Certified barriers for Turán's tetrahedron problem: exact ceilings at levels 6 and 7

David E. Thompson

SIMARC — Simulation Architecture & Research Computing

simarc.org · david@simarc.org

August 2026

Abstract

This is a companion to our exactly certified bound $\pi(K_4^{(3)}) \leq 951/1697 = 0.5604007071302298$. Here we map, with exact certificates wherever possible, why none of the routes examined here continues from that bound to the conjectured value $\frac{5}{9}$. On the flag-algebra side we prove: the level-6 family ceiling satisfies $\lambda_{\mathrm{fam}}(6) \geq 56096898007/100000000000 = 0.560968980070$; the level-7 family ceiling satisfies $\lambda_{\mathrm{fam}}(7) \geq 278405889284/499999999985 = 0.556811778584704$ over *every* admissible type set, the even types being exactly redundant; and, adjoining the 63 order-5 regularity rows that power the current record, still $\lambda^{\mathrm{reg}}_{\mathrm{fam}}(7) \geq 22251/40000 = 0.556275 > \frac{5}{9}$. Each is certified by an explicit integer ray verified in exact rational arithmetic, with no floating point in any verification step. These theorems forbid a *plain* (uncoloured, regularity-free) level-6 proof of the conjecture, and a level-7 proof either plain or carrying the 63 regularity rows; they do *not* forbid record improvements, and we compute exactly how much room they leave (85.15% of the remaining gap, for regularity-carrying level-7 certificates). Beyond flag algebras we show that each of the five standard alternatives—stability, containers/regularity, conditional (selector) certificates, spectral methods, and entropy methods—fails here for a specific, checkable reason; in particular the deletion set of any "closing selector" must meet the extremal family itself. Finally we locate the obstruction: Turán's own construction contributes *exactly zero* of it ($\lambda(\text{Turán-only}, 6) = \frac{5}{9}$ exactly), a single Fon-der-Flaass sign pattern contributes 69.28% of it, and the extremal continuum contributes nothing (its tightness footprint saturates at 45 classes and rank 8). The measured deficit–distance dichotomy—quadratic only along reweighting directions, linear with uniform constant in all first-order structural directions—motivates a quotient-stability conjecture that we state and place under adversarial test. We prove that conjecture on Ch_q, the sharpest adversarial family we know: a $K_4^{(3)}$-free family whose density approaches $\frac{5}{9}$ with deficit $D_q \sim \frac{4}{81q}$. Its closed-form codegree kernel turns stationarity into a finite semialgebraic problem, and we show $\delta_Q(\mathrm{Ch}_q) = O(D_q)$, so this family cannot be a counterexample. We then add three obstructions of a different kind, found by asking which *certificate classes* rather than which methods can see the extremal object: an extremal object at an interior zero defeats every Pólya–Handelman certificate but not sum-of-squares or copositive ones; the classically abundant extremal patterns at density exactly $\frac{5}{9}$ [8, 5, 6, 7] contain non-isomorphic pairs sharing their entire codegree distribution, their codegree-kernel spectrum, and all induced densities on four vertices, so the ℓ_2 method cannot separate them; and the $\ell_2 \to \ell_1$ transfer provably cannot reach $\frac{5}{9}$, since equality in Cauchy–Schwarz needs a constant codegree while Turán's is two-valued. The certificate class that survives—copositivity of the link matrix $C(p)$ on the nonnegative orthant—settles all 41,138 isomorphism classes of maximal $K_4^{(3)}$-free patterns on six colours (29,099,574 labelled) unconditionally, and by itself: no critical-value elimination enters. Finally we state what is left: after complete-apex domination and the closure of the no-clean-collision branch, every

configuration at density exactly $\frac{5}{9}$ that we know lies in the surviving branch, and the principal remaining object is an adaptive completion-closure statement. We record a caution with it: three independent reductions here terminate in a statement whose supremum over its own class is still exactly $\frac{5}{9}$, so a correct reduction is not automatically progress.

1 Introduction

Turán's 1941 conjecture that $\pi(K_4^{(3)}) = \frac{5}{9}$ remains open; the lower bound is his construction, and the upper bound is the entire problem; see Keevash [12] for the surrounding theory. Every upper bound since Razborov's flag-algebra method [24] has been computer-assisted, the best previously published certificate being Baber's [2], of certified value 0.5614656206270598. In a companion manuscript [1] we proved the exactly certified bound

$$\pi(K_4^{(3)}) \;\leq\; \frac{951}{1697} \;=\; 0.5604007071302298, \tag{1}$$

leaving a gap of $74/15273 = 4.845152 \times 10^{-3}$ to $\frac{5}{9}$. That remaining gap is 81.98% of the gap that separated Baber's certificate from $\frac{5}{9}$; equivalently, (1) closes 18.0186% of it.[1] The present paper answers the natural next question: *what stands between* (1) *and the conjecture?* The answer, assembled here with exact certificates where we have them and explicit labels where we do not, is that every route on the current map is either capped by a theorem, dead for a named structural reason, or open only in a direction nobody knows how to enter. We consider this negative map a result in itself: the value of knowing where a proof cannot come from is that the search can concentrate on where it still might.

Three conventions from [1] are kept throughout. Every number is labelled *proved* (exact rational certificate, verified without floating point), *measured* (numerical, with stated controls), or *cited*. Since $\pi(K_4^{(3)}) \geq \frac{5}{9}$ is a construction, every valid relaxation value satisfies $\lambda \geq \frac{5}{9}$; any computation suggesting otherwise indicates a wrong encoding, never progress. And no trend is extrapolated: ratios of successive gaps carry no information about later ones.

Summary of results. Section 2 proves three exact ceilings for family-tight flag-algebra programs (Theorems 1–3) and computes what they do and do not forbid. Section 3 reports the numerical capping of the colour-refined level-6 program, the two artefact controls it passed, and precisely what remains before it is a theorem. Section 4 treats the five standard methods outside flag algebras and gives, for each, the specific obstruction; one (Proposition 4) is settled by a proof. Section 5 locates the obstruction inside the extremal family: concentration (Theorem 5), saturation, and the measured rank law. Section 6 presents the deficit–distance dichotomy and the quotient-stability conjecture it motivates. Section 7 adds three obstructions indexed by certificate class rather than by method, together with two structural results that remove configurations from the search. Section 8 states the surviving frontier and the reformulation caution that goes with it. Section 9 collects limitations, including the incompleteness theorem of Hatami and Norine [10] which bounds what any programme of this kind can hope for.

[1] The bound (1) supersedes weaker values circulated while this paper and its companion were in preparation. Each was exactly certified against the same 49,242-row program, and the lineage is given in full in [1].

2 Exact ceilings at levels 6 and 7

For a level-N flag-algebra program in the sense of Razborov [24], write S_N for the set of N-vertex $K_4^{(3)}$-free isomorphism classes realised with positive density by the known extremal family (Turán's construction together with the Fon-der-Flaass variants [21] and all their reweightings); $|S_4|, \ldots, |S_8| = 4, 12, 45, 170, 649$. The *family ceiling* is

$$\lambda_{\mathrm{fam}}(N) \;=\; \min_{Q\succeq 0}\; \max_{H\in S_N}\; \big[\, d(H) + \langle Q, M(H)\rangle \,\big],$$

the best value any level-N certificate can attain on the classes the extremal family forces it to respect. Because every extremal object has density exactly $\frac{5}{9}$, a level-N proof of the conjecture requires $\lambda_{\mathrm{fam}}(N) = \frac{5}{9}$ exactly (Section 5 makes this tightness-forcing precise). The theorems of this section show it fails at 6 and 7 by explicit margins.

All three certificates below share one verification discipline: the certifying object is an *integer* ray $z \geq 0$ on S_N; the moment matrices are exact integers; positive semidefiniteness is verified by rational LDL^{T} elimination with explicit pivot tests; and no floating-point number appears in any verification step. Floating point is used only to *search* for the ray, never to check it.

