Beyond Bottom-k: Structural Nullspaces in Spectral Clustering of Sparse Heterogeneous Graphs

Anonymous authors
Paper under double-blind review

Abstract

Spectral clustering is often reported to underperform on sparse heterogeneous graphs, motivating increasingly sophisticated spectral operators. We show that on several standard benchmarks the classical Laplacian is not the main obstacle: literal bottom-k selects structural zero modes created by disconnected components before it reaches the nonzero low-frequency spectrum. We prove that when these zero modes fill the k-dimensional embedding and the evaluated vertices lie in one component, row normalization collapses the representation and removes all within-component geometry available to k-means. We also quantify the partial-distortion regime below capacity. This suggests a minimal correction: skip the structurally known zero modes and use the first k nonzero eigenvectors, leaving the graph, Laplacian, normalization, and clustering algorithm unchanged. Across the sparse benchmarks we study, this corrected classical baseline reaches the highest mean AMI (0.62) in a strictly protocol-matched comparison against the sophisticated operators proposed for this regime (regularized spectral clustering, Bethe–Hessian, and Di-SIM). Because the nullspace is known from connected components, the correction can also be implemented directly in a sparse eigensolver without computing the discarded modes. Our results suggest that spectral benchmarks should verify which part of the spectrum a baseline actually extracts before attributing performance gains to a better operator.

1 Introduction

Spectral clustering represents a graph through low-frequency Laplacian eigenvectors and clusters the resulting node embeddings (Shi and Malik, 2000; Ng et al., 2001; von Luxburg, 2007). Its poor performance on sparse, heterogeneous graphs has motivated a succession of more elaborate operators: regularized Laplacians (Qin and Rohe, 2013; Zhang and Rohe, 2018), graph powering (Abbe et al., 2020), non-backtracking and Bethe–Hessian constructions (Krzakala et al., 2013; Saade et al., 2014; Dall'Amico et al., 2021; Stephan and Zhu, 2025), and methods designed for directed graphs (Rohe et al., 2016; Sevi et al., 2026). We revisit a more basic question: before replacing the operator, is the classical baseline actually extracting the spectral subspace it is intended to use?

On several standard benchmarks, the answer is no. By literal bottom-k we mean the k eigenvectors associated with the k smallest eigenvalues, including exact zeros. Sparse graph construction and preprocessing can leave many disconnected components, each contributing one structural zero mode to the normalized Laplacian. If their number is d_0, the spectrum begins as

$$0 = \lambda_1 = \cdots = \lambda_{d_0} \;<\; \lambda_{d_0+1} \le \lambda_{d_0+2} \le \cdots ,$$

so when $d_0 \ge k$ the literal readout $[v_1, \ldots, v_k]$ and the first k non-null modes $[v_{d_0+1}, \ldots, v_{d_0+k}]$ do not overlap at all: the entire embedding budget is spent on structural zero modes before the algorithm reaches the positive spectrum. On the four citation benchmarks that motivate our study, d_0 runs into the hundreds while the requested embed-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

ding dimension is a single digit. The useful low-frequency eigenvectors have not disappeared; the baseline simply never selects them (Figure 1).

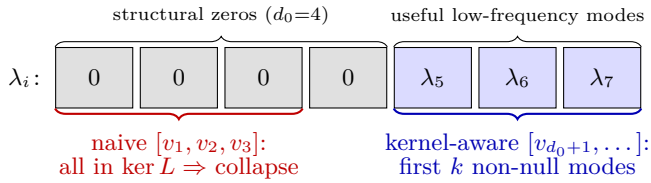

Figure 1: The indexing correction, at a glance. A sparse graph whose symmetric support has $d_0 = 4$ connected components and a $k = 3$ embedding budget. Literal bottom-k selects $[v_1, v_2, v_3]$, all inside $\ker L$; when the evaluated core lies in one component this embedding collapses to a single point after row normalization (Proposition 2). Kernel-aware indexing instead selects the first k non-null modes $[v_{d_0+1}, \dots, v_{d_0+k}]$. Same graph, same Laplacian, same k, same row normalization, same k-means: only the selected indices change.

This indexing error has a precise geometric consequence. For the generalized spectral-clustering family considered here, connectivity fixes the nullity and the component-support decomposition of the zero eigenspace exactly, and the operator weights σ then fix its normalized basis (the $\sqrt{\sigma}$-weighted component indicators); the component partition itself a single connected-components computation recovers (with no random-graph assumption, eigengap, or eigendecomposition). When the evaluated vertices lie in one connected component and $d_0 \geq k$, the bottom-k embedding has rank at most one on those vertices, so row normalization maps all nonzero rows to the same point and no within-component structure survives for k-means. We then treat the partial regime $d_0 < k$, where a basis-invariant null-occupancy statistic $\eta_i \in [0, 1]$ measures how strongly the structural modes still distort each row. Both results point to the same minimal correction: skip the structurally known nullspace and cluster the first k non-null eigenvectors, changing only which modes are selected (Figure 1). On the benchmarks where literal bottom-k is degenerate, this indexing change alone restores strong classical clustering, robustly across standard external metrics. It also does not require computing the discarded d_0 modes: because the nullspace is known from connected components, a component-wise projector constrains a sparse eigensolver to its orthogonal complement and returns only the k non-null modes, in $O(Nk)$ auxiliary working memory beyond the sparse operator.

Contributions. Our contributions are threefold.

- Diagnosis and theory (§2). We characterize the structural nullspace exactly from graph connectivity, prove complete within-component embedding collapse when it fills the spectral budget (evaluated vertices in one component), and quantify partial distortion below capacity through a basis-invariant null-occupancy statistic.
- Benchmark correction and reassessment (§§3, 5). The diagnosis yields a minimal same-operator correction (K_0: skip the structural zero modes, take the first k non-null modes, nothing else changed). Controlled experiments (capacity sweeps, factorized graph controls, and directed and undirected real networks) give evidence that structural nullity, rather than the tested alternatives (directionality, heterophily, reciprocity), drives the observed gap; under a strictly protocol-matched comparison K_0 then explains most of the apparent gap to more sophisticated operators, which shrink to small, graph-dependent refinements.
- Scalable computation (§4). Because the nullspace is known from connected components, it need not be computed: we constrain a sparse eigensolver directly to the non-null subspace, avoiding the cost of the discarded zero modes.

The component–nullspace identity itself is classical. The contribution of this paper is to show that, in modern sparse-graph benchmarks, failing to account for it can dominate the measured performance of the classical baseline and distort comparisons with more sophisticated spectral methods; Section 6 places this finding in the broader literature.

2 When Structural Nullspaces Break Bottom-k

The introduction argued that the apparent failure of bottom-k originates in the structural nullspace rather than in the positive spectrum. We now formalize this mechanism in three steps: we identify the structural nullspace exactly (§2.1), show that it forces complete geometric collapse once it fills the embedding budget (§2.2), and quantify the partial distortion that remains when some non-null modes survive (§2.3).

2.1 The structural nullspace is known from graph topology

The key fact is familiar for the ordinary symmetric normalized Laplacian $L_{\text{sym}} = I - D^{-1/2}WD^{-1/2}$: each connected component contributes exactly one zero eigenvector, the degree-weighted indicator of that component. We first show that the same structural statement (nullity equals the number of components, with component-indicator null vectors) holds verbatim for the generalized spectral-clustering (GSC) operator our benchmarks use (Sevi et al., 2026), including on directed graphs. A reader interested only in L_{sym} may read $\bar{L}_\nu$ as L_{sym} throughout this section and lose nothing essential.

To state the result in general form, let $\mathcal{G} = (V, E, W)$ be a directed graph on N vertices with nonnegative weights $W \in \mathbb{R}^{N\times N}_{\geq 0}$ and random-walk matrix $P = D_{\text{out}}^{-1}W \geq 0$ (so $P\mathbf{1} = \mathbf{1}$). As in the benchmark implementations, we form $W^+ = W + \text{diag}(s)$ with unit loops $s_i = 1$ on every vertex of zero in- or out-degree and recompute all degrees, $P = D_{\text{out}}^{-1}W^+$, p_{out}, and p_{in} from W^+: a loop on a zero-out-degree vertex makes P row-stochastic, while one on a zero-in-degree vertex keeps the endpoint degree measures entrywise positive. This may leave several disconnected components in the symmetric support graph $A_{\text{sym}} = \mathbf{1}\{W^+ + (W^+)^\top > 0\}$ (taken after the loop additions). For a positive measure $\nu \in \mathbb{R}^N_{>0}$ let $\xi = P^\top\nu$, $\sigma = \nu + \xi$, and

$$L_\nu \;=\; D_\sigma - D_\nu P - P^\top D_\nu, \qquad \bar{L}_\nu \;=\; D_\sigma^{-1/2}\, L_\nu\, D_\sigma^{-1/2}. \tag{1}$$

On an undirected graph with $\nu \propto d$ (the degree vector; any positive scalar rescaling cancels in $\bar{L}_\nu$), $\bar{L}_\nu$ reduces to the familiar symmetric normalized Laplacian $I - D^{-1/2}WD^{-1/2}$. For each connected component C_c of A_{sym}, define the $\sqrt{\sigma}$-weighted indicator

$$z_c = \sigma^{1/2}\,\mathbf{1}_{C_c} \,/\, \|\sigma^{1/2}\,\mathbf{1}_{C_c}\|_2, \qquad Z = [z_1, \ldots, z_r] \in \mathbb{R}^{N\times r}, \quad Z^\top Z = I_r. \tag{2}$$

Theorem 1 (Exact structural nullspace). For every $\nu \in \mathbb{R}^N_{>0}$, $\ker \bar{L}_\nu = \text{range}(Z)$, and hence $\dim\ker \bar{L}_\nu = r$, the number of connected components of A_{sym}.

Proof sketch. The Dirichlet form $y^\top L_\nu y = \sum_{ij} \nu_i P_{ij}\,(y_i - y_j)^2 \geq 0$ shows $L_\nu, \bar{L}_\nu \succeq 0$ (so every non-null eigenvalue is strictly positive) and vanishes iff y is constant along every edge, hence constant on each connected component of A_{sym}; normalization turns those constants into the weighted indicators Z. The two-inclusion proof is in Appendix A. □

Thus graph connectivity determines the nullity and the component supports of a canonical zero-mode basis before any eigendecomposition: a standard connected-components computation (e.g. union–find) recovers this structure, after which the operator weights σ determine the normalized basis Z. An arbitrary basis of the degenerate nullspace may rotate these component indicators, but the subspace itself is exact, with no random-graph assumption or eigengap. Storage of Z and its numerical-nullity checks across thresholds are deferred to Appendices B and C. Theorem 1 identifies exactly the subspace from which the structural modes selected by literal bottom-k can come; we next ask what selecting them does to the embedding geometry.

2.2 Saturated regime: exact collapse

Let (λ_j, v_j) be the eigenpairs of $\bar{L}_\nu$ in non-decreasing order and $d_0 = |\{j : \lambda_j = 0\}| = r$. The naive (literal bottom-k) and kernel-aware readouts are

$$U_{\text{naive}} = [v_1, \ldots, v_k], \qquad U_{\text{KA}} = [v_{d_0+1}, \ldots, v_{d_0+k}]. \tag{3}$$

If $d_0 \geq k$, every column of U_{naive} lies in $\ker \bar{L}_\nu = \text{range}(Z)$, and no non-null mode enters the readout at all. The mechanism is then simple. Inside a single component every selected null vector is the same $\sqrt{\sigma}$-weighted indicator up to a scalar, so two vertices' embedding rows differ only by a positive per-vertex scale; row normalization divides that scale out and maps every vertex to one shared point. The collapse is therefore caused not by rank at most one alone but by the positive-scalar factorization followed by row normalization.

Proposition 2 (Naive-embedding collapse). Let $C \subseteq C_c$ for some component C_c of A_{sym}. If $d_0 \geq k$, then

$$U_{\text{naive}}\big|_C = \big(D_\sigma^{1/2}\mathbf{1}_C\big)\, a^\top \quad \text{for some } a \in \mathbb{R}^k, \tag{4}$$

so $\text{rank}\big(U_{\text{naive}}\big|_C\big) \leq 1$. The factorization has a positive scalar per row, $(D_\sigma^{1/2}\mathbf{1}_C)_i = \sqrt{\sigma_i} > 0$; hence if $a \neq 0$, row-normalization sends every row of $U_{\text{naive}}\big|_C$ to the same unit vector $a/\|a\|$, and if $a = 0$ every core row is zero (row-normalized to 0 under the convention $\mathcal{R}(0) = 0$). In either case all row-normalized core rows coincide.

Proof sketch. When $d_0 \geq k$, all selected eigenvectors lie in range(Z). Restricted to $C \subseteq C_c$, every component indicator vanishes except that of C_c, so all core rows are positive scalar multiples of a common vector; row normalization then makes them identical. The full proof is in Appendix A. □

Consequently the k-means objective contains no within-C graph geometry; any resulting partition is determined by numerical or tie-breaking behavior rather than the null embedding itself.

In short,

$$\boxed{d_0 \geq k \implies U_{\text{naive}}\big|_C = (\sqrt{\sigma_i})_{i\in C}\, a^\top \xrightarrow{\mathcal{R}} \text{identical rows} \implies \text{no within-}C\text{ geometry.}}$$

The non-null low-frequency eigenvectors have not vanished: the operator still carries them at positions $d_0+1, \ldots, d_0+k$; the algorithm simply never reaches them. This gives an exact failure criterion when the nullspace fills the budget, but it does not explain graphs such as PolBlogs, where $d_0 < k$ and yet structural modes still dominate the embedding. For those we must measure how much of each row lies in the nullspace.

2.3 Below capacity: null occupancy

Nullity alone quantifies the distortion only at saturation; in general what matters is the fraction of each row contributed by the nullspace. Let $\Pi_0 := ZZ^\top$ be the projector onto $\ker \bar{L}_\nu$ and $U = U_{\text{naive}}$. For any embedding U and each vertex i with nonzero row, define the null-occupancy

$$\eta_i \;:=\; \frac{\|e_i^\top \Pi_0 U\|_2^2}{\|e_i^\top U\|_2^2} \;\in\; [0,1], \tag{5}$$

so that $\eta_i = 1$ means row i is entirely null-supported and $\eta_i = 0$ means it has no nullspace contribution. η_i is invariant to orthogonal rotations of the nullspace basis and to rotations of the non-null embedding block: it depends only on the subspaces range(Π_0) and range(U), not on the returned eigenvector coordinates.

Proposition 3 (Below-capacity distortion bound). Assume $d_0 < k$. Let $C \subseteq C_c$ for some connected component C_c of A_{sym}, and let $U = U_{\text{naive}}$ be the literal bottom-k embedding, restricted to vertices $i \in C$ with $\|e_i^\top U\| > 0$. Up to an orthogonal change of embedding coordinates, the row-normalized core embedding of such a vertex decomposes as

$$\hat{u}_i \;=\; \big(\; \underbrace{\sqrt{\eta_i}\, b}_{\substack{\text{shared null}\\ \text{direction}}} \;,\; \underbrace{\sqrt{1-\eta_i}\, s_i}_{\substack{\text{vertex-specific}\\ \text{non-null geometry}}} \;\big), \quad \|b\| = \|s_i\| = 1, \tag{6}$$

where $b \in \mathbb{R}^{d_0}$ is common to every $i \in C$ and $s_i \in \mathbb{R}^{k-d_0}$ is the unit direction of the non-null block for vertex i. Fix $\varepsilon \in [0,1]$. If $\eta_i, \eta_j \geq 1-\varepsilon$ for any two such vertices $i, j \in C$, then

$$\|\hat{u}_i - \hat{u}_j\|_2 \;\leq\; 2\sqrt{\varepsilon}. \tag{7}$$

Proof sketch. Rotate the embedding coordinates so that the first d_0 axes span the structural nullspace. Within a component the null block points in one common direction b, and only the non-null block s_i varies across vertices; the bound follows from the worst case, in which the two non-null directions oppose each other. The full argument (an orthogonal completion $UR = [Z, V]$ and a principal-angle bound) is in Appendix A. □

Thus η_i controls how much of the normalized row is forced into the shared direction b: if every nonzero row of C has $\eta_i \geq 1 - \varepsilon$, the embeddings lie in a ball of radius $2\sqrt{\varepsilon}$ around $u_0 = (b, 0)$ and k-means sees essentially one direction (whether that non-null direction encodes labels is a separate, empirical question). Proposition 2 is the saturated limit: at $d_0 \geq k$ every nonzero-row vertex has $\eta_i = 1$ exactly, and the ball collapses to a point.

Combining Theorem 1 with Propositions 2 and 3 gives three regimes:

$d_0 \geq k$	$\Longrightarrow$	exact structural collapse,
$d_0 < k,\ \eta \approx 1$	$\Longrightarrow$	approximate collapse (ball of radius $2\sqrt{1 - \min \eta}$),
$d_0 < k,\ \eta \approx 0$	$\Longrightarrow$	structural block is geometrically weak.

Whether the remaining non-null modes encode useful clustering information in the third regime is an empirical question that Section 3 tests on real graphs.

3 The Missing Baseline: Kernel-Aware Bottom-k

Section 2 identifies exactly what literal bottom-k can get wrong. We now ask the simplest empirical question: what happens if we change only the spectral indices and leave everything else untouched? Define the corrected classical readout

$$K_0 \;:=\; \left[v_{d_0+1}, \ldots, v_{d_0+k}\right] \quad \text{of } L_{\text{sym}} \text{ on the binary symmetric support } A_{\text{sym}}, \tag{8}$$

where $d_0 = |\text{components}(A_{\text{sym}})|$ is known from graph topology without any eigendecomposition (this requires $1 \leq k \leq N - d_0$ non-null modes to exist, which holds on every graph we study). When $k < N - d_0$ and the cutoff is strict, $\lambda_{d_0+k} < \lambda_{d_0+k+1}$, K_0 is the unique k-dimensional invariant subspace spanned by the first k non-null modes (strict on all six headline graphs; cutoff-gap audit in Appendix C). At a tied cutoff $\lambda_{d_0+k} = \lambda_{d_0+k+1}$ the k-plane is not unique, so K_0 is set-valued; our fixed numerical implementation returns one representative k-plane. When $k = N - d_0$, K_0 is the entire non-null positive subspace, with no upper cutoff. Only the selected eigenvectors change.

This correction should not be confused with restricting the graph to its largest connected component. LCC preprocessing changes the graph by deleting vertices; K_0 keeps the graph fixed and changes only which part of its spectrum is read. The two coincide only under particular component-spectrum orderings (§6 and Appendix C).

We measure agreement with the benchmark labels by Adjusted Mutual Information (AMI; Vinh et al., 2009), which corrects mutual information for agreement expected by chance; Appendix C confirms that NMI and ARI lead to the same qualitative conclusions. Table 1 compares literal bottom-k with K_0 under the same graph, operator, embedding dimension, row normalization, and k-means protocol: only the spectral indices differ.

	k	d_0	d_0/k	median η	Naive AMI	K_0 AMI	Δ AMI
PolBlogs	2	1	0.50	1.000	0.001	0.839	+0.838
Email-Eu	42	1	0.02	0.035	0.701	0.698	−0.003
Cora	7	78	11.14	1.000	0.153	0.505	+0.352
Citeseer	6	438	73.00	1.000	0.344	0.546	+0.202
Cora-ML	7	61	8.71	1.000	0.005	0.466	+0.461
Citeseer-full	6	515	85.83	1.000	0.118	0.691	+0.573

Table 1: Same-operator indexing correction. Naive uses literal bottom-k; K_0 skips the d_0 structural zero modes. Both columns use the same core-restricted clustering protocol (five-seed k-means, $n_{\text{init}} = 200$), so only the spectral indices differ; AMIs are five-seed means (across-seed SD ≤ 0.005; per-seed values saved by the reproduction harness). Median η is null occupancy on the core. Protocol variants (incl. full-embedding naive = 0.000), metric (NMI/ARI) robustness, and baseline comparators are in Appendix C.

The six graphs separate naturally into two regimes. On Cora, Citeseer, Cora-ML, and Citeseer-full, $d_0 \geq k$ and the evaluated core lies in one connected component, so Proposition 2 predicts degeneracy before any labels are inspected: the literal bottom-k core embedding has rank ≤ 1 and the implementation returns a degenerate partition (full-embedding AMI = 0.000; core-restricted tie-breaking 0.005–0.344). Re-indexing the same Laplacian to the first k non-null modes restores AMI between 0.47 and 0.69. The signal was not destroyed by sparsity; literal bottom-k simply never reached it.

