A Projective Theory of Local Stability for Generalized Spectral Clustering

Anonymous authors
Paper under double-blind review

Abstract

Generalized Spectral Clustering (GSC) on directed graphs depends on a tunable vertex measure, but it is unclear when changing this measure can change the resulting partition. We show that a naïve perturbation analysis can be misleading: positive rescaling of the measure rescales the GSC Laplacian but leaves its eigenspaces (and, whenever the bottom-k target is well defined, the clustering) unchanged. We therefore study stability after quotienting out this irrelevant scale. Before measuring stability, we check whether the target bottom-k eigenspace is uniquely defined, and give a cheap, linear-time structural test that detects an important class of non-identifiable cases. When the target is identifiable, we derive a local anchored-rounding stability radius that factors into three interpretable terms: how strongly the measure moves the operator, how well the target eigenspace is spectrally isolated, and how robust the embedding is to rounding. Local certificates then compose through continuation. Experiments on directed graphs show that removing scale changes the inferred stability by an order of magnitude, that clustering quality and stability need not coincide, that the radius tracks directly verified finite perturbations, and that local certificates compose through continuation into reachable regions more than 58× larger.

1 Introduction

Generalized Spectral Clustering (GSC; Sevi et al. (2026)) extends spectral clustering by introducing a positive vertex measure ν as an algorithmic degree of freedom, so that a family of Laplacians L_ν arises from a single graph as ν is tuned. A canonical two-parameter family, indexed by a diffusion time $t \in \{1, 2, \ldots\}$ and a sharpness exponent $\alpha > 0$, provides a common parameterised family covering several spectral clustering constructions: unnormalized, random-walk, and diffusion-weighted operators arise at specific choices of ν (Sevi et al., 2026). The original GSC work showed empirically that different measure choices can lead to either identical or different partitions but did not characterise theoretically when this choice matters. This paper addresses that gap: how does the positive vertex measure affect the clustering produced by GSC?

It is useful to split this into three questions. (Q1) Which changes of ν are meaningful, which can move the partition, and which are pure gauge? (Q2) When a meaningful change occurs, is the spectral object it perturbs identifiable? (Q3) When it is, what graph properties determine how much measure motion the clustering tolerates? Answering these gives a stability theory of the measure-to-partition map,

$$\vartheta \longrightarrow \nu_\vartheta \longrightarrow L_{\nu_\vartheta} \longrightarrow U_\vartheta \longrightarrow A_\vartheta,$$

whose behaviour is graph-dependent: on some graphs many ϑ recover the same partition; on others, small changes produce different partitions.

Two features shape how the question must be studied. Measures that differ only by positive scale, $\nu \sim c\nu$, produce the same clustering, so stability lives on measure rays rather than raw measures: we work on the projective quotient of the nonnegative measure cone. And because many parameter values can produce the same partition, the object of interest is not

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.

a single best ϑ but a region of parameter space over which the partition remains unchanged. Our goal is to prove that every parameter in a neighbourhood of a reference point produces the same partition under anchored rounding; we call such a neighbourhood a certified stable region, and we certify local ones rather than merely observing them (the precise stability set is deferred to §2).

Two issues prevent a standard perturbation argument from applying directly. The first is scale: L_ν is linear in ν with $L_{c\nu} = c\,L_\nu$ (Sevi et al., 2026), so rescaling leaves the eigenspaces and partition unchanged. Yet the raw family $\nu_{t,\alpha}$ has total mass varying by three orders of magnitude across our grid; a theory built on L_ν would miscount this motion as sensitivity. Even after removing this irrelevant degree of freedom, a second issue remains: the target eigenspace itself may fail to be unique. If P has k or more zero-incoming vertices, they contribute k null directions in addition to the constant null mode $\mathbf{1}$, so L^π has at least $k+1$ zero eigenvalues (Prop 1); then $\lambda_k = \lambda_{k+1} = 0$ and no unique bottom-k subspace (hence no canonical radius, absent a subspace-selection rule) exists at any (t, α).

These observations set the structure of our analysis:

remove irrelevant scale $\longrightarrow$ check identifiability $\longrightarrow$ quantify stability radius $\longrightarrow$ compose local certificates

We quotient out scale through the projective gauge-fixed operator $L^\pi_\vartheta := L_{\pi(\nu_\vartheta)}$ with $\pi(\nu) = \nu/\mathbf{1}^\top\nu$; test identifiability with the linear-time structural condition of Prop 1; quantify how much operator motion the partition tolerates with the partition-preservation theorem (Thm 1) and the three-factor stability radius it yields (Cor 1); and compose verified local steps into a reachable region (Lem 1).

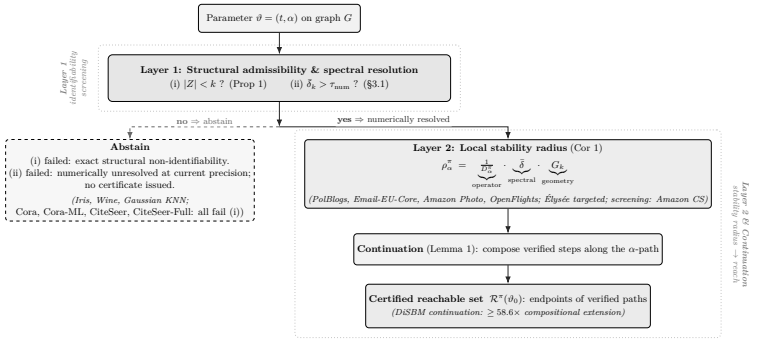

Figure 1: The four-step framework. Layer 1 screens identifiability with two gates: the structural gate $|Z| < k$ (its failure is a linear-time sufficient condition for exact non-identifiability, Prop 1) and the numerical resolution gate $\bar{\delta}_k > \tau_{\text{num}}$; graphs passing both are numerically resolved at the tested anchors, not proved exactly identifiable. On screened-in graphs the projective radius $\rho^\pi_{\alpha,\text{oracle}}$ decomposes into operator, spectral, and geometric factors, and continuation composes local certificates into a reachable set $\mathcal{R}^\pi(\vartheta_0)$ (§2).

Contributions

A projective stability theory with a three-factor radius. We formulate the sensitivity theory on the projective quotient of the nonnegative measure cone. A local partition-preservation theorem on L^π (Thm 1) yields the projective first-order local stability radius $\rho^\pi_{\alpha,\text{oracle}} = \bar{\delta} \cdot G_k/D^\pi_\alpha$ (Cor 1), which decomposes into operator response, spectral isolation, and embedding geometry.

Identifiability before perturbation. A zero-incoming null-mode proposition (Prop 1) gives a linear-time structural sufficient condition, $|Z| \geq k$, under which the bottom-k spectral subspace, hence its projector, is not uniquely defined, so no canonical certificate exists. Combined with a numerical spectral-resolution test, this layer runs before any perturbation analysis.

Practical consequences. Clustering quality and the stability radius are distinct, graph-dependently related objectives; and directly verified local certificates compose through continuation (Lem 1) into reachable regions more than 58× larger ($\geq$ 58.6× on the DiSBM continuation). The mechanism decomposition and the continuation gain are reported quantitatively in §4.

Scope. The theorems certify a chosen reference partition, not an intrinsic target. They combine standard perturbation and rounding tools (Davis–Kahan, Procrustes, Voronoi-margin rounding) with GSC's projective structure behind an explicit identifiability-then-stability ordering. The experiments are a diagnostic on four Stage-2 graphs, a targeted case study, and one synthetic family, not a population-level generality claim.

2 Setup

Generalized spectral clustering. Let $G = (V, E, W)$ be a weighted directed graph with $N = |V|$ vertices, and let $k \in \{2, \ldots, N-1\}$ be the desired number of clusters. Write $d_{\text{out}}(i) = \sum_j W_{ij}$ for the out-degree and $P_{ij} = W_{ij}/d_{\text{out}}(i)$ for the row-stochastic transition matrix, with the self-loop convention $P_{i\cdot} = e_i^\top$ (i.e. $P_{ii} = 1$) for any dangling row $d_{\text{out}}(i) = 0$, so that P is row-stochastic on every row (Appendix A.11); all reported Z, q_t, and screening verdicts are computed from this P. Let $\mu = N^{-1}\mathbf{1}$. For a diffusion time $t \in \{1, 2, \ldots\}$ and exponent $\alpha > 0$, define the diffusion intermediate and the vertex measure

$$q_t \;=\; (P^t)^\top \mu, \qquad \nu_{t,\alpha} \;=\; q_t^{\odot\alpha}, \tag{1}$$

where $x^{\odot\alpha}$ denotes elementwise exponentiation. Following Sevi et al. (2026), the GSC Laplacian of a nonnegative measure $\nu \in \mathbb{R}^N_{\geq 0}$ ($\mathbf{1}^\top\nu > 0$), with incoming measure $\xi := P^\top\nu$, is

$$L_\nu \;=\; D_{\nu+\xi} - D_\nu P - P^\top D_\nu, \tag{2}$$

symmetric PSD and linear in ν, with $L_\nu\mathbf{1} = 0$ (used in Prop 1).

Three related partition maps appear in the sequel and should be kept distinct (Table 6 in App. E): fresh k-means $A_{\text{km}}(\vartheta)$ on the rows of U_ϑ (the standard pipeline, no theoretical guarantee); anchored rounding $A_{\text{anc}}(\vartheta; \vartheta_0)$ to the nearest reference centroid after optimal Procrustes alignment to U_{ϑ_0} (certified by Thm 1); and continuation rounding A_{cont}, anchored sequentially along a discrete path (certified by Lem 1). Where the rule is immaterial we write $A(\vartheta)$; the empirical comparison between fresh k-means and anchored rounding on the three graphs included in that comparison is reported in Appendix E.

Projective gauge-fixed operator. Throughout we write $\vartheta = (t, \alpha)$, $\nu_\vartheta := \nu_{t,\alpha}$, and $s_\vartheta := \mathbf{1}^\top\nu_\vartheta$ for the total mass. Because $L_{c\nu} = c\,L_\nu$ for every $c > 0$, the GSC partition depends on ν only through its ray $\{c\nu : c > 0\}$. The projective quotient is thus $\mathbb{R}^N_{\geq 0} \setminus \{0\}$ modulo positive scaling, and each ray meets the simplex $\Delta^{N-1} = \{\nu \geq 0 : \mathbf{1}^\top\nu = 1\}$ in exactly one point; we fix the scale by evaluating the operator at that simplex representative,

$$\boxed{L^\pi_\vartheta \;:=\; L_{\nu_\vartheta/s_\vartheta} \;=\; \tfrac{1}{s_\vartheta} L_\nu}. \tag{3}$$

L^π_ϑ has the same eigenspaces as the unnormalized L_ν (and hence the same embedding and partition whenever the bottom-k target is well-defined), but with the scale removed, so all sensitivity analysis is developed on L^π (raw L_ν appears only to make gauge motion visible); computing it is division by a scalar, and the map is well-defined on the projective cone, with differential claims taken on a fixed-support face where the theory extends continuously for $\alpha > 0$ (Appendix A.11).

Having fixed the operator coordinate system, we now define the reference partition and the geometric quantities needed to ask how far that partition can move. Fix $\vartheta_0 = (t_0, \alpha_0)$, a nonempty reference partition $A^\star = \{A^\star_1, \ldots, A^\star_k\}$ ($|A^\star_a| \geq 1$), and let U_0 be an orthonormal basis of the bottom-k eigenspace of $L^\pi_{\vartheta_0}$. Write $\hat{u}_i$ for row i of U_0. The centroids of $A^\star$ under this basis are

$$c_a \;:=\; \tfrac{1}{|A^\star_a|} \sum_{i \in A^\star_a} \hat{u}_i,$$

and the Voronoi (embedding) margin at ϑ_0 is

$$\gamma_{\vartheta_0} := \min_i \min_{b \neq a^\star(i)} \big(\|\hat{u}_i - c_b\|_2 - \|\hat{u}_i - c_{a^\star(i)}\|_2 \big).$$

Three parameter sets recur: the anchored stability set $\Sigma_{\mathrm{anc}}(A^\star; \vartheta_0)$ on which anchored rounding returns $A^\star$; the continuation-reachable set $\mathcal{R}^\pi(\vartheta_0)$ of endpoints joined to ϑ_0 by a finite path of verified steps (Lem 1); and, when labels are available, the near-optimal quality region Q_q ($q = 0.95$). Their set-builder definitions and the path-independence of $\mathcal{R}^\pi$ are in Appendix A.1.

These definitions reduce the local stability problem to four scalars, all evaluated on L^π, that recur throughout the theory:

- Projective spectral gap $\bar{\delta}_\vartheta := \lambda_{k+1}(L^\pi_\vartheta) - \lambda_k(L^\pi_\vartheta)$, gauge-invariant; the gap above the k-th eigenvalue, where k is the fixed cluster count, suppressed in the subscript (written $\bar{\delta}_k$ where the k-dependence is emphasised).
- Voronoi margin γ_ϑ (above; schematic in Fig 3).
- Projective operator response $D^\pi_\alpha(\vartheta) := \|\partial_\alpha L^\pi_\vartheta\|_2$, gauge-invariant, and controlled by an analytic upper bound Ψ^π_α (§3).
- Geometric factor $G_k(\vartheta) := \gamma_\vartheta / (2\sqrt{2} + \gamma_\vartheta)$.

We denote $\vartheta = (t, \alpha)$ throughout; where notation is used only once it is defined in place.

3 Projective Local Stability Theory

A parameter change $\vartheta \to \vartheta'$ affects the GSC partition through one causal chain,

$$\underbrace{\Delta\nu}_{\text{measure}} \longrightarrow \underbrace{\Delta L^\pi}_{\text{operator}} \longrightarrow \underbrace{\Delta U}_{\text{embedding}} \longrightarrow \underbrace{\Delta A}_{\text{partition}},$$

and everything in this section controls one of these arrows.

Table 1 is the reader's map: it fixes, in one place, both the vocabulary and the epistemic status of each object before it is derived. Several of these objects are mathematically close but different in what they license (identifiable $\neq$ resolved; $\rho^\pi_{\alpha,\mathrm{bound}} \leq \rho^\pi_{\alpha,\mathrm{oracle}}$ are scales, not a certificate; anchored $\neq$ fresh rounding); the map keeps them apart, and the rest of the section derives the results in causal order.

Table 1: Reader map. The recurring terms and results, what each establishes, and the neighbouring notion it must not be confused with. Continuation (Lem. 1) chains the $h^\star_{\mathrm{cert}}$ row; the β-axis certificate (App. F) is its exact analogue on an affine path.

Term / result	What it establishes	Not to be confused with
target status	exact: unique bottom-k target if $\lambda_k < \lambda_{k+1}$; Prop. 1 proves failure if $\lvert Z\rvert \geq k$	resolved (its numerical test)
resolved	numerical: $\bar{\delta}_k > \tau_{\mathrm{num}}$ at the current precision	stable
Thm. 1	small $\|\Delta L^\pi\|_2$ preserves the anchored partition	a theorem for fresh k-means A_{km}
$\rho^\pi_{\alpha,\mathrm{bound}} \leq \rho^\pi_{\alpha,\mathrm{oracle}}$	first-order stability scale (computable $\leq$ mechanistic)	a finite certificate
$h^\star_{\mathrm{cert}}$	a directly verified finite admissible step	the whole equivalence region

To quantify the measure's effect we follow its consequences along this chain, one arrow per subsection. A perturbation argument, however, only makes sense around a uniquely defined spectral target, so we begin by asking when the bottom-k eigenspace is identifiable at all.

3.1 Identifiability: when the target eigenspace is not uniquely determined

Proposition 1 (Structural non-identifiability from zero-incoming vertices). Assume P is row-stochastic. Let $Z := \{i : P_{\cdot,i} = 0\}$ be the zero-in-degree set of P. For every $t \geq 1$ and $\alpha > 0$, both $\mathbf{1} \in \ker L^{\pi}_{t,\alpha}$ and $e_i \in \ker L^{\pi}_{t,\alpha}$ for each $i \in Z$; hence, provided $|Z| < N$,

$$\dim \ker L^{\pi}_{t,\alpha} \;\geq\; |Z| + 1.$$

If $|Z| \geq k$ then $\lambda_k(L^{\pi}_{t,\alpha}) = \lambda_{k+1}(L^{\pi}_{t,\alpha}) = 0$ and $\bar{\delta}_k(\vartheta) = 0$ for every such ϑ.

Sketch. For each $i \in Z$ the i-th row and column of L_ν vanish, so $\ker L^{\pi} \supseteq \operatorname{span}\{\mathbf{1}\} \cup \{e_i\}_{i \in Z}$ has dimension $\geq |Z| + 1$; $|Z| \geq k$ then forces $\lambda_k = \lambda_{k+1} = 0$ (Appendix A.2). □

Interpretation: enough zero-incoming vertices create extra exact null directions, so there is no unique bottom-k object whose stability could be certified. This is a persistent zero-support mechanism noted by Sevi et al. (2026) (there patched by a small positive value): Prop 1 is a structural sufficient condition from P, not an exhaustive characterization of all zero-support configurations.

Prop 1 detects this degeneracy from P alone, without any eigencomputation. It has a numerical counterpart, because a computed positive eigengap is only meaningful if it exceeds the numerical uncertainty of the two eigenpairs: even when the exact gap is zero, a finite-precision eigensolver may return a small apparent positive gap. We therefore call the target numerically resolved only when the observed gap is comfortably larger than the backward-error scale of the two eigenpairs,

$$\bar{\delta}_k(\vartheta) \;>\; \tau_{\mathrm{num}}(\vartheta), \qquad \tau_{\mathrm{num}}(\vartheta) := 10 \cdot \max\big(r_k(\vartheta), r_{k+1}(\vartheta), \epsilon_{\mathrm{mach}} \|L^{\pi}_{\vartheta}\|_2\big), \tag{4}$$

where $r_k = \|L^{\pi} u_k - \lambda_k u_k\|_2$ is the eigenpair residual (near machine precision for an accurate dense solve, but not identically zero, and appreciably larger for iterative solvers), and $\epsilon_{\mathrm{mach}} \|L^{\pi}\|_2$ ($\epsilon_{\mathrm{mach}} = 2.22 \times 10^{-16}$) is the floating-point scale of the operator; the factor 10 is a conservative safety factor, not a theorem. Structural and numerical tests together form the identifiability layer: fail either and no perturbative certificate is issued (all reported verdicts use (4); the two dense Stage-2 grids pass at all 42 anchors by large margins and Amazon Photo at 41/42, Appendix B.2).

Suppose now the target passes these structural and numerical checks. The next question is how rapidly changing α moves the operator, and because raw operator motion contains an irrelevant scaling component, this derivative must be taken in the projective gauge.

3.2 The gauge decomposition of the operator derivative

Differentiating $L_\nu = s(\alpha)\, L^{\pi}$ along the α-path gives the gauge decomposition

$$\partial_\alpha L_{\nu_\alpha} \;=\; \underbrace{s'(\alpha)\, L^{\pi}_{\alpha}}_{\text{radial (pure scale)}} \;+\; \underbrace{s(\alpha)\, \partial_\alpha L^{\pi}_{\alpha}}_{\text{projective (potentially partition-relevant)}} . \tag{5}$$

The first term is a scalar multiple of L^{π}: it preserves eigenspaces and cannot rotate the isolated bottom-k subspace; only the second, projective term can. A sensitivity theory built on $\partial_\alpha L_\nu$ conflates the two, exactly the confusion we remove by working on L^{π} (formal statement in Appendix A.4). The exact response $D^{\pi}_{\alpha}(\vartheta) := \|\partial_\alpha L^{\pi}_{\vartheta}\|_2$ drives the mechanism analysis, while an analytic upper bound $\Psi^{\pi}_{\alpha}(\vartheta) \geq D^{\pi}_{\alpha}(\vartheta)$ (Appendix A.5) gives a computable, conservative surrogate; both are gauge-invariant, and we assess the bound's tightness empirically in §4.

The response D^{π}_{α} tells us how fast the operator moves, but not how much motion the clustering can tolerate; that tolerance is set jointly by the target's spectral isolation and the geometry of the reference embedding.