Theorem 1 (Level-6 family ceiling)**.** $\lambda_{\mathrm{fam}}(6) \;\geq\; \dfrac{56096898007}{100000000000} \;=\; 0.560968980070 \;>\; \frac{5}{9}.$

This hardens the earlier certified value 0.560968171518778 by 8.086×10^{-7} and was reproduced by two independently written from-the-definition instruments whose moment matrices agree with the original computation entrywise (6 blocks $\times$ 45 classes).

Theorem 2 (Level-7 family ceiling, all type sets)**.** *Over every admissible type/flag-size pair* (k, m) *at* $N = 7$*—all nine pairs with* $3 \leq m$*,* $k < m$*,* $2m - k \leq 7$*—*

$$\lambda_{\mathrm{fam}}(7) \;\geq\; \frac{278405889284}{499999999985} \;=\; 0.556811778584704 \;>\; \tfrac{5}{9}.$$

Moreover the even-size types are exactly redundant: adjoining them changes $\lambda_{\mathrm{fam}}(7)$ *by zero, identically, because their moment matrices are a fibre-sum of the odd-type matrices and the corresponding lift* $V^{\mathsf{T}}QV$ *of any feasible* Q *is again positive semidefinite.*

The redundancy claim closes what had been the last "type-side" escape: a hope that a wrong-parity type could behave differently at level 7. It cannot; the corollary *no plain level-*7 *flag certificate over any admissible type set proves* $\pi(K_4^{(3)}) \leq \frac{5}{9}$ is unconditional. The certificate was verified by integer Bareiss elimination on all 25 live blocks, and the theorem was independently reproduced by two further implementations, written from the statement alone and importing nothing from the original code, with the boundary behaving exactly as it must: every stored ray with target ≤ 0.5568 verifies as positive semidefinite, and every ray with target ≥ 0.55682 fails with a negative eigenvalue.

The current record (1), however, does not live in the plain cone: it uses colour refinement and, critically, regularity rows. Neither device is ours: the coloured variant of the flag hierarchy is Baber's [2, Sec. 5.1], and the order-5 regularity rows are the Hladký–Kráľ–Norin constraints [23], in the form Baber adapts them [2, Sec. 5.1.1]. Measured directly, the record is *smaller* than $\lambda_{\mathrm{fam}}(6)$ by 5.682729×10^{-4}—the certificate class that holds the record demonstrably escapes the plain family ceiling, and the escape mechanism is the regularity family. The next theorem shows the escape is only partial.

Theorem 3 (Level-7 ceiling with regularity rows)**.** *Let A_j, B_j $(j = 1, \dots, 63)$ be the order-5 regularity rows on S_7, normalised so that $\sum_j B_j = 1$ and $\sum_j A_j(H) = d(H)$ for every $H \in S_7$. Then for every $Q \succeq 0$ and every $\mu \in \mathbb{R}^{63}$, with $\rho = 22251/40000$,*

$$\max_{H\in S_7}\Big[\, d(H) + \langle Q, M(H)\rangle + \textstyle\sum_j \mu_j \left(A_j(H) - \rho\, B_j(H)\right)\Big] \;\geq\; \rho \;=\; 0.556275 \;>\; \tfrac{5}{9}.$$

Proof sketch. The certificate is an integer ray $z \in \mathbb{Z}^{170}_{\geq 0}$ (entries up to 65 digits) with $(A_j - \rho B_j)\cdot z = 0$ exactly for all 63 rows, $(d \cdot z)/(\mathbf{1} \cdot z) = \rho$ exactly, and $S_b(z) = \sum_H z_H M_b(H)$ positive semidefinite on all 26 blocks (rational LDL^{T}, integer data). Averaging the bracket against $z/(\mathbf{1} \cdot z)$ gives $\rho + \langle Q, S(z)\rangle/(\mathbf{1} \cdot z) + 0 \geq \rho$, using $\langle Q, S(z)\rangle \geq 0$ for two positive semidefinite matrices. No duality theorem and no closedness assumption is used. The construction is exact by design: 137 ray coordinates were rounded from a floating-point search at denominator 10^{14} and the 33 pivot coordinates were then solved exactly in rational arithmetic, so the ray lies in the kernel of the row system identically. □

Two further targets were certified independently (2781/5000 and 1389/2500), and the pipeline run *above* the transition fails in two independent ways (negative ray entries, and a negative exact pivot), so the control that should fail does fail. A soundness control at $\rho = \frac{5}{9}$ recovers the Turán ray. The 63 rows were re-derived from their specification by an independent canonical form, with the identity gates $\sum_j B_j = 1$, $\sum_j A_j = d$ holding exactly on all 170 classes. Theorem 3 was also confirmed by a third instrument built for a different purpose (Section 3), which reproduced ρ to 1.3×10^{-12} and the regularity gain to 3.7×10^{-8}.

What the ceilings forbid, exactly. Theorems 1–3 forbid $\lambda = \frac{5}{9}$ at levels 6 and 7; they do *not* forbid record improvements, and it is worth computing exactly how much room they leave. Writing g for the gap record $- \frac{5}{9} = 74/15273 = 4.845152 \times 10^{-3}$:

certificate class	floor	unreachable share of g	reachable share
plain level 6	0.560968980070	111.73%	none
plain level 7	0.556811778585	25.93%	74.07%
level 7 + regularity	0.556275000000	14.85%	85.15%

The first row is a consistency check: the bound of (1) already lies below the plain level-6 ceiling. That it can do so is exactly the effect of colour refinement and the regularity rows. The last row is the honest state of the art: 85.15% of the remaining gap is not priced by any ceiling we can prove, and the 14.85% that is priced is Theorem 3. The regularity rows remove at most 5.367786×10^{-4} of the level-7 obstruction; 57.27% of it survives them, now as a theorem.

3 The coloured level-6 program is capped (numerical)

The colour-refined level-6 program that produced (1) has an optimum λ_{col} known a priori to lie in $[\frac{5}{9}, 0.5604007071302298]$. Three independent cutting-plane feasibility runs at targets $\rho = 0.5560,\ 0.5580,\ 0.5595$ each returned a stable positive infeasibility certificate value t^*; since the working-set problem is a relaxation, $t^*(W) \leq t^*(\text{all})$, a positive relaxation value proves infeasibility with no convergence requirement. Conclusion (*measured*):

$$\lambda_{\text{col}} \;>\; 0.5595, \qquad \text{hence} \qquad \lambda_{\text{col}} \in (0.5595,\ 0.5604007071302298],$$

a 5.4-fold narrowing of the bracket. On the same measured footing, the coloured program can improve the record by at most 9.007×10^{-4} and does not reach $\frac{5}{9}$.

Two artefact mechanisms were eliminated. The trace cap was inactive at the deciding iterations (trace at 0.8% of its bound), and a paired control re-running the decisive target with the denominator floor 100× smaller left t^* positive and essentially unchanged ($+1.606 \times 10^{-4}$ at $\rho = 0.5595$). The verdicts held for three consecutive iterations under a pre-registered stability rule.

This result is labelled *measured*, not proved, and we state exactly what is owed. The exact certificate for an infeasibility verdict is a Farkas ray $y \geq 0$ with $S_b(y)$ positive semidefinite, where $S_b(y)$ is determined by y (no Gram factorisation is available, unlike the feasible direction, where exact certification of a bound costs seconds). Four of the five certificate conditions already pass exactly on the extracted ray; the fifth—positive semidefiniteness—we can localise only to a window of perturbation margin ($\approx 3 \times 10^{-14}$, $\approx 5 \times 10^{-9}$), which is five orders of magnitude wide, but available solvers return points with residual $\approx 10^{-9}$, at the top of the window. The blocker is solver precision, not exact arithmetic; we expect the certificate to be completed and will report it separately.

