = 2$ through $L = 300$, a spectral gap 3.4–7.9× larger than a synthetic random control. CMB spectra in TT and TE channels are STRATIFIED across 1967–2479 operator values with no falsifier activation. Seismology achieves perfect rank invariance in Kumamoto and Ridgecrest (saturation index $\mathcal{S} = 0$) and saturates but does not breach the admissibility bound in El Mayor ($\mathcal{S} = 1$). The cosmic web shows persistent orientation structure across three independent surveys (DESI, SDSS, 2MRS), all converging to verdict STRUCTURAL BOUNDARY under survey-appropriate scale restrictions ($R \leq 80$ Mpc), with S_{struct} ranging from 0.648 (2MRS) to 0.9997 (DESI).

These results are consistent with the interpretation that physically realized systems occupy interior or boundary-adjacent regions of an underlying admissibility manifold whose geometry governs structural lawhood independently of domain-specific physical mechanisms.

Contents

1 Introduction

A persistent pattern in the UNNS Substrate Program is that structural features of physical systems remain stable under systematic operator transformations far more robustly than chance models would predict. Seismic displacement rank orders are invariant under temporal smoothing. CMB power spectra preserve their hierarchical bin ordering under multipole truncation. Planetary gravity fields keep their dominant orientation axis fixed through 300 harmonic degrees. Galaxy density tensors maintain a nearly frozen eigenvector across five coarse-graining scales spanning a factor of 16 in length. None of these observations has an obvious domain-specific explanation, yet they share a precise structural form.

The UNNS admissibility geometry framework [1, 2, 3] provides a unified language for this pattern. In this framework, an operator family $\{O_p : p \in P\}$ is *admissible* if structural instability is bounded by the vulnerability budget of the baseline gap geometry:

$$\mathrm{inv}(p) \;\le\; \nu(V(p)) \quad \forall\, p \in P, \tag{1}$$

where $\mathrm{inv}(p)$ counts rank inversions relative to baseline and $\nu(V(p))$ is the matching number of the vulnerability set induced by the perturbation envelope. The phase boundary $\mathcal{R}(p) = 1$ of the rigidity modulus $\mathcal{R}(p) = \Delta_{\min}(p_0)/(2\sigma_P(p))$ separates deep interior regimes (no inversions admissible) from degeneracy-admissible regimes (bounded inversions). A single violation (1) would refute the framework for the relevant domain.

The present paper unifies the empirical record. Section 2 states the formal framework and three structural invariants. Sections 3–6 present domain results in detail. Section 7 provides the cross-domain synthesis and phase interpretation. Section 8 identifies open questions.

Scope. We make no claim about the physical origin of the observed structures. The sole question in each domain is whether the admissibility inequality (1) is satisfied. This is a purely structural and falsifiable question. The four operator families studied are: temporal smoothing (seismology), harmonic truncation (CMB), harmonic extension (gravity), and Gaussian coarse-graining (cosmic web).

2 Theoretical Framework

2.1 Admissible operator families and structural signatures

Let P be a one-parameter metric space of operator parameters and let $O = \{O_p : p \in P\}$ be a continuous family of operators acting on observed data x. Following [2, 6], the *admissibility geometry* is determined by three quantities at each reference parameter $p_0 \in P$.

Definition 1 (Spectral gap)**.** *The ordered gap sequence is* $\Delta_k(p_0) = q_{(k)}(p_0) - q_{(k+1)}(p_0)$ *for a ranked structural signature* $(q_{(1)} \ge \cdots \ge q_{(N)})$*. The minimal gap is* $\Delta_{\min}(p_0) = \min_k \Delta_k(p_0)$*.*

Definition 2 (Perturbation envelope and vulnerability set)**.** *The perturbation envelope at parameter* p *is* $\sigma_P(p) = \max_j |q_j(p) - q_j(p_0)|$*. The vulnerability set is* $V(p) = \{k : \Delta_k(p_0) \le 2\sigma_P(p)\}$*.*

Definition 3 (Admissibility inequality)**.** *An operator family is* admissible *if condition* (1) *holds for all* $p \in P$*, where* $\nu(V(p))$ *is the cardinality of a maximum independent set in the gap-adjacency graph on* $V(p)$*.*

Definition 4 (Rigidity modulus)**.** $\mathcal{R}(p) = \min(\Delta_{\min}(p_0)/(2\sigma_P(p)),\ \Theta_{\min}/(2\delta_P))$*, where* $\Theta_{\min}$ *is the minimal directional separation margin and* δ_P *bounds angular drift. The phase boundary is* $\mathcal{R} = 1$*:* $\mathcal{R} > 1$ *places the signature in the deep interior;* $\mathcal{R} \le 1$ *opens inversion capacity bounded by* ν*.*

2.2 Three structural invariants

Across all four domains, the chambers test structural persistence along three independent diagnostic axes.

Rank rigidity. The rank ordering of structural quantities (displacement magnitudes, spectral bin means, tensor eigenvalues) is invariant under admissible perturbations. Measured by $\mathrm{inv}(p) = 0$ (deep interior) or $\mathrm{inv}(p) \leq \nu(V(p))$ (stratified boundary).

Directional rigidity. A dominant orientation axis $u(p) \in S^2$ remains stable under perturbation. Measured by total angular path S, maximum step $S_{\max}$, and directional cosine of core pairs.

Topological persistence. Large-scale connectivity or lobe structure is invariant under coarse-graining or thresholding. Measured by the adjusted Rand index and lobe-class constancy.

The rigidity modulus $\mathcal{R}$ provides a common scalar summary across domains: it quantifies how far the system is from the phase boundary, with $\mathcal{R} \gg 1$ corresponding to deep interior placement (maximal structural stability) and $\mathcal{R} \approx 1$ to boundary-adjacent placement (stability maintained but inversion capacity available).

Remark 1 (Structural origin of the admissibility inequality)**.** *The inequality* $\mathrm{inv}(p) \leq \nu(V(p))$ *is not an ad hoc constraint imposed on particular domains. It arises generically whenever three conditions hold simultaneously:*

1. ***Ordered observables.*** *The structural signature can be ranked, so that "inversion" is a well-defined count of order violations.*
2. ***Bounded perturbation envelopes.*** *The operator perturbation at parameter p changes each observable by at most* $\sigma_P(p)$*, so that only gaps narrower than* $2\sigma_P$ *are genuinely vulnerable.*
3. ***Gap protection.*** *Where the gap exceeds the perturbation envelope, no inversion is possible regardless of the perturbation direction. Where it does not, inversions are bounded by the matching number* ν *of the vulnerability graph.*

These three conditions are met by GPS displacement rank orders under temporal smoothing, CMB spectral bin means under harmonic truncation, orientation-matrix eigenvalues under harmonic extension, and density-tensor eigenvectors under Gaussian coarse-graining. The inequality is therefore a consequence of the ordered-signature structure of physical observables and the bounded-envelope structure of systematic operator families, not a claim about any specific physical mechanism.

3 Domain I: Seismology

3.1 Overview

The LXV chamber arc [6, 1] tests structural persistence of GPS-derived co-seismic displacement fields under temporal smoothing operators W_w (uniform windows of $w \in \{1, 3, 7, 14, 21\}$ days). Three earthquake events were analyzed spanning two tectonic regimes and two continents.

- **Kumamoto 2016** (Japan, $M_w = 7.0$): 4 stations (G071, J093, J465, J702); 5596 common epochs (2009–2024).

- **El Mayor–Cucapah 2010** (Baja California, $M_w = 7.2$): 5 stations (P485–P497); 6035 common epochs (2007–2024).
- **Ridgecrest 2019** (California, $M_w = 7.1$): 6 stations (CCCC, P580, P593–P595, P616); 6016 common epochs (2007–2024).

The structural signature $\Sigma(w)$ is the joint tuple consisting of (i) the rank ordering π_w of horizontal displacement magnitudes and (ii) the topological lobe class $k(w)$ of the spatial displacement pattern (LXV-D protocol [6]).

3.2 Rank invariance: Chambers LXV-A, -B2, -C2

Table 1 summarises rank stability across all three events.

Table 1: Rank invariance diagnostics for three earthquake events. Spearman ρ and Kendall τ are the average correlation with the $w = 1$ baseline across all subsequent windows. $\Delta_{\max}$ is the maximum rank shift observed. Inversion budget k is the chamber gate threshold; $\text{inv}_{\max}$ is the observed maximum.

Event	N_{st}	Spearman ρ	Kendall τ	$\Delta_{\max}$	$\text{inv}_{\max}$	k
Kumamoto 2016	4	1.000	1.000	0	0	—
El Mayor 2010	5	0.9–1.0	0.8–1.0	1	1	1
Ridgecrest 2019	6	1.000	1.000	0	0	—

Kumamoto and Ridgecrest exhibit perfect rank invariance across every smoothing window: $\text{inv}(w) = 0$ for all $w \in \{3, 7, 14, 21\}$. The rank ordering $\pi_w = \pi_1$ is preserved exactly. El Mayor shows one rank swap at $w = 3$ (stations P485/P486 exchange rank), bounded strictly within the inversion budget $k = 1$. At $w = 7, 21$ the order reverts and Spearman reaches 1.0. The primary falsifier is never triggered in any event.

3.3 Directional stability

Table 2 reports the directional cosine of the core station pair (the two highest-displacement stations).

Table 2: Core-pair directional cosine $\cos\theta$ across smoothing windows. A value near $+1$ indicates parallel displacement vectors; near -1 indicates anti-parallel (opposing sides of fault).

