Lyapunov Stability of Structural Density and Spectral Gap in a Sheaf-Governed Graph Dynamical System with Cohomological Feedback

Jason Volk
Invariant Research
jason@invariant.pro
https://invariant.pro

May 2026

Abstract

We prove conditional Lyapunov stability results for a sheaf-governed graph dynamical system whose topology evolves under cohomological feedback. The first result shows that the structural density $r(t) = |E(t)|/|V(t)|$ converges to a system-specific equilibrium r^* (approximately 3.0 for the dynamics studied here) under local restorative drift and bounded-update assumptions. The second result shows that the sheaf-Laplacian spectral gap $\lambda_2(L_\mathcal{F}(t))$ converges to a positive equilibrium $\lambda_2^* = 0.530 > 0$ under empirically verified stationary spectral drift and bounded perturbation assumptions. Both results are stated as conditional stability theorems rather than universal guarantees for arbitrary sheaves. Using the Borkar stochastic approximation framework [Borkar, 2008, 2026], we show both quantities converge to $O(1/n)$-neighborhoods of their attractors. The coupled result yields a sheaf-expansion result for the restricted class of Purity-Gate-certified near-identity sheaves, avoiding the scope of the First-Kaufman impossibility for arbitrary sheaves [First and Kaufman, 2024] by restricting attention to Purity-Gate-certified sheaves. We further introduce the Coherence Margin, a spectral-topological leading indicator derived from the coupled stability that forecasts every major H^1 spike with 1–70 cycle lead time; the Coherence Margin remains positive throughout the 350-cycle attractor window, indicating that the system's measured contradiction-resolution capacity exceeds its contradiction-creation pressure at all times. These results provide a formal stability analysis for autonomously growing sheaf-governed graph state under cohomological feedback.

We validate both theorems on a 500-cycle run (seed 42, ARPACK-corrected eigensolver) starting from a 150-vertex seed and growing to 515 vertices. The ratio $r(t)$ locks to 3.0 by cycle 150 and remains there for 350 consecutive cycles. The spectral gap λ_2 converges to 0.530 with ADF-confirmed stationarity ($p = 0.016$). Multi-seed validation across four seeds (42, 72, 137, 271) at 300 cycles each confirms seed-invariant convergence with cross-seed coefficient of variation CV $= 1.9\%$. We further validate the architecture at scale via cellular decomposition, demonstrating constant per-vertex eigensolve cost (3.55 ms/vertex) from 50,000 to 1,000,000 vertices on commodity hardware. We also report a streaming incremental update architecture with lazy evaluation that achieves edit-path $O(1)$ cost with deferred exact reconciliation, with measured scaling exponent 0.19 ($R^2 = 0.975$) across $V = 21$K to $V = 1$M: 63 microseconds per edit, 13 microseconds per query, zero measured post-flush correctness drift across 8 seeds. To our knowledge, this is one of the largest reported sheaf-based structural verification computations of which we are aware.

1 Introduction

Self-modifying dynamical systems (systems whose state space changes under the dynamics) resist standard stability analysis. In classical settings, the state space is fixed and Lyapunov functions certify convergence within that space. When the space itself grows, the standard framework requires extension.

We study a concrete instance: a finite graph $G(t)$ equipped with a cellular sheaf $\mathcal{F}(t)$, evolving under two competing rate processes. Edge addition enlarges the graph at a bounded average rate. Vertex expansion fires when a sheaf-derived pressure signal exceeds a threshold. The competition between these processes determines the long-run edge-vertex ratio $r(t) = |E(t)|/|V(t)|$.

Our main results are two conditional stability propositions and a predictive observable. Theorem 1 shows that $r(t)$ converges to a stable equilibrium, proved by constructing a Lyapunov function $V(r) = (r - r^*)^2$ and showing it decreases outside a neighborhood of r^*. The proof is conditional on four explicitly stated assumptions (A1–A4) regarding restorative drift and rate balance; Assumption A2 (local restorative density drift) is enforced above the equilibrium by the density-aware governor and empirically verified across the observed attractor neighborhood. Theorem 2 shows that the spectral gap λ_2 of the sheaf Laplacian converges to a positive equilibrium $\lambda_2^* > 0$ via a parallel Lyapunov argument with an eigenvalue perturbation lemma. The proofs leverage the stochastic approximation framework of Borkar [2008] and Robbins–Monro [Robbins and Monro, 1951], showing that the discrete-time trajectories asymptotically track continuous-time ODEs whose unique stable equilibria are r^* and λ_2^* respectively. The coupled result (Corollary, Section 5.4) establishes simultaneous structural-spectral stability for the restricted class of Purity-Gate-certified sheaves, not for arbitrary sheaves. The Coherence Margin, a real-time observable combining the Dirichlet energy residual with the spectral gap, provides 1–70 cycle advance warning of every major H^1 spike, converting the conditional stability results into an operational predictive signal.

Existing Lyapunov results in learning systems [Richards et al., 2018, Revay et al., 2024, Banakar, 2011] operate on fixed-topology architectures. Growing graph models [Barabási and Albert, 1999, Dorogovtsev and Mendes, 2002] study structural evolution without Lyapunov stability guarantees for structural ratios. Adaptive coevolutionary network models [Gross and Blasius, 2008] address coupled dynamics but typically without sheaf structure or formal convergence results for density ratios. Sheaf neural networks [Hansen and Gebhart, 2020, Bodnar et al., 2022, Barbero et al., 2022, Braithwaite et al., 2024b] and sheaf-theoretic knowledge representations [Gebhart et al., 2023] operate on fixed graphs; recent extensions include spectral filtering approaches to sheaf diffusion [Borgi et al., 2025], heterogeneous sheaf structures [Braithwaite et al., 2024a], learned restriction maps via Riemannian geometry [Barbero et al., 2022], and analyses of when learnable sheaf Laplacians are necessary versus sufficient [Hernandez Caralt et al., 2026]. The question of how sheaf Laplacian spectra evolve under structural modification, the central concern of the present work, has not been addressed in this literature. Separately, Di Nino et al. [2025] proposed learning the sheaf Laplacian itself (including topology and restriction maps) from observed data by optimizing restriction maps in closed form; our system inverts this relationship, using the sheaf Laplacian's spectral properties to govern the graph's autonomous structural evolution. Battiloro et al. [2024], Riess and Hale [2025] extended sheaf-theoretic methods to tangent bundle convolutional learning and distributed multi-agent coordination over cellular sheaves, establishing connections between sheaf Laplacians and control-theoretic frameworks that parallel our use of Lyapunov stability for sheaf-governed dynamics. First and Kaufman [2022, 2024] proved that no universal Cheeger inequality extends from graphs to arbitrary sheaves, but established positive results for restricted classes. We are not aware of prior Lyapunov analysis for a sheaf-governed graph-growth system that combines spectral obstruction feedback, convergence of an edge-vertex ratio, and convergence of a sheaf Laplacian spectral gap.

2 Setup and Notation

Let $G(t) = (V(t), E(t))$ be a finite graph at discrete time t, with $|V(t)| = n(t)$ and $|E(t)| = m(t)$. Define $r(t) = m(t)/n(t)$.

2.1 Cellular Sheaf Structure

A cellular sheaf $\mathcal{F}(t)$ over $G(t)$ assigns to each vertex v a stalk $\mathcal{F}(v) = \mathbb{R}^d$ (d fixed) and to each edge $e = (u, v)$ restriction maps $\rho_{u\to e}, \rho_{v\to e}$ ($d \times d$ matrices).

A global section $s(t) \in C^0(G(t); \mathcal{F}(t))$ assigns $s(v) \in \mathcal{F}(v)$ to each vertex.

2.2 Coboundary Operator and Laplacian

The coboundary operator $\delta_0 : C^0 \to C^1$:

$$(\delta_0 s)(e) = \rho_{v\to e}(s(v)) - \rho_{u\to e}(s(u)) \tag{1}$$

The 0-sheaf Laplacian:

$$L_{\mathcal{F}} = \delta_0^\top \delta_0 \quad \text{(positive semidefinite)} \tag{2}$$

Its kernel $\ker(L_{\mathcal{F}}) = \ker(\delta_0) = H^0(G; \mathcal{F})$, the space of global sections. This is *not* H^1.

2.3 Dirichlet Energy

$$\mathcal{E}[s] = \|\delta_0(s)\|^2 = s^\top L_{\mathcal{F}} s \tag{3}$$

Per-edge: $\mathcal{E}_e[s] = \|\rho_{v\to e}(s(v)) - \rho_{u\to e}(s(u))\|^2$.

2.4 Sheaf Cohomology

For a graph (1-dimensional CW complex):

$$H^0(G; \mathcal{F}) = \ker(\delta_0) \tag{4}$$

$$H^1(G; \mathcal{F}) = C^1 / \operatorname{im}(\delta_0) \tag{5}$$

By rank-nullity:

$$\dim H^1 = \dim C^1 - \dim C^0 + \dim H^0 = d \cdot (m - n) + \dim H^0 \tag{6}$$

Equivalently, $\dim H^1 = \dim \ker(\delta_0 \delta_0^\top)$ (the 1-Laplacian kernel).

2.5 Spectral Obstruction Proxy

Our implementation uses a proxy:

$$\mu(t) = |\{\text{eigenvalues of } L_{\mathcal{F}} \text{ below tolerance}\}| - \beta_0 \tag{7}$$

This counts excess near-zero eigenvalues of the 0-Laplacian beyond those explained by connected components. This is *not* $\dim H^1$ from equation (6). It is a coarser spectral proxy that correlates with the number of approximately frustrated graph cycles. See Appendix A for detailed discussion.

The stability argument does not require $\mu = \dim H^1$. It requires that the Dirichlet energy $\mathcal{E}[s(t)]$ be positively correlated with the edge-vertex ratio $r(t)$, which is mediated through the graph-theoretic Betti number (Section 2.6), not through any cohomology dimension.

