arXiv:2610.07049v2 [math.CO] 19 Nov 2026

On possible Turán densities under general uniformity conditions

Heng Li Xizhi Liu

November 20, 2026

Abstract

The Turán problem asks how dense a hypergraph can be while avoiding a prescribed family of subhypergraphs. For a family $\mathcal{S} \subseteq 2^{[r]}$, an r-graph is $\mathcal{S}$-uniformly dense in the sense of Reiher, Rödl, and Schacht if, in every large collection of ordered vertex tuples selected by joint restrictions on the positions in each $S \in \mathcal{S}$, at least a prescribed proportion of the tuples form edges. Ordinary and minimum ℓ-degree Turán densities correspond to $\mathcal{S} = \{\{1\}\}$ and $\mathcal{S} = \{\{1, \dots, \ell\}\}$, respectively, for $1 \le \ell < r$. This survey collects results on possible Turán densities within this framework and gives a complete description of these values for many choices of $\mathcal{S}$. All proofs presented here were generated by GPT.

Contents

1 Introduction 2

2 Parameter identities and shared tools 5

I Ordinary and minimum-degree Turán densities 9

3 Ordinary Turán densities and a shared construction 9

4 The ordinary three-uniform interval 25

5 Minimum-degree densities and the three-uniform case 44

6 The four-uniform minimum 2-degree case 54

7 Minimum 2-degrees in higher uniformities 58

II Uniform Turán densities for complete layers 65

8 Palette tools and finite-family realization 65

9 Finite attainment and nonclosedness 71

10 The upper interval at the highest nontrivial level 79

11 Intermediate uniform densities: realization and construction 93

12 Proofs of the palette realization theorems 97

III Turán densities for general test systems 113

13 Completion degrees and conditional densities 113

14 Three-uniform vertex-pair density 114

15 Classification of general test systems 123

1 Introduction

This survey describes the possible values of Turán densities under general uniformity conditions. It brings together ordinary, minimum-degree, and uniform Turán densities, focusing on their classification, interval constructions, finite attainment, and nonclosedness.

Let $r \geq 2$. An *r-graph* is a collection of r-subsets of a finite vertex set. We identify a hypergraph with its edge set, write $v(H) \coloneqq |V(H)|$, and use $|H|$ for its number of edges.

Let H be an r-graph on a vertex set V of size n, and fix $\mathcal{S} \subseteq 2^{[r]}$. The members of $\mathcal{S}$ specify which sets of positions may be restricted when testing the density of H. For each $S \in \mathcal{S}$, choose an arbitrary set $G_S \subseteq V^S$ of allowed assignments to these positions, and write $\boldsymbol{G} = (G_S)_{S\in\mathcal{S}}$. Let $K_{\mathcal{S}}(\boldsymbol{G})$ consist of the ordered tuples $(v_1, \dots, v_r)$ of distinct vertices satisfying $(v_i)_{i\in S} \in G_S$ for every $S \in \mathcal{S}$.

For $K \subseteq V^r$, let $e_H(K)$ count the tuples in K whose entries form an edge of H. Given $d \in [0,1]$ and $\eta > 0$, we say that H is *$(d, \eta, \mathcal{S})$-dense* if

$$e_H(K_{\mathcal{S}}(\boldsymbol{G})) \geq d|K_{\mathcal{S}}(\boldsymbol{G})| - \eta n^r \qquad \text{for every } \boldsymbol{G}.$$

Let $\mathcal{F}$ be a family of r-graphs, each with at least one edge; the family may be infinite or empty. An r-graph is $\mathcal{F}$-free if it contains no member of $\mathcal{F}$ as a subgraph, where containment need not be induced. The *$\mathcal{S}$-uniform Turán density* $\pi_{\mathcal{S}}(\mathcal{F})$ is the supremum of all $d \in [0,1]$ such that, for every $\eta > 0$, there are arbitrarily large $\mathcal{F}$-free $(d, \eta, \mathcal{S})$-dense r-graphs. We suppress the fixed uniformity r in this notation.

We study the set

$$\Pi^{(r)}_{\mathcal{S},\infty} \coloneqq \{\pi_{\mathcal{S}}(\mathcal{F})\colon \mathcal{F} \text{ is a family of } r\text{-graphs}\}.$$

The empty forbidden family gives density one. There are only countably many finite forbidden families, so realizing a nondegenerate interval requires allowing infinite families.

Reiher, Rödl, and Schacht [43, Proposition 2.4(b)] observed that it suffices to consider *antichains*, that is, families in which no member properly contains another. Indeed, the restriction indexed by A can be absorbed into that indexed by B whenever $A \subsetneq B$. Thus replacing $\mathcal{S}$ by its inclusion-maximal members preserves the $(d, \eta, \mathcal{S})$-dense condition for every fixed H, d, η, and consequently both $\pi_{\mathcal{S}}(\mathcal{F})$ and $\Pi^{(r)}_{\mathcal{S},\infty}$.

The highest nontrivial complete layer, $\mathcal{S} = \binom{[r]}{r-2}$, is described by palette Lagrangians. A finite *k-palette* is a relation $P \subseteq C(P)^k$ on a nonempty finite color set; its tuples are ordered and may repeat colors. Its Lagrangian is

$$\lambda(P) \coloneqq \max\left\{ \sum_{(c_1,\dots,c_k)\in P} \prod_{i\in[k]} x_{c_i} \colon x_c \geq 0,\ \sum_{c\in C(P)} x_c = 1 \right\}.$$

Let $\mathcal{D}_r$ be the set of Lagrangians of finite r-palettes; that is,

$$\mathcal{D}_r \coloneqq \{\lambda(P)\colon P \text{ is a finite } r\text{-palette}\}.$$

Theorem 1.1. *There exist integers $r_0, r_1 \geq 4$ such that the following statements hold.*

(i) *Suppose that $r \geq 3$ and $\mathcal{S}$ is a nonempty antichain of nonempty proper subsets of $[r]$. Then $\Pi^{(r)}_{\mathcal{S},\infty} = [0,1]$ if and only if $|\bigcup \mathcal{S}| \geq 2$ and $\mathcal{S} \neq \binom{[r]}{j}$ for every $1 \leq j < r$.*

(ii) *Suppose that $r \geq r_0$ and that either $\mathcal{S} = \{\varnothing\}$ or $\mathcal{S} = \{\{1\}\}$. Then $\Pi^{(r)}_{\mathcal{S},\infty} = \{0\} \cup [r!/r^r, 1]$.*

(iii) *Suppose that $\mathcal{S} = \{\varnothing\}$ or $\mathcal{S} = \{\{1\}\}$. Then $[4/9, 1] \subseteq \Pi^{(3)}_{\mathcal{S},\infty}$. Consequently, $[2r!/r^r, 1] \subseteq \Pi^{(r)}_{\mathcal{S},\infty}$ for every $r \geq 3$.*

(iv) *Suppose that $r \geq r_1$, $1 \leq j \leq r-3$, and $\mathcal{S} = \binom{[r]}{j}$. Then $\Pi^{(r)}_{\mathcal{S},\infty} = \{0\} \cup [N^{-N}, 1]$, where $N \coloneqq \binom{r}{j+1}$.*

(v) *Suppose that $r \geq 3$ and $\mathcal{S} = \binom{[r]}{r-2}$. Then $\Pi^{(r)}_{\mathcal{S},\infty} = \big(\mathcal{D}_r \cap [0, 4r^{-r})\big) \cup [4r^{-r}, 1]$.*

Remark. Li, Liu, and Mubayi [22, Theorem 1.1] proved Theorem 1.1(ii) with $r_0 = 18$. The finite-palette argument at the end of the proof of Theorem 1.1(iv) then gives $r_1 = 8$, since $\binom{r}{j+1} - 10 \geq \binom{r}{2} - 10 \geq 18$ for $r \geq 8$ and $1 \leq j \leq r-3$. These bounds need not be optimal.

At the top layer, the vanishing result recalled by Reiher, Rödl, and Schacht [43, Section 1.4 and Proposition 2.5] gives $\Pi^{(r)}_{\binom{[r]}{r-1},\infty} = \{0,1\}$. We write $\Pi^{(3)}_{\mathrm{vp},\infty}$ for the set of vertex-pair Turán densities.

