hexagon:2611.01315v1

Least signless Laplacian eigenvalues in graphs of odd girth at least seven

András Pongrácz
University of Debrecen, Institute of Mathematics
ORCID: 0000-0002-2771-8974

Abstract

We prove that every finite simple graph G on $n \geq 1$ vertices containing neither a triangle nor a pentagon satisfies $q_{\min}(G) < n/30$, where $q_{\min}(G)$ is the least eigenvalue of its signless Laplacian. Balanced blow-ups of the folded seven-cube attain $q_{\min}(G) = n/32$. The proof optimizes a unit-vector semidefinite minimum over vertex weights and reduces the problem to a weighted induced subgraph with constant local edge energy. This reduction gives independent-set mass constraints, which we combine with rooted Rayleigh inequalities, cubic walk functions, and positive semidefinite matrices of subgraph densities. Exact independent sampling expresses the resulting inequalities on a complete catalog of 1141 nine-vertex graphs and applies to graphs of every order. The accompanying integer certificates and Python verifier make the finite coefficient verification reproducible.

Mathematics Subject Classification (2020). 05C50; 05C35, 90C22.

Keywords. Signless Laplacian, odd girth, forbidden cycles, semidefinite programming, flag algebras, exact certificates.

1 Introduction

Let G be a finite simple graph with adjacency matrix A and diagonal degree matrix D. Its signless Laplacian is $Q = D + A$, and we denote its least eigenvalue by $q_{\min}(G)$. The variational identity

$$q_{\min}(G) = \min_{x \neq 0} \frac{\sum_{uv \in E(G)} (x_u + x_v)^2}{\sum_{v \in V(G)} x_v^2}$$

connects this parameter with bipartiteness. In particular, it vanishes exactly when G has a bipartite connected component. We study the largest possible ratio $q_{\min}(G)/|V(G)|$ when G contains neither C_3 nor C_5, equivalently when every odd cycle has length at least seven. Standard background on graph spectra is provided by Brouwer and Haemers [4].

Our main result is the following uniform bound.

Theorem 1.1. *Every finite simple graph G on $n \geq 1$ vertices containing neither C_3 nor C_5 satisfies*

$$q_{\min}(G) < \frac{n}{30}.$$

The proof is completed in Section 9. The folded seven-cube gives the coefficient $1/32$ for infinitely many orders; its construction and spectrum are recalled below. Thus the remaining gap

between the upper bound and this lower bound is a factor of 16/15. Whether $q_{\min}(G) \le n/32$ holds for all graphs in this class remains open.

For triangle-free graphs, Balogh, Clemen, Lidický, Norin, and Volec [2] proved the bound $15n/94$. Cai, Chen, and Zhang [5] subsequently established the sharp coefficient 7/50. In the smaller class considered here, Zhou [18, Theorem 1.1] proved the general bound $0.0569n$ by combining local Rayleigh inequalities with exact flag certificates. Our proof retains that spectral and counting framework and strengthens the inequalities available at an extremal vertex weighting.

The distinction between signless Laplacian and adjacency eigenvalues is relevant for this problem. Abiad, Taranchuk, and van Veluw [1] proved $\lambda_1(A) + \lambda_n(A) < 0.0396n$ under the same forbidden-cycle assumptions. This gives a signless Laplacian bound for regular graphs, since then $q_{\min} = \lambda_1 + \lambda_n$. It does not directly give the corresponding assertion for arbitrary irregular graphs. The folded seven-cube, which appears in their work, supplies the natural candidate coefficient 1/32. Zhou [17] established its optimality within the class of generalized odd graphs of diameter three. No such structural restriction is imposed here. The related asymptotic moment problem for adjacency spectra at growing odd girth is studied by Wang [16].

The proof uses two kinds of positivity. First, optimizing a unit-vector semidefinite minimum over vertex weights produces a weighted induced subgraph with constant local edge energy. A flow argument, adapted from Abiad, Taranchuk, and van Veluw, then bounds the weight of every independent set by one half. Second, conditional Rayleigh inequalities, ordinary flag squares, and inequalities for independent supports give positive semidefinite matrices of weighted subgraph densities. Cubic walk kernels provide test functions that are least-eigenvalue functions on the folded seven-cube. The use of a polynomial adapted to the proposed extremizer is related to the construction in [5].

The semidefinite viewpoint is connected with the Max-Cut relaxation of Goemans and Williamson [7]. Spectral comparisons with this relaxation also occur in Coutinho and Oliveira [6, Proposition 5]; their optimization over edge weights differs from the vertex-weight maximization used below. The counting argument uses the positive-square principle of flag algebras [15], but all identities are proved by finite independent sampling. The language of homomorphism densities is standard in graph limit theory [13]. Multiplying positive matrices by nonnegative density polynomials is a moment-localization technique; see Laurent's survey [11] for its broader context.

Minimum-degree hypotheses lead to a different structural approach. Letzter and Snyder [12] showed that an n-vertex $\{C_3, C_5\}$-free graph with minimum degree greater than $n/5$ is homomorphic to a generalized Andrásfai graph. The weighted extremal reduction below supplies positivity constraints without such a minimum-degree assumption on the original graph.

1.1 The folded seven-cube and the candidate constant

We include the spectral calculation to specify the lower bound and the walk functions used later. Let $\mathbb{F}_2^6$ be the six-dimensional vector space over the two-element field, let $e_1, \dots, e_6$ be its standard basis, and put $j = e_1 + \cdots + e_6$. Define F to have vertex set $\mathbb{F}_2^6$, joining x to $x + s$ for $s \in \{e_1, \dots, e_6, j\}$. This is the folded seven-cube.

Proposition 1.2. *The graph F is 7-regular on 64 vertices, has odd girth seven, and has least signless Laplacian eigenvalue 2. For every positive integer t, its balanced blow-up with t vertices in each class is $\{C_3, C_5\}$-free and satisfies $q_{\min} = n/32$, where $n = 64t$.*

Proof. The seven generators are distinct and nonzero. Their only nonempty subset with sum zero is the entire set: if j is absent, independence of the standard basis excludes a relation; if j is present, every basis vector must also be present. An odd closed walk of length less than seven would yield, after cancelling repeated generators in pairs, a nonempty odd subset with zero sum and fewer than seven elements. This is impossible. Using all seven generators gives a cycle of length seven, since no proper partial sum repeats.

For $z \in \mathbb{F}_2^6$ the character $x \mapsto (-1)^{z \cdot x}$ is an adjacency eigenvector. If z has k nonzero coordinates, its eigenvalue is $6 - 2k + (-1)^k$. The resulting eigenvalues are $7, 3, -1, -5$, and hence the least signless eigenvalue is $7 - 5 = 2$. These 64 characters form a basis: the sum of a nonconstant character is zero, so distinct characters are orthogonal.

In the balanced blow-up, the subspace of vectors constant on each class has signless eigenvalues $t(7 + \lambda)$ for adjacency eigenvalues λ of F. Its orthogonal complement consists of vectors whose coordinates sum to zero within every class; the adjacency operator vanishes there and the degree operator is $7tI$. Thus the least signless eigenvalue is $2t = n/32$. A forbidden odd cycle in the blow-up would project to a forbidden odd closed walk in F. □

At uniform weights on F, the weighted adjacency operator $T = A/64$ has spectrum

$$\operatorname{spec}(T) = \{7/64,\ 3/64,\ -1/64,\ -5/64\}.$$

The polynomial

$$p(s) = s(s - 3/64)(s + 1/64) = s^3 - s^2/32 - 3s/4096$$

annihilates the two middle eigenspaces. If Π is the orthogonal projection onto constants, then $p(T) - (7/8192)\Pi$ is nonzero and its image lies in the least adjacency eigenspace. Its kernel relative to the uniform probability measure, as a function of x for a fixed root v, is

$$h_v(x) = t_3(v, x) - \frac{1}{32} t_2(v, x) - \frac{3}{4096} t_1(v, x) - \frac{7}{8192},$$

where explicitly $t_k(v, x) = (A^k)_{vx}/64^{k-1}$; this is the kernel of T^k relative to the uniform probability measure. Consequently $((7/64)I + T)h_v = (1/32)h_v$. This explains the choice of cubic walk features; their Rayleigh inequalities remain valid for all weighted graphs.

