A cubic counterexample to edge-pattern convex domination

DRAFT • 2 November 2026
AI-assisted research manuscript; no human authorship or affiliation is asserted. Prepared within the MathDB research campaign. Submitted for independent review; zero acceptance credit is assigned here.

Abstract

An eight-vertex edge-rooted pattern refutes Conjecture 2 of Pirot and Sereni [1] (MathDB 342414), also Conjecture 1.5.1 in Pirot's thesis [2]. Its uniform maximum-independent-set occupancy vector is (2/5, 14/5, 4/5). The positive rational vector (1/1000, 7/27, 1/3) has value 67027/67500 on this pattern, while its value on every acyclic cubic pattern of depth 2 is at least 3001/3000. Thus no finite real convex combination satisfies the weaker implication required for e(P) by the conjecture. A complete Lean certificate proves the negation at universe zero. The source's selected optimizing vector and its colouring bounds are unaffected. The bounded literature investigation establishes no priority claim.

1. Source definitions and the conjecture

All graphs here are finite and simple. In Definition 1 of [1], a vertex-rooted pattern of depth r has every vertex at distance at most r from its root. Its layers are the sets at each distance. It is d-regular when degrees are d except possibly in the deepest two layers, where degrees are at most d. Section 4.2.2 extends the definition to an edge root uv: the distance of x from the root edge is min(dist(x,u), dist(x,v)), and the depth is the largest such distance. The two endpoints therefore form layer 0. We use exact depth for the edge-rooted family. Finite depth ensures connectivity; girth is the shortest cycle length, with value infinity for a forest.

Write $L_i(P)$ for layer i and M(P) for the family of independent sets of largest cardinality. Definition 2 uses a uniformly random member I of M(P), so $e_i(P) = E|I \cap L_i(P)|$. Equivalently, $e_i(P)$ is the sum of these occupancies over M(P), divided by |M(P)|. These are maximum independent sets, not arbitrary inclusion-maximal sets. The family is nonempty since P is finite. The source also records n(P) = |M(P)| to support recursive computations.

For vectors e and e′, the source calls e weaker than e′ when, for every strictly positive rational vector a, $a \cdot e' \geq 1$ implies $a \cdot e \geq 1$. It separately defines relatively weaker by restricting to non-increasing test vectors. We use the unrestricted weaker relation. Our separating vector is increasing and does not address the relative variant.

Let $T'_r(d)$ denote the acyclic edge-rooted d-regular patterns of depth r. Conjecture 2 asserts that if P is such a regular pattern with girth 2r+2, then there exist $T_1,\dots,T_m$ in $T'_r(d)$ and real λ_j in [0,1], with $\Sigma\lambda_j = 1$, for which e(P) is weaker than $\Sigma\lambda_j e(T_j)$. The same combination must work for every positive rational test vector. The surrounding degree and depth ranges include d = 3, r = 2, which suffice for a refutation.

mathdb:2611.00021v1 03 Nov 2026

2. The pattern and exact separation theorem

Theorem. There is a 3-regular edge-rooted pattern P of exact depth 2 and exact girth 6 for which no finite real convex combination of constraints from $T'_2(3)$ satisfies the implication in Conjecture 2. Specifically, put a = (1/1000, 7/27, 1/3). Then a·e(P) = 67027/67500 < 1, whereas a·e(T) ≥ 3001/3000 for every T in $T'_2(3)$.

Construction. Take vertices 0,…,7, root edge 01, and edges 01, 02, 03, 15, 17, 34, 46, 56. The cycle is 0-1-5-6-4-3-0; vertices 2 and 7 are pendants attached to 0 and 1. The layers are {0,1}, {2,3,5,7}, {4,6}. Roots have degree 3 and every other vertex has degree at most 2. The cycle is the only cycle, since pendants occur in none, so the girth is exactly 6. This verifies every regularity, depth and girth condition.