The two below-capacity graphs show why nullity alone is not enough. PolBlogs and Email-Eu both have $d_0 = 1 < k$, yet the correction is decisive on PolBlogs ($0.001 \to 0.839$) and inert on Email-Eu ($0.701 \to 0.698$). Null occupancy separates the cases: on PolBlogs 95% of evaluated vertices have $\eta \geq 0.999$, so Proposition 3 places almost the entire core within radius $2\sqrt{1 - 0.999} \approx 0.06$ of a common direction, whereas Email-Eu has median $\eta = 0.035$ and the null block contributes little. Section 5 tests this continuous picture under controlled perturbations; we do not claim η as a universal quantitative predictor.

Finally, the correction is structural rather than threshold-based. We obtain d_0 from connected components, not from a floating-point eigenvalue tolerance, and the two agree across thresholds $\varepsilon \in \{10^{-10}, \ldots, 10^{-4}\}$ on every benchmark graph (Appendix C). The dense implementation is therefore a one-line change in spectral indexing, from evecs[:, :k] to evecs[:, r:r+k]. Its only remaining drawback is computational: when d_0 is large, computing $d_0 + k$ eigenvectors merely to discard the first d_0 is wasteful. Section 4 removes that cost.

4 Skipping the Kernel Without Computing It

Section 3 gives the correct indices, but a literal implementation is wasteful: when d_0 is large, computing $d_0 + k$ eigenvectors only to discard the first d_0 can dominate the cost. Because Theorem 1 gives the unwanted subspace range(Z) exactly, we can skip it before eigensolution, and there are two equivalent ways to express this. Conceptually, the kernel-aware modes solve the constrained Rayleigh–Ritz problem

$$\min_{X \in \mathbb{R}^{N \times k}} \ \operatorname{tr}(X^\top L_{\text{sym}} X) \quad \text{s.t.} \quad X^\top X = I_k, \ \ Z^\top X = 0 \tag{9}$$

(for $1 \leq k < N - r$ with $\lambda_{r+k} < \lambda_{r+k+1}$ the minimizer spans $\{v_{r+1}, \ldots, v_{r+k}\}$; the boundary $k = N - r$ needs no upper gap). Computationally, we realize this constraint with a spectral shift rather than a projection: conjugating by $Q = I - ZZ^\top$ leaves $QL_{\text{sym}}Q = L_{\text{sym}}$ and so does not remove the zero modes, whereas, since $L_{\text{sym}}Z = 0$,

$$L_\omega \ = \ L_{\text{sym}} + \omega\, ZZ^\top, \qquad \omega > 2, \tag{10}$$

keeps every positive eigenpair of L_{sym} but lifts the r zero modes to ω, above the $[0, 2]$ spectrum of a normalized Laplacian; the k smallest eigenpairs of L_ω are then exactly $v_{r+1}, \ldots, v_{r+k}$. Constrained eigensolvers are themselves standard (LOBPCG (Knyazev, 2001), for instance, accepts an explicit orthogonality constraint), so the point here is not

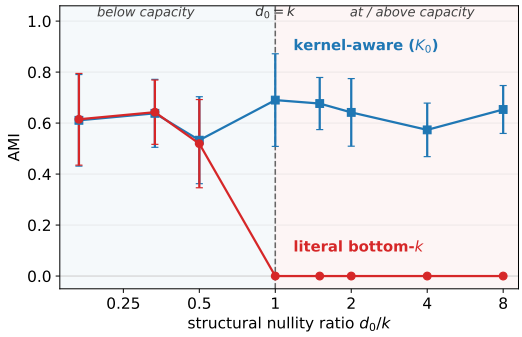

Figure 2: The predicted capacity transition. We hold the community-bearing DC-SBM core fixed and vary only the number of disconnected components. Below $d_0/k = 1$, literal bottom-k retains non-null modes and agrees with K_0; at the threshold $d_0 = k$ its embedding becomes entirely null-supported and AMI collapses to zero, while K_0 remains stable. Error bars: ± 1 s.d. across 25 paired (core-seed, KMeans-seed) runs. Reproducible by activation_sweep_d0_over_k.py.

constrained eigensolving but that the special component-indicator structure of Z lets the constraint be applied implicitly: because its columns have disjoint support, $ZZ^\top = I - Q$ acts in $O(Nk)$ per iteration through $O(N + r)$ component metadata alone, never forming or storing an $N \times r$ basis (Appendix B).

Empirically this removes the cost of the discarded modes. On DAG-chain synthetics at $N = 20{,}000$ (where d_0 ranges from ≈ 1900 to ≈ 4100 across seeds), the shifted-operator solve (LOBPCG, $\omega = 2.1$) returns the corrected embedding in 0.5–0.95 s; where dense eigh was executed ($N \leq 10{,}000$) its recovered positive eigenpairs agree, and across all routes the largest observed AMI difference was 0.007. The large wall-clock gaps we report thus follow from computing only the k wanted modes rather than $d_0 + k$, together with the $O(Nk)$ projector, not from a new eigensolver: dense eigh becomes impractical beyond $N \approx 15{,}000$, and an unconstrained sparse solver that requests $d_0 + k$ modes slows sharply or triggers a dense fallback once $5(k + d_0) > N$. Per-seed timings, the constraint/eigenresidual diagnostics, and the scaling figure are in Appendix B.

5 Stress-Testing the Mechanism

Having established both the corrected readout (§3) and a scalable way to compute it (§4), we now ask whether structural nullity is actually the mechanism behind the observed benchmark failures. Section 3 showed that kernel awareness corrects the six graphs, but not why. Three alternatives remain: (A1) the transition at $d_0 = k$ is accidental; (A2) the effect tracks other sparsity-correlated properties (heterophily, directionality, reciprocity); (A3) kernel awareness is merely another regularizer. We test each, then ask what remains.

5.1 Does failure activate at the predicted capacity?

To isolate the capacity variable we hold the community-bearing graph fixed and vary only the number of disconnected components: into each of five undirected DC-SBM cores ($N_{\text{core}} = 1500$, $k = 6$) we inject $d_0 - 1$ isolated singletons to reach $\gamma := d_0/k \in \{0.17, 0.33, 0.5, 1, 1.5, 2, 4, 8\}$, holding the community-bearing core fixed; only the number of added singleton components varies, so total graph size changes accordingly.

At the first saturated point $d_0 = k$, the empirical collapse predicted by Proposition 2 occurs. Below capacity the naive readout keeps $k - d_0$ non-null directions with enough geometry for k-means, so naive and K_0 agree; at $\gamma = 1$ the bottom-k eigenvectors span $\ker L_{\text{sym}}$, the embedding collapses to rank ≤ 1, and naive AMI drops to exactly 0.000. The sharp step is specific to this controlled construction.

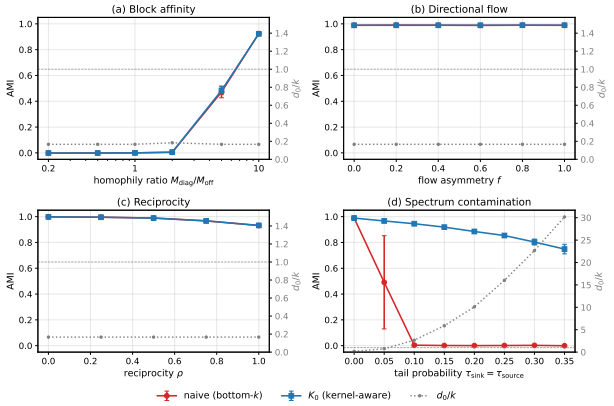

Figure 3: Axis isolation on the factorized DC-SBM testbed. Only Axis 3 (spectrum contamination) produces a material naive-to-K_0 correction: all realizations are above capacity by $\tau = 0.10$, where AMI(naive) collapses while K_0 stays ≥ 0.74. Axes (a) block affinity, (b) directional flow, (c) reciprocity hold $d_0/k \approx 0.17$ fixed, and the readouts agree to $|\Delta\text{AMI}| \leq 0.012$ (we call ≤ 0.02 immaterial). Solid: AMI; dotted (right axis): d_0/k; error bars ± 1 s.d. over 10 seeds. Reproducible by axis_isolation_sweep.py.

5.2 Does the failure track structural nullity rather than common confounders?

We vary four properties one at a time on a factorized DC-SBM testbed (Appendix E): block affinity, directional flow, reciprocity, and structural contamination ($\tau_{\text{sink}} = \tau_{\text{source}}$). The first three introduce no forced isolates and preserve the symmetric block-level support topology relevant to A_{sym}, and empirically their realized d_0/k stays approximately constant; only contamination is designed to systematically enlarge $\ker L_{\text{sym}}$. The test is simple: if structural nullity is the cause, only the axis that changes d_0/k should open a naive-to-K_0 gap.

Only the contamination axis, which changes d_0/k, opens a substantial naive-to-K_0 gap; block affinity, flow, and reciprocity leave d_0/k approximately fixed, stay within $|\Delta\text{AMI}| \leq 0.012$, and keep K_0 strong throughout: controlled evidence that heterophily, directionality, and reciprocity do not by themselves open the gap. At the first mixed point ($\tau = 0.05$), null occupancy η correlates with the gap (Spearman $\rho_s = 0.82$), consistent with but not implied by Proposition 3. As a real-graph negative control, the connected Coauthor CS graph (Shchur et al., 2018) has low occupancy and near-identical naive/K_0 AMI (0.550 vs 0.564): kernel awareness does not generically help when the pathology is absent.

5.3 What remains after fixing the baseline?

Once the baseline is corrected, how much value remains in more sophisticated operators and readouts? We add one form of complexity at a time while retaining kernel awareness. Here, K_0 is the proposed correction; K_1, K_2, and S_3 are diagnostic yardsticks that measure the headroom remaining after the fix. The four-rung ladder is:

$$\underbrace{K_0}_{\text{fix indexing}} \longrightarrow \underbrace{K_1}_{\text{change operator}} \longrightarrow \underbrace{K_2}_{\text{tune its measure}} \longrightarrow \underbrace{S_3}_{\text{change readout}} .$$

Here K_1 is canonical GSC ($\beta = 1/2$), K_2 tunes the measure to a per-graph $\beta^\star$, and S_3 is a Slepian readout on the post-nullspace band, also at $\beta^\star$. Because $\beta^\star$ is supervised and benchmark-tuned, K_2 and S_3 only estimate residual headroom. Even granting these rungs (full ladder in Table 5), the mean AMI above K_0 moves only $\boxed{0.624 \to 0.641\ (+0.017)}$, so the same-operator indexing fix already recovers most of the sophisticated pipeline's advantage. A strictly protocol-matched comparison (same core, same five-seed k-means for every method) confirms this without mixing protocols. Mean AMI over the six graphs:

	SC	reg-SC	BH	Di-SIM$_s$	Di-SIM$_r$	K_0
mean AMI	0.229	0.569	0.558	0.511	0.527	0.624

so K_0 is highest (per-dataset values in Appendix C.4). Nor is the gain from replacing $k-1$ positive modes by k: above capacity, $k-1$, k, and $k+1$ non-null modes all far exceed literal bottom-k (PolBlogs, below capacity, is the exception, needing the full k; Appendix C.5).

Taken together, these experiments (six benchmark graphs, a controlled capacity sweep, factorized synthetic controls, a large connected negative control, a twelve-graph prevalence audit in Appendix C.6, and SNAP audits) support one interpretation: structural nullity is sufficient to reproduce the naive-to-K_0 failure gap and is the only tested axis that opens such a gap, and correcting it narrows the apparent advantage of sophisticated constructions rather than eliminating them.

6 Related Work

Most prior responses to the failure of spectral clustering on sparse, heterogeneous graphs modify the operator itself: regularization (Qin and Rohe, 2013; Le et al., 2017; Zhang and Rohe, 2018), graph powering (Abbe et al., 2020), non-backtracking / Bethe–Hessian rescalings (Krzakala et al., 2013; Saade et al., 2014; Dall'Amico et al., 2021; Stephan and Zhu, 2025), and direction-aware constructions (Leicht and Newman, 2008; Rohe et al., 2016; Sevi et al., 2026; Debaussart-Joniec et al., 2026), all building on classical spectral clustering (Shi and Malik, 2000; Ng et al., 2001; von Luxburg, 2007). We ask an earlier question: before replacing the operator, has the baseline extracted its intended non-null subspace? These mechanisms are complementary: they reposition informative modes (Dall'Amico et al., 2019), whereas we remove an exact structural nullspace at the nominal bottom-k positions.

Handling disconnected components is itself well known: Dall'Amico et al. (2021) detect components from the zero eigenspace and treat each independently, and Zhang and Rohe (2018) delete everything outside the largest connected component. Our point is not that this structure is unknown, but that a common literal bottom-k implementation does not apply it, changing measured outcomes. K_0 differs from such preprocessing in changing only the spectral search space, not deleting vertices (they coincide only under special component-spectrum orderings; Appendix C). Related methods alter the readout not the operator (Lin and Cohen, 2010; Van De Ville et al., 2017; Tremblay et al., 2016), but here the excluded subspace is known exactly from topology.

Our contribution is closest to a third line, benchmark diagnosis: consistency results characterize when non-null eigenvectors recover latent structure (Lei and Rinaldo, 2015; Cape et al., 2019) and the connected-components/nullspace identity is classical (Chung, 1997, Ch. 1), but our finite-graph observation is that failing to account for this kernel can dominate benchmark outcomes, a demystification in the spirit of Arnaiz-Rodriguez and Errica (2026).

7 Discussion

Spectral benchmarks conflate two questions: whether an operator carries useful low-frequency structure, and whether the implementation extracts it. Literal bottom-k can fail the second, so a method gap need not signal an operator-level advantage; kernel awareness is a target-aware audit, not a blanket rule to discard zero modes.

Our analysis has boundaries. We treat exact structural zeros; weakly coupled regions give only near-zero modes without an exact projector. Removing the kernel is not universally beneficial, and below capacity low η need not yield good clustering. We assume k known, and the supervised $\beta^\star$ of K_2/S_3 only bounds residual headroom, claiming no label-free gain beyond K_0.

Use of Large Language Models

In accordance with the ICLR 2027 policy on the use of AI, we disclose the following. The research questions, the structural-nullspace characterization (Theorem 1) and its corollaries, the embedding-collapse and below-capacity propositions (Propositions 2 and 3), the K_0 correction and its implicit sparse-solver form, and the overall proof strategies were conceived by the authors. We did not use generative AI tools to develop these theoretical models or conceptual frameworks, to formulate the mathematical claims or supply critical ingredients for their proofs, to propose or refine hypotheses, or to interpret the results, and we did not use them to produce any reported result that the authors did not independently verify. Machine translation and qualitative or thematic data analysis are not applicable to this work.

We used a large language model (LLM)-based assistant in the following supporting roles, each under author supervision and verification:

- Writing. Editing exposition for readability and hierarchy; drafting prose, figure and table captions (including the schematic in Figure 1) and, from author-specified strategies, the write-up of the proofs; consistency-checking cross-references, notation, and claims across the manuscript; and formatting the references.
- Methods and research execution. Implementing, editing, and debugging the experimental and verification code and the author-specified synthetic generators; running and re-running the experiments and scaling measurements; and assembling the one-command reproduction harness and checksum manifest so that the reported quantities are reproducible.
- Verification and repair. Checking derivations and cross-checking reported numbers against the code and saved result files, which led us to report backward-error-sensitive quantities as configuration-independent bounds rather than exact digits.

We have reviewed all AI-assisted work: every LLM-assisted proof write-up, code artifact, experiment, and reported number was independently checked and approved by the authors. We take responsibility for the final content of this work, including all text, claims, and artifacts produced with the aid of generative AI.

Ethics statement

This paper studies benchmark protocol and eigenvalue-indexing for sparse spectral clustering on standard public graph datasets. No human subjects are involved. The authors are aware of no ethical concerns regarding the datasets, methods, or conclusions of this work.

Reproducibility statement

Every empirical claim in the paper is reproducible from the scripts we cite in situ: the six-graph headline (naive/K_0 AMI and median η in Table 1) from frozen_headline.py; the kernel-aware ladder (Table 5) from kernel_aware_ladder.py; the phase diagram (Fig. 5) from synthetic_phase_diagram.py; the undirected phase diagram and Coauthor CS control (Fig. 6) from undirected_dcsbm_phase.py and coauthor_cs_test.py; the activation curve (Fig. 2) from activation_sweep_d0_over_k.py; the axis-isolation sweep (Fig. 3) from axis_isolation_sweep.py, and the below-capacity seed split (Table 16) from axis_isolation_below_capacity.py; the NMI/ARI metric-robustness table (Table 10) and the NMI/ARI Spearman correlations from metric_robustness.py; the classical small-network pilot (Table 12) from pilot_four_networks.py; the SNAP audit (Table 13) from snap_audit.py; the structural-Q scaling table (Table 2) from scaling_structural_projector.py and scaling_structural_tail.py; and the Core-first appendix (Table 20) from corefirst_5seed.py. All scripts use fixed seeds $\{42, 43, 44, 45, 46\}$ at $n_{\text{init}} = 200$ where applicable. Theorem 1's proof is given in full in Appendix A.

Protocol note for the naive readout. Two protocol variants of "naive bottom-k" agree on the operator and the eigenvectors but differ in when the labeled core C enters: (a) fit k-

means on the full N-vertex row-normalized embedding and evaluate the resulting labels on C (full-embedding); (b) restrict the row-normalized embedding to C and fit k-means on $|C|$ rows (core-restricted). Under Proposition 2 both variants partition C without graph information, but the specific AMI number they return differs. Table 1 is core-restricted for both the naive and the K_0 columns (from frozen_headline.py), so only the spectral indices differ between them; its naive AMIs are $\{0.001, 0.701, 0.153, 0.344, 0.005, 0.118\}$. The full-embedding naive values, which collapse to exactly 0.000 on the above-capacity graphs, reproducing the benchmark degeneracy whose interpretation we reassess, are reported separately in Appendix C (run_all_baselines_n200.py); the $\eta \equiv 1$ diagnosis and the selected modes are identical under both protocols, but the fitted K_0 AMI is not (full-embedding fitting moves the six-graph mean from 0.624 to 0.605), so we use the core-restricted protocol for the headline comparison. Both the "naive AMI" figure inherited from prior benchmarks and the fitted K_0 AMI are thus protocol-dependent; the nullspace mechanism is not. The complete source code and the one-command reproduction harness described below will be provided as anonymous supplementary material during the review period and released publicly upon publication under a license permitting research use.

Reproduction harness and manifest. The harness is a one-command script (reproduce.sh) with requirements.txt, a two-group MANIFEST.sha256, and the vendored input graphs (data/) and modules needed to run from a clean checkout. Its default run (six gated stages, $\approx$ 6 min, BLAS threads pinned) asserts: (1) the pinned environment; (2) artifact integrity, SHA-256 hard-checking the 19 bit-stable files (raw input graphs, SNAP edge lists, vendored code) and reporting the result files as an informational snapshot; (3) the Table 1 correction, re-running the pipeline and checking each per-graph K_0 against the paper, the PolBlogs $0.001 \to 0.839$ jump, the mean $K_0 = 0.624$, and the structural nullity $d_0 = \#\text{components}(A_{\text{sym}})$ vector $[1, 1, 78, 438, 61, 515]$; (4) the SNAP audit, verifying Fiedler mass on the LCC = 1, cut-membership Jaccard = 1, and a positive simple gap on all six graphs; (5) the scaling constrained solve, checking $\text{rank}(Z)$ equals the numerical nullity, $\|Z^\top U\| < 10^{-4}$, eigenresidual $< 10^{-4}$, and a $> 10\times$ speedup over dense on a fast reduced N-grid; and (6) regeneration of the activation figure with its $\gamma = 1$ collapse. Setting REPRO_FULL=1 additionally re-runs the reassessment ladder ($\approx$ 20 min; verified to reproduce $0.624 \to 0.641$ over all six graphs) and the full scaling N-grid (whose $N = 20{,}000$ dense fallback is multi-hour). Following the harness's honesty convention, bit-stable inputs are hash-checked, whereas everything downstream of a dense eigensolver is asserted by configuration-independent bounds and invariants rather than exact trailing digits, and figures are gated by valid generation rather than byte-identity. The one stated exception is Figure 1, a TikZ diagram built by LaTeX (not by the harness); the axis-isolation sweep of Appendix E is gated under REPRO_FULL=1, its generator synthetic_4axis.py now vendored in the bundle. We ran the default six-stage harness end to end on this snapshot and it passes (exit 0); the reassessment ladder is the seventh, REPRO_FULL-gated stage.