3.3 The partition-preservation theorem

Our goal is to keep the partition fixed. Under anchored rounding this is guaranteed if every aligned embedding row stays inside the Voronoi cell of its reference centroid; a sufficient way to ensure this is that each row moves by less than half the embedding margin γ. The certificate is therefore a chain from operator motion to unchanged labels,

$$\underbrace{\|L^\pi_{\vartheta'} - L^\pi_\vartheta\|_2}_{\text{operator motion}} \overset{\text{Davis–Kahan}}{\Longrightarrow} \underbrace{\|\tilde{U} - U_0 O^\star\|_2}_{\text{subspace motion}} \Longrightarrow \underbrace{\max_i \|\tilde{u}_i - (O^\star)^\top \hat{u}_i\|_2}_{\text{row motion}} \overset{<\gamma/2}{\Longrightarrow} \underbrace{A_{\text{anc}}(\vartheta';\vartheta) = A^\star}_{\text{same partition}}, \tag{6}$$

which we read right-to-left as a requirement. Writing $\epsilon = \|L^\pi_{\vartheta'} - L^\pi_\vartheta\|_2$, the operator-norm Davis–Kahan bound gives $\|\sin\Theta\|_2 \le \epsilon/(\bar{\delta} - \epsilon)$, the polar Procrustes factor $O^\star$ gives $\|\tilde{U} - U_0 O^\star\|_2 \le \sqrt{2}\,\|\sin\Theta\|_2$, and each row's displacement is at most the operator norm, $\max_i \|\tilde{u}_i - (O^\star)^\top \hat{u}_i\|_2 \le \|\tilde{U} - U_0 O^\star\|_2$. A sufficient condition for every row to stay in its Voronoi cell is therefore $\sqrt{2}\,\epsilon/(\bar{\delta} - \epsilon) < \gamma/2$, and solving for ϵ gives the local partition-preservation threshold

$$\tau_\vartheta \;:=\; \frac{\gamma_\vartheta\,\bar{\delta}_\vartheta}{2\sqrt{2} + \gamma_\vartheta}. \tag{7}$$

Thus τ_ϑ is not an arbitrary formula but the exact solution of this sufficient inequality, the operator-motion budget it certifies, free of an explicit $\sqrt{k}$ dimension penalty, and it is the canonical threshold used throughout the α-axis theory and the radius experiments (Appendix G separately reports a conservative Frobenius t-axis diagnostic). A cruder Frobenius route bounds the subspace motion by the dimension-dependent $\sqrt{2k}\,\epsilon/(\bar{\delta} - \epsilon)$, giving the weaker $\tau^{\text{F}}_\vartheta = \gamma\bar{\delta}/(2\sqrt{2k} + \gamma) \le \tau_\vartheta$, recorded in Appendix A.6 only as the dimension-dependent alternative. (γ, and hence τ_ϑ and $\rho^\pi_{\alpha,\text{oracle}}$, is invariant to the orthonormal basis of the unique bottom-k subspace; Appendix A.6.)

Theorem 1 (Local partition-preservation on L^π). Let ϑ be a reference point with partition $A^\star$, bottom-k embedding U_ϑ, centroids $\{c_a\}$, and Voronoi margin γ_ϑ as defined in §2. Assume $\bar{\delta}_\vartheta > 0$ and $\gamma_\vartheta > 0$. Then, for any ϑ' with

$$\|L^\pi_{\vartheta'} - L^\pi_\vartheta\|_2 \;<\; \tau_\vartheta \;=\; \frac{\gamma_\vartheta\,\bar{\delta}_\vartheta}{2\sqrt{2} + \gamma_\vartheta},$$

anchored rounding of the bottom-k eigenvectors of $L^\pi_{\vartheta'}$ against $(U_\vartheta, \{c_a\})$ preserves the partition: $A_{\text{anc}}(\vartheta';\vartheta) = A^\star$. No claim is made about the fresh k-means partition $A_{\text{km}}(\vartheta')$.

Interpretation: if operator motion is small relative to the spectral separation and the embedding margin, no anchored cluster assignment can change. The threshold τ_ϑ is exactly the derivation above; the full proof and the dimension-dependent Frobenius variant τ^{F} are in Appendix A.6.

The theorem certifies anchored rounding; agreement with fresh k-means is examined empirically in Appendix E.

3.4 The projective local stability radius and its factorisation

We now have the two ingredients for a parameter-space stability scale: a tolerance τ_ϑ in operator space (§3.3) and the local speed D^π_α at which α traverses that space (§3.2). Their ratio is the local stability radius.

Corollary 1 (Projective first-order local stability radius). At ϑ with $\bar{\delta}_\vartheta > 0$, $\gamma_\vartheta > 0$, and $D^\pi_\alpha(\vartheta) > 0$, the oracle first-order local stability radius on the α-axis (the linearised scale at which the leading-order operator motion crosses τ_ϑ) is

$$\boxed{\rho^\pi_{\alpha,\text{oracle}}(\vartheta) \;:=\; \frac{\tau_\vartheta}{D^\pi_\alpha(\vartheta)} \;=\; \underbrace{\frac{1}{D^\pi_\alpha(\vartheta)}}_{\text{operator}} \cdot \underbrace{\bar{\delta}_\vartheta}_{\text{spectral}} \cdot \underbrace{G_k(\vartheta)}_{\text{geometry}}}, \tag{8}$$

with $G_k(\vartheta) := \gamma_\vartheta/(2\sqrt{2} + \gamma_\vartheta)$. The computable analytic surrogate is the bound radius $\rho^\pi_{\alpha,\mathrm{bound}}(\vartheta) := \tau_\vartheta/\Psi^\pi_\alpha(\vartheta)$, and $D^\pi_\alpha \le \Psi^\pi_\alpha$ gives $\rho^\pi_{\alpha,\mathrm{bound}} \le \rho^\pi_{\alpha,\mathrm{oracle}}$ (Appendix A.5). This first-order object is a diagnostic, not a finite-step guarantee: the finite certificate is a directly verified step $\|L^\pi_{\alpha+h} - L^\pi_\alpha\|_2 < \eta\,\tau_\vartheta$ (§3.4, below).

Proof. Linearising $L^\pi_{\alpha+h}$ gives operator motion $|h|\,D^\pi_\alpha(\vartheta) + o(|h|)$, which crosses τ_ϑ at $\rho^\pi_{\alpha,\mathrm{oracle}} = \tau_\vartheta/D^\pi_\alpha$; substituting (7) yields the factorisation (degenerate cases in Appendix A.5). □

Interpretation: tuning tolerance is largest (the partition least sensitive to measure changes) when the operator moves slowly (D^π_α small), the target eigenspace is spectrally isolated ($\bar\delta$ large), and the clusters are well separated (G_k large).

These three quantities play complementary roles (the $\rho^\pi_{\alpha,\mathrm{bound}}$, $\rho^\pi_{\alpha,\mathrm{oracle}}$, $h^\star_{\mathrm{cert}}$ rows of Table 1): $\rho^\pi_{\alpha,\mathrm{bound}} \le \rho^\pi_{\alpha,\mathrm{oracle}}$ are, respectively, the a priori scale used for anchor selection (it needs only Ψ^π_α, not the true derivative) and the mechanistic diagnostic that factorises (8), while the finite certificate is a directly verified step $h^\star_{\mathrm{cert}}$: any h whose operator motion stays within a fraction $\eta \in (0,1)$ of the threshold, $\|L^\pi_{\alpha+h} - L^\pi_\alpha\|_2 < \eta\,\tau_\vartheta$. Since the admissible set is open we report $h^\star_{\mathrm{cert}}$ as the largest feasible step located by the bracket search (each point individually verified), not a global supremum; numerical details, non-monotone brackets, and the chaining by continuation (§3.5) are in Appendix A.3.

3.5 Continuation

This radius is local: a single reference point may certify only a short interval even when the partition stays stable much farther away. To extend the guarantee, we update the reference after each verified step and chain the local certificates.

Lemma 1 (Finite-step continuation). Consider a discrete α-path $\vartheta_0 \to \cdots \to \vartheta_m$ (fixed t, each $\alpha_j > 0$) with reference partition $A^\star$. If every step $\vartheta_j \to \vartheta_{j+1}$ satisfies Thm 1 using the embedding and centroids at ϑ_j as reference, that is, $\bar\delta_{\vartheta_j}, \gamma_{\vartheta_j} > 0$ and $\|L^\pi_{\vartheta_{j+1}} - L^\pi_{\vartheta_j}\|_2 < \eta\,\tau_{\vartheta_j}$ for some $\eta \in (0,1]$, then continuation rounding preserves $A^\star$ along the entire path, $A_{\mathrm{cont}}(\vartheta_{j+1} \mid \vartheta_0, \ldots, \vartheta_{j+1}) = A^\star$ for all j, and every $\vartheta_j \in \mathcal{R}^\pi(\vartheta_0)$.

The result follows by induction on Thm 1 (Appendix A.7), turning local certificates into a certified reachable region whose reach depends on the evolution of the local threshold τ_ϑ and the finite operator motion along the path; locally, D^π_α diagnoses the latter. These are sufficient-condition certificates, not calibrated empirical boundaries: an unresolved anchor withholds the certificate without concluding exact non-identifiability (unless Prop 1 fires).

4 Experiments

We evaluate the framework by asking four questions in the order the theory asks them: does removing gauge motion matter (RQ1); when is the requested stability problem identifiable (RQ2); when it is, how large is the radius and what limits it (RQ3); and do local certificates compose (RQ4)?

We use a two-stage experimental protocol. Stage 1 (identifiability): if $|Z| \ge k$ abstain by Prop 1; otherwise check (4) at test anchors and abstain if it fails. Stage 2 (Stage-1-resolved graphs only): a 42-point (t, α) sweep ($t \in \{1, 2, 3, 5, 8, 12\}$, $\alpha \in \{0.1, 0.2, 0.3, 0.5, 0.7, 1.0, 1.3\}$) computing $\bar\delta, \gamma, G_k, \Psi^\pi_\alpha, D^\pi_\alpha, \rho^\pi_\alpha$, AMI by dense eigh on $\frac{1}{2}(L^\pi + (L^\pi)^\top)$ (Appendix B.6 on sparse-solver pitfalls). The reference partition is the standard GSC k-means output at ϑ_0 (not ground truth), and the radius certifies its preservation under anchored rounding; ground truth enters only through the AMI quality gate, never as the estimand. The suite: four full Stage-2 graphs (PolBlogs, Email-EU-Core, Amazon Photo at $N = 6105$ via a validated sparse eigensolver, OpenFlights), Élysée2017fr as a targeted case study (Appendix B.8), a further Stage-1-resolved graph (Amazon CS), seven structurally

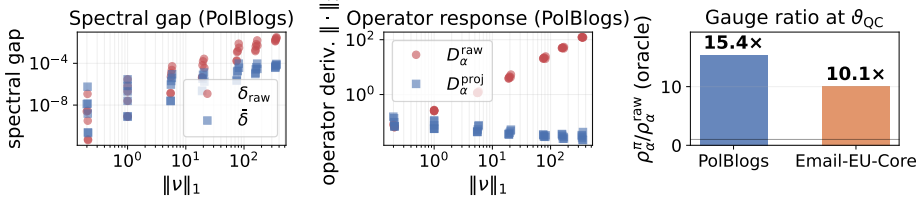

Figure 2: Effect of projective normalization on the stability radius. Left, middle: on PolBlogs, raw $\delta^{\text{raw}}, D^{\text{raw}}$ track $\|\nu\|_1$ across the grid while the projective $\bar{\delta}, D_\alpha^\pi$ stay stable; the raw motion is gauge, not perturbation. Right: same-anchor oracle ratio $\rho^\pi_{\alpha,\text{oracle}}/\rho^{\text{raw}}_{\alpha,\text{oracle}}$ at ϑ_{QC} on the two dense Stage-2 graphs.

inadmissible benchmarks ($|Z| \geq k$), a synthetic DiSBM continuation case, and two UCI intervention-panel datasets; per-graph statistics in Appendix B.2.

4.1 (RQ1) Does removing gauge motion matter?

Removing gauge motion changes the conclusion. On the two dense Stage-2 graphs, the raw and projective first-order radii disagree by an order of magnitude at the same quality-conditioned anchor $\vartheta_{\text{QC}} = \arg\max_{\vartheta \in Q_{0.95}} \rho^\pi_{\alpha,\text{bound}}(\vartheta)$: the same-anchor oracle-radius ratio $\rho^\pi_{\alpha,\text{oracle}}/\rho^{\text{raw}}_{\alpha,\text{oracle}}$ (the radius computed on L^π vs. on the raw L_ν) is **15.4**× (PolBlogs) and **10.1**× (Email-EU-Core). The raw total mass s_ϑ varies by more than 1800× across the grid; by the gauge decomposition (5) this is scale motion that cannot rotate the eigenspace, yet raw perturbation analysis counts it as sensitivity. Removing gauge can even change which parameter is selected: on PolBlogs the raw anchor is $(1, 0.20)$ and the projective anchor $(2, 0.10)$, both in $Q_{0.95}$, with the projective choice giving a 1.35× higher bound radius $\rho^\pi_{\alpha,\text{bound}}$ (Fig 2). Removing gauge presumes a well-defined target, though, which is not automatic; we ask next when the bottom-k target is identifiable at all.

4.2 (RQ2) When is the stability problem identifiable?

Identifiability fails structurally on seven of the eleven stated graph views, detected in linear time before any eigendecomposition: they satisfy $|Z| \geq k$ and are ruled out by Prop 1 (Appendix B.2), with dense eigendecomposition confirming $\bar{\delta}_k$ at the backward-error scale on every abstained graph. The zero-in-degree vertices are points that are no one's K-nearest neighbour in directed k-NN graphs ($|Z| = 5, 7, 18$ vs. $k = 3$) and uncited papers in citation graphs ($|Z| \in [407, 1168]$, up to 67.5% of N). Abstention is the theorem's verdict that no unique bottom-k target exists, not a failure to cluster: fresh k-means can still score nonzero AMI, but on a basis of a non-unique subspace, and representative spectra show the bottom $k{+}1$ eigenvalues collapsing exactly as Prop 1 predicts (Appendices B.9, B.5, B.6, C, Fig 5). On graphs that pass this screen, we quantify the radius itself.

4.3 (RQ3) What mechanism controls the radius, and how large is it?

The radius is controlled by different mixtures of the operator, spectral, and geometric factors: the per-seed decomposition attributes each intervention's effect to the factors that actually move.

We read the three-factor identity $\Delta \log \rho^\pi_{\alpha,\text{oracle}} = -\Delta \log D^\pi_\alpha + \Delta \log \bar{\delta} + \Delta \log G_k$ on a controlled DC-SBM generator ($N = 400$, $k = 4$, 30 seeds, anchor $(1, 0.10)$) with three levers each nominally aligned to one factor: node-measure heterogeneity $\to D^\pi_\alpha$, community contrast $\to \bar{\delta}$, and block-density noise $\to \gamma$ (Appendix J). The levers are not isolating (each moves more than one factor), so this is a mechanism accounting test: which factors explain each intervention's change in $\rho^\pi_{\alpha,\text{oracle}}$? The reference partition is the planted labels.

The factorization (8) is algebraic, so its per-seed closure to machine precision (residual $\sim 10^{-16}$ across all 30 seeds) is a computational sanity check, not evidence for the theory;

the substantive result is which factor dominates each intervention. The operator sweep is dominated by its nominal factor (D_α^π; 30/30 surviving seeds), but the spectral sweep (nominally $\bar{\delta}$) is dominated by the geometry factor G_k (17/30), as is the geometry sweep (22/30): G_k is the most responsive lever, and $\bar{\delta}$ is never the single dominant factor, so the interventions are not factor-isolating. The verified finite step $h^\star_{\mathrm{cert}}$ tracks $\eta\rho^\pi_{\alpha,\mathrm{oracle}}$ to within $\sim 1\%$ in $\Delta\log$ on the tested brackets, confirming the first-order radius predicts it (full decompositions, the $\gamma \le 0$ dropout, and planted-vs-fresh agreement in Appendix J).

On the two dense Stage-2 graphs (42-point grid, (8)), spectral conditioning $\bar{\delta}$ is the dominant single-factor lever, and the headroom jointly attainable within $Q_{0.95}$ is modest (1.30× PolBlogs, 1.16× Email-EU-Core): the quality-conditioned anchors are nearly Pareto-cornered (Appendix K). In the limit, this dominance means separation controls measure-sensitivity: a controlled sweep (Appendix A.10) shows the partition becoming measure-invariant across the grid once the target is sufficiently separated, with D_α^π nearly flat.

A second distinction: clustering quality and stability need not rank settings similarly. The AMI–$\rho^\pi_{\alpha,\mathrm{bound}}$ Spearman correlation is -0.06 (PolBlogs), -0.77 (Email-EU-Core), -0.61 (Amazon Photo, 41 resolved anchors), and -0.64 (OpenFlights), so we maximise stability only within $Q_{0.95}$ rather than trading one for the other (Fig 7, Appendix K). Everything so far certifies a single anchor; we close by asking whether local certificates compose into a larger region.

4.4 (RQ4) Do local certificates compose?

Directly verified local certificates compose into a much larger reachable region ($\ge 58.6\times$ before the step cap).

On chain DiSBM ($\rho = 0.05$, $N = 600$, $k = 3$) we run adaptive continuation from $(1, 0.1)$, fixing the reference partition at ϑ_0 and certifying each step by direct operator evaluation. A single directly verified single-reference step reaches $h_0 \approx 6.35 \times 10^{-3}$; 80 accepted steps reach $\Delta\alpha_+ = 0.3725$, at least $\mathbf{58.6}\times$ that step, terminating on the step cap rather than certificate failure. At every accepted step all three rounding rules agree and AMI against ground truth stays 1.0. The reach is near-linear because $\bar{\delta}, \gamma$ stay bounded from zero and Ψ^π_α bounded above along the path (an observation, not a theorem; Appendix D).

5 Discussion

The experiments trace the theory's dependency chain end to end; for a new graph the recipe is three steps: (i) check identifiability (abstain if $|Z| \ge k$, Prop 1, or if the gap fails (4)); (ii) compute $\rho^\pi_{\alpha,\mathrm{oracle}} = \bar{\delta}\, G_k / D^\pi_\alpha$ and read its factors to see which limits tuning; (iii) compose verified steps by continuation for more reach. Whether this stability picture transfers to other parameterised operator families is open.

Limitations and scope. The guarantees are sufficient, not necessary, and Thm 1 certifies the anchored partition, not fresh k-means. Since ϑ_{QC} maximises $\rho^\pi_{\alpha,\mathrm{bound}}$ over the ground-truth-AMI region, it is a retrospective oracle, not a deployable unsupervised rule, and the Spearman summaries are descriptive. Full Stage-2 covers four real graphs across four edge mechanisms (N up to 6105) plus Élysée2017fr (Appendix B.8); broader coverage is future work. Main-text results are on the α-axis (Appendix G).

6 Related Work

Spectral clustering partitions graphs via Laplacian eigenvectors (Shi & Malik, 2000; Ng et al., 2001; Von Luxburg, 2007); Sevi et al. (2026) generalise it with a tunable vertex measure ν (family $\nu_{t,\alpha}$). Prior stability work perturbs the graph (Andreotti et al., 2021; Peng & Yoshida, 2020) or bounds eigenspace/rounding error (Davis & Kahan, 1970; Yu et al., 2015; Boedihardjo et al., 2021); we instead perturb the measure, with sensitivity in operator/spectral/geometric factors.

Ethics Statement

This is a methodological study of the stability of a clustering operator; it involves no human subjects, no new data collection, and no user interaction. All datasets are pre-existing and publicly available under their stated licenses, used only as fixed graph structures. One benchmark, the Élysée retweet network (Zenodo record 5535333, CC-BY-4.0), derives from public social-media activity. We use it solely as an assortative network testbed: we reindex vertices to consecutive integers and retain only the anonymized edge structure, discarding raw account identifiers, handles, message content, and all user attributes, so no individual is identifiable from the shipped data. The community labels are coarse candidate-support blocks used only to score partition quality; we make no individual-level inference or prediction about any account, and the method outputs group structure, not claims about people. The remaining benchmarks (OpenFlights routes under ODbL; PolBlogs; Email-EU-Core with pre-anonymized identifiers; Amazon co-purchase; citation networks; and the Wikipedia-derived Squirrel/Chameleon graphs) are standard public benchmarks with no personal content. We see no additional societal-harm or dual-use concerns specific to this work beyond those general to unsupervised graph clustering, and we comply with the ICLR Code of Ethics.

Reproducibility Statement

Complete proofs of all formal claims (Proposition 1, Theorem 1, Corollary 1, and Lemma 1) are given in Appendix A.6 and the surrounding subsections. The Stage-2 estimand and numerical configuration (Appendix B.1), benchmark provenance and Stage-1 verdicts (Appendix B.2), sparse-solver validation (Appendix B.7), and the seven-stage hard-gated harness (Appendix B) are documented in full. The code, datasets, per-anchor CSVs, a pinned environment, and a checksum manifest (MANIFEST.sha256), together with a single command (reproduce.sh) that re-runs the experiments and asserts the reported verdicts, bounds, and invariants end to end, will be released with the camera-ready version.

Use of Large Language Models

In accordance with the ICLR 2027 policy on AI assistance, we disclose the following. The research questions, the projective gauge-fixing framework, the identifiability result (Proposition 1), the partition-preservation theorem (Theorem 1) and its corollaries, the continuation lemma (Lemma 1), and the overall proof strategies were conceived by the authors, who verified every mathematical claim and take full responsibility for the 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 clarity and hierarchy; drafting prose, figure/table captions (including the Figure 1 schematic, the reader map Table 1, and the terminology summary), and, from author-specified strategies, the write-up of the proofs; and consistency-checking cross-references, notation, and claims across the manuscript.
- Research execution. Debugging the experimental and verification code, running experiments, and assembling the one-command reproduction harness and checksum manifest (Appendix B) so that reported quantities are reproducible.
- Verification and repair. Checking derivations and cross-checking reported numbers against the code and datasets, which led us to report backward-error-sensitive quantities as configuration-independent bounds rather than exact digits.