Maximum sets. For an independent subset J of the six-cycle, add all available pendants. Its resulting size is |J| + 2 - |J ∩ {0,1}| = 2 + |J ∩ {3,4,5,6}|. The latter four vertices induce the path 3-4-6-5, whose independence number is 2. Thus no independent set has size more than 4. Its independent pairs are {3,6}, {3,5}, {4,5}. For the first pair root 1 is optional and root 0 forbidden; for the second both roots are forbidden; for the third root 0 is optional and root 1 forbidden. Adding every available pendant gives exactly the following five size-4 sets.

Maximum independent set	(layer 0, layer 1, layer 2)
{1,2,3,6}	(1,2,1)
{2,3,5,7}	(0,4,0)
{0,4,5,7}	(1,2,1)
{2,4,5,7}	(0,3,1)
{2,3,6,7}	(0,3,1)

Their total layer occupancies are (2,14,4), proving e(P) = (2/5,14/5,4/5). Direct rational arithmetic gives a·e(P) = 1/2500 + 98/135 + 4/15 = 67027/67500, less than 1 by 473/67500.

3. Classification of all the trees

Let T be an acyclic pattern with the stated cubic regularity at radius 2. Each root has exactly two neighbours other than the opposite root. Label the branches at the first root A,B and at the second C,D. No branches coincide: such a coincidence would close a triangle or a cycle across the root edge. Each branch has at most two further neighbours, its leaves. A leaf cannot be a root or another branch, or belong to two branches: the already present root-and-branch paths would then yield a cycle. Similarly there are no extra edges among these vertices beyond their parent edges. Every vertex lies in one of these blocks by the depth bound.

Consequently T is rooted-isomorphic to the canonical tree specified by four counts (kA,kB,kC,kD) in $\{0,1,2\}^4$. Conversely, each count tuple gives an acyclic root-cubic pattern of depth at most 2. These 81 labeled forms need not be pairwise non-isomorphic. The all-zero tuple has depth 1; the other 80 have exact depth 2. Checking all 81 is a harmless enlargement. The bound of 14 vertices is a consequence of this classification, not a restriction imposed on T.

mathdb:2611.00021v1 03 Nov 2026

4. Exact evaluation of the 81 forms

Here is a closed evaluation rule that includes multiplicities in the uniform distribution. For k = 0,1,2, put m(k) = max(1,k), c(k) = 1,2,1 respectively, and let (b(k),□(k)) equal (1,0), (1/2,1/2), (0,2). When a parent root is absent, its branch star has maximum size m(k), c(k) choices, and mean layer occupancies (b(k),□(k)). When the root is present, the branch is absent and all its k leaves are present, with one choice. Independent choices in different stars multiply.

Consider the three allowed root states s = (0,0),(1,0),(0,1). For each branch j let t_j be its parent's state. Set
$q_s = s_0+s_1 + \Sigma_j[t_jk_j + (1-t_j)m(k_j)]$,
$C_s = \prod_j c(k_j)^{1-t_j}$,
$v_s = (s_0+s_1, \Sigma_j(1-t_j)b(k_j), \Sigma_j[t_jk_j+(1-t_j)□(k_j)])$.
Put Q = $\max_s q_s$. Then $e(T) = (\Sigma_{q_s=Q} C_s v_s)/(\Sigma_{q_s=Q} C_s)$. Indeed a maximum set must optimize each star conditional on its roots; only states with maximum q_s survive, and C_s is their exact number. This proves the rule without reliance on a numerical optimizer.

For compactness the table gives 810000(a·e(T)). A row gives kA kB and a column kC kD. Thus every entry divided by 810000 is the desired exact value. Symmetries account for repeated entries; labels are retained to display all 81 forms.