The familiar parameters and their possible values are summarized in Table 1. Here $r \geq 3$, $2 \leq \ell < r$, $1 \leq j < r$, and $N = \binom{r}{j+1}$; the last two rows concern $r = 3$. The parameter identities are asymptotic and need not hold for the density conditions on a fixed finite hypergraph. We justify them in Section 2.

Test system $\mathcal{S}$	Corresponding parameter	$\Pi^{(r)}_{\mathcal{S},\infty}$
$\{\varnothing\}$ or $\{\{1\}\}$	Ordinary density $\pi = \gamma^{(r)}_1$	$\{0\} \cup [\frac{r!}{r^r}, 1]$ for $r \geq r_0$; $\supseteq [\frac{4}{9}, 1]$ for $r = 3$; $\supseteq [\frac{2r!}{r^r}, 1]$ for $r \geq 3$
$\{[\ell]\}$	Minimum ℓ-degree density $\gamma^{(r)}_\ell$	$[0,1]$
$\{\{1\}, \dots, \{\ell\}\}$	Independent vertex-set density	$[0,1]$
$\binom{[r]}{j}$	Conventional j-uniform density $\pi^{(r)}_j$	$\{0\} \cup [N^{-N}, 1]$ for $r \geq r_1$ and $j \leq r-3$; $\big(\mathcal{D}_r \cap [0, 4r^{-r})\big) \cup [4r^{-r}, 1]$ for $j = r-2$; $\{0,1\}$ for $j = r-1$
$\{\{1,2\},\{2,3\}\}$	Cherry density	$[0,1]$
$\{\{1\},\{2,3\}\}$	Vertex-pair density π_{vp}	$[0,1]$

Table 1: Special cases of $\mathcal{S}$-uniform Turán density and their possible values.

The classical extremal problem goes back to Turán's 1941 paper [50]. It asks for the largest number $\mathrm{ex}(n, \mathcal{F})$ of edges in an $\mathcal{F}$-free r-graph on n vertices. Katona, Nemetz, and Simonovits [16] showed by averaging over induced subgraphs that the normalized extremal numbers are nonincreasing. Their limit is the *ordinary Turán density*

$$\pi(\mathcal{F}) \coloneqq \lim_{n\to\infty} \frac{\mathrm{ex}(n, \mathcal{F})}{\binom{n}{r}}.$$

The degree version asks how many extensions every ℓ-set can be required to have. Fix $1 \leq \ell < r$. For $A \in \binom{V(H)}{\ell}$, put $d_H(A) \coloneqq |\{e \in H \colon A \subseteq e\}|$ and $\delta_\ell(H) \coloneqq \min_{A \in \binom{V(H)}{\ell}} d_H(A)$. The *minimum ℓ-degree Turán density* is

$$\gamma^{(r)}_\ell(\mathcal{F}) \coloneqq \lim_{n\to\infty} \max\left\{ \frac{\delta_\ell(H)}{\binom{n-\ell}{r-\ell}} \colon H \text{ is } \mathcal{F}\text{-free},\ v(H) = n \right\}.$$

Mubayi and Zhao [32, Proposition 1.2] studied the codegree case $\ell = r-1$ and proved convergence for arbitrary forbidden families. Lo and Markström [28, Proposition 1.5] treated general ℓ and proved convergence in the same setting. For $\ell = 1$, averaging degrees and deleting vertices of small degree recover the ordinary density; for larger ℓ, the minimum-degree parameters give different extremal problems.

A further direction imposes lower density bounds throughout the hypergraph. Erdős and Sós [7] considered extremal questions requiring every sufficiently large vertex subset to span many edges. Mubayi and Rödl [31, Definition 1.1] distinguished a one-sided lower bound from the two-sided condition of uniform edge distribution and studied how these conditions pass to induced subhypergraphs. For the hierarchy considered by Reiher, Rödl, and Schacht [43, Definition 1.6], fix $1 \leq j < r$ and let $\mathcal{K}_r(G)$ be the r-sets whose j-subsets all belong to a j-graph G on $V(H)$. The condition

$$|H \cap \mathcal{K}_r(G)| \geq d|\mathcal{K}_r(G)| - \eta n^r \qquad \text{for every } G$$

defines the conventional j-uniform Turán density $\pi^{(r)}_j$ using the same quantifiers on η and the host order as above. When $j = 1$, G is a vertex subset and $\mathcal{K}_r(G)$ consists of all its r-subsets, giving the usual uniformly dense condition.

These tests also occur in hypergraph regularity. Frankl and Rödl [11] proved a regularity lemma for 3-graphs. Peng, Rödl, and Skokan [38] gave a proof of the corresponding counting lemma by sampling

dense substructures that retain regularity. Rödl and Skokan [47] extended the regularity lemma to arbitrary uniformities, and Nagle, Rödl, and Schacht [33] proved the accompanying counting lemma. Gowers [13] obtained regularity and counting lemmas by a different approach.

In the formulation of Rödl and Schacht [45, Section 2.1], the relative density of a 3-graph is measured on the triangles supported by 3 bipartite graphs between 3 disjoint vertex classes. Regularity requires this density to change little when the bipartite graphs are replaced by subgraphs that still support a sufficiently large proportion of the original triangles. These tests restrict the 3 pairs of positions, as in $\mathcal{S} = \binom{[3]}{2}$; in higher uniformities, the corresponding tests use hypergraphs on proper subsets of the positions. Regularity uses two-sided relative estimates, whereas $(d, \eta, \mathcal{S})$-density requires a lower bound for every test, with error ηn^r.

Hypergraph quasirandomness has several inequivalent forms. Lenz and Mubayi [19] studied the implications among these properties. Towsner [49] organized two-sided quasirandomness conditions using antichains of coordinate sets and their associated σ-algebras, and characterized them through subhypergraph counts. Lenz and Mubayi [20, Section 1.1] defined layouts indexed by antichains and imposed one-sided discrepancy bounds in their study of perfect packings, explicitly relating their definitions to Towsner's. Their layouts use unordered hypergraphs to select ordered vertex tuples.

Reiher, Rödl, and Schacht [43, Sections 1.4–2] studied a common extension of Mantel's theorem and the uniformly dense 3-uniform problem. Their proof uses separately chosen test relations on the $(r-2)$-subsets of the vertex positions, a formulation convenient for applying hypergraph regularity. In their Definitions 2.1–2.2, they formulated $(d, \eta, \mathcal{S})$-density using arbitrary ordered relations $G_S \subseteq V^S$ and introduced the associated Turán parameter $\pi_{\mathcal{S}}(F)$. We adopt this ordered formulation for arbitrary forbidden families. Restricting the tests to injective tuples changes their counts by at most $O_r(n^{r-1})$ and leaves the limiting density unchanged.

For $\mathcal{S} = \{\varnothing\}$, the only nonempty test consists of all ordered r-tuples of distinct vertices. A single face $\mathcal{S} = \{[\ell]\}$ instead selects an arbitrary relation of ℓ-roots and sums their numbers of extensions. Replacing that face by $\mathcal{S} = \{\{1\}, \dots, \{\ell\}\}$ restricts the root relation to a product of separately chosen vertex subsets. In particular, testing one vertex position gives the ordinary parameter, whereas testing all r positions independently gives the conventional uniformly dense parameter. The single-face identity is proved in Lemma 2.1. The complete-layer identity is due to Reiher, Rödl, and Schacht [43, Proposition 2.5]; Section 2 explains its extension to arbitrary forbidden families.

Write $\Pi_\infty^{(r)}$ and $\Gamma_\ell^{(r)}$ for the sets of values of π and $\gamma_\ell^{(r)}$ over arbitrary forbidden families, and put $\Pi_3 \coloneqq \Pi_\infty^{(3)}$ and $\alpha_r \coloneqq r!/r^r$. Erdős's theorem [5] gives the ordinary gap $(0, \alpha_r)$. Pikhurko [40] proved that $\Pi_\infty^{(r)}$ is closed and has cardinality continuum for $r \geq 3$, and Liu and Pikhurko [27, Theorem 1.1] proved that it contains a terminal interval. The proof of Theorem 1.1(ii) starts with the initial interval in Theorem 3.1, valid for every $r \geq 10$.