4 Beyond flag algebras: five methods, five obstructions

We examined five standard alternatives. One is settled by a proposition proved here; for each of the others we record the specific obstruction we met, so that a later proposal can be checked against it. Throughout, the useful filter was a single explicit family: colour 0 joined to every pair, one arc $0 \to 1$, and a transitive tournament on $\{1, \ldots, q-1\}$. This pattern is $K_4^{(3)}$-free at every q, has none of the structural properties (degeneracy, Turán triples, profile coincidences) that proposals typically exclude, and its maximum density is $\frac{5}{9} - \frac{4}{81q} + o(1/q)$ with limit equal to Turán's construction. Consequently *no criterion that fails to exclude this family can bound anything away from* $\frac{5}{9}$; every proposal below was run against this family first.

Proposition 4 (Closing selectors must cut the extremal family)**.** *Call a predicate H on 3-graphs a closing selector if the level-6 program restricted to classes satisfying H certifies $\frac{5}{9}$. Every closing selector is false on a class that occurs with positive density inside an extremal object of density exactly $\frac{5}{9}$.*

Proof. The value $\max_{H \in S} [\, d(H) + \langle Q, M(H) \rangle \,]$ is monotone in the class set S: enlarging S can only enlarge the maximum, so a program restricted to $S' \supseteq S$ has value at least that of S. By Theorem 1 the level-6 program already has value $\lambda_{\mathrm{fam}}(6) \geq 0.560968980070 > \frac{5}{9}$ on the family classes S_6 alone. Hence a closing selector cannot retain all of S_6: if it did, its value would be at least $\lambda_{\mathrm{fam}}(6)$ and it would not certify $\frac{5}{9}$. So it deletes some $c \in S_6$, and $H(c)$ is false.

The point is *which* class it must delete. By the definition of S_6, c occurs with positive density in some member O of the extremal family, and every such O has $d(O) = \frac{5}{9}$. So the deletion set of any closing selector meets the extremal family itself: H is false on a configuration that genuinely occurs inside an extremal object. A theorem conditioned on H therefore excludes part of the extremal set, so the hypothesis cannot simply be dropped: the residual case left by any closing selector still contains genuine density-$\frac{5}{9}$ configurations. □

Quantitatively the failure is not marginal: each of the five closing-selector candidates we tested is violated by 56 of 70 extremal-family members of density exactly $\frac{5}{9}$, with violating class densities as large as 103112/253125. The vacuous selector deletes nothing and correctly fails to close, so the test discriminates.

Stability methods. Exact-result stability arguments require an essentially unique (or finite, non-degenerate) extremal configuration. For $K_4^{(3)}$ the extremal set is infinite even modulo isomorphism, and the natural stability number is infinite; the standard two-step architecture cannot start. (Section 6 proposes what we believe is the correct repair: quotient first.) Notably the obstruction is *not* the reweighting continuum—see Theorem 5.

Containers and hypergraph regularity. We are not aware of a published instance of either method improving an upper bound on a Turán density; in container arguments the extremal number is an *input* (the container exponent is $ex(n, F)$ itself), and removal-type consequences of regularity likewise consume rather than produce π. On that evidence — no instance known to us, together with the direction in which both methods consume rather than produce π — we do not see a route to $\pi(K_4^{(3)}) \leq \frac{5}{9}$ here.

Spectral and tensor methods. For every $p > 1$ the p-spectral Turán density coincides with the Turán density itself [19]; the Lagrangian—which this project already computes exactly—is the entire algebraic content. Spectral relaxation is not a distinct method here.

Entropy and the local lemma. The local lemma and entropy compression certify *existence* and so can only serve the (closed) lower half. A genuine entropy method for hypergraph Turán *upper* bounds does exist [4]; evaluated on the tetrahedron, its entropic density coincides with the Lagrangian, so the reformulation is exact, and its one independently computable coordinate has a floor of $2^{4/5}/3 = 0.5803670\ldots$, above the pre-2012 records.

5 Where the obstruction lives

A valid certificate at $\frac{5}{9}$ must be *tight* on every extremal object: averaging $d(H) + \langle Q, M(H)\rangle \leq \frac{5}{9}$ against an extremal ϕ (density exactly $\frac{5}{9}$) forces both $\langle Q, S(y_\phi)\rangle = 0$—hence $Q\, S_b(y_\phi) = 0$, a kernel condition—and $d(H) + \langle Q, M(H)\rangle = \frac{5}{9}$ for every $H \in S_N$. The certificate is thus confined to the orthogonal complement of the span of the extremal moment ranges, and each new extremal object can only shrink that space. The natural expectation, voiced already in [9], is that the sheer multiplicity of extremal configurations is the obstruction. Our measurements contradict this in an instructive way.

Theorem 5 (Turán alone is certifiable)**.** *Restricted to the classes realised by Turán's construction alone, the level-6 program attains* $\frac{5}{9}$ *exactly:* $\lambda(\text{Turán only}, 6) = \frac{5}{9}$*, certified in both directions. By contrast, adjoining the classes of a single Fon-der-Flaass sign pattern already forces*

$$\lambda \geq \frac{279653}{500000} = 0.559306,$$

which is 69.28% *of the entire level-6 family excess.*

Turán's construction—the conjectured extremiser—contributes exactly zero of the obstruction. The obstruction is *concentrated* in the non-Turán extremal objects, and mostly in one of them.

Observation 6 (Saturation: the continuum contributes nothing)**.** Over 210 extremal objects spanning the full reweighting continuum of [21] (including 120 non-uniformly weighted members, each of density exactly $\frac{5}{9}$), the level-6 tightness footprint is constant: $|S_6| = 45$ classes and profile rank 8, with no growth after the 54th object, and none from an exhaustive 4-colour census. The *number* of extremal configurations is the wrong currency; the tightness constraints saturate at a small finite set.

Observation 7 (Rank law)**.** The rank of the family profile span at level N satisfies $\operatorname{rank}(N) \leq 2^{N-3}$, and this upper bound is proved for all N; equality is verified exactly at $N = 4, \ldots, 8$, the case $N = 8$ being the pre-registered blind prediction $\operatorname{rank}(8) = 32$, which was confirmed. A power of two indicates the family's genuine degrees of freedom are the binary sign choices of the Fon-der-Flaass recursion, not continuum parameters. The support ladder $4, 12, 45, 170, 649$ matches no catalogued integer sequence (checked against OEIS with a positive control).

6 The reweighting dichotomy and a quotient conjecture

Write $\operatorname{deficit}(O) = \frac{5}{9} - d(O)$ and, in the pattern formalism, $\deg_i(w) = \frac{1}{3}\,\partial d/\partial w_i$ and $\operatorname{dist}(O) = \sum_i w_i\,|\deg_i - \frac{5}{9}|$, the weight-weighted L^1 degree deviation; the box distance dominates $\operatorname{dist}/2$. This metric reproduces, exactly, the recorded witness formulas: at the Turán reweighting $(\frac{1}{3}+e, \frac{1}{3}-e, \frac{1}{3})$, $\operatorname{deficit} = 2e^2(1+\frac{3}{2}e)$ and $\operatorname{dist} = \frac{4}{9}e(1+\frac{3}{2}e)^2$; and along the sliver edge-removal family, $\operatorname{deficit} \equiv \operatorname{dist} = \frac{5}{3}t - t^2$ identically, whence the constant 2 below.

An earlier stability clause failed on exactly this asymmetry: the deficit is *quadratic* along reweighting directions but distance is linear, so no linear stability constant survives. The refutation was directional, and the following measurement shows how directional.

Observation 8 (Uniform transverse linearity at Turán; exact)**.** Among all fifteen single-orbit sliver deletions at Turán's construction, deficit and distance are both linear in the sliver mass with *equal* leading coefficients: $\operatorname{deficit} = \operatorname{dist} + O(t^2)$ in every direction. The quadratic-deficit valley occurs in the two reweighting directions and only there, and a mixed perturbation retains its linear term. (First-order sweep; three second-order directions have both leading terms zero and are not decided by it.)

