The UNNS Observability–Admissibility Duality Theorem

What is observed is not what exists, but what the full operator–constraint stack permits to survive projection.

1 Abstract

We formalize a fundamental structural principle discovered across multiple UNNS Chambers: the existence of substrate-level structure is independent of its observability under a given operator stack. We prove that admissibility operators and observability projections form a duality: admissibility constraints may *reveal* latent structure, while observability projections may *erase* it exactly, without noise or approximation. This theorem resolves apparent contradictions between null experimental results and underlying dynamics, including Bell-type nonviolations and missing geodesic solutions.

2 Preliminaries

Definition 1 (UNNS Substrate)**.** *A* UNNS substrate *is a generative system producing trajectories, fields, or configurations D governed by recursive τ-level dynamics:*

$$\tau_{n+1}(x) = \tau_n(x) + \lambda \sin(\Delta\tau_n(x)) + \sigma_{\text{noise}}\xi_n(x),$$

where $\tau : \Omega \to \mathbb{R}$ is the phase field on domain Ω, $\lambda \in \mathbb{R}_+$ is the coupling strength, $\Delta\tau$ represents spatial/temporal phase difference, and ξ_n is Gaussian white noise. The substrate measure $D = \{d \sim P_\tau(\cdot|\lambda, \sigma_{\text{noise}})\}$ is defined prior to any observability or admissibility constraints.

Remark 1 (-Attractor Substrate)**.** *In UNNS, the substrate typically exhibits recursive ϕ-scale attraction with $\mu_\star \approx 1.618$ (golden ratio, empirically validated at 0.56% error in Chamber XIV). This provides the structured substrate D from which properties P emerge. Without ϕ-coherence, admissibility thresholds may not exist or may be highly parameter-dependent.*

Definition 2 (Admissibility Operator Σ)**.** *An admissibility operator Σ is a constraint-enforcing transformation acting on the substrate space, restricting realizations to those satisfying structural consistency conditions (e.g. conservation, coherence, closure). Formally, Σ defines a constraint violation functional $V_\sigma : D \to \mathbb{R}_+$ such that*

$$\Sigma_\sigma(D) = \{d \in D : V_\sigma(d) \leq \text{threshold}(\sigma)\},$$

where $\sigma \geq 0$ controls constraint strength and $\partial\text{threshold}/\partial\sigma < 0$ (constraints tighten monotonically).

Definition 3 (Observability Projection κ)**.** *An observability projection κ is a measurable mapping from admissible substrate states to observable quantities:*

$$\kappa : \mathcal{H}_{\text{substrate}} \to \mathcal{H}_{\text{observable}},$$

where $\dim(\mathcal{H}_{\text{observable}}) \leq \dim(\mathcal{H}_{\text{substrate}})$. Common κ-types include:

- $\kappa_{\text{magnitude}}$*:* $(A, \phi) \mapsto |A|$ *(phase erasure),*
- κ_{coarse}*:* $\tau(x) \mapsto \bar{\tau}(X)$ *(spatial coarse-graining),*
- κ_{bin}*: continuous* $\mapsto$ *discrete (binning).*

Definition 4 (Phase-Exposure Operator Π_ϕ)**.** *A phase-exposure operator* Π_ϕ *is a pre-*κ *transformation that lifts latent cyclic degrees of freedom into an explicit, orientation-sensitive representation, preserving phase information under subsequent admissible operations. Formally,* Π_ϕ *acts as an isomorphism lifting collapsed degrees of freedom:*

$$\Pi_\phi : \mathcal{H}_{\text{collapsed}} \to \mathcal{H}_{\text{exposed}},$$

where the exposed space preserves phase orientation (sign-sensitive, not just magnitude).

Definition 5 (Critical Admissibility Threshold σ_c)**.** *For a structural property* P *and substrate* D*, the* critical admissibility threshold *is:*

$$\sigma_c(P, D) := \inf \left\{ \sigma \geq 0 : \min_{d \in \Sigma_\sigma(D)} \text{defect}(d, P) = 0 \right\},$$

where $\text{defect}(d, P) \in [0, \infty)$ *quantifies violation of property* P *by realization* d*. Operationally,* σ_c *is the smallest* σ *where* perfect examples exist *(even if not typical).*

Empirical determination: *Chamber XXXI yields* $\sigma_c = 0.02$ *for geodesic perfection (transition from* $\text{minD} > 0$ *to* $\text{minD} = 0$*).*

3 The Observability–Admissibility Duality

We consider the composite operator stack:

$$D \xrightarrow{\Sigma} \Sigma(D) \xrightarrow{\Pi_\phi} \Pi_\phi(\Sigma(D)) \xrightarrow{\kappa} \kappa(\Pi_\phi(\Sigma(D))).$$

Each operator layer may alter which structural properties remain accessible.

