Convex Sign Preservation and the Calibration-Floor Obstruction in Recursive Curvature Dynamics: Chamber Evidence, Methodological Pipeline, and Outlook

Abstract

We analyze recursive curvature dynamics under convex kernel recombination and gain scaling, drawing on a five-chamber empirical program (Chambers LVI–LX) executed under fully locked Protocol v1.0.0.

The program establishes a threefold structural result. First, a consistent operator family hierarchy is identified across all chambers: family V2×V3 exhibits structural compounding; V2×V4 is near-isometric (trivial); and V6×V7 occupies an intermediate regime stabilizable exclusively under depth-indexed bias damping (Chamber LVII, $\alpha^* = 0.5$). Second, across all simplex-based chambers (LVIII–LX), no certified negative-curvature region was detected despite the presence of strongly negative point estimates ($b_{\min} \approx -0.30$). Third, this absence is not a failure of detection but a structural consequence of a *calibration-floor obstruction*: the negative-curvature region is confined to a high-variance boundary stratum of the operator simplex where the non-degeneracy gate cannot be cleared.

We formalize the Convex Sign Preservation Conjecture and provide a partial theorem-level statement under submultiplicativity and convexity axioms. We also present the methodological pipeline — certified partition invariance, per-cell degeneracy filtering, bracket-based sign localization, variance-calibrated obstruction detection, and basis-robust topology verification — as a reusable protocol for studying recursive operator families. The results reveal a geometric–statistical rigidity principle governing recursive operator dynamics, with conceptual parallels in convex ensemble learning, discrete flow stability, and quantum channel operator theory.

1 Setting

Let $\mathcal{K} = \{K_0, K_1, K_2\}$ be a fixed operator basis. Consider convex mixtures on the 2-simplex

$$K(\beta, \gamma) = (1 - \beta - \gamma)K_0 + \beta K_1 + \gamma K_2,$$

with $\beta, \gamma \geq 0$ and $\beta + \gamma \leq 1$.

Let $\kappa_n(K)$ denote the recursive curvature functional evaluated at depth n. Define the log-slope statistic

$$b(K) = \text{bootstrap estimate of } \frac{1}{3}\sum_{n=1}^{4} \log \frac{\kappa_{n+1}(K)}{\kappa_n(K)}.$$

Certification requires passing a non-degeneracy gate:

$$L_2(K) > 5\sigma_F(K),$$

where $\sigma_F(K)$ is the calibration floor estimated from null runs.

2 Empirical Observations Across Chambers LVI–LX

The following results are drawn from five independently certified chambers executed under Protocol v1.0.0. All runs passed sanity checks and partition-invariance validation. Statistical thresholds were preregistered: compounding threshold $\delta_{\text{comp}} = \log(1.15) \approx 0.1398$, saturation threshold $\delta_{\text{sat}} = \log(1.05) \approx 0.0488$, non-degeneracy gate $L_2 > 5\sigma_F$, and 99% bootstrap confidence intervals throughout.

Chamber LVI — Recursive Growth Classification

Chamber LVI tested operator families V2×V3, V2×V4, and V6×V7 against Theorem 2 (exponential compounding: $\text{lower}_{99}(\bar{b}) > \delta_{\text{comp}}$) and Theorem 1 (saturation: $0 \in \text{CI}_{95}(\bar{b})$ and $|\bar{b}| \leq \delta_{\text{sat}}$). The three families resolved to distinct growth regimes:

- V2×V3: **COMPOUNDING** ($\bar{b} = 0.279$, $\text{CI}_{99} = [0.243, 0.315]$, non-degenerate). Structural exponential growth confirmed.
- V2×V4: **TRIVIAL_OR_WEAK** (non-degeneracy gate failed: L_2 exceeded $5\sigma_F$, indicating near-isometric behaviour; depth ratios $r_2 = 1.22$, $r_3 = 1.14$). Dismissed from compounding analysis.
- V6×V7: **INCONCLUSIVE** ($\bar{b} = 0.074$, non-degenerate but $\bar{b} < \delta_{\text{comp}}$, $\text{CI}_{99} = [0.027, 0.122]$). Neither saturating nor compounding at preregistered thresholds.

No intrinsic stable β^* was identified. Partition invariance was confirmed throughout (both runs, reproducible to within $\Delta\bar{b} < 0.001$).