2.6 Graph-Theoretic Quantities

$$\beta_1(t) = m(t) - n(t) + \beta_0(t) \quad \text{(first Betti number)} \tag{8}$$

We have $\dim H^1(G; \mathcal{F}) \le d \cdot \beta_1$.

3 The Dynamical System

At each discrete step $t \to t+1$, two additive processes modify $G(t)$.

3.1 Edge and Vertex Creation

Each cycle t contributes increments $\Delta_m(t) \geq 0$ (new edges) and $\Delta_n(t) \in \{0, 1\}$ (new vertex). Edges come from two sources: periodic co-observation of existing vertex pairs (subject to a per-batch cap and duplicate rejection) and neighbor connections for newly added vertices. The per-cycle totals $\Delta_m(t)$ and $\Delta_n(t)$ are the operationally observable quantities.

We do not decompose Δ_m into sub-sources for the stability argument. The proposition requires only the aggregate rates and their dependence on r.

3.2 Feedback Mechanism

The pressure functional $\Omega(t)$ aggregates the per-vertex Dirichlet energy $\varepsilon(t) = \mathcal{E}[s(t)]/n(t)$, the spectral obstruction proxy $\mu(t)$, a spectral connectivity measure, and a rate-of-change term into a single bounded signal in $[0, 1]$. Vertex expansion fires when Ω exceeds a threshold, subject to a cooldown timer.

The key coupling: higher $r \to$ higher β_1 (Euler characteristic) $\to$ more graph cycles available for frustration $\to$ higher Dirichlet energy (more frustrated restriction maps) $\to$ higher $\Omega \to$ expansion fires more frequently $\to$ more vertices added relative to edges $\to$ r decreases. This is the negative feedback loop that stabilizes r.

3.3 Heat Flow

$$\frac{ds}{dt} = -L_{\mathcal{F}} s \tag{9}$$

drives the section toward the nearest harmonic configuration after structural modifications. This process reduces Dirichlet energy without modifying the graph structure.

3.4 Purity Gate (Algebraic Certification)

All restriction maps satisfy a bounded spectral radius condition enforced by deterministic rescaling at every construction site:

$$\sigma_{\max}(\rho) \leq 0.99 \tag{10}$$

After SVD-conditioned construction of each restriction map, the matrix is uniformly rescaled by $\min(1, 0.99/\sigma_{\max})$ if $\sigma_{\max}$ exceeds the bound. This preserves singular vector direction (and thus sheaf topology) while guaranteeing strict contractivity. The bound is satisfied by construction, not verified post-hoc. Initialization maps (scale 0.68) satisfy the bound without rescaling; growth maps (empirical $\sigma_{\max} = 1.0$ pre-enforcement) are rescaled by 1%.

This places the sheaf in the near-identity regime of First and Kaufman [2024]. The Purity Gate provides intrinsic stability: strict contractivity ($\sigma_{\max} \leq 0.99$) introduces minimum topological friction on every edge, ensuring Dirichlet energy rises monotonically with density. This naturally couples edge accumulation to vertex expansion pressure.

The system additionally tracks diagnostic indicators for projector idempotency and nilpotency. In a growing substrate with mixed-scale restriction maps (initialization scale 0.68, growth scale 0.99), the idempotency error is $O(|s_{\text{init}} - s_{\text{growth}}|)$, reflecting the dissipative cost of mixing two scale regimes. These diagnostic indicators are not load-bearing for the stability proof, which depends solely on the spectral radius bound.

3.5 Density-Aware Spectral Governor

A density-aware spectral governor provides extrinsic stability as a formal Lyapunov safety net. The Ω functional includes a linear density penalty:

$$\text{density\_penalty} = w \cdot \max\left(0,\ \frac{r(t) - r^*}{r^*}\right) \tag{11}$$

where $w = 1.5$ is a tunable weight and r^* is the target attractor. The equilibrium r^* is a property of the system's mutation dynamics (edge batch size, vertex expansion threshold, co-occurrence saturation), not a universal constant; it is measured empirically at approximately 3.0 for the dynamics studied here.

When E/V exceeds r^*, the penalty pushes Ω toward the warning threshold (0.70), triggering vertex expansion that dilutes E/V back toward r^*. The governor guarantees a restorative force with constant positive derivative $b'(r) = w/r^* = 0.5$ for $r > r^*$, satisfying Assumption A2 by construction.

The linear form is chosen over quadratic because quadratic penalty has $b'(r^*) = 0$ (degenerate fixed point, polynomial convergence), whereas linear maintains exponential convergence rate. With $w = 1.5$ and $r^* = 3.0$, the trigger point is $E/V = r^* \cdot (1 + 0.12/w) = 3.24$, providing 8% overshoot tolerance before intervention.

The system thus exhibits two-tiered stability: the Purity Gate (Section 3.4) provides intrinsic stability sufficient for some seeds, while the density governor provides the extrinsic safety net required for seed-invariant convergence across all seeds.

3.6 Update Law

$$m(t+1) = m(t) + \Delta_m(t) \tag{12}$$

$$n(t+1) = n(t) + \Delta_n(t) \tag{13}$$

$$r(t+1) = m(t+1)/n(t+1) \tag{14}$$

4 Lyapunov Function and Convergence

4.1 Equilibrium Characterization and Stochastic Framework

Define the time-averaged rates over the attractor regime (a window of N consecutive cycles where $r(t)$ remains within a bounded band):

$$\bar{m} = \frac{1}{N} \sum_t \Delta_m(t) \quad \text{(mean edges per cycle)} \tag{15}$$

$$\bar{n} = \frac{1}{N} \sum_t \Delta_n(t) \quad \text{(mean vertices per cycle)} \tag{16}$$

Over the attractor regime, the equilibrium is characterized by the stationary rate balance:

$$r^* = \bar{m}/\bar{n} \tag{17}$$

Both $\bar{m}$ and $\bar{n}$ depend on r through the feedback mechanism (Section 3.2). The equilibrium r^* is the value where the total edge creation rate equals r^* times the vertex creation rate, as measured within the stationary window.

Measured values (attractor window, cycles 150–500 of primary run, ARPACK-corrected eigensolver, seed 42): $\bar{m} = 940/350 = 2.686$ edges/cycle, $\bar{n} = 314/350 = 0.897$ vertices/cycle, $r^* = 2.686/0.897 = 2.994$. This is consistent with the observed ratio of 3.0 to within rounding precision. The near-exact match between $\bar{m}/\bar{n}$ and the instantaneous r confirms that the system is in true stationary equilibrium throughout the attractor window.

Source of expectations. The system is deterministic given a fixed seed and initial state. The expectations $\mathbb{E}[\Delta_m|r]$ and $\mathbb{E}[\Delta_n|r]$ are taken over the pseudo-random mutation process, conditioned on the current graph state. The Dirichlet energy threshold gates whether a mutation can fire; the specific mutation type and target are drawn from a pseudo-random distribution. Over the 350-cycle attractor window, these draws generate a well-defined empirical distribution. Because the Augmented Dickey–Fuller test confirms stationarity of the trajectory within the attractor window ($p = 0.016$), the ergodic interpretation is justified: the time average converges to the ensemble average.

Stochastic approximation framework. The one-step update of $r(t)$ has the canonical Robbins–Monro form:

$$r(t+1) = r(t) + \frac{1}{n(t)}\big[h(r(t)) + \xi(t)\big] \tag{18}$$

where $h(r) = \bar{m}(r) - r \cdot b(r)$ is the mean drift function and $\xi(t)$ is the deviation from the mean, bounded by $|\xi(t)| \leq C_m/n$ via assumption (A4). The step size $1/n(t)$ is decreasing because $n(t)$ is non-decreasing. By the fundamental theorem of stochastic approximation [Borkar, 2008, Chapter 2], the discrete-time trajectory $r(t)$ asymptotically tracks the solutions of the continuous-time ODE:

$$\frac{dr}{dt} = h(r) = \bar{m}(r) - r \cdot b(r) \tag{19}$$

The Lyapunov candidate $V(r) = (r - r^*)^2$ can then be differentiated continuously:

$$\frac{dV}{dt} = 2(r - r^*) \cdot h(r) \tag{20}$$

Since $h(r^*) = 0$, $h(r) < 0$ for $r > r^*$, and $h(r) > 0$ for $r < r^*$ (by the monotonicity assumptions below), we have $dV/dt < 0$ for all $r \neq r^*$. This establishes asymptotic stability of r^* in the ODE limit. The discrete-time convergence inherits this stability via the stochastic approximation tracking result.

4.2 Feedback Structure and Phase Transition

The stabilizing mechanism requires that the mean drift function $h(r) = \bar{m}(r) - r \cdot b(r)$ has restorative sign in a neighborhood of r^*: $h(r) < 0$ for $r > r^*$ and $h(r) > 0$ for $r < r^*$ (Assumption A2).

$$r > r^* \;\Rightarrow\; b(r) > b(r^*) \;\Rightarrow\; \text{more vertex creation} \;\Rightarrow\; r \text{ decreases}$$
$$r < r^* \;\Rightarrow\; b(r) < b(r^*) \;\Rightarrow\; \text{less vertex creation} \;\Rightarrow\; r \text{ increases}$$

The chain of causation: r increases $\rightarrow$ $\beta_1 = (r-1)n + 1$ increases $\rightarrow$ more graph cycles available for restriction map frustration $\rightarrow$ higher minimum-energy section residual (Dirichlet energy) $\rightarrow$ Ω exceeds expansion threshold more often $\rightarrow$ expansion frequency $b(r)$ increases $\rightarrow$ r decreases (restoring force).