The LLM was not used to originate the research ideas, to generate data, or to produce any reported result that the authors did not independently verify. All theoretical claims, proofs, experimental designs, reported numerical results, and final manuscript content were independently checked and approved by the authors, who take full responsibility for the work.

References

Lada A. Adamic and Natalie Glance. The political blogosphere and the 2004 US election: divided they blog. In Proceedings of the 3rd International Workshop on Link Discovery, pp. 36–43, 2005.

Edoardo Andreotti, Dominik Edelmann, Nicola Guglielmi, and Christian Lubich. Measuring the stability of spectral clustering. Linear Algebra and its Applications, 610:673–697, 2021.

March Boedihardjo, Shaofeng Deng, and Thomas Strohmer. A performance guarantee for spectral clustering. SIAM Journal on Mathematics of Data Science, 3(1):369–387, 2021.

Aleksandar Bojchevski and Stephan Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. In International Conference on Learning Representations (ICLR), 2018.

Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. III. SIAM Journal on Numerical Analysis, 7(1):1–46, 1970.

Dheeru Dua and Casey Graff. UCI machine learning repository. http://archive.ics.uci.edu/ml, 2019.

Ophélie Fraisier, Guillaume Cabanac, Yoann Pitarch, Romaric Besançon, and Mohand Boughanem. #élysée2017fr dataset. Zenodo, 2018a. CC-BY-4.0.

Ophélie Fraisier, Guillaume Cabanac, Yoann Pitarch, Romaric Besançon, and Mohand Boughanem. #élysée2017fr: The 2017 French presidential campaign on Twitter. In Proceedings of the International AAAI Conference on Web and Social Media (ICWSM), volume 12, 2018b.

Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. Advances in Neural Information Processing Systems, 14, 2001.

Jouko Patokallio. OpenFlights: Airport, airline and route database. https://openflights.org/data.html, 2014. Open Database License (ODbL).

Pan Peng and Yuichi Yoshida. Average sensitivity of spectral clustering. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), pp. 1132–1140, 2020.

Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In International Conference on Learning Representations (ICLR), 2023.

Benedek Rozemberczki, Carl Allen, and Rik Sarkar. Multi-scale attributed node embedding. Journal of Complex Networks, 9(2), 2021.

Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. AI Magazine, 29(3):93–106, 2008.

Harry Sevi, Gwendal Debaussart-Joniec, Malik Hacini, Matthieu Jonckheere, and Argyris Kalogeratos. Generalized Dirichlet energy and graph Laplacians for clustering directed and undirected graphs. Transactions on Machine Learning Research, 2026. URL https://openreview.net/forum?id=AA6D7fJ9PN. arXiv:2203.03221.

Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph neural network evaluation. In Relational Representation Learning Workshop, NeurIPS, 2018.

Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(8):888–905, 2000.

Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. Journal of Machine Learning Research, 11:2837–2854, 2010.

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

Hao Yin, Austin R. Benson, Jure Leskovec, and David F. Gleich. Local higher-order graph clustering. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 555–564, 2017.

Yi Yu, Tengyao Wang, and Richard J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. Biometrika, 102(2):315–323, 2015.

A Extended derivations of the theory

Roadmap. This appendix supplies the formal details behind §3. Appendix A.1 formalizes the stability sets; A.2 proves the structural identifiability result; A.3 gives the finite-step numerical procedure; A.4–A.5 derive the gauge decomposition and the operator-response bound; A.6 proves the partition-preservation theorem; A.7 proves continuation; A.8–A.10 collect secondary consequences (the discrete t-axis certificate, perfect separation, and a controlled separation sweep); and A.11 extends the construction to nonnegative measures on fixed support.

A.1 Stability sets and quality region (formal)

The two anchor-relative stability sets of §2 are

$$\Sigma_{\mathrm{anc}}(A^\star;\vartheta_0) := \{\vartheta : A_{\mathrm{anc}}(\vartheta;\vartheta_0) = A^\star\},$$
$$\mathcal{R}^\pi(\vartheta_0) := \{\vartheta_m : \exists\, \vartheta_0 \to \cdots \to \vartheta_m \text{ finite, with every step verified (Lem 1)}\}.$$

Σ_{anc} is a fixed-reference anchored stability set (a fibre over $(\vartheta_0, A^\star)$), a function of the fixed reference data $(\vartheta_0, A^\star)$, not of ϑ_0 alone. $\mathcal{R}^\pi$ is the union over all verified paths: membership $\vartheta_m \in \mathcal{R}^\pi$ is witnessed by the existence of some verified path from ϑ_0 to ϑ_m, so as a set it depends only on $(\vartheta_0, A^\star)$; a particular realized continuation endpoint, by contrast, is history-dependent through the specific path and its accepted step sizes. We suppress the $A^\star$ argument in $\mathcal{R}^\pi(\vartheta_0)$ where the reference partition is understood, and write $\Sigma(A^\star)$ for $\Sigma_{\mathrm{anc}}(A^\star;\vartheta_0)$ with the anchor understood from context. Here $\mathrm{AMI}(\vartheta)$ is the adjusted mutual information (Vinh et al., 2010) between $A_{\mathrm{km}}(\vartheta)$ and the ground-truth labels, a chance-corrected agreement score with AMI $= 1$ for identical partitions and ≈ 0 for independent ones, used purely as an external quality measure, never as the estimand. The near-optimal quality region is $Q_q := \{\vartheta : \mathrm{AMI}(\vartheta) \geq q \cdot \mathrm{AMI}_{\max}\}$ with $q = 0.95$ throughout. All reported Stage-2 graphs satisfy $\mathrm{AMI}_{\max} > 0$ (AMI can be negative in general); Q_q is defined under this condition.

A.2 Proof of Proposition 1

Row-stochasticity gives $L_\nu \mathbf{1} = D_{\nu+\xi}\mathbf{1} - D_\nu P\mathbf{1} - P^\top D_\nu \mathbf{1} = (\nu+\xi) - \nu - \xi = 0$, so $\mathbf{1} \in \ker L_\nu$. For $i \in Z$, a zero i-th column of P forces a zero i-th column of P^t for every $t \geq 1$, so $q_t(i) = ((P^t)^\top \mu)(i) = \sum_j (P^t)_{ji}\mu_j = 0$; hence $\nu_{t,\alpha}(i) = q_t(i)^\alpha = 0$ and $\xi(i) = (P^\top \nu)(i) = 0$. Row i of $L_\nu = D_{\nu+\xi} - D_\nu P - P^\top D_\nu$ then vanishes and, by symmetry, column i vanishes, so $L_\nu e_i = 0$. When $|Z| < N$, the vector $\mathbf{1}$ is independent of $\{e_i : i \in Z\}$, so $\ker L_\nu \supseteq \mathrm{span}\{\mathbf{1}\} \cup \{e_i : i \in Z\}$ has dimension $\geq |Z| + 1$. Since $L^\pi = L_\nu / s_\vartheta$ with $s_\vartheta > 0$, the same holds for $\ker L^\pi$; $|Z| \geq k$ then forces the bottom $k+1$ eigenvalues to be 0. □

A.3 Computing the verified step $h^\star_{\mathrm{cert}}$

The admissible set is $H_\vartheta(\eta) := \{h \in \mathbb{R} : \alpha + h > 0,\ \|L^\pi_{\alpha+h} - L^\pi_\alpha\|_2 < \eta\, \tau_\vartheta\}$ (the $\alpha + h > 0$ constraint keeps the step inside the theory's domain). It is open by continuity of L^π in α, so

its supremum $h^{\sup}_{\mathrm{cert}}$ is generally not attained. We therefore report $\widehat{h}_{\mathrm{cert}}(\vartheta)$, the largest feasible point of $H_\vartheta(\eta) \cap [0, h_{\max}]$ located by bisection to a stated absolute tolerance $\epsilon_h = 10^{-6}$ (safety factor $\eta \in (0,1)$); bisection returns the first crossing only under the assumption that $h \mapsto \|L^\pi_{\alpha+h} - L^\pi_\alpha\|_2$ is monotone in $|h|$ on the search bracket, so $\widehat{h}_{\mathrm{cert}}$ is a feasible lower endpoint, not a proven global maximum, and each returned point is verified directly by evaluating the operator inequality. On brackets where monotonicity is not verified, the integral bound $\int_0^{|h|} \Psi^\pi_\alpha(\alpha + \mathrm{sgn}(h)s)\, ds < \eta\, \tau_\vartheta$ also certifies the step (it retains the same safety factor η) and is manifestly monotone in $|h|$.

A.4 Gauge decomposition (formal statement and proof)

Proposition A.1 (Gauge decomposition). For any smooth measure path ν_α with fixed support and $\pi(\nu_\alpha)$ smooth in $\alpha > 0$,

$$\partial_\alpha L_{\nu_\alpha} \;=\; s'(\alpha)\, L^\pi_\alpha \;+\; s(\alpha)\, \partial_\alpha L^\pi_\alpha.$$

The first term is a scalar multiple of L^π_α and thus commutes with L^π_α; it preserves the eigenspaces and, when the bottom-k subspace of L^π_α is spectrally isolated ($\bar{\delta}_k > 0$), cannot rotate that subspace. The second term is the only component that can rotate the eigenspace under that isolation assumption.

Proof. Differentiate $L_{\nu_\alpha} = s(\alpha)\, L^\pi_\alpha$ by the product rule. The commutation $s'(\alpha) L^\pi_\alpha \cdot L^\pi_\alpha = s'(\alpha)(L^\pi_\alpha)^2$ shows the first term preserves eigenspaces. □

The unnormalized map $\nu \mapsto L_\nu$ is linear in ν on the nonnegative cone; the projective map $\nu \mapsto L^\pi_\nu$ is degree-zero homogeneous (invariant under $\nu \mapsto c\nu$, $c > 0$). The simplex $\{\nu : \mathbf{1}^\top \nu = 1\}$ is an affine slice of the cone, not a linear subspace, so "linear on the simplex" is a shorthand for the restriction of the ambient linear map $\nu \mapsto L_\nu$ to that slice; equivalently, along the simplex, $L^\pi_{\lambda\nu_1 + (1-\lambda)\nu_2} = \lambda L^\pi_{\nu_1} + (1-\lambda) L^\pi_{\nu_2}$ for $\lambda \in [0,1]$, i.e. L^π is affine under convex mixtures of simplex points. Sensitivity claims in the main text always refer to α-derivatives at fixed support on this simplex section.

A.5 Direction-aware operator-perturbation bound

Proposition A.2 (Direction-aware operator bound). For any nonnegative measures $\nu, \nu + \Delta\nu$ with $\nu + \Delta\nu \geq 0$,

$$\|L_{\nu+\Delta\nu} - L_\nu\|_2 \;\leq\; \big(1 + r_\xi(\Delta\nu) + 2\|P\|_2\big)\, \|\Delta\nu\|_\infty, \tag{9}$$

where $r_\xi(\Delta\nu) := \|P^\top \Delta\nu\|_\infty / \|\Delta\nu\|_\infty$ for $\Delta\nu \neq 0$, and $r_\xi(0) := 0$; $r_\xi \leq \kappa_c := \|P\|_1$.

Proof. By linearity of $\nu \mapsto L_\nu$ and with $\Delta\xi := P^\top \Delta\nu$, $L_{\nu+\Delta\nu} - L_\nu = D_{\Delta\nu+\Delta\xi} - D_{\Delta\nu} P - P^\top D_{\Delta\nu}$. The diagonal term $\|D_{\Delta\nu+\Delta\xi}\|_2 = \|\Delta\nu + \Delta\xi\|_\infty \leq (1 + r_\xi(\Delta\nu))\|\Delta\nu\|_\infty$. The two off-diagonal terms each satisfy $\|D_{\Delta\nu} P\|_2 \leq \|\Delta\nu\|_\infty \|P\|_2$. Triangle inequality gives (9). Strict positivity is not required. □

The bound $D^\pi_\alpha(\vartheta) \leq \Psi^\pi_\alpha(\vartheta)$ used in the main text is (9) applied to $\Delta\nu = h\, g^\pi_\alpha$ where $g^\pi_\alpha := \partial_\alpha \pi(\nu_\alpha)$ is the projective tangent (see Appendix A.11). Explicitly,

$$\Psi^\pi_\alpha(\vartheta) \;:=\; \big(1 + r_\xi(g^\pi_\alpha) + 2\|P\|_2\big)\, \|g^\pi_\alpha\|_\infty, \tag{10}$$

so $\|h\, g^\pi_\alpha\|_\infty = |h|\, \|g^\pi_\alpha\|_\infty$ and the projective operator-motion bound is $\|L^\pi_{\alpha+h} - L^\pi_\alpha\|_2 \leq |h|\, \Psi^\pi_\alpha(\vartheta) + o(|h|)$. Because g^π_α is a simplex tangent, $\pi(\nu_\alpha) + h\, g^\pi_\alpha$ stays on the fixed-support nonnegative face for all sufficiently small $|h|$ (Appendix A.11), so (9) applies with nonnegative endpoints; dividing by $|h|$ and letting $h \to 0$ yields $D^\pi_\alpha(\vartheta) = \|\partial_\alpha L^\pi_\alpha\|_2 \leq \Psi^\pi_\alpha(\vartheta)$. Empirically $\Psi^\pi_\alpha / D^\pi_\alpha \in [2.08, 4.73]$ on the two dense Stage-2 grids (PolBlogs, Email-EU-Core); the ratio is materially larger on other graphs (e.g. Amazon Photo screening anchors reach $\Psi^\pi_\alpha / D^\pi_\alpha \approx 9.8$, Appendix B.3). The bound is loose because it treats arbitrary nonnegative-cone perturbations and does not exploit the zero-sum simplex-tangent constraint $\mathbf{1}^\top g^\pi = 0$; sharpening it for the simplex tangent is a natural theoretical direction.

Degenerate cases in Cor 1. If $D^{\pi}_{\alpha}(\vartheta) = 0$ the first-order operator motion vanishes and $\rho^{\pi}_{\alpha,\text{oracle}}$ is extended by convention to $+\infty$ as a tangent statement; higher-order motion may still cross τ_{ϑ}, so this is not a finite guarantee. If $\Psi^{\pi}_{\alpha}(\vartheta) = 0$ the analytic bound radius $\rho^{\pi}_{\alpha,\text{bound}}$ is trivial in the same sense. In both cases a genuine finite guarantee comes only from evaluating $\|L^{\pi}_{\alpha+h} - L^{\pi}_{\alpha}\|_2$ directly (the verified step $h^{\star}_{\text{cert}}$) and applying Thm 1.

A.6 Proof of Theorem 1

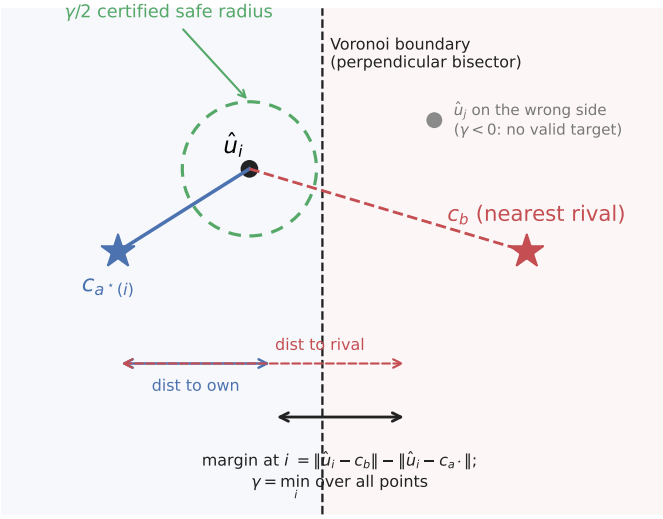

Figure 3: The Voronoi (embedding) margin γ and the $\gamma/2$ safe-move region (schematic). Each embedded row $\hat{u}_i$ is assigned to its nearest reference centroid; its margin is the distance to the nearest rival centroid minus the distance to its own, and $\gamma := \min_i$ over all points. Theorem 1 uses the reverse-triangle bound: perturbing a row by at most η can reduce its distance advantage over a rival centroid by at most 2η, so $\eta < \gamma/2$ is a sufficient condition for preserving every nearest-centroid assignment and leaving the anchored partition unchanged. Geometrically, $\gamma/2$ is a conservative lower bound on the Euclidean distance to the relevant Voronoi boundary (perpendicular bisector), not in general the exact boundary distance. A negative margin ($\hat{u}_j$) means the point already lies on the wrong side, so the chosen partition is not a strict nearest-centroid partition of the embedding and the anchored certificate is inapplicable ($\gamma < 0$); this is the positive-margin boundary of Appendix A.10, not a spectral-identifiability boundary.

Anchored rounding, defined. Given a perturbed basis U and the reference $(U_0, \{c_a\})$, let

$$O^{\star}(U, U_0) \;\in\; \arg\min_{O \in \mathrm{O}(k)} \|U - U_0 O\|_F \tag{11}$$

be the optimal orthogonal (Procrustes) factor for this objective, computed as $O^{\star} = VW^{\top}$ from the SVD $U^{\top}U_0 = W\Sigma V^{\top}$ (equivalently the polar factor of $U_0^{\top}U$): minimizing $\|U - U_0O\|_F^2 = 2k - 2\,\mathrm{tr}(U^{\top}U_0\,O)$ maximizes $\mathrm{tr}(W\Sigma V^{\top}O)$, attained at $V^{\top}OW = I$, i.e. $O^{\star} = VW^{\top}$. Anchored rounding assigns row i to the nearest Procrustes-aligned reference centroid,

$$A_{\text{anc}}(U; U_0, A^{\star})_i \;=\; \arg\min_{a} \|u_i - O^{\star\top} c_a\|_2,$$

ties broken by smallest index a. Under Theorem 1's strict bound every principal angle between the subspaces is below $\pi/2$, so $U^{\top}U_0$ is nonsingular and $O^{\star}$ is unique; this is the map the theorem certifies. Continuation rounding A_{cont} applies the same rule stepwise against the current reference $(U_j, \{c_a^{(j)}\})$.

Let $\epsilon := \|L^{\pi}_{\vartheta'} - L^{\pi}_{\vartheta}\|_2 < \tau_{\vartheta}$, $\bar{\delta} := \bar{\delta}_{\vartheta}$, $\gamma := \gamma_{\vartheta}$. Note $\gamma \le 2$ since $\|\hat{u}_i\|_2, \|c_a\|_2 \le 1$.

Proof. Step 1 (Davis–Kahan, operator norm). $\epsilon < \tau_\vartheta = \gamma\bar{\delta}/(2\sqrt{2}+\gamma) \le 2\bar{\delta}/(2\sqrt{2}+2) < \bar{\delta}/2$. By Weyl's inequality the perturbed eigengap satisfies $\tilde{\lambda}_{k+1} - \tilde{\lambda}_k \ge \bar{\delta} - 2\epsilon > 0$, so the perturbed bottom-k subspace is uniquely defined; the cross-separation used by the residual form of Davis–Kahan satisfies $\tilde{\lambda}_{k+1} - \lambda_k \ge (\lambda_{k+1} - \epsilon) - \lambda_k = \bar{\delta} - \epsilon > 0$. Hence the residual form of the Davis–Kahan $\sin\Theta$ theorem (Davis & Kahan, 1970), with residual $R = (L^\pi_{\vartheta'} - L^\pi_\vartheta)U_0$ and $\|R\|_2 \le \epsilon$, gives $\|\sin\Theta\|_2 \le \|R\|_2/(\bar{\delta} - \epsilon) \le \epsilon/(\bar{\delta} - \epsilon)$, where Θ collects the principal angles between the two bottom-k subspaces. Writing $\theta_{\max}$ for the largest principal angle (with $\theta_{\max} < \pi/2$ since $\|\sin\Theta\|_2 < 1$), the CS decomposition gives $\min_{O\in\mathrm{O}(k)} \|\tilde{U} - U_0 O\|_2 = 2\sin(\theta_{\max}/2)$, and this minimum is attained by the polar Procrustes factor $O^\star = VW^\top$ of (11), the alignment used by A_{anc} itself. Therefore $\|\tilde{U} - U_0O^\star\|_2 = 2\sin(\theta_{\max}/2) \le \sqrt{2}\,\sin\theta_{\max} = \sqrt{2}\,\|\sin\Theta\|_2 \le \sqrt{2}\,\epsilon/(\bar{\delta} - \epsilon)$, the middle inequality using $2\sin(\phi/2) \le \sqrt{2}\,\sin\phi$ for $0 \le \phi < \pi/2$.

Step 2 (operator norm to row). $\eta_{\mathrm{row}} := \max_i \|\tilde{u}_i - (O^\star)^\top \hat{u}_i\|_2 = \max_i \|e_i^\top(\tilde{U} - U_0O^\star)\|_2 \le \|\tilde{U} - U_0O^\star\|_2 \le \sqrt{2}\,\epsilon/(\bar{\delta} - \epsilon)$.