2 Weighted operators and a semidefinite minimum

All graphs are finite, simple, and undirected. A graph is *admissible* if it contains neither a triangle nor a cycle of length five; these are not required to be induced subgraphs. Any odd closed walk contains an odd cycle of no greater length. Indeed, a closed walk with a repeated internal vertex splits into two shorter closed walks, one of which has odd length; repeating this operation ends at an odd cycle. A loopless graph has no closed walk of length one. Thus admissibility also excludes closed walks of lengths three and five, even when vertices repeat.

Let H have vertex set $[m] = \{1, \dots, m\}$ and adjacency matrix A. A probability weight vector is $w = (w_i)_{i=1}^m$ with $w_i \geq 0$ and $\sum_i w_i = 1$. Unless stated otherwise, zero-weight vertices are deleted. For positive weights put $W = \operatorname{diag}(w)$, $d = Aw$, and

$$\langle f, g \rangle_w = \sum_i w_i f(i) g(i), \qquad (Tf)(i) = \sum_j A_{ij} w_j f(j), \qquad (Df)(i) = d_i f(i).$$

The operators T and $\mathcal{Q} = D + T$ are self-adjoint on the real Hilbert space $L^2(w) = \mathbb{R}^m$ with this inner product. The map $f \mapsto W^{1/2} f$ is an isometry onto Euclidean space, under which $\mathcal{Q}$ becomes

$$Q(w) = \operatorname{diag}(d) + W^{1/2} A W^{1/2}.$$

Write $F_H(w)$ for its least eigenvalue. Its quadratic form is

$$\langle f, \mathcal{Q} f \rangle_w = \frac{1}{2} \sum_{i,j} A_{ij} w_i w_j (f(i) + f(j))^2. \tag{1}$$

In particular, $F_H(w) \geq 0$. At uniform weights, $F_H(\mathbf{1}/m) = q_{\min}(H)/m$, where $q_{\min}(H)$ is the least eigenvalue of the usual signless Laplacian of H. Matrix inequalities and trace identities below are over the reals; the Rayleigh principle and positive semidefinite Gram factorizations are used in their usual finite-dimensional form; see [9].

Let

$$\mathcal{E}_m = \{X \in \mathbb{R}^{m \times m} : X = X^{\mathsf{T}},\ X \succeq 0,\ X_{ii} = 1\ (i \in [m])\}$$

be the set of correlation matrices, and define

$$S_H(w) = \min_{X \in \mathcal{E}_m} \sum_{i,j} A_{ij} w_i w_j (1 + X_{ij}). \tag{2}$$

All sums over i, j are ordered sums. The set $\mathcal{E}_m$ is compact: it is closed, and its two-by-two principal minors give $|X_{ij}| \leq 1$. Consequently the minimum is attained and S_H is continuous on the closed probability simplex. The same formula defines S_H for arbitrary positive vectors w, without the normalization $\sum_i w_i = 1$.

This is the unit-vector semidefinite formulation associated with the Max-Cut relaxation of Goemans and Williamson [7]. Precisely, if H is unweighted and

$$\operatorname{SDPcut}(H) = \max_{X \in \mathcal{E}_m} \frac{1}{2} \sum_{\{i,j\} \in E(H)} (1 - X_{ij}),$$

then $S_H(\mathbf{1}/m) = 4(|E(H)| - \operatorname{SDPcut}(H))/m^2$. No rounding theorem is used here.

Lemma 2.1. *For every real symmetric matrix C of order m, the minimum $p = \min_{X \in \mathcal{E}_m} \operatorname{tr}(CX)$ is attained, and there is $y \in \mathbb{R}^m$ such that*

$$C - \operatorname{diag}(y) \succeq 0, \qquad \sum_i y_i = p.$$

For every minimizing X and every such y, $(C - \operatorname{diag}(y))X = 0$ and $y_i = (CX)_{ii}$. In particular, the maximizing dual vector y is unique.

Proof. We include the duality argument to specify the precise optimization fact needed later; it is also a special case of semidefinite strong duality and complementary slackness [3, Chapter 5]. For $b \in [0, \infty)^m$ define

$$\phi(b) = \min\{\operatorname{tr}(CZ) : Z \succeq 0,\ \operatorname{diag}(Z) = b\}.$$

Every feasible Z has the form $\operatorname{diag}(\sqrt{b}) X \operatorname{diag}(\sqrt{b})$ with $X \in \mathcal{E}_m$. For a zero coordinate of b, the corresponding row of Z is zero, and a correlation matrix on its positive support extends

by an identity block. This proves the assertion also at the boundary. Compactness of $\mathcal{E}_m$ now shows that ϕ is finite and continuous on the nonnegative orthant. It is convex, because convex combinations of feasible matrices are feasible for the corresponding combinations of diagonals.

The closed epigraph of ϕ has a supporting hyperplane at $(\mathbf{1}, \phi(\mathbf{1}))$. Since $\mathbf{1}$ is interior to its domain and ϕ is locally bounded there, the coefficient of the vertical coordinate in this hyperplane is nonzero. Normalizing it gives a vector y with

$$\phi(b) \geq \phi(\mathbf{1}) + y \cdot (b - \mathbf{1}) \quad (b \geq 0).$$

Here only the finite-dimensional supporting-hyperplane separation theorem is used. The identity $\phi(t\mathbf{1}) = t\phi(\mathbf{1})$ for $t > 0$, applied for t on both sides of 1, yields $y \cdot \mathbf{1} = \phi(\mathbf{1}) = p$. Thus $\operatorname{tr}(CZ) \geq y \cdot \operatorname{diag}(Z)$ for every $Z \succeq 0$. Taking $Z = zz^{\mathsf{T}}$ for arbitrary z proves $C - \operatorname{diag}(y) \succeq 0$.

If X is minimizing, the two positive semidefinite matrices $P = C - \operatorname{diag}(y)$ and X have $\operatorname{tr}(PX) = 0$. Since this trace is the squared Frobenius norm of $P^{1/2}X^{1/2}$, we have $PX = 0$. Its ith diagonal entry gives $(CX)_{ii} - y_i X_{ii} = 0$, hence $y_i = (CX)_{ii}$. Fixing one minimizing X determines every optimal dual vector by this formula, proving uniqueness. □

Lemma 2.2. *On the positive orthant, S_H is differentiable. If X minimizes (2), then*

$$\frac{\partial S_H}{\partial w_i} = 2d_i + 2\sum_j A_{ij} w_j X_{ij}. \tag{3}$$

The right-hand side is the same for all minimizing matrices X.

Proof. Put $C = WAW$. The part of (2) that depends on X is $\operatorname{tr}(CX)$. By Lemma 2.1, all minimizing matrices have the same values $(CX)_{ii} = y_i$. The gradient of the objective for a fixed X is $2d_i + 2(CX)_{ii}/w_i$, proving that all active gradients coincide.

For completeness, this equality of gradients implies differentiability without a differentiability assumption on the optimizer. Denote the objective by $g(w, X)$ and its common active gradient by v. For a sufficiently small vector h, choose an optimizer X_h at $w + h$. Compactness and continuity imply that every limit point of X_h as $h \to 0$ is an optimizer at w. Therefore $\nabla_w g(w, X_h) = v + o(1)$, uniformly over these choices of optimizers. Otherwise a sequence violating this assertion would have a convergent subsequence contradicting equality of the active gradients. The polynomial Taylor remainder for $g(w+h, X)$ is $O(\|h\|^2)$ uniformly on the compact set $\mathcal{E}_m$. Using a fixed optimizer at w gives an upper bound $S_H(w+h) \leq S_H(w) + v \cdot h + O(\|h\|^2)$. Using $g(w, X_h) \geq S_H(w)$ gives the reverse bound with error $o(\|h\|)$. Together these prove the assertion. □

Theorem 2.3 (Weighted extremal reduction)**.** *For every finite simple graph H,*

$$\max_w F_H(w) = \max_w S_H(w),$$

where w ranges over probability vectors and $F_H(w)$ is evaluated on the positive support of w. At a maximizing vector for S_H, after zero-weight vertices are deleted, there are positive weights w, a number $q = F_H(w) = S_H(w)$, and a correlation matrix X such that

$$Q(w) \succeq qI, \qquad E_{ij} = A_{ij}(1 + X_{ij}) \geq 0, \qquad Ew = q\mathbf{1}. \tag{4}$$

In particular, for every unweighted H such a weighted induced subgraph has $q \geq q_{\min}(H)/|V(H)|$.