Phase transition in growth dynamics. The monotonicity conditions (A1) and (A2) are local properties, holding in a neighborhood $[r^* - \epsilon, r^* + \epsilon]$ of the attractor. The system exhibits a phase transition between two regimes:

Growth phase ($r \sim 1.6$ to 2.3, cycles 0–100): $\bar{m}'(r) > 0$. The graph is sparse. Dirichlet energy is low and the energy landscape is easily traversable. Most proposed topological mutations produce valid, novel edges. Co-occurrence finds fresh vertex pairs without saturation. Vertex expansion is infrequent: vertex rate $15/50 = 0.30$/cycle.

Attractor-bound phase ($r \sim 3.0$, cycles 150–500): $\bar{m}'(r) \leq 0$. Contradiction density has increased and the restriction maps heavily veto new edges via duplicate rejection and co-occurrence capping. The graph has transitioned from an exploratory state to a regulated equilibrium. Vertex expansion is frequent: vertex rate $314/350 = 0.90$/cycle.

The transition occurs as r crosses the saturation boundary where the sign of $\bar{m}'(r)$ flips. This is the point at which the negative feedback loop engages.

Verification from trajectory data: at $r \sim 2.3$–2.8 (cycles 50–100, below attractor), vertex rate is $15/50 = 0.30$/cycle with expansion infrequent and edge creation dominating. At $r \sim 3.0$ (cycles 150–500, in attractor), vertex rate is $314/350 = 0.90$/cycle with rates balanced. The monotonicity is verified across the observed range $[2.3, 3.0]$.

4.3 Theorem 1 (Stability of the Edge-Vertex Ratio)

Let $G(t)$ be a growing graph governed by the dynamics of Section 3. Assume:

Assumption 1 (A1: Bounded edge rate, monotone saturation)**.** There exists $M > 0$ and a Lipschitz function $\bar{m}(r) \in [0, M]$ with $\bar{m}'(r) \leq 0$ for r in a neighborhood of r^*. (Due to duplicate-edge saturation: denser graphs reject more proposed edges.) *Verified:* edge saturation engages above $r \approx 2.8$.

Assumption 2 (A2: Local restorative density drift)**.** There exist constants $\epsilon > 0$ and $\kappa > 0$ such that, for all r with $0 < |r - r^*| < \epsilon$, the mean drift

$$h(r) = \bar{m}(r) - rb(r)$$

satisfies

$$(r - r^*)h(r) \leq -\kappa(r - r^*)^2.$$

Equivalently, the expected update points toward r^* on both sides of the equilibrium. For $r > r^*$, the density-aware governor supplies the restorative component through the linear penalty in Section 3.5. For $r < r^*$, recovery is supplied by reduced expansion frequency together with continued edge creation in the unsaturated regime. *Enforced above r^* by the density-aware governor; empirically verified across the observed attractor neighborhood.*

Assumption 3 (A3: Stationary rate balance and uniqueness)**.** There exists unique $r^* > 0$ solving $r^* = \bar{m}(r^*)/b(r^*)$. *Measured:* $\bar{m}/\bar{n} = 2.686/0.897 = 2.994$ (approximately 3.0).

Assumption 4 (A4: Small single-step updates)**.** For all t, $\Delta_m(t) \leq C_m$, $\Delta_n(t) \leq 1$, and $n(t)$ large enough so single-step change in r is $O(1/n)$. *Verified:* $C_m = 12$. With $n \geq 201$ in the attractor window, the maximum single-step ratio change is $|\Delta r|_{\max} \leq 12/201 = 0.060$.

Remark 1 (Empirical regression functions). The functions $\bar{m}(r)$ and $b(r)$ are defined as conditional expectations of per-cycle increments. Their Lipschitz continuity follows from the smoothness of the energy-threshold gating function and the bounded variation of the mutation distribution. Individual cycles produce discrete, integer-batch increments; the smooth dependence on r emerges in the conditional mean over many cycles. In the Borkar stochastic approximation framework, the drift function $h(r)$ need only be Lipschitz. The discrete nature of the mutations is fully absorbed by the bounded martingale difference noise term $\xi(t)$, whose magnitude is bounded by C_m/n via (A4).

Theorem 1 (Structural Density Stability)**.** *Under Assumptions A1–A4, the edge-vertex ratio $r(t)$ converges to an $O(1/n)$-neighborhood of r^*.*

Proof. We verify the conditions and construct the Lyapunov function.

(i) $V(r) = (r - r^*)^2 \geq 0$ with $V(r^*) = 0$.

(ii) One-step drift derivation. Starting from the update law $r(t{+}1) = (m + \Delta_m)/(n + \Delta_n)$, factor out n:

$$r(t{+}1) = \frac{r(t) + \Delta_m/n}{1 + \Delta_n/n}$$

By (A4), $\Delta_n/n = O(1/n)$, so applying the geometric series expansion $1/(1+x) = 1 - x + O(x^2)$:

$$r(t{+}1) = \big[r(t) + \Delta_m/n\big]\big[1 - \Delta_n/n + O(1/n^2)\big] = r(t) + \frac{\Delta_m - r(t)\Delta_n}{n} + O(1/n^2)$$

Taking conditional expectations given $r(t) = r$:

$$\mathbb{E}[r(t{+}1) - r(t) \,|\, r(t) = r] = \frac{\bar{m}(r) - r \cdot b(r)}{n} + O(1/n^2)$$

Define the drift function $h(r) = \bar{m}(r) - r \cdot b(r)$. At the equilibrium r^*, $h(r^*) = 0$ by (A3): $\bar{m}(r^*) = r^* \cdot b(r^*)$.

For $r > r^*$: $\bar{m}(r) \leq \bar{m}(r^*)$ by (A1) and $b(r) > b(r^*)$ by (A2). Therefore $\bar{m}(r) - r \cdot b(r) < \bar{m}(r^*) - r^* \cdot b(r^*) = 0$, giving $h(r) < 0$. The expected drift is negative.

For $r < r^*$: $\bar{m}(r) \geq \bar{m}(r^*)$ by (A1) and $b(r) < b(r^*)$ by (A2). Therefore $h(r) > 0$. The expected drift is positive.

This establishes that $h(r) < 0$ for $r > r^*$ and $h(r) > 0$ for $r < r^*$, so r^* is a globally attracting fixed point of the ODE $dr/dt = h(r)$ in the neighborhood where (A1)–(A2) hold.

(iii) For $r(t) > r^*$: $b(r) > b(r^*)$ by (A2), so vertex creation is elevated. Meanwhile $\bar{m}(r) \leq \bar{m}(r^*)$ by (A1). By the drift derivation, $\mathbb{E}[r(t{+}1) - r(t)] < 0$. The per-cycle ratio update tends to decrease r because the vertex denominator grows faster than the edge numerator relative to equilibrium. $\Delta V < 0$.

(iv) For $r(t) < r^*$: $b(r) < b(r^*)$ by (A2), so vertex creation is suppressed. Edge creation continues at rate $\bar{m}(r) \geq \bar{m}(r^*)$ by (A1). On cycles without expansion ($\Delta_n = 0$), r increases by $\Delta_m/n > 0$. On the infrequent expansion cycles, the ratio change is small relative to the accumulation from non-expansion cycles. Net effect: r increases. $\Delta V < 0$.

(v) Attractor width and topological latency. The attractor width $\delta(n)$ has two components: $\delta(n) = \delta_{\text{instant}} + \delta_{\text{latency}}$.

The instantaneous perturbation bound is:

$$\delta_{\text{instant}} \sim \frac{C_m}{n \cdot |h'(r^*)|}$$

where $|h'(r^*)|$ is the slope of the drift function at equilibrium. From trajectory data (vertex rate 0.30/cycle at $r \sim 2.5$ versus 0.90/cycle at $r = 3.0$), $|h'(r^*)| \sim 1.2$, giving $\delta_{\text{instant}} \sim 12/(400 \times 1.2) \sim 0.025$.

The topological latency component accounts for the delay between structural perturbation and governor response. When edge additions increase β_1, the new independent cycles must propagate through the restriction map landscape before registering as elevated H^1 in the spectral proxy. This topological latency τ ranges from 1 to 70 cycles in the experimental data (13 distinct spike events with $H^1 > 50$). During this latency window, $r(t)$ continues to drift, producing overshoot beyond the instantaneous bound:

$$\delta_{\text{latency}} \sim \tau \cdot |\bar{m}(r^*) - r^* \cdot b(r^*)|/n$$

Since $h(r^*) = 0$, the latency drift is proportional to the second-order deviation from equilibrium and is also $O(1/n)$.

The observed macroscopic attractor width is $r \in [2.9, 3.1]$, or ± 0.1 around $r^* = 3.0$. This is approximately $4\times$ the instantaneous bound, consistent with the topological latency explanation: the system overshoots before the governor snaps it back, producing damped oscillation around the attractor.

Both components are $O(1/n)$, so $\delta(n) \to 0$ as $n \to \infty$. Single-step updates modify r by $O(1/n)$, and the restoring force is bounded below by a function of $|r - r^*|$, so the ratio is trapped in ever-narrower bands as the graph grows.

Combining (iii) and (iv): $V(r(t))$ is decreasing outside the δ-neighborhood of r^*, and $r(t)$ converges to $[r^* - \delta, r^* + \delta]$.

Via the stochastic approximation tracking theorem [Borkar, 2008, Theorem 2.1], the discrete trajectory $r(t)$ converges to the equilibrium r^* of the ODE $dr/dt = h(r)$ with probability 1. □

Remark 2 (Status of assumptions). Assumption A2 (local restorative density drift) is enforced above the equilibrium by the density-aware spectral governor (Section 3.5) and empirically verified across the observed attractor neighborhood. The remaining assumptions (A1, A3, A4) are verified empirically on the primary trajectory. Assumption B1 (restorative spectral drift, Theorem 2) remains empirically verified but not derived from first principles; deriving it requires analyzing the chain $r \to \beta_1 \to$ frustration $\to$ energy $\to$ spectral gap as a composition of monotone maps, which is left for future work.