Together with Theorem 5 and the saturation observation, this suggests the extremal landscape is a discrete, tree-structured family (the sign patterns) surrounded by a benign continuum (the reweightings, which preserve density exactly and carry every measured pathology but none of the obstruction). The natural repair of stability is then to quotient the continuum away:

Conjecture 9 (Quotient stability)**.** *Let $\mathcal{T}'$ be the closure of the tower of Fon-der-Flaass constructions under measure reweighting, and $\delta_Q(O) = \delta_\square(O, \mathcal{T}')$. There is a constant $c \in (0, 2]$ such that every $K_4^{(3)}$-free limit object O satisfies*

$$\operatorname{deficit}(O) \;\geq\; c\,\delta_Q(O).$$

The witness that refuted the un-quotiented clause has $\delta_Q = 0$ and no longer applies; the sliver family shows $c \leq 2$. We have placed the conjecture under a four-flank adversarial battery with the decisive flank declared in advance: the recorded deeper-rim valleys (bridge macro-shift, intra-pair shift, k=3 class shift) must each have $\delta_Q = 0$, i.e. must themselves be reweightings; a single such move with linear δ_Q refutes Conjecture 9 exactly as its predecessor was refuted. Results will be reported regardless of direction. If the conjecture survives, it supplies the missing first half of a two-step exact-result architecture whose second half—behaviour on and near the family—is precisely what Section 2 quantifies.

6.1 The conjecture holds for the sharpest adversarial family

The most dangerous candidate counterexample we know is the family Ch_q: colour 0 joined to every pair, one arc $0 \to 1$, and a transitive tournament on $\{1, \ldots, q-1\}$. It is $K_4^{(3)}$-free at every q by Lemma 10 below, its limit is Turán's construction, and its deficit is $D_q = \frac{5}{9} - d(\mathrm{Ch}_q) \sim \frac{4}{81q}$, i.e. it approaches the conjectured density along a one-parameter path. If quotient stability fails anywhere near the extremal locus, this family is where it should fail.

Lemma 10 (Ch_q is $K_4^{(3)}$-free, all q)**.** *Writing a triple by the multiset of parts its vertices occupy, the edges of* Ch_q *are exactly* $(0,0,1)$*;* $(0,j,k)$ *with* $1 \leq j < k$*; and* (i,i,k) *with* $1 \leq i < k$*. No four parts* $a \leq b \leq c \leq d$ *have all four triples edges.*

Proof. The edge list is read off the kernel. Suppose $a \leq b \leq c \leq d$ spans a $K_4^{(3)}$. If $a = b = 0$, then $(0,0,c)$ and $(0,0,d)$ are edges only for $c = d = 1$, but $(0,1,1)$ is not an edge. If $a = 0 < b$, then $(0,b,c)$, $(0,b,d)$ and $(0,c,d)$ force $b < c$, $b < d$ and $c < d$, hence $b < c < d$; but a triple of three distinct parts ≥ 1 is never an edge. If $a \geq 1$, every triple has the form (i,i,k) with $i < k$: from (a,b,c) we get $a = b < c$, while from (a,c,d) we get $a = c$, contradicting $a < c$. No case survives. $\square$

Proposition 11 (Quotient stability on Ch_q)**.** $\delta_Q(\mathrm{Ch}_q) = O(D_q)$*. Equivalently there is* $c_{\mathrm{Ch}} > 0$ *with* $D_q \geq c_{\mathrm{Ch}}\, \delta_Q(\mathrm{Ch}_q)$ *for all* q*. Hence* Ch_q *does not refute Conjecture 9.*

The proof is a finite semialgebraic argument on the maximising weights. Its codegree kernel is closed-form, so stationarity at the maximiser becomes two polynomial equations; eliminating the multiplier leaves an exact quadratic relating $a = w_0 - \frac{1}{3}$, w_3 and the ratio $u = w_3/w_2$. Positivity of H_2 and of $\varepsilon = w_0 - M$ confine u to $(1/\sqrt{2}, 1)$. The target $a \leq \frac{39}{100} w_3$ is then established by contradiction: its negation forces $u \geq u_c = (-83 + \sqrt{387289})/634$, and that region is covered by two disjoint contradictions—one from $S_2 = \sum_{j \geq 3} w_j^2 > 0$, one from the tail-mass bound $M \geq \frac{17}{675}$—each certified by exact root counting and Bernstein positivity on rational boxes. The bound $w_3 \leq 0.055$ needed to close the region is itself unconditional for $q \geq 45$, bootstrapped from the stationarity recursion alone; finitely many small q are absorbed into c_{Ch}. Propagating $a = O(w_3)$ through the recursion gives $w_i = \Theta(w_3 i^{-1/3})$, $w_3 = \Theta(q^{-2/3})$, $S_2 = O(1/q)$ and $E_q \leq 6S_2$, which is the stated bound.

We stress the scope. Proposition 11 eliminates one adversarial family by a structural theorem rather than by computation, and it is the family we judged most likely to break the conjecture. It does *not* prove Conjecture 9, and we do not claim it as evidence of any particular strength: the argument leans on Ch_q's closed-form kernel throughout, and a general near-extremal object has no such kernel. The transferable content is the architecture—stationarity, a compact ratio variable, a finite semialgebraic contradiction—not the computation. The natural next question is the smallest class strictly containing Ch_q that still admits an exact finite-state kernel.

7 Three further obstructions, and a classification of certificate classes

The five obstructions of Section 4 were assembled by asking which published *methods* could reach $\frac{5}{9}$. This section reports three further obstructions found by asking a different question — which *certificate classes* can see the extremal object at all — together with the certificate class that survives. Each check reported below as *verified* was carried out in exact rational or integer arithmetic; the provenance labels (P1)–(P3) are the convention stated in §7.5 below.

7.1 The interior-zero obstruction

Observation 12 (Interior zeros kill positive-coefficient certificates)**.** Let F be a form that is nonnegative on the simplex and attains 0 at an interior point. Then no Pólya–Handelman certificate for $F \geq 0$ exists: such certificates require strict positivity, whereas sum-of-squares and copositive certificates may vanish.

This is immediate, and we claim no novelty for it: strict positivity on the simplex is the hypothesis of Pólya's theorem [17] (and of Handelman's representation [18]), and the converse direction is a one-line evaluation — if $(\sum_i x_i)^N F$ has only nonnegative coefficients and F vanishes at an interior point $p > 0$, then evaluating at p makes a sum of nonnegative terms zero, so every coefficient vanishes and $F \equiv 0$. We record it because of what it excludes below, not because it is new.

This is elementary, but it is the operative constraint here, and it appeared three times independently before it was recognised as one mechanism.

1. For the level-6 form of Section 2, $(x_0 + x_1 + x_2)^N(27F - 5s^4)$ has a negative coefficient for every $N \leq 25$, with magnitudes *growing* (-60, -1720, -437976, $-1.21 \cdot 10^9$, $-1.54 \cdot 10^{12}$ at $N = 0, 4, 10, 18, 25$). The cause is that $F - 5/27$ vanishes at the interior point $(1/3, 1/3, 1/3)$. (That $F \geq 5/27$ on the simplex, with equality only there — the hypothesis of Observation 12 — follows from the integer identity $135\,[\,F-(5/27)s^4\,] = \sum_i (3x_i - s)^2\big(14x_i^2 + 44x_{i-1}^2 + 10x_{i+1}^2 + 112\,x_i x_{i-1}\big)$, in which every summand is a square times a form with nonnegative coefficients. The expanded right-hand side itself does have negative coefficients — necessarily, since it equals $5\,(27F - 5s^4)$, whose least coefficient is the -60 recorded above.)