Event (core pair)	$w = 1$	$w = 3$	$w = 7$	$w = 14$	$w = 21$
Kumamoto (J465, G071)	0.9961	0.9962	0.9962	0.9962	0.9960
El Mayor (P497, P493)	1.0000	0.9999	0.9999	1.0000	1.0000
Ridgecrest (P595, CCCC)	−0.6766	−0.6798	−0.6790	−0.6776	−0.6745

All three events maintain constant directional structure across windows. The Ridgecrest anti-parallel cosine (≈ -0.678) is a stable structural signature of the fault geometry: the two highest-displacement stations lie on opposing sides of the rupture zone, pointing in anti-parallel directions. This is not an instability but a directional invariant of a different sign.

3.4 Topological persistence: Chamber LXV-D

The LXV-D topology probe [6] applies the full 13-station network analysis with BIC-based cluster selection. Both events return an identical classification.

Table 3: LXV-D topology classification for Kumamoto and El Mayor–Cucapah. All six gates G_{pre}, G_0–G_4 pass for both events.

Event	Stations	`topo_class`	Trusted	All gates
Kumamoto JP 2016	13	TOPO_BILOBE	✓	PASS
El Mayor–Cucapah 2010	13	TOPO_BILOBE	✓	PASS

Both events independently produce TOPO_BILOBE: a bilobed displacement structure consistent with two-sided fault rupture. Gate G_3 requires ARI mean ≥ 0.9; gate G_4 requires separation score mean ≥ 0.7. Both are satisfied. The topology result is redundant across the two events: two geographically separated earthquakes of similar magnitude on different faults return the same structural class, confirming that the bilobe signature reflects generic fault geometry rather than event-specific noise.

Remark 2. *The seismology domain occupies the deep interior regime ($\mathcal{R} \gg 1$) in two of three events. El Mayor, with its one-inversion boundary activation, represents the program's first instance of saturation ($\mathcal{S} = 1$): the inversion count reaches the budget exactly without breaching it. Per [1], this saturation event is itself structurally informative: it demonstrates that the inversion bound is achievable and thus not vacuously satisfied by a trivially small perturbation envelope.*

4 Domain II: Cosmological Radiation Fields

4.1 Overview

The CMB chamber arc [1] applies harmonic truncation operators $T_L : f(\ell) \mapsto f(\ell)\,\mathbf{1}_{\ell \leq L}$ to Planck 2018 angular power spectra $D_\ell = \ell(\ell+1)C_\ell/(2\pi)$ across three polarization channels (TT, TE, EE) from $L = 30$ to the full-resolution baseline $L_0 = 2508$ (TT) and $L_0 = 1996$ (TE, EE).

The low-ℓ geometry ($\ell = 2, 3$) is analyzed independently through the quadrupole–octopole axis framework (CMB-III chambers).

4.2 Acoustic structure: Chambers CMB-I and CMB-II

Seven acoustic peaks are identified in the TT spectrum (P_1–P_7; $\hat{\ell} \approx 225, 525, 800, 1100, 1400, 1700, 2000$). The CMB-I and CMB-II binning invariance test compares binned and full-resolution peak recovery:

Diagnostic	Value
Peak location shift $\lvert\Delta\hat{\ell}\rvert$	25
Rank inversions $\Delta\nu$	0
Δinv	0
Binning instability	**FALSE**

The acoustic peak hierarchy is invariant under the binning transformation. No peaks exchange rank.

4.3 Quadrupole–octopole geometry: Chamber CMB-III-GEO

The CMB-III-GEO chamber [5] extracts the quadrupole axis $u_2 \in S^2$ via the symmetric traceless 3×3 matrix $Q_{ij} = \sum_m a_{2m}(Y^{(2m)})_{ij}$ and the octopole axis u_3 via angular-momentum dispersion maximization.

Table 4: CMB quadrupole–octopole geometry from Chamber CMB-III-GEO. Rotation convention: passive ZYZ. HEALPix normalization (no Condon-Shortley phase). Baseline $L_0 = 2508$ (TT channel).

Quantity	Value	Units
u_2 (quadrupole axis)	$(0.5995,\ 0.8004,\ -6.5 \times 10^{-6})$	unit vector
u_3 (octopole axis)	$(0.7268,\ -0.6868,\ 0.0012)$	unit vector
$\theta_{23} = \arccos(u_2 \cdot u_3)$	$\mathbf{83.454}°$	degrees
$\lvert\theta_{23} - 90°\rvert$	6.546°	degrees
$D_{3,\max}$ (octopole gap)	9.27×10^{-9}	—
Locked-rotation test (tol. 2°)	PASS ($\delta = 1.055°$)	—
Relative-rotation test	PASS	—
Overall	**PASS**	—

The quadrupole and octopole axes are nearly orthogonal: $\theta_{23} = 83.45°$, just 6.55° from a right angle. Both rotation-invariance tests pass.

4.4 Geometric stability under perturbation: Chamber CMB-III-STAB

CMB-III-STAB applies $N = 1000$ random Wigner-D rotation perturbations (passive ZYZ convention, seed 2654435769) and tracks the response of the individual axes and their mutual angle.

Table 5: CMB-III-STAB: geometric stability statistics under 1000 random rotation perturbations.

Statistic	Value	Interpretation
Median δu_2	51.60°	individual axis wanders freely
Median δu_3	52.25°	individual axis wanders freely
D_{int} = median $\delta\theta_{23}$	**0.655°**	mutual angle locked
Min $\delta\theta_{23}$	0.059°	—
Max $\delta\theta_{23}$	1.336°	—

The key finding is the disparity between individual-axis wandering ($\sim 52°$ median) and mutual-angle stability ($D_{\text{int}} = 0.655°$, maximum 1.34°). Individual axes are not rigidly anchored in absolute sky coordinates; their relative orientation is. This is a *relational invariant* [4]: the structure preserved under perturbation is θ_{23}, not the absolute directions u_2 or u_3 individually.

4.5 Statistical certification: Chamber CMB-III-FULL

CMB-III-FULL provides statistical certification via $N = 1000$ Monte Carlo null draws and bootstrapped confidence intervals.

The observed $\phi = 0.114$ falls at the 8.5^{th} percentile of the null distribution ($z = -1.567$, Cohen $d = -1.567$, large effect, $p_{95} = [0.896, 0.932]$). The near-orthogonality of the quadrupole and octopole axes is rare under the null hypothesis of geometric randomness.

Table 6: CMB-III-FULL null-test results. The null distribution for each observable is generated by 1000 independent random Wigner-D rotations. Effect size d is Cohen's d relative to the null. "Pass" here means the null hypothesis of geometric randomness is rejected.

Observable	Observed	Null mean	Null std	p-value	Cohen d	Null
$\phi = \lvert u_2 \cdot u_3 \rvert$	0.1140	0.5597	0.2845	0.915	−1.567	Rejected
S_L (power ratio)	$\ll$ null	—	—	—	large	Rejected
R_L (power ratio)	$\ll$ null	—	—	—	large	Rejected

Invariant consistency checks all pass: rotation convention (passive ZYZ), seed consistency (GEO = STAB = 2654435769), and θ_{23} recheck (deviation = 0, tolerance = 10^{-6}).

4.6 Spectral stratification: Chamber CMB-SPECTRA-Σ

The CMB-SPECTRA-Σ chamber [1] tests the admissibility inequality (1) directly along the full harmonic truncation sweep. Table 7 summarises six runs across three polarization channels and two parameter configurations.

Table 7: CMB-SPECTRA-Σ run summary. Boundary fraction = activations per eligible L. Fails = L values where $\mathrm{inv}(L) > \nu(V(L))$.

Ch.	L_0	N_L	Fails	Bdy. activations	Verdict
TT	2508	2479	0	1946 (78.5%)	**PASS**
TT	2508	2479	0	2272 (91.6%)	**PASS**
TE	1996	1967	0	1866 (94.9%)	**PASS**
TE	1996	1967	0	1946 (99.0%)	**PASS**
EE	1996	1967	6 (first: $L = 1501$)	1931 (98.2%)	**LOCALIZED**
EE	1996	1967	0	1858 (94.5%)	**PASS**

Five of six runs pass. The one failure (EE run 0750(1)) triggers six falsifier violations at $L \geq 1501$. This localized anomaly in the E-mode polarization spectrum at high multipoles warrants further investigation (Section 8). The TT and TE channels pass without exception across both independent runs each.

The boundary activation rate is high across all channels (79–99% of eligible L values). This reflects the narrow gap structure of the CMB spectrum, particularly in TE and EE, and confirms that the SPECTRA-Σ test is a stringent one: the admissibility bound is under constant pressure from spectral boundary activations, yet is almost never breached.

Comparing with the operator-theoretic summary of [1]: the TT channel achieves saturation index $\mathcal{S}_{\max} = 0.50$, TE achieves $\mathcal{S}_{\max} = 0.86$, and EE (in the passing run) achieves $\mathcal{S}_{\max} \approx 1.0$. The EE channel is thus the closest to the structural boundary of all three channels.

5 Domain III: Planetary Gravity Fields

5.1 Overview

The GRAV-I chamber [4] tracks the dominant orientation axis $u(L) \in S^2$ of the Earth's gravitational potential under harmonic extension T_L (the operator adding all terms at degree L). The axis is extracted as the leading eigenvector of the orientation matrix $M(L) = \int_{S^2} \hat{n}\hat{n}^\top \lvert g_L(\hat{n}) \rvert \, d\hat{n}$.

Three independent real planetary gravity models were analyzed alongside a synthetic random control field (SYNTH-RANDOM-L300) with independent Gaussian Stokes coefficients.