Remark 3 (Generality). The proposition does not compute r^* in closed form. The equilibrium depends on system parameters through the feedback chain in complex ways. Different parameter choices would yield different attractors. The result guarantees convergence to *some* r^*, given the feedback structure.

Remark 4 (Parameterization by r^*). The theorem is parameterized by r^*, so different system configurations yield different attractor values. The ARPACK-corrected run converges to $r^* = 3.0$; a previous run with a buggy eigensolver fallback converged to $r^* = 3.7$. Both are valid instances of the theorem. The shift in r^* from 3.7 to 3.0 is itself informative: the buggy eigensolver miscounted H^1, which altered governor dynamics, which changed the vertex expansion frequency $b(r)$, which shifted the equilibrium. The stability mechanism is invariant; only the equilibrium value changes.

5 Spectral Gap Stability and Coupled Result

We now establish a second stability result: the spectral gap of the sheaf Laplacian converges to a positive equilibrium. As with Theorem 1, this is conditional on explicitly stated assumptions (B1–B4) that are verified empirically or by design. Combined with Theorem 1, the result yields a coupled structural-spectral stability proposition for the restricted class of Purity-Gate-certified sheaves.

5.1 Assumptions

The spectral gap $\lambda_2(t) = \lambda_2(L_{\mathcal{F}}(t))$ is the smallest nonzero eigenvalue of the sheaf Laplacian (the Fiedler eigenvalue). The following assumptions parallel A1–A4 of Theorem 1.

Assumption 5 (B1: Restorative spectral drift)**.** The drift function $h_\lambda(x) = \mathbb{E}[\Delta\lambda_2 \cdot n \,|\, \lambda_2 = x]$ satisfies $h_\lambda(x) > 0$ for $x < \lambda_2^*$, $h_\lambda(x) < 0$ for $x > \lambda_2^*$, and $h'_\lambda(\lambda_2^*) < 0$. *Verification:* $h'_\lambda(\lambda_2^*) = -75.31$ ($p = 0.002$). Recovery rate after below-equilibrium events: 86.5% (45/52 events). Note: structural event *rate* is flat (∼27%/cycle); the restoring force operates through the *effect* of mutations on λ_2, not their triggering rate.

Assumption 6 (B2: Bounded spectral perturbation)**.** $|\Delta\lambda_2(t)| \leq C_\lambda/n(t)$ for some finite C_λ. *Verification:* $C_\lambda = 73.21$ (empirical maximum of $|\Delta\lambda_2 \cdot n|$ across 350 attractor observations).

Assumption 7 (B3: Bounded spectral radius (Purity Gate))**.** For all edges: $\sigma_{\max}(\rho) \leq 0.99$, enforced by deterministic rescaling at every restriction map construction site (Section 3.4).

Assumption 8 (B4: Stationary spectral drift equilibrium)**.** The scaled spectral drift

$$h_\lambda(x) = \mathbb{E}[\Delta\lambda_2 \cdot n \mid \lambda_2 = x]$$

is locally Lipschitz in an attractor neighborhood and has a unique zero $\lambda_2^* > 0$. In this neighborhood, $h_\lambda(x) > 0$ for $x < \lambda_2^*$ and $h_\lambda(x) < 0$ for $x > \lambda_2^*$. Empirically, the attractor-window trajectory is stationary (ADF $p = 0.016$, KPSS $p = 0.100$) with estimated $\lambda_2^* = 0.530$.

Assumption 9 (B5: Local spectral-density regularity)**.** In the attractor regime, the sheaf Laplacian spectrum has bounded local density near λ_2^*: for fixed rank perturbation q,

$$\lambda_{2+q}(L_{\mathcal{F}}) - \lambda_2(L_{\mathcal{F}}) = O(q/(nd)).$$

This condition is empirically supported by the observed bounded rank perturbations and measured $|\Delta\lambda_2|n$ values, but is not implied by Weyl's inequality alone.

5.2 Eigenvalue Perturbation Lemma

The mathematical core distinguishing Theorem 2 from Theorem 1. For E/V, the $O(1/n)$ perturbation scaling is algebraic ($\Delta(m/n)$ has a $1/n^2$ denominator). For λ_2, no such identity exists.

Lemma 2 (Per-edge bound)**.** *For any edge e with Purity-Gate-certified restriction maps ρ_u, ρ_v, the per-edge Laplacian contribution $L_e = B_e^\top B_e$ satisfies $\|L_e\|_{\mathrm{op}} \le 2\sigma_{\max}^2 = 1.9602$.*

Proof. The operator $B_e B_e^\top = \rho_u\rho_u^\top + \rho_v\rho_v^\top$ acts on the edge stalk space. Each restriction map satisfies $\sigma_{\max}(\rho) \le 0.99$, so $\|\rho_u\rho_u^\top\|_{\mathrm{op}} \le 0.99^2$ and $\|\rho_v\rho_v^\top\|_{\mathrm{op}} \le 0.99^2$. By the triangle inequality, $\|B_e B_e^\top\|_{\mathrm{op}} \le 2 \times 0.99^2 = 1.9602$. Since L_e and $B_e B_e^\top$ share nonzero eigenvalues, $\|L_e\|_{\mathrm{op}} \le 1.9602$. □

Lemma 3 (Batch bound)**.** *Adding $k \le 14$ Purity-Gate-certified edges produces a PSD perturbation ΔL with: (a) $\|\Delta L\|_{\mathrm{op}} \le k \cdot 1.9602 \le 27.44$ (Weyl bound); (b) $\mathrm{rank}(\Delta L) \le 2kd \le 224$; (c) $\mathrm{tr}(\Delta L) \le k \cdot 2d \cdot 0.99^2 \le 217.24$.*

The rank bound provides the qualitative $O(1/n)$ scaling: a rank-224 PSD perturbation to an nd-dimensional spectrum ($nd \ge 1616$ in the attractor) shifts λ_2 by at most the gap between λ_2 and λ_{226}. By Weyl's interlacing inequalities for rank-r PSD perturbations, $\lambda_2(L_{\mathcal{F}}) \le \lambda_2(L_{\mathcal{F}} + \Delta L) \le \lambda_{2+r}(L_{\mathcal{F}})$, and the gap $\lambda_{2+r} - \lambda_2$ scales as $O(r/(nd)) = O(1/n)$ for Laplacians with smooth spectral density near λ_2.

5.3 Theorem 2 (Spectral Gap Stability)

Theorem 4 (Spectral Gap Stability)**.** *Let $\{G(t)\}$ be a growing graph governed by the dynamics of Section 3. Under Assumptions B1–B5:*

(a) $\lambda_2(L_{\mathcal{F}}(t))$ converges to a neighborhood of $\lambda_2^* = 0.530 > 0$.

(b) The attractor width is $O(1/n) \to 0$.

(c) In the attractor regime, $\lambda_2(t) > 0$ conditional on graph connectivity (zero partitions across all runs) and Purity-Gate-bounded frustration.

Proof. The one-step update has the Robbins–Monro form:

$$\lambda_2(t{+}1) = \lambda_2(t) + \frac{1}{n}\big[h_\lambda(\lambda_2(t)) + \xi_\lambda(t)\big]$$

with $|\xi_\lambda| \leq 73.21$ (B2) and step size $1/n(t)$ satisfying the Borkar conditions [Borkar, 2008, 2026]. The Lyapunov function $W = (\lambda_2 - \lambda_2^*)^2$ has orbital derivative

$$\frac{dW}{ds} = 2(\lambda_2 - \lambda_2^*) \cdot h_\lambda(\lambda_2) < 0$$

for $\lambda_2 \neq \lambda_2^*$ by B1. The Borkar tracking theorem gives convergence to a neighborhood of λ_2^*. The attractor width $\delta \sim C_\lambda/(n \cdot |h'_\lambda(\lambda_2^*)|) = 73.21/(n \cdot 75.31) = O(1/n)$. Positivity follows from graph connectivity (zero partitions across all runs) and the BSS Cheeger inequality for connection Laplacians with Purity-Gate-bounded frustration [Bandeira et al., 2013]. □

Remark 5 (Non-Markovian dynamics). The graph state at time t depends on the entire mutation history. Borkar [2026] extends the SA tracking theorem to non-Markovian, non-ergodic environments under an asymptotic stationarity condition. Our ADF test ($p = 0.016$) verifies this condition in the attractor window.

5.4 Corollary (Coupled Structural-Spectral Stability)

Corollary 5. *Under Assumptions A1–A4 and B1–B5, certified dynamical sheaves (connected graphs, Purity-Gate-certified restriction maps, $E/V \geq 2$) satisfy in the attractor regime:*

(i) $|E|/|V| \to r^* \sim 3.0$ *with width* $O(1/n)$ *(Theorem 1)*

(ii) $\lambda_2(L_\mathcal{F}) \to 0.530 > 0$ *with width* $O(1/n)$ *(Theorem 4)*

(iii) $\beta_1/n(t) \to r^* - 1 \sim 2.0$ *(Euler characteristic)*

(iv) $\dim H^1 \leq d \cdot \beta_1 \to 2d \cdot n(t) = 16n(t)$ *(rank-nullity)*

(v) $\lambda_2 \geq 0.530 - O(1/n) > 0$ *(positivity)*

This is a sheaf expansion result for the certified class. First and Kaufman [2022, 2024] proved no universal sheaf Cheeger inequality exists. The Corollary establishes that the specific restricted class defined by Purity Gate algebraic certification maintains a positive spectral gap under autonomous growth, conditional on the verified assumptions A1–A4 and B1–B5.