Step 3 (Geometric partition criterion). For $i \in A^\star_a$ and $b \ne a$, triangle inequality gives $\|\tilde{u}_i - (O^\star)^\top c_b\|_2 - \|\tilde{u}_i - (O^\star)^\top c_a\|_2 \ge \|\hat{u}_i - c_b\|_2 - \|\hat{u}_i - c_a\|_2 - 2\eta_{\mathrm{row}} \ge \gamma - 2\eta_{\mathrm{row}}$. Preservation requires $\eta_{\mathrm{row}} < \gamma/2$, i.e., $\sqrt{2}\,\epsilon/(\bar{\delta} - \epsilon) < \gamma/2$, which rearranges to $\epsilon < \tau_\vartheta$. □

Remark (Basis invariance of γ, τ, $\rho^\pi_{\alpha,\mathrm{oracle}}$). The margin γ is defined from an orthonormal basis U_0 of the unique bottom-k subspace. Replacing U_0 by U_0Q for any $Q \in \mathrm{O}(k)$ sends every embedded row $\hat{u}_i \mapsto \hat{u}_iQ$ and every centroid $c_a \mapsto c_aQ$, so $\|\hat{u}_iQ - c_aQ\|_2 = \|\hat{u}_i - c_a\|_2$ and γ is unchanged. Hence γ, and therefore G_k, τ_ϑ, and $\rho^\pi_{\alpha,\mathrm{oracle}}$, do not depend on the arbitrary orthonormal basis chosen for the bottom-k target whenever that subspace is unique.

Remark (Frobenius variant and the role of k). The proof factors into a geometric partition criterion (Step 3, requiring $\eta_{\mathrm{row}} < \gamma/2$) and a spectral backend (Steps 1–2, converting ϵ into η_{row}). The canonical threshold τ_ϑ carries no explicit $\sqrt{k}$ dimension penalty: the operator-norm backend introduces no dependence on k beyond that already present through $\bar{\delta}_k$ and γ_k. A cruder Frobenius backend instead bounds $\|\tilde{U} - U_0O\|_F \le \sqrt{2k}\,\epsilon/(\bar{\delta} - \epsilon)$ and $\eta_{\mathrm{row}} \le \|\cdot\|_F$, yielding the weaker, dimension-dependent

$$\tau^{\mathrm{F}}_\vartheta \;=\; \frac{\gamma_\vartheta\,\bar{\delta}_\vartheta}{2\sqrt{2k} + \gamma_\vartheta} \;\le\; \tau_\vartheta.$$

Since $\tau^{\mathrm{F}}_\vartheta \le \tau_\vartheta$, every step certified under the Frobenius threshold is also certified under the canonical τ_ϑ; the gap $\tau_\vartheta/\tau^{\mathrm{F}}_\vartheta = (2\sqrt{2k} + \gamma_\vartheta)/(2\sqrt{2} + \gamma_\vartheta)$ is the dimension penalty the operator-norm route removes (it depends on γ_ϑ, so it is not a single graph-wide constant). All radii, anchors, and factor rankings reported in the paper use the canonical τ_ϑ; we retain $\tau^{\mathrm{F}}_\vartheta$ only to name the Frobenius variant.

A.7 Proof of Lemma 1

Stepwise notation. Fix t and a discrete path $\vartheta_0 = (t, \alpha_0) \to \cdots \to \vartheta_m = (t, \alpha_m)$ ($\alpha_j > 0$). At each step j, let U_j be an orthonormal basis of the bottom-k eigenspace of $L^\pi_{\vartheta_j}$ and define the stepwise centroids and margin

$$c^{(j)}_a := \tfrac{1}{|A^\star_a|}\sum_{i\in A^\star_a}(U_j)_i, \qquad \gamma_{\vartheta_j} := \min_i \min_{b\ne a^\star(i)} \big(\|(U_j)_i - c^{(j)}_b\|_2 - \|(U_j)_i - c^{(j)}_{a^\star(i)}\|_2\big),$$

so that $\{c^{(j)}_a\}$ determines γ_{ϑ_j}, hence τ_{ϑ_j} and the numerator of $\rho^\pi_{\alpha,\mathrm{oracle}}(\vartheta_j)$, from the continuation membership $A^\star$ (not a fresh clustering at ϑ_j), while $D^\pi_\alpha(\vartheta_j) = \|\partial_\alpha L^\pi_{\vartheta_j}\|_2$ depends only on the operator path.

Proof. Induction on j. Base case $j = 0$: by hypothesis $\gamma_{\vartheta_0} > 0$, so every vertex $i \in A^\star_a$ is strictly closer to its own centroid $c^{(0)}_a$ than to any $c^{(0)}_b$, $b \ne a$; hence anchored rounding of

$L^{\pi}_{\vartheta_0}$ against $(U_0, \{c_a^{(0)}\})$ returns exactly $A^\star$, i.e. $A_{\mathrm{cont}}(\vartheta_0 \mid \vartheta_0) = A^\star$ (this holds whether $A^\star$ is planted or derived from a same-anchor clustering, because $\gamma_{\vartheta_0} > 0$ makes $A^\star$ the nearest-centroid partition of its own centroids). Inductive step: assume ϑ_j's continuation-rounded partition equals $A^\star$. The step $j \to j+1$ is verified, so $\|L^{\pi}_{\vartheta_{j+1}} - L^{\pi}_{\vartheta_j}\|_2 < \eta\,\tau_{\vartheta_j} \le \tau_{\vartheta_j}$; Thm 1 applied with reference ϑ_j gives that anchored rounding at ϑ_{j+1} against $(U_j, \{c_a^{(j)}\})$ preserves ϑ_j's partition, which equals $A^\star$ by IH. □

A.8 Discrete t-axis certificate

The same architecture certifies discrete t-moves: if $\epsilon^{\pi}_t(\vartheta) := \|L^{\pi}_{t+1,\alpha} - L^{\pi}_{t,\alpha}\|_2 < \tau_\vartheta$, the one-step t-move at fixed α preserves the anchored partition. This is Thm 1 instantiated at $\vartheta' = (t+1, \alpha)$. Empirical validation is in Appendix G.

A.9 Perfect-separation reference case

Proposition A.3 (Partition invariance under perfect separation). Assume P is row-stochastic, $\nu > 0$ everywhere, and $V = V_1 \sqcup \cdots \sqcup V_k$ decomposes into exactly k weakly connected components with no inter-component edges. Under a global k-means optimum, the GSC partition of L_ν recovers exactly $\{V_1, \ldots, V_k\}$.

Proof. The energy identity $x^\top L_\nu x = \sum_{i,j} \nu_i P_{ij}(x_i - x_j)^2$ (row-stochasticity gives symmetric quadratic form) implies $x \in \ker L_\nu$ iff $x_i = x_j$ whenever $\nu_i P_{ij} > 0$. Strict positivity of ν makes this "constant on every edge," so $\ker L_\nu = \mathrm{span}\{\mathbf{1}_{V_1}, \ldots, \mathbf{1}_{V_k}\}$. Let $H = [\mathbf{1}_{V_1}, \ldots, \mathbf{1}_{V_k}] \in \mathbb{R}^{N \times k}$. Since $\mathrm{col}(U) = \mathrm{col}(H)$ with both of rank k, there is an invertible $R \in \mathbb{R}^{k \times k}$ with $U = HR$; hence every vertex in V_a has embedding row equal to $R_{a,:}$, and the k rows of R are distinct because R is invertible. There are thus exactly k distinct embedded points, one per component, so global k-means with k nonempty clusters attains the zero objective only by placing each in its own cluster, recovering $\{V_1, \ldots, V_k\}$ up to a label permutation. □

Prop 1 is not the WCC $= k$ setting of Prop A.3. Perfect separation: $\dim\ker L^\pi = k$ and the desired partition is represented exactly. Excess nullity from zero-incoming vertices (Prop 1): $\dim\ker L^\pi > k$ and the bottom-k embedding is not identifiable.

A.10 Controlled chain-DiSBM separation sweep

Proposition A.3 describes the perfectly separated endpoint, where invariance is exact for every $\nu > 0$. It does not by itself imply that an arbitrary increase in separation raises $\bar{\delta}$ and γ monotonically, nor that $\rho^{\pi}_{\alpha,\mathrm{oracle}} = \bar{\delta}\, G_k / D^{\pi}_{\alpha}$ increases monotonically: the denominator D^{π}_{α} may also vary, and eigengaps and embedding margins need not be monotone under edge-weight changes. We therefore test empirically, on a controlled chain-DiSBM contrast sweep, whether $\bar{\delta}$ and γ grow without a commensurate increase in D^{π}_{α}. No general monotonicity follows from Proposition A.3.

Setup. On chain DiSBM ($N = 300$, $k = 3$, 10 seeds) we sweep the intra/inter contrast p_{in}/q (with $\rho = p_{\mathrm{out}} = q$; smaller q is better separation). At each contrast we evaluate the 42-anchor $\nu_{t,\alpha}$ grid and record the grid-mean factors and the measure-invariance: the mean pairwise AMI between the 42 resulting partitions, which is 1 iff every measure yields the same clustering (exp_separation_invariance.py).

Result. Fig 4 matches the expectation. Below a threshold at $\gamma = 0$ (contrast ≈ 5) the Voronoi margin is negative: the planted partition is not a strict nearest-centroid partition of its own embedding, so Theorem 1 and the reported radius do not apply and $\rho^{\pi}_{\alpha,\mathrm{oracle}}$ is undefined; the bottom-k eigenspace itself remains well defined, so this is a rounding-margin boundary, not spectral non-identifiability. Above it, measure-invariance rises monotonically over this sweep from 0.70 to 1.000 ± 0.000 (at contrast 50, all 42 measures give the identical partition), tracking AMI against the planted target. Across the sweep D^{π}_{α} grows only $\approx$

1.3× while $\bar{\delta}$ triples and γ swings from negative to +0.09: separation acts through the numerator, so the radius grows with separation and the measure choice becomes irrelevant over the tested grid once clusters are well separated. This is the finite-separation companion to Prop A.3 (which gives the universal statement over all positive measures at perfect disconnection), and its onset is the positive-margin anchored-rounding boundary ($\gamma = 0$) for the planted partition.

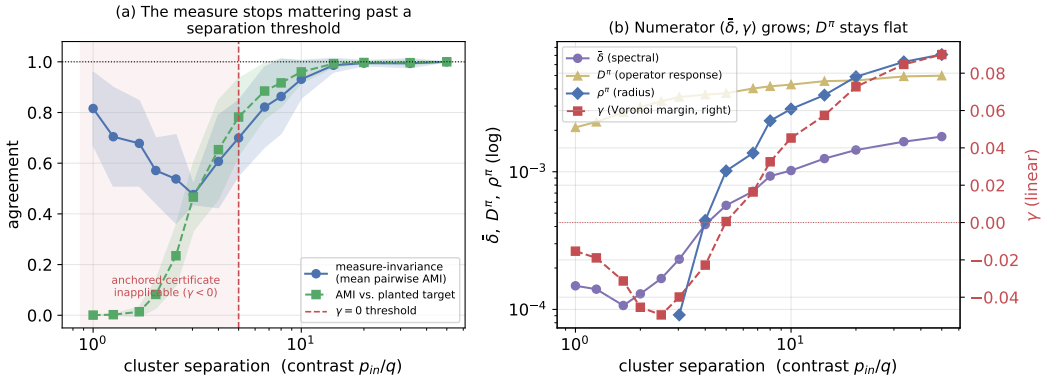

Figure 4: Separation makes the measure irrelevant. Chain DiSBM ($N = 300$, $k = 3$, 10 seeds), swept over intra/inter contrast p_{in}/q; bands are ±1 s.d. across seeds. (a) Mean pairwise AMI across the 42-measure ν-grid (measure-invariance) rises to 1 past the $\gamma = 0$ threshold; the shaded region ($\gamma < 0$) marks where the planted partition is not a strict nearest-centroid partition, so the anchored certificate is inapplicable (the bottom-k eigenspace is still well defined, and the pairwise-AMI metric remains valid there). (b) The factors of $\rho^{\pi}_{\alpha,\text{oracle}} = \bar{\delta}\, G_k / D^{\pi}_{\alpha}$: $\bar{\delta}$ and γ (numerator) grow while D^{π}_{α} (denominator) stays nearly flat, so $\rho^{\pi}_{\alpha,\text{oracle}}$ grows with separation.

A.11 Nonnegative measures and continuous extension

Dangling-node convention. A dangling row ($d_{\text{out}}(i) = 0$) is completed with a self-loop, $P_{i\cdot} = e_i^{\top}$, so P is row-stochastic on every row and Prop 1's premise holds; Z, q_t, and all screening verdicts (Tables in Appendix B) are computed from this completed P. A self-loop adds no inter-vertex transitions, so it does not merge weak components: the support-graph reading of Prop A.3 is unaffected.

We formulate the theory on nonnegative measures with fixed support. Fix t and let $S_t := \{i : q_t(i) > 0\}$. For every $\alpha > 0$, $\nu_{t,\alpha}(i) = q_t(i)^{\alpha}$ on S_t and 0 elsewhere. The expectation $\mathbb{E}[\log q_t]$ below is taken restricted to S_t, so no $0 \log 0$ terms arise:

$$\bar{\ell}_{\alpha} \;:=\; \sum_{j \in S_t} \pi_j(\nu_{t,\alpha}) \log q_t(j).$$

Because $x^{\alpha} \log x \to 0$ as $x \to 0^{+}$ for $\alpha > 0$, the projective tangent

$$g^{\pi}_{\alpha,i} \;=\; \begin{cases} \pi_i(\nu_{\alpha})\big(\log q_t(i) - \bar{\ell}_{\alpha}\big), & i \in S_t, \\ 0, & i \notin S_t, \end{cases}$$

admits the continuous extension $g^{\pi} = 0$ on S_t^{c}, still satisfying $\mathbf{1}^{\top} g^{\pi} = 0$. The bound of Appendix A.5 is algebraic and does not require strict positivity. The one special point is $\alpha = 0$ (where 0^0 requires a convention and support can jump); we scope the differential theory to $\alpha > 0$.

Why L^{π}, not the normalized variant. Sevi et al. (2026)'s normalized Laplacian is $L^{\text{norm}}_{\nu} := D^{-1/2}_{\nu+\xi}\, L_{\nu}\, D^{-1/2}_{\nu+\xi}$ (where $D_{\nu+\xi}$ is invertible, or under the zero-degree convention of Sevi et al., 2026). Both L^{norm}_{ν} and L^{π} are invariant to positive rescaling $\nu \mapsto c\nu$, so the distinction is

not scale-invariance per se. Rather, L_ν^{norm} generally changes the eigenspaces of the unnormalized GSC operator L_ν we study, it is a different clustering operator, whereas our gauge-fixed $L^\pi = L_\nu / s_\vartheta$ evaluates the same unnormalized operator at the ray representative, preserving the eigenvectors (and hence the partition) of L_ν exactly while removing only the ray scale. Moreover L^π inherits the affine-in-ν parent map $\nu \mapsto L_\nu$, the structure the gauge decomposition (5) rests on.

B Screening tables and structural-abstention verification

Purpose. This is the reproducibility appendix for the two-stage protocol of §4. It serves three roles: (i) fix the exact benchmark, estimand, and solver conventions (B.1); (ii) record the Stage-1 verdicts and the residual-aware numerical-resolution evidence (B.2–B.4); and (iii) document the structurally abstained graphs and the sparse-eigensolver failure mode (B.5–B.9).

Reproduction harness and manifest. The code, the datasets, and the intermediate-result CSVs, together with a one-command harness (reproduce.sh), a pinned environment (requirements.txt: numpy 1.26.4, scipy 1.13.1, scikit-learn 1.8.0), and a checksum manifest (MANIFEST.sha256), will be released with the camera-ready version. The harness is a hard gate (it exits non-zero if any check fails) with seven stages. (i) It checks the environment. (ii) It verifies the shipped benchmark and case-study .npz datasets against the manifest by SHA-256 (these are bit-stable, so this is a cross-machine integrity check). (iii) It re-runs the Stage-2 residual test (4) (verify_stage2_eq6.py) and asserts the configuration-independent verdicts, 42/42 anchors resolved on both graphs, $\min_{\mathrm{grid}} \bar{\delta}_k / \tau_{\mathrm{num}} > 10^4$ (PolBlogs) and $> 10^9$ (Email-EU-Core), worst-case pair residual ≤ 10. (iv) It re-runs from scratch the DC-SBM decomposition (§4.3) using a generator vendored into the repository (src/graphs/synthetic_4axis.py) and asserts the per-seed paired log-identity of (8) over 30 seeds: matched-cohort survivor counts 30/17/22, per-seed residual $< 10^{-9}$ (machine precision), and the dominant factor of each sweep (operator $\to 1/D_\alpha^\pi$ nominal; G_k dominating both the spectral and geometry sweeps), reproducing the Appendix J table. (v) It re-runs from scratch the separation sweep (in-repo generator) and asserts measure-invariance: mean pairwise AMI $\to 1$ at high separation, a bounded (nearly flat) D_α^π, and the $\gamma \leq 0$ dropout at low separation. (vi) It re-generates the data figures from the verified numbers (gen_projective_figures.py and fig_voronoi_margin.py; the synth-decomposition and separation figures are produced by stages (iv) and (v)) and asserts each is a valid, non-trivial PDF; Figure 1 (the framework schematic) is a TikZ diagram built from source by the paper itself and is reproduced by the paper build rather than this harness. (vii) It runs the anchored-rounding vs. fresh-k-means comparison on the three real Stage-2 graphs (exp_anchored_vs_fresh_stage2.py) under the canonical fresh-k-means configuration and asserts, without hard-coding any disagreement rate, that the per-anchor artifacts are produced and that each graph's ϑ_{QC} is resolved and lies in Σ_{anc} (Appendix E; PolBlogs and Email-EU-Core run by default, the Amazon full sweep under GSC_FULL= 1). Stage (iii) pins the BLAS thread count (GSC_BLAS_THREADS = 1) so the dense-eigh eigenpair residuals ($\sim 10^{-13}$) are deterministic; the intermediate CSVs carry these backward-error-dependent trailing digits, and Matplotlib PDFs embed timestamps and font subsets, so manifest hashes are a snapshot-integrity check and stage (vi) gates figure generation, not byte-identity; the reproduced scientific content is the verdicts, bounds, and invariants of stages (iii)–(v). We ran all seven stages end to end on a clean checkout and they pass.

Compute. All experiments were run on a single Apple M4 Pro (12 cores, 24 GB RAM) under macOS 15.1 with Python 3.12; no GPU is used. The reproduction environment is pinned in requirements.txt (numpy 1.26.4, scipy 1.13.1, scikit-learn 1.8.0); BLAS threads are pinned to 1 (GSC_BLAS_THREADS= 1) for deterministic eigenpair residuals. The full harness (including the Amazon Photo sparse sweep under GSC_FULL= 1) completes in well under an hour on this machine; the long poles are the 30-seed DC-SBM re-run (stage iv) and the Amazon anchored-vs-fresh sweep (stage vii), each ~ 10 minutes.

Experimental axes vs. tuned hyperparameters. The (t, α), K, teleportation, and synthetic-intervention grids are objects of study (analysis axes swept to probe the theory), not hyperparameters optimised against test performance; their complete tested ranges are reported with each experiment. No analysis axis is tuned to maximise AMI; the reported ϑ_{QC} is nevertheless label-conditioned through $Q_{0.95}$ (§4.1) and is used only as a retrospective oracle diagnostic. The only configurations governing stochastic optimisation (k-means initialisation and restarts) are fixed globally (init="k-means++", $n_{\mathrm{init}} = 30$, random_state= 0; Appendix B.1) and used identically in every experiment, so there is no per-benchmark tuning to disclose.

B.1 Stage-2 estimand and numerical implementation

For the real Stage-2 graphs the reference partition $A^\star$ is the GSC partition itself (fresh k-means on the bottom-k embedding U_{ϑ_0} at the reference anchor ϑ_0), not the ground-truth labelling. Ground truth enters only through the quality gate: AMI(ϑ) against the planted labels defines the near-optimal region $Q_{0.95}$ used to select ϑ_{QC}, and is reported as an external quality measure. Thus the reference is GSC's own partition, not a supervised target; the radius certifies that this reference is preserved under anchored rounding as ϑ moves; it does not predict what a fresh k-means rerun outputs at ϑ' (their empirical agreement is Appendix E). For reference selection each candidate anchor is scored with its own fresh-k-means partition; the anchored-stability, reachability, and anchored-vs-fresh experiments then fix $A^\star = A_{\mathrm{km}}(\vartheta_0)$ at the chosen ϑ_0 and hold it along the path. At each anchor γ_ϑ and A_{anc} are computed against this GSC-derived $A^\star$ in that anchor's bottom-k embedding U_ϑ. The operator response is computed analytically, not by finite differences: since L_ν is linear in ν, $D^\pi_\alpha(\vartheta) = \|\partial_\alpha L^\pi_\vartheta\|_2 = \|L_{g^\pi_\alpha}\|_2$ with the closed-form projective tangent $g^\pi_\alpha = \partial_\alpha \pi(\nu_\alpha)$ ($g^\pi_{\alpha,i} = \pi_i(\log q_t(i) - \sum_j \pi_j \log q_t(j))$; Appendix A.5), and the operator 2-norm is the leading singular value from the dense symmetric eigendecomposition (numpy.linalg.eigh on $\frac{1}{2}(L^\pi + (L^\pi)^\top)$). Fresh k-means uses scikit-learn's KMeans with $n_{\mathrm{init}} = 30$, random_state= 0, labels matched to $A^\star$ by the Hungarian rule for AMI; $\mathrm{AMI}_{\max}$ (used in $Q_{0.95}$) is the maximum over the same 42-anchor grid under this configuration. Each Σ_{anc} row in Appendix B.2 uses $(\vartheta_0, A^\star) = (\vartheta_{\mathrm{QC}}, A_{\mathrm{km}}(\vartheta_{\mathrm{QC}}))$. All scripts and per-anchor CSVs are in the release.