Compute environment. All experiments were run on a single machine (Apple M4 Pro, 12 cores, 24 GB RAM, macOS 15.1) with no GPU, using Python 3.12.2, NumPy 1.26.4, SciPy 1.13.1, scikit-learn 1.8.0, NetworkX 3.3, and Matplotlib 3.9.2 (pinned in requirements.txt); BLAS threads are pinned for the timing and invariant checks. The wall-clock scaling numbers of Section 4 are specific to this hardware, whereas the reported verdicts, bounds, and invariants are version-independent.

References

H. Sevi, G. Debaussart-Joniec, M. Hacini, M. Jonckheere, A. Kalogeratos. Generalized Dirichlet energy and graph Laplacians for clustering directed and undirected graphs. Transactions on Machine Learning Research (TMLR), 2026.

L. Dall'Amico, R. Couillet, N. Tremblay. Revisiting the Bethe-Hessian: improved community detection in sparse heterogeneous graphs. NeurIPS, 2019.

L. Dall'Amico, R. Couillet, N. Tremblay. A unified framework for spectral clustering in sparse graphs. Journal of Machine Learning Research, 22(217): 1–56, 2021.

C. M. Le, E. Levina, R. Vershynin. Concentration and regularization of random graphs. Random Structures & Algorithms, 51(3): 538–561, 2017.

K. Rohe, T. Qin, B. Yu. Co-clustering directed graphs to discover asymmetries and directional communities. Proceedings of the National Academy of Sciences, 113(45): 12679–12684, 2016.

E. A. Leicht, M. E. J. Newman. Community structure in directed networks. Physical Review Letters, 100(11): 118703, 2008.

T. Qin, K. Rohe. Regularized spectral clustering under the degree-corrected stochastic blockmodel. NeurIPS, 2013.

Y. Zhang, K. Rohe. Understanding regularized spectral clustering via graph conductance. NeurIPS, 2018.

E. Abbe, E. Boix-Adsera, P. Ralli, C. Sandon. Graph powering and spectral robustness. SIAM J. Math. Data Sci., 2(1): 132–157, 2020. Earlier version: arXiv:1809.04818, 2018.

L. A. Adamic, N. Glance. The political blogosphere and the 2004 U.S. election: divided they blog. Proceedings of the 3rd International Workshop on Link Discovery (LinkKDD), 36–43, 2005.

A. Bojchevski, S. Günnemann. Deep Gaussian embedding of graphs: unsupervised inductive learning via ranking. International Conference on Learning Representations (ICLR), 2018.

F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborová, P. Zhang. Spectral redemption in clustering sparse networks. PNAS 110(52): 20935–20940, 2013.

A. Saade, F. Krzakala, L. Zdeborová. Spectral clustering of graphs with the Bethe Hessian. NeurIPS, 2014.

L. Stephan, Y. Zhu. Community detection with the Bethe-Hessian. COLT, 2025. Full version: arXiv:2411.02835.

F. Lin, W. W. Cohen. Power iteration clustering. ICML, 2010.

G. Debaussart-Joniec, H. Sevi, M. Jonckheere, A. Kalogeratos. Parametrized power-iteration clustering for directed graphs. Forty-third International Conference on Machine Learning (ICML), 2026.

N. Tremblay, G. Puy, R. Gribonval, P. Vandergheynst. Compressive spectral clustering. ICML, 2016.

J. Lei, A. Rinaldo. Consistency of spectral clustering in stochastic block models. Annals of Statistics, 43(1): 215–237, 2015.

J. Cape, M. Tang, C. E. Priebe. The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics. Annals of Statistics, 47(5): 2405–2439, 2019.

A. Arnaiz-Rodriguez, F. Errica. Oversmoothing, "Oversquashing", Heterophily, Long-Range, and more: Demystifying Common Beliefs in Graph Machine Learning. ICLR, 2026.

F. R. K. Chung. Spectral graph theory. American Mathematical Society, CBMS 92, 1997.

A. V. Knyazev. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. SIAM J. Sci. Comput., 23(2): 517–541, 2001.

C. Davis, W. M. Kahan. The rotation of eigenvectors by a perturbation. III. SIAM J. Numer. Anal., 7(1): 1–46, 1970.

D. Van De Ville, R. Demesmaeker, M. G. Preti. When Slepian meets Fiedler: Putting a focus on the graph spectrum. IEEE Signal Processing Letters, 24(7): 1001–1004, 2017.

N. X. Vinh, J. Epps, J. Bailey. Information theoretic measures for clusterings comparison: is a correction for chance necessary? ICML, 2009.

O. Shchur, M. Mumme, A. Bojchevski, S. Günnemann. Pitfalls of graph neural network evaluation. Relational Representation Learning Workshop, NeurIPS 2018. Coauthor CS dataset: https://github.com/shchur/gnn-benchmark.

W. W. Zachary. An information flow model for conflict and fission in small groups. Journal of Anthropological Research, 33(4): 452–473, 1977.

V. Krebs. Political books co-purchase network. Unpublished, distributed with M. E. J. Newman's network data collection, 2004.

M. Girvan and M. E. J. Newman. Community structure in social and biological networks. PNAS, 99(12): 7821–7826, 2002.

D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, S. M. Dawson. The bottlenose dolphin community of Doubtful Sound features a large proportion of long-lasting associations: can geographic isolation explain this unique trait? Behavioral Ecology and Sociobiology, 54(4): 396–405, 2003.

J. Shi and J. Malik. Normalized cuts and image segmentation. IEEE Trans. PAMI, 22(8): 888–905, 2000.

A. Y. Ng, M. I. Jordan, Y. Weiss. On spectral clustering: Analysis and an algorithm. NeurIPS, 2001.

U. von Luxburg. A tutorial on spectral clustering. Statistics and Computing, 17(4): 395–416, 2007.

J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/data, June 2014.

Appendix roadmap. Appendix A contains the theoretical proofs (Theorem 1 and Propositions 2–3). Appendix B gives the scalable eigensolver and its scaling experiments. Appendix C collects benchmark protocols, robustness checks, and comparisons against more sophisticated baselines. Appendix D reports undirected controls, and Appendix E defines the factorized synthetic experiment that isolates structural nullity from competing graph properties. Appendix F documents the calibrated DAG-chain generator, and Appendix G reports the band-width sensitivity sweep that fixes the exploratory S_3 default. Appendix H compares kernel awareness with clustering the labeled core alone, and Appendix I gives the auxiliary stability theory for the exploratory S_3 readout.

A Proofs of the Structural-Nullspace Results

This appendix proves the three structural results of Section 2: the exact nullspace identity (Theorem 1), the saturated collapse (Proposition 2), and the below-capacity distortion bound (Proposition 3). Theorem 1 is the load-bearing structural fact behind the kernel-aware readout of Section 3; we give its setup and proof first, then the two propositions. (The auxiliary stability theory for the exploratory S_3 readout is deferred to Appendix I, as it concerns the diagnostic pipeline rather than the K_0 correction.)

A.1 Setup

Let $\mathcal{G} = (V, E, W)$ be a directed graph with N vertices and nonnegative weights $W \in \mathbb{R}_{\geq 0}^{N \times N}$, self-loop-padded at every vertex with $d_{\text{in}}(i) = 0$ or $d_{\text{out}}(i) = 0$ to keep $P = D_{\text{out}}^{-1} W \geq 0$ well-defined (Section 2). Fix a measure $\nu \in \mathbb{R}_{>0}^{N}$ and let $\xi = P^{\top}\nu$, $\sigma = \nu + \xi$. Denote by $L_\nu = D_\sigma - D_\nu P - P^{\top} D_\nu$ the unnormalized generalized graph Laplacian and by $\bar{L}_\nu = D_\sigma^{-1/2} L_\nu D_\sigma^{-1/2}$ its symmetric normalization (1).

Let $A_{\text{sym}} = \mathbf{1}\{W + W^{\top} > 0\}$ be the binary symmetric hull of the (padded) adjacency and let $\{C_1, \ldots, C_r\}$ be its connected components. For each c define $z_c \in \mathbb{R}^N$ by

$$z_c(i) \;=\; \begin{cases} \sigma(i)^{1/2} \,/\, \|(\sigma(j)^{1/2})_{j \in C_c}\|_2 & \text{if } i \in C_c \\ 0 & \text{otherwise,} \end{cases} \tag{11}$$

and stack the columns into $Z \in \mathbb{R}^{N \times r}$. Columns have disjoint support, so $Z^{\top} Z = I_r$.

A.2 Proof of Theorem 1

Recall the claim: for every $\nu \in \mathbb{R}^N_{>0}$, $\ker \bar{L}_\nu = \mathrm{range}(Z)$, where the columns of Z are the $\sqrt{\sigma}$-weighted indicators of the connected components of A_{sym} (the symmetric hull of W after self-loop padding).

We prove the two inclusions separately.

Proof of $\mathrm{range}(Z) \subseteq \ker \bar{L}_\nu$. Fix a component C_c and its associated column z_c. Because C_c is a connected component of A_{sym}, no edge of W crosses the cut, so $P_{ij} = P_{ji} = 0$ for every $i \in C_c$, $j \notin C_c$. Consequently $L_\nu = D_\sigma - D_\nu P - P^\top D_\nu$ acts block-diagonally with C_c as one block. Inside that block, the usual identity gives

$$\begin{aligned}(L_\nu \mathbf{1}_{C_c})_i &= \sigma(i) - \nu(i)\,(P\mathbf{1}_{C_c})_i - (P^\top D_\nu \mathbf{1}_{C_c})_i \\ &= \sigma(i) - \nu(i) - \xi(i) = 0 \qquad \text{for } i \in C_c,\end{aligned}$$

using $(P\mathbf{1}_{C_c})_i = 1$ (closure under P) and $(P^\top D_\nu \mathbf{1}_{C_c})_i = \sum_{j \in C_c} P_{ji}\nu_j = \xi(i)$ (again by closure, since $P_{ji} = 0$ for $j \notin C_c$, $i \in C_c$). Since $D_\sigma^{-1/2} z_c \propto \mathbf{1}_{C_c}$, we conclude $\bar{L}_\nu z_c = D_\sigma^{-1/2} L_\nu (D_\sigma^{-1/2} z_c) = 0$. Applying this to every column of Z gives $\bar{L}_\nu Z = 0$. □

Proof of $\ker \bar{L}_\nu \subseteq \mathrm{range}(Z)$. Let $x \in \ker \bar{L}_\nu$ and set $y := D_\sigma^{-1/2} x$. Then $L_\nu y = D_\sigma^{1/2} \bar{L}_\nu D_\sigma^{1/2} y = 0$ (one factor of $D_\sigma^{1/2}$ cancels with the $D_\sigma^{-1/2}$ in y's definition, the other with x's membership in $\ker \bar{L}_\nu$). The Dirichlet identity for L_ν is

$$y^\top L_\nu y \;=\; \sum_{i,j} \nu_i\, P_{ij}\, (y_i - y_j)^2, \tag{12}$$

which is a direct expansion using $\sigma_i = \nu_i + \xi_i$, $\nu_i = \nu_i \sum_j P_{ij}$ (row-stochasticity of P), and $\xi_i = \sum_j \nu_j P_{ji}$: the two cross-terms combine into the symmetric squared-difference sum. Because $\nu > 0$ and $P \geq 0$, every term in (12) is non-negative. Since $L_\nu y = 0$ implies $y^\top L_\nu y = 0$, every term must vanish: $y_i = y_j$ whenever $P_{ij} > 0$. Applying the same identity after the dummy-index swap $i \leftrightarrow j$ gives $y_j = y_i$ whenever $P_{ji} > 0$. Together, $y_i = y_j$ whenever $\{i, j\}$ is an edge of A_{sym}. Because A_{sym}'s connected components partition V, y is constant on each component: $y = \sum_c \alpha_c \mathbf{1}_{C_c}$ for some scalars α_c. Converting back, $x = D_\sigma^{1/2} y = \sum_c \alpha_c D_\sigma^{1/2} \mathbf{1}_{C_c}$, so $x \in \mathrm{span}\{D_\sigma^{1/2} \mathbf{1}_{C_c}\}_c = \mathrm{range}(Z)$. □

Combining the two inclusions gives $\ker \bar{L}_\nu = \mathrm{range}(Z)$, hence $\dim \ker \bar{L}_\nu = r = \mathrm{rank}(Z)$. Because $\bar{L}_\nu Z = 0$ and $\bar{L}_\nu$ is symmetric, $(I - ZZ^\top)\,\bar{L}_\nu (I - ZZ^\top) = \bar{L}_\nu$ follows.

Remark. The proof uses only the closed-block property of connected components under the symmetric hull, plus the Dirichlet-form identity. Both are preserved under self-loop padding, which keeps W nonnegative, keeps P row-stochastic after normalization, and adds no cross-component edge; the energy identity (12) therefore still vanishes exactly on functions constant per component, so the nullspace characterization is unchanged on the augmented graph. (A zero-out-degree sink has an empty P-row, so padding it adds only the diagonal $P_{ii} = 1$; the row-rescaling case is a zero-in-degree vertex that already has positive out-degree, whose existing off-diagonal P-row entries are rescaled when the loop enlarges its normalizer. Either way the claim is not that only diagonal terms change, only that the three structural properties above survive.) The rank r of Z therefore equals the number of A_{sym}-components on the augmented graph, computable by one union-find pass in $O(N + \mathrm{nnz})$ time.

A.3 Proof of Proposition 2 (saturated collapse)

Proof. When $d_0 \geq k$, the k smallest eigenpairs of $\bar{L}_\nu$ all lie in $\mathrm{range}(Z)$. Each nullspace vector, restricted to $C \subseteq C_c$, is a scalar multiple of $D_\sigma^{1/2} \mathbf{1}_C$: writing the vector as $\sum_{c'} \alpha_{c'} z_{c'}$, only the c-th term has nonzero support on C, and it restricts to $(\alpha_c / \|\sigma^{1/2} \mathbf{1}_{C_c}\|)\, D_\sigma^{1/2} \mathbf{1}_C$.

Stacking k such restrictions gives the rank-≤ 1 factorization; row-normalization then collapses every nonzero row to $a/\|a\|$ (and every zero row to 0 under the stated convention), so all row-normalized core rows coincide. □

A.4 Proof of Proposition 3 (below-capacity distortion)

Proof. Since $d_0 < k$ and range$(Z) \subseteq$ range(U), put $A = U^\top Z \in \mathbb{R}^{k\times d_0}$, which has orthonormal columns ($A^\top A = I_{d_0}$), and extend it to an orthogonal $R = [A, A_\perp] \in O(k)$. Then $UR = [Z, V]$ with $V = UA_\perp$ spanning the non-null block, which supplies the coordinate decomposition used in the main text. For $i \in C_c$ write $\alpha_i = \sqrt{\sigma_i}/\sqrt{\sum_{j\in C_c}\sigma_j}$ and $v_i = e_i^\top V$; then $\eta_i = \alpha_i^2/(\alpha_i^2 + \|v_i\|^2)$, and in the rotated frame $\hat{u}_i R = (\sqrt{\eta_i}\, e_c,\ \sqrt{1-\eta_i}\, s_i)$ with $s_i = v_i/\|v_i\|$ (any unit vector if $v_i = 0$, whose coefficient is then zero). Because $C \subseteq C_c$, the first d_0 coordinates share the common unit direction $b = e_c$. Set $\sqrt{\eta_i} = \cos\theta_i$; then $\theta_i \leq \arcsin\sqrt{\varepsilon}$ whenever $\eta_i \geq 1-\varepsilon$, and $\langle \hat{u}_i, \hat{u}_j\rangle \geq \cos(\theta_i+\theta_j)$ (worst case: the non-null directions oppose), so $\|\hat{u}_i - \hat{u}_j\|_2^2 = 2 - 2\langle \hat{u}_i, \hat{u}_j\rangle \leq 4\sin^2((\theta_i+\theta_j)/2) \leq 4\varepsilon$. □

B Scaling details

Section 4 claims that the known structural nullspace can be skipped computationally rather than explicitly computed and discarded. This appendix validates that claim: we describe three implementations that recover the same non-null eigenspace, spell out the implicit projector and shifted operator that skip the kernel with $O(Nk)$ projector cost per block application and $O(Nk)$ auxiliary working memory, and compare their wall-clock scaling. Throughout, all operators are built with the self-loop padding of Theorem 1, and on every run $r = \text{rank}(Z) = \#\text{components}(A_{\text{sym}}) = \dim\ker L_{\text{sym}}$ (numerical nullity at tolerance 10^{-6}) held exactly, so all three methods skip the same structural kernel.

B.1 Three computational routes

Does the route taken to the corrected embedding change the answer, and if not, what does it change? All three target the same k-dimensional non-null invariant subspace of L_{sym} and differ only in how they obtain it: (A) dense eigh then discard the first d_0 eigenvectors; (B) unconstrained sparse LOBPCG that requests d_0+k eigenpairs (SciPy falls back to dense eigh when $5(k+d_0) > N$); (C) the shifted-operator solve (10). The three routes differ by at most 0.007 AMI across the twelve trials, so the route is a pure cost choice, not an accuracy trade-off.

Route (C) runs LOBPCG on the shifted operator $L_\omega = L_{\text{sym}} + \omega(I - Q)$ (eq. 10, $\omega = 2.1$); each matvec is $L_{\text{sym}}x + \omega(x - Qx)$, applying $ZZ^\top = I - Q$ through the same $O(N)$ component metadata described below. The shifted solve is equivalent to the constrained solve: across the twelve trials the returned modes satisfy $\|Z^\top U\|_F \leq 5.1\times 10^{-6}$ and eigenresidual $\max_j \|L_{\text{sym}}u_j - \widehat{\lambda}_j u_j\|_2 \leq 8.4\times 10^{-6}$ (at the LOBPCG tolerance 10^{-5}), with recovered positive eigenvalues matching dense eigh to $\leq 5\times 10^{-6}$. Wall-clock is measured via time.time().

B.2 Implicit structural projector

Why can Q be applied without ever forming an $N \times r$ matrix? Because the columns of Z have disjoint support (one per component of A_{sym}), applying $Q = I - ZZ^\top$ needs no $N \times r$ matrix: with $S_c = \sum_{j\in C_c}\sigma_j$, for $i \in C_c$,

$$(QX)_{i,:} = X_{i,:} - \frac{\sqrt{\sigma_i}}{S_c}\sum_{j\in C_c}\sqrt{\sigma_j}\, X_{j,:} \tag{13}$$

(rescale each row by $1/\sqrt{\sigma_i}$, subtract the σ-weighted component mean, rescale back). Each projection of an $N \times k$ block costs $O(Nk)$; the projector is $O(N + r)$ metadata (a vertex→component index and per-vertex $\sqrt{\sigma}$ weights). Route (C) uses exactly this compact representation (component-ID map plus per-vertex weights).

B.3 Scaling results

How does each route scale with graph size? Figure 4 and Table 2 report wall-clock on DAG-chain-with-tails synthetics ($\tau = 0.3$). The bracketed "[dense fb]" rows are where the sparse driver detects $5(k + d_0) > N$ and explicitly runs dense eigh rather than LOBPCG; the sharp jump at $N = 20{,}000$ comes from one seed whose nullity $d_0 = 4143$ crossed that threshold (detailed in the table caption).

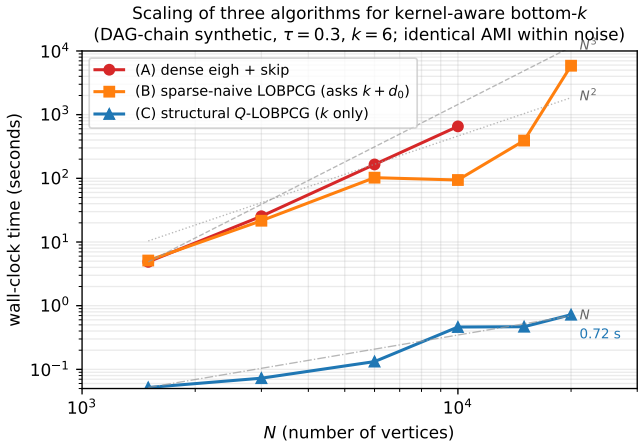

Figure 4: Wall-clock scaling of the three routes on DAG-chain-with-tails synthetics ($\tau = 0.3$, $k = 6$). The structural-Q shifted-operator solve (blue) requests only the k wanted modes and stays sub-second through $N = 20{,}000$; dense eigh (red) and unconstrained sparse LOBPCG (orange) instead grow with the nullity d_0. The plotted N, N^2, N^3 lines are reference slopes, not fitted complexities; full timings, fallbacks, and residuals are in Table 2.