The result holds because: (a) the Purity Gate bounds restriction map norms, placing the sheaf in a near-constant regime; (b) the graph is connected with $E/V \sim 3.0$; (c) the spectral governor maintains λ_2 via a feedback loop with verified restorative drift. Without any of these, the result fails.

Note on three levels of contradiction bounding. The Corollary's part (iv) involves three distinct quantities that must not be conflated: $\beta_1 = m - n + 1$, the graph-theoretic first Betti number, an exact combinatorial invariant determined by E/V alone; $\dim H^1(G; \mathcal{F}) \leq d \cdot \beta_1$, the sheaf cohomological upper bound, an algebraic inequality from rank-nullity of δ_0; and $\mu(t)$, the spectral obstruction proxy, a numerical quantity computed from near-zero eigenvalues of $L_\mathcal{F}$ via ARPACK with adaptive k-doubling (Section 2.5, Appendix A). The stability argument depends on β_1 (exact) and the Euler characteristic (exact). The sheaf bound $\dim H^1 \leq d \cdot \beta_1$ is algebraic. The proxy μ is operational. The three agree in direction but differ in precision and epistemic status.

5.5 Comparison of Theorems 1 and 2

Theorem 2 has a steeper restoring force (75 vs 1.2) compensating for larger noise (73 vs 12), producing a narrower attractor (0.005 vs 0.025). The spectral gap is more tightly regulated than the structural density.

Table 1: Comparison of the two stability results.

	Theorem 1	Theorem 2
Subject	E/V	$\lambda_2(L_{\mathcal{F}})$
Equilibrium	$r^* \sim 3.0$	$\lambda_2^* = 0.530$
ADF p-value	0.0003	0.016
Noise bound C	12	73.21
Drift slope $\lvert h'\rvert$	~1.2	~75.3
Attractor width at onset	~0.025	~0.005
$O(1/n)$ source	Algebraic	Spectral (rank arg.)

6 Empirical Validation

6.1 Setup

Cellular sheaf with stalk dimension $d = 8$. Initial: 150-vertex seed, 240 edges ($r_0 = 1.6$), 10 domains, 91 baked-in graph cycles. Primary validation: seed 42, 500 cycles (canonical run). Multi-seed validation: seeds 42, 72, 137, 271, 300 cycles each (Phase 3 runs with Purity Gate + density governor active).

6.2 Trajectory (Primary Run, Seed 42, 500 Cycles)

Table 2: Trajectory checkpoints for the canonical 500-cycle run (seed 42).

Cycle	$\lvert V\rvert$	$\lvert E\rvert$	$r(t)$	H^1	Ω	λ_2	Phase
0	150	240	1.60	7			Initial
50	150	340	2.27	0	0.286	0.396	Growth
100	165	466	2.82	6	0.463	0.682	Growth
150	201	604	3.00	12	0.540	0.413	Transition
200	253	746	2.95	24	0.566	0.533	Attractor
250	293	882	3.01	38	0.568	0.555	Attractor
300	336	1006	2.99	21	0.567	0.655	Attractor
350	387	1152	2.98	27	0.574	0.458	Attractor
400	431	1280	2.97	19	0.559	0.541	Attractor
450	470	1410	3.00	71	0.590	0.626	Attractor
500	515	1544	3.00	73	0.623	0.445	Attractor

6.3 Attractor Statistics (Primary Run)

Attractor window: cycles 150–500 (350 observations). Mean r: 3.0, std: 0.02, range: $[2.9, 3.1]$.

Measured stationary rates (attractor window): total edge growth 940 edges $\to \bar{m} = 2.686$/cycle; total vertex growth 314 vertices $\to \bar{n} = 0.897$/cycle; stationary rate ratio $\bar{m}/\bar{n} = 2.994$.

Spectral data (attractor window): λ_2 mean 0.540 (2nd half) vs 0.484 (1st half); λ_2 range $[0.205, 0.689]$; zero-gap cycles: 0/500.

Governor data (full run): Ω mean 0.504, std 0.099, peak 0.625; warning threshold (0.70) never reached.

Coherence Margin (attractor window): CM mean 3.99, CM min 0.23; stress events ($\mathrm{CM}_{\mathrm{ema}} < 0$): 0; lead times before H^1 spikes: 1 to 70 cycles; 13 distinct spike events ($H^1 > 50$).

6.4 Phase Transition Evidence

The trajectory data exhibits the growth-to-attractor phase transition described in Section 4.2.

Growth phase (cycles 0–100, $r = 1.6$ to 2.8): vertex rate $15/100 = 0.15$/cycle; edge rate $(466 - 240)/100 = 2.26$/cycle; edge creation dominates; r climbing steadily.

Transition (cycles 100–150, $r = 2.8$ to 3.0): vertex rate accelerates as Ω crosses threshold; edge saturation begins to engage.

Attractor phase (cycles 150–500, $r = 3.0 \pm 0.1$): vertex rate $314/350 = 0.90$/cycle; edge rate $940/350 = 2.69$/cycle; rates balanced at $r^* = 2.994$.

The vertex expansion rate increases by a factor of $6\times$ (0.15 to 0.90) as r transitions from the growth phase to the attractor phase, confirming the monotonicity of $b(r)$.

6.5 Multi-Seed Validation

Four seeds (42, 72, 137, 271) were run for 300 cycles each with the Purity Gate ($\sigma_{\max} \leq 0.99$) and density governor ($w = 1.5$) active. All restriction maps were enforced by deterministic rescaling at construction time.

Table 3: Multi-seed validation at 300 cycles.

Seed	$\|V\|$	$\|E\|$	E/V	$\bar{\lambda}_2$	Gov. fired	Crashes
42	319	964	3.0	0.40	Yes	0
72	334	1000	3.0	0.40	No	0
137	309	968	3.1	0.46	Yes	0
271	308	968	3.1	0.54	Yes	0
Mean / CV			3.05		3/4	CV = 1.9%

Cross-seed E/V: mean 3.05, std 0.058, CV = 1.9%. Zero crashes across 1200 total cycles. Zero partitions. Zero zero-gap cycles. λ_2 positive throughout all runs.

Three of four seeds required the density governor (Ω exceeded the 0.70 warning threshold, triggering vertex expansion). Seed 72 converged to $E/V = 3.0$ with the Purity Gate alone (Ω peaked at 0.624, below threshold). Seeds 42, 137, 271 overshot to $E/V = 3.4$–3.5 before the governor corrected the trajectory back toward 3.0–3.1.

This confirms the two-tiered stability architecture (Section 3.5): the Purity Gate provides intrinsic stability sufficient for some seeds, while the density governor provides the extrinsic safety net required for seed-invariant convergence.

6.6 Topological Latency

The lead times between structural perturbation and H^1 response (1 to 70 cycles across 13 spike events in the primary run) provide direct evidence of the topological latency described in Section 4.3(v). The causal chain operates as follows: (1) edge additions increase $\beta_1 = m - n + 1$; (2) the graph undergoes topological latency (1–70 cycles) as the new independent cycles propagate through the restriction map landscape; (3) restriction map frustration reaches critical mass, registering as an H^1 spike in the spectral proxy; (4) elevated H^1 raises Ω above the expansion threshold; (5) vertex expansion fires ($\Delta_n = 1$), acting as the restoring force.

This latency explains the observed attractor width of ± 0.1 (versus the instantaneous perturbation bound of ± 0.025): the system overshoots the instantaneous bound during the latency window before the governor response corrects the trajectory. The overshoot is a structural feature of the system's temporal architecture, not a deficiency of the stability mechanism.

6.7 Ablation: Two-Tiered Stability

An ablation study isolates the contribution of each stability mechanism across the multi-seed validation.

Without the Purity Gate ($\sigma_{\max} = 1.0$, no enforcement): the edge-vertex ratio drifts above the attractor ($E/V = 3.52$ at 300 cycles for seed 42 in Phase 2 baseline), confirming that the zero-resistance regime permits unchecked edge accumulation.

With both tiers active: two distinct behaviors emerge. Seed 72 converged to $E/V = 3.0$ without governor intervention (Ω peak $0.624 < 0.70$ threshold); the Purity Gate alone was sufficient. Seeds 42, 137, 271 overshot to $E/V = 3.4$–3.5, triggering the governor (Ω peaks 0.80, 0.83, 0.80 respectively) and executing corrective vertex expansion. All four seeds converged to $E/V \in [3.0, 3.1]$ with CV = 1.9%.

Both tiers are necessary for seed-invariant convergence. Removing the Purity Gate allows unchecked drift. Removing the governor leaves 3/4 seeds without correction for overshoots.

6.8 Scale Validation via Cellular Decomposition

The monolithic sheaf architecture described in Sections 1–5 operates on a single sheaf Laplacian $L_{\mathcal{F}}$. For graphs exceeding approximately $V = 1{,}000$ vertices, the dense eigensolve required for H^1 computation scales as $O(N^3)$, creating a practical barrier.

To validate the stability results at enterprise-relevant scale, we implemented a cellular decomposition that partitions the sheaf into self-contained cells, each independently computing local cohomology and spectral data. Global cohomology is recovered via Mayer–Vietoris assembly and a Čech spectral sequence computed from the nerve complex of the cell decomposition. Two algorithmic optimizations proved critical: (1) shift-invert eigsh as the primary Fiedler solver, targeting eigenvalues near zero via LU factorization and converging in 10–50 iterations regardless of graph size; and (2) $O(\text{boundary\_pairs})$ boundary map initialization replacing the previous $O(\text{cells}^2)$ nested loop.

All restriction maps in the cellular architecture satisfy the same Purity Gate constraint (B3): $\sigma_{\max}(\rho) \le 0.99$, enforced by deterministic rescaling at every construction site. The cellular decomposition preserves the sheaf structure on which Theorems 1 and 2 are stated.