Earlier work on non-jumping numbers studied where ordinary Turán densities accumulate from above. Peng [34, 35] used hypergraph Lagrangians to construct further infinite sequences of non-jumps, and Peng and Zhao [39] gave a recursive procedure for generating further examples from known ones.

Theorem 1.1(iii) proves the upper-interval assertion of the non-jump conjecture of Liu and Mubayi [26]. Its proof starts with sparse hosts whose local cone optimization gives an interval at $4/9$, then uses substitutions and a link operation to extend it to one.

Mubayi and Zhao [32, Theorem 1.6] proved that the sets of codegree Turán densities are dense in $[0, 1)$ for every $r \geq 3$, and Lo and Markström [28, Theorem 1.7] extended this to every $r > \ell \geq 2$. Gao, Pikhurko, Rong, and Sun [12] realized every rational codegree value by a finite family. Lin and Zhou [25] independently proved the corresponding classification for $\ell > 2$, with the additional requirement that every finite subfamily has strictly larger density. These results exclude the empty forbidden family and hence the value one.

For the simultaneous statement below, a *monotone class* is closed under isomorphisms and taking subgraphs and contains every edgeless hypergraph. It is the class avoiding its finite nonmembers, so the density notation also applies to such classes.

Theorem 1.2. *For every $r \geq 3$, $2 \leq \ell < r$, and $\alpha \in [0, 1]$, there is a monotone class $\mathscr{P}$ with $\gamma_\ell^{(r)}(\mathscr{P}) = \pi_{\{\{1\},\dots,\{\ell\}\}}(\mathscr{P}) = \alpha$. In particular, $\Gamma_\ell^{(r)} = [0, 1]$ for every $r > \ell \geq 2$.*

The second parameter tests independent vertex subsets. Any system supported on $\ell < r$ positions lies between this parameter and minimum ℓ-degree, so the same class realizes all such systems simultaneously. This proves the unused-position cases of Theorem 1.1(i). The assertion does not include the additional strict finite-subfamily requirement in the conjecture of Mubayi and Zhao [32, Conjecture 1.7].

Li and Liu [21] proved the following finite-attainment and nonclosedness results. Combining finite attainment with the construction in Section 10 gives Theorem 1.1(v) for every $r \geq 3$.

Write $\Pi^{(r)}_{j,\infty}$ for the set of possible values of $\pi^{(r)}_j$. To state finite attainment, order the vertices of an r-graph and color its $(r-1)$-subsets by $C(P)$. An edge $\{v_1 < \cdots < v_r\}$ is admissible when its face-color tuple, in the order of deletion of $v_r, \ldots, v_1$, belongs to P. A palette is *$\mathcal{F}$-deficient* if no member of $\mathcal{F}$ has such a coloring. We use the palette characterizations of Lamaison [18] and Lin, Sun, Wang, and Zhou [24].

Theorem 1.3. *The following statements hold.*

(i) *For every $r \geq 3$, the values in $\Pi^{(r)}_{r-2,\infty} \cap [0, 4r^{-r}]$ form a countable set of algebraic numbers; each is a finite r-palette Lagrangian and is realized by a finite forbidden family. If $\pi^{(r)}_{r-2}(\mathcal{F}) < 4r^{-r}$, there is a finite $\mathcal{F}$-deficient r-palette P with $\lambda(P) = \pi^{(r)}_{r-2}(\mathcal{F})$.*

(ii) *For every $r \geq 3$, the set $\Pi^{(r)}_{r-2,\infty}$ omits continuum many of its limit points below $4r^{-r}$. For all sufficiently large r, $\overline{\Pi^{(r)}_{r-2,\infty}} = \overline{\mathcal{D}_r} = \{0\} \cup [r^{-r}, 1]$.*

Finite attainment comes from excluding long ladders and then extracting a finite palette that attains the limiting optimum. Above $4r^{-r}$, compatible palette refinements give exact realization. Combining finite attainment with the ordinary interval gives the missing limit points in Theorem 1.3(ii).

The initial-interval theorem in Theorem 11.2 applies to a wider parameter range than the classification in Theorem 1.1(iv). The construction uses the same 10-uniform optimization as the ordinary initial interval, with 3 indexed coordinates. Finite and countable palette realization theorems turn its values into densities. Finite palette products extend the interval, and the numerical gap in Proposition 11.1 excludes the remaining positive values.

For systems covering every position, the test system determines which palette coordinates can be prescribed jointly. These restrictions define the corresponding conditional palette capacities. A combinatorial analysis of these coordinates completes the classification for $r \geq 4$. When $r = 3$, symmetric vertex templates also settle cherry density, while vertex-pair density requires a directional construction and cyclic substitutions.

Section 2 gives the common parameter identities and limit and interval principles. Part I treats ordinary and minimum-degree densities. In the minimum 2-degree case, symmetric ternary templates precede the constructions in higher uniformities. Part II treats complete layers; at intermediate levels we present the realization statements and density constructions before their regularity proofs. Part III constructs completion capacities, proves the vertex-pair case, and concludes with Theorem 1.1(i). The verification scripts and exact rational certificate data are available on GitHub.[1] Densities defined by a single forbidden hypergraph and structural characterizations of zero density are outside the scope of this paper.

2 Parameter identities and shared tools

We establish the parameter identities in Table 1, finite-family approximation for compatible limits, and the interval and polynomial identities used throughout the constructions. All density identities apply to arbitrary forbidden families.

First, suppose that $A, B \in \mathcal{S}$ and $A \subsetneq B$. Replacing G_B by $\{\boldsymbol{v} \in G_B \colon \boldsymbol{v}|_A \in G_A\}$ absorbs the restriction on A. Repeating this operation leaves only inclusion-maximal faces and preserves the possible test sets. This is the antichain reduction of Reiher, Rödl, and Schacht [43, Proposition 2.4(b)]. Relabeling the positions also leaves $\pi_{\mathcal{S}}$ unchanged. More generally, if every member of $\mathcal{S}$ is contained in a member of $\mathcal{T}$, assigning each restriction to a containing face gives $\pi_{\mathcal{T}} \leq \pi_{\mathcal{S}}$.

[1] https://github.com/xliu2022/xliu2022.github.io/tree/main/turan-density-verification

2.1 Parameter identities

The following lemma gives the single-face identity. Its reverse implication uses an induced sample to eliminate the exceptional roots allowed by the error term in a uniformity test.

Lemma 2.1. *For every family $\mathcal{F}$ of r-graphs and every $1 \le \ell < r$,*

$$\pi_{\{[\ell]\}}(\mathcal{F}) = \gamma_\ell^{(r)}(\mathcal{F}).$$

Proof. Summing a minimum ℓ-degree bound over the allowed roots proves $\gamma_\ell^{(r)}(\mathcal{F}) \le \pi_{\{[\ell]\}}(\mathcal{F})$. For the reverse inequality, fix $0 < d < \pi_{\{[\ell]\}}(\mathcal{F})$ and $0 < \varepsilon < d/2$, and put $p := r - \ell$. Choose an arbitrarily large integer m such that

$$\binom{m}{\ell} \exp\left(-\frac{\varepsilon^2(m-\ell)}{2p^2}\right) < \frac{1}{4},$$

and then choose $\eta > 0$ with $2\eta\binom{m}{\ell}/\varepsilon < 1/4$. There is an arbitrarily large $\mathcal{F}$-free $(d, \eta, \{[\ell]\})$-dense r-graph H on n vertices. Let $\mathcal{B}$ consist of those $A \in \binom{V(H)}{\ell}$ with $d_H(A) < (d-\varepsilon)\binom{n-\ell}{p}$. Testing all orderings of the members of $\mathcal{B}$ gives

$$\varepsilon \ell!\, |\mathcal{B}|\, p! \binom{n-\ell}{p} \le \eta n^r, \qquad \frac{|\mathcal{B}|}{\binom{n}{\ell}} \le \frac{2\eta}{\varepsilon}$$

for sufficiently large n.