		(A) dense	(B) sparse-naive	(C) struct. Q	speedup of (C) vs	
N	r	time (s)	time (s) [fallback]	time (s)	(A)	(B)
1500	514	4.88	5.12 [dense fb]	0.05	94×	99×
3000	780	25.39	21.59 [dense fb]	0.07	349×	297×
6000	890	165.12	102.41	0.13	1247×	773×
10000	1166	654.26	94.12	0.46	1415×	204×
15000	1422	skipped	391.76	0.47	n/a	838×
20000	3007	skipped	5868.74‡	0.72	n/a	8134×

Table 2: Scaling comparison of three algorithms that target the same $k = 6$ non-null invariant subspace of L_{sym} on DAG-chain-with-tails synthetics ($\tau = 0.3$); two-trial means, with r the mean structural nullity. All methods agree on AMI to within 0.007 on every row. Bracketed "[dense fb]" flags rows where the sparse driver detects $5(k + d_0) > N$ and explicitly runs dense eigh rather than LOBPCG. ‡: at $N = 20{,}000$ the two seeds differ sharply because their nullities do ($d_0 = 1871$ and 4143): the first ran genuine LOBPCG in 342 s, the second crossed the fallback threshold and took 11,395 s (a dense solve), giving the reported mean; the shifted-operator structural-Q solve took 0.50 and 0.95 s on the same two graphs.

Dense eigh was independently rerun at $N = 15{,}000$ on four fresh seeds and showed high seed-to-seed variance (7740, 7782, 15,265, 42,556 s with peak dense memory 8.4 GB measured standalone), consistent with a hardware memory-regime transition rather than a change in the algorithm's asymptotic complexity; we therefore skip dense for $N > 10{,}000$ and do not extrapolate it.

Memory. We report wall-clock time only. Because all three methods run in one process and resource.getrusage returns the process-lifetime maximum RSS, it cannot isolate a per-

method peak (the heaviest method sets a high-water mark the others inherit), so we do not tabulate empirical per-method memory. The relevant contrast is structural rather than measured: structural Q-LOBPCG uses $O(Nk)$ auxiliary working memory (an $N \times k$ block plus $O(N + r)$ projector metadata, a few MB at $N = 20{,}000$, $k = 6$) on top of the shared sparse operator, whereas dense eigh forms a full $N \times N$ double matrix (≈ 3.2 GB at $N = 20{,}000$), which is what makes it impractical at large N.

Reproducibility. All results are reproducible by scaling_structural_projector.py (the scaling sweep) and make_scaling_figure.py (Figure 4).

C Full benchmark and per-graph diagnostics

C.1 Benchmark data and evaluation protocol

This subsection fixes the datasets, evaluated cores, and clustering protocol shared by every experiment in this appendix. Unless stated otherwise, all methods run k-means with $n_{\text{init}} = 200$ and a fixed random_state, and clustering quality is scored on the labeled core (Section 3).

Datasets, sources, and preprocessing. Table 3 lists the six benchmark graphs with their sizes, evaluated core sizes, cluster counts, sources, and core construction. All are public. The further datasets used only in the prevalence audit (Appendix C.6) are public node-classification benchmarks: cora-full, dblp, and pubmed are the citation graphs of Bojchevski and Günnemann (2018), while acm (a co-authorship graph), blogcatalog (a blog-community social network), and uai (UAI2010) are standard attributed-graph clustering benchmarks.

Dataset	N	$\lvert C\rvert$	k	Source	Core preprocessing
PolBlogs	1222	713	2	Adamic and Glance (2005)	mass-pruned largest SCC
Email-Eu	986	718	42	SNAP (Leskovec and Krevl, 2014)	mass-pruned largest SCC
Cora	2708	923	7	Bojchevski and Günnemann (2018)	degree-based LWCC
Citeseer	3312	415	6	Bojchevski and Günnemann (2018)	degree-based LWCC
Cora-ML	2995	1294	7	Bojchevski and Günnemann (2018)	degree-based LWCC
Citeseer-full	4230	215	6	Bojchevski and Günnemann (2018)	degree-based LWCC

Table 3: The six benchmark datasets: analyzed graph size N (the largest weakly connected component for PolBlogs and Email-Eu; see Core construction), evaluated core size $|C|$, number of clusters k (the benchmark's own class count), source, and how the evaluated core C is extracted (mass-pruned largest strongly connected component for the two social graphs; the degree-based largest weakly connected component for the four citation graphs).

Core construction. Both preprocessing steps are unsupervised (labels are used only for evaluation, never for vertex selection). The two social graphs are first restricted to the largest weakly connected component of their raw release (PolBlogs: $1490 \to 1222$; Email-Eu: $1005 \to 986$), the analyzed size N of Table 3. The evaluated core C is then extracted deterministically. Mass-pruned largest SCC (PolBlogs, Email-Eu): with $\nu = (p_{\text{out}} + p_{\text{in}})/2$ and $\sigma = \nu + P^{\top}\nu$, keep the vertices whose σ is at or above its 20th percentile and have both $d_{\text{in}} \geq 1$ and $d_{\text{out}} \geq 1$, then take the largest strongly connected component of the induced subgraph. Degree-based LWCC (the four citation graphs): remove degree-0 vertices, then take the largest weakly connected component. The reproduction harness saves the vertex-ID list of every analyzed graph and core.

Structural vs numerical nullity (threshold audit). The paper takes $d_0 = \#\text{components}(A_{\text{sym}})$ from a union-find pass rather than by thresholding near-zero eigenvalues. The two agree: sweeping the numerical zero threshold $\varepsilon \in \{10^{-10}, 10^{-8}, 10^{-6}, 10^{-4}\}$ returns the same integer nullity as the component count on every benchmark graph, because the gap after the d_0 structural zeros is wide. On Citeseer-full (the largest structural nullspace, $d_0 = 515$) $K_0 = 0.6908$ under all four thresholds. We do not treat

$(I - ZZ^\top)L_{\text{sym}}(I - ZZ^\top)$ as an independent check: since $L_{\text{sym}}Z = 0$ it equals L_{sym} exactly, so it is the same operator, not a deflation. The K_0 cutoff itself is strict on all six headline graphs: the gap $\delta_{K_0} = \lambda_{d_0+k+1} - \lambda_{d_0+k}$ between the last selected non-null mode and the first discarded one is positive throughout, ranging from 2.4×10^{-4} (Citeseer) to 9.9×10^{-2} (PolBlogs), so the reported K_0 subspace is unique up to an internal orthogonal rotation of the selected k-plane (k0_cutoff_gap.py).

Hyperparameter provenance. Table 4 consolidates every hyperparameter of the kernel-aware pipeline (with its final value, the range searched, the selection rule, and whether it uses labels) together with the fixed, literature-prescribed defaults of the comparator baselines. The K_0 correction itself uses no supervised hyperparameter (k is the benchmark's class count, the zero threshold is invariant over six orders of magnitude, and the seeds and restart count are fixed) so only the exploratory K_2/S_3 rungs ($\beta^\star$, α) are benchmark-tuned.

Parameter	Final value	Values tried	Selection rule	Sup.?
k (embedding dim = #clusters)	Table 3	n/a	benchmark class count	n/a
k-means n_{init}	200	n/a	fixed (stability)	no
seeds	$\{42, \dots, 46\}$	n/a	fixed	no
zero threshold ε (diagnostic)	10^{-6}	$\{10^{-10}, \dots, 10^{-4}\}$	invariant across (above)	n/a
shift ω (scalable solve)	2.1	any $\omega > 2$	lifts nulls above $[0, 2]$	no
GSC measure $\beta^\star$ (K_2, S_3)	Table 5	$\{0.0, 0.1, \dots, 1.0\}$	single-seed AMI-smoothed (§C.10)	yes
band cap α (S_3)	3	$\{2, 3, 4, 5, 6, 8\}$	benchmark AMI (App. G)	yes
Comparator baselines: literature-prescribed, not tuned				
Bethe–Hessian r_{BH}	$\sqrt{\bar{d}}$	n/a	Dall'Amico et al. (2021)	no
Di-SIM τ	$\bar{d}_{\text{out}}$	n/a	Rohe et al. (2016)	no
reg-SC τ	$\bar{d}$	n/a	Qin and Rohe (2013)	no
Louvain resolution	1.0	n/a	default	no

Table 4: Hyperparameter provenance. Only the exploratory K_2/S_3 rungs use a supervised, benchmark-tuned choice ($\beta^\star$, α); the K_0 correction and its scalable solver use none.

C.2 Kernel-aware reassessment ladder

We define the four kernel-aware stages and report the reassessment ladder here; the next subsection compares them against prior baselines.

The kernel-aware ladder, defined. The stages K_0, K_1, K_2 take the first k non-null (positive) modes of a graph operator, then row-normalize and k-means; S_3 is also kernel-aware but instead takes a top-k Slepian subspace inside the band of the first m positive modes, where $m = \min\{\alpha k,\, N - r - 1\}$; S_3 is defined when $N - r \geq k + 1$, so that $k \leq m$:

$$
\begin{aligned}
K_0 &: \text{kernel-aware } L_{\text{sym}} = I - D^{-1/2} A_{\text{sym}} D^{-1/2}, \\
K_1 &: \text{kernel-aware } \bar{L}_{\nu_{1/2}} \text{ (GSC at canonical } \beta = \tfrac{1}{2}), \\
K_2 &: \text{kernel-aware } \bar{L}_{\nu_{\beta^\star}} \text{ (GSC at the selected } \beta^\star), \\
S_3 &: \text{top-}k \text{ Slepian modes of the first } \alpha k \text{ positive modes of } \bar{L}_{\nu_{\beta^\star}}.
\end{aligned}
$$

Here $\bar{L}_{\nu_\beta}$ is the GSC operator of Section 2 at measure $\nu_\beta = \beta\, p_{\text{out}} + (1 - \beta)\, p_{\text{in}} > 0$, where $p_{\text{out}}, p_{\text{in}} \in \mathbb{R}^N_{>0}$ are the (padded) out- and in-degree vectors, $p_{\text{out}}(i) = \sum_j W_{ij}$ and $p_{\text{in}}(i) = \sum_j W_{ji}$ (so $\beta = \frac{1}{2}$ is the symmetric measure and the undirected L_{sym} is the β-independent special case). For S_3: let Π_B project onto the first $m = \alpha k$ positive (non-null) modes of $\bar{L}_{\nu_{\beta^\star}}$ (when the m-th and $(m{+}1)$-th positive eigenvalues are tied, Π_B takes a fixed dimension-completing subspace of the tied eigenspace, matching the K_0 cutoff convention of Section 3), let P_C be the diagonal projector onto the evaluated core C, and let $K = \Pi_B P_C \Pi_B$ be the concentration operator; S_3 takes the top-k eigenvectors of K (the Slepian modes), expressed in the ambient space via Π_B. (The implementation caps the band at $m = \min\{\alpha k,\, N - r - 1\}$ so it never exceeds the available non-null capacity; whenever $\alpha k < N - r$ this reduces to the stated $m = \alpha k$.) The band cap is fixed at $\alpha = 3$ (Appendix G, a benchmark-selected

exploratory choice). $\beta^\star \in \{0.0, 0.1, \ldots, 1.0\}$ is a fixed, supervised selector value inherited from the single-seed pipeline (see Provenance of $\beta^\star$ in §C.10; it is not a pointwise AMI oracle), so K_2 and S_3 are supervised, benchmark-tuned values. The K_0, K_1, K_2 columns of the reassessment ladder (Table 5) and the full benchmark (Table 6) are generated from kernel_aware_ladder.py; the S_3 column is generated from slepian_ablation_ladder.py (row S3) at the tabulated $\beta^\star$, since kernel_aware_ladder.py does not itself build the Slepian concentration operator.

	$\beta^\star$	K_0	K_1	K_2	S_3	$\Delta_{\text{op}} = K_1 - K_0$	$\Delta_{\text{Slep}} = S_3 - K_2$
PolBlogs	0.5	0.839	0.807	0.807	0.832	−0.032	+0.025
Email-Eu	1.0	0.698	0.586	0.644	0.703	−0.112	+0.059
Cora	0.9	0.505	0.458	0.497	0.539	−0.048	+0.042
Citeseer	1.0	0.546	0.545	0.553	0.546	−0.001	−0.007
Cora-ML	1.0	0.466	0.423	0.405	0.573	−0.044	+0.168
Citeseer-full	0.8	0.691	0.573	0.678	0.654	−0.118	−0.024
mean		0.624	0.565	0.597	0.641	−0.059	+0.044

Table 5: Reassessment ladder (§5.3). Every stage retains kernel awareness. The per-graph $\beta^\star$ (shared by K_2 and S_3) is a fixed, supervised selector value (locally AMI-smoothed, single-seed; see Provenance of $\beta^\star$), not a pointwise AMI oracle, so the K_2 and S_3 columns are supervised, benchmark-tuned scores rather than achievable label-free ones. Bold marks the per-graph best. K_0, K_1, K_2 reproducible by kernel_aware_ladder.py; S_3 by slepian_ablation_ladder.py (row S3) at the tabulated $\beta^\star$.

Per-graph notes. Cora. $K_0 = 0.505$, $K_2 = 0.497$, $S_3 = 0.539$ at supervised $\beta^\star = 0.9$; the Slepian readout adds $\Delta_{\text{Slep}} = S_3 - K_2 = +0.042$, and S_3 exceeds the kernel-aware baseline by $S_3 - K_0 = +0.034$. Cora-ML. $K_0 = 0.466$, $K_2 = 0.405$, $S_3 = 0.573$ at supervised $\beta^\star = 1.0$; the largest Slepian readout gain in the benchmark ($\Delta_{\text{Slep}} = S_3 - K_2 = +0.168$; equivalently $S_3 - K_0 = +0.107$ over the kernel-aware baseline). Citeseer. $K_0 = 0.546$, $K_2 = 0.553$ at supervised $\beta^\star = 1.0$; K_0 and S_3 tie. Citeseer-full. $r = 515$, $\gamma = d_0/k = 85.83$: heaviest structural nullspace. $K_0 = 0.691$ is outright best; S_3 loses by -0.037. PolBlogs. Below capacity ($\gamma = 0.5$), yet naive $\to K_0$ changes $0.001 \to 0.839$: below-capacity effects can still be severe. Email-Eu. $\gamma = 0.02$; naive and K_0 tie.

C.3 Full benchmark against prior baselines

We now place the kernel-aware variants alongside the prior operators they are meant to correct. Bottom-k normalized SC (sym): bottom-k eigenvectors of $L_{\text{sym}} = I - D^{-1/2} A_{\text{sym}} D^{-1/2}$ on the symmetrized adjacency, no null skipping (the "naive" rule). Bottom-k GSC ($\beta = 1/2$): bottom-k of $\bar{L}_{\nu_{1/2}}$, no null skipping. Bethe-Hessian (sym): $H_{r_{\text{BH}}} = (r_{\text{BH}}^2 - 1)I + D - r_{\text{BH}} A$ at $r_{\text{BH}} = \sqrt{\bar{d}}$ following Dall'Amico et al. (2021). Di-SIM: SVD of the regularized adjacency at $\tau = \overline{d_{\text{out}}}$, with sender (left) and receiver (right) clusterings reported separately. Directed Louvain: modularity maximization at resolution 1.0. All methods use $n_{\text{init}} = 200$ and identical random_state across baselines.

C.4 Strictly protocol-matched comparison

Table 6 mixes evaluation protocols across row groups, a transparent but imperfect comparison. We therefore add a strictly protocol-matched table: every method produces a k-dimensional node embedding of the same graph, and we apply one identical readout to all of them, restrict to the same evaluated core C, row-normalize, and cluster with five-seed k-means ($n_{\text{init}} = 200$, seeds 42–46). Only the embedding differs across columns. We add regularized spectral clustering (Qin and Rohe, 2013) (adjacency regularization $A_\tau = A + (\tau/N)\,\mathbf{1}\mathbf{1}^\top$, $\tau = \bar{d}$), which removes disconnection by construction, alongside Bethe–Hessian, Di-SIM (sender/receiver), and the naive/GSC bottom-k readouts. The selected spectral indices are: naive SC and reg-SC take the k eigenvectors of smallest eigenvalue (including any trivial constant or null mode); Bethe–Hessian takes the k most negative

Method	PolBlogs	Email-Eu	Cora	Citeseer	Cora-ML	Citeseer-full
Bottom-k SC sym (naive)	0.001	0.704	0.000	0.000	0.000	0.000
Bottom-k GSC ($\beta = 1/2$, naive)	0.001	0.561	0.000	0.000	0.000	0.000
Bethe-Hessian sym	0.832	0.700	0.284	0.286	0.129	0.351
Di-SIM sender	0.738	0.658	0.366	0.266	0.515	0.516
Di-SIM receiver	0.590	0.700	0.427	0.268	0.437	0.462
Directed Louvain	0.817	0.629	0.373	0.371	0.437	0.476
K_0 (kernel-aware sym; ours)	0.839	0.698	0.505	0.546	0.466	0.691
K_2 (kernel-aware GSC, supervised $\beta^\star$; ours)	0.807	0.644	0.497	0.553	0.405	0.678
S_3 (Slepian on GSC, supervised $\beta^\star$; ours)	0.832	0.703	0.539	0.546	0.573	0.654

Table 6: Full six-graph benchmark against four prior baselines plus our three variants. Protocols differ by row group and are not mixed within a group: the two naive bottom-k rows and our $K_0/K_2/S_3$ rows are five-seed KMeans means at $n_{\text{init}} = 200$ (std ≤ 0.010); the four prior baselines are a single run at random_state $= 42$, $n_{\text{init}} = 200$. Unlike Table 1 (core-restricted), the two naive rows here use the full-embedding protocol (fit k-means on all N vertices, evaluate on the core), so they empirically report exactly 0.000 on the four above-capacity graphs; Proposition 2 establishes that their evaluated core embeddings contain no within-core graph information, but does not by itself determine the resulting AMI. The corresponding core-restricted SC/GSC embeddings instead score small, unstable tie-breaking values (0.005–0.46, numerically and tie-breaking dependent under the degenerate readout) (§3, Protocol note). Our $K_0/K_2/S_3$ rows use core-restricted fitting (matching Table 5); K_2 and S_3 use the supervised $\beta^\star$ and are supervised, benchmark-tuned scores. Bold: per-graph best. Reproducible by run_all_baselines_n200.py (prior baselines), full_embed_naive.py (the two full-embedding naive rows), kernel_aware_ladder.py (K_0, K_2), and slepian_ablation_ladder.py (row S3, for S_3).

eigenvectors of $H_{r_{\text{BH}}}$; Di-SIM takes the top-k left (sender) and right (receiver) singular vectors of the regularized adjacency; and K_0 takes the first k non-null eigenvectors of L_{sym}.

	SC	reg-SC	BH	Di-SIM$_s$	Di-SIM$_r$	GSC	K_0
PolBlogs	0.001	0.776	0.818	0.812	0.832	0.001	0.839
Email-Eu	0.701	0.722	0.729	0.667	0.699	0.569	0.698
Cora	0.153	0.471	0.461	0.418	0.443	0.163	0.505
Citeseer	0.442	0.357	0.359	0.181	0.215	0.353	0.546
Cora-ML	0.056	0.513	0.491	0.489	0.451	0.022	0.466
Citeseer-full	0.020	0.577	0.488	0.499	0.525	0.311	0.691
mean	0.229	0.569	0.558	0.511	0.527	0.237	0.624

Table 7: Strictly protocol-matched comparison: same graph, same evaluated core, same row-normalized five-seed k-means ($n_{\text{init}} = 200$) for every method. Only the embedding differs. K_0 has the highest mean AMI and is best on four of six graphs, ahead of regularized SC (Qin and Rohe, 2013), Bethe–Hessian, and Di-SIM. The naive SC/GSC columns are degenerate rank-≤ 1 readouts on the above-capacity graphs (Proposition 2), so their values are numerically and tie-breaking dependent under the degenerate readout and need not match the core-restricted naive column of Table 1. On a single-component core Proposition 2 forces every admissible null k-plane to constant normalized rows (an orthogonal change $U \mapsto UR$ leaves row norms, pairwise distances, and the k-means objective invariant), so the residual variation reflects finite-precision, eigensolver, and k-means tie-breaking rather than a mathematical dependence on the choice of basis. The other columns use non-degenerate embeddings and are stable. Reproducible by protocol_matched.py.