Scale validation was performed on a single commodity workstation (Intel i9-13900H, 64 GB RAM, no GPU acceleration) across four scales in a single overnight run (total wall time: 14 hours 38 minutes, zero crashes):

Table 4: Cellular decomposition scale validation.

V	Cells	Partition	Eigensolve	ms/vertex	Memory
50,000	431	59.8 s	192.9 s	3.858	1.6 GB
100,000	890	238.3 s	337.4 s	3.374	3.1 GB
250,000	2,200	940.7 s	913.6 s	3.655	7.4 GB
1,000,000	8,663	23,283.5 s	3,550.3 s	3.550	28.6 GB

The per-vertex eigensolve cost is effectively constant: mean 3.55 ms/vertex with coefficient of variation 5.9% across a 20× range of graph sizes. This confirms the cellular architecture achieves its design goal: decomposing the monolithic $O(V^3)$ eigenproblem into $O(V/v_{\max})$ independent $O(v_{\max}^3)$ eigenproblems, with bounded $v_{\max}$ (200 vertices per cell) producing constant per-vertex cost.

The nerve complex of the cellular decomposition maintains maximum simplex dimension 4 from $V = 100{,}000$ through $V = 1{,}000{,}000$, confirming that the partitioner produces a locally finite cover without combinatorial explosion. Simplex count grows linearly with cell count ($\sim$4.6 simplices per cell).

At $V = 1{,}000{,}000$, the sheaf Laplacian has dimension $8{,}000{,}000 \times 8{,}000{,}000$ (stalk dimension $d = 8$). The cellular decomposition reduces this to 8,663 independent eigenproblems of dimension ∼1,000 each, processed without memory overflow or solver crashes. Memory peaked at 28.6 GB on 64 GB RAM (44% utilization), indicating room for further scaling on the same hardware.

To our knowledge, this is one of the largest reported sheaf-based structural verification computations of which we are aware. Prior implementations of sheaf Laplacian spectral analysis typically operate at scales below 5,000 vertices. The 200× extension to 1,000,000 vertices with constant per-vertex cost validates the practical applicability of the stability results established in Theorems 1 and 2.

6.9 Power-Law Network Validation (Enron Email Corpus)

The geometric random graphs of Section 6.8 have bounded degree variance. Real-world knowledge graphs exhibit power-law degree distributions where a small number of hub vertices dominate the spectrum. These hubs cause Fiedler vector degeneracy, breaking naive spectral bisection. We validated the cellular architecture on two power-law networks: the Stanford SNAP Enron email corpus and a synthetic Barabási–Albert graph.

Enron Email Network. The Enron email dataset (36,692 accounts, 183,831 edges) was filtered to its 3-core (21,309 vertices, 166,039 edges, $E/V = 7.79$, degree coefficient of variation 2.609, maximum degree 1,141). The sheaf Laplacian has dimension $170{,}472 \times 170{,}472$.

An eight-layer decomposition pipeline was developed to handle the power-law topology: (1) hub extraction (549 vertices with degree exceeding mean $+2\sigma$, threshold 97); (2) hub neighbor freeze (14,593 vertices, 70.3% of non-hub, adjacent to any hub); (3) connected component detection (1,209 components after hub removal); (4) BFS balanced partitioning for the giant component ($V > 5{,}000$); (5) adaptive $v_{\max}$ computed from local E/V density; (6) post-partition merge of undersized partitions; (7) k-doubling cap with graph Laplacian fallback; (8) clean split with runtime merge size check.

Results (Run #23, i9-13900H workstation, no GPU): Phase 2 (decomposition) 3.3 s (639 cells, mean $V = 34$); Phase 3 (eigensolves) 12.9 s (0.606 ms/vertex); nerve complex 796 simplices, max dimension 2; peak memory 762 MB; total pipeline 30.3 s.

Multi-seed validation (4 seeds): partition structure identical across all seeds (3,760 cells, 254 nerve edges, max dim 2). The partition depends only on graph topology, not sheaf data.

Barabási–Albert Scale-Free Graph. BA(100,000, 8): $V = 100{,}000$, $E = 799{,}936$, $E/V = 8.00$, degree CV = 1.464, max degree 1,474. Hub extraction: 1,756 hubs (degree > 63). Hub neighbor freeze: 82,858 vertices (84.3% of non-hub). Phase 2: 31.8 s. Phase 3: 59.9 s (0.599 ms/vertex). Total: 2 min 44 s, 1,254 MB peak.

Scaling curve across topologies: $V = 21{,}309$ (Enron, power-law) 0.606 ms/vertex; $V = 100{,}000$ (BA, power-law) 0.599 ms/vertex; $V = 100{,}000$ (geometric) 0.511 ms/vertex; $V = 1{,}000{,}000$ (geometric) 0.845 ms/vertex. The per-vertex cost is sub-millisecond across both geometric and power-law topologies from 21K to 1M vertices, confirming that the cellular decomposition achieves $O(n)$ total cost independently of the degree distribution.

6.10 Sublinear Streaming via Lazy Evaluation (B5 Architecture)

The batch cellular decomposition of Section 6.8 computes H^1 from scratch. For streaming applications (continuous edge insertions with intermittent queries), this is wasteful: most edits affect only a small fraction of cells. We implemented a streaming architecture with three key properties.

(a) $O(1)$ per-edit cost. Each edge insertion performs: one dict lookup (find host cell), one set membership check (is cell already dirty?), one dict write and set add (mark cell dirty). No

eigensolves, no matrix operations, no traversal. The per-edit cost is bounded by $v_{\max}$ and stalk dimension s, independent of V.

(b) Deferred exact reconciliation. Dirty cells accumulate during streaming. At query time, `flush_dirty_cells()` recomputes H^1 only for cells marked dirty since the last flush. After flush, the global H^1 is identical to a full-graph eigensolve (drift = 0 across all measured seeds and scales).

(c) $O(1)$ query via nerve tree. A hierarchical dendrogram over the cell partition provides a Lowest Common Ancestor (LCA) oracle for $O(1)$ topological distance queries between arbitrary vertex pairs.

Two optimizations reduced per-edit cost by 9× (cProfile verified): (1) lazy-lazy design deferring all cache invalidation to flush time (prior implementation consumed 62% of streaming cost on per-edit deallocation); (2) restriction pool of 1,024 pre-computed random orthogonal 8×8 matrices, replacing per-edit QR decomposition + SVD conditioning (previously 24% of cost).

Measured performance (DR-1 canonical run, April 17, 2026):

Table 5: Streaming edit and query performance.

V	Mean/edit	p99/edit	Drift	Seeds
21,000	0.031 ms	0.081 ms	0	3
100,000	0.046 ms	0.099 ms	0	3
250,000	0.051 ms	0.152 ms	0	1
1,000,000	0.063 ms	0.143 ms	0	1

Scaling exponent: 0.1851 ($R^2 = 0.975$, four-point log-log fit). Baseline (flat nerve, no B5 optimizations): slope 0.5532. Throughput at $V = 1\text{M}$: ∼15,873 edits/second/core (single-threaded).

Query latency via nerve tree LCA oracle: at $V = 1\text{M}$, p99 = 0.013 ms, representing a 10,504× speedup over flat baseline. Flush cost: $O(\text{dirty\_cells} \cdot v_{\max}^3 \cdot s^3)$. After 1,000 streaming edits at $V = 1\text{M}$, 749 of 4,611 cells (16%) were dirtied. Flush completed in 48 seconds single-threaded (parallelizable to ∼5 seconds on 20 cores, as each cell's eigensolve is independent).

Memory: 40.7 GB at $V = 1\text{M}$ (dense Laplacian storage in the Python prototype). A sparse-storage implementation is projected to reduce memory 5–10×.

All restriction maps in the streaming architecture satisfy the same Purity Gate constraint (B3): $\sigma_{\max}(\rho) \leq 0.99$, enforced by the restriction pool (all 1,024 pre-computed matrices are rescaled to satisfy the bound). The drift-zero fix ensures that cross-cell edge insertions mark both affected cells dirty, preventing boundary-induced drift.

The streaming architecture preserves the formal properties required by Theorems 1 and 2: after flush, the sheaf Laplacian and its spectral data are identical to the batch computation. The deferred consistency model (stale between edits, exact after flush) is analogous to LSM-tree write-optimized databases and incremental SAT solvers with checkpointing.

Hardware: Intel i9-13900H, 64 GB RAM, no GPU. All measurements on a single commodity workstation, single-threaded.

7 Discussion

7.1 Limitations

(a) Assumption B1 (restorative spectral drift, Theorem 2) is verified empirically but not derived from the sheaf dynamics. This is the main gap between the current propositions and self-contained theorems. Assumption A2 (local restorative density drift, Theorem 1) is enforced above the equilibrium by the density-aware governor (Section 3.5) and empirically verified across the observed attractor neighborhood. For B1, the restoring force operates through the *effect* of mutations on λ_2, not their triggering rate. This was discovered during verification and required reformulating B1.

(b) The equilibria r^* and λ_2^* are characterized via stationary rate balance but not computed in closed form.

(c) The spectral obstruction proxy μ is not the standard sheaf-cohomological $\dim H^1$ (Section 2.5, Appendix A). The stability argument does not depend on this distinction.

(d) Multi-seed validation uses 300-cycle runs across four seeds. The canonical 500-cycle trajectory (seed 42) provides the primary attractor-phase statistics. The cross-seed coefficient of variation (CV = 1.9%) confirms seed-invariant convergence to $r^* \approx 3.0$.

(e) The $O(1/n)$ scaling of the spectral perturbation (Theorem 2) is supported by the rank-vs-dimension argument (Section 5.2) and verified empirically, but the theoretical bound from Weyl gives only $O(1)$. The gap between theoretical $O(1)$ and empirical $O(1/n)$ is explained but not closed by a tight theoretical bound.