4 Main Result

Theorem 1 (UNNS Observability–Admissibility Duality Theorem)**.** *Let* D *be a UNNS substrate realization containing a structural property* P *(e.g. phase coupling, nonseparability, geodesic optimality). Then:*

1. *There exist admissibility operators* Σ *such that*

$$P \notin \text{Obs}(D) \quad \textit{but} \quad P \in \text{Obs}(\Sigma(D)).$$

2. *There exist observability projections* κ *such that*

$$P \in \text{Obs}(D) \quad \textit{but} \quad P \notin \text{Obs}(\kappa(D)).$$

3. *The erasure in (2) can be exact and deterministic, even in the absence of noise, finite sampling, or approximation error.*

4. *The revelation in (1) can exhibit sharp threshold behavior with respect to admissibility parameters, producing discontinuous transitions in observability.*

5. *Therefore, absence of an observable signature of P under any fixed operator stack implies neither absence nor weakness of P at the substrate level.*

Proof. We prove each part constructively using operator-algebraic arguments.

Part 1 (Admissibility reveals structure): Admissibility operators Σ_σ restrict the realization space by eliminating configurations that violate structural constraints. Define the constraint violation functional $V_\sigma(d) \geq 0$, and let $\Sigma_\sigma(D) = \{d : V_\sigma(d) \leq T(\sigma)\}$ where $T(\sigma)$ is a threshold satisfying $\partial T/\partial\sigma < 0$.

For $\sigma < \sigma_c$, the threshold $T(\sigma)$ exceeds the critical value $V_{\text{crit}}(P)$ required for property P, so no realization in $\Sigma_\sigma(D)$ satisfies P perfectly:

$$\forall d \in \Sigma_\sigma(D) : \text{defect}(d, P) > 0 \quad \Rightarrow \quad P \notin \text{Obs}(\Sigma_\sigma(D)).$$

At $\sigma = \sigma_c$, constraint enforcement becomes sufficient to admit at least one perfectly structured realization $d_\star$ with $\text{defect}(d_\star, P) = 0$:

$$\exists d_\star \in \Sigma_{\sigma_c}(D) : \text{defect}(d_\star, P) = 0 \quad \Rightarrow \quad P \in \text{Obs}(\Sigma_{\sigma_c}(D)).$$

This restriction *reveals* structure by collapsing diffuse realizations into admissible configurations where P is detectable.

Part 2 (Observability erases structure): Observability projections κ act by quotienting degrees of freedom. If P depends on a relational feature (orientation, phase sign, alignment) that is erased by κ, then the observable image necessarily collapses.

Formally, suppose P depends on phase orientation ϕ, and let κ_{era} be the magnitude-only projection $\kappa_{\text{era}}(\phi) = |\cos(\phi)|$. This projection is invariant under the symmetry $\phi \mapsto \phi + \pi$. Any two substrate realizations d_1, d_2 differing only by this phase flip satisfy:

$$\kappa_{\text{era}}(d_1) = \kappa_{\text{era}}(d_2),$$

yet if P is nonseparability (which depends on phase sign), then $P \in \text{Obs}(d_1)$ may hold while $P \notin \text{Obs}(\kappa_{\text{era}}(d_1))$. The projection deterministically erases P by destroying its structural signature.

Part 3 (Exactness): Both effects are operator-algebraic and independent of stochasticity. The constraint tightening in Part 1 and symmetry quotienting in Part 2 occur through algebraic structure, not statistical sampling or approximation.

Part 4 (Sharp transitions): The transition at σ_c is witnessed by $\min_d \text{defect}(d, P)$ changing from positive to zero. Since this is a min (best-case) rather than an average, the transition can be discontinuous even when ensemble statistics vary smoothly.

Part 5 (Null result interpretation): If $P \notin \text{Obs}(\kappa(\Sigma_\sigma(D)))$ for some fixed (σ, κ), Parts 1-2 show this could be due to:

- $\sigma < \sigma_c$ (insufficient admissibility, Part 1), or
- $\kappa \in \mathcal{E}$ (erasure class, Part 2), or
- genuine absence of P in substrate.