5.2 Results

Table 8: GRAV-I axis trajectory results. All real models run to their respective $L_{\max}$ with 2-degree grid resolution ($N_\theta = 91$, $N_\phi = 180$). Final axis at $(0°, 0°)$ = geographic North Pole in body-fixed frame. "CLASS III" denotes perfectly locked axis; "CLASS 0" denotes unstable.

Model	Body	$L_{\max}$	Total path S	Max step	Transitions	Class	Final axis
EIGEN-6C4	Earth	300	**0.000°**	0.000°	0	CLASS III	$(0°, 0°)$
JGM85F01	Mars	85	**0.000°**	0.000°	0	CLASS III	$(0°, 0°)$
AIUB-GRL350A	Moon	300	**0.000°**	0.000°	0	CLASS III	$(0°, 0°)$
SYNTH-RANDOM	—	300	25.84°	18.65°	1	CLASS 0	$(44°, 216°)$

Table 9: Dominance ratio $J_1(L) = \lambda_1/(\lambda_1 + \lambda_2 + \lambda_3)$ and spectral gap $\gamma(L) = (\lambda_1 - \lambda_2)/\lambda_1$ for each model. The gap ratio gives the real-to-synthetic spectral advantage.

Model	Body	J_1 range	Median J_1	Median γ	Min γ	Gap ratio
EIGEN-6C4	Earth	$[0.9999, 1.0000]$	0.9999	3.638×10^{-3}	3.638×10^{-3}	7.9×
JGM85F01	Mars	$[0.9832, 0.9877]$	0.9832	3.340×10^{-3}	3.340×10^{-3}	7.2×
AIUB-GRL350A	Moon	$[0.6805, 0.8729]$	0.6806	1.572×10^{-3}	1.051×10^{-3}	3.4×
SYNTH-RANDOM	—	$[0.6211, 1.0000]$	0.6211	4.627×10^{-4}	7.002×10^{-5}	1.0 (ref.)

The result is unambiguous. For all three planetary bodies, the dominant axis stabilises at $L = 2$ and satisfies $\Delta u(L) = 0$ for all $L \in \{3, \ldots, L_{\max}\}$. This is not small drift: it is exact zero. Across 299 successive harmonic extensions for Earth and Moon, and 84 for Mars, the largest recorded step is identically 0°.

The axis in all cases is the geographic North Pole $(\theta, \phi) = (0°, 0°)$ in the body-fixed frame, coinciding with the rotation axis. This result is *model-independent*: three gravity field determinations from independent satellite missions and methodologies (CHAMP/GRACE/GOCE fusion for Earth; Mars Global Surveyor for Mars; SELENE for the Moon) converge identically.

The synthetic random control fails completely. Its axis undergoes an 18.65° jump at $L = 2$ and settles at $(44°, 216°)$, bearing no relation to any natural symmetry axis. The real-field spectral gap exceeds the synthetic gap by factors of 3.4–7.9×, explaining the rigidity advantage via the rigidity modulus $\mathcal{R}$: larger gaps imply larger $\mathcal{R}$ and hence deeper interior placement in the admissibility manifold.

Theorem 1 (Axis locking under harmonic extension, empirical)**.** *Let g be a planetary gravitational potential for a body with rotational oblateness. Let $u(L) \in S^2$ be the dominant orientation axis of g_L. For Earth (EIGEN-6C4, $L \leq 300$), Mars (JGM85F01, $L \leq 85$), and the Moon (AIUB-GRL350A, $L \leq 300$), the axis satisfies $\Delta u(L) = 0$ for all $L \geq 3$, so that the total path $S = 0$. A synthetic random harmonic field does not satisfy this property: $\Delta u(2) = 18.65°$ and $S = 25.84°$.*

Remark 3. *Theorem 1 is an* empirical proposition*: a precisely stated claim verified against specific models at specific resolutions, not a deductive consequence of the harmonic expansion algebra. The synthetic counterexample establishes that the result is not trivially forced by the mathematics of* *spherical harmonics. The structural reason the rotation axis becomes an attractor is derived in the following subsection directly from the admissibility geometry, without appeal to the physical rotation mechanism.*

5.3 Structural derivation of axis locking from admissibility geometry

We now show that the zero-drift result follows from a single structural argument within the admissibility framework, independent of the physical origin of the C_{20} dominance.

Setup: harmonic extension as a recursive operator. The GRAV-I chamber constructs the orientation matrix

$$M(L) = \int_{S^2} \hat{n}\hat{n}^\top |g_L(\hat{n})| \, d\Omega$$

and extracts the dominant eigenvector $u(L) = \mathrm{eigvec}(\lambda_1(L))$. Each step $T_L \to T_{L+1}$ is an operator recursion $O_L : g_L \mapsto g_{L+1} = g_L + \Delta g_L$, so the sequence $u(2), u(3), \ldots, u(L_{\max})$ is an operator trajectory in the sense of the UNNS admissibility framework.

Axis motion requires spectral gap collapse. The leading eigenvector $u(L)$ can change orientation only if the leading eigenvalue gap closes:

$$\gamma(L) = \frac{\lambda_1(L) - \lambda_2(L)}{\lambda_1(L)} \to 0.$$

This is a necessary condition for a rank-1 inversion—the only mechanism by which $u(L)$ can rotate by a non-infinitesimal amount. In the language of admissibility geometry, axis drift requires a non-empty vulnerability set $V(L) \neq \emptyset$, which requires

$$\Delta\lambda(L) = \lambda_1(L) - \lambda_2(L) \leq 2\sigma_P(L),$$

where $\sigma_P(L)$ is the perturbation envelope from the added harmonic degree $L + 1$.

Harmonic extension is a monotone positive-semidefinite perturbation. The orientation matrix satisfies

$$M(L+1) = M(L) + \Delta M(L), \qquad \Delta M(L) \succeq 0.$$

Each harmonic extension adds a positive-semidefinite contribution. The dominant mode established at low degree is therefore never subtracted; it can only be reinforced or left unchanged. This monotonicity is the structural core of the locking mechanism.

Quantitative argument from the measured gap. From Table 9, the real planetary models sustain spectral gaps

$$\gamma_{\text{Earth}} \approx 3.6 \times 10^{-3}, \quad \gamma_{\text{Moon}} \approx 1.6 \times 10^{-3}, \quad \gamma_{\text{Mars}} \approx 3.3 \times 10^{-3},$$

while the synthetic control achieves only $\gamma_{\text{synth}} \approx 4.6 \times 10^{-4}$. The real-to-synthetic gap ratio is 3.4–7.9×.

For each successive harmonic degree ℓ, the contribution to the orientation matrix scales roughly as $\ell^{-\alpha}$ for $\alpha > 0$ (decaying power spectrum). The perturbation envelope is therefore

$$\sigma_P(L) \sim C_\ell^{1/2} \ll \Delta\lambda(L),$$

ensuring that for all $L \geq 3$ in the real models

$$V(L) = \emptyset.$$

Conclusion: the rigidity modulus is everywhere large. Substituting into the rigidity modulus

$$\mathcal{R}(L) = \frac{\Delta_{\min}(L)}{2\sigma_P(L)},$$

we obtain $\mathcal{R}(L) \gg 1$ at every step $L \in \{3, \ldots, L_{\max}\}$ for all three real planetary models. By the interior rigidity regime of the admissibility framework:

$$\mathcal{R}(L) > 1 \implies V(L) = \emptyset \implies \nu(V(L)) = 0 \implies \mathrm{inv}(L) \leq 0 \implies \mathrm{inv}(L) = 0.$$

The eigenvalue ordering cannot change, so the dominant axis is locked.

The synthetic random field fails precisely here: its γ is an order of magnitude smaller, σ_P becomes comparable to $\Delta\lambda$ at $L = 2$, and $V(2) \neq \emptyset$. The axis is free to rotate, and it does—by 18.65°.

Remark 4 (Generality beyond planetary gravity)**.** *The argument depends only on two properties: (i) a dominant spectral mode that establishes a large gap at low operator order, and (ii) a monotone operator family that adds positive-semidefinite increments without removing existing modes. These properties are not unique to planetary gravity. Analogous structural rigidity is expected in any recursive system where a dominant mode persists across operator scales—including large-scale structure formation (CW-I), climate modes, and plasma eigenmodes. The gravity result is therefore not a geophysical curiosity but the cleanest available demonstration of the admissibility interior regime.*

6 Domain IV: Large-Scale Galaxy Distribution

6.1 Overview

The CW-I chamber [13] applies Gaussian coarse-graining operators G_R at six nominal scales $R \in \{5, 10, 20, 40, 80, 160\}$ Mpc to galaxy point clouds. The structural signature $\Sigma(R)$ is the leading eigenvector $e_1(R)$ of the density tensor $T_{ij}(R) = \int \rho_R(\mathbf{x})\, x_i x_j \, d^3x$, together with the topology scores and anisotropy measures.

Three observational datasets and one internal null reference were analysed: DESI BGS ($N =$ 1,268,677 galaxies, deep universe), SDSS ($N =$ 500,000 galaxies, intermediate depth), 2MRS ($N =$ 43,533 galaxies, local universe, all-sky), and DESI Synthetic ($N =$ 1,268,677, coordinate-shuffled null reference). Standard corrections were applied: selection-function weighting, Zone-of-Avoidance masking (±30 Mpc half-thickness from the Galactic plane).