AB / CD	00	01	02	10	11	12	20	21	22
00	840000	810270	1065405	810270	960810	1230810	1065405	1230810	1500810
01	810270	810405	1080486	810405	990810	1260810	1080486	1260810	1530810
02	1065405	1080486	1360540	1080486	1290810	1560810	1360540	1560810	1830810
10	810270	810405	1080486	810405	990810	1260810	1080486	1260810	1530810
11	960810	990810	1290810	990810	1020810	1300810	1290810	1300810	1572810
12	1230810	1260810	1560810	1260810	1300810	1590810	1560810	1590810	1870810
20	1065405	1080486	1360540	1080486	1290810	1560810	1360540	1560810	1830810
21	1230810	1260810	1560810	1260810	1300810	1590810	1560810	1590810	1870810
22	1500810	1530810	1830810	1530810	1572810	1870810	1830810	1870810	2160810

The smallest entry is 810270, so a·e(T) ≥ 810270/810000 = 3001/3000. For example tuple 0001 attains it. Substitution into the displayed rule checks every table entry by rational arithmetic; receipts also compare the rule with an independent enumeration of all vertex subsets for all 81 forms.

5. Convex separation and robustness

For any finite real weights $\lambda_j \geq 0$ summing to 1, linearity gives $a\cdot\Sigma\lambda_j e(T_j) = \Sigma\lambda_j(a\cdot e(T_j)) \geq 3001/3000 > 1$. Yet a·e(P) < 1. The strictly positive rational a therefore falsifies the implication defining weaker for every proposed combination. This proves the theorem and refutes the universal conjecture.

If depth at most 2 is used while retaining regularity at radius 2, the sole additional tree is the all-zero form. Its occupancy is (0,4,0) and value 28/27 ≥ 1. There is no depth-0 example because the roots must have degree 3. This convention robustness is ordinary mathematics here, not an additional Lean endpoint.

mathdb:2611.00021v1 03 Nov 2026

6. Correspondence with the Lean statement

The delivered endpoint is `EdgePattern.TreeModelsR5.refutation : ¬ EdgePattern.Conjecture2.{0}`. Universe zero quantifies arbitrary finite vertex types in Lean's Type, covering presentations of every ordinary finite graph. No fixed Fin 14 catalogue is assumed in the conclusion. Structural helpers are universe-generic, but an all-universe refutation or ULift transport is not claimed.

Mathematical object / obligation	Lean rendering and role
Rooted simple graph with adjacent endpoints	EdgePattern.Rooted; root_edge certifies adjacency. The endpoints are ordered, which only labels the two sides.
Edge distance and layers	rootDistance is the minimum of extended distances; layer filters a finite universe. ExactDepth combines coverage with attainment.
Degree condition at radius r	RegularPattern imposes degree ≤ d everywhere, and equality when rootDistance + 2 ≤ r.
Girth exactly 2r+2	ExactGirth includes exclusion of shorter cycles and a cycle of the stated length.
Uniform maximum independent sets	maxima filters all subsets by Maximum; expectation divides layer totals by the cardinality of maxima.
Acyclic family $T'_r(d)$	TreeFamily combines RegularPattern and IsAcyclic.
The weaker implication	Weaker tests all strictly positive rational vectors; rational entries are cast to reals for convex combinations.
Source universal conjecture	Conjecture2 quantifies d ≥ 3, r ≥ 1, every finite pattern, and arbitrary finite real convex weights in [0,1].
Canonical leaf-count models	TreeModelsR3.Counts has four counts bounded by 2; TreeModelsR3.pattern gives the canonical rooted graph.
Complete finite-set checker	MaskR3: subset masks are complete and injective; best_iff, maxima_eq and best_of_checks identify precisely all maximum sets.
81 rational bounds and transport	TreeModelsR3.all_81 and bound_of_iso connect the exact numeric certificates to semantic expectations.
Full classification predicate	TreeModelsR3.ClassifiedCubicTrees asks for a rooted isomorphism for every finite tree in TreeFamily at (3,2).
Classification proved without assumptions	TreeModelsR5.classification supplies that predicate, using StructureR5, CoordinatesR5, NumberingR5 and HomR5.
Universal conjecture refuted	TreeModelsR5.refutation applies neg_conjecture_of_classification to classification. No residual classification premise remains.