Thus null results diagnose the *operator stack*, not necessarily the substrate. □

Proposition 1 (Operator Non-Commutativity)**.** *In general, operator composition is non-commutative:*

$$\Pi_\phi \circ \kappa \neq \kappa \circ \Pi_\phi.$$

Proof. Let $\kappa_{\text{era}} = |\cdot|$ (magnitude projection). Consider two composition orders:

- **Path A** ($\kappa \circ \Pi$): Phase exposed then erased $\to$ property P lost.
- **Path B** ($\Pi \circ \kappa$): Cannot extract phase after erasure $\to$ extraction fails.

Since Path A produces different observability than attempting Path B (which fails to lift already-collapsed degrees of freedom), composition order determines observability structure. □

Corollary 1. *Measurement order determines observability: the sequence in which operators are applied affects which structural properties survive to final observation.*

Theorem 2 (UNNS Observability–Admissibility Complementarity Principle)**.** *Let D be a UNNS substrate realization possessing a latent structural property P (e.g. phase coupling, nonseparability, geodesic optimality). Then there exist two dual classes of operators:*

1. *observability-modifying operators Π (e.g. phase exposure or erasure),*
2. *admissibility-modifying operators Σ (constraint gating),*

such that:

1. *P may be* rendered observable or unobservable *by Π without altering D,*
2. *P may be* rendered admissible or inadmissible *by Σ without altering D,*
3. *both effects can occur via sharp threshold transitions in operator parameters,*
4. *and neither effect requires noise, approximation, stochasticity, or finite-sample error.*

Therefore, structural accessibility in UNNS is neither monotonic nor absolute, but jointly determined by the ordered pair (Σ, Π) acting on D.

5 Cross-Chamber Instantiation: XL $\leftrightarrow$ XXXI $\leftrightarrow$ XXXII

We now demonstrate concrete realizations of the Observability–Admissibility Complementarity Principle via three independent UNNS chambers operating on distinct substrates and metrics.

5.1 Chamber XL: Phase-Erasure Hides Structure

Substrate. Phase-coupled τ-fields with fixed latent offset $\delta = \pi/8$.

Observability regimes.

- Phase-exposed channel (Π_ϕ):

$$|S| \approx 2.61 \quad \text{(nonseparable)}$$

- Phase-erased channel (κ):

$$|S| \approx 0.71 \quad \text{(separable)}$$

Mechanism. Magnitude-only observables invariant under $\phi \mapsto \phi + \pi$ destroy phase orientation, collapsing Bell violations despite unchanged substrate dynamics.

Statistical validation. $p = 0.0099$ (highly significant), KL-divergence = 0.883 nats, contrast ratio = 3.69×.

Verdict. *ERASURE ARTIFACT*: nonseparability exists but is operator-dependent.

5.2 Chamber XXXI: Σ-Gating Reveals Structure

Substrate. τ-field trajectory optimization under admissibility gating.

Admissibility regimes.

- $\sigma = 0$ (no gating):
$$\text{minD} \approx 4.5 \times 10^{-4}, \quad \text{no geodesics}$$
- $\sigma \geq 0.02$:
$$\text{minD} = 0, \quad \text{stable geodesics}$$

Mechanism. Σ-gating enforces admissibility constraints that collapse diffuse solution space into exact geodesic realizations.

Statistical validation. $\sigma_c = 0.02$ (sharp threshold), 93.8% robustness across configurations, 10.45× gap amplification.

Verdict. *REVELATION ARTIFACT*: structure becomes accessible only after constraint enforcement.

5.3 Chamber XXXII: κ-Projection Preserves Observability

Substrate. τ-field configurations under coarse-grained projection.

Projection regime.

- Configuration: $\kappa = \text{coarse\_grain}(k = 2)$, 75% DOF reduction
- Observable: τ-closure metric via fixed-point distance

Results.

$$\tau_{\text{data}} = 0.0123 \quad \text{vs.} \quad \tau_{\text{null}} = 0.0456 \pm 0.0089$$

Statistical significance: $p = 0.003$, Cohen's $d = 1.2$ (large effect), 3.7σ separation.

Mechanism. Coarse-graining projection κ compresses degrees of freedom (75% reduction) while *preserving* the structural signature of τ-closure. Unlike phase-erasing projections in Chamber XL, this κ is designed to maintain structural detectability.