Choose U uniformly from $\binom{V(H)}{m}$. The probability that U contains a member of $\mathcal{B}$ is at most $2\eta\binom{m}{\ell}/\varepsilon < 1/4$. For a fixed $A \notin \mathcal{B}$, conditional on $A \subseteq U$, the set $U \setminus A$ is a uniform $(m-\ell)$-subset of $V(H) \setminus A$, and

$$\mathbb{E}\left[\frac{d_{H[U]}(A)}{\binom{m-\ell}{p}} \,\middle|\, A \subseteq U\right] = \frac{d_H(A)}{\binom{n-\ell}{p}} \ge d - \varepsilon.$$

Expose $U \setminus A$ in a uniform random order. Replacing one sampled vertex changes $d_{H[U]}(A)$ by at most $\binom{m-\ell-1}{p-1}$. Coupling the remaining samples by a transposition gives the same bound on each increment of the exposure martingale. The bounded-differences inequality [29] therefore bounds the conditional probability that $d_{H[U]}(A) < (d-2\varepsilon)\binom{m-\ell}{p}$ by $\exp(-\varepsilon^2(m-\ell)/(2p^2))$. Summing over A, with $\mathbb{P}[A \subseteq U] = \binom{m}{\ell}/\binom{n}{\ell}$, shows that with positive probability $\delta_\ell(H[U]) \ge (d-2\varepsilon)\binom{m-\ell}{p}$. The induced subgraph $H[U]$ remains $\mathcal{F}$-free. Since m can be arbitrarily large, $\gamma_\ell^{(r)}(\mathcal{F}) \ge d - 2\varepsilon$. Letting $\varepsilon \to 0$ and then $d \uparrow \pi_{\{[\ell]\}}(\mathcal{F})$ proves the claim; the case $\pi_{\{[\ell]\}}(\mathcal{F}) = 0$ follows from the first inequality. □

Taking $\mathcal{S} = \{\varnothing\}$ gives the set of all ordered r-tuples of distinct vertices as the only nonempty test. Since each edge contributes $r!$ tuples, $\pi_{\{\varnothing\}} = \pi$. To identify $\gamma_1^{(r)}$ with π, we use the following pruning lemma.

Lemma 2.2. *Let $s \ge 2$, and let $\mathscr{Q}$ be a monotone class of s-graphs with ordinary Turán density β. For every $\varepsilon > 0$, there are arbitrarily large $H \in \mathscr{Q}$ such that*

$$|H| \ge (\beta - \varepsilon)\binom{v(H)}{s}, \qquad \text{and} \qquad \delta_1(H) \ge (\beta - 2\varepsilon)\binom{v(H)-1}{s-1}.$$

Proof. If $\beta = 0$, take edgeless s-graphs of arbitrary order. Suppose that $\beta > 0$ and set $\eta := \min\{\varepsilon/2, \beta/4\}$. Choose $c \in (0,1)$ so that $(1 - \beta + 2\eta)c^s < \eta$. For every sufficiently large N, choose an N-vertex $G \in \mathscr{Q}$ with $|G| \ge (\beta - \eta)\binom{N}{s}$.

While the current s-graph has order m, delete a vertex of degree smaller than $(\beta - 2\eta)\binom{m-1}{s-1}$, if one exists. If the process reaches order m, telescoping gives $|G_m| \ge (\beta - 2\eta)\binom{m}{s} + \eta\binom{N}{s}$. If $m \le cN$, then $\binom{m}{s} \le c^s\binom{N}{s}$, so the right-hand side exceeds $\binom{m}{s}$, which is impossible. Thus the process stops with $m > cN$.

At that point its minimum vertex degree is at least $(\beta - 2\eta)\binom{m-1}{s-1}$, while the displayed edge bound and $m \le N$ give edge density at least $\beta - \eta$. Since $\eta \le \varepsilon/2$, both claimed bounds follow, and $m \to \infty$ as $N \to \infty$ through sufficiently large integers. □

Averaging vertex degrees gives $\gamma_1^{(r)} \le \pi$, and Lemma 2.2 gives the reverse inequality. Together with Lemma 2.1, this proves the parameter identities in the ordinary and minimum-degree rows of Table 1.

Reiher, Rödl, and Schacht [43, Proposition 2.5] prove the complete-layer identity in Table 1. For $j \ge 2$, testing separate ordered relations on all j-faces is equivalent to testing one unordered j-graph, after changing the error parameter. For $j = 1$, their converse passes to an induced subhypergraph of positive proportional order. Both operations preserve avoidance of every member of an arbitrary family $\mathcal{F}$, so the same proof gives the stated identity for arbitrary forbidden families. Ordered and unordered counts differ by the factor $r!$, which only rescales the permitted error.

2.2 Finite approximation, intervals, and polynomial identities

We next record the approximation that allows compatible finite constructions to realize one limiting density.

Proposition 2.3. *For every $\mathcal{S} \subseteq 2^{[r]}$ and every forbidden family $\mathcal{F}$,*

$$\pi_{\mathcal{S}}(\mathcal{F}) = \inf\{\pi_{\mathcal{S}}(\mathcal{F}') \colon \mathcal{F}' \subseteq \mathcal{F},\ \mathcal{F}' \text{ finite}\}. \tag{2.1}$$

Consequently, for increasing families $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$,

$$\pi_{\mathcal{S}}\left(\bigcup_n \mathcal{F}_n\right) = \inf_n \pi_{\mathcal{S}}(\mathcal{F}_n). \tag{2.2}$$

Proof. The strong removal lemma of Rödl and Schacht [46, Theorem 6] says that, for every $\varepsilon > 0$, some finite subfamily $\mathcal{F}'$ has the following property. Every sufficiently large $\mathcal{F}'$-free r-graph can be made $\mathcal{F}$-free by deleting at most $\varepsilon\binom{n}{r}$ edges. Each deletion removes at most $r!$ tuples from any test, so the loss is at most εn^r simultaneously for every test.

Fix $d > \pi_{\mathcal{S}}(\mathcal{F})$ with $d \le 1$. There are $\eta_0 > 0$ and n_0 excluding $(d, \eta_0, \mathcal{S})$-dense $\mathcal{F}$-free hosts of order at least n_0. Apply removal with $\varepsilon = \eta_0/4$. If $\pi_{\mathcal{S}}(\mathcal{F}') > d$, arbitrarily large $\mathcal{F}'$-free $(d, \eta_0/4, \mathcal{S})$-dense hosts exist. After deletion they are $\mathcal{F}$-free and $(d, \eta_0/2, \mathcal{S})$-dense, a contradiction. Thus the infimum in Equation (2.1) is at most every $d > \pi_{\mathcal{S}}(\mathcal{F})$. The reverse inequality is monotonicity. Every finite subfamily of $\bigcup_n \mathcal{F}_n$ lies in one $\mathcal{F}_n$, giving Equation (2.2). □

We apply this proposition to decreasing classes of colorable hypergraphs, whose canonical forbidden families form an increasing sequence.

For continuous functions A, B on a nonempty compact set K, put $F(t) := \max_{x \in K}\{A(x) + tB(x)\}$. If x_s, x_t maximize at $s < t$, respectively, then

$$(t - s)B(x_s) \le F(t) - F(s) \le (t - s)B(x_t). \tag{2.3}$$

Indeed, evaluate x_s at t and x_t at s. Bounds on B at these maximizers therefore bound the change in the maximum; a bound on $|B|$ throughout K also gives continuity of F.

A continuous parameter can give an exact target while an auxiliary parameter is restricted to a dense set. Suppose D is dense in an interval J and $\Phi \colon [a, b] \times J \to \mathbb{R}$ is continuous. If $t_0 \in J$ satisfies

$$\Phi(a, t_0) < \alpha < \Phi(b, t_0), \tag{2.4}$$

then some $t \in D$ and $s \in (a, b)$ satisfy $\Phi(s, t) = \alpha$. Indeed, continuity preserves the strict inequalities for t sufficiently close to t_0. Choose such a $t \in D$, fix it, and apply the intermediate value theorem to $\Phi(\cdot, t)$.

The next two lemmas fill intervals by overlapping ranges and iteration. The first concerns successive prefixes.