Proof. First $F_H(w) \le S_H(w)$ for every probability vector. Indeed, for any $X \in \mathcal{E}_m$, the matrix $R = W^{1/2} X W^{1/2}$ is positive semidefinite with trace one, and

$$\operatorname{tr}(Q(w)R) = \sum_{i,j} A_{ij} w_i w_j (1 + X_{ij}).$$

The Rayleigh principle gives $F_H(w) \le \operatorname{tr}(Q(w)R)$, and minimization proves the claim. This argument applies on the positive support when some weights are zero.

Continuity gives a maximizing vector w for S_H. Restriction to its positive support changes neither its objective nor its maximality on that face: every correlation matrix on the support extends by an identity block. We may therefore assume $w_i > 0$ and work at an interior maximum on its probability simplex. Lemma 2.2 and the feasible directions $e_i - e_j$ imply that all partial derivatives have a common value γ. Multiplying (3) by w_i and summing shows $\gamma = 2S_H(w)$. Put $q = S_H(w)$. Lemma 2.1 gives

$$y_i = w_i(q - d_i), \qquad C - \operatorname{diag}(y) = W^{1/2}(Q(w) - qI)W^{1/2} \succeq 0.$$

Invertibility of $W^{1/2}$ implies $Q(w) \succeq qI$, so $F_H(w) \ge q$. Combined with $F_H(w) \le S_H(w) = q$, this proves equality. It also proves equality of the two maxima, without requiring continuity of F_H across faces of the simplex. Finally, (3) with common value $2q$ gives $Ew = q\mathbf{1}$. The inequality $E_{ij} \ge 0$ follows from $|X_{ij}| \le 1$. □

3 Independent supports at an extremal weight vector

The next argument adapts the weighted edge counting in Abiad, Taranchuk and van Veluw [1, Lemma 9]. Their statement uses squared Perron-vector weights of an adjacency matrix. Here the symmetric flow $w_i w_j E_{ij}$ supplied by (4) has marginals qw_i, so the same counting argument applies without assuming ordinary degree regularity or connectedness.

Lemma 3.1. *Suppose $w > 0$, $\sum_i w_i = 1$, $q > 0$, and a symmetric nonnegative matrix E is supported on the edges of H and satisfies $Ew = q\mathbf{1}$. For every independent set I of H,*

$$w(I) \le w(N_H(I)) \le 1 - w(I), \qquad w(I) \le \tfrac{1}{2}.$$

Here $w(S) = \sum_{i \in S} w_i$ and $N_H(I)$ is the union of the neighborhoods of vertices in I.

Proof. Writing $J = N_H(I)$, symmetry and nonnegativity give

$$qw(I) = \sum_{i \in I, j \in J} w_i w_j E_{ij} \le \sum_{j \in J, i \in [m]} w_j w_i E_{ji} = qw(J).$$

Since I is independent, $I \cap J = \varnothing$ and $w(J) \le 1 - w(I)$. Division by $q > 0$ proves both conclusions. □

For an admissible graph let $N_v = N_H(v)$ and put

$$P_v = \{x : \text{there is a walk of length two from } v \text{ to } x\}.$$

The set P_v includes v whenever v has a neighbor. The sets N_v and P_v are disjoint and independent: an intersection gives a closed walk of length three, an edge inside N_v gives a triangle, and

an edge inside P_v gives a closed walk of length five. Also $N_H(N_v) = P_v$. Consequently at a core from Theorem 2.3 with $q > 0$,

$$q \le d_v = w(N_v) \le p_v := w(P_v) \le \tfrac{1}{2}. \tag{5}$$

The first inequality is the diagonal-entry consequence of $Q(w) \succeq qI$; the other inequalities follow from Lemma 3.1. These restrictions concern an extremal weighted induced subgraph, not arbitrary weights on the original graph.

4 A density restriction from explicit unit vectors

Proposition 4.1. *For every admissible graph with probability weights, put $a = \sum_{i,j} A_{ij} w_i w_j$. For every $1 < u \le 2$,*

$$S_H(w) \le f_u(a) := (1 - 2/u)a + \frac{(1 + 2/u)a}{1 + u^3(u + 2)a}. \tag{6}$$

Proof. We construct unit vectors from the disjoint neighborhood classes, then bound their total edge energy using only a. Put

$$s = u + 1, \quad R = u^2 + u - 1, \quad A_0 = R^2 - 1, \quad B_0 = s^2 - 1, \quad k = A_0 + B_0 = u^3(u + 2).$$

Both A_0 and B_0 are positive. For each vertex i, define a vector $f_i \in L^2(w)$ with values $-R$ on N_i, s on P_i, and -1 elsewhere. Its squared norm is $Z_i = 1 + A_0 d_i + B_0 p_i \ge 1$. If ij is an edge, then

$$N_i \subseteq P_j, \quad N_j \subseteq P_i, \quad N_i \cap N_j = P_i \cap P_j = \varnothing.$$

The final intersection would produce a closed walk of length five; the other assertions follow from a walk of length two or from the absence of triangles. Summing the coordinate products over $N_i, N_j, P_i \setminus N_j, P_j \setminus N_i$ and their complement gives

$$\langle f_i, f_j \rangle_w = 1 - s(R - 1)(d_i + d_j) - (s + 1)(p_i + p_j) = 1 + 2/u - (Z_i + Z_j)/u.$$

For the last equality use $A_0/u = s(R - 1)$ and $B_0/u = s + 1$. The vectors $f_i/\sqrt{Z_i}$ have a feasible correlation matrix. The arithmetic–geometric mean inequality therefore bounds their edge contribution by

$$1 + \frac{\langle f_i, f_j \rangle_w}{\sqrt{Z_i Z_j}} \le 1 - 2/u + \frac{1 + 2/u}{\sqrt{Z_i Z_j}}.$$

On an edge, $p_i \ge d_j$ and $p_j \ge d_i$, whence

$$\begin{aligned} Z_i Z_j &\ge (1 + A_0 d_i + B_0 d_j)(1 + A_0 d_j + B_0 d_i) \\ &\ge (1 + k d_i)(1 + k d_j). \end{aligned}$$

The difference in the second inequality is $A_0 B_0 (d_i - d_j)^2$. Applying the arithmetic–geometric mean inequality to reciprocals and then summing the ordered edge weights gives

$$\sum_{i,j} \frac{A_{ij} w_i w_j}{\sqrt{Z_i Z_j}} \le \sum_i \frac{w_i d_i}{1 + k d_i} \le \frac{a}{1 + ka}.$$

The last step is Jensen's inequality for the concave function $t \mapsto t/(1+kt)$ on $[0, \infty)$. Substitution in the preceding edge-energy estimate proves (6). □

Corollary 4.2. *At an admissible weighted core with $q \geq 1/30$, the ordered edge density satisfies*

$$\frac{1}{28} \leq a < \frac{77}{200}.$$

Proof. Since $q = S_H(w)$, Proposition 4.1 at $u = 2$ gives $q \leq 2a/(1+32a)$. Thus $q \geq 1/30$ implies $1 + 32a \leq 60a$. For $u = 31/20$, direct rational evaluation gives

$$f_u(77/200) = \frac{1287125609}{38973479400} < \frac{1}{30}, \qquad f_u'(77/200) = -\frac{8679216829741951}{37973302413555609} < 0.$$

The first strict inequality has difference $11990371/38973479400$. Differentiating (6) gives $f_u''(a) = -2k(1 + 2/u)/(1 + ka)^3 < 0$. Hence f_u' stays negative to the right of $77/200$ and $f_u(a) < 1/30$ throughout that range, contradicting $q \leq f_u(a)$ if $a \geq 77/200$. □

5 Exact finite sampling, including repeated classes

We use a finite form of the averaging and positive-square principles underlying flag algebras [15]; spectral rooted inequalities in the present problem also appear in Zhou [18]. The following argument establishes the identities directly, without a limiting graph sequence or a distinctness assumption on the sampled base vertices.

Fix a positive integer t. Independently sample $X_1, \ldots, X_t$ from the vertices of (H, w). On the distinct formal positions $[t]$ put an edge ij exactly when $A_{X_iX_j} = 1$. Denote the resulting random simple graph by J_t. If two sampled classes coincide, their formal positions are still different and are nonadjacent. An odd cycle of length three or five in J_t would give an odd closed walk of the same length in H. Thus J_t is admissible whenever H is admissible.

Lemma 5.1. *Let $\mathcal{H}_t$ contain one representative of every admissible isomorphism class of order t, and let $y_J = \Pr(J_t \cong J)$. For any real function Ψ of a labeled simple graph on $[t]$,*

$$\mathbb{E}\Psi(J_t) = \sum_{J \in \mathcal{H}_t} y_J \frac{1}{t!} \sum_{\pi \in S_t} \Psi(J^\pi), \qquad y_J \geq 0, \quad \sum_J y_J = 1. \tag{7}$$

Here J^π is the adjacency matrix of J pulled back along the bijection π from the formal positions to its vertices.