Chamber LVII — Bias-Damped Recursion and $\alpha^*(F)$

Chamber LVII introduced depth-indexed damping $\alpha \in \{1.0, 0.95, 0.90, 0.85, 0.80, 0.70, 0.60, 0.50\}$, applied as $L_{\text{eff}}(n) = \alpha^{n-1} \cdot L(n)$. The calibration floor σ_F was computed at $\alpha = 1$ and locked across all α levels, preventing calibration artifacts.

Family	$\alpha^*(F)$	$\alpha = 1$ label	Notes
V2×V3	None	COMPOUNDING	Compounding at all $\alpha \geq 0.5$
V2×V4	None	TRIVIAL_OR_WEAK	Trivial at all α levels
V6×V7	**0.5**	INCONCLUSIVE	SAT/BND first achieved at $\alpha = 0.5$

The existence summary yields ExistsStableDamped = TRUE, with minimal $\alpha^* = 0.5$ for family V6×V7. This is the *only positive existence result* in the five-chamber program: damped stabilization is achievable for V6×V7 but not for V2×V3 (curvature-intrinsic compounding) or V2×V4 (trivial at all depths). The trajectory of $\bar{b}(\alpha)$ for V6×V7 declines monotonically from 0.074 ($\alpha = 1$) to -0.008 ($\alpha = 0.5$), crossing zero between $\alpha = 0.6$ and $\alpha = 0.5$.

Chamber LVIII — Intrinsic Stability Scan ($\mathbf{K}_A/\mathbf{K}_B$ Hull)

Chamber LVIII performed a canonical intrinsic stability existence test without damping: does a recursively stable family F_β exist within the K_A/K_B convex hull? Two protocol variants were certified.

Version v1.0.0 scanned $\beta \in \{0.00, 0.05, \ldots, 1.00\}$ (21 cells) with $K_A = \text{V6}\times\text{V7}$, $K_B = \text{V2}\times\text{V3}$, global $\sigma_F = 0.0917$. Version v1.0.0-B introduced four mandatory fixes: deterministic convex mixture, per-β $\sigma_F(\beta)$ calibration, per-seed amplification ratios $r_2(s)/r_3(s)$ with bootstrap CI, and a saturation CI gate.

Both versions agree:

$$\text{ExistsIntrinsicStable} = \text{FALSE}.$$

No β achieves SATURATING_OR_BOUNDED. The label distribution shifts between versions (v1.0.0: 7 COMPOUNDING, 12 INCONCLUSIVE, 2 TRIVIAL; v1.0.0-B: 5 COMPOUNDING, 6 INCONCLUSIVE, 10 TRIVIAL) as a consequence of the tighter per-β noise floor, but the existence verdict is invariant.

Chamber LVIX — 2D $\hat{\kappa}$ Sign Map on the Kernel Simplex

Chamber LVIX extended the search to the full 2D kernel simplex (β, γ) with $\beta, \gamma \geq 0$, $\beta + \gamma \leq 1$, grid step $\Delta = 0.05$ (231 cells). Kernels: $K_0 = \text{V6}\times\text{V7}$, $K_1 = \text{V2}\times\text{V3}$, $K_2 = \text{V2}\times\text{V4}$, deterministic convex mixture. Per-cell pipeline: partition invariance $\rightarrow$ per-cell σ_F calibration $\rightarrow$ non-degeneracy gate $\rightarrow$ log-slope $\bar{b}$ $\rightarrow$ CI_{99} sign certification $\rightarrow$ bracket detection $\rightarrow$ bisection refinement ($\text{I}_{\text{max}} = 12$).

Results (both submitted files are identical exports of the same run):

- 231 cells in domain, all valid.
- 6 non-degenerate; 4 CERT_POS; 0 CERT_NEG; 0 certified brackets.
- ExistsIntrinsicStableAtκZero = FALSE.

CERT_POS cells are localized in the interior of the simplex (low-$(\beta + \gamma)$ region). No sign-change bracket was found, so no θ^* candidate entered bisection.