Under this single controlled protocol K_0 has the highest mean AMI (0.624) and is best on four of the six graphs, ahead of regularized SC (0.569), Bethe–Hessian (0.558), and Di-SIM (0.51–0.53). Regularized SC (the natural competing fix for the multi-component pathology) is the runner-up and edges K_0 on Cora-ML (0.513 vs 0.466); Bethe–Hessian edges it on the low-occupancy Email-Eu, where the single structural null direction ($d_0 = 1$) contributes little to the naive embedding. The same-operator indexing correction is thus competitive with,

and on average better than, substantially more elaborate operators when all are compared fairly.

C.5 Embedding-dimension ablation

K_0 keeps the embedding dimension at k and replaces every structural zero direction by a positive one, so one may ask whether its advantage is escaping the nullspace or merely using k positive modes (one more informative direction than the $k-1$ that a connected-graph NJW readout effectively carries). We disentangle these by clustering, under the same core-restricted protocol, four readouts of the same L_{sym}: literal bottom-k $[v_1, \ldots, v_k]$; the $(k-1)$ non-null readout $[v_{d_0+1}, \ldots, v_{d_0+k-1}]$; $K_0 = [v_{d_0+1}, \ldots, v_{d_0+k}]$; and the $(k+1)$ non-null readout.

	d_0	naive	$(k-1)$-nn	K_0	$(k+1)$-nn
PolBlogs	1	0.001	−0.000	0.839	0.839
Email-Eu	1	0.701	0.695	0.698	0.704
Cora	78	0.153	0.509	0.505	0.501
Citeseer	438	0.344	0.510	0.546	0.539
Cora-ML	61	0.005	0.458	0.466	0.445
Citeseer-full	515	0.118	0.720	0.691	0.691

Table 8: Embedding-dimension ablation (core-restricted, five-seed). On all four above-capacity graphs the $(k-1)$ non-null readout already recovers essentially the full K_0 correction, and the result is stable across $(k-1)$, k, and $(k+1)$ non-null modes, all far above literal bottom-k. The gain is therefore escaping the structural nullspace, not the extra positive dimension. PolBlogs (below capacity) is the exception: it needs the full k-dimensional positive embedding (the single mode $[v_{d_0+1}]$ alone gives 0, and literal $[v_1, v_2]$ also fails because the null mode dominates). Reproducible by emb_dim_audit.py.

Because $(k-1)$, k, and $(k+1)$ non-null modes all recover the correction on the above-capacity graphs, K_0 is a dimension-matched, kernel-free readout rather than a device that smuggles in an extra positive direction.

C.6 Prevalence: a broad labeled-graph audit

To show when the pathology occurs, we audit twelve real labeled graphs (the six of the main text plus six further citation/social networks), reporting the capacity ratio d_0/k, median occupancy, and the naive-to-K_0 gap under the core-restricted protocol. On every graph we verified that the evaluated core lies in a single connected component, so Proposition 2 applies wherever $d_0 \geq k$.

Across the audited graphs, $d_0/k \geq 1$ together with the single-component-core condition of Proposition 2 identifies the saturated regime in which the evaluated embedding is structurally degenerate. Below capacity, d_0/k alone is insufficient: null occupancy distinguishes severe cases such as PolBlogs ($d_0/k = 0.5$, $\eta \approx 1$, $0.001 \to 0.839$) from weakly affected ones such as Email-Eu ($d_0/k = 0.02$, $\eta = 0.035$). Three of the above-capacity graphs (acm, cora-full, dblp) lie outside the original four citation benchmarks.

C.7 Robustness to metric and fitting protocol

The structural diagnosis should not depend on the external clustering metric, while the fitted clustering score may depend on where k-means is fit; we audit both sensitivities here, and report the null-occupancy distribution that underlies the below-capacity bound.

We first ask whether the correction depends on the external clustering metric. Table 10 reuses the core-restricted protocol of Table 1 and holds each predicted partition fixed while changing only the external validity metric: AMI, Normalized Mutual Information (NMI), and Adjusted Rand Index (ARI); only the validity index changes across the three columns of a partition. The naive-to-K_0 movement is qualitatively identical under all three metrics

dataset	N	k	d_0	d_0/k	med η	naive	K_0 (Δ)
acm	3025	3	717	239.0	1.000	0.142	0.375 (+0.233)
cora	2708	7	78	11.1	1.000	0.153	0.505 (+0.352)
cora-ml	2995	7	61	8.7	1.000	0.005	0.466 (+0.461)
citeseer	3312	6	438	73.0	1.000	0.344	0.546 (+0.202)
citeseer-full	4230	6	515	85.8	1.000	0.118	0.691 (+0.573)
cora-full	19793	70	364	5.2	1.000	0.000	0.532 (+0.532)
dblp	17716	4	589	147.2	1.000	0.011	0.287 (+0.276)
PolBlogs	1222	2	1	0.50	1.000	0.001	0.839 (+0.838)
Email-Eu	986	42	1	0.02	0.035	0.701	0.698 (−0.003)
blogcatalog	5196	6	1	0.17	0.478	0.179	0.230 (+0.051)
uai	3067	19	8	0.42	0.734	0.263	0.296 (+0.034)
pubmed	19717	3	1	0.33	0.688	0.157	0.174 (+0.017)

Table 9: Broad audit over twelve labeled graphs (above the rule: $d_0/k \geq 1$; below: $d_0/k < 1$). Above capacity (median $\eta = 1$, single-component core) the naive readout is a degenerate rank-≤ 1 collapse by Proposition 2 (verified directly, e.g. its core rank is 1 on cora-full and dblp, so the naive entries there are tie-breaking ≈ 0) and K_0 recovers a large gap, on a new dataset (acm, $d_0/k = 239$) as well as the citation benchmarks. Below capacity the gap covaries with occupancy (consistent with, though not implied by, Proposition 3): decisive on PolBlogs ($\eta \approx 1$), negligible on Email-Eu ($\eta = 0.035$). For the three graphs with $N \approx 20{,}000$, occupancy and K_0 use the scalable structural-Q solver of Section 4 (broad_audit_large.py) and the naive collapse is confirmed from the exact component-indicator null basis; the rest use dense eigendecomposition (broad_audit.py).

on every graph: no metric flips the ordering or changes the sign of Δ, and the six-graph mean of Δ is +0.404 (AMI), +0.399 (NMI), +0.405 (ARI), matching the core-restricted K_0 column of Table 1; under the alternative full-embedding protocol (Table 6) the four citation-graph naive scores empirically collapse to ≈ 0 under all three metrics, consistently with Proposition 2, which establishes graph-independence of the embedding rather than the value of any particular partition-quality index. All three metrics identify the same qualitative regimes (large positive Δ on the four above-capacity graphs, near-zero on Email-Eu, decisive on PolBlogs).

We next test whether K_0 itself is sensitive to where k-means is fit. The naive collapse is protocol-independent (the core embedding is rank-≤ 1 under either fit), but the K_0 scores are not. Refitting k-means on the full row-normalized embedding of all N vertices (still evaluating on the core), rather than on the core alone, moves the six-graph mean K_0 from 0.624 (core-restricted, used in Table 1) to 0.605, with the largest per-graph shifts on Cora ($0.505 \to 0.459$) and Citeseer ($0.546 \to 0.438$) and the rest moving by ≤ 0.03. We report the core-restricted numbers throughout because full-embedding K_0 is itself unstable on these small labeled cores (independent reruns of Citeseer differed by ~ 0.05, as the full-embedding k-means fits mostly unlabeled vertices), whereas core-restricted K_0 is stable and matches the reassessment ladder (Table 5). Reproducible by m1_verify.py.

Finally, we inspect the full null-occupancy distribution underlying Proposition 3. Its ball bound is stated in terms of $\min_i \eta_i$, whereas Table 1 reports the median for compactness. Table 11 gives the full distribution on the evaluated core. The four above-capacity citation graphs have $\eta = 1$ exactly at every quantile (Proposition 2). PolBlogs is the only graph where the median-vs-minimum distinction matters: its 5th percentile is 0.999 (so at least 95% of the core lies in a ball of radius $2\sqrt{1 - 0.999} \approx 0.06$), but its minimum is 0.054, so a small tail of vertices retains non-null geometry and the all-rows guarantee does not cover the entire core. Email-Eu is uniformly low-occupancy.

		AMI		NMI		ARI	
	d_0	naive	K_0	naive	K_0	naive	K_0
PolBlogs	1	0.001	**0.839**	0.003	**0.839**	0.001	**0.907**
Email-Eu	1	0.701	0.698	0.775	0.773	0.496	0.493
Cora	78	0.153	**0.505**	0.163	**0.511**	0.071	**0.441**
Citeseer	438	0.344	**0.546**	0.359	**0.556**	0.355	**0.564**
Cora-ML	61	0.005	**0.466**	0.013	**0.471**	−0.003	**0.326**
Citeseer-full	515	0.118	**0.691**	0.140	**0.699**	0.072	**0.693**

Table 10: Metric robustness of the naive-to-K_0 correction. For each graph and each method the same five-seed clustering (as in Table 1, core-restricted fitting) is scored by AMI, NMI, and ARI and averaged over the five seeds: only the external validity index changes, so this isolates metric choice from the fitting protocol. (The naive columns therefore report the core-restricted tie-breaking values, not the full-embedding 0.000 of Table 6; the qualitative naive-to-K_0 movement is the same under all three metrics.) Bold: outright per-metric best. Reproducible by metric_robustness.py.

Graph	d_0/k	min η	p5 η	median η	mean η
PolBlogs	0.50	0.054	0.999	1.000	0.998
Email-Eu	0.02	0.001	0.005	0.035	0.056
Cora	11.14	1.000	1.000	1.000	1.000
Citeseer	73.00	1.000	1.000	1.000	1.000
Cora-ML	8.71	1.000	1.000	1.000	1.000
Citeseer-full	85.83	1.000	1.000	1.000	1.000

Table 11: Null-occupancy η distribution on the evaluated core. The d_0/k, min, median, and mean columns are from frozen_headline.py; the 5th-percentile column is computed from the per-vertex η saved by occupancy_diagnostic.py (occupancy_per_vertex.json). On the four above-capacity graphs every core vertex is fully null-occupied, so all η statistics equal 1.000.

C.8 External real-network controls

Beyond the six main benchmarks, we test the mechanism on further real graphs: three classical labeled networks, then a six-graph SNAP audit that reaches the multi-component regime the labeled benchmarks cannot.

Classical small-network pilot. As a first step toward broader real-network breadth, we run a three-protocol comparison on three canonical small networks from the prior literature: Karate (Zachary, 1977), PolBooks (Krebs, 2004), and Football (Girvan and Newman, 2002). Dolphins (Lusseau et al., 2003) is omitted because Newman's distributed GML does not ship with community labels. For each network we run

- A: full graph, naive bottom-k;
- B: full graph, K_0 (kernel-aware);
- C: largest connected component, naive bottom-k.

Table 12 reports d_0, median η, and AMI/NMI/ARI (5-seed KMeans at $n_{\text{init}} = 200$) for each protocol on the labeled core.

Every graph in this pilot is connected ($d_0 = 1$). Two consequences. (i) The multi-component pathology of Proposition 2 does not apply here, and the LCC preprocessing used by Dall'Amico et al. (2019; 2021) is a no-op (protocol A ≡ protocol C on every metric). (ii) The below-capacity mechanism of Proposition 3 can nevertheless still operate: $d_0 = 1 < k$ leaves the naive readout $[v_1, \ldots, v_k]$ with one null-supported direction and $k - 1$ non-null directions, and the geometry of the resulting core embedding is governed by η. Karate has the largest occupancy among the three pilot graphs ($\eta = 0.544$) and also shows the largest K_0 gain (+0.107 AMI, +0.105 NMI, +0.110 ARI); PolBooks is intermediate

				AMI		NMI		ARI		
	N	k	med η	A	B	A	B	A	B	ΔAMI
Karate	34	2	0.544	0.726	**0.833**	0.732	**0.837**	0.772	**0.882**	+0.107
PolBooks	105	3	0.429	0.533	**0.558**	0.542	**0.567**	**0.645**	0.588	+0.025
Football	115	12	0.088	**0.899**	0.867	**0.924**	0.900	**0.897**	0.824	−0.032

Table 12: Classical small-network pilot: A = full graph, naive bottom-k; B = full graph, K_0. All three graphs are connected ($d_0 = 1$), so the LCC-only protocol C is exactly equal to A on every metric and is omitted from the table. Reproducible by pilot_four_networks.py.

($\eta = 0.429$, +0.025 AMI); Football has very low occupancy ($\eta = 0.088$) and the correction slightly decreases AMI (−0.032). At very low occupancy the structural null direction contributes little to the embedding, and replacing it need not improve the partition. This ordering is consistent with the below-capacity geometry of Proposition 3, but three networks are insufficient to establish a quantitative predictor. Together with Email-Eu and PolBlogs, this pilot expands the real below-capacity evidence from two graphs to five, complemented by seven controlled synthetic realizations at $\eta \in [0.21, 0.69]$. Football also serves as a useful control: K_0 is not a generic performance-enhancing regularizer.

SNAP audit: kernel-aware Fiedler vs. LCC preprocessing on unlabeled graphs. The classical small-network pilot leaves the LCC-preprocessing comparison unresolved because all three analyzed graphs are singly connected. To reach the multi-component regime we run a targeted audit on six SNAP graphs (Leskovec and Krevl, 2014) chosen to overlap the Zhang and Rohe (2018) test set: p2p-Gnutella08, wiki-Vote, ca-HepTh, ca-HepPh, ca-AstroPh, and email-Enron. We measure graph conductance (ϕ) via the standard sweep cut on v_2, evaluated on the raw symmetric hull (matching Zhang-Rohe's protocol). Three protocols per graph: A = full graph, naive bottom-2, when $d_0 > 1$ a literal numerical bottom-2 is an arbitrary, solver-dependent rotation within the d_0-dimensional kernel, so we freeze it to a deterministic worst-case null vector, the smallest-component indicator (an explicitly adversarial, constructed element of $\ker L_{\text{sym}}$); B = full graph, kernel-aware Fiedler (first non-null mode of L_{sym}, obtained by Z-constrained LOBPCG); C = largest connected component of the padded graph, naive bottom-2. Protocol B is the $K = 2$ analogue of the K_0 indexing principle of Section 3 under Zhang-Rohe's sweep-cut protocol; it uses one non-null eigenvector rather than k because the sweep cut is defined on a single vector.

C.9 Relation to LCC preprocessing

When does kernel awareness coincide with simply restricting to the largest connected component, and when does it differ? The SNAP audit above answers this, in three observations.

(i) B and C agree here under a component-spectrum condition. Because A_{sym} is padded with self-loops on isolated vertices, L_{sym} is block-diagonal over connected components of the padded graph, and its positive spectrum is the disjoint union of the positive spectra of all non-trivial components. In general, the smallest positive eigenvalue of the full L_{sym} comes from $c^\star = \arg\min_{c:\,|C_c|>1} \mu_c$, where $\mu_c := \min(\text{spec}(L_c) \setminus \{0\})$ is the smallest positive eigenvalue of a non-singleton component c (a padded singleton has $\text{spec}(L_c) = \{0\}$ and is excluded), not necessarily the LCC. Three cases arise. If $\mu_{\text{LCC}} < \min_{c \neq \text{LCC}} \mu_c$, the first positive eigenspace of the full L_{sym} is supported on the LCC, and if moreover μ_{LCC} is simple the kernel-aware Fiedler of protocol B equals the zero-padded LCC Fiedler up to sign. If $\mu_{\text{LCC}} > \min_{c \neq \text{LCC}} \mu_c$, the first positive mode is supported on another component. Equality is a genuine degenerate third case: the global first-positive eigenspace then spans several components and the returned vector depends on the eigensolver's choice within it. We check the favorable case directly rather than assume it: on all six audited SNAP graphs (i) v_B carries ≥ 0.99999 of its ℓ_2 mass on the LCC, and (ii) the LCC-local gap $\lambda_3(L_{\text{LCC}}) - \lambda_2(L_{\text{LCC}}) > 0$ (1.0×10^{-3} to 6.8×10^{-2}) shows that μ_{LCC} is simple and, with (i), gives strong numerical evidence (short of a proof, as we do not bound every other component's μ_c) that it is the global smallest positive eigenvalue. We then recompute the

	N	d_0	$\|LCC\|/N$	ϕ_A (full-naive)	$\phi_B = \phi_C$	small-side
p2p-Gnutella08	6301	2	0.9997	0.0000	0.0476	11
wiki-Vote	7115	24	0.9931	0.0000	0.1111	5
ca-HepTh	9877	429	0.8746	–	0.0036	24
ca-HepPh	12008	278	0.9330	–	0.0024	21
ca-AstroPh	18772	290	0.9537	–	0.0101	23
email-Enron	36692	1065	0.9183	0.0000	0.0045	28

Table 13: Six-graph SNAP audit. Protocol A is a deterministic, adversarially-chosen null vector: the smallest-component indicator, a constructed element of $\ker L_{\text{sym}}$, not a solver output. Its sweep is degenerate: on three graphs it collapses to a component boundary ($\phi_A = 0$, isolating a tiny component, a trivial split), and on the other three (dashes) the null-supported sweep yields only a meaningless near-balanced cut ($\phi_A \approx 0.50$); neither carries within-LCC community structure (the $\phi_A \approx 0.50$ dashes are an artifact of the sweep's tie-handling: the chosen smallest component is a raw-isolated singleton with zero conductance volume, so the sweep skips it and default-orders the large tied zero-block instead). Protocols B (kernel-aware Fiedler) and C (LCC-naive Fiedler, recomputed independently on the LCC) produce the identical cut on every graph, so we report them as a single column $\phi_B = \phi_C$ with a single small-side count: the two agree not merely in conductance and cardinality but in the smaller-side vertex sets themselves (Jaccard 1.0000 on all six, snap_verify.py). This holds because B's Fiedler places ≥ 0.99999 of its ℓ_2 mass on the LCC and the LCC-local gap $\lambda_3(L_{\text{LCC}}) - \lambda_2(L_{\text{LCC}}) > 0$ (1.0×10^{-3} to 6.8×10^{-2}) certifies a simple smallest positive eigenvalue on the LCC (observation (i)). Reproducible by snap_audit.py (protocols A/B/C) and snap_verify.py (Fiedler support, component gap, and cut-membership comparison).

LCC Fiedler independently (protocol C) and compare the actual smaller-side vertex sets of the two sweep cuts: they are identical (Jaccard 1.0000; snap_verify.py saves both), so B and C return the same partition, not merely equal conductance and cut size. This is an empirical property of this family, verified per graph, not a mathematical identity in general.

(ii) Kernel awareness and LCC preprocessing can coincide but are not the same operation. On this SNAP family, kernel awareness delivers the same partition-quality as the widely-used "take the largest connected component" recipe of Dall'Amico et al. (2021) and Zhang and Rohe (2018), and reaches that partition without an explicit preprocessing step. The two procedures are different in kind, however: LCC preprocessing commits a priori to a single component and deletes every other vertex, while kernel-aware indexing removes only the zero-energy ambiguity and lets the positive spectrum decide which component(s) contribute the lowest-energy modes. Under the condition of (i) they select the same LCC mode; when the smallest positive eigenvalue lives elsewhere, they do not.

(iii) The SNAP audit does not test the label-preservation value of K_0. Conductance is unsupervised; the paper's citation-graph benchmarks (Table 1) already showed the labeled case where K_0 recovers 0.47–0.69 AMI on the labeled core after evaluating on the full graph. On those graphs the labeled core is entirely inside the LCC (verified directly: $|C \setminus \text{LCC}| = 0$ on all six benchmarks), so LCC preprocessing does not delete any labeled vertex. K_0 also differs from re-computing the operator on the labeled core $G[C]$ alone: the Core-first comparison of Appendix H rebuilds L_{sym} on $G[C]$ and thereby deletes unlabeled but coupled residual vertices whose edges enter the spectral geometry of the core; the observed gap (up to 0.120 AMI on Citeseer-full, Table 20) is consistent with that loss, though it does not by itself isolate the coupling channel. Appendix C.6 provides an initial extension beyond the six headline benchmarks; a substantially broader survey across graph domains remains useful for characterizing the prevalence of the pathology.