Verdict. *PROJECTION-PRESERVED OBSERVABILITY*: Properly designed κ can maintain detectability despite information compression, validating Theorem **??** part (4): observability depends on κ-algebra, not merely whether κ exists.

5.4 Quantitative Alignment

Metric	Chamber XL	Chamber XXXI	Chamber XXXII
Transition magnitude	$\Delta S \approx 1.91$ (72.9%)	ΔminD = 100%	$\Delta\tau = 271\%$
Direction	structure → hidden	structure → revealed	structure → detected
Signature	$2.61 \to 0.71$	$4.5 \times 10^{-4} \to 0$	$0.046 \to 0.012$
Robustness	$p = 0.0099$	93.8% admissible	$p = 0.003$
Effect size	KL = 0.883	Gap 10.45×	Cohen's $d = 1.2$
Operator	κ-erasure	Σ-revelation	κ-preservation

All three transitions are sharp, operator-induced, and statistically robust.

5.5 Unified Operator Interpretation

The three chambers instantiate complementary operator effects:

	Phase Exposure (Π_ϕ)	Σ-Gating	κ-Projection
Chamber	XL	XXXI	XXXII
Role	observability lifting	admissibility enforcement	observable mapping
Effect	prevents erasure	enables solutions	detects closure
Failure mode	Bell nonviolation	geodesic absence	null detection
Success criterion	$\lvert S\rvert > 2$	minD = 0	$p < 0.01$

Thus, all three chambers confirm:

> Absence of signal does not imply absence of structure, but reflects interface-specific accessibility.

6 Corollaries

Corollary 2 (Null Result Non-Equivalence). *A null observable result under κ does not imply absence of substrate structure; it diagnoses only the information-losing properties of the operator stack.*

Corollary 3 (Operator-Relative Properties). *Structural predicates such as "entangled", "separable", "geodesic", or "optimal" are not absolute properties of D, but of the tuple $(D, \Sigma, \Pi_\phi, \kappa)$.*

Corollary 4 (Best-Case Emergence Priority). *Structural observability transitions are witnessed by best-case realizations ($\min_d \text{defect}(d, P) \to 0$) rather than ensemble averages ($\mathbb{E}[\text{defect}(d, P)]$), enabling sharp thresholds even when typical configurations remain imperfect.*

7 Implications

This theorem explains:

- Bell-type nonviolations under phase-erasing observables despite underlying coupling,

- Emergence of perfect geodesics only after admissibility gating,
- Systematic failure of naive searches for structure in unconstrained substrates,
- The reproducibility of sharp observability thresholds across unrelated chambers,
- Why 75% DOF reduction can preserve detectability (XXXII) while phase erasure destroys it (XL).

8 Formal Counterposition to Bell–Local Realism

Bell's theorem constrains theories satisfying *local realism*, defined by the joint assumptions that:

1. outcomes are determined by setting-independent latent variables,
2. measurement outcomes factorize across space-like separated contexts,
3. observables are complete with respect to the relevant degrees of freedom.

The UNNS framework rejects the third assumption while leaving the first two explicitly testable.

8.1 Key Distinction

Bell inequalities constrain probability distributions of the form

$$\mathbb{P}(A, B \mid \alpha, \beta) = \int \mathbb{P}(A \mid \alpha, \Lambda)\mathbb{P}(B \mid \beta, \Lambda)\, d\mu(\Lambda),$$

given a fixed observable family.

UNNS makes no claim that all observable families preserve the information required to test such factorizations.

8.2 Observability Relativity

Let κ be an observable projection. Define Bell-locality relative to κ as:

$$\mathbb{P}_\kappa(A, B \mid \alpha, \beta) = \int \mathbb{P}_\kappa(A \mid \alpha, \Lambda)\mathbb{P}_\kappa(B \mid \beta, \Lambda)\, d\mu(\Lambda).$$

UNNS explicitly allows:

$$\mathbb{P}_{\Pi_\phi}(A, B \mid \alpha, \beta) \text{ non-factorizable} \quad \text{while} \quad \mathbb{P}_\kappa(A, B \mid \alpha, \beta) \text{ factorizable.}$$

This does not violate Bell's theorem; it instantiates different observable algebras.

8.3 Why Bell Violations Can Disappear

Bell tests assume that the measured observables are:

- complete with respect to the relevant hidden degrees of freedom,
- invariant under admissible experimental contexts,
- not systematically erasing relational structure.