2. In the $q \leq 6$ pattern census, no Pólya–Handelman certificate for max $\leq 5/9$ can exist for any class containing a Turán triple: restricting to the triple's face, nonnegative coefficients together with an interior zero force every coefficient to vanish. That is 32,564 of 41,138 classes, 79.2%.

3. Conversely, Balogh–Clemen–Lidický's ℓ_2 semidefinite certificate *does* close at level 6 despite Turán being an interior extremum, because sum-of-squares certificates are permitted to vanish.

Remark 13. The practical consequence is a positive one. Writing $M(p)_{jk} = \sum_i p_i A[i, j, k]$ for the link matrix of a pattern, copositivity of $C(p) = (5/9)J - M(p)$ on the nonnegative orthant certifies $d \leq 5/9$ at every full-support critical point: there $a_i(w) = \mu$ for every colour i, so for any real p with $\sum_i p_i = 1$ one has $w^\mathsf{T} M(p)w = \sum_i p_i a_i(w) = \mu = d(w)$, and copositivity gives $0 \leq w^\mathsf{T} C(p)w = 5/9 - \mu$. Nonnegativity of p is never used; the gradient ideal enters only through *where* the conclusion is drawn, not as a weakening of the test. Because copositive certificates are permitted to vanish — and 14,652 of the ones below do vanish — Observation 12 does not apply. This certificate settles all 41,138 isomorphism classes of maximal $K_4^{(3)}$-free patterns on six colours *by itself*: for each class there is a single exact rational vector p with $\sum_i p_i = 1$ for which $C(p)$ is copositive, verified in exact arithmetic. No critical-value eliminant, resultant, rational-univariate representation or Sturm–Tarski layer enters the base case at any point. In 41,081 of the 41,138 classes p may be taken from the fixed list of 63 vectors $p = \mathbf{1}_S/|S|$, $\emptyset \neq S \subseteq [6]$, known in advance; the remaining 57 use a general p, of which 19 have a negative entry. Throughout this paper a count of patterns is a count of *labelled* patterns unless the words *class* or *isomorphism class* appear: the 41,138 here are S_6-isomorphism classes at exactly six colours, corresponding to 29,099,574 labelled patterns. Together with the induction on the number of colours — a maximiser of non-full support is a maximiser of the induced pattern on fewer colours, and the maxima at one and two colours are 0 and 4/9 —

$$\text{every } K_4^{(3)}\text{-free pattern on at most six colours has Lagrangian } \le 5/9.$$

The extension of this statement to all q is *equivalent* to Turán's conjecture, by disjoint union with a fresh Turán triple; so it is a base case and not the first rung of a ladder.

7.2 Codegree-isospectral extremal objects

The ℓ_2 method of [3] bounds σ by a linear combination of densities of 4-vertex subgraphs. That places a hard limit on what it can distinguish.

That many non-isomorphic $K_4^{(3)}$-free configurations attain density $\frac{5}{9}$ is classical, and is emphatically not our observation. Kostochka [8] showed that if Turán's conjecture holds there are at least 2^{k-2} pairwise non-isomorphic extremal 3-graphs on $3k$ vertices; further families are due to Brown [5] and Fon-der-Flaass [6], and Frohmader [7] gives 6^{k-1} on $3k+2$ vertices. That abundance is the recognised reason the problem admits no stability theorem. What we add is a sharpening in a different direction: the members are not merely non-isomorphic, they are *indistinguishable by the statistics the ℓ_2 method reads.*

Observation 14 (Codegree-isospectral extremal pairs)**.** Among those configurations there are pairs at density exactly $\frac{5}{9}$ which, beyond being non-isomorphic, share (i) their entire codegree distribution and (ii) the spectrum of their codegree kernel, and agree on *all* induced densities on at most four vertices, first separating at five.

An explicit pair is $S = (-2,-1,7)$ and $S = (-7,-5,3)$ in the Fon-der-Flaass parametrisation [21]: both $K_4^{(3)}$-free, both of density exactly $\frac{5}{9}$, with common codegree distribution — throughout, the distribution of the codegree $d(u,v)$ over all *ordered* pairs of atoms, $u = v$ included, at uniform layer masses — $\{1/3 \mapsto 5/27,\ 4/9 \mapsto 4/27,\ 5/9 \mapsto 4/27,\ 2/3 \mapsto 14/27\}$ and common kernel spectrum $\{\lambda_-^{(2)}, \lambda_+^{(2)}, 0^{(4)}, 5\}$, where $\lambda_\pm = -\frac{1}{2} \pm \frac{\sqrt{33}}{18}$ are the roots of $x^2 + x + \frac{4}{27}$, numerically $-0.819142\ldots$ and $-0.180858\ldots$; non-isomorphism was established by exhausting all $9! = 362{,}880$ bijections and independently by a separating invariant (cyclic versus transitive triangle counts).

The phenomenon is not sporadic. In the sign-matrix parametrisation [21] the negative region is a self-conjugate partition inside the $k \times k$ box, so the level-k family is indexed by the set $A \subseteq \{0,\ldots,k-1\}$ of its Frobenius arm lengths and has exactly 2^k members. Reversal $\rho(A) = \{k-1-a : a \in A\}$ preserves the entire codegree distribution, and also the codegree kernel spectrum (checked exhaustively for $k \le 9$ and $k \le 7$ respectively). It fails to deliver a mate in exactly two ways: on the $2^{\lceil k/2 \rceil}$ *palindromic* A it is the identity, and on the *antipalindromic* ones — those with $\{0,\ldots,k-1\} \setminus A = \rho(A)$, of which there are $2^{k/2}$ for even k and none for odd k — it coincides with complementation $A \mapsto \{0,\ldots,k-1\} \setminus A$, which is the substitution $S \mapsto -S$ and is therefore always an isomorphism, realised by the layer reflection $(a,x) \mapsto (-a,-x)$. Off those two sets ρ is never an isomorphism for $k \le 9$, so it supplies a genuine mate to a fraction $1 - 2^{-\lfloor (k-1)/2 \rfloor} \to 1$ of the members (measured $50, 50, 75, 75, 87.5, 87.5, 93.75$ per cent at $k = 3,\ldots,9$). Reversal is not the only mechanism: the fraction of members admitting *some* non-isomorphic partner of equal codegree distribution is larger still, measured $50, 50, 75, 81.25, 87.5, 90.625, 93.75$ per cent at $k = 3,\ldots,9$. At $k = 6$, for instance, the four arm sets $\{0,1,3\}, \{0,3,4\}, \{1,2,5\}, \{2,4,5\}$ carry one common codegree distribution but split into two isomorphism classes, each of which ρ preserves; the mates there are found by a coincidence the reversal rule does not see.

Corollary 15. *No functional of the codegree distribution, of the codegree kernel spectrum, or of the 4-vertex profile separates such a pair. In particular the ℓ_2 method and Razborov's G_1 separator — which is one entry of that profile — are blind to them. The blindness is specific to 4-vertex functionals: these pairs do separate at level five and above, so this obstructs the ℓ_2 route and the G_1 separator, not the flag hierarchy as such.*

Turán's own fibre is nevertheless clean: over all oriented graphs on nine vertices with no prefilter, exactly 560 objects carry Turán's codegree distribution, and all are blow-ups. The uniqueness conjecture of [3] is therefore untouched by this obstruction; what fails is *identification*, not *selection*.

7.3 The $\ell_2 \to \ell_1$ transfer cannot reach $\frac{5}{9}$

The bound itself is not ours. Balogh, Clemen and Lidický prove $\pi(H)^2 \le \sigma(H) \le \pi(H)$ for every k-uniform H [3, Prop. 1.11(i)], by the Cauchy–Schwarz step given in their Section 4.3, and they determine $\sigma(K_4^{(3)}) = 1/3$ [3, Thm 1.2]. Together these give

$$\pi(K_4^{(3)}) \;\le\; \sqrt{\sigma(K_4^{(3)})} \;=\; \sqrt{1/3} \;=\; 0.57735\ldots$$

What we add is only the residual observation, which is what the present map needs.