Lemma 2.4. *Let $D = \{d_-, d_- + 1, \dots, d_+\}$, where $d_- < d_+$ are integers, and give $D^{\mathbb{N}}$ its product topology. Suppose $F \colon D^{\mathbb{N}} \to \mathbb{R}$ is continuous and nondecreasing in the coordinatewise order. If, for every finite word p and $d_- \le d < d_+$,*

$$F\big(p(d+1)d_-^{\infty}\big) \le F\big(pd\,d_+^{\infty}\big),$$

then $F(D^{\mathbb{N}}) = [F(d_-^{\infty}), F(d_+^{\infty})]$.

Proof. For a finite word p, put $I_p := [F(pd_-^\infty), F(pd_+^\infty)]$. Monotonicity and the assumed overlap give $I_p = \bigcup_{d\in D} I_{pd}$. Uniform continuity on the compact space $D^{\mathbb{N}}$ makes the diameters of these intervals tend uniformly to zero as the prefix length increases. Given $x \in I_\varnothing$, choose successive child intervals containing x. Their prefixes determine a sequence ω, and $F(\omega)$ belongs to every chosen interval by monotonicity. Hence $F(\omega) = x$. Monotonicity also gives the reverse inclusion. □

The second lemma propagates a realized interval under a continuous monotone operation.

Lemma 2.5. *Let $a < b < t$ and $A \subseteq \mathbb{R}$. Suppose $T\colon [a,t] \to [a,t]$ is continuous and nondecreasing, $T(A \cap [a,t]) \subseteq A$, $T(t) = t$, and $T(x) > x$ for $b \le x < t$. If $[a,b] \subseteq A$ and $T(a) \le b$, then $[a,t) \subseteq A$. The same conclusion holds if $[a,b) \subseteq A$ and $T(a) < b$.*

Proof. First suppose that $[a,b] \subseteq A$. Put $b_0 := b$ and $b_{j+1} := T(b_j)$. These endpoints increase to t, since continuity makes their limit a fixed point and $T(x) > x$ excludes every fixed point in $[b,t)$. If $[a,b_j] \subseteq A$, then $T([a,b_j]) = [T(a), b_{j+1}] \subseteq A$. This interval overlaps $[a,b_j]$, so induction gives $[a,b_j] \subseteq A$ for every j. Their union contains $[a,t)$.

Now suppose that $[a,b) \subseteq A$ and $T(a) < b$. Continuity and $T(b) > b$ allow $b' < b$ sufficiently close to b that $T(a) < b'$ and $T(x) > x$ for $b' \le x \le b$. Apply the first case with b' in place of b. □

For a homogeneous polynomial, the degrees of its terms in each group of variables determine the group weights at a maximum. Partition the variables into $V_1, \dots, V_m$ and write $F = \sum_{\boldsymbol{a}} F_{\boldsymbol{a}}$, where $F_{\boldsymbol{a}}$ has nonnegative coefficients and is homogeneous of degree a_i in group V_i, with $\sum_i a_i = r$. At a positive-value maximum over nonnegative weights of total weight $s > 0$, put $z_i = \sum_{v\in V_i} x_v$. Then

$$\frac{z_i}{s} = \sum_{\boldsymbol{a}} \frac{F_{\boldsymbol{a}}(\boldsymbol{x})}{F(\boldsymbol{x})} \frac{a_i}{r}. \tag{2.5}$$

Indeed, stationarity and Euler's identity give $\partial_v F = rF/s$ on the positive support. Multiplying by x_v and summing over V_i gives $(rF/s)z_i = \sum_{\boldsymbol{a}} a_i F_{\boldsymbol{a}}$; zero coordinates contribute zero.

The same degrees give the completion counts for a prescribed root. Let $V_1, \dots, V_m$ partition an n-vertex set, and let H consist of all r-sets with a_i vertices in each V_i, where $\sum_i a_i = r$. Put $z_i = |V_i|/n$. Its normalized edge-count polynomial is $M_{\boldsymbol{a}}(\boldsymbol{z}) := r! \prod_i z_i^{a_i} / \prod_i a_i!$. For an ℓ-set S with $b_i = |S \cap V_i| \le a_i$ and $1 \le \ell < r$, counting the remaining vertices in each part gives

$$\frac{d_H(S)}{\binom{n-\ell}{r-\ell}} = \frac{(r-\ell)!}{\prod_i (a_i - b_i)!} \prod_i z_i^{a_i - b_i} + O_{r,m}\left(\frac{1}{n}\right) = \frac{(r-\ell)!}{r!} \frac{\partial^\ell M_{\boldsymbol{a}}}{\partial z_1^{b_1} \cdots \partial z_m^{b_m}}(\boldsymbol{z}) + O_{r,m}\left(\frac{1}{n}\right). \tag{2.6}$$

If some $b_i > a_i$, both the completion count and the derivative are zero. These formulas add over disjoint allowed profiles.

We also use the following polynomial bound. If $f(x) = \sum_{j=0}^k a_j x^j$ on an interval $[u,v]$ with $u < v$, put

$$q_i := \sum_{j=i}^k a_j \binom{j}{i} u^{j-i} (v-u)^i, \qquad b_\ell := \sum_{i=0}^\ell q_i \frac{\binom{\ell}{i}}{\binom{k}{i}}.$$

Then $f(u + (v-u)z) = \sum_{\ell=0}^k b_\ell \binom{k}{\ell} z^\ell (1-z)^{k-\ell}$. The Bernstein basis functions are nonnegative and sum to one for $0 \le z \le 1$, so

$$\min_{0\le\ell\le k} b_\ell \le f(x) \le \max_{0\le\ell\le k} b_\ell \qquad \text{for every } x \in [u,v]. \tag{2.7}$$

All coefficient calculations are rational when the interval endpoints and the coefficients of f are rational. Applying the expansion in each variable gives the corresponding bound on a rational box.

Part I

Ordinary and minimum-degree Turán densities

3 Ordinary Turán densities and a shared construction

The ordinary construction begins with the following initial interval of Li, Liu, and Mubayi.

Theorem 3.1. *There exists a constant $\varepsilon > 0$ such that, whenever $r \geq 10$ and either $\mathcal{S} = \{\varnothing\}$ or $\mathcal{S} = \{\{1\}\}$, one has $[\alpha_r, (1+\varepsilon)\alpha_r] \subseteq \Pi^{(r)}_{\mathcal{S},\infty}$.*

We use a tree-polynomial construction to obtain a common initial interval for ordinary and intermediate uniform Turán densities. A fixed list of cubic components supplies the variable part of a 10-uniform construction. Finite restrictions give ordinary densities; ordered versions will give the intermediate uniform densities in Section 11.2.

Liu and Pikhurko [27, Theorem 1.1] proved the following terminal-interval theorem.

Theorem 3.2. *For every integer $r \geq 3$, there is a constant $c_r < 1$ such that $[c_r, 1] \subseteq \Pi^{(r)}_{\infty}$.*

Their result answers the interval question of Frankl, Peng, Rödl, and Talbot [9, Question 6.1]. The proof of the large-uniformity classification uses only finite Lagrangians and their limits. Brown and Simonovits [4, Theorem 4] proved that finite normalized Lagrangians belong to $\Pi^{(r)}_{\infty}$ and are dense in it. Pikhurko [40, Proposition 1] proved that $\Pi^{(r)}_{\infty}$ is closed. Consequently every limit of finite normalized Lagrangians is an ordinary Turán density.

Lagrangians encode the limiting edge densities of blowups, an approach surveyed by Peng [37]. For a finite r-graph H, its edge polynomial and normalized Lagrangian are

$$\Lambda_H(\boldsymbol{x}) \coloneqq \sum_{e \in H} \prod_{v \in e} x_v, \qquad \lambda(H) \coloneqq r! \max_{\boldsymbol{x} \in S(V(H))} \Lambda_H(\boldsymbol{x}),$$

where $S(V)$ is the probability simplex on V. A weighting of total weight p belongs to $pS(V)$. The *relative Lagrangian* $\lambda(H)/\alpha_r$ is the maximum of Λ_H at total weight r. All vector and operator norms below are Euclidean.