C.10 Provenance of the exploratory $\beta^\star$ (supervised selector, not an oracle)

The K_2 and S_3 rungs depend on a per-graph $\beta^\star$; this subsection documents exactly how it was chosen and why it is supervised rather than a pointwise oracle. The per-graph $\beta^\star = (0.5, 1.0, 0.9, 1.0, 1.0, 0.8)$ in Table 5 are not pointwise AMI oracles. They are the fixed, supervised, locally-AMI-smoothed selector values inherited from the earlier single-

seed pipeline (figure1_reformulated.py, saved in figure1_reformulated_data.json as selector_beta): for each graph, on one k-means seed, that selector maximizes the mean AMI over the $\beta \pm 0.1$ window rather than the AMI at β itself. It is therefore supervised (it uses the labels) but it does not maximize the AMI: on the same single-seed sweep the pointwise maximizer (oracle_beta $= (0.4, 1.0, 0.8, 0.9, 1.0, 0.7)$) differs on four of six graphs; e.g. on Cora the AMI peaks at $\beta = 0.8$ (0.545) while the selected $\beta^\star = 0.9$ gives 0.539, and on Citeseer-full it peaks at $\beta = 0.7$ (0.672) while $\beta^\star = 0.8$ gives 0.660. The same $\beta^\star$ is reused for both K_2 and S_3, although these are different maps of β whose maximizers need not coincide. Consequently K_2 (GSC at $\beta^\star$) and S_3 (Slepian readout at $\beta^\star$) are supervised, benchmark-tuned values, and we report them only to estimate the residual headroom above K_0, not as certified oracle upper bounds and not as attainable without labels. A genuinely label-free, per-method selector is left to future work; one candidate is $\beta^\star = \arg\max_\beta \mathrm{Stab}_\Delta(\beta)$, with $\mathrm{Stab}_\Delta(\beta)$ the canonical correlation between the ambient S_3 output subspaces $U(\beta)$ and $U(\beta \pm \Delta)$ on the post-nullspace band, which Proposition 4 in Appendix I shows is locally Lipschitz in β under band-separation and Slepian-eigengap assumptions. That guarantee motivates but does not by itself deliver such a selector.

D Undirected controls: details

Section 5 gives summary undirected results. This appendix contains the full DC-SBM phase diagram across three operators, the fine-grained activation curve, and the Coauthor CS diagnostic in full.

D.1 Controlled synthetic: the pathology across the collapse regime

Goal. Test whether the structural-nullspace mechanism persists across Laplacian variants and locate the operating regime in which it fires: the six-graph correction exposes the naive-vs-kernel-aware gap on real graphs but not the regime, and a controlled synthetic sweep delineates it.

Design. We vary the tail-vertex fraction τ on the DAG-chain generator of Appendix F, which changes the empirical nullity d_0 while holding $N = 1{,}500$ and $k = 6$ fixed, and record AMI of naive and kernel-aware bottom-k on three operators: L_{sym} (symmetric normalized), L_{comb} (combinatorial), and $\bar{L}_{\nu_{1/2}}$ (directed GSC at the canonical measure).

Common regime coordinate. All three operators share the same off-diagonal hull and component partition of A_{sym} and differ only in their diagonal self-loop padding (L_{sym} and L_{comb} pad vertices isolated in the symmetric hull, while the GSC construction pads every zero-in- or zero-out-degree vertex), but diagonal loops do not change connected components, so this padding difference leaves the support partition intact and the common d_0/k axis is the structural nullity ratio of each operator. On every realization we verified $d_0 = \#\mathrm{components} = \dim\ker L_{\mathrm{sym}}$ numerically.

The sweep sits entirely above capacity. With the self-loop padding convention a sparse DAG-chain graph at this density fragments into many singleton and small components even before tail injection, so the structural nullity ratio is $d_0/k \in [30.7, 137.7]$ across the sweep and never drops below the capacity boundary $d_0 = k$. Realistic sparse construction at $N = 1{,}500$ therefore lands deep in the above-capacity regime; the capacity transition itself is isolated instead by the controlled activation sweep of §5.1, which injects a prescribed number of isolates into a clean core. Because τ also shrinks the labeled core as it enlarges d_0, we read d_0/k here as a regime coordinate, not a monotone difficulty knob.

Findings. Throughout this regime, naive L_{sym} bottom-k averages $\overline{\mathrm{AMI}} = 0.30$ while kernel-aware L_{sym} averages 0.68 (mean gap +0.38). Naive L_{comb} averages 0.52 vs. kernel-aware L_{comb} at 0.66 (mean gap +0.14); naive $\bar{L}_{\nu_{1/2}}$ averages 0.42 vs. 0.66 (mean gap +0.24). Kernel awareness closes most of the naive-vs-kernel-aware gap on all three operators, and the three kernel-aware curves converge to $\overline{\mathrm{AMI}} \approx 0.66$–$0.68$ regardless of operator.

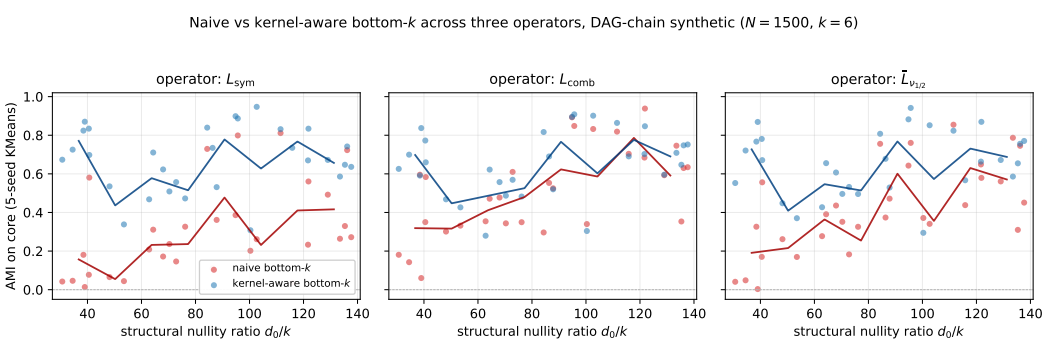

Figure 5: AMI of naive vs. kernel-aware bottom-k against the structural nullity ratio d_0/k, on three operators. DAG-chain-with-tails synthetic, $N = 1{,}500$, $k = 6$, $\tau \in \{0, 0.05, \ldots, 0.45\}$ (3 trials each, 30 realizations). Under the self-loop padding convention every realization sits well above capacity ($d_0/k \in [31, 138]$), so the horizontal axis spans the heavy-nullity regime rather than a transition. Kernel-aware curves (blue) stay in the 0.6–0.8 AMI band for all three operators. Naive L_{sym} (left) is worst, averaging 0.30; naive L_{comb} (center) and naive $\bar{L}_{\nu_{1/2}}$ (right) score higher but remain below their kernel-aware counterparts. The operator-dependent differences among the naive curves are numerical tie-breaking artifacts on a degenerate embedding, not an operator advantage (see text). Reproducible by synthetic_phase_diagram.py and make_phase_diagram_figure.py.

Two mechanistic observations. First, the naive readouts differ across operators, but not for a reason that reflects an operator advantage. When $d_0 \geq k$ and the evaluated core lies in one support component, Proposition 2 shows that every exact null-supported bottom-k readout is constant on the core after row normalization, independent of which k-plane is drawn from the d_0-dimensional kernel; the positive per-row scale $\sqrt{\sigma_i}$ cancels ($U_i/\|U_i\| = R_{c,:}/\|R_{c,:}\|$). The nonzero and operator-dependent naive scores in this regime are therefore not evidence of a meaningful "different null basis": they arise from finite-precision leakage of near-zero positive modes into the numerically computed kernel and from k-means tie-breaking on near-degenerate rows, both implementation- and operator-specific (the same numerical/tie-breaking mechanism noted for the real benchmarks in §3, Protocol note). Second, the kernel-aware curves converge across operators: switching from L_{sym} to L_{comb} or $\bar{L}_{\nu_{1/2}}$ under kernel-aware indexing moves the mean AMI by at most ≈ 0.02 on this synthetic, consistent with the six-graph reassessment ladder of Section 5.3 where $K_1 - K_0$ is uniformly negative but small in magnitude. Under kernel-aware indexing, the operator choice matters less than kernel awareness itself.

Bridge to the real benchmark and the capacity boundary. Using the authoritative (d_0, r) values of Table 1, the four above-capacity citation graphs have $d_0/k \in \{8.71, 11.14, 73.00, 85.83\}$ (Cora-ML, Cora, Citeseer, Citeseer-full). The synthetic sweep's $d_0/k \in [31, 138]$ overlaps the heaviest of these (Citeseer, Citeseer-full) and extends beyond, corroborating the pathology across the heavy-nullity regime; it does not reach the lighter citation graphs or the capacity boundary, which are covered instead by the activation sweep and by the two below-capacity real graphs. PolBlogs ($d_0/k = 0.5$) and Email-Eu ($d_0/k \approx 0.024$) both sit below the structural capacity boundary $d_0 = k$ of Section D.2 and are discussed per-graph in Section 5. The four above-capacity graphs each satisfy Proposition 2's core-in-one-component hypothesis, so their naive readouts are predicted to give rank-≤ 1 core embeddings; Proposition 2 fixes this core geometry exactly, while the resulting naive AMI is implementation- and tie-breaking-dependent (the core-restricted values of Table 1). The controlled experiment therefore locates the correction in a mathematically-characterized operating regime, not a benchmark accident.

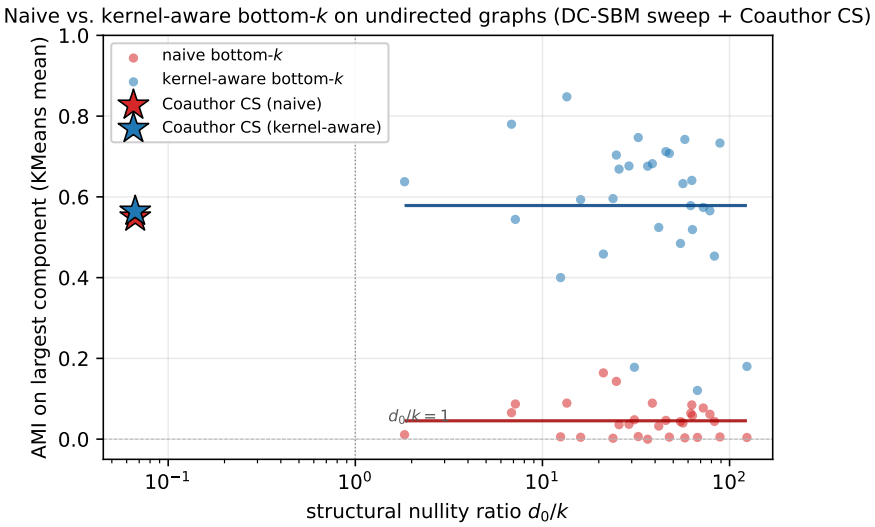

Figure 6: Undirected phase diagram: AMI vs. structural nullity ratio d_0/k on a sparse undirected DC-SBM (30 realizations, $N = 1{,}500$, $k = 6$, $\sigma_\theta = 0.7$, $\sigma_{\text{block}} = 1.5$). Kernel-aware bottom-$k$ (blue) recovers AMI in $[0.12, 0.85]$ (mean 0.579) across the sweep; naive bottom-k (red) collapses to AMI ≤ 0.17 throughout. Mean gap $+0.53$ over the 30 trials. Stars mark Coauthor CS ($N = 18{,}333$, $k = 15$), an external undirected control where $d_0/k = 0.07$: naive and kernel-aware essentially coincide at ~ 0.55, a below-capacity negative control (though $d_0/k < 1$ alone does not predict inertness; cf. PolBlogs). The x-axis is logarithmic to span both regimes; horizontal lines mark the sweep means (naive 0.045, kernel-aware 0.579). Reproducible by undirected_dcsbm_phase.py, coauthor_cs_test.py, and make_undirected_figure.py.

D.2 Undirected control: the mechanism is not directional

The mechanism of Section 3 is stated in the GSC framework, which encompasses undirected graphs via $\nu \propto d$ on the symmetric normalized Laplacian. We test this scope with two experiments.

Undirected DC-SBM phase diagram. We generate a sparse undirected degree-corrected SBM at $N = 1{,}500$, $k = 6$, with log-normal degree heterogeneity ($\sigma_\theta = 0.7$), log-normal block-pair density heterogeneity ($\sigma_{\text{block}} = 1.5$), and a tail-vertex fraction $\tau \in \{0, 0.05, \ldots, 0.45\}$ of vertices with expected degree ~ 1 that become isolated after self-loop padding. For each τ we run 3 trials and report AMI on the largest connected component of the padded A_{sym}. Figure 6 plots naive and kernel-aware bottom-k against the observed d_0/k.

Findings. Across the 30 undirected DC-SBM trials, mean naive AMI is 0.045 (range $[0.00, 0.16]$) and mean kernel-aware AMI is 0.579 (range $[0.12, 0.85]$), giving mean gap $+0.53$ AMI, positive on every trial. On every realization $\text{rank}(Z)$ matches the numerical d_0 exactly, as it does on the citation benchmarks (§2), confirming that Theorem 1's topological characterization of the nullspace is a saturated identity on this generator. The phase diagram matches the directed case (Figure 5) qualitatively at every d_0/k: kernel-aware curves stay high, naive curves collapse.

Coauthor CS: an undirected real-graph external control. Coauthor CS (Shchur et al., 2018) is an undirected coauthorship network with $N = 18{,}333$ authors, 81,894 edges, and 15 field-of-study classes (chosen because coauthorship is a naturally symmetric relation, unlike citation). It has $\text{rank}(Z) = 1$ and $d_0 = 1$ (single connected component of A_{sym}), i.e. $d_0/k = 0.07$, well below the capacity boundary. Naive bottom-k SC and kernel-aware

bottom-k SC give AMI 0.5499 and 0.5644 respectively, a small gap of +0.014. Coauthor CS provides an external below-capacity negative control: its nearly identical naive and kernel-aware AMI is consistent with the inert behavior observed below capacity in our controlled isolate sweep, but our theory does not predict inertness for arbitrary graphs when $d_0 < k$ (PolBlogs is the counterexample, Section 3).

Scaling on a real undirected graph. As an incidental scaling check, the Z-constrained LOBPCG of Section 4 returned the $k = 15$ kernel-aware eigenpairs of L_{sym} on Coauthor CS in 554 s using a sparse solver; unconstrained sparse LOBPCG requesting $k + r + 10$ eigenpairs took 2,625 s. We do not treat Coauthor CS as a scaling benchmark (its $d_0 = 1$ makes the sparse-naive method not much worse off than the constrained one; the sparse-naive's larger cost here comes from LOBPCG's finite eigengap convergence at $k+r+10 = 26$ eigenpairs rather than from wasted work on nullspace directions), but the numbers confirm that Z-constrained LOBPCG is the correct default when r is not known in advance: building Z takes seconds and either delivers the informative subspace directly (when r is large) or reduces to a small constraint (when r is small).

The structural capacity boundary $d_0 = k$. The general mathematical statement is precise but modest:

$$d_0 \geq k \quad \Longrightarrow \quad \text{the literal bottom-}k \text{ eigenspace lies entirely in } \ker L$$

(in exact arithmetic; what remains optional is only whether the component structure is useful for the label target). This does not universally imply AMI = 0 at $d_0 \geq k$: when the zero eigenvalue is tied beyond the requested dimension, the bottom-k prescription does not select a unique k-dimensional subspace (a right-orthogonal basis rotation $U \mapsto UR$ cannot change row-normalized distances or k-means, but a different admissible k-plane inside the larger zero eigenspace can), and different admissible subspaces may retain different component-level information. What $d_0 \geq k$ certifies is a capacity threshold, past which the conventional bottom-k prescription loses the guarantee of containing any non-null, potentially informative direction.

To measure how the capacity boundary plays out empirically, we run a fine-grained sweep in which we hold community difficulty fixed (each of five undirected DC-SBM core realizations is held fixed across the entire d_0 trajectory) and inject a controlled number of isolated single-vertex components as extra graph vertices, giving $d_0 \in \{1, 2, 3, 6, 9, 12, 24, 48\}$ ($d_0/k \in \{0.17, 0.33, 0.5, 1, 1.5, 2, 4, 8\}$) with $k = 6$. The injected null modes are supported on singleton residual vertices and carry no core-label information by construction. We use the five paired core realizations, with five KMeans seeds per trial.

Below capacity: some naive modes are informative. For $d_0 \leq k - 1$ the naive readout retains at least one informative non-null direction and produces AMI comparable to kernel-aware ($0.62, 0.64, 0.52$ at $d_0 = 1, 2, 3$, matching kernel-aware within KMeans noise, gap in $[-0.005, +0.014]$).

At and above capacity: bottom-k is fully null-supported. At $d_0 = k = 6$ the k smallest eigenvectors of L_{sym} span all of $\ker L_{\text{sym}} = \text{range}(Z)$. The evaluated core lies within one connected component of A_{sym}, so Proposition 2 applies: the core restriction $U_{\text{naive}}\big|_C = (D_\sigma^{1/2}\mathbf{1}_C)\, a^\top$ has rank ≤ 1, and after row-normalization every nonzero row coincides with $a/\|a\|$. KMeans therefore has no within-core information and returns AMI = 0.000 to five decimal places. Kernel-aware bottom-k takes the next k non-null eigenpairs and returns 0.690 ± 0.182. Above $d_0 = k$ the same collapse mechanism applies to whichever k-subset of $\ker L_{\text{sym}}$ the eigensolver picks, so the naive embedding stays rank ≤ 1 and graph-uninformative (adding more null directions does not help), with numerical AMI near zero and tie-breaking dependent, while kernel-aware stays in 0.57–0.68.

Two scope notes. First, whenever the evaluated core lies within a single connected component of A_{sym} and $d_0 \geq k$, Proposition 2 forces the exact core embedding to have rank ≤ 1: null vectors associated with other connected components have identically zero support

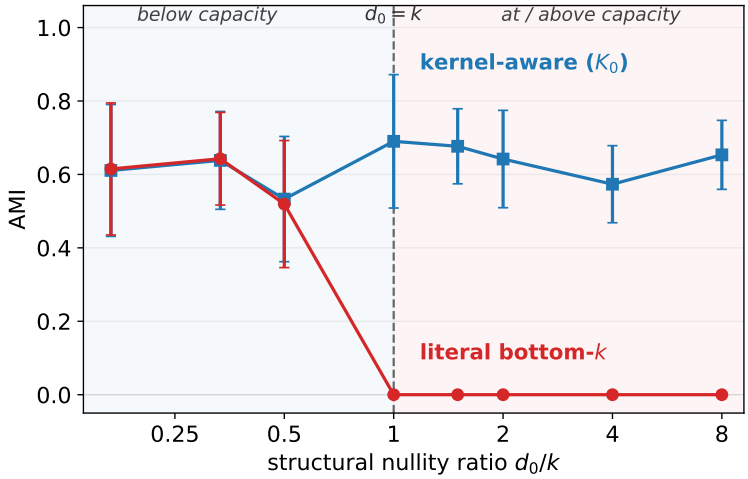

Figure 7: Activation curve of the kernel-aware correction. Same core DC-SBM at every d_0/k value; only the number of injected isolated single-vertex components varies. The evaluated core lies within a single connected component of A_{sym}, so Proposition 2 applies whenever $d_0 \geq k$: the naive core embedding has rank ≤ 1 independently of eigensolver basis choice. Error bars: ± 1 s.d. over $5 \times 5 = 25$ (core-seed, KMeans-seed) pairs. The vertical dashed line marks the capacity boundary $d_0 = k$; naive AMI collapses to 0.000 there and kernel-aware stays at 0.57–0.69. Reproducible by activation_sweep_d0_over_k.py and make_activation_figure.py.

on the core. The absence of within-core graph information is therefore basis-invariant. The particular partition returned by k-means, and hence its AMI, can still depend on numerical and tie-breaking behavior. The non-zero naive AMI $\in [0.00, 0.17]$ we observe in the DC-SBM sweep at $d_0 \geq k$ is not a graph signal from the null modes (Proposition 2 rules that out) but a combination of finite-precision eigensolver tie-breaking on the degenerate zero-eigenspace and k-means's arbitrary partitioning on constant-row embeddings, which need not evaluate to exactly 0.000 on a specific realization. Second, the general result is about capacity, not about AMI: $d_0 = k$ is a bound on how much of the naive embedding can be null. Whether the informal informativeness caveat matters at all is subsumed by whether the evaluated set spans multiple components; under Proposition 2's single-component hypothesis it does not.