Chamber LX — Basis-Robust Modular $\hat{\kappa}$ Sign Map

Chamber LX added two orthogonal axes of expressivity over LVIX. Module A (basis swap): three kernel configurations B_0 ($K_0 = \text{V6}\times\text{V7}$, $K_1 = \text{V2}\times\text{V3}$, $K_2 = \text{V2}\times\text{V4}$), B_1 (K_1/K_2 swapped), B_2 ($K_0 = \text{V2}\times\text{V4}$ anchor). Module B (damping gain): $g \in \{1.25, 1.50\}$ applied to K_1, K_2 with weight renormalization $w_0'/Z \cdot K_0 + (g\beta)/Z \cdot K_1 + (g\gamma)/Z \cdot K_2$. Protocol used 20 of 30 preregistered seeds and 150 bootstrap resamples. Total: 6 slabs $\times$ 231 cells = **1,386 evaluated cells**.

Cross-slab summary:

Slab	Kernels	g	Non-deg	CERT_POS
B_0	V6×V7 / V2×V3 / V2×V4	1.25	9	6
B_0	V6×V7 / V2×V3 / V2×V4	1.50	8	3
B_1	V6×V7 / V2×V4 / V2×V3	1.25	9	6
B_1	V6×V7 / V2×V4 / V2×V3	1.50	8	4
B_2	V2×V4 / V2×V3 / V6×V7	1.25	9	6
B_2	V2×V4 / V2×V3 / V6×V7	1.50	13	7

CERT_NEG = 0 and brackets = 0 in all slabs. ExistsIntrinsicStableAtκZero = FALSE across all (B_i, g) combinations. Negative point estimates ($b \approx -0.31$ near $(\beta, \gamma) \approx (0.55, 0.45)$) persist across all six slabs but remain in a high-σ_F stratum ($\sigma_F \in [0.15, 0.22]$, 5σ threshold ≈ 1.05) that fails the non-degeneracy gate in every configuration.

Cross-Chamber Summary

Chamber	Existence verdict	Key finding
LVI	No stable β^*	V2×V3: COMPOUNDING; V2×V4: TRIVIAL; V6×V7: INCO
LVII	**ExistsStableDamped = TRUE**	$\alpha^* = 0.5$ (V6×V7 only)
LVIII	No stable β^*	ExistsIntrinsicStable = FALSE (v1.0.0 and v1.0.0-B)
LVIX	No stable $(\beta, \gamma)^*$	4 CERT_POS; 0 CERT_NEG; 0 brackets
LX	No stable $(\beta, \gamma)^*$	Basis-robust; CERT_POS topology preserved

Structural hierarchy. The family ordering V2×V3 $\succ_{\text{comp}}$ V6×V7 $\succ_{\text{comp}}$ V2×V4 is consistent across all five chambers: V2×V3 remains compounding regardless of damping or simplex position; V6×V7 is intermediate and uniquely stabilizable via bias damping; V2×V4 is universally trivial (near-isometric). This hierarchy is basis-robust: it persists across all three kernel permutations of LX.

The LVII positive result in context. Chamber LVII is the only chamber in this batch that yields an affirmative existence result (ExistsStableDamped = TRUE). The correct global statement is therefore not that everything is null, but that *intrinsic* stability is absent while *damped* stability is achievable for the intermediate family. This distinction matters for the interpretation of the LVIII–LX null results: those chambers test undamped convex mixing, which is precisely the regime where the LVII result predicts instability.

Non-certifiable negative curvature (LVIX/LX). The strongly negative point estimates localized near high-$(\beta + \gamma)$ are not absent from the geometry — they survive across all six LX slabs unchanged. Their non-certifiability is exclusively a consequence of elevated σ_F, not of their magnitude. Section 7 formalizes this as the Calibration-Floor Obstruction Principle.

3 Structural Asymmetry

Positive-curvature cells exhibited:

$$b \in [0.046, 0.138],$$

with narrow bootstrap variance and successful non-degeneracy certification.

Negative-curvature cells exhibited:

$$b \approx -0.30,$$

but large seed-to-seed variance prevented certification.

This yields a geometric asymmetry:

> Positive curvature occupies low-variance interior strata. Negative curvature is confined to high-variance boundary strata.

4 Convex Sign Preservation Conjecture

Conjecture 1 (Convex Sign Preservation). *Let κ_n satisfy admissible recursion and submultiplicative growth bounds. Suppose $b(K_i) > 0$ for all extreme kernels K_i in a basis.*

Then for all convex mixtures $K(\beta, \gamma)$ lying in the interior non-degenerate region of the simplex,

$$b(K(\beta, \gamma)) > 0.$$

Interpretation

Convex recombination and scalar gain scaling preserve curvature polarity in the interior of the operator manifold.