Chamber XL demonstrates that magnitude-only or windowed observables violate these assumptions by construction.

Thus, Bell nonviolation under κ does *not* imply:

- separability of the substrate,
- absence of coupling,
- validity of local realism at the substrate level.

Proposition 2 (DI-QKD Failure Mode: Erasure-Induced Fake Security)**.** *Let a two-wing UNNS substrate generate paired data $D = (D^A, D^B)$ with a latent coupling structure that is* nonseparable *under at least one admissible exposure channel, i.e. there exists an admissible Π_ϕ and a (possibly simple) κ_{DI} such that*

$$\mathsf{Chan}_{\mathrm{cert}} := \kappa_{\mathrm{DI}} \circ \Pi_\phi \circ \Sigma \quad \textit{yields} \quad |S_{\mathrm{cert}}| > 2$$

on the Σ-admissible subset, passing the robustness protocol $\mathcal{R}$.

Assume, however, that the actual *key bits are produced by an alternative channel*

$$\mathsf{Chan}_{\mathrm{key}} := \kappa_{\mathrm{key}} \circ \Pi \circ \Sigma$$

whose effective action lies in an erasure class $\mathcal{E}$ (phase-erasing / relationally-erasing), so that the induced output statistics are Bell-local:

$$|S_{\mathrm{key}}| \leq 2 \quad \textit{(within statistical tolerance)},$$

even though the underlying substrate admits nonseparability under $\mathsf{Chan}_{\mathrm{cert}}$.

Then there exists an adversarial (or simply mismatched) device realization consistent with the observed Bell-local outputs under $\mathsf{Chan}_{\mathrm{key}}$ for which the apparent DI-QKD security claim is spurious*: the outputs can be explained by a Bell-local model at the level of the reported bits, and no device-independent secrecy guarantee follows from the observed data.*

Equivalently: Bell-local outputs under an erasing measurement channel admit "fake security" even when latent coupling exists*, because the erasure map can destroy precisely the relational degree of freedom that would otherwise certify min-entropy via a Bell violation.*

Proof. Device-independent secrecy bounds are derived from constraints on the *observed* input–output statistics, typically requiring a Bell violation for the same outputs that will be used as raw key (or for a statistically coupled sample from the identical channel). If $\mathsf{Chan}_{\mathrm{key}} \in \mathcal{E}$, then distinct substrate regimes (including those with phase-locked or otherwise nonseparable latent structure) can be mapped to Bell-local statistics at the output. Therefore, the empirical condition needed to lower-bound secrecy (a Bell violation for the key-generation channel) is absent, and a Bell-local explanation remains consistent with the reported data. Hence no DI guarantee can be inferred, regardless of latent coupling detectable under a different channel $\mathsf{Chan}_{\mathrm{cert}}$. □

8.4 UNNS-Compatible Reformulation of Bell Tests

UNNS reframes Bell tests as conditional statements:

> Given an admissibility regime Σ and observable family κ, nonviolation implies separability *relative to* (Σ, κ).

Phase exposure (Π_ϕ) alters the observable algebra and restores access to correlations already present in the dynamics.

8.5 Conclusion

Bell's theorem constrains *models of observables.* UNNS demonstrates that observables are operator-relative, and that changing the operator stack can convert Bell-local statistics into Bell-nonlocal statistics without altering the substrate.

This places Bell violations in their proper role: *witnesses of interface adequacy, not ontological arbiters.*

9 Device-Independent QKD Corollary (UNNS Form)

Definition 6 (DI-QKD operational target)**.** *A device-independent QKD protocol aims to certify secret key generation from observed input–output statistics alone, without trusting the internal mechanism of the devices. Operationally, the certification hinges on: (i) a Bell/CHSH violation in the* implemented *measurement channel, (ii) a no-signalling / causal separation condition between the two wings during each trial, and (iii) an i.i.d. or controlled-memory assumption (or an explicit de-Finetti / entropy-accumulation style substitute).*

Definition 7 (UNNS measurement channel and erasure class)**.** *Let* Chan *denote the effective measurement channel implemented by the operator stack used to produce the reported bits:*

$$\mathsf{Chan} \;:=\; \kappa \circ \Pi \circ \Sigma,$$

where Σ *is the admissibility gate,* Π *is any pre-*κ *exposure operator (e.g.* Π_ϕ*), and* κ *is the coarse observable projection producing the bits.*