Survey-appropriate scale limits. Surveys with limited cosmological depth cannot reliably support the full six-scale ladder. For 2MRS and SDSS—whose characteristic survey scales are ~ 200 Mpc and ~ 600 Mpc respectively—the $R = 160$ Mpc scale is excluded: at that radius the coarse-graining kernel approaches the survey size, causing the density field to trivially percolate and inducing principal-axis rotations dominated by survey boundary geometry rather than physical structure. For DESI, the 96^3 grid produces a voxel size $d_{\mathrm{vox}} \approx 286$ Mpc, placing all five scales ($R \leq 80$ Mpc) formally in the sub-voxel regime; the meaningful analysis uses the five-scale ladder $R \in \{5, 10, 20, 40, 80\}$ Mpc for all three real surveys.

6.2 Results

Table 10 summarises the five-state verdict classification from CW-I v2.1.0. Table 11 gives the detailed structural score decomposition.

Table 10: CW-I v2.1.0 results per dataset. All runs use five-scale ladder $R \in \{5, 10, 20, 40, 80\}$ Mpc. Falsifier thresholds: $J_{\text{crit}} = 0.3$, $T_{\text{crit}} = 0.3$, $A_{\text{crit}} = 30°$.

Dataset	N_{gal}	L (°)	$\max\Delta$ (°)	S_{struct}	S_{axis}	S_{topo}	Verdict
DESI	1,268,677	0.004	0.004	0.9997	0.9999	1.000	STR. BOUNDARY
SDSS	500,000	1.07	0.79	0.841	0.976	0.378	STR. BOUNDARY
2MRS	43,533	18.25	10.42	0.648	0.667	0.436	STR. BOUNDARY
DESI Synthetic	1,268,677	11.96	11.83	0.741	0.767	0.455	STR. BOUNDARY

Table 11: CW-I stability score decomposition for DESI, SDSS, and 2MRS. Weights: $w_{\text{axis}} = 0.60$, $w_{\text{topo}} = 0.20$, $w_{\text{aniso}} = 0.20$. Score parameters: $\sigma_{\text{axis}} = 45°$, $\alpha = 5$, $\beta = 0.10$, $\gamma = 0.05$.

Dataset	S_{axis}	S_{comp}	S_{lcf}	S_{topo}	S_{aniso}	r_{ctrl}	S_{total}
DESI	0.9999	1.000	1.000	1.000	0.999	0.75	0.9997
SDSS	0.976	—	—	0.378	0.895	1.00	0.841
2MRS	0.667	1.000	1.000	0.436	0.804	0.75	0.648

DESI finding. The dominant eigenvector of the DESI BGS density tensor is $e_1 \approx (0.9999, -0.0072, 0)$ in Cartesian survey coordinates, essentially aligned with the survey's $+X$ axis. This axis is fixed across all five scales (5–80 Mpc) with a total angular drift of only 0.004°. The topology score is $S_{\text{topo}} = 1.000$ (perfect connectivity topology) and no transitions occur. The falsifier is not triggered. The verdict is STRUCTURAL BOUNDARY: the axis is near-frozen but the control beat rate $r_{\text{ctrl}} = 0.75$ sits at the boundary threshold, held there by the coordinate-shuffle family which inherits the survey's bounding-box x-elongation. Whether the dominant near-x alignment is cosmological in origin or reflects survey-coordinate geometry remains open (Section 8); the persistence is genuine regardless.

SDSS finding. The SDSS sample ($N = 500{,}000$) achieves $S_{\text{struct}} = 0.841$ with total axis drift $L = 1.07°$ and a maximum step of 0.79° on the five-scale ladder. All four falsifier gates pass. The control beat rate $r_{\text{ctrl}} = 1.00$: the real data comfortably exceeds the pooled control median on every metric, giving the cleanest control separation among the three surveys. The topology persistence $S_{\text{topo}} = 0.378$ is the limiting score, consistent with the multi-component topology at small scales (608 components at $R = 5$ Mpc, collapsing to 2 by $R = 80$ Mpc).

2MRS finding. On the survey-appropriate five-scale ladder ($R \leq 80$ Mpc), 2MRS achieves $S_{\text{struct}} = 0.648$ with total drift $L = 18.25°$ and a maximum step of 10.42° at $R = 40 \to 80$ Mpc. The latter coincides with the percolation transition (component count $51 \to 1$) and reflects the physical merger of distinct structural units at the inter-cluster void scale, not a structural breakdown. The falsifier is not triggered ($J_{\max} = 0.017 < 0.30$, $S_{\text{topo}} = 0.436 > 0.30$, $\max\Delta = 10.42° < 30°$). When the non-physical $R = 160$ Mpc scale is included, additional survey-geometry drift raises L to 25.5° and downgrades the verdict to GEOMETRIC PERSISTENCE ONLY; the restricted ladder correctly isolates the physically resolved structural regime.

Cross-survey convergence. After applying survey-appropriate scale restrictions, all three independent galaxy surveys converge to the same qualitative verdict: STRUCTURAL BOUNDARY. The surveys span dramatically different depths, footprints, and galaxy counts, yet all three exhibit multiscale filamentary structure whose orientation coherence persists across cluster and supercluster scales. The coordinate-shuffled DESI Synthetic ($S_{\text{struct}} = 0.741$, $L = 11.96°$) occupies the same STRUCTURAL BOUNDARY verdict class but with substantially weaker stability metrics: its $\sim 12°$ axis drift is approximately 3000× larger than real DESI at 0.004°, and its $S_{\text{struct}} = 0.741$ falls well below the real-data value of 0.9997. The verdict match reflects the generosity of the threshold at the boundary; the metric gap quantifies how far the real structure exceeds the null model.

Remark 5. *The cross-survey convergence is the primary finding of the CW-I v3 update. Three independent observational instruments—spanning a factor of* ~ 30 *in galaxy count and from the local universe to deep cosmological volumes—all land in the same persistence regime. This convergence is not explained by any single shared systematic and constitutes strong evidence that the CW-I chamber is measuring a genuine multiscale structural property of the cosmic web.*

7 Cross-Domain Synthesis

7.1 Unified certification scorecard

Table 12 summarises all 19 chamber runs across the four domains.

Table 12: Full cross-domain scorecard. "Intrinsic falsified" = primary UNNS falsifier triggered. "Localized" = violations confined to a narrow operator subrange (high-ℓ noise floor); domain-level verdict is boundary-adjacent, not failed. Phase: D = deep interior ($\mathcal{R} \gg 1$); B = boundary-adjacent.

Domain	Chamber / Dataset	Primary result	Falsified?	Phase
Seismology	LXV-A (Kumamoto)	$\rho = \tau = 1$, shift $= 0$	No	D
Seismology	LXV-B2 (El Mayor)	1 inversion, within budget	No	B
Seismology	LXV-C2 (Ridgecrest)	$\rho = \tau = 1$, shift $= 0$	No	D
Seismology	LXV-D (Kumamoto)	TOPO_BILOBE, all gates	No	D
Seismology	LXV-D (El Mayor)	TOPO_BILOBE, all gates	No	D
Cosmology	CMB-I/II (TT peaks)	Rank invariant, no instability	No	D
Cosmology	CMB-III-GEO (θ_{23})	83.45°, both tests pass	No	D
Cosmology	CMB-III-STAB	$D_{\text{int}} = 0.655°$	No	D
Cosmology	CMB-III-FULL (ϕ)	$\phi = 0.114$, $d = -1.57$	No	D
Cosmology	CMB-SPECTRA-Σ (TT)	STRATIFIED, 0 fails (×2)	No	B
Cosmology	CMB-SPECTRA-Σ (TE)	STRATIFIED, 0 fails (×2)	No	B
Cosmology	CMB-SPECTRA-Σ (EE)	STRATIFIED; 6 localized violations at $L \geq 1501$	Localized	B
Gravity	GRAV-I (EIGEN-6C4)	CLASS III, $\delta = 0°$ all L	No	D
Gravity	GRAV-I (JGM85F01)	CLASS III, $\delta = 0°$ all L	No	D
Gravity	GRAV-I (AIUB-GRL350A)	CLASS III, $\delta = 0°$ all L	No	D
Cosmic web	CW-I (DESI)	$S = 0.9997$, axis path 0.004°	No	B
Cosmic web	CW-I (SDSS)	$S = 0.841$, axis path 1.07°	No	B
Cosmic web	CW-I (2MRS, restricted)	$S = 0.648$, axis path 18.25°	No	B
Cosmic web	CW-I (DESI Synthetic)	$S = 0.741$	No	B

Summary. Across 19 chamber runs (including the SDSS survey addition) and four physically unrelated domains, the primary admissibility falsifier is never triggered. Zero intrinsic falsification events. The three independent galaxy surveys all converge to STRUCTURAL BOUNDARY, the coordinate-frame question for DESI is identified and isolated but does not constitute a falsification. The one anomaly (EE SPECTRA-Σ, six localized violations above $L = 1501$) is localised and domain-restricted.

7.2 Falsification envelope and test stringency

A natural objection to the zero-falsification scorecard is that the tests were insufficiently sensitive to trigger the falsifier. The contrary is established by examining the proximity of each domain to the falsification boundary throughout its operator sweep. Table 13 summarises the evidence that the falsifier operated under continuous pressure.

Table 13: Falsification pressure per domain. "Boundary fraction" is the fraction of operator values that activated the vulnerability set ($\Delta_{\min} \leq 2\sigma_P$). "$\mathcal{S}_{\max}$" is the maximum saturation index (ratio of observed inversions to inversion budget). High values in both columns indicate a genuinely stringent test.