Proof. The law of $(X_1, \ldots, X_t)$ is invariant under all permutations of its positions. Conditional on an unlabeled isomorphism class of positive probability, all labeled realizations are therefore equally likely. Every labeled realization occurs exactly $|\operatorname{Aut}(J)|$ times in the sum over S_t. Conditional averaging proves (7); classes of probability zero contribute zero. □

For a graph F on at most t formal vertices, define its weighted homomorphism density to be the expectation of the product of its required edge indicators. Isolated vertices have no effect on this expectation. Padding F to order t, Lemma 5.1 gives its coefficient on J as the number of edge-preserving bijections from the padded F to J, divided by $t!$. Nonedges impose no conditions in this statement. Products of required edge and required nonedge indicators have the same averaging formula, and describe induced flags.

To spell out positivity, fix a labeled root graph on k positions and a finite collection of rooted induced flags with s free positions. Conditional on the root classes, let p be the vector of the corresponding flag probabilities, taking $p = 0$ when the root type is wrong. Sampling

two independent free groups gives the matrix $\mathbb{E}(pp^{\mathsf{T}}) \succeq 0$ on $k+2s$ formal positions. The independence is conditional on the classes, and remains true when sampled classes in different groups coincide.

For the spectral version, conditional on root classes let $z(v,x)$ be a vector of real functions of an output class x; the parameter v denotes the entire sampled root tuple. If $Q(w) \succeq rI$, then

$$M_{ij} = \mathbb{E}_v\big[\langle z_i(v,\cdot), \mathcal{Q} z_j(v,\cdot)\rangle_w - r\langle z_i(v,\cdot), z_j(v,\cdot)\rangle_w\big], \qquad M \succeq 0. \tag{8}$$

Indeed, contraction with a real vector is the Rayleigh inequality for its linear combination of the coordinates, averaged over roots. The scalar energy is exactly (1), with two independent output classes. An output class may equal a root class; the function need not vanish there. If these functions are conditional flag probabilities or rooted homomorphism densities, independent private samples expand their products into indicators on distinct formal positions, so Lemma 5.1 applies exactly.

Coefficients of smaller order lift without approximation. If $b(F)/s!$ is an order-s coefficient and $s \le t$, its order-t numerator is

$$(t-s)! \sum_{\substack{U \subseteq V(J) \\ |U|=s}} b(J[U]). \tag{9}$$

To verify the formula, average over a uniformly chosen s-subset of the t formal positions. Its classes still have the original independent law, and $t!/(\binom{t}{s}s!) = (t-s)!$. This is a statement about positions, not an injective sample of base classes.

6 One positive semidefinite interval relaxation

Lemma 6.1. *Let $a \le b$ be real numbers. Suppose symmetric matrices, possibly depending on the same weighted graph, satisfy $M = M_0 + aH \succeq 0$ and $H \succeq 0$. Then*

$$M_0 + bH = M + (b-a)H \succeq 0.$$

Moreover, for every $X \succeq 0$,

$$\operatorname{tr}\big(X(M_0 + bH)\big) = \operatorname{tr}(XM) + (b-a)\operatorname{tr}(XH) \ge 0.$$

Proof. The first assertion is the sum of two positive semidefinite matrices. For the second, $\operatorname{tr}(XH) = \operatorname{tr}(X^{1/2}HX^{1/2}) \ge 0$, and the same argument applies to M. □

This lemma permits either of two identical presentations: use the relaxed matrix directly, or use its actual Rayleigh matrix together with a nonnegative upper-density multiplier. The latter has order nine in the applications below, whereas the former uses patterns of order at most eight. There is no change in the underlying inequality.

6.1 Root-centered walk kernels

Continue with the weighted operators from Section 2. Write $g_j = T^j\mathbf{1}$ for $j \ge 0$, $m_j = \langle \mathbf{1}, T^j\mathbf{1}\rangle_w$, and $a = m_1$. Self-adjointness gives $\langle g_i, g_j\rangle_w = m_{i+j}$. For $k \ge 1$ let $t_k(v,x)$ be the kernel of T^k

relative to w: $(T^k f)(v) = \sum_x w_x t_k(v,x) f(x)$. Put $h_k(v,x) = t_k(v,x) - g_k(v)$; its weighted mean in x is zero. Consider the seven functions of the output x, indexed by a root v,

$$g_0(v), g_1(v), g_2(v), g_3(v), h_1(v,x), h_2(v,x), h_3(v,x).$$

Let B, K be their root-averaged norm and $\mathcal{Q}$-energy matrices, as in (8). Index the first group by $i, j \in \{0,1,2,3\}$ and the second by $k, l \in \{1,2,3\}$. Their entries are

$$\begin{aligned} B_{ij} &= m_{i+j}, & K_{ij} &= 2a m_{i+j}, \\ B_{ik} &= 0, & K_{ik} &= 2m_{i+k+1} - 2a m_{i+k}, \\ B_{kl} &= \operatorname{tr}(T^{k+l}) - m_{k+l}, & K_{kl} &= \operatorname{tr}(DT^{k+l}) + \operatorname{tr}(T^{k+l+1}) - 4m_{k+l+1} + 2a m_{k+l}. \end{aligned} \tag{10}$$

For the first row use $\langle \mathbf{1}, \mathcal{Q}\mathbf{1} \rangle_w = 2a$ and $\langle g_i, g_j \rangle_w = m_{i+j}$. The second row follows from the zero output mean and $\mathcal{Q}\mathbf{1} = 2d$: specifically, $\langle \mathbf{1}, \mathcal{Q} h_k(v,\cdot) \rangle_w = 2g_{k+1}(v) - 2a g_k(v)$. For the final row let $\Pi = \mathbf{1} \otimes \mathbf{1}$ be the orthogonal projection onto the constants. The kernel operator of h_k is $T^k(I - \Pi)$. Cyclicity of trace gives

$$B_{kl} = \operatorname{tr}((I-\Pi)T^{k+l}), \qquad K_{kl} = \operatorname{tr}((I-\Pi)\mathcal{Q}(I-\Pi)T^{k+l}).$$

Expanding the second expression, each of its two mixed terms is $2m_{k+l+1}$, since $\mathcal{Q}\mathbf{1} = 2T\mathbf{1}$; its final term is $2a m_{k+l}$. This proves all the stated entries.

In (10), the coefficient of $2a$ in K is

$$G = \mathbb{E}_v\, \xi(v)\xi(v)^{\mathsf{T}} \succeq 0, \qquad \xi(v) = (g_0(v), g_1(v), g_2(v), g_3(v), -g_1(v), -g_2(v), -g_3(v)).$$

Thus replacing a by an upper endpoint b adds precisely $2(b-a)G$ to $K - rB$. Lemma 6.1 proves the resulting matrix inequality whenever $Q(w) \succeq rI$ and $a \le b$. The three-function family (h_1, h_2, h_3) and the four-function family (g_0, h_1, h_2, h_3) are principal submatrices of the same construction. This covers all three sizes without a separate positivity argument.

The moments m_j are densities of paths on $j+1$ vertices. The traces $\operatorname{tr}(T^s)$ and $\operatorname{tr}(DT^s)$ are densities of a cycle and a cycle with an attached leaf; for $s = 2$ the repeated edge reduces these to the one-edge graph and the three-vertex path. The exact products $a m_6$ use nine formal positions. Replacing a by a fixed upper endpoint leaves only patterns on at most eight positions. The explicit Rayleigh-plus-localizer interpretation above nevertheless expresses the same relaxed inequality exactly on the order-nine catalog.

6.2 Adding one root-only cubic feature

Let $z_F(v,x)$ range over all two-rooted admissible graphs F on four formal vertices, with roots v, x and two private sampled classes; z_F is the expectation of the product of their edge indicators. There are 25 such features up to permutations of the private vertices. Indeed, inclusion–exclusion on the four possible triangles gives $64 - 4\cdot 8 + 6\cdot 2 - 4 + 1 = 41$ admissible labeled graphs on four vertices. Exactly nine are fixed by swapping the private vertices: two have the private edge, four have neither the private nor the root edge, and three have the root edge but not the private edge. The orbit count is therefore $(41+9)/2 = 25$. Append the root-only function $g_3(v) = T^3\mathbf{1}(v)$, constant in x. The norm and energy matrices of these 26 functions satisfy $K - rB \succeq 0$ whenever $Q(w) \succeq rI$.