Coauthor CS as an external below-capacity control. Coauthor CS has $d_0/k = 1/15 \approx 0.067$, i.e. it sits deep in the below-capacity regime. Naive and kernel-aware bottom-k give nearly identical AMI (0.550 vs. 0.564). This is consistent with the inert behavior our controlled isolate sweep exhibits below capacity, but our theory does not predict inertness for arbitrary graphs when $d_0 < k$ (PolBlogs is the counterexample, Section 3). What Coauthor CS establishes is a real-graph negative control against the reading that the correction is a regularizer that nudges results upward regardless of graph structure.

What the three controls together establish. (a) The nullspace pathology is not intrinsically directed (undirected DC-SBM sweep, mean gap +0.53). (b) $d_0 = k$ is a structural capacity boundary with a clean mathematical interpretation: at $d_0 \geq k$, literal bottom-k lies entirely in the nullspace ($\text{range}(U_{\text{naive}}) \subseteq \ker L$ in exact arithmetic). The controlled uninformative-null experiment saturates this bound into an empirical AMI collapse at $d_0 = k$. (c) Below capacity the effect of retaining structural null modes is graph-dependent: PolBlogs demonstrates it can be severe, Coauthor CS demonstrates it can be nearly inert. Together this makes the correction a mechanism with a mathematically characterized saturation regime and graph-dependent behavior below capacity, not a universally-applied benchmark boost.

E Axis-isolation experiment: generator and full sweep

This appendix isolates a mechanism rather than proposing a universal graph generator: the factorized testbed below exists only to show that the naive-to-K_0 AMI gap tracks structural nullity and not the three competing axes, a scope we make precise in the closing subsection ("What the experiment does and does not show").

E.1 The factorized DC-SBM testbed

Directed graph-clustering benchmarks routinely conflate several structural phenomena; each has been used to motivate a distinct family of spectral methods, and the effects overlap on real graphs. To disentangle nullspace contamination from three other candidate sources of spectral difficulty, block affinity (homophilic vs. heterophilic structure), directional flow (channel split), and reciprocity (directedness), we use a factorized DC-SBM testbed whose four axes are swept one at a time (they are controlled separately, though not fully distributionally independent; see the reciprocity note below). The construction is standard: log-normal degree correction plus per-block-pair multipliers plus a directed edge-probability matrix decomposed into symmetric and antisymmetric parts. We use it here strictly for confounder isolation on axis 3 (spectrum contamination); we do not claim the four-axis parameterization itself as a contribution of this paper. Vanilla DC-SBM, bipartite SBM, and the calibrated DAG-chain used elsewhere in the paper are recovered as parameter settings.

Table 14 summarizes the four axes before the formal construction: each row names the parameter swept along that axis, the graph property that parameter perturbs, and the property the sweep deliberately holds fixed so the perturbation stays confined to a single axis.

Axis	Parameter	What it changes	What stays fixed
Block affinity (Axis 0)	$M_{\text{diag}}/M_{\text{off}}$	homophily vs. heterophily	nullity d_0
Flow (Axis 1)	f	directionality	symmetric block-level support
Reciprocity (Axis 2)	ρ	paired-edge correlation	expected edge count
Contamination (Axis 3)	τ	structural nullity d_0	the core model

Table 14: The four factorized axes of the DC-SBM testbed. Each axis is swept one at a time with all other parameters held at their base-configuration values; the last column names the property the sweep holds approximately fixed, so that the observed change in the naive-to-K_0 AMI gap is attributable, empirically, to the swept axis. In the sweep tables τ abbreviates the contamination parameters τ_{sink} and τ_{source}.

Parameters. Fix a vertex count N and block count k; blocks are balanced. Two per-vertex log-normal degree-correction factors $\theta_i^{\text{out}}, \theta_i^{\text{in}} \sim \text{LogNormal}(0, \sigma_\theta)$ are drawn independently and mean-normalized to 1. Tail-vertex masks then set $\theta_i^{\text{out}} \leftarrow 0$ with probability τ_{sink} and $\theta_i^{\text{in}} \leftarrow 0$ with probability τ_{source}. A $k \times k$ log-normal block-pair multiplier matrix $b_{ab} \sim \text{LogNormal}(0, \sigma_{\text{block}})$ is drawn once per graph and mean-normalized (we write it b_{ab}, not β_{ab}, to avoid collision with the operator measure parameter β).

The per-direction edge probability decomposes as

$$P_{\text{dir}}[a, b] = S[a, b] + f \cdot M_{\text{off}} \cdot A[a, b], \qquad 0 \le f \le 1, \tag{14}$$

where S is a symmetric block-affinity base ($S[a, a] = M_{\text{diag}}$, $S[a, b] = M_{\text{off}}$ for $a \neq b$; homophilic when $M_{\text{diag}} > M_{\text{off}}$, heterophilic otherwise) and A is an antisymmetric flow base with entries in $\{-1, 0, +1\}$; f = flow_asymmetry $\in [0, 1]$ controls the flow-channel scale. Our experiments use the chain pattern $A[a, b] = \text{sign}(b - a)$; hub and cyclic patterns are also implemented but not needed here. Writing c_i for the block label of vertex i (not $\sigma(i)$, again to avoid collision with the operator weight σ), the per-ordered-pair rate is

$$\lambda_{ij} = P_{\text{dir}}[c_i, c_j] \cdot \theta_i^{\text{out}} \cdot \theta_j^{\text{in}} \cdot b_{c_i c_j}, \tag{15}$$

clipped to $[0, 1]$.

Reciprocity coupling. For each unordered pair $\{i, j\}$ with $i < j$, draw $u_{\text{couple}} \sim \text{Uniform}[0, 1]$. If $u_{\text{couple}} < \rho$ (coupled case), draw a single $u_c \sim \text{Uniform}[0, 1]$ and place both directed edges $i \to j$ and $j \to i$ iff $u_c < \frac{1}{2}(\lambda_{ij} + \lambda_{ji})$. Otherwise draw independent $u_f, u_b \sim \text{Uniform}[0, 1]$ and place $i \to j$ iff $u_f < \lambda_{ij}$ and independently $j \to i$ iff $u_b < \lambda_{ji}$. Thus $\rho = 1$ produces a functionally undirected graph and $\rho = 0$ a fully independent directed sampling. This coupling preserves the expected total edge count ($\mathbb{E}[X + Y] = \lambda_{ij} + \lambda_{ji}$) but not the directed marginals: $\Pr(i \to j) - \Pr(j \to i) = (1 - \rho)(\lambda_{ij} - \lambda_{ji})$, so raising ρ also attenuates directional asymmetry when $\lambda_{ij} \neq \lambda_{ji}$. The reciprocity axis is therefore not perfectly independent of the flow axis; the experiment shows only that varying this coupling produces no naive-to-K_0 gap, not that reciprocity and flow are distributionally orthogonal. The coupling parameter ρ is also not the realized edge reciprocity, which depends on the generated $(\lambda_{ij}, \lambda_{ji})$ distribution.

Mapping to the four axes. The controlling parameters per axis are those of Table 14, with two clarifications. Axis 1 also selects the flow pattern (chain here), and Axis 3 (spectrum contamination) sweeps only the tail-vertex masks $\tau_{\text{sink}} = \tau_{\text{source}} = \tau$, with σ_{block} and σ_θ held at their base values: the double-masked vertices it creates enlarge $\ker L_{\text{sym}}$ and raise d_0. Axis 3 is thus a robustness sweep in a realistic contamination regime; the clean one-variable isolation of d_0 is the singleton-injection sweep of Section 5 (Figure 2), which varies only the number of injected isolated components. The other three axes leave the block-topological connectivity of the padded symmetric hull intact.

E.2 Base configuration

Unless otherwise stated, every sweep uses $N = 1500$, $k = 6$ (balanced blocks of 250), $M_{\text{diag}} = 0.03$, $M_{\text{off}} = 0.003$, $\rho = 0.5$, $\sigma_\theta = 0$, $\tau_{\text{sink}} = \tau_{\text{source}} = 0$, $\sigma_{\text{block}} = 0$, and $f = 0$ with the chain flow pattern. Each swept parameter is varied along its axis with all other parameters held at these values. For the block-affinity axis we vary the ratio $M_{\text{diag}}/M_{\text{off}}$ while holding the unweighted block-matrix base rate $(M_{\text{diag}} + 5M_{\text{off}})/6 = 0.005$ fixed, so the homophily sweep holds the mean density fixed up to a finite-size diagonal correction (within-block pairs number $n(n-1)$ rather than n^2) rather than confounding affinity with density. For the contamination axis, we report the measured d_0 (connected components of the padded symmetric hull) at each point rather than a masking probability: a one-sided zero-degree mask (probability $\tau_{\text{sink}} + \tau_{\text{source}} - \tau_{\text{sink}}\tau_{\text{source}}$) does not by itself isolate a vertex, since a pure sink still has incoming edges; only the double mask (probability $\tau_{\text{sink}}\tau_{\text{source}}$) forces isolation, and additional fragmentation arises from sparse Bernoulli sampling. Ten graph seeds per point; five k-means seeds at $n_{\text{init}} = 200$ per graph seed. K_0 and naive AMI are computed on the largest connected component of the padded symmetric hull, matching the evaluation protocol of Sections 2–3.

E.3 Sweep results

Full per-point statistics for the four axes appear below. Each row is the mean ± standard deviation over 10 graph seeds; each AMI number is itself the mean over five KMeans seeds at $n_{\text{init}} = 200$ inside the fit for that graph. d_0/k is the ratio of the number of connected components of the padded symmetric hull to the block count; median η is the projector-based null-occupancy of Proposition 3 evaluated on the nonzero-row vertices of the largest connected component. $\Delta\text{AMI} = \text{AMI}(K_0) - \text{AMI}(\text{naive})$. We verified the algebraic invariant $\|\Pi_0 U - U\|_F / \|U\|_F$ per seed whenever $d_0 \geq k$: on every graph seed of every sweep it stayed $\leq 5 \times 10^{-15}$, confirming $\text{range}(U) \subseteq \text{range}(\Pi_0)$. On rows entering that regime $\eta = 1$ exactly on every nonzero row (rows with $\|e_i^\top U\| = 0$ are excluded from the median, per the definition of Proposition 3); the reported "$1.000^\dagger$" cells denote these exact-collapse points.

E.4 Seed split at $\tau = 0.05$: below-capacity validation

The aggregate $\tau = 0.05$ point of Table 15 is a mixed regime: three of the ten graph seeds satisfy $d_0 \geq k$ and are governed by Proposition 2, while the remaining seven have $d_0 < k$ and are governed by Proposition 3. We split the ten seeds and report each subset separately.

Axis	Parameter	d_0/k	median η	AMI(naive)	AMI(K_0)	ΔAMI
Axis 0 ($M_{\text{diag}}/M_{\text{off}}$)	0.2	0.17	0.192	-0.002 ± 0.001	-0.002 ± 0.001	$+0.000 \pm 0.001$
	0.5	0.17	0.200	-0.001 ± 0.001	-0.001 ± 0.001	-0.000 ± 0.001
	1	0.17	0.208	$+0.000 \pm 0.001$	$+0.000 \pm 0.001$	$+0.000 \pm 0.001$
	2	0.18	0.209	$+0.004 \pm 0.002$	$+0.006 \pm 0.003$	$+0.001 \pm 0.002$
	5	0.17	0.192	$+0.472 \pm 0.045$	$+0.484 \pm 0.032$	$+0.012 \pm 0.019$
	10	0.17	0.176	$+0.925 \pm 0.009$	$+0.923 \pm 0.009$	-0.002 ± 0.004
Axis 1 (flow f)	0	0.17	0.172	$+0.989 \pm 0.005$	$+0.989 \pm 0.004$	-0.000 ± 0.002
	0.2	0.17	0.172	$+0.990 \pm 0.005$	$+0.989 \pm 0.004$	-0.002 ± 0.002
	0.4	0.17	0.172	$+0.990 \pm 0.003$	$+0.989 \pm 0.003$	-0.001 ± 0.002
	0.6	0.17	0.173	$+0.988 \pm 0.003$	$+0.989 \pm 0.004$	$+0.000 \pm 0.001$
	0.8	0.17	0.173	$+0.989 \pm 0.005$	$+0.989 \pm 0.005$	-0.000 ± 0.002
	1	0.17	0.173	$+0.989 \pm 0.005$	$+0.989 \pm 0.005$	$+0.001 \pm 0.002$
Axis 2 (ρ)	0	0.17	0.171	$+0.999 \pm 0.002$	$+0.998 \pm 0.002$	-0.000 ± 0.001
	0.25	0.17	0.172	$+0.996 \pm 0.003$	$+0.995 \pm 0.003$	-0.000 ± 0.001
	0.5	0.17	0.172	$+0.989 \pm 0.005$	$+0.989 \pm 0.004$	-0.000 ± 0.002
	0.75	0.17	0.174	$+0.967 \pm 0.011$	$+0.967 \pm 0.009$	-0.000 ± 0.002
	1	0.17	0.176	$+0.931 \pm 0.013$	$+0.932 \pm 0.013$	$+0.001 \pm 0.003$
Axis 3 (τ)	0	0.17	0.172	$+0.989 \pm 0.005$	$+0.989 \pm 0.004$	-0.000 ± 0.002
	0.05	0.75	0.612	$+0.491 \pm 0.361$	$+0.967 \pm 0.008$	$+0.475 \pm 0.361$
	0.1	2.67	$1.000^{\dagger}$	$+0.004 \pm 0.011$	$+0.945 \pm 0.008$	$+0.941 \pm 0.016$
	0.15	5.88	$1.000^{\dagger}$	$+0.001 \pm 0.003$	$+0.919 \pm 0.008$	$+0.917 \pm 0.008$
	0.2	10.10	$1.000^{\dagger}$	$+0.000 \pm 0.003$	$+0.885 \pm 0.014$	$+0.885 \pm 0.014$
	0.25	16.02	$1.000^{\dagger}$	$+0.001 \pm 0.003$	$+0.854 \pm 0.011$	$+0.852 \pm 0.011$
	0.3	22.67	$1.000^{\dagger}$	$+0.002 \pm 0.003$	$+0.803 \pm 0.023$	$+0.801 \pm 0.023$
	0.35	30.17	$1.000^{\dagger}$	-0.001 ± 0.002	$+0.749 \pm 0.037$	$+0.749 \pm 0.037$

Table 15: Full per-point statistics for the axis-isolation sweep of Figure 3. Only Axis 3 (spectrum contamination) drives d_0/k above the capacity boundary; only Axis 3 activates a material ΔAMI. Homophily, flow, and reciprocity keep d_0/k, η, and ΔAMI near their baseline values. †: at $d_0 \geq k$, in exact arithmetic Proposition 2 gives $\Pi_0 U = U$ and hence $\eta_i = 1$ on every nonzero core row; numerically we verify relative residual $\|\Pi_0 U - U\|_F / \|U\|_F \leq 5 \times 10^{-15}$, so the reported 1.000 is exact up to floating-point error.

Subset	n	d_0 range	median η	AMI(naive)	AMI(K_0)	ΔAMI
All 10 seeds	10	$[2, 7]$	0.612 ± 0.303	$+0.491 \pm 0.361$	$+0.967 \pm 0.008$	$+0.475 \pm 0.361$
Below capacity ($d_0 < k$)	7	$[2, 5]$	0.445 ± 0.197	$+0.700 \pm 0.203$	$+0.966 \pm 0.007$	$+0.267 \pm 0.200$
At/above capacity ($d_0 \geq k$)	3	$[6, 7]$	$1.000^{\dagger}$	$+0.004 \pm 0.004$	$+0.967 \pm 0.010$	$+0.963 \pm 0.014$

Table 16: Split of the $\tau = 0.05$ Axis-3 seeds by capacity. Reproducible by axis_isolation_below_capacity.py.

Within the below-capacity subset the seven per-seed pairs $(\eta, \Delta\text{AMI})$ are:

Seed	107	109	102	101	106	103	100
d_0	2	2	3	4	5	5	5
median η	0.212	0.210	0.276	0.458	0.607	0.662	0.692
ΔAMI	+0.008	+0.027	+0.092	+0.342	+0.494	+0.464	+0.438

η and ΔAMI show a strong positive rank association (the sequence is not strictly monotone): Spearman $\rho_s = 0.82$; exact two-sided permutation $p = 0.034$ over the $7! = 5040$ rank permutations (Pearson $r = 0.97$). With $n = 7$ we treat this as exploratory. The same rank correlation and exact p hold under ΔNMI and ΔARI: rerunning the seven below-capacity seeds with all three metrics (metric_robustness.py) yields $\rho_s(\eta, \Delta\text{AMI}) = \rho_s(\eta, \Delta\text{NMI}) = \rho_s(\eta, \Delta\text{ARI}) = 0.821$, so the below-capacity positive rank association is metric-robust. Combining with the two real benchmarks gives three qualitatively distinct below-capacity data points:

- Email-Eu (median $\eta = 0.035$, $\Delta\text{AMI} = -0.003$): non-null block dominates the row norm; correction inert.
- Axis-3 $\tau = 0.05$ below-capacity seeds (median $\eta \in [0.21, 0.69]$, $\Delta\text{AMI} \in [+0.01, +0.49]$): partial null occupancy; ΔAMI grows with η.
- PolBlogs (median $\eta = 1.000$, $\Delta\text{AMI} = +0.838$): approximate collapse; the correction is decisive.

Proposition 3 bounds the row-normalized embedding to a ball of radius $2\sqrt{1 - \min \eta}$; the monotone empirical relationship above is consistent with, but does not prove, the stronger statement that η predicts ΔAMI across arbitrary graph families. Broader external validation remains open.

E.5 What the experiment does and does not show

The experiment isolates a mechanism, not a universal law. It shows that within a factorized DC-SBM parameter family calibrated to citation-network densities, the naive-to-K_0 AMI gap is insensitive to three plausible confounders (block affinity, flow, reciprocity) and activates only through the axis that grows the structural nullspace. The generator's other reductions (bipartite SBM, calibrated DAG-chain) instantiate the same insensitivity in different corners of the parameter space. We do not claim that no graph family exists on which the other axes activate the correction; we claim only that within a natural realistic regime the correction's activation is not confounded with directionality or heterophily.

F Synthetic DAG-chain DC-SBM with tail vertices

F.1 Purpose and calibration targets

The synthetic generator of Section D.1 exists to reproduce, in a controlled setting, the naive-to-corrected gap observed on the real citation graphs. A vanilla directed DC-SBM does not: without the two structural ingredients below, all spectral methods saturate together near AMI 1.0 across the full reciprocity range, so the failure mode never appears. Reproducing it requires both (i) tail vertices, which create a non-trivial residual zero-mode kernel after self-loop padding, and (ii) per-block-pair density heterogeneity. Both are calibrated from statistics measured on the four benchmark citation networks (Cora, Citeseer, Cora-ML, Citeseer-full) before any reciprocity sweep was run. This calibration is necessary to reproduce the empirical regime of interest, because vanilla directed DC-SBMs do not exhibit the same failure mode.

F.2 Generator

Block structure. Vertices are partitioned into k blocks of equal size. Edge probabilities depend on block-pair direction:

$$\bar{M}_{a,b} \;=\; \begin{cases} p_{\text{in}} & \text{if } a = b, \\ p_{\text{fwd}} & \text{if } b > a, \\ p_{\text{back}} = p_{\text{fwd}}/R & \text{if } b < a. \end{cases}$$

The chain structure $0 \to 1 \to \cdots \to k-1$ is encoded by $p_{\text{fwd}} \gg p_{\text{back}}$ (high-DAG regime, $R = 25$).

Block-pair density heterogeneity. Real citation graphs have substantial heterogeneity in per-block-pair edge density: some block pairs cite each other heavily, others sparsely. We multiply each entry of $\bar{M}$ by an independent log-normal factor with scale σ_{block}:

$$M_{a,b} \;=\; \bar{M}_{a,b} \cdot e^{\eta_{a,b}}, \qquad \eta_{a,b} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2_{\text{block}}),$$

with the resulting $\{e^{\eta_{a,b}}\}$ renormalized to unit mean, which preserves the calibration-target density in expectation before clipping (the realized graph-wide mean differs slightly through

the base-rate weights $n_a n_b \bar{M}_{a,b}$ over distinct vertex pairs and any probability clipping). Setting $\sigma_{\text{block}} = 0$ removes this ingredient, at which point bottom-k SC on the symmetrized adjacency saturates near AMI 1.0 across the entire reciprocity range and hides the failure mode.