Sign reversal, if it exists, must either:

- Occur in boundary strata dominated by calibration variance, or
- Require nonlinear operator coupling beyond convex mixing.

5 Convex Preservation of Asymptotic Curvature Sign

The results in this section are purely analytic. They do not depend on bootstrap certification, finite-depth estimators, or the non-degeneracy gate. Separation from the statistical framework of Section 2 is intentional: what follows is a theorem, not a summary of empirical observations.

5.1 Standing Assumptions

Let $\kappa_n(K) > 0$ denote the recursive curvature functional evaluated at depth n for an admissible kernel K.

(A1) Positivity. For all admissible K and all $n \geq 1$, $\kappa_n(K) > 0$.

(A2) Submultiplicativity (up to constant). There exists a constant $C \geq 1$ such that for all admissible K and all integers $m, n \geq 1$,

$$\kappa_{m+n}(K) \leq C\,\kappa_m(K)\kappa_n(K). \tag{1}$$

Equivalently, $a_n(K) := \log \kappa_n(K) + \log C$ satisfies $a_{m+n}(K) \leq a_m(K) + a_n(K)$ (exact subadditivity).

(A3) Convexity in the kernel argument. For each fixed n, the map $K \mapsto \kappa_n(K)$ is convex along convex combinations:

$$\kappa_n\big((1-t)K + tK'\big) \;\leq\; (1-t)\kappa_n(K) + t\,\kappa_n(K'), \qquad t \in [0,1]. \tag{2}$$

5.2 Asymptotic Growth Exponent

Define the asymptotic curvature growth rate

$$\lambda(K) := \limsup_{n\to\infty} \frac{1}{n} \log \kappa_n(K). \tag{3}$$

Lemma 1 (Existence of asymptotic growth rate)**.** *Assume* (A1)–(A2)*. Then the* lim sup *in* (3) *is a true limit, and*

$$\lambda(K) = \inf_{n\geq 1} \frac{1}{n}\big(\log \kappa_n(K) + \log C\big).$$

Proof. Set $a_n(K) = \log \kappa_n(K) + \log C$. Assumption (A2) gives $a_{m+n}(K) \leq a_m(K) + a_n(K)$, so $\{a_n(K)\}_{n\geq 1}$ is subadditive. By the classical subadditive lemma, $\lim_{n\to\infty} a_n(K)/n$ exists and equals $\inf_{n\geq 1} a_n(K)/n$. Subtracting $\log C/n$ (which vanishes as $n \to \infty$) yields the claim for $\lambda(K) = \lim_{n\to\infty} \frac{1}{n} \log \kappa_n(K)$. □

5.3 Convex Upper Bound on Asymptotic Growth

Theorem 1 (Convex upper bound)**.** *Assume* (A1)–(A3)*. Let K_A, K_B be admissible kernels and set*

$$K_t := (1-t)K_A + tK_B, \qquad t \in [0,1].$$

Then

$$\lambda(K_t) \;\le\; \max\{\lambda(K_A), \lambda(K_B)\}. \tag{4}$$

Proof. Fix $t \in [0,1]$ and $n \ge 1$. By (2),

$$\kappa_n(K_t) \le (1-t)\kappa_n(K_A) + t\kappa_n(K_B) \le \max\bigl\{\kappa_n(K_A), \kappa_n(K_B)\bigr\}.$$

Taking logarithms and dividing by n,

$$\frac{1}{n}\log \kappa_n(K_t) \;\le\; \max\left\{\frac{1}{n}\log \kappa_n(K_A),\; \frac{1}{n}\log \kappa_n(K_B)\right\}.$$

Taking $n \to \infty$ (using Lemma 1) yields (4). □

Interpretation. Theorem 1 shows that convex mixing cannot produce asymptotic amplification beyond the larger endpoint growth rate. In particular, no "free sign reversal" can be created by mixing two kernels that both compound.

5.4 Interior Sign Preservation Under Uniform Non-Collapse

Convexity alone does not bound $\lambda(K_t)$ away from $-\infty$; mixing can reduce the growth rate. The following theorem establishes a positive lower bound under an additional regularity hypothesis that is separate from, and not implied by, (A1)–(A3).