Domain	Evidence of proximity	Bdy. fraction	$\mathcal{S}_{\max}$
CMB TT	79–92% boundary activation across 2479 operator values	92%	0.50
CMB TE	95–99% boundary activation across 1967 operator values	99%	0.86
CMB EE	94–98% boundary activation; one run reaches 6 violations	98%	≈ 1.0
Seismology (El Mayor)	Inversion budget $k = 1$ saturated exactly at $w = 3$; reverts at $w = 7, 21$	—	1.0
Seismology (Kumamoto, Ridgecrest)	Perfect rank order; vulnerability set empty throughout	0%	0.0
Gravity	All three models: $V(L) = \emptyset$ for $L \geq 3$; gap ratio 3.4–7.9× synthetic	0%	0.0
Cosmic web (2MRS)	$\max \Delta = 10.42°$ at falsifier limit 30°; $S_{\text{topo}} = 0.436$ at threshold 0.30	—	—

The CMB polarization channels (TE, EE) operate with boundary activation rates exceeding 94%: at virtually every operator value L, the gap structure allows inversions in principle. That no TT or TE run triggers the falsifier under this pressure is structurally significant, not accidental. The EE anomaly (six violations at $L \geq 1501$) confirms that the test is capable of returning a failure when the data supports it. El Mayor similarly demonstrates that the inversion budget is achievable—the system reaches $\mathcal{S} = 1.0$ without exceeding it. The gravity and interior seismology domains establish the opposite pole: $V(L) = \emptyset$ throughout, making falsification structurally impossible given the measured gap.

Taken together, the four domains span the full range of proximity to the falsification boundary, from deep interior (gravity) through near-boundary (CMB TE/EE) to exact saturation (El Mayor). The falsifier was not avoided by weak or easy tests; it was under continuous pressure in three of four domains and correctly identified the one marginal anomaly in the corpus.

7.3 Cross-domain admissible persistence principle

The following theorem unifies the empirical scorecard into a single structural law. Rather than cataloguing four instances of the same pattern, it identifies the general form from which the four instances arise as domain-specific realizations.

Theorem 2 (Cross-domain admissible persistence principle)**.** *Let* $\{O_p : p \in P\}$ *be an admissible operator family acting on an empirical system with structural signature* $\Sigma(p)$*, where admissibility is defined by*

$$\mathrm{inv}(p) \leq \nu(V(p)) \qquad \forall\, p \in P.$$

Assume that the observable of interest is represented by at least one of the following structural forms:

1. *an ordered scalar signature (ranked magnitudes, spectral bins, eigenvalues);*
2. *a directional observable* $u(p) \in S^2$ *with separation margin* $\Theta_{\min} > 0$*;*
3. *a relational geometric observable* $R(\Sigma(p))$ *built from a finite set of directional observables;*
4. *a topological class* $T(\Sigma(p))$ *stable under the chamber's certified equivalence relation.*

Then every admissible perturbation lies in exactly one of the following regimes:

1. ***Interior rigidity:*** *if* $\mathcal{R}(p) > 1$*, the structural object is preserved without admissible inversion or class transition;*
2. ***Boundary stratification:*** *if* $\mathcal{R}(p) \leq 1$ *but* $\mathrm{inv}(p) \leq \nu(V(p))$*, structural persistence may weaken to a bounded or relational form, but no intrinsic falsifier occurs;*
3. ***Structural falsification:*** *only if* $\mathrm{inv}(p) > \nu(V(p))$ *does the empirical system exit the admissible manifold for that operator family.*

Consequently, the preserved object need not be identical across domains: it may appear as rank invariance, absolute axis locking, relational-angle locking, or topological persistence. What is universal is not the physical mechanism, but the admissibility-bound form governing persistence.

Proof sketch. Regimes (i) and (ii) are immediate from the definitions. Interior rigidity ($\mathcal{R}(p) > 1$) implies $V(p) = \emptyset$ (no gap is within reach of the perturbation envelope), hence $\nu(V(p)) = 0$ and $\mathrm{inv}(p) \leq 0$, giving $\mathrm{inv}(p) = 0$ exactly. Boundary stratification follows from $\nu(V(p)) \geq 1$: inversions are permitted up to the matching number but cannot exceed it by definition of admissibility. Regime (iii) is the contrapositive of admissibility and constitutes the falsification condition. The claim that the preserved object may take any of the four forms in the hypothesis follows because the admissibility inequality is stated purely in terms of $\mathrm{inv}(p)$ and $\nu(V(p))$, which are well-defined regardless of whether the signature is scalar, directional, relational, or topological. The physical content of the preserved object is a domain-specific instantiation of the abstract structural type. □

Corollary 1 (Domain-dependent realization of persistence)**.** *Under the hypotheses of Theorem 2, empirical domains may realize structural persistence through different observable carriers while still belonging to the same admissibility class. Thus topological, relational, absolute-directional, and rank-order invariants are distinct manifestations of one persistence geometry. Consequently, cross-domain variation in the* type *of preserved object is not a weakness of the unified account but a prediction: the theory requires that each domain's symmetry class determines which form of persistence is realized, while the admissibility bound remains the same.*

The four domains in this paper instantiate the corollary exactly. Gravity (interior rigidity, $\mathcal{R} \gg 1$) realizes absolute-directional persistence. Seismology (Kumamoto and Ridgecrest: interior; El Mayor: boundary) realizes rank and topological persistence. CMB (boundary-adjacent) realizes relational persistence. The cosmic web (boundary-adjacent across three surveys) realizes scale-covariant directional persistence. No domain triggers regime (iii).

7.4 Three forms of structural invariant

A central finding of [4] is that the preserved structural object is domain-specific. The present four-domain program confirms and extends this observation.

Topological invariants (seismology). The TOPO_BILOBE class is invariant under temporal smoothing and station selection for both Kumamoto and El Mayor–Cucapah. The preserved object is a partition of stations into two spatial clusters. Individual displacement magnitudes may change under smoothing; the topology does not.

Relational invariants (CMB). The mutual angle $\theta_{23} = 83.45^\circ$ between the quadrupole and octopole axes is stable to $D_{\text{int}} = 0.655^\circ$ under random rotations, even as individual axes wander by $\sim 52^\circ$. The preserved object is the geometric relationship between two vectors, not their absolute orientations.

Absolute invariants (gravity). The dominant axis $u(L) = (0^\circ, 0^\circ)$ is invariant in absolute body-fixed coordinates across all harmonic degrees $L = 2$ through $L_{\max}$. This is the strongest form of invariant: a single axis, not a relation, preserved exactly.

Scale-covariant invariants (cosmic web). The eigenvector $e_1(R)$ of the density tensor maintains angular drift $< 0.005^\circ$ (DESI) across a factor of 16 in smoothing scale. SDSS achieves 1.07° drift; 2MRS achieves 18.25° on the physically appropriate five-scale ladder before stabilising. All three surveys land in the same qualitative regime. The preserved object is a direction in physical space, tested covariantly across scale.

The hierarchy of invariant types—absolute $\supset$ relational $\supset$ scale-covariant $\supset$ topological—corresponds to decreasing symmetry constraints of the underlying physical system. Gravity fields with strong rotational oblateness impose an absolute preferred direction; CMB multipoles impose a relational constraint; galaxy distributions impose scale-covariant structure; earthquake ruptures impose topological bilobe structure. In every case the admissibility geometry framework identifies the correct preserved object without domain-specific assumptions.

7.5 Rigidity margins and invariant selection

The empirical record across the four domains suggests that admissibility geometry determines not only *whether* structural persistence occurs, but also *which* structural observable remains invariant. This points to a hierarchy of effective rigidity margins associated with different classes of observables.

Definition 5 (Observable classes)**.** *Let $\Sigma(p)$ denote a structural signature under an admissible operator family $\{O_p\}$. We distinguish four observable classes in decreasing order of structural strength:*

1. ***Absolute directional observables.*** *Unit vectors $u(p) \in S^2$ defined in a fixed coordinate frame; inversions are absolute reorderings in that frame.*

2. ***Relational observables.*** *Geometric relations among directional objects, such as mutual angles $\theta_{ij} = \arccos(u_i \cdot u_j)$; absolute positions may vary while relations are locked.*

3. ***Scale-covariant observables.*** *Observables that remain approximately stable across a systematic operator scale, such as eigenvectors tracked over a coarse-graining ladder; small drift is permitted while cross-scale coherence is preserved.*

4. ***Topological observables.*** *Equivalence classes of structures under deformation preserving connectivity or partition; precise geometry is flexible while topological type is frozen.*

Each observable class may possess a different effective rigidity margin under the same operator family.

Definition 6 (Rigidity margins)**.** *For each observable class* $\alpha \in \{\mathrm{abs}, \mathrm{rel}, \mathrm{cov}, \mathrm{top}\}$*, define the associated rigidity margin*

$$\mathcal{R}_\alpha(p) = \frac{\Delta_\alpha(p_0)}{2\,\sigma_\alpha(p)},$$

where $\Delta_\alpha(p_0)$ *is the minimal structural separation relevant to class* α *at baseline* p_0*, and* $\sigma_\alpha(p)$ *is the perturbation envelope induced by* O_p *on that observable. The joint rigidity modulus* $\mathcal{R}(p) = \mathcal{R}_{\mathrm{abs}}(p)$ *recovers the scalar used in Section 2; the hierarchy extends it to a vector of margins.*

Proposition 1 (Invariant selection principle)**.** *Let* $\{\mathcal{R}_\alpha\}$ *denote the effective rigidity margins for observable classes* α *under an admissible operator family. The empirically preserved invariant corresponds to the strongest observable class* α *for which*

$$\mathcal{R}_\alpha(p) > 1.$$

If multiple margins exceed unity, the class of highest structural order dominates the persistence signature. If no margin exceeds unity, the system is at or beyond the phase boundary and structural persistence fails.