For $r \geq 3$ and $b, c \in [0,1]$, put

$$\Theta_r(b,c) \coloneqq \max_{0 \leq x \leq 1} \{bx^r + rc(1-x)x^{r-1}\}. \tag{3.1}$$

Lemma 3.3. *For finite r- and $(r-1)$-graphs G, F, there is a finite r-graph with normalized Lagrangian $\Theta_r(\lambda(G), \lambda(F))$. The map Θ_r is continuous and nondecreasing in both arguments, and*

$$\Theta_r(b,c) = \begin{cases} c\left(\dfrac{(r-1)c}{rc-b}\right)^{r-1}, & 0 \leq b < c, \\ b, & b \geq c. \end{cases}$$

If $b < c$, then $b < \Theta_r(b,c) < c$. If b and c are limits from below of finite normalized Lagrangians in their respective uniformities, the same holds for $\Theta_r(b,c)$. In particular, $\Theta_r(b,c) \in \Pi^{(r)}_{\infty}$ whenever $b \in \Pi^{(r)}_{\infty}$ and $c \in \Pi^{(r-1)}_{\infty}$.

Proof. Take the vertex set $V(G) \times V(F)$ and a new vertex z, adding isolated vertices to G, F if needed. An r-set in the product is an edge when its first-coordinate projection is an edge of G. An r-set containing z is an edge when the second-coordinate projection of its other vertices is an edge of F.

If the product set has total weight x, its row and column marginals bound the normalized contributions by $\lambda(G)x^r$ and $r\lambda(F)(1-x)x^{r-1}$. Product weights xp_vq_w, with $\boldsymbol{p}, \boldsymbol{q}$ optimal for G, F, attain both bounds simultaneously. This proves the Lagrangian formula. The case $c = 1$ was considered by Grosu [14, Theorem 44].

Differentiating the objective gives $rx^{r-2}((r-1)c+(b-rc)x)$, which proves the explicit formula and the strict inequalities. Continuity follows by maximization over $[0,1]$, and monotonicity follows from the nonnegative coefficients of b, c. Applying the construction to approximating finite graphs proves the assertion about limits from below. Arbitrary finite approximations and closedness of $\Pi_\infty^{(r)}$ give the final assertion. □

Figure 1 illustrates the two edge types and the interval propagation of Lemma 2.5 for the map $\Theta_r(\cdot, c)$.

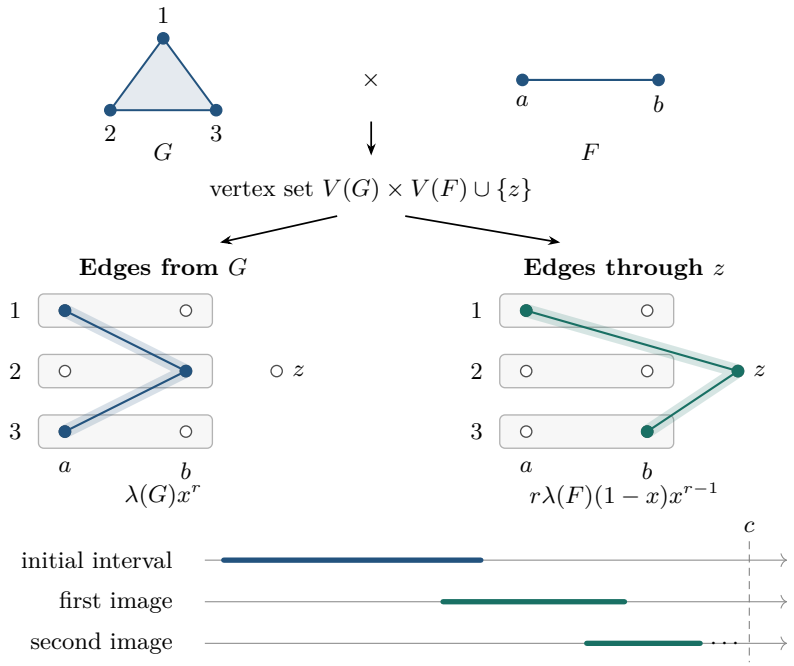


Figure 1: The link construction for $r = 3$. The top row shows the inputs, with $\{1,2,3\} \in G$ and $\{a,b\} \in F$. Each band joins the vertices of one representative edge; the grid has total weight x. Below, successive images under $\Theta_r(\cdot, c)$ overlap and approach c when the hypotheses of Lemma 2.5 hold. The intervals are schematic.

Write $p_r := \alpha_r/\alpha_{r-1} = (1-1/r)^{r-1}$. Then $\Theta_r(0,c) = p_r c$, and iteration of $\Theta_r(\cdot, c)$ from zero increases to c, since no smaller value is a fixed point. The values p_r decrease to e^{-1}, since the derivative of $(t-1)\log(1-1/t)$ is $\log(1-1/t) + 1/t < 0$. In particular, $e^{-1} < p_r \le 4/9$. Closedness therefore gives

$$p_r \Pi_\infty^{(r-1)} \cup \Pi_\infty^{(r-1)} \subseteq \Pi_\infty^{(r)}.$$

3.1 A common 10-uniform construction

Fix integers $Q \ge 10^9$ and $n \ge 2Q$. The auxiliary vertices are α_i, β_i for $i \in \mathbb{Z}/n\mathbb{Z}$, together with d_* and d. We use each auxiliary vertex label also for its nonnegative weight. Put $\mathsf{A} := \sum_i \alpha_i$, $\mathsf{B} := \sum_i \beta_i$, and $S := \mathsf{A} + \mathsf{B} + d_* + d$. Define the auxiliary polynomials

$$U_1 := \sum_i \alpha_i(\beta_i + d_*)(\mathsf{B} - \beta_i + d), \qquad U_2 := \sum_i \beta_i(\alpha_{i+1} + d_*)(\mathsf{A} - \alpha_{i+1} + d).$$

These are simple 3-graph polynomials. The 2 parenthesized factors in each summand of U_1 have sum $S - \mathsf{A}$, so AM–GM gives $U_1 \le \mathsf{A}(S-\mathsf{A})^2/4 \le S^3/27$. Similarly, $U_2 \le \mathsf{B}(S-\mathsf{B})^2/4 \le S^3/27$. Both polynomials attain this bound at

$$\alpha_i = \beta_i = \frac{S}{3n}, \quad d_* = \frac{S}{3}\left(1 - \frac{1}{n}\right), \quad d = \frac{S}{3n}.$$

The four edge types in each auxiliary polynomial are shown in Figure 2.

For $1 \le j \le Q$, the blocks of 4 auxiliary vertices with weights $\alpha_{2j-1}, \beta_{2j-1}, \alpha_{2j}, \beta_{2j}$ are pairwise disjoint. Write $z_j := \alpha_{2j-1}\beta_{2j-1}\alpha_{2j}\beta_{2j}$, where subscripts are modulo n.

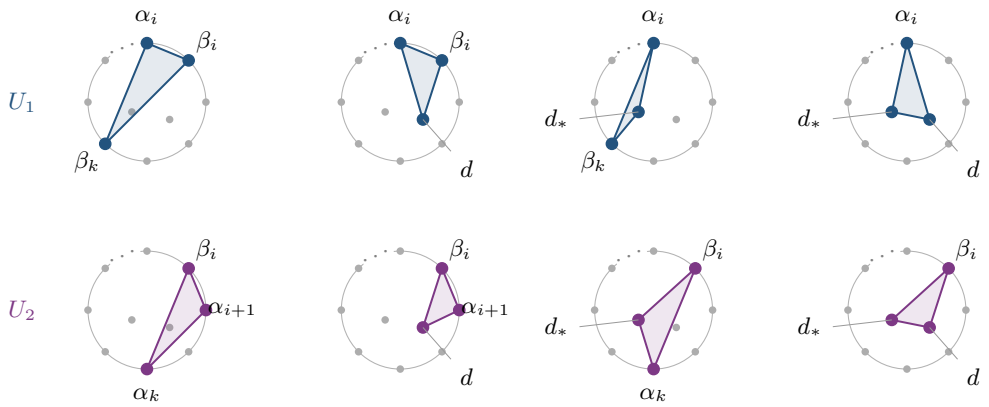


Figure 2: The four edge types in each auxiliary polynomial. In the first row $k \neq i$; in the second row $k \neq i+1$. Colored vertices form the displayed edge. The cyclic order and the positions of d_*, d are the same in all panels.