B.2 Stage-1 verdicts on all benchmarks

Table 2: Stage-1 verdicts and full per-graph statistics across the eleven labelled benchmarks plus the two directed Stage-2 additions (OpenFlights, Élysée2017fr). Wine's z-scored construction is discussed in Appendix C.

Graph	Modality	View	N	k	$\lvert Z\rvert$ / %	Structural + numerical verdict
PolBlogs	political-blog cit.	SCC	793	2	0 / 0.0%	admissible; $\min_{\mathrm{grid}} \bar\delta_k/\tau_{\mathrm{num}} > 10^4$ (full Eq. (4), 42/42); resolved
Email-EU-Core	communication	SCC	803	40	0 / 0.0%	admissible; $\min_{\mathrm{grid}} \bar\delta_k/\tau_{\mathrm{num}} > 10^9$ (full Eq. (4), 42/42); resolved
Amazon Photo	co-purchase	SCC	6105	8	0 / 0.0%	admissible; full Stage-2 (sparse, 41/42 anchors resolved; Appendix B.7)
Amazon CS	co-purchase	SCC	11058	10	0 / 0.0%	admissible; $\bar\delta_k/\tau_{\mathrm{num}} > 10^{10}$ (3 pre-screen anchors, full Eq. (4)); resolved
OpenFlights	air routes	SCC	3041	6	0 / 0.0%	admissible; full Stage-2 (42/42 anchors resolved)
Élysée2017fr	retweets	SCC	12246	5	0 / 0.0%	admissible; targeted Stage-2 (34 resolved / 8 Eq. (4)-unresolved; Appendix B.8)
Cora	citation	WCC	2485	7	407 / 16.4%	inadmissible ($\lvert Z\rvert \geq k$)
Cora-ML	citation	WCC	2810	7	1168 / 41.6%	inadmissible
CiteSeer	citation	WCC	2110	6	521 / 24.7%	inadmissible
CiteSeer-Full	citation	WCC	1681	6	1135 / 67.5%	inadmissible
Iris	directed k-NN	full	150	3	5 / 3.3%	inadmissible
Wine	directed k-NN	full	178	3	7 / 3.9%	inadmissible
Gaussian (d=10, sep = 5)	directed k-NN	full	300	3	18 / 6.0%	inadmissible

Dataset provenance and availability. All benchmarks are public. PolBlogs (Adamic & Glance, 2005), Email-EU-Core (Yin et al., 2017), Amazon Photo and Amazon CS (Shchur et al., 2018), Cora and CiteSeer (Sen et al., 2008), and Cora-ML (Bojchevski & Günnemann, 2018) are standard graph benchmarks; Iris and Wine are from the UCI repository (Dua & Graff, 2019) (bundled with scikit-learn), and the Gaussian mixture is generated in-repo. The two directed Stage-2 additions are the OpenFlights air-route database (Patokallio, 2014) (Open Database License) and the Élysée2017fr retweet network (Fraisier et al., 2018b;a) (Zenodo record 5535333, CC-BY-4.0); both have documented community structure (geographic and political, respectively) and are used here as assortative directed testbeds for GSC measure-perturbation stability. We redistribute the exact preprocessed

.npz graphs used here, together with their construction scripts and manifest hashes; redistributed datasets remain under their respective source terms (the Élysée graph is reindexed and carries no raw user identifiers), and our code and author-generated artifacts are released under the MIT license.

Graph-view dependence of the structural verdicts. The Stage-1 verdicts in Table 2 are graph-view dependent, not intrinsic dataset properties. We use the giant strongly connected component when it retains a meaningful benchmark graph (PolBlogs, Email-EU-Core, the Amazon co-purchase networks); citation networks are nearly acyclic and have no comparably useful giant SCC, so we retain their largest weakly connected component, which deliberately keeps source (zero-in-degree) vertices. Proposition 1 therefore diagnoses the operator actually supplied to GSC on the stated view, not a property of the dataset independent of preprocessing; a different admissible view (e.g. an SCC of a citation graph) would yield a different verdict.

Proposition 1 as a structural sanity check: Squirrel and Chameleon. Platonov et al. (2023) observe that the directed Wikipedia benchmarks Squirrel and Chameleon (Rozemberczki et al., 2021) contain large groups of duplicate nodes (identical regression target and out-neighbourhood) whose edges are all outgoing, inducing train/test leakage; they publish filtered versions with those duplicates removed. A node with only outgoing edges has zero in-degree, i.e. it lies in the GSC identifiability set $Z = \{i : \text{column } i \text{ of } P \text{ is } \mathbf{0}\}$, so Proposition 1 should localize the same pathology from an independent direction. It does (Table 3): on the raw directed graphs $|Z| = 2996$ (Squirrel) and 1413 (Chameleon), within 2% of the 2978 and 1387 duplicate counts reported independently by Platonov et al. (2023), and both are structurally inadmissible ($|Z| \gg k = 5$). Their leakage filter, which removes exactly the zero-incoming duplicates, collapses $|Z|$ by over 99% (to 21 and 27), though a few ordinary source nodes remain, so the filtered full graphs are still inadmissible for $k = 5$. The largest SCC is identical for the raw and filtered graphs ($N = 2154$ and 794) and has $|Z| = 0$: a nontrivial SCC cannot contain a zero-incoming vertex, so the entire pathology is peripheral to the strongly connected core, a clean instance of the graph-view dependence above (raw full graph inadmissible, its SCC admissible).

Table 3: Proposition 1 on the directed Squirrel/Chameleon benchmarks (Platonov et al., 2023) ($k = 5$). $|Z|$ on the raw full graph tracks the independently reported duplicate counts (2978, 1387); the leakage filter removes $> 99\%$ of it; the largest SCC is identical raw-vs-filtered and is admissible. A structural check only, no clustering quality is claimed.

Graph (directed)	N	k	$\lvert Z\rvert$	verdict	SCC N / $\lvert Z\rvert$
Squirrel (raw)	5201	5	2996	inadmissible	2154 / 0
Squirrel (filtered)	2223	5	21	inadmissible	2154 / 0
Chameleon (raw)	2277	5	1413	inadmissible	794 / 0
Chameleon (filtered)	890	5	27	inadmissible	794 / 0

This experiment is a structural sanity check for Proposition 1, not an evaluation of GSC clustering quality. Squirrel and Chameleon are heterophilous node-classification benchmarks, and their class labels need not correspond to clusters recoverable by the homophily-oriented GSC objective; we therefore report no AMI or quality region for them. The verdicts are computed in linear time by exp_prop1_squirrel_chameleon.py on the shipped directed graphs (redistributed under their MIT source terms and hashed in the manifest).

The reported ratios are the full residual-based Eq. (4) values, $\bar{\delta}_k/\tau_{\text{num}}$ with $\tau_{\text{num}} = 10\max(r_k, r_{k+1}, \epsilon_{\text{mach}}\|L^\pi\|_2)$: for PolBlogs and Email-EU-Core we rebuilt L^π at every one of the 42 grid anchors, computed both eigenpair residuals by dense eigh, and evaluated τ_{num} per anchor (verify_stage2_eq6.py); all 42/42 anchors are resolved on each graph, with grid-minimum $\bar{\delta}_k/\tau_{\text{num}} > 10^4$ (PolBlogs) and $> 10^9$ (Email-EU-Core) and per-anchor residual ratios $\max(r_k, r_{k+1})/(\epsilon_{\text{mach}}\|L^\pi\|_2) \le 10$. We report these as bounds rather than exact figures: the verdict and its several-order-of-magnitude margin are invariant, but the trailing digits of the minima and residual ratios depend on the dense-eigensolver backward error at the 10^{-13} level, which is not bit-reproducible across BLAS thread counts; the script

pins the thread count for a deterministic artifact (see below). For Amazon Photo and CS the pre-screen scripts likewise evaluate the full Eq. (4) threshold with both residuals; the ratios quoted are those pre-screen values (5 anchors for Photo, 3 for CS). These two verification CSVs are the canonical full-Eq. (4) evidence. A faster raw-margin analysis flag ($\bar{\delta}_k/(\epsilon_{\text{mach}}\|L^\pi\|_2) > 100$, no residual recompute) is retained in the analysis/figure scripts for the printed resolution counts and regime labels (not for anchor selection, which uses quality and radius); it agrees with the full test at every anchor, since $\max(r_k, r_{k+1}) \le 10\,\epsilon_{\text{mach}}\|L^\pi\|_2$ puts $\tau_{\text{num}} \le 100\,\epsilon_{\text{mach}}\|L^\pi\|_2$, so any anchor passing the screen ($\bar{\delta}_k > 100\,\epsilon_{\text{mach}}\|L^\pi\|_2$) also satisfies $\bar{\delta}_k > \tau_{\text{num}}$. Takeaway: every Stage-2 claim sits many orders of magnitude above the residual-aware numerical floor, so the resolution verdicts are not precision-sensitive.

B.3 Amazon Photo SCC ($N = 6105$, $k = 8$): early pre-screen

These five pre-screen anchors predate, and are superseded by, the full 42-anchor sparse Stage-2 sweep of Appendix B.7; they are retained as an additional dense-vs-sparse consistency point (the $\bar{\delta}_8$ values here match the full sweep at the shared anchors).

(t, α)	s_ϑ	γ	$\bar{\delta}_8$	D^π_α	Ψ^π_α	$\rho^\pi_{\alpha,\text{bound}}$	AMI
$(1, 0.10)$	2.36×10^3	1.07×10^{-4}	8.54×10^{-7}	2.99×10^{-3}	1.49×10^{-2}	7.67×10^{-10}	0.573
$(1, 0.20)$	9.32×10^2	2.26×10^{-5}	6.73×10^{-7}	2.83×10^{-3}	2.00×10^{-2}	9.50×10^{-11}	0.587
$(1, 0.30)$	3.74×10^2	7.53×10^{-5}	4.29×10^{-7}	3.24×10^{-3}	2.80×10^{-2}	1.44×10^{-10}	0.586
$(1, 0.50)$	6.34×10^1	3.55×10^{-5}	3.59×10^{-7}	5.75×10^{-3}	5.63×10^{-2}	2.83×10^{-11}	0.615
$(1, 0.70)$	1.15×10^1	1.88×10^{-5}	1.18×10^{-6}	1.33×10^{-2}	1.11×10^{-1}	2.50×10^{-11}	0.611
$(1, 1.00)$	1.00×10^0	1.78×10^{-4}	1.55×10^{-8}	3.35×10^{-2}	2.69×10^{-1}	1.28×10^{-12}	0.307
$(2, 0.10)$	..... 5-anchor pre-screen				,	$\bar{\delta}_8 = 8.56 \times 10^{-7}$	0.587
$(3, 0.10)$	..... 5-anchor pre-screen				,	$\bar{\delta}_8 = 8.33 \times 10^{-7}$	0.587
$(2, 0.50)$	..... 5-anchor pre-screen				,	$\bar{\delta}_8 = 3.96 \times 10^{-8}$	0.544
$(5, 0.50)$	..... 5-anchor pre-screen				,	$\bar{\delta}_8 = 4.65 \times 10^{-8}$	−0.000

Five of these anchors, $(1, 0.10), (2, 0.10), (3, 0.10), (2, 0.50), (5, 0.50)$, were run through the full Eq. (4) pre-screen (both eigenpair residuals), giving $\bar{\delta}_8/\tau_{\text{num}} > 10^9$ at every one, all resolved (we report a bound because τ_{num}'s trailing digits are dense-eigh backward-error dependent at $\sim 10^{-13}$ and vary across BLAS thread counts; the pre-screen script pins the thread count, and exact per-anchor values are in the artifact). The remaining five rows ($t = 1$, $\alpha \in \{0.2, 0.3, 0.5, 0.7, 1.0\}$) come from the full (t, α) sweep, which logs $\bar{\delta}_8$ but not the eigenpair residuals or τ_{num}; their $\bar{\delta}_8$ spans 1.55×10^{-8} to 1.18×10^{-6} (overlapping the pre-screened anchors' range), but a gap alone does not establish Eq. (4) resolution, so we do not claim a verified ratio for them here.

B.4 Amazon CS SCC ($N = 11058$, $k = 10$): screening anchors

(t, α)	excess nullity beyond $\mathbf{1}$	$\bar{\delta}_{10}$	$\bar{\delta}_{10}/\tau_{\text{num}}$	AMI
$(1, 0.10)$	0	1.365×10^{-6}	$> 10^{10}$	0.4694
$(2, 0.10)$	0	1.200×10^{-6}	$> 10^{10}$	0.4680
$(3, 0.10)$	0	7.706×10^{-7}	$> 10^{10}$	0.4684

Extending to $\alpha = 0.5$ was compute-limited (dense 11058×11058 eigh + operator-2-norm calls at ~ 3 h each). The three-anchor result confirms that Amazon CS SCC is numerically resolved at the three tested anchors. As above, the ratios are reported as bounds: the gaps $\bar{\delta}_{10}$ are eigenvalue-stable, but τ_{num} (hence the exact ratio, here ~ 2.5–3.3×10^{10} in the pinned-thread run) depends on the dense-eigh backward error at the 10^{-17} level; the pre-screen script pins the thread count and logs exact per-anchor values.

B.5 Structural-abstention numerical verification

Dense direct eigh on L^π_ϑ at a representative anchor per graph; the count of eigenvalues satisfying $|\lambda_i| \le \theta_{\text{ker}}$, with $\theta_{\text{ker}} := 10 \cdot \epsilon_{\text{mach}}\|L^\pi\|_2$ (i.e. ten backward-error units, matching the numerator of τ_{num}), meets or exceeds Prop 1's lower bound $\dim\ker L^\pi \ge |Z| + 1$. The backward-error scale itself spans $\sim 10^{-14}$ to $\sim 10^{-20}$ across the seven inadmissible graphs

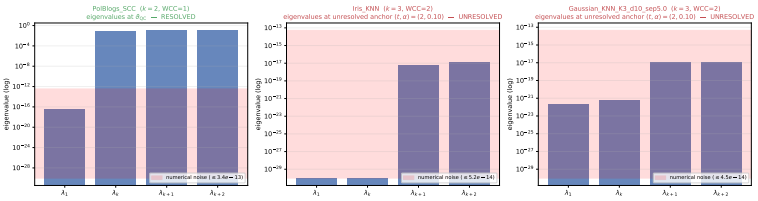

Figure 5: Bottom eigenvalues of L^π on three archetypal graphs. Left: PolBlogs at $\vartheta_{\mathrm{QC}} = (2, 0.10)$ (Stage-2 resolved), a clean gap between λ_k and λ_{k+1}. Middle, right: Iris and Gaussian at an unresolved anchor (directed k-NN, structurally inadmissible), where the bottom $k{+}1$ eigenvalues sit at the backward-error scale ($\sim\pm 10^{-19}$); the log axis shows $|\lambda_j|$ and the shaded band marks the noise floor $\theta_{\mathrm{ker}} = 10\,\epsilon_{\mathrm{mach}}\|L^\pi\|_2$, exactly as Prop 1 predicts. Wine and the four citation benchmarks behave identically at larger $|Z|$.

(CiteSeer at the loose end, Cora-ML at the tight end). Citation-graph norms are reported below alongside the KNN graphs; the analysis is on the largest WCC. The exact values (dense eigh on $\frac{1}{2}(L^\pi + (L^\pi)^\top)$ at $\vartheta = (1, 0.10)$) are also in the release CSV.

Graph	$\|L^\pi\|_2$	$\|Z\|{+}1$ (Prop lower bound)	# eig. with $\|\lambda_i\| \le \theta_{\mathrm{ker}}$	$\bar{\delta}_k$ (backward-error scale)
Iris	$\sim 3 \times 10^{-2}$	6	7	$\sim 10^{-19}$
Wine	$\sim 5 \times 10^{-2}$	8	≥ 8	$\sim 10^{-19}$
Gaussian	$\sim 4 \times 10^{-2}$	19	≥ 21	$\sim 10^{-19}$
Cora	$\sim 2.8 \times 10^{-3}$	408	721	$\sim 10^{-19}$
Cora-ML	$\sim 1.2 \times 10^{-3}$	1169	1559	$\sim 10^{-20}$
CiteSeer	$\sim 3.1 \times 10^{-4}$	522	1198	$\sim 10^{-14}$
CiteSeer-Full	$\sim 6.5 \times 10^{-4}$	1136	1457	$\sim 10^{-18}$

The "≥ 8" / "≥ 21" entries for Wine and Gaussian record a soft ceiling: the eigenvalues immediately above rank $|Z| + 1$ sit within a factor of two of θ_{ker}, so the count is truncated at the last value that is unambiguously below the threshold. The exact counts (which are 9 and 23 respectively when the softer threshold $50\,\epsilon_{\mathrm{mach}}\|L^\pi\|_2$ is used) are in the release CSV. Excess nullity beyond $|Z|{+}1$ arises from weak-connectivity structure and other mechanisms; we make no additivity claim, only the lower bound.

B.6 Sparse vs. dense eigensolver on Wine

At $\vartheta = (8, 0.7)$ on Wine KNN with target $k = 3$, the exact gap $\bar{\delta}_3 = \lambda_4 - \lambda_3$ is zero by Prop 1 (Wine's z-scored construction has $|Z| = 7 \ge k$). Any nonzero value returned by an eigensolver is therefore roundoff or spurious Ritz behaviour, not a numerical estimate of a true nonzero gap. Software and settings: numpy 1.26.4, scipy 1.13.1; sparse solver scipy.sparse.linalg.eigsh with which='SM', k= $k + 2 = 5$, tol= 0 (machine default), ncv $= \max(2k + 1, 20) = 20$, maxiter $= 10\,000$, random v0 at fixed seed 0; the sparse call converged (ArpackNoConvergence not raised) but its reported eigenvalues sit above the true kernel by $\sim 10^{-5}$ and its per-Ritz-pair residuals $\max(r_3, r_4) \sim 10^{-5}$ inflate τ_{num} so that the (4) screen correctly rejects them. Dense solver numpy.linalg.eigh on $\frac{1}{2}(L^\pi + (L^\pi)^\top)$, LAPACK driver dsyevd, 64-bit floats; residual $r_i := \|L^\pi u_i - \lambda_i u_i\|_2 \le 6 \times 10^{-16}\,\|L^\pi\|_2$ for every reported eigenpair ($i \le k + 1$).

Method	λ_3	λ_4	Ritz $\hat{\delta}_3$	$\max(r_3, r_4)$	τ_{num}
eigsh (sparse)	$\approx 2 \times 10^{-5}$	$\approx 5 \times 10^{-5}$	$\approx 3 \times 10^{-5}$	$\sim 10^{-5}$	$\sim 3 \times 10^{-4}$
eigh (dense)	$\approx -1.6 \times 10^{-18}$	$\approx -1.0 \times 10^{-18}$	$\approx 6 \times 10^{-19}$	$\sim 10^{-17}$	$\sim 10^{-16}$

Both rows return the same Eq. (4) verdict, $\hat{\delta}_3 < \tau_{\mathrm{num}}$: unresolved. For the sparse solver this correctly rejects a spurious positive Ritz gap (its inflated residual $\max(r_3, r_4) \sim 10^{-5}$ raises τ_{num} above the apparent gap); for the dense solver it is machine-precision noise around a true zero gap.

Thirteen orders of magnitude separate the two Ritz gaps, but both are correctly declared unresolved by the residual-aware screen (4). The scientific point is that a naive gap-only check ($\hat{\delta}_3 > 0$?) would have accepted the sparse output as a positive gap; the residual-aware τ_{num} inflates proportionally to the eigenpair residual and rejects it. Independently, Prop 1 rules out any positive $\bar{\delta}_3$ structurally, before any eigencomputation.

B.7 Sparse-eigensolver validation and a third full Stage-2 graph

The certificate is solver-agnostic: it needs the bottom-k eigenspace and the gap $\bar{\delta}_k$, however obtained. But a sparse eigensolver must be validated along two distinct axes before its verdict is trusted: (i) eigenpair accuracy: the returned Ritz pairs have small residuals $r_j = \|L^\pi u_j - \lambda_j u_j\|_2$; and (ii) spectral completeness around the target index: no eigenvalue between λ_k and λ_{k+1} was missed. Small residuals control (i) but do not imply (ii): a solver could return λ_k and λ_{k+2} (both with tiny residuals) while skipping λ_{k+1}, inferring a spuriously large gap. The residual-aware screen (4) handles (i); we establish (ii) by oversampling (requesting $k+p$ eigenpairs, $p = 5$) and checking the bottom-$(k+1)$ eigenvalues against a dense reference where feasible.