Theorem 2 (Interior asymptotic sign preservation)**.** *Assume* (A1)–(A3)*. Let Ω be a convex subset of the interior of the kernel simplex. Suppose:*

(i) (Uniform non-collapse) $\inf_{K\in\Omega} \kappa_n(K) > 0$ *for each fixed n.*

(ii) (Positive endpoint growth) $\lambda(K_i) > 0$ *for each extreme kernel K_i generating Ω.*

Then $\lambda(K) > 0$ for all $K \in \Omega$.

Proof. Fix $K \in \Omega$. By hypothesis (ii) and Lemma 1, there exists $\varepsilon > 0$ such that $\kappa_n(K_i) \ge e^{\varepsilon n - \log C}$ for each extreme kernel K_i and all sufficiently large n. Since K lies in the convex hull of the extreme kernels, convexity (2) gives

$$\kappa_n(K) \;\ge\; \frac{1}{M}\sum_{i=1}^{M} \kappa_n(K_i) \;\ge\; \frac{e^{\varepsilon n - \log C}}{M},$$

where M is the number of extreme kernels. Therefore $\lambda(K) \ge \varepsilon > 0$. □

What is proved and what is not. Theorem 1 is unconditional under (A1)–(A3). Theorem 2 is conditional on hypothesis (i) (uniform non-collapse on Ω), a substrate-specific regularity condition not derivable from the axioms. The empirical chambers do not verify (i) directly; they observe that the certified interior stratum (cells passing $L_2 > 5\sigma_F$) exhibits positive curvature, which is consistent with (i) holding on that stratum. The boundary stratum, where (i) fails empirically (elevated σ_F prevents certification), is precisely where Theorem 2 makes no claim.

6 Partial Monotonicity Under Submultiplicativity

We now prove a monotonicity statement that is useful even without full sign-preservation: submultiplicativity forces the depth-to-depth log-increments to stabilize, and under mild regularity this implies one-sided monotonic behavior of the finite-depth slope statistic.

Define the depth increment:

$$\Delta_n(K) := \log \kappa_{n+1}(K) - \log \kappa_n(K).$$

Lemma 2 (Increment upper bound from submultiplicativity)**.** *Assume (A2). Then for all $n \geq 1$,*

$$\Delta_n(K) \leq \log \kappa_1(K) + \log C. \tag{5}$$

Proof. Apply (1) with $m = n$ and $n = 1$:

$$\kappa_{n+1}(K) \leq C\,\kappa_n(K)\kappa_1(K).$$

Take logs and subtract $\log \kappa_n(K)$ to obtain (5). □

Lemma 3 (Asymptotic monotonic stabilization of averaged increments)**.** *Assume (A2). Then the averaged increment*

$$\bar{\Delta}_N(K) := \frac{1}{N}\sum_{n=1}^{N} \Delta_n(K) = \frac{1}{N}\big(\log \kappa_{N+1}(K) - \log \kappa_1(K)\big)$$

converges to $\lambda(K)$ as $N \to \infty$.

Proof. This is a telescoping identity plus the definition of $\lambda(K)$:

$$\bar{\Delta}_N(K) = \frac{1}{N}\log \kappa_{N+1}(K) - \frac{1}{N}\log \kappa_1(K) \to \lambda(K).$$

□

Connection to the chamber statistic b. Your chamber's b is a finite-depth proxy for an averaged increment of $\log \kappa_n$ (implemented via the log-slope regression across depths). The lemma implies: if $\lambda(K)$ is strictly positive (or strictly negative) and the variance gates certify non-degeneracy, then b will converge to the sign of $\lambda(K)$ as depths increase.

Proposition 1 (One-sided sign monotonicity on a certified interior stratum)**.** *Assume (A1)–(A2). Fix a certified interior region Ω of kernels such that:*

- *(Non-degeneracy) $L_2(K) > 5\sigma_F(K)$ for all $K \in \Omega$,*
- *(Uniform variance control) the bootstrap CI width of $b(K)$ is uniformly bounded on Ω.*

If $\lambda(K) \geq \lambda_0 > 0$ for all $K \in \Omega$, then there exists a depth cutoff n_0 such that for all $K \in \Omega$ the finite-depth slope statistic computed over depths $\geq n_0$ is strictly positive (hence no interior certified sign reversal).