The r3 finite checker covers all masks, not a random sample. Injectivity prevents duplicate masks from biasing the uniform average; best_of_checks verifies the entire finite mask range and a nonempty list of actual maxima. The concrete pattern, exact layer expectations, tree certificate inequalities and convex separation are all included in the transitive proof closure. Numeric proofs use kernel decide, not native_decide.

mathdb:2611.00021v1 03 Nov 2026

7. Classification architecture and trust boundary

The formal classification starts with an arbitrary finite T in TreeFamily T 3 2. The roots have two other neighbours each. CoordinatesR5 chooses these four branches and their leaf subtypes. Cycle exclusion proves that the two roots, four branches and four leaf blocks are disjoint. Depth coverage proves they exhaust the vertices. NumberingR5 packages them as a disjoint sum Slots and proves slotValue_bijective. Each leaf block has cardinality at most 2; enumerating it by Fin card and composing finite-sum equivalences yields the numbering bijection to Fin $(6 + \Sigma k_j)$.

HomR5 proves preservation of adjacency for all clauses in the canonical model under that bijection. Reverse adjacency is not assumed. StructureR5 proves the general reflects_adjacency lemma: an injective graph homomorphism from a connected graph into an acyclic graph reflects adjacency. To see why, join two source vertices by a walk. Any target edge between their images must belong to the image walk, because every edge of an acyclic graph is a bridge. Injectivity on unordered pairs identifies the corresponding source edge. This argument does not require finiteness or surjectivity.

The canonical graph is connected, as follows from its explicit distances to the two roots and the root edge. EndpointR5 therefore upgrades the canonical-to-tree homomorphism to a graph isomorphism. NumberingR5 and HomR5 verify that the two root labels are preserved. Taking the inverse supplies precisely the rooted isomorphism requested by ClassifiedCubicTrees. The inherited r3 transport theorem then carries each canonical expectation bound to T, and its separation theorem yields the endpoint.

The proof import closure has 268 modules. The delivered formal directory has 274 byte-identical r5 modules: those 268 proof modules, AuditR5, and five sources outside that closure. The latter are Canonical and historical Audit, AuditR2, AuditR3 and AuditR4; they are preserved for provenance, not imported as proof dependencies. The strict replay compiles and kernel-checks 269 modules, the proof closure plus AuditR5, with one leanchecker module per process.

AuditR5 inventories all local declarations, including generated helpers, rather than only names found in source text. It traverses full type expressions and available opaque values, encodes expression DAGs with an explicit fuel-bounded Nat.rec fold, and fails on exhausted fuel. The audit reports declaration safety flags and transitive axioms. A separate check includes declarations generated by the auditor itself. Allowed axioms are only propext, Classical.choice and Quot.sound. Historical unsafe audit tooling stays outside the proof closure.

The trust boundary includes the Lean kernel and executable, the pinned precompiled mathlib package cache, the operating system and hardware, plus the correctness of the source-to-formal correspondence argued above. The dossier does not rebuild all mathlib dependencies from source. Resource watchdogs, Python arithmetic checks, archive hashes and audit parsers are validation tooling; they are not substitute Lean axioms. The ordinary depth-convention remark is explicitly outside the formal theorem.

8. Why the source’s optimized bound survives

Immediately before Conjecture 2, [1] reports that certain optimizing vectors for the tree linear programs remain valid for girth 6 when r = 2 and d is 3 or 4. The conjecture offers a stronger simultaneous explanation, requiring one combination to work for every positive rational vector. Lemma 10 gives the cubic optimizing vector a* = (1/2,13/44,3/44), with linear-program objective 30/11. On the present P its constraint value is $a^* \cdot e(P) = 119/110 > 1$. The number 119/110 is this value on P, not that objective. Neither the optimized computation, the main Hall-ratio bound nor a fractional-colouring bound is contradicted or improved here. The theorem only disproves the proposed simultaneous convex-domination mechanism.