Let $D_1, \dots, D_Q$ be simple 4-graph polynomials on a common vertex set disjoint from the auxiliary vertices, each of maximum one at total weight 4. Denote the total weight on this common set by Y. For $\boldsymbol{\theta} \in \mathbb{R}^Q_{\geq 0}$, set

$$B(\boldsymbol{\theta}) := \sup \left\{ \sum_{j \in [Q]} \theta_j D_j(\boldsymbol{y}) \colon \boldsymbol{y} \geq 0,\ \sum y = 4 \right\}.$$

The common vertex set may be countable. In that case each edge polynomial is defined by its absolutely convergent edge sum on summable weightings, and the displayed supremum may equivalently be taken over finite-support weightings. Necessarily $1 \leq B(\mathbf{1}) \leq Q$.

Introduce 5 further vertices $t_1, \dots, t_5$ and 2 vertices u_1, u_2. Define the simple 10-graph polynomial, with P_{base} and P_{add} denoting the first and second terms on the right, respectively, by

$$P := P_{\text{base}} + P_{\text{add}} = \left(\prod_{a=1}^{5} t_a \right) Y(u_1 U_1 + u_2 U_2) + u_1 u_2 \sum_{j=1}^{Q} z_j D_j(\boldsymbol{y}). \tag{3.2}$$

The degrees are $5+1+1+3$ and $2+4+4$. The presence of the t-vertices separates the 2 edge families; u_b distinguishes the 2 base edge types, and the auxiliary block associated with z_j distinguishes the list entries in the additional part. Thus repeated components in the list do not create multiple 10-edges. The two edge families are shown in Figure 3.

Proposition 3.4. *Let v_n be the supremum of (3.2) at total weight 10. The following statements hold.*

(i) *The supremum satisfies*

$$0 \leq v_n - 1 - \frac{B(\mathbf{1})}{4^5 n^4} \leq \frac{9Q}{n^5}.$$

(ii) *In the finite case, every maximizing weighting satisfies*

$$3 \leq S \leq 4, \quad 1 \leq Y \leq 4, \quad \left| \frac{(3n)^4 z_j}{S^4} - 1 \right| \leq \frac{130\sqrt{Q}}{n}.$$

For countable components, the finite-dimensional outer objective obtained by taking the supremum over the vertex weights $\boldsymbol{y}$ has a maximizer satisfying the same bounds.

These conclusions are uniform over all component lists of length Q.

Proof. We first assume the components are finite. The test weighting below gives a value greater than one; comparing the loss in the base term with the possible gain in the additional term will force α_i, β_i close to $S/(3n)$. We then identify the leading correction involving $B(\mathbf{1})$ and pass to countable components.

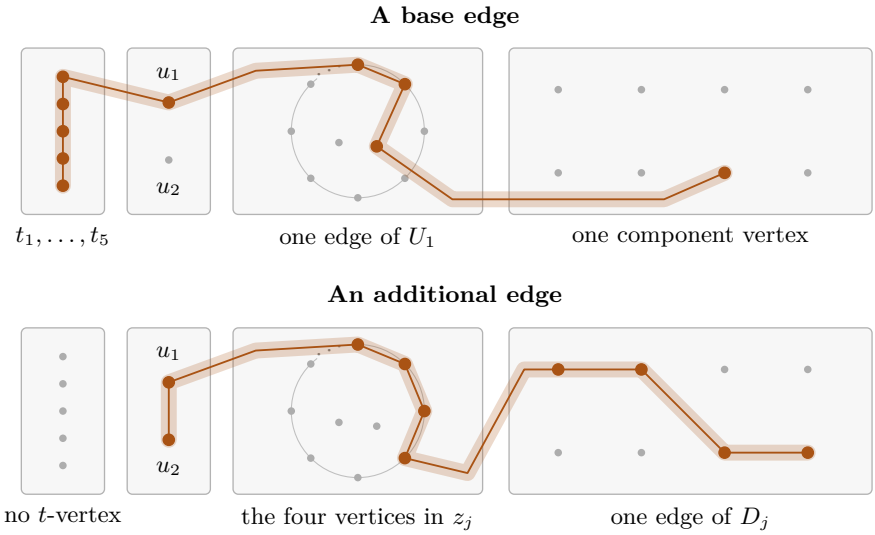


Figure 3: The two edge families in Equation (3.2), with the same vertex groups in both panels. Each orange band selects one 10-edge: $5+1+3+1$ vertices in the base family and $2+4+4$ in the additional family. Base edges also use u_2 with an edge of U_2. The rightmost box is the common vertex set of $D_1, \dots, D_Q$.

A test weighting has $t_a = Y = 1$, $u_1 = u_2 = 1/2$, $\alpha_i = \beta_i = d = 1/n$, $d_* = 1 - 1/n$, and vertex weights $\boldsymbol{y}$ of total one on the common component vertex set for which $\sum_j D_j(\boldsymbol{y}) = B(\mathbf{1})/4^4$. It gives

$$v_n \geq 1 + \frac{B(\mathbf{1})}{4^5 n^4} > 1. \tag{3.3}$$

Put $\rho_A := (\mathsf{A} + d - d_*)/2$, $\rho_B := (\mathsf{B} + d - d_*)/2$, and $\Delta := 2S^3/27 - U_1 - U_2$. Completing the square in each summand of U_1 gives

$$U_1 = \frac{\mathsf{A}(S-\mathsf{A})^2}{4} - \sum_i \alpha_i(\beta_i - \rho_B)^2$$

and the analogous formula for U_2. For $0 \leq x \leq S$,

$$\frac{S^3}{27} - \frac{x(S-x)^2}{4} = \frac{1}{4}\left(x - \frac{S}{3}\right)^2\left(\frac{4S}{3} - x\right) \geq \frac{S}{12}\left(x - \frac{S}{3}\right)^2.$$

Applying this with $x = \mathsf{A}$ and $x = \mathsf{B}$, and using the definition of Δ, gives

$$\sum_i \left[\alpha_i(\beta_i - \rho_B)^2 + \beta_i(\alpha_{i+1} - \rho_A)^2\right] \leq \Delta, \qquad \left(\mathsf{A} - \frac{S}{3}\right)^2 + \left(\mathsf{B} - \frac{S}{3}\right)^2 \leq \frac{12\Delta}{S} \quad \text{for } S > 0. \tag{3.4}$$

To bound the products z_j in terms of the auxiliary deficit Δ, put $h := \max\{0, \rho_A, \rho_B\}$.

Claim 3.5. *The products z_j satisfy*

$$\sum_j z_j \leq 80S\Delta + 26QS^4n^{-4}.$$

If $\Delta \geq h^3/28$, the term $26QS^4n^{-4}$ may be omitted. If $\Delta < h^3/28$, then $S > 0$ and $S/(27n) < \alpha_i, \beta_i \leq 9S/(4n)$ for every i.

Proof of the claim. In the jth auxiliary block, write its 4 weights as $x_0 := \alpha_{2j-1}$, $x_1 := \beta_{2j-1}$, $x_2 := \alpha_{2j}$, and $x_3 := \beta_{2j}$. For $1 \leq k \leq 3$, the pair (x_{k-1}, x_k) then corresponds to a squared-difference term on the left side of the first inequality in (3.4). If $\max\{x_1, x_2, x_3\} \leq 5h/4$, z_j is at most $125x_0h^3/64$. Otherwise choose a largest factor x_k among x_1, x_2, x_3, so $x_k > 5h/4$. Since $x_k - \rho \geq x_k - h > x_k/5$ for $\rho \in \{\rho_A, \rho_B\}$, the

squared-difference term corresponding to (x_{k-1}, x_k) is at least $x_{k-1}x_k^2/25$, while z_j is at most $Sx_{k-1}x_k^2$. Distinct auxiliary blocks use distinct such terms and $\sum x_0 \le S$, so that inequality gives

$$\sum_j z_j \le 25S\Delta + \frac{125Sh^3}{64}.$$

This is at most $80S\Delta$ when $\Delta \ge h^3/28$.