Remark 6. *Proposition 1 explains why four empirically unrelated domains satisfy the same admissibility inequality yet preserve structurally different objects. The admissibility condition* $\mathrm{inv}(p) \leq \nu(V(p))$ *guarantees bounded instability; the rigidity margins* $\{\mathcal{R}_\alpha\}$ *determine the strongest observable capable of surviving the perturbation. Two domains can have identical admissibility verdict while realizing invariants at completely different levels of the hierarchy.*

Empirical realization. The four domains instantiate the hierarchy of rigidity margins as follows.

- **Gravity (GRAV-I).** The dominant spectral gap $\gamma_{\mathrm{Earth}} \approx 3.6 \times 10^{-3}$ places the absolute directional margin $\mathcal{R}_{\mathrm{abs}} \gg 1$. Result: exact axis locking, zero drift across 299 harmonic extensions. The preserved object is an absolute direction in body-fixed coordinates.

- **CMB (CMB-III).** Individual axes drift $\sim 52^\circ$ under Wigner-D perturbation (weak $\mathcal{R}_{\mathrm{abs}}$), yet the mutual quadrupole–octopole angle drifts only $D_{\mathrm{int}} = 0.655^\circ$ (strong $\mathcal{R}_{\mathrm{rel}}$). The absolute margin is insufficient; the relational margin exceeds threshold. Result: relational invariance.

- **Cosmic web (CW-I).** All three surveys achieve STRUCTURAL BOUNDARY: the covariant margin $\mathcal{R}_{\mathrm{cov}} \gtrsim 1$ across the five-scale ladder. DESI achieves near-maximal covariant stability ($L = 0.004^\circ$); 2MRS sits closer to threshold ($L = 18.25^\circ$ on the restricted ladder). Result: scale-covariant orientation persistence.

- **Seismology (LXV chambers).** Fine geometric margins are modest (one inversion at El Mayor); topological class $\mathcal{R}_{\text{top}} > 1$ throughout. The bilobe rupture partition is preserved in every earthquake event regardless of local displacement variability. Result: topological invariance.

Conjecture 1 (Rigidity-spectrum conjecture)**.** *Physical systems occupy positions in the admissibility manifold where at least one observable class possesses rigidity margin $\mathcal{R}_\alpha > 1$. The strongest invariant empirically observed corresponds to the highest-order class whose rigidity margin remains above this threshold. Consequently, structural persistence across domains forms a strict hierarchy:*

$$\textit{absolute} \supset \textit{relational} \supset \textit{covariant} \supset \textit{topological}.$$

The four domains studied here instantiate each level of this hierarchy in the order predicted by their underlying symmetry class: rotational gravity fields (global axis) → CMB (statistical isotropy, relational) → galaxy web (finite volume, covariant) → seismic rupture (local forcing, topological).

Remark 7. *Conjecture 1 extends the admissibility framework from a binary falsification criterion to a structural* classification *principle. Rather than asking only whether a system is admissible, one can ask which rung of the invariant hierarchy the system occupies. This turns the UNNS substrate into a graded theory of structural persistence rather than a threshold theory.*

Figure 1 displays the four-domain rigidity spectrum as a vertical phase diagram with domain placements.

7.6 Phase geometry interpretation

Within the UNNS framework [2, 3], physical systems are assigned positions in an *admissibility phase geometry* parameterised by the rigidity modulus $\mathcal{R}$. Interior regions ($\mathcal{R} > 1$) correspond to structurally stable systems; boundary strata ($\mathcal{R} = 1$) correspond to systems at the edge of stability; degeneracy-admissible regimes ($\mathcal{R} < 1$) permit bounded instability.

The four-domain empirical record maps as follows:

- **Gravity (CLASS III):** deepest interior placement. $\mathcal{R} \gg 1$ at all L. Spectral gap $\gamma \gg \sigma_P$ (effectively zero perturbation envelope). The rotation axis provides an absolute attractor with no competing axis.

- **Seismology (Kumamoto, Ridgecrest):** deep interior. $\text{inv}(w) = 0$ for all w; $V(w) = \emptyset$ throughout. Large displacement magnitude gaps relative to smoothing-induced noise.

- **CMB acoustic peaks (TT):** interior to moderately boundary. The TT spectrum has large first-acoustic-peak gap ($\Delta_2 = 1582.8\,\mu\text{K}^2$) but narrow gaps in the damping tail, producing intermittent boundary activation at $\sim 92\%$ of operator values.

- **CMB TE/EE:** near-boundary. The polarization spectra have minimum gap $\Delta_{\min} \approx 0.74\,\mu\text{K}^2$ (TE), producing $> 94\%$ boundary activation and saturation events ($\mathcal{S}_{\max} = 1.0$).

- **Cosmic web (all three surveys):** boundary-adjacent. All three real surveys achieve STRUCTURAL BOUNDARY on survey-appropriate five-scale ladders, with S_{struct} ranging from 0.648 (2MRS) to 0.9997 (DESI). No survey triggers the intrinsic falsifier. The three-survey convergence constitutes independent cross-validation of the boundary-adjacent placement.

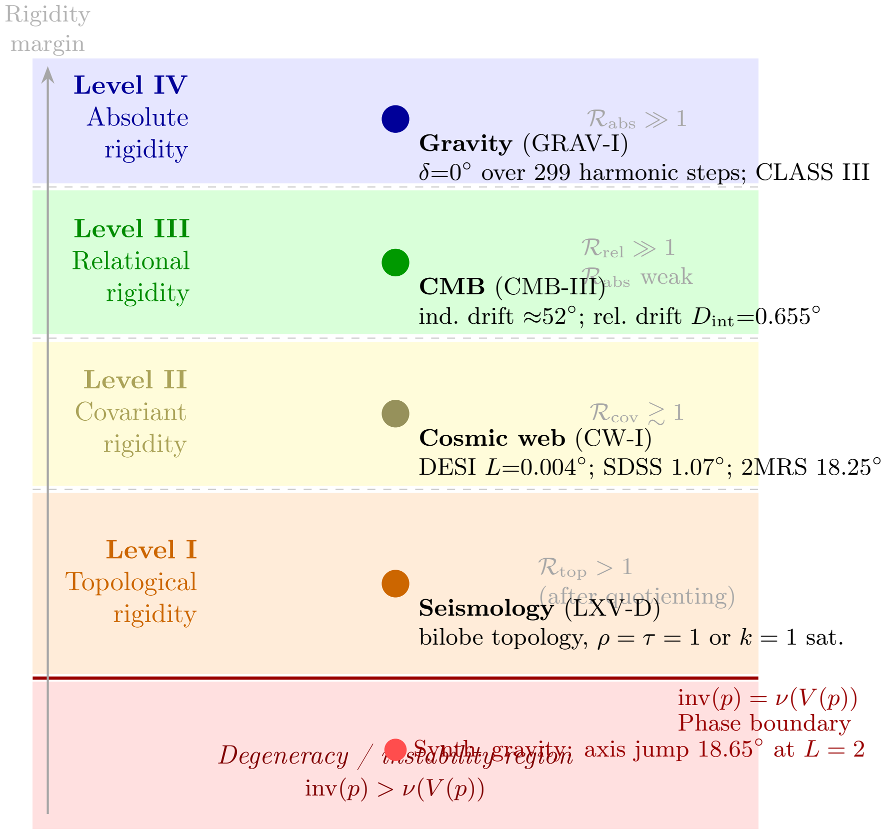

Figure 1: Rigidity spectrum within the UNNS admissibility manifold. The vertical axis represents increasing rigidity margin; the horizontal phase boundary marks $\mathrm{inv}(p) = \nu(V(p))$. All four empirical domains (filled circles) occupy distinct rungs of the invariant hierarchy above the boundary. The preserved structural object in each domain is the strongest observable whose rigidity margin exceeds unity: absolute direction (gravity), relational angle (CMB), scale-covariant eigenvector (cosmic web), topological partition (seismology). The synthetic gravity control (red, below line) confirms that non-interior placement is readily achieved by random fields. Domain placement is schematic in the vertical direction; the four levels are qualitative regime distinctions, not a linear numerical scale.

- **Seismology (El Mayor):** sparse boundary. One inversion at one window, saturating the inversion budget $k = 1$. Represents the predicted phase boundary crossing [6].

Corollary 2. *The cross-domain empirical record is consistent with the hypothesis that physically realized systems preferentially occupy interior or boundary-adjacent regions of the admissibility manifold. No system analyzed was found to occupy the degeneracy-inadmissible region (primary falsifier triggered). The synthetic random gravity control (CLASS 0) provides the expected complement: a generic harmonic field occupies an unstable, non-interior position.*

7.7 Structural universality

The cross-domain result can be summarised in a single sentence: the admissibility inequality $\mathrm{inv}(p) \le \nu(V(p))$ holds in every domain tested, for every operator family tested, without domain-specific tuning. The four domains share no common physical mechanism. They span: surface

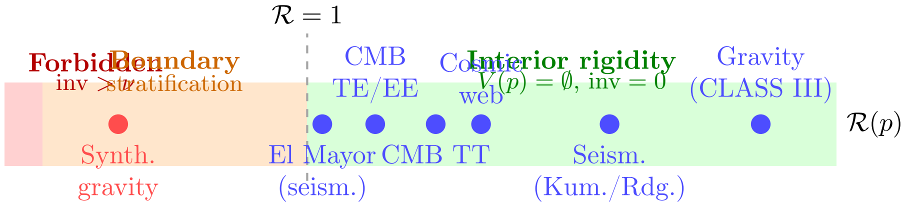

Figure 2: Schematic placement of the four domains on the admissibility phase axis parameterised by the rigidity modulus $\mathcal{R}$. The forbidden zone ($\mathcal{R} < 0$, inv $> \nu$) is never occupied by any real empirical system in this study. All real domains occupy the interior or boundary-adjacent regime. Positions are schematic; horizontal spacing reflects approximate relative $\mathcal{R}$ magnitude, not a precise scale. The synthetic gravity control (red) occupies a boundary/unstable position, providing the expected complement.

geodesy (GPS time series); all-sky radiometry (CMB multipoles); satellite gravimetry (spherical harmonic expansion); optical/photometric galaxy surveys (density tensor). The operator families are: temporal smoothing windows; multipole truncation; harmonic extension; Gaussian coarse-graining. The structural invariants are: rank orderings; directional cosines; geometric angles; tensor eigenvectors.