Solver. We use ARPACK (scipy.sparse.linalg.eigsh) in shift-invert mode with a small negative shift $\sigma = -10^{-6}$: since $L^\pi \mathbf{1} = 0$, exactly $\sigma = 0$ would be singular, so the shift both regularises the factorisation and targets the near-zero cluster (the constant null mode is returned as $\lambda_1 \approx 0$ and kept). SciPy's own guidance is that which='SM' can behave anomalously and shift-invert is the recommended remedy for small eigenvalues, consistent with the which='SM' spurious gap in Appendix B.6. PRIMME and SLEPc provide mature high-performance alternatives; we did not need them here.

Validation against dense references. On PolBlogs and Email-EU-Core, at all 42 anchors, shift-invert reproduces the dense eigh result on both axes (Table 4): identical resolved/unresolved verdicts, and bottom-$(k+1)$ eigenvalues matching to $\leq 10^{-15}\|L^\pi\|_2$ (complete), with maximum relative gap error 3.7×10^{-7} (PolBlogs) and 1.3×10^{-12} (Email-EU). As an independent method, unpreconditioned LOBPCG attains completeness but larger residuals, so the residual screen conservatively marks some borderline anchors unresolved (never falsely resolved), a concrete demonstration that the two axes are distinct. Because a single dense eigh at $N = 6105$ is feasible ($\sim$75s) even though the 42-anchor dense sweep is not, we also validate directly at scale: on five representative Amazon Photo anchors, shift-invert matches dense to $\leq 5 \times 10^{-10}$ relative gap error with completeness to $\sim 10^{-16}\|L^\pi\|_2$, at $\sim 56\times$ lower wall time.

Table 4: Sparse (ARPACK shift-invert, oversampled to $k+5$) vs. dense eigh. Completeness err $= \max_{j \leq k+1} |\lambda_j^{\text{sp}} - \lambda_j^{\text{de}}| / \|L^\pi\|_2$.

Graph (N)	verdict agree	completeness err	max rel. gap err
PolBlogs (793)	42/42	7×10^{-16}	3.7×10^{-7}
Email-EU (803)	42/42	8×10^{-16}	1.3×10^{-12}
Amazon Photo (6105, 5 anchors)	5/5	$\sim 10^{-16}$	5×10^{-10}

Amazon Photo, full Stage-2. With the validated pipeline we run the full 42-anchor sweep on Amazon Photo (co-purchase, $N = 6105$, $k = 8$): 41/42 anchors resolve (residual ratios ≤ 2.7; AMI up to 0.61), with quality-conditioned anchor $\vartheta_{\text{QC}} = (1, 0.7)$. The single exception is the extreme corner $(t, \alpha) = (12, 1.3)$, where the measure is maximally concentrated and $\bar{\delta}_k$ sits at the machine-zero floor: the iterative solver does not converge there and the pipeline abstains rather than report a gap it could not compute, the sound behaviour, and the same conservative direction as the numerical resolution gate. This is a third full Stage-2 real graph, spanning a new (co-purchase) domain at roughly $8\times$ the size of the two dense graphs, obtained without any change to the certificate itself.

B.8 Two additional directed Stage-2 graphs: OpenFlights and Élysée2017fr

These two graphs were selected on pre-GSC structural properties only (natural edge directionality, assortative community structure, an admissible strongly connected component with full label coverage, and manageable size), never on AMI; whatever clustering quality GSC then attains is a reported result, not a selection criterion.

OpenFlights (full Stage-2). The OpenFlights database (Patokallio, 2014) is a public, community-maintained record of the world's scheduled commercial air routes (released under the Open Database License). We build the directed graph whose vertices are airports and whose edge $i \to j$ marks a scheduled route from i to j (load_openflights; binary route-existence, largest SCC $N = 3041$, from routes.dat), with each airport's continent (derived from country metadata) as an external $k = 6$ partition fixed before any GSC computation. Directionality is intrinsic (routes are one-way objects), so no symmetrisation is imposed. Air-transport networks are established community-detection testbeds with geographic community structure; here we use OpenFlights to probe GSC's measure-perturbation stability rather than to detect communities de novo. The graph is strongly assortative (edge homophily 0.86, adjusted homophily 0.80). All 42 anchors resolve; over the grid AMI ranges $[0.00, 0.55]$ (median 0.11), peaking in a low-α band. The quality-conditioned anchor is $\vartheta_{\mathrm{QC}} = (1, 0.5)$ with $\mathrm{AMI}(\vartheta_{\mathrm{QC}}) = 0.54$ ($\mathrm{AMI}_{\max} = 0.55$), and the AMI–$\rho^{\pi}_{\alpha,\mathrm{bound}}$ Spearman correlation is -0.64, matching the pattern of the other full graphs. This is a fourth full Stage-2 real graph in a new (transportation) domain.

Élysée2017fr (targeted case study). The Élysée2017fr corpus (Fraisier et al., 2018b;a) collects 22,853 Twitter profiles active during the 2017 French presidential campaign, with each profile manually annotated by political affiliation among five parties (FI, PS, EM, LR, FN), one of the few political-Twitter resources offering a non-binary, human-curated stance labelling. We use its retweet network: a directed edge $i \to j$ means profile i retweeted j, and we keep single-affiliation profiles as the $k = 5$ ground truth (load_elysee; largest SCC $N = 12246$; redistributed reindexed under CC-BY-4.0 from Zenodo record 5535333, exposing no raw user identifiers). Its retweet graph has documented political-community structure (communities align with the five parties, with preferential same-party attachment; Fraisier et al., 2018b), which is precisely what makes it an apt assortative GSC testbed; we use it to study measure-perturbation stability, not to establish that communities exist. It is strongly assortative (edge homophily 0.84, adjusted homophily 0.79). The 42-anchor accounting separates three distinct notions the paper insists are not interchangeable:

- Numerical resolution (Eq. (4)). From the dense eigenvalue gap compared against the machine-precision floor $\epsilon_{\mathrm{mach}}\|L^{\pi}\|_2$, the dominant term of Eq. (4) for a backward-stable dense decomposition, whose eigenpair residuals are themselves at that scale: 34 resolved, 8 unresolved ($\bar{\delta}_k \le \tau_{\mathrm{num}}$, all at $\alpha \in \{1.0, 1.3\}$, $t \ge 3$), with no solver-budget failures (the dense decomposition always returns). On every anchor carried forward to eigenspace validation the actual residuals r_k, r_{k+1} ($\sim 10^{-18}$) were computed and confirm the floor.
- Eigenspace trustworthiness (for clustering). On resolved anchors we validate the bottom-k eigenspace by residual and a two-solver bottom-k projector agreement $\|U_1 U_1^{\top} - U_2 U_2^{\top}\|_2$; anchors on which the sparse solver cannot resolve the eigenvectors within a fixed wall-clock budget are recorded as computationally unresolved (solver budget), kept distinct from Eq. (4) abstentions.
- Clustering quality. Where the eigenspace is validated, AMI is computed against the affiliations.

The result is a clean three-phenomena illustration. At $\alpha = 0.1$ (every diffusion time t) the target is resolved, the eigenspace validates, and AMI ≈ 0.71; this is the entire quality region $Q_{0.95}$ ($0.95 \cdot \mathrm{AMI}_{\max} = 0.95 \cdot 0.712 = 0.676$). At $\alpha \in \{0.2, 0.3\}$ the target is still resolved and the eigenspace still validates, yet AMI ≈ 0.00: a trustworthy embedding that simply no longer aligns with parties. Farther into high α the low spectrum becomes strongly clustered and resolution/eigenspace computation itself degrades. The quality-conditioned anchor is $\vartheta_{\mathrm{QC}} = (3, 0.1)$, with $\mathrm{AMI}(\vartheta_{\mathrm{QC}}) = 0.71 \approx \mathrm{AMI}_{\max} = 0.712$, the stability-conditioned choice

sacrifices essentially no external quality. We report Élysée2017fr as a targeted case study rather than forcing the pathological corner through an eigensolver; it reinforces that clustering quality, spectral identifiability, numerical resolvability, and local stability are genuinely different properties. Both graphs and the generating scripts ship with the release and are hashed in the manifest.

B.9 AMI in the degenerate regime

Fresh k-means on the numerically-computed bottom-k eigenvectors of the seven inadmissible graphs can return nontrivial AMI on the KNN cases (Iris $\mathrm{AMI}_{\max} = 0.71$, Wine 0.43, Gaussian 0.65) and essentially zero on the citation graphs (0.004–0.05). We do not interpret any of these as evidence of spectral clustering: when $\dim\ker L^{\pi} > k$, the bottom-k subspace is not uniquely defined, so any specific basis returned by the eigensolver depends on numerical implementation. A good external clustering score from one numerical basis does not imply spectral identifiability of the underlying subspace.

C Teleportation-modified UCI panel: from abstention to resolution

Question and result. Can the obstruction diagnosed by Prop 1 be removed by modifying the graph construction? Yes: both teleportation and a larger neighbourhood K can restore numerical resolution, but clearing the structural test ($|Z| < k$) does not by itself guarantee it, as the Iris K-sweep makes explicit (resolution appears only at a still-larger K). The exact construction, software, and k-means conventions are stated next, followed by the results and the trade-off interpretation.

C.1 Reproducible graph construction

Panel. We test three UCI classification datasets that abstain via Prop 1 on the raw directed k-NN construction with $K = \lceil \log N \rceil$: Iris ($N = 150$, $k = 3$), Breast Cancer Wisconsin (diagnostic; $N = 569$, $k = 2$), and Digits ($N = 1797$, $k = 10$). Anchor $(t, \alpha) = (2, 0.10)$; fresh k-means labels are Hungarian-mapped to ground truth only for per-vertex comparisons downstream; AMI itself is permutation-invariant.

Raw (unstandardised) features via sklearn.datasets; Euclidean distance. The final directed adjacency is built directly from the lexicographically sorted cdist neighbours: we form the full distance matrix with scipy.spatial.distance.cdist, exclude the self-distance, sort each row by (distance, index) so equidistant candidates go to the smallest row index, and keep the first K; edge orientation $i \to j$ iff x_j is one of x_i's K nearest neighbours. This deterministic rule replaces sklearn.neighbors.kneighbors_graph (whose default algorithm='auto' may select ball_tree, kd_tree, or brute-force and does not guarantee a fixed tie-break), so the graph is reproducible independently of any internal neighbour-search choice. The graph construction and the AMIs of this appendix were computed under scikit-learn 1.6.1, numpy 1.26.4, scipy 1.13.1, an earlier environment than the pinned reproduction harness (scikit-learn 1.8.0, Appendix B); the reported verdicts are version-invariant, and the shipped .npz graphs are hashed in the manifest. The natural log gives K_{base}. Every AMI reported in this appendix (base, teleportation, and K-sweep) uses sklearn.cluster.KMeans with n_init = 30, init = "k-means++", random_state = 0, max_iter = 300 (the sklearn default), and default tolerance, applied to the bottom-k eigenvectors. Wine under this raw-feature construction yields $|Z| = 1$, whereas the $|Z| = 7$ value reported for Wine in Table 2 arises under z-scored features (see the footnote to that table for the precise StandardScaler settings). The two values correspond to two different graph-construction pipelines, not to conflicting reports about a single pipeline. This appendix uses the raw-feature construction throughout for consistency with Iris/Breast Cancer/Digits, and Wine is excluded from this panel.

C.2 Teleportation intervention

Reference: §4.2. Prop 1 rules out perturbative certification whenever $|Z| \geq k$. Teleportation regularisation

$$\hat{P} \;:=\; (1-\varepsilon)\,P \;+\; \varepsilon\,\frac{1}{N}\,\mathbf{1}\mathbf{1}^{\top}, \qquad \varepsilon \in (0,1), \tag{12}$$

enforces $|Z(\hat{P})| = 0$ by construction: every column of $\hat{P}$ has entries $\geq \varepsilon/N > 0$, so no vertex has zero in-degree. Prop 1's hypothesis is neutralised structurally; whether the target eigenspace is then numerically resolved ($\bar{\delta}_k > \tau_{\text{num}}$) is a separate empirical question.

Dataset	N	k	$\lvert Z\rvert_{\text{base}}$	Base verdict	Base AMI	ε=0.02 verdict	ε=0.02 AMI
Iris	150	3	5	abstain (Prop 1)	0.674	resolved	0.803
Breast Cancer	569	2	10	abstain (Prop 1)	0.000	resolved	0.489
Digits	1797	10	38	abstain (Prop 1)	0.000	resolved	0.900

Findings.

- Teleportation removes the zero-incoming obstruction on all three benchmarks; the eigenspace is then numerically resolved at the tested anchor. Structural admissibility is guaranteed by construction; at $\varepsilon = 0.02$ and $(t, \alpha) = (2, 0.10)$, the numerical eigengap $\bar{\delta}_k$ jumps from the backward-error scale ($\sim 10^{-20}$–10^{-14}) to $\sim 10^{-4}$–10^{-6}, resolving by $\bar{\delta}_k/\tau_{\text{num}} \sim 10^{10}$–$10^{12}$ above the residual-based threshold Eq. (4) (per-case breakdown in Table 5). Removing Prop 1's structural hypothesis does not by itself prove identifiability across all (t, α); we report the tested anchor.
- Reported fixed-seed AMI is higher after modification. Because the base regime is non-identifiable, its AMI is not a valid baseline for a causal comparison: it is either near zero (Breast Cancer, Digits, where any fresh k-means output reflects a numerical basis of a degenerate subspace, not clustering of an identifiable eigenspace) or moderate (Iris 0.67, an incidental alignment despite non-identifiability). After regularisation the reported fixed-seed AMI is 0.803 (Iris), 0.489 (Breast Cancer), and 0.900 (Digits) at $\varepsilon = 0.02$; we report this as a description of the modified construction, not as a causal estimate of improvement.
- Digits is the largest resolved point-cloud case in our experiments ($N = 1797$, $k = 10$; AMI 0.900 at $\varepsilon = 0.02$), on a data modality (a handwritten-digit point cloud) distinct from the communication networks of §4.1–§4.3. Amazon Photo ($N = 6105$) is now resolved over the full 42-anchor Stage-2 grid (41/42; sparse eigensolver, Appendix B.7), and Amazon CS ($N = 11058$) is resolved at its 3 tested pre-screen anchors (Appendix B).
- Resolution holds at the tested ε values. Repeating at $\varepsilon \in \{0.05, 0.10\}$ preserves resolution and holds AMI within ± 0.05 of the $\varepsilon = 0.02$ value, so the effect is not a fragile artefact of the regularisation constant.

C.3 Increasing the neighbourhood K

Teleportation is not the unique way to remove the zero-incoming obstruction. For directed k-NN graphs, increasing the neighbourhood size K reduces $|Z|$ while preserving sparsity and the local graph construction. We repeat the identifiability check while varying $K \in \{6, 8, 10, 12, 15, 20, 30, 50\}$ on the same three datasets. Importantly, passing the structural test is not sufficient for spectral identifiability. On Iris, the strengthened Prop 1 trigger $|Z| \geq k = 3$ fires up to $K = 10$ ($|Z| = 3$ at $K = 8, 10$); at $K = 12$, $|Z| = 2 < k$ clears the structural test, but the target gap $\bar{\delta}_3$ still remains at the backward-error scale ($\sim 10^{-18}$). The eigenspace becomes numerically resolved only at $K = 15$ ($\bar{\delta}_3 = 4.4 \times 10^{-4}$). This is a clean empirical illustration of the one-way nature of Prop 1: $|Z| \geq k$ guarantees non-identifiability, while $|Z| < k$ only removes this particular obstruction and must be followed by the spectral-resolution test.

Dataset	N	k	default K	first resolved K / best K	best K-sweep AMI	teleportation AMI (ε=0.02)
Iris	150	3	6	15 / 20	0.783	0.803
Breast Cancer	569	2	7	30 / 50	0.413	0.489
Digits	1797	10	8	12 / 30	0.846	0.900

The first-resolved-K column is the smallest $K \in \{6, 8, 10, 12, 15, 20, 30, 50\}$ for which both gates pass at the tested anchor $(t, \alpha) = (2, 0.10)$ under the exact software, LAPACK/BLAS, and tie-handling of this appendix (scikit-learn 1.6.1 as noted above, numpy 1.26.4, scipy 1.13.1, numpy.linalg.eigh on $\frac{1}{2}(L^\pi + (L^\pi)^\top)$, LAPACK driver dsyevd, OpenBLAS 0.3.21, and $\theta_{\ker} = 10\,\epsilon_{\text{mach}}\|L^\pi\|_2$). At Digits $K = 12$, $|Z| = 9 < k = 10$ clears Prop 1 and $\bar{\delta}_{10} = 2.85 \times 10^{-6}$; the residual-based ratio $\bar{\delta}_{10}/\tau_{\text{num}} > 10^{11}$ (Eq. (4), Table 5). Fresh-k-means AMI at this anchor is 0.00 under KMeans(n_init=30, random_state=0) (the resolved eigenspace does not match the digit labels at this coarse anchor). The best-K column reports the maximum fresh-k-means AMI over the swept K values, obtained on the same stack (0.846 at $K = 30$ under the lexicographic neighbour rule of this appendix). Any small run-to-run variation in the AMI comes from the upstream computed labels (neighbour tie resolution, eigensolver output, k-means initialisation), not from AMI itself, which is deterministic once the two label vectors are fixed.

Table 5: Residual-based numerical-resolution ratios (Eq. (4)) at anchor $(t, \alpha) = (2, 0.10)$, computed with the lexicographic neighbour rule of this appendix and dense numpy.linalg.eigh on $\frac{1}{2}(L^\pi + (L^\pi)^\top)$. Here $\tau_{\text{num}} = 10\max(r_k, r_{k+1}, \epsilon_{\text{mach}}\|L^\pi\|_2)$ with $r_i := \|L^\pi u_i - \lambda_i u_i\|_2$; on these dense solves the eigenpair residuals sit at the backward-error floor, so τ_{num} is dominated by $\max(r_k, r_{k+1})$ rather than by the third term $\epsilon_{\text{mach}}\|L^\pi\|_2$ (the $\|L^\pi\|_2$ column lets the reader check all three terms of the maximum). The teleportation rows use $\varepsilon = 0.02$. The $r_k, r_{k+1}, \tau_{\text{num}}$ columns are representative values on the stack named above (their trailing digits are dense-eigh backward-error dependent at $\sim 10^{-17}$); the $\bar{\delta}_k/\tau_{\text{num}}$ column is reported as a robust decade lower bound, which every stack we tested clears with margin. All ratios use the residual-aware τ_{num} of Eq. (4).

Case	$\lvert Z\rvert$	$\bar{\delta}_k$	$\|L^\pi\|_2$	r_k	r_{k+1}	τ_{num}	$\bar{\delta}_k/\tau_{\text{num}}$
Iris (teleport)	0	3.05×10^{-4}	2.65×10^{-2}	2.20×10^{-17}	2.46×10^{-17}	2.46×10^{-16}	$> 10^{12}$
Breast Cancer (teleport)	0	3.75×10^{-6}	6.70×10^{-3}	6.43×10^{-18}	6.23×10^{-18}	6.43×10^{-17}	$> 10^{10}$
Digits (teleport)	0	2.18×10^{-6}	2.68×10^{-3}	3.78×10^{-18}	3.85×10^{-18}	3.85×10^{-17}	$> 10^{10}$
Digits K=12 (raw)	9	2.85×10^{-6}	2.94×10^{-3}	2.02×10^{-18}	2.53×10^{-18}	2.53×10^{-17}	$> 10^{11}$

C.4 Interpretation and scope

The panel converts Prop 1's abstention output from "this graph cannot be certified" into "this graph cannot be certified as constructed, but the theory suggests a principled modification." Concretely,

$$\text{diagnose an exact structural obstruction} \longrightarrow \text{modify the construction}$$
$$\longrightarrow \text{retest structural admissibility and numerical spectral resolution}$$
$$\longrightarrow \text{proceed with the stability analysis.}$$

The identifiability layer thus becomes a design tool as well as a diagnostic, for point-cloud data where the raw directed k-NN construction is inadmissible, a small teleportation constant is one such modification.

Scope and caveats. This panel establishes single-anchor numerical-resolution recovery at the tested settings, not exact identifiability of the modified graph and not full parameter-space stability: we ran Stage-1 structural admissibility (by construction) plus a single-anchor spectral-resolution check (Eq. 4, a residual-aware numerical heuristic) plus a single-anchor AMI report. The three regularised cases are therefore best viewed as members of the screening tier alongside Amazon CS, not as Stage-2 deep-analysis cases. The regularisation ε trades numerical resolution against departure from the original directed structure; $\varepsilon = 0.02$ suffices in the three cases here, but the appropriate value in general is graph-dependent.

We do not extend the analysis to Wine because this appendix uses the raw-feature k-NN construction (under which $|Z| = 1$ and Prop 1 does not fire), while Table 2 reports Wine under a z-scored construction ($|Z| = 7$); the two pipelines are simply different, and we do not compare them here.