(f) The cellular decomposition (Section 6.8) addresses the $O(N^3)$ eigensolver barrier that previously limited practical application to $V < 1{,}000$. With constant per-vertex cost demonstrated to $V = 1{,}000{,}000$, the stability results of Theorems 1 and 2 are now validated at scales relevant to enterprise knowledge bases. The streaming architecture (Section 6.10) extends this further: per-edit cost of 63 microseconds at $V = 1\text{M}$ with measured scaling exponent 0.19 enables real-time incremental maintenance of the sheaf-governed dynamics.

(g) Theorem 2 applies to certified sheaves specifically (Purity Gate constraints, connected graphs, $E/V \geq 2$). It does *not* generalize to arbitrary sheaves, consistent with the First–Kaufman impossibility result [First and Kaufman, 2022, 2024].

(h) Idempotency and nilpotency diagnostics show scale-mismatch artifacts from mixed initialization/growth regimes (idempotency error up to $O(|s_{\text{init}} - s_{\text{growth}}|)$). These are diagnostic indicators, not load-bearing for the stability proof, which depends solely on the spectral radius bound (B3).

7.2 Domain Dependence of r^*

The equilibrium r^* is determined by the system's mutation dynamics, specifically the balance between edge creation rate $\bar{m}(r)$ and vertex expansion rate $b(r)$. Different knowledge domains (e.g., legal citation graphs, medical ontologies) will exhibit different natural densities and thus different r^* values. The stability machinery (Theorems 1–2, Corollary) transfers unchanged; only the equilibrium value and associated constants require re-measurement. Domain pilots to characterize r^* for specific verticals are planned as future work.

7.3 Topological Latency as a Structural Property

The topological latency τ (ranging from 1 to 70 cycles in our experiments) is a fundamental property of the feedback architecture. It arises because the coboundary operator δ_0 must be recomputed after structural modifications, and the heat flow $ds/dt = -L_{\mathcal{F}}s$ requires multiple cycles to propagate energy changes through newly created graph cycles to the spectral proxy.

This latency has a direct analogue in control theory: the system has a transport delay between the plant (graph structure) and the controller (governor). Classical results on delayed feedback systems (Smith predictor, Artstein reduction) suggest that the attractor width could be tightened by incorporating predictive elements into the governor. This is left for future work.

The observation that topological latency produces bounded overshoot but does not destabilize the attractor is itself a robustness result: the system tolerates delays of up to 70 cycles without losing structural coherence.

7.4 Extensions

The Lyapunov framework extends to higher-dimensional sheaves, sheaves over hypergraphs, and alternative pressure functionals, provided the feedback monotonicity (A2) holds.

The cellular decomposition (Section 6.8) extends the practical reach of the stability results by three orders of magnitude (from $V \sim 1{,}000$ to $V = 1{,}000{,}000$). The streaming architecture (Section 6.10) adds a fourth mode: incremental maintenance with $O(1)$ per-edit cost and measured $V^{0.19}$ overhead, enabling the sheaf-governed dynamics to run in real time on streaming data. Both preserve all algebraic properties required by Theorems 1 and 2: Purity Gate enforcement, connected subgraphs, and bounded restriction map norms. The Mayer–Vietoris assembly and Čech spectral sequence recover exact global cohomology from the local cell data, ensuring that the stability guarantees transfer from the monolithic formulation to the cellular and streaming implementations.

The stochastic approximation framing (Section 4.1) opens a natural path to multi-timescale analysis. If the system were extended with a second slow variable (e.g., stalk dimension $d(t)$ or restriction map complexity), the Borkar two-timescale framework [Borkar, 2008] would apply, with r converging on the fast timescale and the slow variable evolving on a separate ODE.

7.5 The Coherence Margin as Predictive Observable

The Coherence Margin (CM) is defined as the ratio of contradiction-resolution capacity to contradiction-creation pressure, combining the spectral gap λ_2, the Dirichlet energy residual, and the current H^1 obstruction count into a single real-valued signal. The coupled stability of Theorems 1 and 2 provides the formal basis for CM's predictive power: because both the structural density and the spectral gap are attracted to stable equilibria, deviations from equilibrium are self-correcting, and the CM detects these deviations before they manifest as H^1 spikes.

In the 500-cycle authoritative run, the CM remained positive throughout the entire 350-cycle attractor window (CM mean: 3.99, CM min: 0.23), providing 1 to 70 cycles of advance warning before all 13 major H^1 spike events ($H^1 > 50$). Zero stress events ($\mathrm{CM}_{\mathrm{ema}} < 0$) were recorded. This means the system's resolution capacity exceeded its creation pressure at every point in the attractor regime; a direct observable consequence of the coupled Lyapunov stability.

The CM converts the abstract stability guarantees of Theorems 1 and 2 into an operational signal: a positive CM certifies that the substrate is within its stability envelope. A declining CM provides advance warning that the system is approaching the boundary of its attractor, enabling preemptive intervention before structural coherence degrades.

False positive immunity. The CM does not trigger false alarms during the natural ± 0.1 macroscopic E/V overshoot (Section 4.3(v)) for four reasons: (i) E/V does not enter the CM formula directly; the overshoot reaches the CM only through the multi-step causal chain $E/V \rightarrow \beta_1 \rightarrow$ frustration $\rightarrow H^1 \rightarrow$ rates, mediated by topological latency; (ii) double EMA smoothing ($\alpha = 0.1$ on both inputs and output) creates a low-pass filter with $\sim$20-cycle effective window, attenuating transient perturbations; (iii) λ_2 is $20\times$ more tightly regulated than E/V (attractor width 0.005 vs 0.1 per Section 5.5), stabilizing the CM's capacity term during structural oscillations; and (iv) the minimum observed CM (0.23) provides ample headroom above the stress threshold ($\mathrm{CM}_{\mathrm{ema}} < 0$). The topological latency that causes the E/V overshoot simultaneously prevents the overshoot from propagating to the CM inputs, acting as a natural band-limiting filter between the structural plant and the spectral indicator.

7.6 The Coupled Result

Theorems 1 and 2 together establish that the system simultaneously stabilizes two independent structural invariants: the edge-vertex ratio (a combinatorial quantity) and the sheaf Laplacian spectral gap (a spectral quantity). These operate on different mathematical substrates: one is a simple ratio, the other an eigenvalue of a growing block matrix, yet both converge via the same Borkar stochastic approximation framework.

The coupling is not coincidental. Both invariants are governed by the same spectral governor (Ω), which receives inputs from both the structural density (via energy pressure) and the spectral gap (via the gap component). The governor's weighted response creates the restorative drift for *both* variables simultaneously. This shared governance mechanism is why the feedback loop stabilizes two quantities at once, rather than trading one off against the other.

The coupled stability has a practical consequence: the system's Reasoning Assurance Layer operates on a substrate with formal stability results covering both its topological structure (bounded density, bounded complexity) and its spectral properties (positive connectivity, bounded obstruction count), conditional on the verified assumptions.

7.7 Connection to Spectral Sheaf Theory

Several recent results in spectral sheaf theory bear on the properties observed in our system.

Sheaf expansion and the Cheeger inequality. First and Kaufman [2022, 2024] proved that no universal Cheeger-type inequality exists relating graph expansion to sheaf coboundary expansion for arbitrary sheaves. However, they established positive results for restricted classes: constant sheaves recover the full classical Cheeger bound, and near-constant sheaves inherit expansion from the underlying graph. Our system's restriction maps are generated by the Algebraic Purity Gate with bounded spectral radius and near-identity structure. This places our sheaves in the near-constant regime where First–Kaufman positive results apply. The observed spectral gap maintenance (λ_2 bounded away from zero during H^1 spikes of magnitude 76) is consistent with Cheeger-type bounds holding for this restricted class, though we do not prove this.

Over-squashing and over-smoothing. Giraldo et al. [2023] proved that for standard graph neural networks, the spectral gap of the graph Laplacian simultaneously controls both over-squashing and over-smoothing, creating an inevitable trade-off. Bamberger et al. [2025] showed that bundle neural networks (a special case of sheaf neural networks with flat vector bundles) admit equilibria with distinct node representations even with a positive spectral gap; the fixed-point manifold permits non-constant solutions under non-trivial bundles. Our attractor-window data (λ_2 mean 0.530, non-trivial H^1 preserved for 350 cycles) provides empirical evidence of this decoupling under autonomous dynamics. To our knowledge, this is the first observation of spectral gap / harmonic space coexistence over extended autonomous graph growth.

Persistent sheaf Laplacians. Wei and Wei [2025] introduced persistent sheaf Laplacians combining cellular sheaves with filtration structure. The spectral stability of persistent Laplacian eigenvalues under simplex insertions has been established by Le Vu Anh et al. [2025] for the standard (non-sheaf) case, but the sheaf extension remains open. Our system's incremental rebuild mechanism generates eigenvalue trajectories under sequential structural modifications with certified restriction maps, providing a natural experimental platform for persistent sheaf Laplacian stability analysis.

Spectral filtering and polynomial diffusion on sheaves. Borgi et al. [2025] introduced polynomial spectral filtering on cellular sheaves, applying polynomial filters directly to the sheaf

Laplacian spectrum rather than learning restriction maps end-to-end. This spectral filtering perspective is relevant to our stability analysis: the spectral proxy $\mu(t)$ (Section 2.5) operates as a spectral filter on $L_{\mathcal{F}}$, and the stability of the filtered signal under autonomous graph growth is precisely what Theorem 2 establishes.

Sheaf Laplacian learning and restriction map inference. Di Nino et al. [2025] proposed learning the sheaf Laplacian, including both graph topology and restriction maps, from observed data by minimizing total variation, with restriction map optimization resolved in closed form. Our system inverts this relationship: rather than inferring the sheaf from data, we use the sheaf Laplacian's spectral properties to govern the graph's structural evolution. The stability of this inverse process (using the sheaf to control the graph rather than learning the sheaf from the graph) is the novel contribution of the present work. Hernandez Caralt et al. [2026] examined when learning restriction maps provides genuine benefit over identity sheaves, finding that with proper normalization and residual connections, identity sheaves can match learned sheaves on several benchmarks. Our Purity Gate constraint (B3) places restriction maps in a near-identity regime, consistent with their finding that the near-identity class retains favorable spectral properties.