Proposition 16 (Residual to [3, Prop. 1.11(i)])**.** *No sharpening of σ improves this route to $\frac{5}{9}$.*

Proof. Equality in the Cauchy–Schwarz step requires the codegree to be constant. Turán's codegree kernel is two-valued, taking $1/3$ with weight $1/3$ and $2/3$ with weight $2/3$; and reaching $\frac{5}{9}$ by this route would require $\sigma \le (5/9)^2 = 25/81 = 0.30864\ldots$, whereas $\sigma = 1/3$ exactly and Turán's construction attains it. ☐

The structural loss is $\sqrt{1/3} - 5/9 = 2.179 \cdot 10^{-2}$, and it has the same source as Observation 12: the extremal object is two-valued where the method is sharp only at constants.

7.4 Summary of the certificate landscape

Certificate class	Sees an interior zero?	Status here
Pólya / Handelman	no	dead on 79.2% of the census
Sum of squares (flag algebra)	yes	closes ℓ_2; provably incomplete at ℓ_1
Copositive link certificate	yes	settles $q \le 6$ (all 41,138 classes)
Functionals of the 4-profile	—	blind by Observation 14

7.5 Two structural results from the minimal-counterexample side

Provenance convention. The remainder of this paper draws on a parallel line of work in this programme, and we label every claim from here on by *how it was established*, since the distinction is not visible from the statements themselves. We state plainly what that parallel line is: unpublished work of the author, not contained in the companion [1] and not yet written up, so a (P3) claim below cannot at present be checked against any document. Such claims should be read as announcements accompanied by first-party arithmetic checks, not as results a reader can verify from the literature; nothing in this paper's own conclusions depends on them, and none of them bounds $\pi(K_4^{(3)})$.

(P1) Proved here. A complete proof appears in this paper.

(P2) Verified here. Checked by an implementation written independently from the prose, in exact arithmetic, over the range stated: an exhaustive enumeration, or a re-check of every coefficient of a supplied certificate. The verification is first-party; the argument being verified need not be.

(P3) Imported. Taken from the parallel line. Its arithmetic is checked here, but the argument establishing it is not re-derived in this paper, and the claim should be read as conditional on that source.

A (P2) label over a finite range never upgrades an all-q assertion: where a statement is verified on a range and asserted beyond it, both labels are given. Words like "closes" and "verifies" below are always qualified by one of these three.

The obstructions above say what cannot certify. Two positive structural statements, obtained in a parallel line of the same project and re-verified here from their statements rather than their code, say something about where a counterexample could live. Both concern the family $\mathrm{Ch}_{q,T}$: the pattern with edges $(0,0,1)$, $(0,i,j)$ for all $1 \le i < j$, and (i,i,j) exactly when $i \to j$ in a tournament T on $\{1,\dots,q-1\}$. Lemma 10 is the transitive member.

Theorem 17 (Complete-apex domination; all-q form imported)**.** *If a $K_4^{(3)}$-free pattern has a colour* 0 *for which every transversal* $(0,i,j)$ *is an edge, then its Lagrangian is at most* $\frac{5}{9}$*. Consequently every* $\mathrm{Ch}_{q,T}$ *is dominated, with maximum exactly* $\frac{5}{9}$*; and in a minimal counterexample every apex must have an incomplete link.*

Provenance: (P2) for $q \le 5$; (P3) for all q. The all-q statement rests on an edge-polynomial argument from the parallel line which is not reproduced here, and a reader who wants this paper self-contained should take the theorem only in the verified range. That range is genuinely exhaustive: we re-verified by enumeration with an independent encoder, over *all* $K_4^{(3)}$-free complete-apex patterns there are 18 at $q = 3$ and 241 at $q = 4$, and 285 maximal ones at $q = 5$; the maximum Lagrangian is 0.555555555556 in every case, attained and never exceeded. The argmax collapses to $(1/3, 1/3, 0, \dots, 0, 1/3)$, i.e. Turán's construction. The statement must exclude $r = 1$, since Ch_1 is not $K_4^{(3)}$-free: $(0,0,1)$ and $(0,1,1)$ complete the 4-multiset $(0,0,1,1)$.

Remark 18 (Terminology, and why the hypothesis must stay a pattern hypothesis)**.** The structure in Theorem 17 is standard and carries published names that we adopt here: a vertex joined to every pair of the remainder is the *apex* of a *cone* [13, Sec. 2] (see also [12, p. 13]), and in the multiset-pattern framework the same hypothesis reads "colour 0 has *complete shadow*" [13].

The hypothesis must be read on patterns, not on simple 3-graphs. On simple 3-graphs the statement is folklore *and has a different constant*: $K_4^{(3)}$-freeness together with a complete link forces every triple avoiding the apex to be a non-edge, so the 3-graph is exactly the star cone($\emptyset$), whose Lagrangian is 4/9 and not $\frac{5}{9}$. That reading is a special case of [13, Thm. 3.1] ($\lambda(\mathrm{cone}(Q)) \le 4/9$ for sparse Q) and of the structural results of [22]. What the star reduction cannot see, and what Theorem 17 is about, are the repeated-colour edges (i,i,j), $(0,i,i)$ and $(0,0,s)$: they survive the hypothesis, they are where the content sits, and $\frac{5}{9}$ is sharp on the pattern class, being attained — the argmax above is Turán's construction. A literature sweep for this revision found no prior statement of the pattern form under any of those names, nor under the contrapositive framings, nor under the complementary $(3,4)$-covering convention. That is a negative on the *indexed* literature only: the $(3,4)$-construction papers [8, 5, 6] are not full-text indexed, so we claim no more than that we found no prior statement.

Proposition 19 (A support ladder with a proved top; moment inequality imported)**.** *In the branch where the minimal counterexample admits no clean witness collision, a moment inequality of the form*

$$\frac{5q}{81}+\frac{11}{36q}\ \le\ 1$$

is necessary; its left side is increasing for $q \ge 3$ and exceeds 1 *from $q = 16$ onward. Hence that branch is confined to $q \le 15$.*

Checked exactly: the left side takes the values 0.287037 at $q = 3$, 0.886023 at $q = 14$, 0.946296 at $q = 15$ and 1.006752 at $q = 16$, and the increments are positive throughout.

Remark 20 (Why the contrast matters)**.** Proposition 19 should be read against Remark 13. There, extending the $q \le 6$ certification to all q was shown to be *equivalent* to Turán's conjecture, so that ladder provably has no top. Here a different formulation — restricted to one branch of a clean-collision dichotomy — has a ladder that terminates at $q = 15$ for a reason internal to the moment method. Two ladders over the same parameter, one with no top and one with a proved top, is the sharpest available illustration that the choice of formulation, not the parameter, is what decides whether a support induction can close. We note the corresponding caution: the branch so terminated is the one that provably contains *no* extremal configuration, since the balanced 2+2+2 pullback at density exactly $\frac{5}{9}$ admits a clean collision.