The two interventions expose a trade-off rather than a universally preferred repair. Increasing K retains a sparse, locality-based graph but blurs community boundaries as neighbourhoods grow; teleportation preserves 98% of the original transition mass at $\varepsilon = 0.02$ but introduces weak global transitions and a dense support. In our three cases teleportation yields the higher AMI, but we describe this as single-anchor numerical-resolution recovery at $(2, 0.10)$, not proof of exact identifiability on the modified graph. The broader principle is that Prop 1 diagnoses a graph-construction failure mode and motivates modifying the construction before re-testing structural admissibility and numerical spectral resolution. Teleportation is one such modification, not a prerequisite of GSC and not the uniquely prescribed repair.

D Continuation: local quantities along the DiSBM path

Purpose. This appendix supports RQ4 by showing how a sequence of individually verified local steps produces the reported continuation reach. We first specify the step-search procedure, then report the path and the evolution of the quantities entering the certificate.

Reference $\vartheta_0 = (1, 0.1)$ on chain DiSBM ($\rho = 0.05$, $N = 600$, $k = 3$). Local quantities at ϑ_0: $\bar{\delta}_0 = 8.40 \times 10^{-4}$, $\gamma_0 = 6.30 \times 10^{-2}$, $D_0^\pi = 1.98 \times 10^{-3}$, $\Psi_0^\pi = 4.53 \times 10^{-3}$, $\tau_0 = 1.07 \times 10^{-5}$, $\mathrm{AMI}(\vartheta_0) = 1.0$.

Bisection protocol (continuation). This continuation search is a separate protocol from the single-anchor $h^\star_{\mathrm{cert}}$ search (Appendix A.3, which uses an absolute tolerance $\epsilon_h = 10^{-6}$; §3.4); we state both explicitly to avoid conflating them. The continuation single-shot verified step h_0 is the largest h located by bisection within the search bracket for which the strict predicate $\|L^\pi_{\alpha_0+h} - L^\pi_{\alpha_0}\|_2 < \eta\,\tau_{\alpha_0}$ holds ($\eta = 0.9$); it is the largest feasible point located by the bracket search, not a proven global maximum. The routine starts from an analytic estimate $h_{\mathrm{init}} = h_{\mathrm{dir}} = \eta\,\tau_{\alpha_0}/\Psi^\pi_\alpha(\vartheta_0) \approx 2.12 \times 10^{-3}$, doubles the upper end while the predicate still holds (each doubling capped at the domain limit $h_{\max} = \min(h_{\mathrm{bound}}, 0.5)$, where h_{bound} is the distance to the directional parameter boundary, $\alpha_{\max} - \alpha_j$ increasing or $\alpha_j - \alpha_{\min}$ decreasing, with the code's bounds $\alpha_{\min} = 0$, $\alpha_{\max} = 3$), then bisects the resulting feasible/infeasible bracket $[h^\downarrow, h^\uparrow]$ for at most 20 iterations, stopping when the relative bracket width $(h^\uparrow - h^\downarrow)/\max(h^\uparrow, 10^{-13})$ falls below $\mathrm{tol}_{\mathrm{rel}} = 0.05$. At ϑ_0 this terminates in 7 operator-norm evaluations (1 initial test, 2 expansion tests, 4 bisection tests; the routine returns a cached feasible endpoint, with no additional endpoint re-check) at the terminal bracket $[h^\downarrow, h^\uparrow] = [6.3541 \times 10^{-3}, 6.6188 \times 10^{-3}]$ (relative width $0.040 < 0.05$), with $h_0 := h^\downarrow = 6.354 \times 10^{-3}$; the operator distance there is $\|L^\pi_{\alpha_0+h_0} - L^\pi_{\alpha_0}\|_2 = 0.867\,\tau_{\alpha_0}$, i.e. $0.033\,\tau_{\alpha_0}$ below the $\eta\,\tau_{\alpha_0} = 0.9\,\tau_{\alpha_0}$ threshold, so h_0 is feasible. The rounded display value 6.35×10^{-3} elsewhere in the paper refers to the same h_0.

In the increasing-α direction the run took 80 certified accepted steps (the configured max_steps); the certificate holds at every one (termination on max_steps, not on certificate failure), yielding $\Delta\alpha_+ = 0.3725$ ($\alpha : 0.1 \to 0.4725$). In the decreasing direction the implementation produced 16 numerical transitions: the first 15 are certified accepted steps ending at the last positive node $\alpha = 2.4 \times 10^{-3}$; the 16th, a diagnostic transition to $\alpha = 0$, uses the increment $h = -\alpha \notin H_\vartheta(\eta)$ (its endpoint $\alpha + h = 0$ leaves the domain; Eq. 1 and Appendix A.11 scope the theory to $\alpha > 0$), so it is not a certified step. The certified decreasing reach is therefore $\Delta\alpha_- = 0.1 - 2.4 \times 10^{-3} = 0.0976 \approx 0.098$; certifying the $\alpha = 0$ endpoint would require a separately defined $\alpha = 0$ measure for this graph. At the visited nodes (both directions), the quantities entering the certificate stay bounded away from their degenerate limits: $\bar{\delta} \in [4.67 \times 10^{-4}, 9.44 \times 10^{-4}]$ (bounded away from 0; max/min span 2.02×, the widest of the three), $\gamma \in [6.19 \times 10^{-2}, 6.31 \times 10^{-2}]$ (bounded away from 0; span 1.02×, essentially constant), and the operator bound $\Psi^\pi_\alpha \in [3.82 \times 10^{-3}, 4.85 \times 10^{-3}]$ (bounded above; span 1.27×). $\bar{\delta}, \gamma$ bounded away from zero keep τ from collapsing, and

Ψ_α^π bounded above keeps the analytic operator-motion bound controlled; this empirically helps explain the near-linear reach on the computed path; each individual step is independently certified by its directly checked operator inequality, so the linear-reach statement is an observation about the visited nodes, not a theorem (a general implication would require controlling $\bar{\delta}, \gamma, \Psi_\alpha^\pi$ across each inter-node interval, not only at the nodes). We report the range of the bound $\Psi_\alpha^\pi \geq D_\alpha^\pi$; D_α^π itself was not logged, and $D_\alpha^\pi \leq \Psi_\alpha^\pi$ does not by itself imply D_α^π is stationary.

Compositional gain. $\Delta\alpha_+/h_0 = 0.3725/(6.354 \times 10^{-3}) \approx 58.6\times$ relative to the returned feasible endpoint $h_0 = h^\downarrow$ ($0.3725/h^\uparrow \approx 56.3\times$ relative to the bracket's upper endpoint, under the monotonicity assumption). This is a lower bound on the reach, since the run stopped on the step cap, not on certificate failure. Continuation gives an approximately linear compositional gain when the certificate quantities stay bounded, not a qualitative rescue: the scientific content is the conditions under which they do.

E Anchored rounding vs. fresh k-means: empirical agreement

Thm 1 certifies anchored rounding: label vertex i by its nearest reference centroid after optimal (Procrustes) orthogonal alignment (Appendix A.6, (11)). Standard GSC pipelines instead run fresh k-means on the bottom-k eigenvectors at each parameter and match labels with the Hungarian rule against the reference. We do not claim mathematical equivalence between the two rules, but we compare them empirically on all three real Stage-2 graphs (PolBlogs, Email-EU-Core, Amazon Photo).

Table 6: The three rounding maps. Only anchored and continuation rounding carry theoretical guarantees; fresh k-means is the standard pipeline we compare against empirically.

Map	Rule	Status
$A_{\mathrm{km}}(\vartheta)$	fresh k-means on the rows of U_ϑ	empirical (no guarantee)
$A_{\mathrm{anc}}(\vartheta; \vartheta_0)$	nearest reference centroid after optimal (Procrustes) orthogonal alignment to U_{ϑ_0} (App. A.6)	certified by Thm 1
$A_{\mathrm{cont}}(\vartheta_j \mid \vartheta_0, \ldots, \vartheta_j)$	sequential anchored rounding along a discrete path	certified by Lem 1

Disagreement metrics. We report two quantities:

- Anchor-level disagreement rate $d_{\mathrm{anc}}(R) := |\{\vartheta \in R : A_{\mathrm{anc}}(\vartheta; \vartheta_0) \neq A_{\mathrm{km}}(\vartheta)\}|/|R|$: the fraction of anchors in region R whose complete partitions are unequal, after Hungarian relabelling of $A_{\mathrm{km}}(\vartheta)$ to $A_{\mathrm{anc}}(\vartheta; \vartheta_0)$.
- Anchor–vertex disagreement rate $d_{\mathrm{vert}}(R) := \sum_{\vartheta \in R} d_H(\vartheta)/(|R| \cdot N)$ (a rate over anchor–vertex observations, not over pairs of vertices), where $d_H(\vartheta)$ is the Hamming distance between the two relabelled partitions at ϑ.

The two metrics differ by a large factor: a single flipped boundary vertex makes d_{anc} count the whole anchor as a mismatch while d_{vert} counts $1/N$. Fresh k-means uses exactly the canonical Stage-2 configuration (Appendix B.1): sklearn.cluster.KMeans with n_init = 30, init = "k-means++", random_state = 0, and max_iter = 300 (the sklearn default), the same fresh-k-means rule everywhere in the paper. We compare the two rounding maps on the three real Stage-2 graphs over their shared 42-anchor grid ($t \in \{1, 2, 3, 5, 8, 12\}$, $\alpha \in \{0.1, 0.2, 0.3, 0.5, 0.7, 1.0, 1.3\}$), using each graph's established quality-conditioned anchor ϑ_{QC} (§4.1) as the reference and $A^\star = A_{\mathrm{km}}(\vartheta_{\mathrm{QC}})$ as the estimand, exactly as in the main results. Only resolved anchors enter the comparison (Amazon's corner $(12, 1.3)$ abstains at Stage 1 and is excluded, leaving 41). The near-optimal quality region $Q_{0.95}$ and the empirical anchored stability set $\Sigma_{\mathrm{anc}} = \{\vartheta : A_{\mathrm{anc}}(\vartheta; \vartheta_{\mathrm{QC}}) = A^\star\}$ are computed from the run, not fixed in advance. Each cell reports d_{anc} with d_{vert} in parentheses.

Graph	$\lvert R_{\mathrm{res}}\rvert$	full-grid d_{anc} (d_{vert})	$\lvert Q_{0.95}\rvert$: d_{anc} (d_{vert})	$\lvert\Sigma_{\mathrm{anc}}\rvert$: d_{anc}
PolBlogs SCC (N=793, k=2)	42	31/42 = 73.8% (3.6%)	11: 0/11 = 0.0% (0.0%)	10: 0/10 = 0.0%
Email-EU-Core SCC (N=803, k=40)	42	41/42 = 97.6% (8.4%)	3: 3/3 = 100% (14.4%)	1: 0/1 = 0.0%
Amazon Photo SCC (N=6105, k=8)	41	39/41 = 95.1% (1.8%)	14: 13/14 = 92.9% (3.9%)	1: 0/1 = 0.0%

The comparison is strongly graph-dependent. On PolBlogs the two rules are identical on every anchor of both certificate-relevant regions (0/11 on $Q_{0.95}$, 0/10 on Σ_{anc}), and $\Sigma_{\mathrm{anc}} \subseteq Q_{0.95}$ (all 10 anchors coincident); anchored rounding and standard fresh-k-means agree wherever the certificate operates, and disagree only on low-quality anchors far outside. The two larger examples (Email-EU-Core, Amazon Photo) show substantially more disagreement. Exact global-partition equality is a stringent test (a single flipped vertex makes both Σ_{anc}, exact equality to $A^\star$, and the anchor-level d_{anc} all-or-nothing), and there Σ_{anc} collapses to $\{\vartheta_{\mathrm{QC}}\}$ while d_{anc} reaches 92.9%–100% inside $Q_{0.95}$. The vertex-level rate tells the substantive story: on Amazon the two rules agree on $\geq 96\%$ of vertex labels throughout $Q_{0.95}$ ($d_{\mathrm{vert}} = 3.9\%$) and on 98% over the full resolved grid (1.8%); the disagreements are boundary vertices among the 8 clusters (9–396 of 6105 per anchor). Email is the hardest case: its 40 small clusters leave the two rounding rules differing on $\approx 14\%$ of vertices even in $Q_{0.95}$, so we do not claim they are interchangeable there. These three graphs do not isolate k from graph geometry, cluster balance, or N, so we read the pattern as graph dependence rather than a monotone k-effect. The certificate is a guarantee about anchored rounding; empirically it coincides with fresh-k-means exactly on PolBlogs and at the vertex level on Amazon, and diverges materially on Email. Per-anchor artifacts (anchored_vs_fresh_<graph>.csv) and the generating script (exp_anchored_vs_fresh_stage2.py) are gated by the reproduction harness (Appendix B).

Continuation (§4.4) further reports that on the full increasing/decreasing DiSBM path all three rounding rules (anchored, continuation, fresh k-means) produced identical partitions at every certified accepted step, with AMI = 1.0 against ground truth throughout. We read these results as delimiting, not extending, the theorem's scope: the bridge from the anchored certificate to standard fresh-k-means practice is strongly graph-dependent: exact on PolBlogs, close at the vertex level on Amazon Photo ($\geq 96\%$ label agreement in $Q_{0.95}$), and materially different on Email-EU-Core ($\sim 14\%$ of vertices reassigned in $Q_{0.95}$). The two larger-k examples disagree more than PolBlogs, but these three graphs do not isolate k from graph geometry or cluster structure. Anchored and fresh-k-means stability are therefore distinct notions; the extent of their disagreement varies substantially across the tested graphs, and Theorem 1 should be read strictly as an anchored-rounding guarantee rather than a prediction about fresh k-means.

F Stability along the β-axis: directionality interpolation

The main paper (§3.4) analyses tuning along the α-axis of the canonical family $\nu_{t,\alpha}$, where the operator map is nonlinear and Cor 1 is a first-order diagnostic. Here we show that on a distinct canonical measure path (the directionality-interpolation path between the out- and in-degree measures) the operator map is affine in the path parameter and the analogous radius is an exact finite-step certificate, with no first-order approximation.

The directionality path. For a weighted directed graph $G = (V, E, w)$ with total edge weight $M := \sum_i d_{\mathrm{out}}(i) = \sum_i d_{\mathrm{in}}(i) > 0$ (i.e. at least one edge), define the out- and in-degree probability measures

$$p_{\mathrm{out}}(i) := \frac{d_{\mathrm{out}}(i)}{M}, \qquad p_{\mathrm{in}}(i) := \frac{d_{\mathrm{in}}(i)}{M},$$

and the directionality-interpolation path

$$\nu_\beta \;:=\; \beta\, p_{\mathrm{out}} \;+\; (1-\beta)\, p_{\mathrm{in}}, \qquad \beta \in [0,1]. \tag{13}$$

Since $\mathbf{1}^\top p_{\mathrm{out}} = \mathbf{1}^\top p_{\mathrm{in}} = 1$, we have $\mathbf{1}^\top \nu_\beta = 1$ for every β: the entire path lives in the projective simplex, so $\pi(\nu_\beta) = \nu_\beta$ and no gauge motion contributes anywhere along it. The path connects the midpoint $\nu_{1/2}(i) = (d_{\mathrm{out}}(i) + d_{\mathrm{in}}(i))/(2M)$, which is the symmetric degree measure ν_{deg}.

Exact affine operator path. The GSC Laplacian is linear in ν, so

$$L^{\pi}_{\nu_\beta} \;=\; L_{\nu_\beta} \;=\; L_{\nu_{\mathrm{deg}}} \;+\; \big(\beta - \tfrac{1}{2}\big)\big(L_{p_{\mathrm{out}}} - L_{p_{\mathrm{in}}}\big), \tag{14}$$

i.e., $L^{\pi}_{\nu_\beta}$ is affine in β. Consequently $\partial_\beta L^{\pi}_{\nu_\beta} = L_{p_{\text{out}}} - L_{p_{\text{in}}}$ is constant in β and, for every $\beta, \beta + h \in [0, 1]$,

$$\left\| L^{\pi}_{\nu_{\beta+h}} - L^{\pi}_{\nu_\beta} \right\|_2 \;=\; |h| \cdot C_\beta, \qquad C_\beta := \left\| L_{p_{\text{out}}} - L_{p_{\text{in}}} \right\|_2,$$

exactly, with no higher-order remainder.

Exact affine sufficient certificate on β (corollary of Thm 1). At $\beta_0 \in [0, 1]$ with $\bar{\delta}_{\beta_0} > 0$ and $\gamma_{\beta_0} > 0$, applying Thm 1 to the operator path (14) gives the exact affine sufficient certificate

$$\boxed{\rho^{\pi}_{\beta}(\beta_0) \;:=\; \begin{cases} \tau_{\beta_0}/C_\beta \;=\; \bar{\delta}_{\beta_0}\, G_k(\beta_0)/C_\beta, & C_\beta > 0, \\ +\infty, & C_\beta = 0. \end{cases}} \tag{15}$$

Any $\beta \in [0, 1]$ with $|\beta - \beta_0| < \rho^{\pi}_{\beta}(\beta_0)$ preserves the anchored-rounded partition. The affinity of (14) makes the operator-motion side $\|\Delta L^{\pi}\|_2 = |h| C_\beta$ exact (no first-order remainder), so (15) is the exact conversion of Thm 1's sufficient condition. It is not, however, the maximal true partition-stability radius: like Thm 1 it uses the canonical operator-norm constant $2\sqrt{2}$, and its converse is not claimed: partition invariance may extend further along the path than this sufficient certificate reports.

Balanced digraphs. If $d_{\text{out}}(i) = d_{\text{in}}(i)$ for every vertex, then $p_{\text{out}} = p_{\text{in}}$ and $\nu_\beta \equiv \nu_{\text{deg}}$ for all $\beta \in [0, 1]$. Hence $C_\beta = 0$ and $\rho^{\pi}_{\beta}(\beta_0) = +\infty$. Balancedness alone proves only that the operator path is constant; the further conclusion that the entire bounded β-interval belongs to a single anchored stability set $\Sigma(A^\star)$ additionally requires an identifiable, admissible reference at β_0 ($\bar{\delta}_{\beta_0} > 0$ and $\gamma_{\beta_0} > 0$). Under that condition the directionality component of the measure has no effect on the partition when the graph's degree sequence is symmetric.

Remark (Two structural readings, one theory). Prop 1 identifies a structural regime where the target eigenspace is not uniquely determined, so no unique bottom-k projector (hence no canonical radius or certificate) is defined. (15) identifies a distinct structural reading: on balanced digraphs the operator path along β is constant, so $C_\beta = 0$ and $\rho^{\pi}_{\beta} = +\infty$. Balancedness proves operator constancy; whole-interval invariance of a chosen anchored partition additionally requires a well-defined reference subspace and positive rounding margin at β_0. Both readings are available directly from graph structure ($|Z|$ and $d_{\text{out}} - d_{\text{in}}$) without spectral computation.

G Discrete t-axis diagnostic

Purpose and result. The same theorem can certify discrete t-moves, but t has a natural unit spacing that cannot be refined the way α can. On both dense Stage-2 graphs the adjacent unit steps are too large for the certificate: 0/11 forward moves certify. PolBlogs nevertheless stays empirically unchanged, whereas Email-EU-Core does not, so the discrete certificate is sound but highly conservative at unit spacing. The protocol and tables below quantify this.

Reference §5. Fixed $\alpha = 0.1$; scan $t = 1, \ldots, 12$. For each adjacent pair $t \to t + 1$: $\epsilon_t := \|L^{\pi}_{t+1,0.1} - L^{\pi}_{t,0.1}\|_2$; source-point threshold $\tau^{\text{F}}_t := \gamma_t \bar{\delta}_t / (2\sqrt{2k} + \gamma_t)$ (this diagnostic uses the conservative Frobenius bound; the canonical τ_t is at most $\sqrt{2k}/\sqrt{2}\times$ larger and yields the same 0/11 verdict below); $q_t := \epsilon_t / \tau^{\text{F}}_t$; certified iff $q_t < 1$; anchored same checks whether the Procrustes-aligned nearest-centroid partition at $t{+}1$ (vs. reference at t) equals the reference; fresh km same uses Hungarian label-matching. Downward $2 \to 1$ uses τ_2. This is a conservatism diagnostic, not a positive test of "certified ⇒ preserved": at unit t-spacing no adjacent move certifies (0/11), so the certificate is (soundly) silent here rather than confirmed against observed preservation.