Sheaf Laplacians in multi-agent coordination and control. Battiloro et al. [2024] extended sheaf-theoretic methods to tangent bundle convolutional learning, establishing connections between cellular sheaf Laplacians and signal processing on manifolds. Riess and Hale [2025] applied cellular sheaves to distributed multi-agent coordination, using the sheaf Laplacian to enforce inter-agent consistency constraints; a formulation that parallels our use of sheaf cohomology for contradiction detection. Riess and Ghrist [2022] extended sheaf Laplacians beyond vector spaces to partially ordered structures via the Tarski Laplacian, suggesting a path toward applying our stability framework to non-numeric knowledge representations.

Heterogeneous and structured sheaf architectures. Braithwaite et al. [2024a] extended sheaf neural networks to heterogeneous graphs, encoding different node and edge types directly in the sheaf structure rather than in the model architecture. Barbero et al. [2022] proposed computing restriction maps via manifold-aware orthogonal alignment using connection Laplacians, drawing on Riemannian geometry. Both approaches address the question of how to construct appropriate sheaves for structured data, a question our system resolves through the Purity Gate's algebraic constraints rather than through learning.

Code and data. Source code, multi-seed run data, and the SIGMA engine are available at `https://invariant.pro`. Preprint also available on Zenodo (DOI: 10.5281/zenodo.19713406).

References

Jacob Bamberger, Federico Gallicchio, Federico Barbero, and Charlie Sherborne. Bundle neural networks for message diffusion on graphs. In *International Conference on Learning Representations (ICLR)*, 2025. Spotlight. arXiv:2405.15540.

M. Banakar. Lyapunov stability analysis of iterative optimization algorithms. *ISRN Applied Mathematics*, 2011.

Afonso S. Bandeira, Amit Singer, and Daniel A. Spielman. A Cheeger inequality for the graph connection Laplacian. *SIAM Journal on Matrix Analysis and Applications*, 34(4):1611–1630, 2013.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

Federico Barbero, Cristian Bodnar, Haitz Saez de Ocariz Borde, Michael M. Bronstein, Petar Veličković, and Pietro Liò. Sheaf neural networks with connection Laplacians. In *ICML Topological, Algebraic and Geometric Learning Workshop*, 2022. arXiv:2206.08702.

Claudio Battiloro, Zhiyang Wang, Hans Riess, Paolo Di Lorenzo, and Alejandro Ribeiro. Tangent bundle convolutional learning: from manifolds to cellular sheaves and back. *IEEE Transactions on Signal Processing*, 2024. arXiv:2210.15058.

Cristian Bodnar, Francesco Di Giovanni, Benjamin Paul Chamberlain, Pietro Liò, and Michael M. Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Amine Borgi, Fabrizio Silvestri, and Pietro Liò. Polynomial neural sheaf diffusion: A spectral filtering approach on cellular sheaves. *arXiv preprint arXiv:2512.00242*, 2025.

Vivek S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.

Vivek S. Borkar. Stochastic approximation in non-Markovian environments revisited. *arXiv preprint arXiv:2603.21091*, 2026.

Luke Braithwaite, Iulia Cristina Duta, and Pietro Liò. Heterogeneous sheaf neural networks. *arXiv preprint arXiv:2409.08036*, 2024a.

Luke Braithwaite, Iulia Cristina Duta, and Pietro Liò. Sheaf hypergraph networks. *arXiv preprint arXiv:2409.02616*, 2024b.

Luca Di Nino, Sergio Barbarossa, and Paolo Di Lorenzo. Learning sheaf Laplacian optimizing restriction maps. *arXiv preprint arXiv:2501.19207*, 2025.

Sergey N. Dorogovtsev and José F. F. Mendes. Evolution of networks. *Advances in Physics*, 51 (4):1079–1187, 2002.

Uriya First and Tali Kaufman. The cheeger inequality and coboundary expansion: Beyond constant coefficients. *arXiv preprint arXiv:2208.01776*, 2022. Updated April 2025.

Uriya First and Tali Kaufman. Cosystolic expansion of sheaves on posets with applications to good 2-query locally testable codes and lifted codes. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, 2024.

Thomas Gebhart, Paul Schrater, and Alan Hylton. Knowledge sheaves: A sheaf-theoretic framework for knowledge graph embedding. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023.

Jhony H. Giraldo, Konstantinos Skianis, Thierry Bouwmans, and Fragkiskos D. Malliaros. On the trade-off between over-smoothing and over-squashing in deep graph neural networks. In *ACM International Conference on Information and Knowledge Management (CIKM)*, 2023.

Thilo Gross and Bernd Blasius. Adaptive coevolutionary networks: A review. *Journal of the Royal Society Interface*, 5(20):259–271, 2008.

Jakob Hansen and Thomas Gebhart. Sheaf neural networks. *arXiv preprint arXiv:2012.06333*, 2020.

Ferran Hernandez Caralt, Neil Sherborne, Andrés Sherding, and Max Sherborne. On the necessity of learnable sheaf Laplacians. *arXiv preprint arXiv:2603.05395*, 2026.

D. Le Vu Anh, M. Dik, and D. Nguyen. Lipschitz bounds for persistent Laplacian eigenvalues under one-simplex insertions. *arXiv preprint arXiv:2506.21352*, 2025.

Max Revay, Ruigang Wang, and Ian R. Manchester. Lyapunov-stable deep equilibrium models. In *AAAI Conference on Artificial Intelligence*, 2024.

Spencer M. Richards, Felix Berkenkamp, and Andreas Krause. The Lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems. In *Conference on Robot Learning (CoRL)*, 2018.

Hans Riess and Robert Ghrist. Cellular sheaves of lattices and the tarski Laplacian. *Homology, Homotopy and Applications*, 2022. arXiv:2007.04099.

Hans Riess and Matthew Hale. Distributed multi-agent coordination over cellular sheaves. *arXiv preprint arXiv:2504.02049*, 2025.

Herbert Robbins and Sutton Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3):400–407, 1951.

Xiaoqi Wei and Guo-Wei Wei. Persistent sheaf Laplacians. *Foundations of Data Science*, 7(2), 2025.

A Relationship Between μ and $\dim H^1$

The spectral obstruction proxy $\mu(t)$:

$$\mu(t) = |\{\text{eigenvalues of } L_{\mathcal{F}} \text{ below tolerance}\}| - \beta_0$$

is *not* the standard $\dim H^1$ from equation (6):

$$\dim H^1 = d \cdot (m - n) + \dim H^0$$

For identity restriction maps on a connected graph, $\ker(L_{\mathcal{F}})$ has dimension d (one constant section per stalk coordinate), giving $\mu = d - 1$, while $\dim H^1 = d \cdot \beta_1$. For generic (non-identity) restriction maps, $\ker(L_{\mathcal{F}})$ is typically much smaller.

The distinction does not affect the stability argument. The proof requires that Dirichlet energy $\mathcal{E}[s(t)]$ correlates positively with $r(t)$, mediated by $\beta_1 = (r - 1) \cdot n + 1$ (purely graph-theoretic, exact). The specific obstruction count (whether μ or $\dim H^1$) enters the pressure functional but does not appear in the Lyapunov function or the equilibrium condition (17).

B Notation Summary

Theorem 1 notation. $G(t)$: finite graph at time t. $n(t), m(t)$: $|V(t)|, |E(t)|$. $r(t)$: edge-vertex ratio $m(t)/n(t)$. r^*: attractor equilibrium $= \bar{m}/\bar{n}$ ($\sim$3.0). $\mathcal{F}(t)$: cellular sheaf over $G(t)$. d: stalk dimension ($d = 8$ in implementation). δ_0: coboundary operator. $L_{\mathcal{F}}$: 0-sheaf Laplacian $\delta_0^\top \delta_0$. $\mathcal{E}[s]$: Dirichlet energy. H^0, H^1: sheaf cohomology groups. $\mu(t)$: spectral obstruction proxy (not $\dim H^1$). β_0, β_1: Betti numbers. $\Delta_m(t), \Delta_n(t)$: per-cycle edge/vertex increments. $\bar{m}, \bar{n}$: time-averaged stationary rates. $b(r)$: expansion frequency (function of r). $h(r)$: drift function $\bar{m}(r) - r \cdot b(r)$. $V(r)$: Lyapunov function $(r - r^*)^2$. C_m: maximum single-cycle edge increment (empirically 12). τ: topological latency (1–70 cycles). $\xi(t)$: martingale difference noise term. w: density governor weight (1.5).

Theorem 2 notation. $\lambda_2(t)$: spectral gap (Fiedler eigenvalue) of $L_{\mathcal{F}}$. λ_2^*: spectral gap equilibrium (0.530). $h_\lambda(x)$: drift function for λ_2. $\xi_\lambda(t)$: spectral noise term. C_λ: maximum scaled spectral perturbation (73.21). $W(x)$: Lyapunov function $(x - \lambda_2^*)^2$. $\rho_{\max}$: Purity Gate spectral radius bound (0.99). L_e: per-edge Laplacian contribution. B_e: per-edge coboundary operator. ΔL: batch perturbation to $L_{\mathcal{F}}$.

Coherence Margin. $\mathrm{CM}(t)$: ratio of resolution capacity to creation pressure, derived from λ_2, Dirichlet energy, and H^1 count. $\mathrm{CM} > 0$ certifies stability. $\mathrm{CM}_{\mathrm{ema}}(t)$: exponential moving average of CM.