Remark 21 (Status)**.** Observation 12 and Proposition 16 are proved. Observation 14 is verified exhaustively on the pairs and families stated; its genericity claim is measured to $k = 9$ and is not proved for all k. The $q \le 6$ certification is exact and exhaustive. Theorem 17 is re-verified here exhaustively for $q \le 4$ and over maximal patterns at $q = 5$; its all-q form rests on the edge-polynomial argument of the parallel line, which we have not re-derived. Proposition 19's arithmetic is checked exactly; the moment inequality it rests on is imported, not re-proved here. *No statement in this section bounds $\pi(K_4^{(3)})$ below the value reported in* (1), and novelty is claimed for none of them. A literature sweep for this revision located the prior art for Observation 14's ambient fact — the exponentially many non-isomorphic extremal configurations of [8, 5, 6, 7] — now cited in place; the residual claim there is only the statistical indistinguishability, for which we found no prior statement. The novelty of Theorem 17 has since been priced, in Remark 18: its simple-3-graph reading is folklore and carries the sharper constant 4/9 [13, 22], while for the pattern form we found no prior statement — a negative on the indexed literature, not a novelty claim. The Goodman link-equality material behind the theorem is *not* new and is used here as a cited step rather than as a classification: the vertex-weighted inequality is Goodman's [11], in this form due to Moon and Moser [16, p. 285] and independently to Nordhaus and Stewart [20, Eq. (5)] (restated as Eq. (1.4) of [15]), with equality exactly when the edge count is that of a balanced complete multipartite graph on k parts with k dividing n [15, Eq. (1.4)]. The exact minimisation beyond that bound is the subject of Lovász and Simonovits [14] and of [15].

8 The surviving frontier: adaptive completion closure

The barriers of Sections 4 and 7 say what cannot certify $\frac{5}{9}$; Theorem 17 and Proposition 19 remove two families of configuration. This section states what is left, and gives the one measurement that says where the difficulty now sits.

8.1 The unconditional six-colour base case

We restate Remark 13 with the emphasis it deserves. The copositive link certificate settles *all* 41,138 isomorphism classes of maximal $K_4^{(3)}$-free patterns on six colours, and it does so *with no collision hypothesis and no branch restriction.* It therefore supersedes both branch-restricted small-support results — the exhaustive elimination in the no-clean-collision branch, which reaches $q \leq 5$, and any corresponding statement on the complementary branch — and stands as the unconditional base case at six colours. Together with Theorem 17, a minimal counterexample must have at least seven colours and an incomplete link at every apex.

8.2 Two branches, and where equality lives

Fix the dichotomy on whether the counterexample admits a clean witness collision. Proposition 19 and the associated link-capacity certificates are asserted in the parallel line to close the *no-clean-collision* branch. We state exactly what we checked, since "closes" and "verifies" carry different weights here. The ladder has nine rungs: $q = 3, 4$ and $q = 5$ by exhaustive elimination; combined, refined and full link capacity at $q = 6, 7, 8$; diamond/star hybrids at $q = 9, 10$; the $q = 11$ negative box with 137 Bernstein capacity branches; and the $q \geq 12$ weight cutoff. We re-ran all nine verifiers end to end in exact arithmetic and all nine terminate reporting success — a (P2) check of the supplied certificates, not a (P1) proof of the branch closure. Two rungs are stronger than that: at $q = 4, 5$ we reproduced the pattern counts 68 and 2,304 from an encoder written independently from the prose, and the $q \geq 12$ cutoff we re-derived by hand (Proposition 19). For the remaining rungs we checked the certificates' arithmetic without re-deriving the arguments that generate them. The branch closure as a whole is therefore (P3).

The measurement that matters is the size of the branch so removed. Writing $\mathcal{A}_q$ for the maximal $K_4^{(3)}$-free patterns on q colours with every apex incomplete — the live universe, by Theorem 17 — and $\mathcal{N}_q \subseteq \mathcal{A}_q$ for those with no clean collision:

q	$\lvert\mathcal{N}_q\rvert$	$\lvert\mathcal{A}_q\rvert$	share closed
4	68	112	60.7%
5	2,304	26,164	8.8%
6	8,880	29,067,138	0.0305%

The $q = 4$ row is exhaustive over all 2^{20} patterns and we re-verified it independently (5,688 nonempty $K_4^{(3)}$-free patterns, 224 maximal, 112 apex-incomplete; counts of patterns exclude the empty pattern throughout). The $q = 6$ row is likewise complete: all 29,099,574 labelled maximal $K_4^{(3)}$-free patterns on six colours were generated and tested, of which 29,067,138 have every apex incomplete.

Observation 22 (Equality lives on the collision side)**.** The balanced 2+2+2 pullback — density exactly $\frac{5}{9}$ — admits a clean witness collision, as does every equality configuration we have examined. The branch closed above therefore contains *no* extremal configuration.

This is not a defect of the argument; it is why that branch was closable. But it locates the problem: the surviving branch is the one containing every object at density $\frac{5}{9}$ known to us, and its share of the live universe is increasing in q.

8.3 The remaining problem

Problem 23 (Adaptive completion closure)**.** Show that for a minimal counterexample admitting a clean witness collision, injective maximal completion closure forces one of: owned capacity, a cheaper clean pair, an addable edge, or a closure-dependent nonnegative tangent direction.

Every term in this statement is defined in the parallel line of work described in the provenance convention above, and is not redefined here; we state the problem in that vocabulary so that it can be quoted exactly. A reader who wants the present paper self-contained should read Problem 23 as a pointer rather than as a statement to be attacked from this text alone.

Two features make this the principal remaining object rather than one lemma among several. First, by Observation 22 it is where every extremal configuration lives. Second, the natural per-collision strengthening — that *every* clean collision localises — is false: an explicit five-colour counterexample at $w = (1, 2, 2, 1, 1)/7$ refutes it under full support, exact degree regularity, strict tangent stability and edge maximality simultaneously. Only an existential, minimum-selecting form survives, and the tangent direction must adapt to the full completion closure rather than being fixed in advance.

8.4 A caution on reformulation

Three independent reductions in this programme have terminated in a statement whose supremum over its own class is still exactly $\frac{5}{9}$: the selector reduction of Proposition 4; the all-q extension of the six-colour certification (Remark 13); and the minimal-counterexample apex lemma that would close Problem 23. Each reduction is correct. None is progress.

Remark 24 (A test for candidate lemmas)**.** If the supremum of the target functional over the class quantified by a proposed lemma is still $\frac{5}{9}$ *by construction*, then proving the lemma proves the conjecture, and the reduction has moved the difficulty without reducing it. We have found this the single most useful filter to apply before committing effort: it is cheap, and it retired several attractive reformulations here.

8.5 Feeder routes

Two further lines remain open and feed Problem 23 rather than replacing it: completing the compact Bernstein cover in the local tournament-barrier analysis, and the descent below the support bound of Proposition 19 on the complementary branch. Neither is expected to close the problem alone; both narrow the configuration space that Problem 23 must handle.

Remark 25 (Status of Section 8)**.** In the labels of the provenance convention: the six-colour base case is (P1) with a machine-checked exact certificate over all 41,138 classes, produced and verified in this programme. The nine-rung ladder is (P2) as a check of the supplied certificates and (P3) as a branch closure, as detailed above; within it, the $q \leq 5$ eliminations and the $q \geq 12$ cutoff are additionally (P2) and (P1) respectively, and the 137 Bernstein capacity branches at $q = 11$ are (P2) — we confirm that all 137 verify as supplied and have not re-derived the tail-incidence restriction that generates them. The counts $|\mathcal{N}_q|$ are (P2) from our own encoder, as is $|\mathcal{A}_6|$: the $q = 6$ enumeration was run to completion, giving 29,099,574 labelled maximal patterns of which 29,067,138 are apex-incomplete. The five-colour refutation of the universal per-collision form is (P2): we re-checked the witness $w = (1, 2, 2, 1, 1)/7$ and its violation margin $5/882 > 0$ in exact arithmetic against all four stated hypotheses. Observation 22 is (P2) over the configurations examined and is explicitly not a claim about all equality objects. Problem 23 is open. As in

Remark 21, *nothing in this section bounds* $\pi(K_4^{(3)})$ *below* (1), and novelty is asserted for no statement here.