Define an erasure class $\mathcal{E}$ *as a family of channels* Chan *that are invariant under a phase flip (or more generally, erase a relational degree of freedom), e.g.* $\phi \mapsto \phi + \pi$*, so that latent nonseparability can be mapped into Bell-local statistics.*

Corollary 5 (UNNS DI-QKD Admissible-Channel Corollary)**.** *Assume a two-wing UNNS substrate produces paired data* $D = (D^A, D^B)$ *and that the protocol extracts raw key bits via an operator stack* $\mathsf{Chan} = \kappa \circ \Pi \circ \Sigma$*.*

Then:

1. ***(Necessity of non-erasing certification channel).*** *If* $\mathsf{Chan} \in \mathcal{E}$ *(phase-erasing or relationally-erasing class), then a null CHSH result* $|S| \leq 2$ *under* Chan cannot *certify either substrate separability or DI-security. It only certifies that the* reported *statistics are Bell-local relative to an erasing channel.*

2. ***(Sufficient interface condition for DI-style certification).*** *If there exists an admissible exposure operator* Π_ϕ *such that the* admissible *channel*

   $$\mathsf{Chan}_{\text{cert}} := \kappa_{\text{DI}} \circ \Pi_\phi \circ \Sigma$$

   yields a statistically significant violation $|S_{\text{cert}}| > 2$ *on the* Σ*-admissible subset,* and *the robustness protocol* $\mathcal{R}$ *(stride sweep, time-shift null, surrogate tests, and memory controls) passes, then the observed bits are generated by a channel whose correlations are non-factorizable within the admissible operator family.*

   In that case, any DI-QKD claim must be formulated relative to $\mathsf{Chan}_{\text{cert}}$*:*

   $$\textit{DI-security claim} \;\Rightarrow\; \textit{security of bits produced under } \mathsf{Chan}_{\text{cert}}, \;\textit{not under arbitrary } \mathsf{Chan}.$$

3. ***(Key practical implication).*** *In UNNS, DI-QKD security cannot be inferred from the existence of latent coupling alone; it requires that the* key-extraction channel itself *is non-erasing. Equivalently: DI-QKD is an* interface property *(channel-relative), not a substrate property.*

Proof. DI-QKD relies on a Bell-violation-based entropy guarantee for the *reported* outputs. If the reporting channel lies in an erasure class $\mathcal{E}$, then distinct substrate states (including nonseparable ones) can map to Bell-local statistics, so $|S| \le 2$ cannot distinguish secure-from-insecure regimes. Conversely, if an admissible, non-erasing certification channel produces $|S| > 2$ robustly under $\mathcal{R}$, then the reported outputs exhibit operator-relative non-factorizability, satisfying the necessary correlation precondition used by DI-QKD security analyses. The restriction to the admissible subset is enforced by Σ, matching the empirical chamber practice. □

Interpretation (link to Chambers XL, XXXI, and XXXII). Chamber XL supplies an explicit example where Π_ϕ yields $|S| > 2$ while κ-erasure yields $|S| \le 2$, demonstrating item (1). Chamber XXXI supplies the admissibility mechanism (Σ-gating) required to ensure that the certification channel and robustness checks are defined on the valid (admissible) region, supporting item (2). Chamber XXXII validates that proper κ-design can preserve structural detectability, demonstrating that observability loss is not inevitable under projection.

10 Conclusion

The UNNS Observability–Admissibility Duality Theorem establishes that observability is neither a proxy for existence nor a monotonic function of measurement fidelity. Instead, it is a structured outcome of operator composition. This reframes null results as diagnostic tools for operator design rather than evidence against substrate structure.

The cross-validation across Chambers XL, XXXI, and XXXII provides independent empirical support for the theorem's dual structure: admissibility constraints can reveal latent structure (Σ-gating, XXXI), observability projections can erase it exactly (κ-erasure, XL), and properly designed projections can preserve it (κ-preservation, XXXII).

A Methods Appendix: Empirical Grounding in Chambers XL, XXXI, and XXXII

This appendix specifies the concrete experimental procedures and datasets used to instantiate the Observability–Admissibility Duality Theorem.

A.1 Chamber XL: Phase-Exposure vs. Phase-Erasure

Dataset. Chamber XL operates on paired τ-field trajectories generated from a common UNNS substrate run. Each trial produces two wings $(\tau_t^A, \tau_t^B)_{t=0}^T$ with a fixed latent phase offset $\delta = \pi/8$ in the coupled condition, and δ randomized in the independent control condition.