Proof sketch. From the previous lemma, the averaged increments converge to $\lambda(K) \geq \lambda_0$. Uniform CI control plus non-degeneracy allows swapping the asymptotic limit with certified finite-depth estimates: beyond some depth window, the estimator concentrates around $\lambda(K)$, yielding positive slope with high confidence uniformly on Ω. □

Why this matches LVIX/LX exactly. LVIX/LX show: the only region where negative point estimates appear lies in a corner with inflated σ_F and huge CI widths, which violates the uniform variance control needed above. Hence: no CERT_NEG and no brackets is consistent with the theorem attempt; it is not contradictory evidence.

7 Calibration-Floor Obstruction Principle

Proposition 2 (Calibration-Floor Obstruction)**.** *Let $K(\beta, \gamma)$ lie in a boundary stratum where*

$$\sigma_F(K) \textit{ is locally elevated.}$$

Even if $b(K) < 0$ in point estimate, if $L_2(K) \leq 5\sigma_F(K)$, then K is structurally non-certifiable. Thus absence of CERT_NEG does not imply absence of negative curvature.

This explains the LVIX+LX null result mechanistically.

8 Gain Scaling and Basis Swap

Let g scale K_1, K_2 weights prior to convex renormalization.

Empirically:

- Basis swaps (B0,B1,B2) preserve the geometry of negative region.
- Gain scaling $g \leq 1.5$ increases positive certification counts slightly.
- Negative-curvature region remains variance-dominated in all slabs.

Thus:

Convex geometry, not basis orientation, determines certification structure.

9 Geometric Consequence

Non-degenerate cells cluster away from the high-$(\beta + \gamma)$ boundary.

The negative-curvature manifold lies precisely in the high-noise corner.

Therefore:

Recursive curvature sign structure is stratified by variance geometry.

10 Implications

The chamber pipeline establishes:

1. Recursive compounding is statistically robust for interior kernels.
2. Convex recombination does not generate interior sign reversal.
3. Negative curvature is confined to boundary strata with elevated calibration floor.
4. Stability, if intrinsic, cannot emerge via convex mixture of this basis.

11 Relation to Broader Theory

The observed phenomenon aligns with:

- Monotonicity of submultiplicative growth functionals.
- Convex preservation principles in operator theory.
- Boundary-layer instability in dynamical systems.
- Variance amplification near degenerate strata.

12 Methodological Contribution

Beyond the specific existence results, the five-chamber arc constitutes a reusable methodological pipeline for studying recursive operator families in any setting where depth-indexed curvature functionals are measurable. We articulate five distinct methodological contributions.

1. Certified partition invariance. Each chamber verifies, prior to any statistical analysis, that the recursive operator family preserves an underlying admissibility partition across depth increments. This is not a statistical test but a structural identity check: the partition geometry that makes downstream certification meaningful must be stable under the operations being studied. Chambers that fail partition invariance are invalid and produce no epistemic output. All five chambers in this program passed, confirming that the measured curvature statistics reflect genuine recursive dynamics rather than partition drift.

2. Per-cell degeneracy filtering. Certification is applied at the resolution of individual parameter cells, not at the level of aggregate statistics. Each cell in the simplex grid receives its own calibration floor $\sigma_F(\theta)$, estimated from null runs at that θ, and must independently satisfy the non-degeneracy gate $L_2(\theta) > 5\sigma_F(\theta)$. This prevents a single high-variance region from contaminating the certification of low-variance cells, and makes the certified map a spatially resolved characterization of the operator landscape rather than a global pass/fail. The practical consequence is visible in the LX cross-slab matrix: $\approx 96\%$ of the 1,386 evaluated cells are correctly classified as DEGENERATE rather than misattributed to a structural null.

3. Bracket-based sign localization. Sign reversal — the central object of interest for intrinsic stability — is not searched by exhaustive point certification but by *bracket detection*: identifying adjacent cells with opposite certified signs (CERT_POS next to CERT_NEG in the 4-neighborhood of the grid). A bracket is a necessary condition for a certified sign-change interior, and its detection triggers bisection refinement (up to $I_{\max} = 12$ iterations) that localizes θ^* to a resolution of $\Delta\theta \approx 0.001$. The absence of any bracket in LVIX and LX is therefore a structurally strong null result: it is not that negative curvature is absent from the geometry, but that no adjacent cell boundary between certified positive and certified negative regions was found anywhere in 1,386 evaluated cells across six (B_i, g) configurations.