All entries involving the original features use at most eight positions. For a mixed entry, the four-vertex feature and the four-vertex rooted path representing g_3 share their root, using at

most seven positions; an energy term adds at most one further position. The appended diagonal is

$$B_{26,26} = m_6, \qquad K_{26,26} = 2am_6.$$

It is the only potentially nine-position entry. Replacing this energy diagonal by $2bm_6$ adds $2(b-a)m_6 e_{26}e_{26}^{\mathsf{T}} \succeq 0$, because $m_6 = \|T^3\mathbf{1}\|_w^2 \geq 0$. Again this is exactly Lemma 6.1, and its contribution to a Gram contraction can instead be treated as an upper-density localizer. No regularity of the graph or equality of vertex weights is needed.

7 The coefficient criterion

Proposition 7.1 (Coefficient criterion)**.** *Let $r \geq 1/30$, let $[\ell, b]$ be a closed interval, and let t be a positive integer. Suppose $\Phi(H, w) \geq 0$ for every admissible core from Theorem 2.3 satisfying $q \geq r$ and $\ell \leq a \leq b$. Suppose further that exact sampling on t formal positions gives*

$$\Phi(H, w) = \sum_{J \in \mathcal{H}_t} \gamma_J y_J, \qquad \gamma_J < 0 \quad (J \in \mathcal{H}_t).$$

Then no such core has density in $[\ell, b]$. If finitely many intervals with such certificates cover $[1/28, 77/200]$, then every nonempty admissible graph G satisfies $q_{\min}(G) < r|V(G)|$.

Proof. For a core in the stated interval, nonnegativity gives $\Phi \geq 0$. Lemma 5.1 gives $y_J \geq 0$ with sum one, so the displayed expansion gives $\Phi \leq \max_J \gamma_J < 0$, a contradiction. If an unweighted graph violated the final conclusion, Theorem 2.3 would provide an admissible core with $q \geq r$. Corollary 4.2 places its density in the required covered interval, where one of the assumed certificates gives the contradiction. □

Typical terms in Φ are trace pairings of positive semidefinite matrices with Gram factors, and nonnegative scalar or matrix multiples of $a - \ell$, $b - a$, or $(a - \ell)(b - a)$. The proposition requires both an exact coefficient identity and complete interval coverage. Verifying a single proper subinterval establishes only its conditional conclusion.

8 Coordinates and positive matrices

We now specify the matrix inequalities whose nonnegative trace pairings enter Proposition 7.1. Throughout this section, (H, w) is an admissible core from Theorem 2.3, $q = F_H(w) > 0$, and $q \geq r$. In particular,

$$\mathcal{Q} \succeq rI, \qquad r \leq d_v \leq \tfrac{1}{2}, \qquad w(N_v), w(P_v) \leq \tfrac{1}{2}.$$

An interval certificate also assumes $\ell \leq a \leq b$, where $a = \langle \mathbf{1}, T\mathbf{1} \rangle_w$. All samples mentioned below are independent with law w, except for positions explicitly identified as shared. Distinct formal positions are not required to have distinct sampled classes.

8.1 A fixed enumeration convention

For a graph F on the formal positions $\{0, \dots, n-1\}$, order the pairs (i, j) with $i < j$ lexicographically and assign them indices $e = 0, \dots, \binom{n}{2} - 1$. Its integer code is

$$\operatorname{code}(F) = \sum_{e=0}^{\binom{n}{2}-1} 2^e \mathbf{1}_{\{\text{the pair indexed by } e \text{ is an edge}\}}.$$

A *root type* on k positions is an admissible graph whose code is minimum over all permutations of those k positions. Fix such a type τ. A rooted induced flag with s private positions is an admissible graph on $k+s$ positions restricting to τ on $0,\dots,k-1$. Its representative is the minimum code under permutations of the private positions $k,\dots,k+s-1$, with every root fixed individually. The coordinate list $\mathcal{F}(\tau,s)$ consists of these representatives in increasing code order. It is obtained by enumerating the $2^{\binom{k+s}{2}-\binom{k}{2}}$ choices of the remaining edges, rejecting inadmissible choices, and taking these minima.

This convention labels roots even when the root graph has automorphisms. A flag coordinate records an isomorphism class only with respect to its private positions; no division by a root automorphism group is implicit. All factors in a coefficient certificate refer to the full coordinates so defined. A factor may have fewer columns than rows, without imposing any rank condition on the positive matrix it multiplies.

8.2 Ordinary flags and one-output spectral flags

For an ordered root tuple $R=(R_0,\dots,R_{k-1})$, let $\iota_\tau(R)$ be the indicator that its induced graph on the k formal positions is exactly τ. For $F\in\mathcal{F}(\tau,s)$ define

$$p_F(R)=\iota_\tau(R)\Pr\big((R,Y_1,\dots,Y_s)\text{ induces the rooted flag }F\mid R\big).$$

Here $Y_1,\dots,Y_s$ are private samples. Conditional independence of two such groups gives the ordinary matrix

$$O_{\tau,s}=\mathbb{E}_R\big[p(R)p(R)^{\mathsf T}\big]\succeq 0. \tag{11}$$

Its entries use $k+2s$ formal positions. The possible ordinary blocks appearing in the present certificates are the following; a row referring to all types means all admissible root types of that order.

k	s	$k+2s$	root types and coordinate dimensions
1	2	5	one type, dimension 5
3	1	5	three types, dimensions $8,6,5$
1	3	7	one type, dimension 12
3	2	7	three types, dimensions $50,31,23$
5	1	7	all thirteen types
1	4	9	one type, dimension 34
3	3	9	three types, dimensions $293,158,106$
5	2	9	all thirteen types
7	1	9	$3K_2\cup K_1$, dimension 54

The five-position matrices in the first two rows occur multiplied by density localizers specified below. The three root types of order three have codes $0,1,3$, namely the edgeless graph, one edge, and a path of length two. The displayed seven-root type has code 2320; its one-output flags choose at most one endpoint of each matching edge and may choose the isolated root, giving $3^3\cdot 2=54$ coordinates. The other dimensions are determined by the explicit finite enumeration in Section 8.1; they are not assumptions in a positivity argument.

For spectral flags take $s=1$ and let $e(R,x)$ be the vector with a one in the flag coordinate induced by the output class x, multiplied by $\iota_\tau(R)$. Thus it is zero for a wrong root type. Set

$$\begin{aligned}B_\tau&=\mathbb{E}_{R,x}\,e(R,x)e(R,x)^{\mathsf T},\\ K_\tau&=\frac{1}{2}\mathbb{E}_{R,x,y}A_{xy}(e(R,x)+e(R,y))(e(R,x)+e(R,y))^{\mathsf T}.\end{aligned} \tag{12}$$

The Rayleigh principle gives $K_\tau - rB_\tau \succeq 0$, conditional on R and then averaged. Both matrices have expansions on $k+2$ positions; for B_τ the second output position is unused. Equivalently its integrand can be symmetrized as $(e(R,x)e(R,x)^{\mathsf{T}} + e(R,y)e(R,y)^{\mathsf{T}})/2$. This convention fixes the factors of two in the integer coefficients. The certificates use $k = 3, 5, 6$; the same definition is valid for $k = 7$ within the nine-position bound. The three-root spectral blocks occur with quadratic density localization. The five-root and six-root blocks use all admissible types of the relevant order. No value of $e(R,x)$ is suppressed merely because x equals one of the root classes.

8.3 Rooted homomorphism Rayleigh matrices

The preceding spectral construction extends to arbitrary rooted homomorphism features. Fix k parameter-root positions, one output position, and s private positions. A feature is the conditional expectation of a product of required edge indicators, with the parameter roots and output held fixed and the private classes sampled independently. One may also multiply every feature by an induced root-type indicator. Unlike an induced flag, such a product imposes no conditions on other edges. For a feature vector $z(R,x)$ define

$$\begin{aligned} B_{ij} &= \mathbb{E}_R \langle z_i(R,\cdot), z_j(R,\cdot)\rangle_w, \\ K_{ij} &= \mathbb{E}_R \langle z_i(R,\cdot), \mathcal{Q} z_j(R,\cdot)\rangle_w. \end{aligned} \tag{13}$$

Equation (8) gives $K - rB \succeq 0$. The norm entries use $k+1+2s$ positions. A degree-energy term adds a private neighbor of the output; an adjacency-energy term uses two outputs joined by an edge. Both energy terms use at most $k+2+2s$ positions. Fixed linear combinations of the features are allowed, because they give a congruence of these matrices.