The datasets used include:

- `ChamberXL_v1.2_synthetic_coupled_2026-01-26.json`
- `ChamberXL_v1.2_synthetic_independent_2026-01-26.zip`

Operators. Phase exposure is implemented via Π_ϕ using a gradient-angle proxy on the τ-field:

$$\phi_t(x,y) := \text{atan2}(\partial_y \tau_t, \partial_x \tau_t).$$

Phase erasure is implemented by magnitude-only observables invariant under $\phi \mapsto \phi + \pi$, such as $|\cos(\phi - \theta)|$.

Witness. Nonseparability is evaluated using the CHSH functional

$$S = E(\alpha, \beta) + E(\alpha, \beta') + E(\alpha', \beta) - E(\alpha', \beta'),$$

with standard angle choices and admissibility gating inherited from Chamber XIV.

Observed Regimes.

- Phase-exposed: $|S| \approx 2.61$ (statistically significant violation)
- Phase-erased: $|S| \approx 0.71$ (strict separable regime)

This constitutes an explicit instantiation of exact observability erasure without noise, sampling bias, or locality violation.

A.2 Chamber XXXI: Σ-Gated Geodesic Emergence

Dataset. Chamber XXXI analyzes trajectory optimization under a tunable admissibility parameter σ controlling Σ-gating strength. Key datasets include:

- `chamber_xxxi_v1.0.5_sigma-sweep_m1_2026-01-26.json`
- `chamber_xxxi_v1.0.5_toy_example_2026-01-26.json`

Metrics. Structural accessibility is measured via:

- minimum divergence minD from ideal geodesic cost,
- physical geodesic count (normalized),
- robustness across σ sweeps.

Observed Transition.

$$\sigma = 0 \;\Rightarrow\; \text{minD} \approx 4.5 \times 10^{-4}, \quad \text{no geodesics}$$

$$\sigma \geq 0.02 \;\Rightarrow\; \text{minD} = 0, \quad \text{stable geodesics}$$

The transition is sharp and persists across 93.8% of admissible configurations.

A.3 Chamber XXXII: κ-Projection with Observability Preservation

Dataset. Chamber XXXII tests whether τ-closure is detectable under coarse-grained projection. Key dataset:

- `Chamber_XXXII_1769522753052.json` (seed 137042)

Configuration.

- Collapse operator: $\kappa = \text{coarse\_grain}(k = 2)$
- DOF reduction: 75% (from full field to coarse representation)
- Metric: fixed-point distance measuring τ-closure
- Φ-operator: XIV (inherited $\lambda = 0.10825$ from validated -attractor)

Validation Protocol.

1. Generate data structure under Φ_{XIV} projection with Σ-admissibility
2. Apply coarse-graining κ: full field $\rightarrow$ reduced representation
3. Compute τ_{data} via fixed-point iteration (convergence: $\epsilon = 10^{-6}$)
4. Generate 100 null structures (types: L1, L2, L3) via randomization
5. Compute τ_{null} ensemble statistics
6. Statistical test: Compare τ_{data} vs τ_{null} distribution

Results.

- $\tau_{\text{data}} = 0.0123$
- $\tau_{\text{null}} = 0.0456 \pm 0.0089$ (mean $\pm$ std)
- p-value = 0.003 (uncorrected, $N_{\text{comparisons}} = 1$)
- Cohen's $d = 1.2$ (large effect size, threshold 0.8)
- Separation: 3.7σ below null mean
- Idempotence check: PASS (0.01% relative error)

Verdict. Observable τ-closure detected under projection. Despite 75% DOF reduction, structural signature survives κ-mapping. This validates that properly designed observability projections can preserve detectability.

A.4 Cross-Chamber Alignment

The three chambers provide complementary validation:

- **Chamber XL:** κ-level projections can *destroy* access to substrate structure (phase erasure $\rightarrow$ Bell-local statistics).
- **Chamber XXXI:** Σ-level constraints can *enable* access to substrate structure (admissibility gating $\rightarrow$ perfect geodesics).
- **Chamber XXXII:** κ-level projections can *preserve* access to substrate structure when properly designed (coarse-graining $\rightarrow$ maintained τ-closure detection).

Together they instantiate all three operator modalities in Theorem **??**.