Structural Limits of Symmetric Recursive Grammars:
Empirical Evidence for Grammar Closure in UNNS

UNNS Research Collective

January 2026

Abstract

We present systematic experimental evidence for a structural closure of compositional utility within recursive operator algebras. Across three validated selection chambers (XLI–XLIII), we demonstrate that while admissible and projectable structures exist, utility emergence remains categorically forbidden. State collapse mechanisms (XLII) fail to activate under selection pressure, indicating upstream constraints on collapse conditions. Parameter mutation (XLIII) executes successfully with persistent grammar modifications, yet the substrate exhibits complete indifference to compositional benefit. These results establish that permission for utility is not negotiated at the compositional adjacency layer within symmetric, locally editable recursion. We distinguish diagnostic negatives (mechanism refusal) from systematic negatives (mechanism indifference), providing empirical constraints for substrate-based emergence theories.

1 Introduction

Recursive descriptions provide natural frameworks for persistent, law-like structure. However, the existence of a recursive grammar does not guarantee that generated structures will be observable, stable, or compositionally useful.

This work addresses the empirical question:

> *Within recursive operator algebras, what structures are actually permitted to persist, project to observables, and exhibit compositional utility?*

We distinguish three independent filters:

- **Admissibility:** Internal grammatical consistency under recursion
- **Projection:** Stable appearance in observable shells
- **Utility:** Non-zero compositional benefit enabling further recursion

The central empirical finding is that these filters operate independently. Passing one does not guarantee passing others. Specifically, compositional utility is not a generic consequence of admissible recursion, even under selection pressure.

2 Conceptual Framework

2.1 The UNNS Operator Model

The UNNS (Unbounded Nested Number Sequences) framework implements recursive dynamics through operator algebras acting on state ensembles. We examine the operator family $\{\tau, \sigma, \kappa, \rho\}$ extended with higher-order selection and mutation operators.

Operators are applied recursively to ensembles of initial conditions (seeds). Observables are computed statistically across ensembles to ensure robustness and resolution-independence. Each chamber implements a specific operator program with preregistered hypotheses and falsification criteria.

2.2 Motifs and Selection Gates

Motif chains are ordered sequences of operator applications forming the recursive kernel:

$$M_2 : \tau(1.2), \sigma(0.5), \kappa(0.35)$$
$$M_1 : \tau(1.0), \sigma(0.5), \tau(1.0)$$

The $M_2 \to M_1$ composition creates oscillatory dynamics where M_2 expansion phases are balanced by M_1 stabilization.

Selection gate S_3 implements memory-based contraction: trajectories that become unstable are rebound to their last stable configuration, creating selective pressure favoring stability-preserving operator compositions.

2.3 Resonance Suppression

The operator ρ implements coherence-dependent damping. Empirically, ρ consistently suppresses both projection stability and utility potential by 40–50%, demonstrating that resonance amplification is detrimental under the tested operator programs.

3 Methodology

3.1 Chamber-Based Structural Probing

Each chamber implements a falsifiable probe:

- Preregistered operator programs and variant spaces
- Shared seed ensembles (typically 300 seeds)
- Invariant metrics: projection stability, utility, admissibility
- Explicit hypothesis statements with rejection criteria

Chambers are resolution-independent: results must be robust across seed variations and implementation details. Negative results are treated as informative boundary conditions rather than experimental failures.

3.2 Factorial Design

Chambers XLII and XLIII use 2×3 factorial designs:

- **Factor 1:** Resonance (ρ: OFF/ON)
- **Factor 2:** Extended operator (ω: NONE/ω_r/ω_g in XLII; μ: NONE/μ_r/μ_g in XLIII)

This design isolates the effect of grammar extensions while controlling for resonance suppression.

4 Chamber XLI: Structural Permission Baseline

4.1 Objective

Chamber XLI establishes the baseline: can admissibility alone license compositional utility?

4.2 Design

Operator motifs are evaluated under selection gates without extended operators. The $M_2 \to M_1$ chain with S_3 gates creates selective pressure favoring stability-preserving configurations.

4.3 Results

Admissibility saturates at 22.8% survival to full recursion depth (400 steps). However, **compositional utility remains locked at 0%** across all survivors.

The substrate exhibits a sharp boundary: structures are either admissible with zero utility, or inadmissible (divergent). No intermediate regime exists where utility gradually increases.