The principal family takes $(k,s) = (1,3)$: its two distinguished positions are the parameter root $v = 0$ and output $x = 1$, and its private positions are $2, 3, 4$. Enumerate every admissible graph on these five positions and identify permutations of the three private positions, fixing $0, 1$. In increasing representative-code order this gives 89 features $z_F(v,x)$. This finite list is obtained by testing the 1024 edge subsets and all six private-position permutations. Isolated positions are retained, so features depending on fewer classes are included.

There is a particularly direct coefficient construction. For two features F, G, the B entry identifies their parameter roots and their outputs, while giving the two features disjoint private positions. For the degree part of K, add an edge from the common output to one new private position. For the adjacency part, identify only the parameter roots and join the two output positions by an edge. These use at most eight, nine, and nine positions, respectively. Repeated required edges are retained only once, because $A_{uv}^2 = A_{uv}$.

A conditional three-root family uses $(k,s) = (3,2)$. Its parameter roots are positions $0, 1, 2$, its output is position 3, and its private positions are $4, 5$. The root type has code $0, 1$, or 3. Each coordinate is specified by a six-position graph code in the certificate, in the order of its factor rows, and is the corresponding edge-product expectation multiplied by $\iota_\tau(R)$. An arbitrary finite list of such features is legitimate; no completeness of this list is needed for positivity. The norm gluing uses eight positions and either energy gluing uses nine. Required root edges are inserted once, while the at most three forbidden root edges are imposed by expanding $\prod_{ij \notin E(\tau)} (1 - A_{R_i R_j})$. A factor initially given on linear combinations of these coordinates can be expanded into their original coordinates before it is stored; equation (13) then applies without a separate projection argument.

The induced conditional version instead takes every admissible graph on these same six positions with the prescribed three-root type. Canonicalize only the swap of private positions

$4, 5$, keeping all four distinguished positions fixed, and list the representatives in increasing code order. This gives $772,407,268$ coordinates for root types $0, 1, 3$, respectively. A coordinate is the conditional probability of the corresponding induced six-position flag, given the parameter roots and output, multiplied by $\iota_\tau(R)$. Equation (13) and its Rayleigh proof apply to these probability functions directly. Two independent private pairs give eight positions for their norm products and nine for their energy products. The coefficients can therefore be reconstructed by finite induced-flag sampling, without changing to a monomial basis.

Several smaller families are subfamilies or fixed linear combinations of the 89-coordinate family. The 25 two-rooted four-position features are padded by an isolated private position. Appending the root-only function $T^3\mathbf{1}(v)$ gives the 26-function family whose sole upper-endpoint relaxation was proved after Lemma 6.1. A 31-function family consists of these original 25 features, followed by

$$d_v^3, \quad d_v(Td)(v), \quad T(d^2)(v), \quad T^3\mathbf{1}(v), \quad \mathbb{E}_{y,z,t}A_{vy}A_{yz}A_{zt}A_{tv}, \quad c(v,x)^3.$$

Here $c(v,x) = \mathbb{E}_y A_{vy}A_{xy}$. The first five appended features are constant in the output x. Their rooted four-position graphs have codes $7, 13, 25, 26, 30$ when the single root is position zero. To place them in the two-rooted list, map those four positions to $(0, 2, 3, 4)$, leaving output position 1 isolated. The final feature is $K_{2,3}$ with $\{0, 1\}$ its part of size two. These rules fix the order and normalization of every coordinate.

The seven uncentered functions

$$g_0(v), g_1(v), g_2(v), g_3(v), t_1(v,x), t_2(v,x), t_3(v,x)$$

also form a subfamily of the same space. Their centered versions use $h_j = t_j - g_j$ as in (10). The exact nine-position Rayleigh matrices and the three-, four-, and seven-dimensional upper-endpoint relaxations are distinguished by their feature definitions; none assumes regularity. Lemma 6.1 shows that each relaxed matrix is an actual Rayleigh matrix plus an upper-density positive semidefinite localizer.

8.4 Degree and support localizers

We use the following elementary principles, also familiar as positive moment matrices and localizers in polynomial optimization [11].

Lemma 8.1. *If $u(R) \geq 0$ and $f(R)$ is a finite real vector, then $\mathbb{E}_R[u(R)f(R)f(R)^{\mathsf{T}}] \succeq 0$. If every function $z_i(R,\cdot)$ is supported on a set S_R of weighted mass at most $c > 0$, then*

$$\mathbb{E}_R\left[\mathbb{E}_x z(R,x)z(R,x)^{\mathsf{T}} - c^{-1}(\mathbb{E}_x z(R,x))(\mathbb{E}_x z(R,x))^{\mathsf{T}}\right] \succeq 0.$$

Proof. For the first statement, contraction with a vector is the expectation of a nonnegative weight times a square. For the second, a linear combination h of the coordinates satisfies, for each R,

$$(\mathbb{E}_x h(R,x))^2 = \left(\mathbb{E}_x \mathbf{1}_{S_R}(x)h(R,x)\right)^2 \leq w(S_R)\mathbb{E}_x h(R,x)^2 \leq c\mathbb{E}_x h(R,x)^2.$$

This is Cauchy–Schwarz in $L^2(w)$. Averaging proves positivity of every quadratic form of the stated matrix. □

Let $p(v)$ be the vector of the 12 induced four-position flags with one root, sampled on three private positions. Define

$$L_j = \mathbb{E}_v[d_v^j p(v)p(v)^\mathsf{T}] \qquad (j = 0, 1, 2).$$

The core inequalities (5) and Lemma 8.1 give all three degree constraints

$$\tfrac{1}{2}L_0 - L_1 \succeq 0, \qquad L_1 - rL_0 \succeq 0, \qquad (r + \tfrac{1}{2})L_1 - L_2 - \tfrac{r}{2}L_0 \succeq 0. \tag{14}$$

The last multiplier is $(d_v - r)(1/2 - d_v)$. The L_0 product has one root and two disjoint private triples, hence seven positions. Each factor d_v adds an independently sampled neighbor of the root, so L_1, L_2 use eight and nine positions.

The same quadratic interval constraint applies to a three-root flag vector $p(R)$ with two private positions. For a selected labeled root i, use $u(R) = d_{R_i}$ and set $L_j = \mathbb{E}_R[u(R)^j p(R)p(R)^\mathsf{T}]$. The norm product uses seven positions and the two added degree leaves bring the order to nine. For the path root type with center 0 and leaves $1, 2$, one may instead take $u(R) = w(N_{R_1} \cup N_{R_2})$. On this root type, $r \le d_{R_1} \le u(R) \le w(P_{R_0}) \le 1/2$. The vector is zero for a wrong root type, so the quadratic outer product is again positive semidefinite everywhere.

These entries can be counted by ordinary nine-position flags without introducing new graph counts. Let $z(R)$ be the three-root flag vector with three private positions. Define integer matrices P, Q as follows. For each coordinate of z, choose one of its three private positions as a marked leaf and delete it; P_{HF} counts choices yielding the remaining two-private flag F. The entry Q_{HF} counts only choices where the marked leaf is adjacent to the selected root, or to at least one of the two path leaves in the union version. Conditional exchangeability and independence of the marked leaf from the other two samples give

$$p = P^\mathsf{T} z/3, \qquad up = Q^\mathsf{T} z/3.$$

Thus, writing $O = \mathbb{E}_R zz^\mathsf{T}$,

$$L_0 = P^\mathsf{T} OP/9, \quad L_1 = (P^\mathsf{T} OQ + Q^\mathsf{T} OP)/18, \quad L_2 = Q^\mathsf{T} OQ/9.$$

These are exact conditional identities, including repeated sampled classes. Substitution into the last matrix in (14) fixes every normalization factor.

Within the 89 homomorphism features, select those whose pattern contains a two-edge path from position 0 to position 1. There are 22, and their functions vanish outside P_v. Selecting instead the patterns containing edge 01 gives 33 features, vanishing outside N_v. The finite selection rules, not the numerical counts, define these coordinate lists. In either case Lemma 8.1 with $c = 1/2$ gives the covariance constraint

$$B^{\mathrm{sup}} - 2R^{\mathrm{sup}} \succeq 0, \qquad B^{\mathrm{sup}} = \mathbb{E}_{v,x} zz^\mathsf{T}, \quad R^{\mathrm{sup}} = \mathbb{E}_v(\mathbb{E}_x z)(\mathbb{E}_x z)^\mathsf{T}. \tag{15}$$

The B^{sup} gluing identifies both distinguished positions, using at most eight positions. The R^{sup} gluing identifies only the parameter root and uses at most nine. These constraints use Lemma 3.1 at a core; they are not assertions for arbitrary probability weights on an arbitrary admissible graph.