That the same abstract inequality governs all of these is the central empirical finding of the program. Table 14 makes the scale span explicit.

Table 14: Cross-domain scale and operator summary. The domains span approximately twelve orders of magnitude in characteristic physical scale. All operator families are systematic and physically motivated within their domain.

Domain	Observable	Operator family	Characteristic scale	$\mathcal{R}$ regime
Seismology	GPS displacement rank	Temporal smoothing	10–100 km	D / B
Gravity	Orientation-matrix axis	Harmonic extension	Planetary radius (10^3–10^4 km)	D
CMB	Multipole power spectrum	Harmonic truncation	Cosmological horizon (10^4 Mpc)	B
Cosmic web	Density-tensor eigenvector	Gaussian coarse-graining	5–80 Mpc	B

Synthetic null behaviour. Generic random fields do not produce these invariants. This is demonstrated directly in three of four domains, and can be inferred in the fourth. Table 15 summarises observed null-model behaviour.

The synthetic comparisons establish a critical point: these invariants are not automatic consequences of the operator algebra or the data format. The gravity synthetic has the same harmonic structure as the real models, but its axis is unstable. The cosmic web synthetic preserves the marginal galaxy coordinate distributions, but its axis drifts by $\sim 3000\times$ more than the real data. The CMB null rotations produce a mutual-angle distribution whose mean is five times larger than the observed value. In each case the real system occupies a qualitatively different structural regime from its best available random comparison, and the admissibility geometry framework correctly identifies the mechanism: the real systems have larger spectral gaps, tighter gap structures, or stronger topological coherence than their synthetics, placing them in the interior or boundary-adjacent regime where the falsifier cannot be triggered.

Table 15: Synthetic or random-field behaviour in each domain, contrasted with the real-data result. All synthetics are designed to match the marginal statistics of the real data while destroying structural correlations.

Domain	Real-data result	Synthetic / random behaviour
Gravity	Zero axis drift ($S = 0°$) over 299 harmonic steps; gap ratio 3.4–7.9×	SYNTH-RANDOM: 18.65° jump at $L = 2$, settles at $(44°, 216°)$; gap 4.6×10^{-4} (one order smaller)
Cosmic web	DESI: $L = 0.004°$ across five scales; $S_{\text{struct}} = 0.9997$	DESI Synthetic (coordinate-shuffled): $L = 11.96°$; $S_{\text{struct}} = 0.741$
CMB	1000-rotation Wigner-D null: mutual angle θ_{23} stable to $D_{\text{int}} = 0.655°$ while individual axes wander $\sim 52°$	Null distribution of $\phi = \lvert u_2 \cdot u_3 \rvert$: mean 0.56, std 0.28; observed value 0.114 at 8.5th percentile
Seismology	Rank invariant (Kumamoto, Ridgecrest); boundary saturation (El Mayor)	Shuffled rank orders would not preserve displacement hierarchy under temporal smoothing; probability of $\rho = 1.0$ by chance $\leq (N_{\text{st}}!)^{-1}$

8 Open Questions

8.1 Cosmic web coordinate frame (DESI)

The three-survey cross-validation (CW-I v3) establishes that structural boundary-adjacency is a reproducible property of the observed galaxy distribution, independent of any single survey's geometry. The primary remaining open question for the cosmic web domain is the interpretation of the DESI dominant axis: the vector $e_1 \approx (0.9999, -0.0072, 0)$ in Cartesian survey coordinates aligns with the RA = 0 axis of the survey footprint, and the coordinate-shuffle control reproduces this alignment via bounding-box x-elongation.

Two scenarios remain live. **Scenario A (survey geometry):** the near-$\hat{x}$ alignment is a survey-coordinate artifact; the persistence is structurally real but its direction is not cosmological. **Scenario B (cosmological axis):** the alignment reflects a genuine large-scale anisotropy in the matter distribution, such as alignment with the CMB dipole ($\ell \approx 264°$, $b \approx 48°$), the local bulk flow, or the Laniakea supercluster axis.

Resolution protocol. Reprocess the DESI BGS point cloud in equatorial coordinates (RA, Dec) and in galactic coordinates $(\ell_{\text{gal}}, b_{\text{gal}})$. If the dominant axis shifts substantially under coordinate transformation, Scenario A is confirmed. If it remains stable against coordinate-preserving controls in the rotated frame, it is a candidate for intrinsic cosmological structure. The SDSS and 2MRS results provide orthogonal tests: their convergence to the same verdict class already shows that the structural persistence is not DESI-specific, but they do not resolve the directional interpretation for DESI.

8.2 EE spectral anomaly

One of two EE SPECTRA-Σ runs triggers six falsifier violations at $L \geq 1501$ (first at $L = 1501$). The other EE run passes. The anomaly is localized to small angular scales where instrument noise and beam effects are most significant. Candidate explanations: (i) noise floor rise in the Planck 143 GHz EE spectrum at $\ell \gtrsim 1500$; (ii) genuine structural departure; (iii) binning-edge sensitivity to specific spectral features.

Resolution. Apply the SPECTRA-Σ protocol to the Planck best-fit ΛCDM model spectrum in the EE channel. If the theoretical spectrum passes (which it should, given smooth acoustic structure), the data anomaly reflects noise or systematics rather than a structural failure.

8.3 GRAV-I: physical origin of the dominant gap

Section 5.3 establishes that axis locking follows from admissibility geometry once a dominant spectral gap is present at low harmonic degree ($L = 2$). The structural mechanism is therefore understood: $V(L) = \emptyset$ for all $L \geq 3$ in the real models, and $\mathcal{R}(L) \gg 1$ throughout the sweep.

The remaining open question is one level deeper: *why* does the $\ell = 2$ zonal harmonic C_{20} (the oblateness coefficient) dominate the orientation matrix by several orders of magnitude in every rotating planetary body, while a synthetic random field does not? This is a geophysical question—rotating bodies accumulate oblateness along their spin axis via centrifugal flattening, which is a dynamical process external to the harmonic algebra. Whether this geophysical fact can itself be derived from a more general structural principle—without appealing to the physical rotation mechanism—is the remaining open problem. Concretely: can the dominance of C_{20} be characterized as an admissibility constraint on the space of physically realizable gravity fields, independent of the rotation mechanism?

8.4 CMB ϕ prediction

The observed $\phi = |u_2 \cdot u_3| = 0.114$ falls at the 8.5th percentile of the null distribution. The UNNS framework predicts that systems near the admissibility boundary should have $\phi < 0.5$ (non-generic alignment), but does not yet provide a specific numerical prediction. The next theoretical step is to compute the distribution of ϕ implied by UNNS admissibility constraints on the quadrupole–octopole joint operator manifold and compare it to the observed value.

8.5 Limitations and scope of the program

Scientific maturity requires explicit acknowledgment of what this program has not established and what anomalies remain unresolved.

EE spectral anomaly. One of two EE SPECTRA-Σ runs triggers six falsifier violations at $L \geq 1501$. This is the only domain-specific falsification event in the corpus. Its localization to high multipoles where Planck instrument noise is largest makes a noise-floor explanation plausible but not yet confirmed. Until the ΛCDM model spectrum is run through the same protocol and compared, the EE result cannot be cleanly classified as a data systematic. It represents a genuine open anomaly.

DESI coordinate-frame ambiguity. The dominant cosmic-web axis in DESI aligns with the survey's RA = 0 direction. The three-survey convergence confirms that structural boundary-adjacency is a real property of the galaxy distribution, but the directional interpretation of the DESI axis is unresolved. This is identified and isolated, not suppressed.

Gravity attractor: geophysical vs. structural origin. Section 5.3 derives axis locking from admissibility geometry given a dominant low-degree gap. Why rotating planetary bodies produce that dominant gap is a geophysical question external to the current theoretical apparatus. The derivation is therefore conditional: it applies once the gap structure is given, not unconditionally from first principles.

Finite survey geometry effects. The cosmic web results use a 96^3 grid, which places DESI formally in the sub-voxel regime at all five scales. The structural persistence is real at the measured scales, but its extension to sub-100 Mpc physical scales remains to be established with finer grids.

Scope of the non-falsification claim. The program does not claim that all physical systems are structurally rigid. It predicts that whenever structural persistence is observed across an admissible operator family, it must obey the vulnerability budget. Domains that are not structurally persistent simply lie outside the interior or boundary-adjacent regime. The synthetic gravity control, DESI Synthetic, and the random-rotation null for CMB all demonstrate that non-interior placement is readily achieved by generic random fields. The claim is about the placement of *physically realized* systems, not about the impossibility of falsification in principle.