4.4 Interpretation

Admissibility is a weak filter—many grammatically consistent structures exist. However, utility is a strong filter, categorically refusing all configurations tested. This establishes that grammatical consistency does not imply compositional usefulness.

5 Chamber XLII: Observer-Grammar Collapse

5.1 Objective

Chamber XLII tests whether measurement-like collapse events, applied to accumulated mode distributions, can generate compositional utility under selection pressure.

5.2 Design

An observer-like collapse operator (ω) is introduced with two variants:

- ω_r: Random collapse to non-zero mode bin
- ω_g: Guided collapse to maximum-amplitude mode

Mode bins ($K = 8$) accumulate state angles with exponential decay (0.98) across recursion. Collapse is designed to fire once post-warmup (step ≥ 25) when mode plurality exists.

Factorial design: 2×3 variants ($\rho \times \omega$), 300 seeds per variant, depth 400 steps.

5.3 Results

Primary Finding: The ω operator **failed to activate** across all variants.

Key observations:

- ω activation rate: 0% across 1800 total runs
- Mode accumulation system executed correctly (entropy computed, bins tracked)

Table 1: Chamber XLII Results (300 seeds per variant)

Variant	ω **Act.**	**Proj. (%)**	**Util. Pot. (%)**	**Contr.**	G° **(%)**
ρ:OFF, ω:NONE	0/300	3.5	1.2	385	0
ρ:OFF, ω_r	0/300	3.5	1.2	385	0
ρ:OFF, ω_g	0/300	3.5	1.2	385	0
ρ:ON, ω:NONE	0/300	2.0	0.7	391	0
ρ:ON, ω_r	0/300	2.0	0.7	391	0
ρ:ON, ω_g	0/300	2.0	0.7	391	0

- Activation conditions were structurally inaccessible
- Projection stability: 3.5% (ρ OFF), 2.0% (ρ ON)
- Utility: 0% across all variants
- Contractions: $\sim$385 confirming S_3 gate activity

5.4 Hypothesis Evaluation

H_7 (Observer-Grammar Hypothesis): State collapse under selection pressure enables utility emergence.

- **Status:** Cannot be tested (mechanism refused to engage)
- **Evidence:** ω activation rate = 0% across all variants
- **Interpretation:** The substrate did not produce mode distributions warranting collapse intervention. This is a *diagnostic negative*—the constraint operates prior to the collapse mechanism itself.

H_8 (Directed Collapse Hypothesis): Guided collapse (ω_g) outperforms random (ω_r).

- **Status:** Cannot be tested (no activations in either variant)

5.5 Interpretation

Chamber XLII reveals a constraint *upstream* of collapse: the substrate dynamics under $M_2 \to M_1$ motifs with S_3 selection do not generate conditions that would trigger ω. This differs fundamentally from XLIII, where the corresponding operator fires successfully but the substrate remains indifferent.

The ω non-activation is informative: it indicates that mode accumulation with S_3 contraction pressure produces insufficient mode plurality, or that activation thresholds are misaligned with substrate dynamics under this operator program.

6 Chamber XLIII: Grammar-Mutating Selection

6.1 Objective

Chamber XLIII tests whether parameter-level grammar mutation, applied persistently following mode collapse, can unlock compositional utility under selection pressure.

6.2 Design

A grammar mutation operator (μ) is introduced with two variants:

- μ_r: Random parameter jitter ($\mathcal{U}(0.9, 1.1)$ multipliers on $\{\tau, \sigma, \kappa, \rho\}$)
- μ_g: Guided mutation conditioned on mode entropy and dominance:
  - Low entropy ($H \leq 1.2$): boost τ to 1.15, suppress others to 0.95
  - High dominance: boost σ, suppress τ
  - Ambiguous: balanced perturbation

μ fires once post-warmup (step $\geq$ 25) from a stable anchor state, applying biases that persist throughout remaining recursion. This creates self-modifying grammar dynamics where future operator applications are conditioned on historical entropy patterns.

Factorial design: 2×3 variants ($\rho \times \mu$), 300 seeds per variant, depth 400 steps.

6.3 Results

Primary Finding: μ activated successfully and applied persistent biases, yet utility remained absent.