8.5 Common-neighbor union constraints

For an ordered tuple of k root classes R, put

$$c_R = \mathbb{E}_z \prod_{i=0}^{k-1} A_{R_i z}, \qquad u_R = w\left(\bigcup_{i=0}^{k-1} N_{R_i}\right) = \mathbb{E}_y\left[1 - \prod_{i=0}^{k-1}(1 - A_{R_i y})\right].$$

If $c_R > 0$, there is a positive-weight class z adjacent to every root. Each N_{R_i} is then contained in P_z. This independent set has mass at most $1/2$, so $u_R \le 1/2$. If $c_R = 0$, the following product vanishes. In all cases,

$$c_R(1/2 - u_R) \ge 0. \tag{16}$$

This reasoning allows repeated root classes. Multiplying (16) by any vector outer product and averaging in R gives another instance of Lemma 8.1.

For two roots u, v one has $c_R = c(u,v)$ and $u_R = d_u + d_v - c(u,v)$. Take as coordinate functions the 25 two-rooted four-position homomorphism features, and retain the 14 whose pattern stays admissible after adding a common neighbor of the two roots. The resulting matrix is

$$\mathbb{E}_{u,v}\left[c(u,v)(1/2 - d_u - d_v + c(u,v))z(u,v)z(u,v)^{\mathsf{T}}\right] \succeq 0. \tag{17}$$

A removed feature has zero contribution when a common neighbor exists: otherwise its required edges and this neighbor would give a forbidden odd closed walk. The retained subset is specified by the stated admissibility test and increasing code order. Two private pairs for the feature product, a common-neighbor witness, and a union-test position use $2+2+2+1+1 = 8$ formal positions.

For three roots use instead the induced five-position flags with empty root type and two private positions, as in $\mathcal{F}(0,2)$ for $k = 3$. Retain the flags that remain admissible when one new vertex is made adjacent to all three roots. This retains 44 of the 50 flags and defines the vector $p(R)$ in

$$\mathbb{E}_R[c_R(1/2 - u_R)p(R)p(R)^{\mathsf{T}}] \succeq 0. \tag{18}$$

Roots with a common neighbor are necessarily pairwise nonadjacent; otherwise they form a triangle with that neighbor. Thus setting the vector to zero for a nonempty root type discards no nonzero contribution. The formal positions are the three roots, one common-neighbor witness, one union-test position, and two disjoint private pairs, totaling nine. As before, the omitted flags have zero contribution in the presence of the witness. The same construction with a single private position works for $k = 3, 4, 5$ and has $k + 4 \le 9$ positions, although (18) is the larger union-flag block needed here.

8.6 Localizing in the global edge density

If a matrix $M \succeq 0$ has an expansion on at most s positions, then on $\ell \le a \le b$ the matrices

$$(a - \ell)M, \qquad (b - a)M, \qquad (a - \ell)(b - a)M \tag{19}$$

are positive semidefinite. Multiplication by $a = t(K_2)$ adds two independent positions and one required edge, disjoint from the positions of the original expression. Multiplication by a^2 adds two such edges on four new positions. Hence the linear localizers in (19) have order at most $s + 2$, and the quadratic one has order at most $s + 4$.

For the linear localizers take the seven-position ordinary matrix $O_{0,3}$ with one root or the Rayleigh matrix of the 25 two-rooted four-position homomorphism features. For the quadratic localizers take the five-position ordinary matrices with $(k, s) = (1, 2)$ or $(3, 1)$, or the one-output spectral matrices with three roots. All these products have order at most nine. In particular, the Rayleigh entries include r before multiplication: the quadratic multiplier acts on $K - rB$, not on K alone.

The scalar constraints are

$$a - \ell, \quad b - a, \quad (a - \ell)(b - a), \quad (a - \ell)t(P_j), \quad (b - a)t(P_j) \quad (2 \le j \le 7),$$

where P_j denotes the path on j vertices. The path densities are nonnegative, and their products with a are densities of the disjoint union of a path and an edge. Thus these scalar constraints also use at most nine positions. Every interval-dependent expression in this section is a nonnegative scalar times a positive matrix, or the upper-endpoint relaxation of Lemma 6.1; no monotonicity of arbitrary matrix entries is asserted.

9 Finite coefficients and their exact verification

9.1 Completeness of the host catalog

The coefficient hosts are all admissible simple graphs on nine vertices, up to isomorphism. Their completeness can be checked by the following finite inductive algorithm. Start with the single empty graph of order zero. Given a complete list of order j, add a new vertex with each of the 2^j possible neighborhoods. Reject every extension with a triangle or pentagon and identify isomorphic remaining graphs. Every admissible $(j + 1)$-vertex graph occurs: deleting any vertex produces a graph isomorphic to an order-j representative, and the deleted vertex's neighborhood is among those tried. Thus induction proves completeness, independently of any efficiency choices in isomorphism testing.

Admissibility can be tested without enumerating cycles: check for every vertex that its neighborhood and its two-walk support are independent. The absence of triangles and pentagons implies these conditions by the closed-walk argument in Section 2. Conversely, a triangle violates the first condition, and in a pentagon the two vertices at distance two from any chosen cycle vertex are adjacent and violate the second. Graph isomorphism is decided by finite backtracking over vertex bijections, preserving all already exposed adjacency and nonadjacency relations. Degree comparisons and interchangeable vertices can prune this search but do not change its acceptance condition.

The resulting class counts, all reproduced by the standalone verifier, are

j	0	1	2	3	4	5	6	7	8	9
$\lvert\mathcal{H}_j\rvert$	1	1	2	3	7	13	35	89	306	1141.

Consequently every graph sampled on nine formal positions in Lemma 5.1 belongs to exactly one of the 1141 host classes, regardless of the number of classes in the underlying weighted graph.

9.2 Reconstructing the coefficient identity

Fix one interval $[\ell, b]$ and the threshold r. For each chosen positive matrix $M_\alpha(H, w)$, the certificate provides an integer factor F_α whose rows are its stated coordinates. For a positive

integer scale s_0, its positive semidefinite multiplier is

$$X_\alpha = s_0^{-2} F_\alpha F_\alpha^{\mathsf{T}}.$$

A stored integer coordinate scale means that the corresponding row of F_α is first multiplied by that integer; equivalently it can be absorbed into the factor. This is just a diagonal congruence. For each scalar constraint $g_\beta \geq 0$, an integer f_β specifies the multiplier $s_0^{-2} f_\beta^2$. Thus the full expression is

$$\Phi(H, w) = \sum_\alpha \operatorname{tr}(X_\alpha M_\alpha(H, w)) + s_0^{-2} \sum_\beta f_\beta^2 g_\beta(H, w) \geq 0. \tag{20}$$

Positivity follows from the preceding lemmas and integer Gram factorizations; it requires no floating-point eigenvalue tolerance.

Each matrix entry has been specified by products of edge and nonedge indicators on at most nine positions. Pad unused positions and average its defining indicator over all 9! bijections to a host graph J. The result is its exact coefficient on J in Lemma 5.1. A factor $1 - A_{uv}$ may equivalently be expanded by inclusion–exclusion, leaving a signed sum of ordinary homomorphism patterns. A pattern containing a triangle or pentagon has zero density on an admissible graph, by the odd-closed-walk argument, and may be omitted. Such a remaining pattern's numerator is the number of its edge-preserving bijections to J. Flags impose their stated induced conditions; homomorphism features impose only required edges. The two constructions must not be interchanged.

These rules give an exact rational identity

$$\Phi(H, w) = \sum_{J \in \mathcal{H}_9} \gamma_J y_J. \tag{21}$$

A common positive denominator D clears the factors s_0^2, 9!, and the rational parameters appearing in the specified entries. The stored integers c_J represent $\gamma_J = c_J/D$. The verification conditions are that the reconstructed fractions agree with these fractions for every host and that every c_J is strictly negative. As $\sum_J y_J = 1$, strict negativity includes the equality case $q = r$ and interval endpoints.