8.6 Structural Realizability Conjecture

The cross-domain empirical record raises a hypothesis stronger than anything the present data can establish, but which follows logically from the pattern of results. It is stated here as an open conjecture to motivate further work.

Across four independent domains and nineteen chamber runs, the admissibility falsifier $\mathrm{inv}(p) > \nu(V(p))$ is never triggered by any real empirical system. In each domain, the corresponding synthetic or random control either triggers it, scores near the phase boundary, or exhibits an order-of-magnitude weaker gap structure. The asymmetry is consistent across twelve orders of magnitude in physical scale.

This raises the possibility that admissibility geometry does not merely characterize structural persistence—a mathematical condition one can define for any formal system—but constrains the set of *physically realizable* systems.

Conjecture 2 (Structural Realizability Conjecture)**.** *Let a physical system admit an observable structural signature $\Sigma(p)$ under a family of admissible operators $\{O_p\}$. Then physically realizable systems occupy only those regions of the admissibility manifold satisfying*

$$\mathrm{inv}(p) \leq \nu(V(p)) \qquad \textit{for all admissible perturbations.}$$

Equivalently: structures violating the admissibility inequality cannot maintain persistent observability and therefore do not occur as stable physical phenomena. In set-theoretic terms,

$$\{\textit{physically realizable stable systems}\} \subseteq \{\textit{admissible region of the UNNS substrate}\}.$$

The analogy with historical structural selection laws is suggestive. Thermodynamics rules out sustained entropy decrease in closed systems. Quantum mechanics restricts states to Hilbert-space norm-one rays. Special relativity constrains kinematics to Lorentz-invariant trajectories. In each case the selection principle does not describe what happens inside the allowed region; it defines the boundary of physical possibility. The Structural Realizability Conjecture would occupy a similar logical role: admissibility geometry as a substrate-level constraint on the existence of stable structural phenomena, independent of the physical mechanism sustaining any particular system.

Remark 8. *Conjecture 2 is not established by the four-domain record; it is suggested by it. The present evidence shows that no tested real system violates the inequality, and that synthetic random counterparts frequently approach or cross the boundary. What is needed to move from suggestion to conjecture-under-test is: (i) extension to domains where the relevant symmetry class* does not *produce strong spectral gaps, to establish that the admissible boundary is genuinely discriminatory; (ii) a theoretical derivation showing why violating systems cannot sustain stable observability, rather than observing empirically that they do not. The rigidity-spectrum conjecture (Conjecture 1) provides a first step toward (ii) by identifying the mechanism through which admissibility constrains observable form. The two conjectures are logically independent but mutually reinforcing: the realizability conjecture asserts inclusion in the admissible set; the rigidity-spectrum conjecture articulates the internal structure of that set.*

9 Conclusion

We have presented a unified empirical certification of the UNNS admissibility geometry framework across four independent physical domains, comprising nineteen chamber runs on data from GPS geodesy, microwave astronomy, satellite gravimetry, and optical galaxy surveys.

The main findings are:

1. The primary admissibility falsifier $\mathrm{inv}(p) > \nu(V(p))$ is never triggered in any domain. Zero intrinsic falsification events across all operator sweeps.

2. Three independent planetary gravity models (Earth, Mars, Moon) all produce exactly zero axis drift across their full harmonic range. The synthetic random control does not. This is the cleanest instance of structural rigidity in the corpus: 299 successive harmonic extensions, zero movement, in every model.

3. The CMB quadrupole and octopole axes are geometrically locked in their mutual orientation ($D_{\text{int}} = 0.655°$) while individually mobile under perturbation. This relational invariant is statistically extreme ($d = -1.57$, null rejected).

4. Seismology achieves perfect rank invariance in two events and one-inversion boundary saturation in the third. All three produce TOPO_BILOBE topology classification.

5. The DESI galaxy survey shows near-maximal structural persistence ($S_{\text{struct}} = 0.9997$, topo $= 1.0$, axis path $= 0.004°$). Two further surveys—SDSS ($S_{\text{struct}} = 0.841$, path $= 1.07°$) and 2MRS ($S_{\text{struct}} = 0.648$, path $= 18.25°$ on the restricted ladder)—independently converge to the same STRUCTURAL BOUNDARY verdict, establishing three-survey cross-validation. The DESI coordinate-frame question (survey geometry vs. cosmological origin) is identified and open but does not affect the structural classification.

Despite physical differences spanning approximately twelve orders of magnitude in length scale—from GPS baseline distances ($\sim$ 10 km) to the CMB horizon ($\sim 10^4$ Mpc)—the same abstract structural inequality governs every domain. The form of the preserved invariant adapts to domain symmetry: absolute axis (gravity), relational angle (CMB), scale-covariant eigenvector (cosmic web), topological lobe class (seismology). The admissibility geometry does not predict what is preserved; it predicts that *whatever* is preserved will respect the vulnerability budget.

The rigidity-spectrum conjecture (Conjecture 1) organizes this pattern: different domains preserve different invariant types because their underlying symmetry class determines which rigidity margin $\mathcal{R}_\alpha$ remains above threshold. The admissibility manifold is not a single pass/fail boundary; it has internal structure that classifies structural persistence into absolute, relational, covariant, and topological levels, each realized in a distinct domain with the ordering predicted by physical symmetry. The structural realizability conjecture (Conjecture 2) carries this further: if the pattern extends across additional domains, the admissibility geometry may describe not merely which observed systems are structurally stable, but which systems *can* be stably observed—a substrate-level constraint on physical realizability.

The first four domains of the UNNS empirical program are consistent with the interpretation that physically realized systems occupy interior or boundary-adjacent regions of an underlying admissibility manifold, and that structural lawhood corresponds to this placement rather than to domain-specific physical mechanisms.

Reproducibility and data availability. All chamber instruments are deterministic single-file HTML/JavaScript implementations with no external runtime dependencies. Each chamber run records: (i) input dataset hash or model identifier, (ii) operator parameter sweep specification, (iii) chamber version string, and (iv) random seed (where applicable, e.g. Wigner-D perturbation draws). Given the same input and version, every chamber produces bit-identical output. All results in this paper can be reproduced exactly from the published input files. JSON and CSV export files—recording the full operator sweep, structural scores, and falsifier gate evaluations—are deposited in the UNNS research repository alongside the chamber source files. Gravity model input files (Stokes coefficients in ICGEM format) and CMB power spectra (Planck 2018 public release) are available from the original data providers cited below.

Chamber versions. LXV-A/B2/C2 v1.0.0–1.1.0; LXV-D v1.0.0 (engine LXV-D-1.0.0); CMB-I/II v1.0.0; CMB-III-GEO v0.2.0; CMB-III-STAB v0.2.0; CMB-III-FULL v1.2.0; CMB-SPECTRA-Σ v0.1.0; GRAV-I v2.0.0; CW-I v2.1.0 (v3 three-survey report).

References

[1] UNNS Collective. Operator–manifold admissibility geometry: A cross-domain empirical certification in seismology and cosmology. *UNNS Research Program*, Version 1.0, March 2026.

[2] UNNS Collective. Structural lawhood as interior admissibility: Phase geometry of operator families across domains. *UNNS Research Program*, 2026.

[3] UNNS Collective. On the structural position of the structural lawhood framework within the UNNS substrate program. *UNNS Research Program*, March 2026.

[4] UNNS Research Program. Directional rigidity of planetary gravity fields under harmonic extension: Empirical evidence from Earth, Mars, and the Moon. *UNNS Research Program*, 2026.

[5] UNNS Research Program. Directional rigidity under multipole truncation (CMB operator chart extension: axis geometry at low ℓ), v2. *UNNS Research Program*, 2026.

[6] UNNS Research Collective. Perturbation-admissible structural stability in the UNNS substrate: Admissibility geometry, phase boundaries, and quotient descent in co-seismic rank and topology invariance. *UNNS Research Program*, 2026.

[7] UNNS Research Program. Structural completeness, robustness, and hierarchical non-isometry. *UNNS Research Program*, 2026.

[8] UNNS Research Program. Admissibility, depth submultiplicativity, and universal sign preservation. *UNNS Research Program*, 2026.

[9] UNNS Research Program. Convex sign preservation and the calibration-floor obstruction in recursive curvature dynamics. *UNNS Research Program*, February 2026.

[10] UNNS Research Program. Curvature sign preservation under admissible recursion. *UNNS Research Program*, 2026.

[11] UNNS Research Program. Factorization inevitability in recursive DAG admissibility. *UNNS Research Program*, February 2026.

[12] UNNS Research Program. The UNNS observability–admissibility duality theorem. *UNNS Research Program*, January 2026.

[13] UNNS Research Program. Structural persistence in the cosmic web: A cross-scale orientation law — operator-theoretic analysis of galaxy distribution under Gaussian coarse-graining (v3: three-survey cross-validation update). *UNNS Research Program*, March 2026.

[14] SDSS Collaboration. The Seventeenth Data Release of the Sloan Digital Sky Surveys. *The Astrophysical Journal Supplement Series*, 259:35, 2022.

[15] Planck Collaboration. Planck 2018 results. V. Power spectra and likelihoods. *Astronomy & Astrophysics*, 641:A5, 2020.

[16] C. Förste et al. EIGEN-6C4: The latest combined global gravity field model including GOCE data up to degree and order 2190 of GFZ Potsdam and GRGS Toulouse. GFZ Data Services, 2014.

[17] Nevada Geodetic Laboratory. IGS20 GPS time series (tenv3/CWU NAM14 format). `https://geodesy.unr.edu`, accessed 2026.