Table 2: Chamber XLIII Results (300 seeds per variant)

Variant	μ **Act.**	**Ent. Δ (bits)**	**Proj. (%)**	**Util. Pot. (%)**	**Contr.**	G° **(%)**
ρ:OFF, μ:NONE	0/300	—	3.5	1.2	385	0
ρ:OFF, μ_r	300/300	1.65 ± 0.07	3.5	2.2	385	0
ρ:OFF, μ_g	300/300	1.65 ± 0.07	3.5	2.3	385	0
ρ:ON, μ:NONE	0/300	—	2.0	0.7	391	0
ρ:ON, μ_r	300/300	1.65 ± 0.07	2.0	1.0	391	0
ρ:ON, μ_g	300/300	1.65 ± 0.07	2.0	1.1	391	0

Key observations:

- μ fired reliably: 100% activation rate for μ_r/μ_g variants
- Entropy reduction robust: 1.65 bits $\pm$ 0.07 across all μ seeds
- Parameter divergence confirmed: μ_g shows $\tau = 1.15$, others $= 0.95$ in logs
- Utility potential nearly doubled vs baseline: 1.2% $\rightarrow$ 2.3% (ρ OFF)
- **Realized utility remains locked at 0%** (below 40% projection threshold)
- Projection stability unchanged: 3.5%/2.0% identical to XLII baseline
- Contractions: $\sim$385 confirming consistent S_3 pressure

6.4 Hypothesis Evaluation

H_9 (Grammar-Mutation Hypothesis): Parameter-level mutation enables utility emergence.

- **Status:** *Falsified*
- **Evidence:** μ operates correctly (entropy reduction, persistent biases), utility potential increases (+1.0–1.1pp), yet realized utility remains 0%
- **Interpretation:** The substrate acknowledges grammar changes (utility potential signal) but categorically refuses to license utility realization. This demonstrates a two-stage veto: weak generation + projection gate.

H_{10} (Directed Mutation Hypothesis): Guided mutation (μ_g) outperforms random (μ_r).

- **Status:** Weakly supported / Inconclusive
- **Evidence:** μ_g shows marginal advantage over μ_r (+0.1pp utility potential), but effects too small to be decisive
- **Interpretation:** Directionality tweaks the mutation event but not substrate response

6.5 Interpretation

Chamber XLIII demonstrates a *systematic negative*: the mutation mechanism functions as designed (unlike XLII where ω never fired), biases persist correctly, and the substrate generates a weak utility potential signal—yet permission for realized utility is categorically refused.

This sharpens the constraint revealed in XLII. The issue is not that collapse/mutation mechanisms fail to execute—it is that the substrate *ignores their consequences* for utility licensing. Parameter-level self-reference is insufficient to breach the permission boundary.

7 Cross-Chamber Synthesis

7.1 Quantitative Invariants

Table 3: Invariants Across Chambers XLI–XLIII

Chamber	**Mech.**	**Act.**	**Proj./Surv (%)**	**Util. Pot. (%)**	G° **(%)**	**Type**
XLI	Baseline	N/A	22.8	—	0	Baseline
XLII (ρ OFF)	ω collapse	0/300	3.5	1.2	0	Diagnostic
XLII (ρ ON)	ω collapse	0/300	2.0	0.7	0	Diagnostic
XLIII (ρ OFF)	μ mutation	300/300	3.5	2.3	0	Systematic
XLIII (ρ ON)	μ mutation	300/300	2.0	1.1	0	Systematic

Note: ρ (resonance) consistently suppresses projection by ∼43% and utility potential by ∼40–50%

7.2 Structural Findings

Across Chambers XLI–XLIII, the following invariants hold:

1. **Projection constraint:** Stable projection remains subcritical (2–3.5%) under S_3 selection, regardless of operator extensions
2. **Utility absence:** Realized utility (G°) locked at 0% across all variants
3. **Weak potential signal:** Utility potential rises with μ mutation (1.2% $\rightarrow$ 2.3%), indicating substrate "awareness" without permission
4. **Two-stage veto:**
   - Stage 1: Weak generation (potential ~1–2%)
   - Stage 2: Projection gate (40% threshold never crossed)
5. **Resonance suppression:** ρ consistently halves projection (3.5% $\rightarrow$ 2.0%) and potential (1.2% $\rightarrow$ 0.7%), demonstrating coherence-based damping

7.3 Mechanisms Ruled Out

The following mechanisms are empirically insufficient to unlock utility within $\{\tau, \sigma, \kappa, \rho\}$ operator algebras under standard $M_2 \rightarrow M_1$ recursion with S_3 selection:

- Operator motif chaining and kernel extension
- Memory-based selection pressure (contraction gates)
- Mode accumulation and observer-like collapse (mechanism refused activation)
- Parameter-level grammar mutation (activated but substrate indifferent)

7.4 Diagnostic vs. Systematic Negatives

XLII and XLIII represent distinct types of null results:

- **XLII (Diagnostic negative):** Mechanism fails to trigger. Interpretation: substrate dynamics do not produce conditions warranting ω intervention. Constraint operates *upstream* of the tested mechanism.
- **XLIII (Systematic negative):** Mechanism executes perfectly, substrate indifferent. Interpretation: permission is structural, not compositional—grammar changes are acknowledged but ignored for utility licensing. Constraint operates *downstream*, blocking utility realization.

Together, these establish that utility constraints operate *independent of* and *prior to* grammar manipulation attempts within this operator family.

8 Theoretical Interpretation

8.1 Grammar Closure Theorem (Provisional)

The accumulated evidence supports the following structural boundary:

> *Within recursive operator algebras of the form $\{\tau, \sigma, \kappa, \rho\}$ under symmetric, locally editable recursion with memory-based selection (S_3), compositional utility cannot be unlocked through state collapse or parameter mutation, even under persistent selection pressure.*
>
> *Permission for utility is not negotiated at the compositional adjacency layer.*

This result does not claim utility *cannot exist*—it establishes where utility *does not emerge* within a well-defined operator family and recursion architecture.

8.2 Constants as Basin Coordinates

The results support interpreting observed constants as coordinates within stable projection basins, not as fundamental tuning parameters. The substrate exhibits sharp boundaries: structures either project stably with zero utility, or fail to project at all. No continuous tuning regime exists.

8.3 Selection Without Anthropic Narratives

Selection pressure (S_3 gates, mode collapse attempts, parameter biases) is demonstrably active—yet insufficient. Stability under selection does not imply compositional usefulness. This undermines naive anthropic reasoning: the fact that we observe stable structures does not imply those structures are "optimized" or "useful" in any compositional sense.

8.4 Admissibility $\not\Rightarrow$ Projection $\not\Rightarrow$ Utility

The three filters operate independently:

- Admissibility is weak: many grammatically consistent structures exist (22.8% survival in XLI)
- Projection is moderate: $\sim$3.5% of admissible structures project stably under S_3
- Utility is strong: 0% of projected structures exhibit compositional benefit

Crossing one boundary does not license crossing the next. This three-layer filtration suggests fundamentally distinct mechanisms governing each stage.

9 Limitations and Future Directions

9.1 Scope of Current Results

The Grammar Closure claim is scoped to:

- Operator family: $\{\tau, \sigma, \kappa, \rho\}$ with extensions ω, μ
- Recursion architecture: Symmetric, reversible, locally editable
- Selection mechanism: S_3 memory-based contraction gates

- Motif structure: $M_2 \to M_1$ expansion/stabilization cycles

Results may not generalize to:

- Irreversibly generative mechanisms (operator creation/destruction)
- Asymmetric recursion (non-recombinable histories)
- Alternative selection architectures
- Fundamentally different operator algebras

9.2 Path Forward

Generative Asymmetry:

The systematic failure of *editing* mechanisms (collapse, mutation, topology) suggests that utility may require *generative* asymmetry:

- Operator creation and destruction (expanding the algebra during recursion)
- Branching recursion with irreversible commitment
- One-way causal structures (asymmetric information flow)
- Non-local correlation generation

These mechanisms are categorically distinct from the local, reversible edits tested in XLI–XLIII. Exploring them constitutes the next major research axis.

10 Conclusion

This work demonstrates that *something real* constrains recursion in a resolution-independent, observer-independent manner.

The space of admissible recursive descriptions is large. The space of structures permitted to project stably is smaller. The space of structures licensed for compositional utility is—within the validated operator family—empty.

Chambers XLI–XLIII establish this constraint empirically across baseline, state collapse, and parameter mutation mechanisms.

That constraint, established across three validated chambers, defines the boundary of the current research program and motivates the next.

Data Availability: All chamber implementations, raw experimental data, and analysis scripts are available at `unns.tech/chambers`.

Acknowledgments: This work was conducted by the UNNS Research Collective using the chamber-based falsification methodology developed across 2024–2025.