For an order-eight certificate, reconstruct its eight-position entries first and apply (9). In particular, if c_F/D_8 is its coefficient on an eight-vertex host, its nine-position coefficient is

$$\frac{\sum_{v \in V(J)} c_{J-v}}{9 D_8}.$$

The upper-endpoint matrix definitions are retained in this lifting. Alternatively, they can be written as actual nine-position Rayleigh matrices and the upper-density localizers proved in Lemma 6.1. These are equal weighted expressions.

9.3 Standalone verification and the interval requirement

The ancillary proof data consists of a Python verifier, a host/count data file, and one or more explicit JSON files containing integer certificates. If an interval bundle is split into several files, the verifier combines all their intervals before checking global coverage. The verifier uses only the Python standard library. The data file records adjacency encodings and integer homomorphism counts; its counts are not trusted in a full verification. The verifier regenerates the complete host catalog, reconstructs every count from adjacency, and then contracts the exact integer factors

with the defining coefficients. An injective-pattern count may be computed by assigning the pattern vertices to distinct host positions one at a time and checking every required edge when its second endpoint is assigned. If isolated pattern positions are omitted during this search, their assignments contribute the factorial of the number of remaining host positions. This enumerates all edge-preserving bijections exactly.

The adjacency encodings are the elementary graph6 encodings for orders at most nine: subtract 63 from the first character code to obtain the order, then read six bits per remaining character, also offset by 63. These bits specify edges in the order

$$(0,1),(0,2),(1,2),(0,3),(1,3),(2,3),\dots.$$

Unused final padding bits are discarded. This encoding convention is distinct from the integer codes used to order rooted coordinates in Section 8.1.

The check concerns a universal coefficient identity, not eigenvalues of a sample of small graphs. The hosts index the possible outcomes of nine formal samples in an arbitrary weighted graph. Once the analytic positivity and sampling identities are proved, their finite exhaustive coefficient verification applies to graphs of every order. Conversely, a successful coefficient check proves only the stated conditional interval inequality. The analytic core reduction and density pruning are provided by Theorem 2.3 and Corollary 4.2.

For $r = 1/30$, the supplied certificates cover the following nine closed intervals. Their coefficient identities have been reconstructed in integer arithmetic. Each listed rational margin satisfies $\gamma_J \le -\delta_i$ for every one of the 1141 nine-position hosts on the indicated interval; the comparisons are exact, including the lifted order-eight certificates.

	ℓ	b	δ_i
I_1	$1/28$	$1/10$	10^{-9}
I_2	$1/10$	$3/25$	10^{-10}
I_3	$3/25$	$13/100$	10^{-9}
I_4	$13/100$	$3/20$	10^{-8}
I_5	$3/20$	$17/100$	10^{-9}
I_6	$17/100$	$9/50$	10^{-8}
I_7	$9/50$	$1/5$	10^{-7}
I_8	$1/5$	$7/30$	10^{-6}
I_9	$7/30$	$77/200$	10^{-7}

The verifier compares the rational endpoints exactly. Their union is $[1/28, 77/200]$, so no intervening density range or shared endpoint is omitted.

Proof of Theorem 1.1. Suppose that an admissible graph on $n \ge 1$ vertices satisfies $q_{\min}(G) \ge n/30$. Give each vertex weight $1/n$. The weighted extremal reduction of Theorem 2.3 supplies a positive weighted induced subgraph with $q = S_H(w) \ge 1/30$. Corollary 4.2 puts its edge density in $[1/28, 77/200]$. Choose an interval I_i containing that density. All matrix and scalar terms defining its certificate are nonnegative under the proved core hypotheses. Equations (20) and (21), together with the verified margin, give

$$0 \le \Phi(H,w) = \sum_{J \in \mathcal{H}_9} \gamma_J y_J \le -\delta_i \sum_{J \in \mathcal{H}_9} y_J = -\delta_i < 0,$$

a contradiction. The reduction and sampling identities are finite and exact, so this argument has no asymptotic restriction on n. □

The source archive contains `anc/verify.py`, `anc/proof-data.json`, and the two certificate files `anc/certificates-1.json` and `anc/certificates-2.json`. From the archive's root directory the full verification command is

```
python3 anc/verify.py --certificate anc/certificates-1.json \
    anc/certificates-2.json
```

The verifier requires Python 3.10 or later and only its standard library. The proof-data file is located relative to the verifier. Full verification reconstructs the catalog and its counts before checking all coefficients, the margins in the table, and the interval cover.

Numerical semidefinite and linear optimization were used to discover the certificates, using Clarabel [8] and the dual-simplex methods described by Huangfu and Hall [10]. Graph isomorphism tools of McKay and Piperno [14] were used during construction. The final verification does not require these packages or trust their floating-point outputs. It checks integer Gram factors and exact counting identities by the algorithms specified above. The supplied computational check is distinct from formal verification in a proof assistant.

Acknowledgements

The author was supported by the NRDI grant KKP 138270. OpenAI Codex assisted with mathematical development, literature searches, programming, and drafting. The author selected and directed the investigation; independent computational checks and agent reviews were used during preparation. These checks are not formal verification in a proof assistant.

References

[1] A. Abiad, V. Taranchuk, and T. van Veluw, On the sum of the largest and smallest eigenvalues of odd-cycle free graphs, *Electron. J. Combin.* **33** (2026), no. 2, Paper P2.31. doi:10.37236/14427.

[2] J. Balogh, F. C. Clemen, B. Lidický, S. Norin, and J. Volec, The spectrum of triangle-free graphs, *SIAM J. Discrete Math.* **37** (2023), 1173–1179. doi:10.1137/22M150767X.

[3] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, 2004. doi:10.1017/CBO9780511804441.

[4] A. E. Brouwer and W. H. Haemers, *Spectra of Graphs*, Universitext, Springer, New York, 2012. doi:10.1007/978-1-4614-1939-6.

[5] D. Cai, Z. Chen, and X.-D. Zhang, *A sharp bound for the least signless Laplacian eigenvalue of triangle-free graphs*, arXiv:2610.14503v1, 2026. https://arxiv.org/abs/2610.14503.

[6] G. Coutinho and T. Oliveira, A connection between the A_α-spectrum and the Lovász theta number, *Linear Algebra Appl.* **719** (2025), 93–102. doi:10.1016/j.laa.2025.04.007.

[7] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42** (1995), 1115–1145. doi:10.1145/227683.227684.

[8] P. J. Goulart and Y. Chen, *Clarabel: An interior-point solver for conic programs with quadratic objectives*, arXiv:2405.12762, 2024. https://arxiv.org/abs/2405.12762.

[9] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, Cambridge, 2013. doi:10.1017/CBO9781139020411.

[10] Q. Huangfu and J. A. J. Hall, Parallelizing the dual revised simplex method, *Math. Program. Comput.* **10** (2018), 119–142. doi:10.1007/s12532-017-0130-5.

[11] M. Laurent, Sums of squares, moment matrices and optimization over polynomials, in M. Putinar and S. Sullivant (eds.), *Emerging Applications of Algebraic Geometry*, IMA Volumes in Mathematics and its Applications, vol. 149, Springer, New York, 2009, pp. 157–270. https://ir.cwi.nl/pub/13237.

[12] S. Letzter and R. Snyder, The homomorphism threshold of $\{C_3, C_5\}$-free graphs, *J. Graph Theory* **90** (2019), 83–106. doi:10.1002/jgt.22369.

[13] L. Lovász, *Large Networks and Graph Limits*, American Mathematical Society Colloquium Publications, vol. 60, American Mathematical Society, Providence, RI, 2012. doi:10.1090/coll/060.

[14] B. D. McKay and A. Piperno, Practical graph isomorphism, II, *J. Symbolic Comput.* **60** (2014), 94–112. doi:10.1016/j.jsc.2013.09.003.

[15] A. A. Razborov, Flag algebras, *J. Symbolic Logic* **72** (2007), 1239–1282. doi:10.2178/jsl/1203350785.

[16] C. Wang, *Asymptotic spectral bounds under odd-girth constraints*, arXiv:2610.05172v1, 2026. https://arxiv.org/abs/2610.05172.

[17] Q. Zhou, *Spectral bipartiteness in generalized odd graphs of diameter three*, arXiv:2609.11729v1, 2026. https://arxiv.org/abs/2609.11729.

[18] Q. Zhou, *The least signless Laplacian eigenvalue of $\{C_3, C_5\}$-free graphs*, arXiv:2610.00875v1, 2026. https://arxiv.org/abs/2610.00875.