9 Limitations

What is exact and what is not. Theorems 1, 2, 3 and 5 and Proposition 4 are certified in exact arithmetic. The capping of the coloured program (Section 3) is numerical with two artefact controls passed and an explicit account of the missing certificate. The rank-law upper bound is proved for all N; the matching lower bound and the transverse-linearity sweep are exact computations at finitely many points and are labelled measured, not proved. Conjecture 9 is a conjecture; Proposition 11 resolves it only on the family Ch_q. That proposition's certificates are exact, but were produced by a single computer-algebra system and await independent re-verification. Its $K_4^{(3)}$-freeness is proved for all q (Lemma 10).

Incompleteness. By Hatami–Norine [10] the semantic cone of valid density inequalities is not recursively enumerable; no finite hierarchy of the flag-algebra kind is complete. (Changing the norm changes the difficulty entirely: the ℓ_2 analogue of the same question is solved [3], on Turán's own construction.) Our ceilings are consistent with, and quantify an instance of, that phenomenon: they show the incompleteness biting at the two levels anyone can compute, in the exact-tightness form predicted by [9]. Whether *some* finite level certifies $\frac{5}{9}$ is equivalent to the conjecture itself and is untouched by everything here; note the one-sidedness that refuting "no finite level works" would itself prove the conjecture.

Methodology. Every positive claim in this paper carrying a (P1) or (P2) label survived an independent adversarial reproduction; claims labelled (P3) are, by the definition of that label, not re-derived here. Throughout, *independent* means independently implemented, never independently authored. Several claims that did not survive were discarded, including two of our own earlier certificates whose control had been evaluated on their own fit set. We record the practice because it materially changed the paper: claims that did not reproduce were discarded rather than weakened, and the discards outnumbered the survivors among the positive results produced late in this work.

10 Concluding remarks

The map now has the following shape. Below the record (1), the plain flag hierarchy is closed at levels 6 and 7 by exact theorems; the regularity-carrying class is capped at level 7 but retains 85% of the remaining gap as unpriced territory; the coloured level-6 program is capped numerically at 0.5595; and the five standard alternatives outside flag algebras fail for named reasons, one of them provably. The obstruction itself is small and structured: zero of it comes from Turán's construction, most of it from a single sign pattern, none of it from the extremal continuum. One structural theorem has been extracted from that picture: the sharpest adversarial family we know, Ch_q, which approaches $\frac{5}{9}$ with deficit $D_q \sim \frac{4}{81q}$, is eliminated as a counterexample to quotient stability by an exact stationarity argument on its finite-state kernel (Proposition 11); this is a theorem about that family, not about the conjecture.

Two further layers sit on that map. Section 7 indexes the obstruction by *certificate class* rather than by method, which changes what a negative result means: an extremal object at an interior zero defeats every Pólya–Handelman certificate while leaving sum-of-squares and copositive ones intact; the $\ell_2 \to \ell_1$ transfer provably cannot reach $\frac{5}{9}$; and the extremal patterns at density exactly $\frac{5}{9}$ — classically abundant [8, 5, 6, 7] — contain non-isomorphic pairs that no codegree statistic, codegree-kernel spectrum, or four-point density separates. The class that survives the filter — copositivity of the link matrix $C(p)$ — settles all 41,138 isomorphism classes of maximal patterns on six colours with no branch hypothesis and with no critical-value elimination, and that is the strongest unconditional statement here.

Section 8 then says what is left, and this is the conclusion we would most like tested. After complete-apex domination and the closure of the no-clean-collision branch, every configuration at density exactly $\frac{5}{9}$ known to us lies in the single surviving branch, whose share of the live universe grows with q, while the terminated branch contains no equality object at all. The principal remaining object is thus the adaptive completion closure of Problem 23 — carrying the caution of Section 8.4, that three reductions in this programme have terminated in statements whose supremum over their own class is still exactly $\frac{5}{9}$, so that a correct reduction need not be progress. What a proof of the conjecture must supply—exact tightness across a saturated, finite, tree-structured set of constraints, in a calculus that provably cannot enumerate all valid inequalities—is now stated precisely enough that a future proposal can be checked against it without re-deriving the barriers behind it. That, as much as the improved bound, is what we hope the paper contributes.

Acknowledgements

The computations were carried out on the author's machines; the adversarial-verification workflow used large language model agents whose every load-bearing number was re-verified in exact arithmetic as described.

References

[1] D. E. Thompson, *An exactly certified upper bound of* 0.5604007071302298 *for the tetrahedron density, and the structure that stops it*, Zenodo (2026), doi:10.5281/zenodo.22196487.

[2] R. Baber, *Turán densities of hypercubes*, arXiv:1201.3587 (2012).

[3] J. Balogh, F. C. Clemen and B. Lidický, *Solving Turán's tetrahedron problem for the ℓ_2-norm*, arXiv:2108.10408.

[4] T. Chao and H. Yu, *An entropy proof of a Turán-type theorem*, arXiv:2412.08075.

[5] W. G. Brown, *On an open problem of Paul Turán concerning* 3*-graphs*, in: Studies in Pure Mathematics, Birkhäuser (1983), 91–93.

[6] D. G. Fon-der-Flaass, *A method for constructing* $(3,4)$*-graphs*, Mat. Zametki **44** (1988), 546–550; English transl. Math. Notes **44** (1988), 781–783.

[7] A. Frohmader, *More constructions for Turán's* $(3,4)$*-conjecture*, Electron. J. Combin. **15** (2008), #R137; arXiv:0806.4208.

[8] A. V. Kostochka, *A class of constructions for Turán's* $(3,4)$*-problem*, Combinatorica **2** (1982), 187–192.

[9] V. Falgas-Ravry and E. R. Vaughan, *Applications of the semi-definite method to the Turán density problem for* 3*-graphs*, arXiv:1110.1623.

[10] H. Hatami and S. Norine, *Undecidability of linear inequalities in graph homomorphism densities*, J. Amer. Math. Soc. **24** (2011), 547–565.

[11] A. W. Goodman, *On sets of acquaintances and strangers at any party*, Amer. Math. Monthly **66** (1959), 778–783.

[12] P. Keevash, *Hypergraph Turán problems*, in *Surveys in Combinatorics 2011*, 83–139.

[13] X. Liu and D. Mubayi, *The number* $4/9$ *is a non-jump for* 3*-graphs*, arXiv:2605.13567 (2026).

[14] L. Lovász and M. Simonovits, *On the number of complete subgraphs of a graph II*, in: Studies in Pure Mathematics, Birkhäuser (1983), 459–495.

[15] H. Liu, O. Pikhurko and K. Staden, *The exact minimum number of triangles in graphs of given order and size*, Forum Math. Pi **8** (2020), e8.

[16] J. W. Moon and L. Moser, *On a problem of Turán*, Publ. Math. Inst. Hungar. Acad. Sci. **7** (1962), 283–287.

[17] G. Pólya, *Über positive Darstellung von Polynomen*, Vierteljschr. Naturforsch. Ges. Zürich **73** (1928), 141–145.

[18] D. Handelman, *Representing polynomials by positive linear functions on compact convex polyhedra*, Pacific J. Math. **132** (1988), 35–62.

[19] V. Nikiforov, *Analytic methods for uniform hypergraphs*, Linear Algebra Appl. **457** (2014), 455–535.

[20] E. A. Nordhaus and B. M. Stewart, *Triangles in an ordinary graph*, Canad. J. Math. **15** (1963), 33–41.

[21] A. A. Razborov, *On the Fon-der-Flaass interpretation of extremal examples for Turán's* $(3,4)$*-problem*, arXiv:1008.4707 (2010).

[22] A. A. Razborov, *On 3-hypergraphs with forbidden 4-vertex configurations*, SIAM J. Discrete Math. **24** (2010), 946–963.

[23] J. Hladký, D. Kráľ and S. Norin, *Counting flags in triangle-free digraphs*, Combinatorica **37** (2017), 49–76.

[24] A. A. Razborov, *Flag algebras*, J. Symbolic Logic **72** (2007), 1239–1282.