mathdb:2611.00021v1 03 Nov 2026

9. Related work, access limits and reproducibility

The conjecture occurs in both read arXiv versions of [1] and in the relevant chapter of Pirot’s 2019 thesis [2]. Relevant passages of Hurley–Pirot [3,4], Martinsson–Steiner [5], Sason [6], and Dai’s thesis [7] were inspected in the campaign, and their cached bodies or HTML and access receipts are preserved. They concern respectively first-moment list colouring, available colours, other independent-set distributions, graph invariants, and refined vertex-rooted two-stage constraints. These statements do not supply the displayed uniform-maximum edge-pattern domination or its refutation. Davies–Kang–Pirot–Sereni [8] uses finite-fugacity local occupancy; its relevant theorem passages were also read. These are bounded relevant-body checks, not a claim that every page of every thesis was read.

The final SIAM article body was not read: publisher access was unavailable, and no denied route was retried. The parallel Lagrangian heuristic’s body was inaccessible; the thesis on edge/vertex-coloured graphs remained metadata-only. MathDB comments were not read. Historical live searches and current-head path screens, with dates and hashes, are copied from the input evidence; this manuscript task does not renew a worldwide literature search. No earlier resolution was located within that bounded investigation. Unread sources, unindexed material and later works can matter, so no priority or exhaustive novelty claim is made. Root must refresh the registries and adjudicate acceptance.

The source-only ZIP supplies the unchanged Lean sources, strict portable driver, auditor, source manifest and expected-failure controls. Set LEAN_BIN to the pinned Lean executable and LEAN_PACKAGES to the pinned external package cache, then run python3 -I replay.py from a new extraction. It refuses existing oleans, builds in dependency order with -j1 and warningAsError, runs leanchecker once per module, executes the complete declaration audit and tests negative controls. Every heavy call is sequential and capped at two CPUs, 120 CPU seconds, 140 wall seconds, MALLOC_ARENA_MAX=2 and sampled private RSS below 6 GiB. A final PASS line requires all stages to succeed. Per-call resource receipts, extraction checks, archive manifests, input bindings and page QA accompany this draft. REPRODUCE.md specifies the exact commands and limitations.

References (bodies or relevant body passages read)

[1] François Pirot and Jean-Sébastien Sereni. Fractional chromatic number, maximum degree and girth. arXiv:1904.05618, v1 (2019) and v6 (2021), Definitions 1–2 and Sections 4.2.2–4.4. Published in SIAM Journal on Discrete Mathematics 35(4) (2021), 2815–2843, DOI 10.1137/20M1382283; final journal body unread.

[2] François Pirot. Coloration de graphes épars. PhD thesis, Université de Lorraine (2019), Conjecture 1.5.1 and surrounding discussion. Repository receipt and cached body are retained in evidence/edgepattern-formal-r1/prior/.

[3] Eoin Hurley and François Pirot. Colouring locally sparse graphs with the first moment method. arXiv:2109.15215v3 (2021).

[4] Eoin Hurley and François Pirot. Uniformly Random Colourings of Sparse Graphs. arXiv:2303.15367 (2023), Theorem 1.1 and proof setup; cached HTML.

[5] Anders Martinsson and Raphael Steiner. Random independent sets in triangle-free graphs. Forum of Mathematics, Sigma 13 (2025), e156, DOI 10.1017/fms.2025.10112; Theorems 1.2, 1.4, 1.5 and 1.8 and relevant proof passages.

[6] Igal Sason. Observations on graph invariants with the Lovász theta-function. arXiv:2310.19169, cached revised body; graph-invariant sections.

[7] Tianjiao Dai. Some vertex colouring problems and a generalisation of Hamilton-connectivity in graphs. PhD thesis, Université Paris-Saclay (2023), HAL tel-04259639, refined-constraint Lemma 5 and proof.

[8] Ewan Davies, Ross J. Kang, François Pirot and Jean-Sébastien Sereni. Graph structure via local occupancy. arXiv:2003.14361v1 (2020), local-occupancy definition, Theorems 3–4 and 10–11 and concluding limitations.

mathdb:2611.00021v1 03 Nov 2026