Independent in/out heterogeneity. Each vertex i carries two independent degree-correction factors $\theta_i^{\text{out}}, \theta_i^{\text{in}}$ drawn i.i.d. from a log-normal with scale σ_θ, normalized to mean one. The reference Bernoulli probability for $i \to j$ is

$$\lambda_{i\to j} \;=\; \min\big(M_{b(i),b(j)} \cdot \theta_i^{\text{out}} \cdot \theta_j^{\text{in}},\; 1\big).$$

Independence between θ^{out} and θ^{in} produces nodes that are heavy out-emitters but light in-receivers and vice versa, matching the citer-vs-cited heterogeneity of citation networks.

Tail vertices. A fraction τ_{sink} of vertices have $\theta_i^{\text{out}} := 0$ (pure sinks with $d_{\text{out}} = 0$); independently a fraction τ_{source} have $\theta_i^{\text{in}} := 0$ (pure sources with $d_{\text{in}} = 0$). The one-sided zero-degree fraction $\tau_{\text{sink}} + \tau_{\text{source}} - \tau_{\text{sink}}\tau_{\text{source}}$ is not the resulting nullity: a pure sink still receives incoming edges and a pure source still emits, so neither is isolated in the symmetric support; only the double-masked vertices (probability $\tau_{\text{sink}}\tau_{\text{source}}$) are forced isolates, and further fragmentation arises stochastically from sparse Bernoulli sampling. We therefore report the measured structural nullity, $d_0 \in [374, 1078]$ across the 55 realizations of the sweep: the structure the band-limit cap of Section 3 is calibrated for.

Reciprocity coupling. For each unordered pair $\{i, j\}$ a coupling coin Bern(ρ) is drawn. If the coin is 0, the two directions are sampled independently with their respective λ-values. If the coin is 1, the two directions are sampled jointly: a single uniform u is drawn, and both edges exist iff $u < (\lambda_{i\to j} + \lambda_{j\to i})/2$. The coupling rule is chosen so that the expected directed edge count is constant in ρ, isolating the reciprocity effect from edge-density confounds.

F.3 Parameter calibration

Table 17 lists each parameter, the empirical value measured on the citation graphs, and the value used in the generator. All parameters were locked to these targets before the reciprocity sweep was run; $R = 25$ is the only free design choice (high-DAG regime), motivated by the very low reciprocity of citation graphs ($\rho \le 0.06$ on the four benchmark cores), for which a clean forward-flow regime is appropriate.

Property	Empirical (citation graphs)	Synthetic value
Block-pair density log-std σ_{block}	1.56/1.41/1.50/2.12 (mean of first three 1.49)	1.5 (Cora)
Log-degree std σ_θ	0.69 (mean over the four)	0.7
Sink fraction τ_{sink}	0.29	0.30
Source fraction τ_{source}	0.38	0.30
Reciprocity ρ	≤ 0.06	swept $\{0.0, \ldots, 1.0\}$
DAG ratio $R = p_{\text{fwd}}/p_{\text{back}}$	(design choice)	25

Table 17: Calibration of the synthetic generator: each parameter is matched to a statistic of the four benchmark citation networks, measured before any sweep. The reciprocity sweep additionally fixes $N = 1500$, $k = 6$, $p_{\text{in}} = 0.013$, $p_{\text{fwd}} = 0.0007$, and $p_{\text{back}} = p_{\text{fwd}}/25 = 2.8 \times 10^{-5}$, with 5 trials at each of $\rho \in \{0.0, 0.1, \ldots, 1.0\}$ (55 realizations).

F.4 Robustness to the calibration target

To check that the effect is not an artifact of the single calibration choice $\sigma_{\text{block}} = 1.5$ (matched to Cora), we re-ran the sweep at the weaker target $\sigma_{\text{block}} = 1.0$ with all other parameters identical, across $\rho \in \{0.0, 0.1, \ldots, 0.7\}$ with 5 trials per ρ except $\rho = 0.7$ (three trials; $7 \times 5 + 3 = 38$ realizations). The findings of Section D.1 reproduce qualitatively: (i) bottom-k SC on the symmetrized adjacency fails on all 38/38 realizations (AMI ≤ 0.002 every trial), and (ii) the legacy filtered-FIPS variant at the oracle β dominates Bethe–Hessian on 37/38 realizations, with mean AMI margin $+0.113 \pm 0.074$ (vs $+0.19 \pm 0.13$

at $\sigma_{\text{block}} = 1.5$) and mean AMI 0.632. The wedge is smaller at $\sigma_{\text{block}} = 1.0$, as expected (less block-pair heterogeneity is an easier regime for Bethe–Hessian), but the mechanism is robust. The kernel still dominates: mean $d_0 = 532$ over the 38 realizations ($d_0 \in [316, 960]$), comparable to $d_0 \in [374, 1078]$ at $\sigma_{\text{block}} = 1.5$.

F.5 Scope and limitations

Three caveats bound what this synthetic sweep establishes. First, it exercises the legacy filtered-FIPS readout (a prefix-selection rule produced by fips_core), not the Slepian S_3 operator of the main text; we therefore read it as corroborating only the kernel-collapse mechanism and the general value of a concentrated post-nullspace readout, not as validating the S_3 pipeline or a label-free β-selector. Second, the reported d_0 is the component count (the GSC operator's padded nullity), whereas the legacy sweep's bottom-k baseline builds L_{sym} by clamping zero degrees to one without self-loop padding, so that operator's own numerical nullity is smaller (measured 16–36 on a sample; isolated vertices sit at eigenvalue 1 rather than 0). That numerical nullity is nonetheless $\geq k$ on every realization; under the single-component-core condition of Proposition 2, the same rank-≤ 1 collapse then applies to this clamped-degree operator, even though it uses degree-clamping rather than the self-loop padding assumed there. Only the printed d_0 overstates the nullity of that particular operator. Third, we did not test the strongest empirical target $\sigma_{\text{block}} = 2.12$ (Citeseer-full): at $N = 1500$ the log-normal mean-normalization concentrates mass on a few block-pair entries and shrinks the labeled core to ~ 10 vertices; that regime becomes reachable only at larger N and is left to follow-up work.

Reproducibility. The $\sigma_{\text{block}} = 1.5$ sweep is generated by synthetic_dag_chain_tails.py sweep –n_trials 5 and saved in figure5b_data_sigma15.json; the $\sigma_{\text{block}} = 1.0$ robustness run (legacy fips_core readout via run_synthetic_robustness.py) is saved in figure5b_data_sigma10_partial.json.

G Legacy Band-Cap Sensitivity: Provenance of the Exploratory S_3 Default

The Slepian variant S_3 of Section 5.3 uses a post-nullspace band of dimension $m = \alpha k$ with $\alpha = 3$. This value was inherited from an earlier band-cap sensitivity sweep rather than selected independently for the Slepian operator, and we report that sweep here solely to document the provenance of $\alpha = 3$; it should not be read as a direct validation of the Slepian S_3. Two qualifications follow from how the sweep was run. First, it used a capped-FIPS pipeline that selects individual post-nullspace eigenvectors by core support and inverse participation ratio, and does not build the Slepian concentration operator $K = \Pi_B P_C \Pi_B$; it therefore shows how a post-nullspace band cap behaves in general, not how the Slepian band cap behaves in particular. Second, the sweep inspects AMI on the benchmark graphs, so $\alpha = 3$ is a benchmark-selected default, not a label-free constant, and S_3 is best read as an exploratory, benchmark-tuned variant.

Sweep on the binding graphs. On the three graphs where the cap was binding (Email-Eu, Citeseer, Citeseer-full), Table 18 sweeps $\alpha \in \{2, \ldots, 8\}$ against the uncapped bottom-prefix baseline F0 (prefix width $0.95N$): every $\alpha \leq 5$ strictly improves on F0, $\alpha \geq 6$ collapses back toward it, and $\alpha = 3$ is the best single fixed default: the maximum on Email-Eu and Citeseer, and second only to $\alpha = 2$ on Citeseer-full.

Inert verification. On the three benchmark graphs where the cap was inert at $\alpha = 4$ in an earlier sweep (PolBlogs, Cora, Cora-ML), we verified that $\alpha = 3$ does not regress. Cora is not actually inert: $\alpha = 3$ strictly improves AMI by $+0.042$ over $\alpha = 4$ (0.526 vs 0.484 at the supervised $\beta^\star$). PolBlogs and Cora-ML reproduce the uncapped baseline (F0) to within k-means re-initialization noise.

Why a graph-adaptive α is plausible but deferred. The optimal α is mildly graph-dependent: Citeseer-full peaks at $\alpha = 2$ ($+0.10$ above $\alpha = 3$), Email-Eu peaks at $\alpha = 3$ ($+0.013$ above $\alpha = 4$), and Cora gains monotonically as α decreases from 4 to 3. A natural

Graph	F0	$\alpha = 2$	$\alpha = 3$	$\alpha = 4$	$\alpha = 5$	$\alpha = 6$	$\alpha = 8$
Email-Eu	0.562	0.585	0.704	0.693	0.673	0.557	0.557
Citeseer	0.438	0.539	0.546	0.470	0.471	0.450	0.439
Citeseer-full	0.621	0.751	0.651	0.638	0.641	0.633	0.621

Table 18: α-sensitivity of the legacy capped-FIPS pipeline (individual-mode/IPR selection, not the Slepian S_3 operator), on the three binding graphs. Bold marks the per-graph best α; F0 is the uncapped bottom-prefix baseline (prefix width $0.95N$). Reproducible by test_capped_selector.py.

Graph	F0	$\alpha = 3$	Δ
PolBlogs	0.832	0.832	+0.000
Cora	0.467	0.526	+0.059
Cora-ML	0.558	0.559	+0.001

Table 19: Inert-graph verification: $\alpha = 3$ does not regress on graphs where the cap was inert at $\alpha = 4$. Cora gains +0.059 AMI; the others are within numerical noise of F0. These legacy capped-FIPS values (at the supervised $\beta^\star$) differ by ≤ 0.014 from the locked Slepian-S_3 ladder (Table 5, Cora 0.539, Cora-ML 0.573).

graph-adaptive form $m_{\max} = d_0 + \alpha(\text{graph}) \cdot k$ with α set by a label-free spectral-density statistic (e.g., the local DOS slope at $\lambda \sim 0$) is plausible. We do not propose one in this paper because (i) $\alpha = 3$ performs well across the benchmark and is the best single fixed setting among the fully swept binding graphs except Citeseer-full (where the optimum is $\alpha = 2$), and (ii) a graph-adaptive selector would require its own validation with no immediate empirical headroom on this benchmark.

Reproducibility. The band-cap sweep is generated by test_capped_selector.py (over $\alpha \in \{2, 3, 4, 5, 6, 8\}$ on the binding graphs and $\alpha = 3$ on the inert graphs), and the full sensitivity data is recorded in cap_alpha_sensitivity_results.md. The authoritative source for the Slepian S_3 numbers themselves is slepian_ablation_ladder.py (Table 5), not this legacy sweep.

H Core-only Spectral Clustering: A Non-Substitute for Kernel Awareness

K_0 evaluates on a selected labeled core C but computes the bottom-k eigenpairs of L_{sym} on the full padded graph, removing only the structural nullspace. A natural reviewer question is: why not delete everything outside C and spectral-cluster the induced core? Formally, form the induced subgraph $G[C]$ on A_{sym}, its symmetric normalized Laplacian $L_{\text{sym}}(G[C])$, take the first k non-null eigenvectors (the same kernel-aware indexing as K_0, so the comparison isolates vertex deletion rather than the indexing rule), and k-means the rows. We call this Core-first.

Core-first is a strictly different object from K_0. Discarding residual vertices is not equivalent to removing residual null directions, because the two operations couple through the off-diagonal blocks of L_{sym} that link C to R_+ (residual vertices connected to the core). Table 20 makes this concrete.

The scientific implication is conceptual, not numerical: discarding residual vertices $\neq$ removing residual null directions. Core-first changes the problem: it deletes vertices and rebuilds the operator on a smaller graph. K_0 changes only the spectral search space, leaving every vertex and every edge in place and skipping only the structurally-known null eigenvectors. The two operations are not equivalent in general, and the numerical margin between them is graph-specific; we draw no monotone conclusion from Table 20. The purpose of this appendix is to establish that the two are different objects, not that K_0 uniformly dominates Core-first.

Graph	$\lvert C\rvert/N$	Core-first (AMI)	K_0 (AMI)	Δ = Core-first $- K_0$
PolBlogs	0.58	0.832 ± 0.000	0.839	-0.007
Email-Eu	0.73	0.695 ± 0.004	0.698	-0.003
Cora	0.34	0.432 ± 0.000	0.505	-0.073
Citeseer	0.13	0.462 ± 0.000	0.546	-0.084
Cora-ML	0.43	0.524 ± 0.000	0.466	$+0.058$
Citeseer-full	0.05	0.571 ± 0.000	0.691	-0.120

Table 20: Core-first spectral clustering vs. kernel-aware K_0, both using the first k non-null modes so the comparison isolates vertex deletion. 5-seed KMeans means at $n_{\text{init}} = 200$ using the same seeds $\{42, 43, 44, 45, 46\}$ as the K_0 headline (Core-first seed-variance is ≤ 0.004 across all six graphs). Core-first is worse than K_0 on PolBlogs, Email-Eu, Cora, Citeseer, and Citeseer-full ($\Delta \in [-0.120, -0.003]$), and better on Cora-ML ($\Delta = +0.058$). The largest gap is on Citeseer-full ($|C|/N = 0.05$), where core-first discards 95% of the graph; the gap is consistent with a loss of information carried through C–R_+ coupling, though Table 20 alone does not isolate that channel (an edge-removal or coupling-norm ablation would). We do not claim monotonicity in $|C|/N$: PolBlogs at $|C|/N = 0.58$ sits well above Cora at $|C|/N = 0.34$ in the Δ column, consistent with PolBlogs' bulk-community structure being carried mostly by C while more of Cora's appears to run through C–R_+ coupling. Reproducible by iclr2027/corefirst_5seed.py.

I Auxiliary Theory for the Exploratory S_3 Readout

This appendix collects the stability theory behind the exploratory S_3 pipeline of Section 5.3. We keep it separate from the structural-nullspace proofs of Appendix A because it concerns the diagnostic S_3 readout rather than the K_0 correction that is the paper's central claim; K_0 does not use β at all.

The $\beta^\star$ used for S_3 in the main text is a fixed, supervised selector value (a single-seed, locally-AMI-smoothed pick, not a pointwise AMI oracle; Appendix C), so S_3 is reported as a supervised, benchmark-tuned estimate. A natural label-free replacement would choose $\beta^\star = \arg\max_\beta \mathrm{Stab}(\beta; \Delta)$, where $\mathrm{Stab}(\beta; \Delta)$ is the mean canonical correlation between the k-dimensional S_3 output subspace at β and at $\beta \pm \Delta$. The following proposition establishes the local Lipschitz stability of that subspace in β, which is the property such a selector would rely on; it does not evaluate the selector.

The correct invariant object is the output subspace, not any particular eigenvector basis: an eigenvector matrix $U(\beta)$ can jump in sign or rotate within a degenerate block and need not be continuous, whereas the orthogonal projector onto its span can be. We therefore state the stability result for the projector.

Proposition 4 (Local Lipschitz stability of the S_3 subspace). Fix the graph and its self-loop-padded support. Let $J \subseteq [0, 1]$ be a compact interval on which $\beta \mapsto \bar{L}_{\nu_\beta}$ is C^1 with $\nu_\beta \in \mathbb{R}^N_{>0}$ entrywise (so $\sigma_\beta > 0$ and, by Theorem 1, the nullity r is constant on J). Let $m := \alpha k \in \mathbb{N}$ with $1 \leq k < m < N - r$. Assume a uniform separation from zero and a uniform band-boundary gap,

$$\gamma_0 := \inf_{\beta \in J} \min\{\lambda_1^+(\beta),\ \lambda_{m+1}^+(\beta) - \lambda_m^+(\beta)\} > 0,$$

where $\lambda_1^+ \leq \lambda_2^+ \leq \cdots$ are the positive eigenvalues, and a uniform Slepian eigengap

$$g_0 := \inf_{\beta \in J} \big[\tau_k(\beta) - \tau_{k+1}(\beta)\big] > 0,$$

where $\tau_1(\beta) \geq \cdots \geq \tau_N(\beta)$ are the eigenvalues of the concentration operator $K(\beta) := \Pi_B(\beta) P_C \Pi_B(\beta)$, with $\Pi_B(\beta)$ the projector onto the first $m = \alpha k$ positive modes and P_C the fixed diagonal core projector. Let $\Pi_U(\beta)$ be the projector onto the top-k eigenspace of $K(\beta)$ and $C_A := \sup_{\beta \in J} \|\partial_\beta \bar{L}_{\nu_\beta}\|_2$. Then, with derivatives interpreted one-sidedly at the endpoints of J,

$$\|\partial_\beta \Pi_U(\beta)\|_2 \leq \frac{8\, C_A}{\gamma_0\, g_0} =: C_U,$$

so the output subspace projector is Lipschitz on J. Consequently, for $\beta, \beta + h \in J$ with $C_U|h| \leq 1$, the mean canonical correlation between the ambient orthonormal S_3 frames, $\mathrm{corr}(\beta, \beta') := \frac{1}{k}\|U(\beta)^\top U(\beta')\|_*$ (nuclear norm), satisfies

$$\mathrm{corr}(\beta, \beta + h) \geq \sqrt{1 - C_U^2 h^2} \geq 1 - C_U^2 h^2.$$

Proof. Because P is fixed and ν_β is C^1 and entrywise positive, the derived quantities $\xi_\beta = P^\top \nu_\beta$, $\sigma_\beta = \nu_\beta + \xi_\beta$, and hence $\bar{L}_{\nu_\beta}$ are C^1, with $\|\partial_\beta \bar{L}_{\nu_\beta}\|_2 \leq C_A$. Because γ_0 separates the first αk positive modes from both zero and the rest of the positive spectrum, Davis–Kahan (Davis and Kahan, 1970) gives $\|\partial_\beta \Pi_B\|_2 \leq 2C_A/\gamma_0$; then $\|\partial_\beta K\|_2 \leq 2\|\partial_\beta \Pi_B\|_2$, and a second Davis–Kahan under the Slepian gap g_0 gives $\|\partial_\beta \Pi_U\|_2 \leq 8C_A/(\gamma_0 g_0) =: C_U$. Integrating along $[\beta, \beta + h]$ gives $\|\Pi_U(\beta + h) - \Pi_U(\beta)\|_2 \leq C_U|h|$; for $C_U|h| \leq 1$ every canonical correlation then satisfies $\cos\theta_j \geq \sqrt{1 - C_U^2 h^2}$, and the corr bound follows from $\sin\theta_{\max}(\mathcal{E}, \mathcal{F}) = \|\Pi_\mathcal{E} - \Pi_\mathcal{F}\|_2$ for equal-rank subspaces. Orthonormal frames $U(\beta)$ can be made Lipschitz only after a β-dependent orthogonal (gauge) alignment; the statement concerns the subspace before row normalization, and stability of the row-normalized embedding or of discrete k-means labels would require additional row-norm and clustering-margin conditions. □

The proposition guarantees only local smoothness of the output-subspace projector $\Pi_U(\beta)$ under two eigengap assumptions. It does not guarantee that a discrete selector $\beta^\star = \arg\max_\beta \mathrm{Stab}(\beta; \Delta)$ searching over $[0, 1]$ lands in J (a more stable basin outside J could win), and it does not establish that a stability-maximizing β recovers the best achievable AMI: Lipschitz continuity of $\beta \mapsto \Pi_U(\beta)$ constrains how fast the subspace can move, not where its clustering accuracy is maximized. We therefore do not report a label-free selector-vs-oracle gap; the $\beta^\star$ of Section 5.3 are fixed, supervised selector values (not pointwise AMI oracles) and S_3 is a supervised, benchmark-tuned estimate. Either way, this concerns only the S_3 pipeline and does not affect the paper's central claim about K_0, which does not use β at all.