G.1 PolBlogs SCC ($k = 2$, $N = 793$, $\alpha = 0.1$)

step	ϵ_t^π	τ_t	q_t	cert.?	anch same?	fresh km?	AMI(t)/AMI(t+1)
$1 \to 2$	4.36×10^{-3}	1.49×10^{-7}	2.9×10^{4}	no	yes	yes	0.7874/0.7874
$2 \to 3$	1.10×10^{-3}	1.17×10^{-7}	9.5×10^{3}	no	yes	yes	0.7874/0.7874
$3 \to 4$	3.97×10^{-4}	9.51×10^{-8}	4.2×10^{3}	no	yes	yes	0.7874/0.7874
$4 \to 5$	4.70×10^{-4}	7.52×10^{-8}	6.3×10^{3}	no	yes	yes	0.7874/0.7874
$5 \to 6$	1.58×10^{-4}	7.00×10^{-8}	2.3×10^{3}	no	yes	yes	0.7874/0.7874
$6 \to 7$	1.68×10^{-4}	5.72×10^{-8}	2.9×10^{3}	no	yes	yes	0.7874/0.7874
$7 \to 8$	1.05×10^{-4}	5.71×10^{-8}	1.8×10^{3}	no	yes	yes	0.7874/0.7874
$8 \to 9$	1.37×10^{-4}	4.77×10^{-8}	2.9×10^{3}	no	yes	yes	0.7874/0.7874
$9 \to 10$	8.50×10^{-5}	4.84×10^{-8}	1.8×10^{3}	no	yes	yes	0.7874/0.7874
$10 \to 11$	1.14×10^{-4}	4.09×10^{-8}	2.8×10^{3}	no	yes	yes	0.7874/0.7874
$11 \to 12$	6.38×10^{-5}	4.18×10^{-8}	1.5×10^{3}	no	yes	yes	0.7874/0.7874
$2 \to 1$ (τ_2)	4.36×10^{-3}	1.17×10^{-7}	3.7×10^{4}	no	yes	yes	0.7874/0.7874

Among the eleven forward moves $t \to t+1$ ($t = 1, \ldots, 11$), 0/11 certified, 11/11 preserve the anchored partition, 11/11 preserve fresh k-means; the reverse move $2 \to 1$ is reported separately in the last row. The source-anchored certificate is conservative by $\sim 10^3$–$10^4\times$ relative to the observed operator moves at unit t-spacing, not relative to an unknown true partition-transition boundary.

G.2 Email-EU-Core SCC ($k = 40$, $N = 803$, $\alpha = 0.1$)

step	ϵ_t^π	τ_t	q_t	cert.?	anch same?	fresh km?	AMI(t)/AMI(t+1)
$1 \to 2$	6.65×10^{-4}	2.50×10^{-11}	2.7×10^{7}	no	no	no	0.55/0.56
$2 \to 3$	2.66×10^{-4}	1.31×10^{-10}	2.0×10^{6}	no	no	no	0.56/0.53
$3 \to 4$	9.17×10^{-5}	4.55×10^{-11}	2.0×10^{6}	no	no	no	0.53/0.52
$4 \to 5$	8.67×10^{-5}	3.99×10^{-11}	2.2×10^{6}	no	no	no	0.52/0.57
$5 \to 6$	4.74×10^{-5}	1.06×10^{-10}	4.5×10^{5}	no	no	no	0.57/0.56
$6 \to 7$	3.69×10^{-5}	2.55×10^{-10}	1.4×10^{5}	no	no	no	0.56/0.54
$7 \to 8$	2.63×10^{-5}	4.34×10^{-11}	6.1×10^{5}	no	no	no	0.54/0.56
$8 \to 9$	2.31×10^{-5}	3.56×10^{-11}	6.5×10^{5}	no	no	no	0.56/0.53
$9 \to 10$	1.61×10^{-5}	2.02×10^{-10}	8.0×10^{4}	no	yes	no	0.53/0.57
$10 \to 11$	1.36×10^{-5}	2.54×10^{-10}	5.4×10^{4}	no	yes	no	0.57/0.57
$11 \to 12$	9.32×10^{-6}	4.11×10^{-11}	2.3×10^{5}	no	yes	no	0.57/0.54
$2 \to 1$ (τ_2)	6.65×10^{-4}	1.31×10^{-10}	5.1×10^{6}	no	no	no	0.56/0.55

Among the eleven forward moves $t \to t+1$, 0/11 certified, 3/11 preserve the anchored partition, 0/11 preserve fresh k-means; the reverse move $2 \to 1$ is reported separately in the last row. The partition itself is not stable in t at $k = 40$; certificate soundness intact but the underlying phenomenon is absent.

H Bibliography note on GSC citation

The GSC framework we build on is Sevi et al. (2026), published in TMLR (https://openreview.net/forum?id=AA6D7fJ9PN, arXiv:2203.03221); its contribution is the generalized Dirichlet energy and the parametric family $\nu_{t,\alpha}$ we analyse.

I Full 42-point anchor tables (pointer)

Complete per-anchor Stage-2 quantities on the 42-point grid for PolBlogs and Email-EU-Core (the two CSV files scale_audit_merged.csv and projective_anchors.csv) and for Amazon Photo (stage2_sparse_AmazonPhoto.csv), plus the K-sweep and teleportation logs from Appendix C and the DC-SBM mechanism-validation outputs from Fig 8) are provided as supplementary material with the submission, listed in the checksum manifest and recheckable by the reproduction harness of Appendix B (reproduce.sh, MANIFEST.sha256); a stable-URL release with the corresponding scripts and solver/seed settings will accompany a de-anonymised camera-ready version. Key summary values appear in §4.1–§4.3.

J DC-SBM decomposition: full accounting

J.1 Experimental question and target

This appendix expands the mechanism-accounting result of §4.3. The key caution first: the interventions are not factor-isolating, so this experiment tests attribution, not causal isolation. The nominal $\bar{\delta}$- and G_k-levers each move all three factors, so the paired identity supports only the weaker statement that $\log \rho^{\pi}_{\alpha,\text{oracle}}$ decomposes additively across $(D^{\pi}_{\alpha}, \bar{\delta}, G_k)$ per seed (contributions summing to the observed $\Delta \log \rho^{\pi}_{\alpha,\text{oracle}}$ exactly); it does not establish that each intervention isolates its named factor. The remaining subsections give the generator and matched-seed protocol (J.2), the per-sweep attribution (J.3), why the survivor cohort changes (J.4), the finite-step-versus-first-order comparison (J.5), and the meaning of $\gamma \leq 0$ (J.6).

J.2 Generator and matched-seed protocol

The generator is a degree-corrected DiSBM with $N = 400$, $k = 4$, vendored into the repository (src/graphs/synthetic_4axis.py) so the decomposition re-runs from scratch under the harness (Appendix B); each sweep varies one nominal lever (node-measure heterogeneity σ_θ, community contrast $M_{\text{diag}}/M_{\text{off}}$, or block-density noise σ_{block}) over its grid while the other generator parameters are held fixed. A preregistered set of 30 seeds indexes the random draws (the original five are the first entries, so the earlier run is a subset); a seed fixes the same latent draw at both endpoints of a sweep, so the paired (matched-cohort) contrast is within-seed across the two endpoints. Seeds are not coupled by common random numbers across different levers, so cross-sweep comparisons are between independent draws. Full grids, generator parameters, and per-seed contributions are in the release CSVs.

J.3 Factor-attribution results

The headline number for each sweep is the matched-cohort paired fall in $\rho^{\pi}_{\alpha,\text{oracle}}$, the fall over exactly the seeds that survive ($\gamma > 0$, so $\rho^{\pi}_{\alpha,\text{oracle}}$ is defined) at both endpoints, so no seed enters one endpoint without the other; per-factor $\Delta \log$ contributions are in Table 7.

Table 7: Paired per-seed log-decomposition over 30 preregistered seeds; effective n = surviving seeds ($\gamma > 0$ at both endpoints); "±" is 1 sample standard deviation across surviving seeds. Bold marks the largest mean $|\Delta \log|$ contributor per sweep: the operator sweep is dominated by its nominal factor, but G_k dominates both the spectral (nominally $\bar{\delta}$) and geometry sweeps, so the interventions are not factor-isolating.

Sweep (nominal target)	eff. n	$-\Delta \log D^{\pi}_{\alpha}$	$\Delta \log \bar{\delta}$	$\Delta \log G_k$	$\Delta \log \rho^{\pi}_{\alpha,\text{oracle}}$	residual
Operator (D^{π}_{α})	30/30	$\mathbf{-0.895 \pm 0.21}$	-0.506 ± 0.13	-0.231 ± 0.31	-1.632 ± 0.50	$\sim 10^{-16}$
Spectral ($\bar{\delta}$)	17/30	-0.428 ± 0.16	-1.029 ± 0.14	$\mathbf{-2.136 \pm 0.89}$	-3.593 ± 0.95	$\sim 10^{-16}$
Geometry (G_k)	22/30	-0.271 ± 0.26	-0.487 ± 0.35	$\mathbf{-0.625 \pm 0.87}$	-1.383 ± 1.28	$\sim 10^{-16}$

- Operator sweep ($\sigma_\theta : 0 \to 0.5$; nominally D^{π}_{α}; $n = 30/30$ at both endpoints). Matched fall $e^{1.632} = 5.11\times$; dominant contributor D^{π}_{α} ($-\Delta \log D^{\pi}_{\alpha} = -0.895$), as designed, the one sweep whose nominal factor dominates.
- Spectral sweep ($M_{\text{diag}}/M_{\text{off}} : 15 \to 7$; nominally $\bar{\delta}$; $n = 30/30 \to 17/30$). Matched fall $e^{3.593} = 36.3\times$ over the 17 survivors; dominant contributor G_k ($\Delta \log G_k = -2.136$), not the nominal $\bar{\delta}$ (-1.029). Unlike the earlier 5-seed run (one survivor), this off-nominal dominance is now replicated across 17 seeds.
- Geometry sweep ($\sigma_{\text{block}} : 0 \to 0.5$; nominally G_k; $n = 30/30 \to 22/30$). Matched fall $e^{1.383} = 3.99\times$ over 22 seeds; dominant contributor G_k (-0.625), its nominal factor (edging out $\bar{\delta}$, -0.487).

Across the two non-operator sweeps the geometry factor G_k is the single dominant contributor, so $\bar{\delta}$ is never singly dominant despite being a nominal target, direct evidence that

the levers are not factor-isolating. We do not quote ratios of means over changing survivor sets: when the two endpoints have different cohorts (spectral, geometry) those ratios are not seed-matched and need not agree with the paired identity.

J.4 Why the survivor set changes

Because γ can cross zero on some seeds along a sweep, the per-anchor certified cohort ($\gamma > 0$ and $\rho^{\pi}_{\alpha,\mathrm{oracle}} > 0$) varies; the $h^{\star}_{\mathrm{cert}}$ curves in Fig 8 are means over that cohort. The spectral sweep's plotted markers run $15, 12, 10, 9, 8, 7$ with $n_{\mathrm{cert}} = 30, 30, 29, 29, 25, 17$, so the matched-endpoint cohort of the paired identity (contrast 15 to 7) is the 17 seeds surviving at both, matching Table 7. Sweep points where every seed abstains ($n_{\mathrm{cert}} = 0$) are omitted.

In the spectral sweep the largest $|\Delta \log|$ contribution (G_k) is not the nominal target $\bar{\delta}$ (Table 7), consistent with the non-isolation caution of J.1; whether cleaner isolation is achievable in this DC-SBM family or needs a redesigned generator (e.g. decoupling σ_{block} from $\bar{\delta}$) is left for follow-up work.

J.5 Finite step versus first-order radius

The measured $h^{\star}_{\mathrm{cert}}$ tracks $\eta\rho^{\pi}_{\alpha,\mathrm{oracle}}$ within $\sim 1\%$ (in $\Delta \log$) across every sweep point. On a purely affine operator path with $D^{\pi} > 0$ the certified-step supremum equals $h^{\mathrm{sup}}_{\mathrm{cert}} = \eta\tau / D^{\pi} = \eta\rho^{\pi}_{\alpha,\mathrm{oracle}}$ exactly; because the strict certificate $|h| D^{\pi} < \eta\tau$ has an open admissible set, the returned bisection value satisfies $h^{\star}_{\mathrm{cert}} < h^{\mathrm{sup}}_{\mathrm{cert}}$ and approaches it within the bisection tolerance (this ratio $\to \eta = 0.9$). If $D^{\pi} = 0$ the operator is first-order constant (on a genuinely affine path, constant altogether), so the certified step is limited only by the parameter domain. On the tested nonlinear paths, the proximity of $h^{\star}_{\mathrm{cert}} / (\eta\rho^{\pi}_{\alpha,\mathrm{oracle}})$ to one shows the finite-step scalar operator-distance boundary is close to the linearised scalar-norm benchmark; it does not quantify the matrix Taylor remainder $R(h) := L^{\pi}_{\alpha+h} - L^{\pi}_{\alpha} - h\,\partial_{\alpha} L^{\pi}_{\alpha}$. By the reverse triangle inequality the observable scalar discrepancy is only a lower bound on the remainder norm, $\big|\, \|\Delta L^{\pi}(h)\|_2 - |h| D^{\pi}_{\alpha} \,\big| \le \|R(h)\|_2$, so a small discrepancy does not upper-bound $\|R(h)\|_2$ (a large remainder can hide by cancellation or a direction change). Hence this is not independent validation of the factorisation.

J.6 What $\gamma \le 0$ means

Some seeds fall to $\gamma \le 0$ inside a sweep. Spectral identifiability is controlled by $\bar{\delta}$, not γ; $\gamma \le 0$ instead means the evaluated rounding partition is not a strict nearest-centroid partition, so Thm 1's hypothesis fails at that anchor. Concretely, the anchor uses the DC-SBM's planted labels rather than a converged same-anchor k-means partition; when embedding noise pushes a row past a rival centroid, planted-label γ becomes non-positive even though the fresh-k-means anchor A_{km} would still be nearest-centroid optimal. The certificate then abstains for that seed at that anchor; its guarantee is about the planted anchor, not the fresh k-means anchor.

K Visual diagnostics for the radius decomposition

Purpose. This appendix collects the visual diagnostics behind RQ3: Fig 6 separates the three radius factors on the real Stage-2 graphs, Fig 7 shows the quality–stability trade-off, and Fig 8 shows the synthetic mechanism sweeps and the finite-step comparison. The underlying values are tabulated in §4.3.

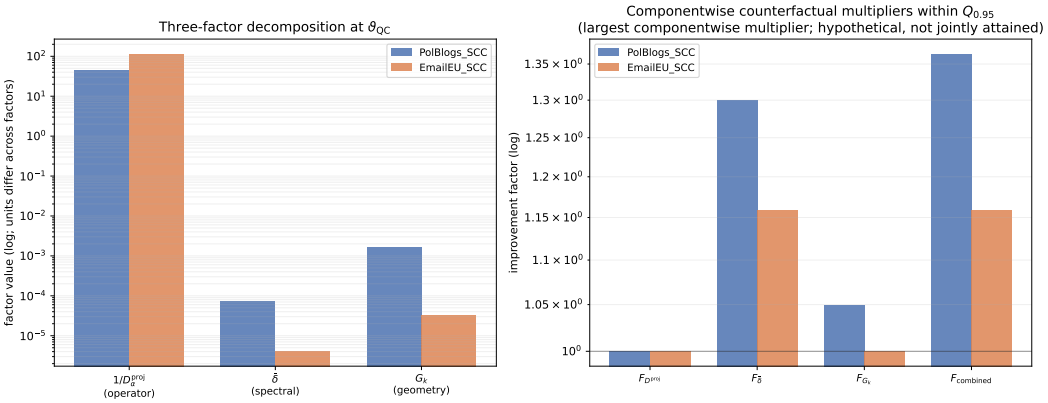

Figure 6: Three-factor decomposition of $\rho^{\pi}_{\alpha,\mathrm{oracle}}$ at ϑ_{QC} and componentwise counterfactual multipliers within $Q_{0.95}$. Spectral conditioning $\bar{\delta}$ is the largest componentwise multiplier on both dense Stage-2 graphs (a hypothetical single-factor replacement, not a jointly attained value); operator response D^{π}_{α} and geometric factor G_k are near-Pareto at the projective anchor. The jointly attained within-region headroom F^{joint} is smaller and is reported alongside in the componentwise-multiplier table below.

Componentwise counterfactual multipliers (§4.3). Each $F_\bullet$ replaces one factor at ϑ_{QC} by its best value across $Q_{0.95}$ (taken at its own anchor, so F_{combined} is a hypothetical joint product); $F^{\mathrm{joint}} := \max_{Q_{0.95}} \rho^{\pi}_{\alpha,\mathrm{oracle}}/\rho^{\pi}_{\alpha,\mathrm{oracle}}(\vartheta_{\mathrm{QC}})$ is the headroom actually attained at a single anchor.

Graph	$F_{D^{\pi}_{\alpha}}$	$F_{\bar{\delta}}$	F_{G_k}	F_{combined}	F^{joint}	Largest componentwise factor in $Q_{0.95}$
PolBlogs	1.00×	**1.30**×	1.05×	1.36×	1.30×	$\bar{\delta}$ (spectral)
Email-EU-Core	1.00×	**1.16**×	1.00×	1.16×	1.16×	$\bar{\delta}$ (spectral)

Figure 7 shows the quality-vs-stability Pareto structure of §4.3.

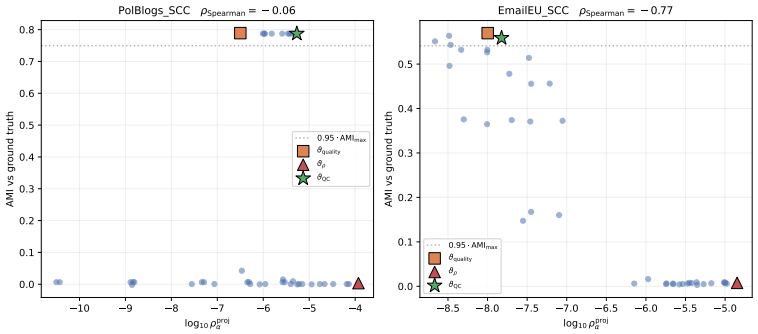

Figure 7: Quality vs. projective local stability radius (bound variant $\rho^{\pi}_{\alpha,\mathrm{bound}} = \tau/\Psi^{\pi}_{\alpha}$) over the 42-point grid. Near-zero rank association on PolBlogs (Spearman -0.06); strong negative association on Email-EU-Core (Spearman -0.77). Each anchor is a single deterministic evaluation. Anchor rules use $\rho^{\pi}_{\alpha,\mathrm{bound}}$: $\vartheta_{\mathrm{quality}} := \arg\max \mathrm{AMI}$; $\vartheta_{\rho} := \arg\max \rho^{\pi}_{\alpha,\mathrm{bound}}$; $\vartheta_{\mathrm{QC}} := \arg\max_{Q_{0.95}} \rho^{\pi}_{\alpha,\mathrm{bound}}$.

Figure 8 visualises the synthetic DC-SBM sweeps of §4.3: the first-order radius $\rho^{\pi}_{\alpha,\mathrm{oracle}}$ and the directly verified finite step $h^{\star}_{\mathrm{cert}}$ against each nominally targeted factor.

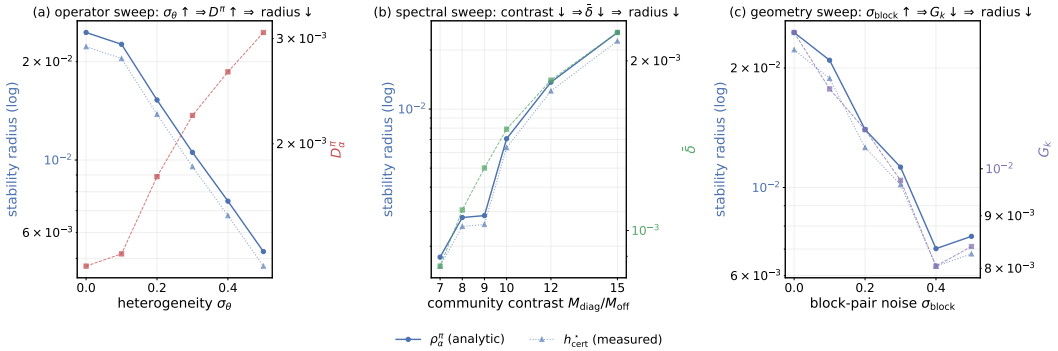

Figure 8: Accounting view of the three-factor decomposition on three DC-SBM interventions. Left-axis: estimated oracle first-order radius $\rho^\pi_{\alpha,\text{oracle}}$ (circles) and $h^\star_{\text{cert}}$ (triangles), a bisection estimate of the supremum of h satisfying $\|L^\pi_{\alpha+h} - L^\pi_\alpha\|_2 < \eta\,\tau$ ($\eta = 0.9$) under assumed operator-norm monotonicity on the bracket. Right-axis: the nominally targeted factor. (a) σ_θ (nominally D^π_α); (b) contrast (nominally $\bar{\delta}$; dominant paired contributor G_k); (c) σ_{block} (nominally G_k; dominant paired contributor G_k). $h^\star_{\text{cert}}$ tracks $\rho^\pi_{\alpha,\text{oracle}}$ within $\sim 10\%$ (the constant-factor gap forced by $\eta = 0.9$; this compares scalar operator norms and does not upper-bound the matrix Taylor remainder). Where Thm 1's hypotheses fail ($\gamma \le 0$ at some seeds) the certificate abstains there, so the certified cohort varies across the sweep (endpoint $n_{\text{cert}} = 17/30$ spectral, $22/30$ geometry, matching Table 7). Curves are arithmetic means over the certified cohort, whose membership changes along the sweep; the mean sequence is therefore not a fixed-cohort longitudinal trajectory, and apparent trends can reflect composition. Pointwise n_{cert} is annotated; per-seed variability is largest near $\gamma = 0$.