4. Variance-calibrated obstruction detection. The distinction between *absent* negative curvature and *non-certifiable* negative curvature is central to the scientific interpretation. This pipeline quantifies the distinction explicitly: for each cell that fails certification, the reason is recorded — whether the point estimate of b is small (genuinely near-zero curvature) or whether σ_F is elevated (structurally non-certifiable despite a meaningful signal). Chamber LX diagnoses that the negative-curvature cells at $(\beta, \gamma) \approx (0.55, 0.45)$ exhibit $b \approx -0.31$ but $\sigma_F \approx 0.22$, placing them firmly in the obstruction regime. The pipeline thus converts a null result into a diagnostic finding about the statistical topology of the operator simplex, which is itself actionable information for protocol design (Paths 1–3 in the LX chamber summary).

5. Basis-robust topology verification. A single kernel assignment could produce certification artifacts specific to that assignment. Chamber LX addresses this by testing three basis permutations (B_0, B_1, B_2) that cover the distinct roles of the three kernel families. Findings that are stable across all three permutations are attributable to the convex geometry of the operator simplex rather than to the indexing convention. The CERT_POS topology (6–7 cells per slab, localized in the low-$(\beta+\gamma)$ interior) and the CERT_NEG absence are both basis-robust, as confirmed by the cross-slab matrix. Basis robustness transforms a result from "possible artifact of a single kernel ordering" to a genuine property of the convex hull, which is the claim required for the Convex Sign Preservation Conjecture to have empirical support.

Taken together, these five contributions define a protocol for recursive operator analysis that is modular, spatially resolved, falsifiable at each step, and capable of distinguishing structural absence from statistical obstruction.

13 Structural Analogies and Potential Cross-Domain Relevance

The following section collects structural analogies only. Each item describes a heuristic correspondence between the phenomena observed in Chambers LVIX and LX and structures appearing in other mathematical domains. None of these correspondences has been formally derived from the chamber results. They are offered as candidate directions for future investigation, not as consequences of the theorems above.

To fix the empirical ground for what follows, we recall three concrete data points from the certified runs:

- *Interior CERT_POS cell (LVIX):* $\bar{b} = 0.1164$, $\sigma_F = 0.1091$, threshold $5\sigma_F = 0.5452$, $L_2^{\mathrm{med}} = 0.5714$, ratio $L_2/(5\sigma_F) = 1.048 > 1$, certified sign $= +$. Positive curvature is certifiable when variance is controlled.
- *Near-boundary SATURATING cell (LVII):* $\bar{b} = 0.0345$, $\mathrm{CI}_{99,\mathrm{lo}} = -0.0129$, $\sigma_F = 0.1066$, threshold $5\sigma_F = 0.5329$, non-degenerate = true, decision = SATURATING_OR_BOUNDED. Even modest slopes are filtered by CI logic independent of the sign gate.
- *Boundary-stratum degenerate cell (LVIX):* $\sigma_F = 0.3058$, threshold $5\sigma_F = 1.529$, $L_2^{\mathrm{med}} = 0.452$, non-degeneracy gate = false. Elevated σ_F raises the 5σ barrier beyond attainable L_2; this is the empirical ground of the calibration-floor obstruction.

What the chambers *actually* establish is therefore:

1. Convex mixtures of the tested kernel families preserve interior curvature polarity under variance control.

2. Negative curvature localizes in high-variance boundary strata where certification is structurally obstructed.

3. Certification failure is geometric–statistical in origin, not algebraic.

4. The obstruction is basis-robust (stable across B_0, B_1, B_2) and gain-insensitive up to $g \leq 1.5$.

The analogies below are heuristic readings of (1)–(4) into other settings.

Structural analogy: convex ensemble methods. *(Conjectural transfer — not derived from the chamber results.)* The convex rigidity phenomenon observed here — interior mixtures bounded in curvature polarity by their extreme kernels — resembles the behavior of ensemble averaging in machine learning, where convex combinations of classifiers or regressors tend to preserve the dominant structural property (e.g., bias direction, margin sign) of the component learners rather than reversing it. Similarly, boundary-localized variance amplification echoes the well-known collapse of signal-to-noise ratio near the boundary of the training distribution's support in high-dimensional kernel methods, where sign-definite estimation can become infeasible not because the signal is absent but because the effective noise floor rises beyond the signal's magnitude. These are structural analogies; no formal derivation linking ensemble theory to the submultiplicative curvature framework is offered here.

Structural analogy: discrete flow stability. *(Heuristic correspondence — not derived from the chamber results.)* Submultiplicativity of the curvature functional (Axiom A2) is structurally reminiscent of Lyapunov growth conditions in discrete-time dynamical systems, and the asymptotic exponent $\lambda(K)$ (equation (3)) plays a role analogous to that of a Lyapunov exponent. Under this reading, the absence of an intrinsic stable fixed point in the convex hull (Chambers LVIII–LX) resembles the absence of an autonomous stable equilibrium when all extreme system parameters yield positive Lyapunov exponents, while the damped stability found in Chamber LVII ($\alpha^* = 0.5$ for V6×V7) resembles dissipation-stabilized trajectories that require explicit regulation to remain bounded. The high-variance boundary stratum may correspond loosely to the region near a saddle point where the local linearization becomes ill-conditioned and reliable stability classification requires more than finite-depth sampling. These are *heuristic correspondences*: the chamber framework does not produce a Lyapunov function, does not verify contraction conditions in a dynamical systems sense, and does not establish any formal equivalence with discrete-flow stability theory.

Structural analogy: operator convexity in quantum channels. *(Conceptual parallel — not derived from the chamber results.)* Completely positive trace-preserving maps form a convex set, and contraction coefficients under channel composition are known to be submultiplicative, paralleling Axiom A2. The question of whether convex mixing of channels can reverse the sign of a capacity or contraction functional has a structural resemblance to the Convex Sign Preservation Conjecture of this paper. Channels near the boundary of the channel simplex (e.g., near the completely depolarizing channel) are known to exhibit inflated variance in capacity estimates due to degeneracy of the output state — a situation that may parallel the elevated σ_F observed in the boundary stratum of the LVIX/LX simplex. The non-degeneracy gate $L_2 > 5\sigma_F$ is conceptually reminiscent of conditions that exclude near-rank-one or depolarizing channels from finite-sample capacity analysis. No quantum channel computation has been performed; this is a conceptual parallel offered as a candidate direction for formalization, not a derived result.

The core finding that warrants cross-domain attention is not the analogy itself but the methodological distinction it rests on: the obstruction is *geometric–statistical*, not algebraic. Negative curvature is present in point estimate ($\bar{b} \approx -0.31$ at $(\beta, \gamma) \approx (0.55, 0.45)$) but cannot be certified because σ_F is structurally elevated in that region, raising the 5σ barrier beyond attainable L_2. Whether a parallel mechanism — signal present but statistically non-certifiable due to boundary-localized variance inflation — operates in the above domains is an open question that would require independent analysis in each setting.

14 Conclusion

The five-chamber program (LVI–LX) yields a layered structural result. Chamber LVI establishes that the three operator families occupy distinct growth regimes (V2×V3: compounding; V2×V4: trivial; V6×V7: inconclusive). Chamber LVII provides the only positive existence result: damped stabilization is achievable for V6×V7 at $\alpha^* = 0.5$, establishing that stability can be attained via explicit depth regulation, while compounding in V2×V3 is curvature-intrinsic and not suppressible by damping alone. Chambers LVIII–LX deliver a convergent structural null: no intrinsic stable operator exists in the K_A/K_B convex hull (LVIII), in the 2D kernel simplex under the canonical B_0 basis (LVIX), or across three basis permutations and two gain levels spanning 1,386 evaluated cells (LX).

The unifying conclusion is:

> Convex kernel recombination and moderate gain scaling preserve curvature polarity in the interior non-degenerate region of the operator simplex. Sign reversal is confined to a boundary stratum with structurally elevated calibration floor, where certification is obstructed by variance rather than by the absence of signal.

The family hierarchy (V2×V3 $\succ$ V6×V7 $\succ$ V2×V4 in compounding order) is basis-robust and persistent across all five chambers. The distinction between intrinsic stability (absent throughout) and damped stability (achievable in LVII) provides the key constraint on future protocol design: stability search must move beyond the convex hull of the current kernel basis, or must operate explicitly in the damped regime.

The methodological pipeline — certified partition invariance, per-cell degeneracy filtering, bracket-based sign localization, variance-calibrated obstruction detection, and basis-robust topology verification — is offered as a reusable protocol for recursive operator analysis. Whether the geometric–statistical obstruction mechanism identified here has formal counterparts in ensemble learning, discrete flow stability, or quantum channel theory is an open question; the structural analogies sketched in Section 13 are heuristic correspondences, not derived results.