In the case $\Delta < h^3/28$, rescale to $S = 1$. We will show that $h/3 < \alpha_i, \beta_i < 3h/2$ for every i; summing then bounds h in terms of n. We have $0 < h \le 1/2$. By (3.4), $(\rho_A - \rho_B)^2 \le 6\Delta < h^2/9$ and $(\mathsf{A} + \mathsf{B} - 2/3)^2 \le 24\Delta < 1/9$. It follows that $2h/3 < \rho_A, \rho_B \le h$ and $\mathsf{A} + \mathsf{B} > 1/3$. The cyclic sequence obtained by alternating α_i/ρ_A and β_i/ρ_B has sum greater than $2/3$ and satisfies

$$\sum_i \left[\frac{\alpha_i(\beta_i - \rho_B)^2}{\rho_A\rho_B^2} + \frac{\beta_i(\alpha_{i+1} - \rho_A)^2}{\rho_B\rho_A^2}\right] \le \frac{27\Delta}{8h^3} < \frac{27}{224} < \frac{1}{8}.$$

A nonnegative cyclic sequence (x_i) with sum at least $1/2$ and $\sum_i x_i(1 - x_{i+1})^2 < 1/8$ has all coordinates in $(1/2, 3/2)$. Indeed, if every coordinate were below $1/2$, then $\sum_i x_i(1 - x_{i+1})^2 > \frac{1}{4}\sum_i x_i \ge 1/8$. If some coordinates were below $1/2$ and others were at least $1/2$, a consecutive pair with $x_i \ge 1/2 > x_{i+1}$ would already contribute more than $1/8$. Once every coordinate is at least $1/2$, each $|x_{i+1} - 1|^2 < 1/4$ follows from its summand. Hence $h/3 < \alpha_i, \beta_i < 3h/2$. Summing over the $2n$ auxiliary coordinates α_i, β_i gives $1/(9n) < h \le 3/(2n)$, proving the coordinate bounds. Rescaling to total weight S gives $z_j < 26S^4n^{-4}$. Summing over the Q auxiliary blocks proves the claim. □

We now apply these estimates at a maximizing weighting. The base term satisfies $P_{\text{base}} \le 1$, and disjointness of the auxiliary blocks gives

$$P_{\text{add}} \le u_1u_2\left(\frac{S}{4}\right)^4\left(\frac{Y}{4}\right)^4 \le u_1u_2\left(\frac{S+Y}{8}\right)^8 \le 1.$$

Both summands and the total weights of all vertex groups are therefore positive at a maximizer.

Fix a maximizing weighting for the finite construction and put $\vartheta := P_{\text{add}}/v_n$. Both summands are positive, so $0 < \vartheta < 1$. Apply Equation (2.5) with degree and total weight both 10. The base and additional terms have respective degrees $(1, 0)$ in each t_a, $(3, 4)$ in the auxiliary group, $(1, 4)$ in the component group, and $(1, 2)$ in the group $\{u_1, u_2\}$. Hence, for $a \in [5]$,

$$t_a = 1 - \vartheta, \qquad S = 3 + \vartheta, \qquad Y = 1 + 3\vartheta, \qquad u_1 + u_2 = 1 + \vartheta.$$

The bound $P_{\text{add}} \le u_1u_2((S + Y)/8)^8$, together with these identities and $v_n > 1$, gives $\vartheta < ((1 + \vartheta)/2)^{10}$. The difference between the right side and ϑ is strictly convex, negative at 10^{-3}, and zero at 1. Hence $0 < \vartheta < 10^{-3}$.

Put $\Gamma_0 := (1 - \vartheta)^5Y\min(u_1, u_2)$. Since $U_i \le S^3/27$ and their deficits sum to Δ,

$$u_1U_1 + u_2U_2 \le (u_1 + u_2)\frac{S^3}{27} - \min(u_1, u_2)\Delta.$$

AM–GM gives $(1 - \vartheta)^5Y(u_1 + u_2)(S/3)^3 \le 1$, since $5(1 - \vartheta) + Y + u_1 + u_2 + S = 10$. Multiplying the bound on $u_1U_1 + u_2U_2$ by $(1 - \vartheta)^5Y$ therefore gives $P_{\text{base}} \le 1 - \Gamma_0\Delta$.

To compare the additional terms with this loss, note that $\frac{u_1u_2SY^4}{\Gamma_0} < 31/10$ and $0 < \Gamma_0 \le 1/2$. Since $D_j(4\boldsymbol{y}/Y) \le 1$, homogeneity and Claim 3.5 give $P_{\text{add}} \le u_1u_2(Y/4)^4(80S\Delta + 26QS^4n^{-4})$. The preceding ratio bound and $S < 3.001$ then yield

$$P_{\text{add}} \le \Gamma_0\left(\frac{31\Delta}{32} + 9Qn^{-4}\right). \tag{3.5}$$

Thus

$$1 < v_n \le 1 + \Gamma_0\left(-\frac{\Delta}{32} + 9Qn^{-4}\right).$$

Consequently $\Delta < 288Qn^{-4}$ and $\vartheta < 288Qn^{-4}$. If $\Delta \geq h^3/28$, Claim 3.5 would eliminate the last term in (3.5), yielding $v_n \leq 1$, a contradiction. Thus the small-deficit case holds, so the bounds in Claim 3.5 apply to every α_i and β_i.

Using the coordinate lower bound $S/(27n)$ in the first inequality of (3.4) and cyclically reindexing gives $\sum_i[(\alpha_i-\rho_A)^2+(\beta_i-\rho_B)^2] \leq 27n\Delta/S$. Centering at the class means $\mathsf{A}/n, \mathsf{B}/n$ can only decrease this sum. Expanding around these means and applying the second inequality in (3.4), with $S \geq 3$ and $n \geq 2$, gives

$$\begin{aligned}\sum_i \left[\left(\alpha_i - \frac{S}{3n}\right)^2 + \left(\beta_i - \frac{S}{3n}\right)^2\right] &= \sum_i \left[\left(\alpha_i - \frac{\mathsf{A}}{n}\right)^2 + \left(\beta_i - \frac{\mathsf{B}}{n}\right)^2\right] \\ &\quad + \frac{1}{n}\left[\left(\mathsf{A} - \frac{S}{3}\right)^2 + \left(\mathsf{B} - \frac{S}{3}\right)^2\right] \leq \frac{27n\Delta}{S} + \frac{12\Delta}{Sn} \leq 10n\Delta.\end{aligned} \tag{3.6}$$

The preceding estimates localize the weights α_i, β_i but bound the additional part using only Q. To recover $B(\mathbf{1})$ in its leading term, keep the vertex weights $\boldsymbol{y}$ of the current maximizer fixed and compare each z_j with $(S/(3n))^4$. The values $D_j(4\boldsymbol{y}/Y)$ lie in $[0,1]$ and have sum at most $B(\mathbf{1})$. Applying the derivative bound for a product of 4 factors and Cauchy–Schwarz to the coordinate differences on the $4Q$ distinct auxiliary vertices gives

$$\sum_j z_j D_j\left(\frac{4\boldsymbol{y}}{Y}\right) \leq B(\mathbf{1})\left(\frac{S}{3n}\right)^4 + \frac{729\sqrt{10}}{32} S^3 \sqrt{Q}\, n^{-5/2}\sqrt{\Delta}.$$

Indeed, along the line segment joining the current vector of coordinates α_i, β_i to the vector with all coordinates $S/(3n)$, every coordinate is at most $9S/(4n)$. Thus each partial derivative of z_j on this segment is at most $(9S/(4n))^3$. The sum of absolute coordinate differences on the $4Q$ selected auxiliary vertices is at most $\sqrt{4Q}\sqrt{10n\Delta}$.

Using $D_j(\boldsymbol{y}) = (Y/4)^4 D_j(4\boldsymbol{y}/Y)$, substitute this estimate into P_{add} and retain the bound $P_{\mathrm{base}} \leq 1 - \Gamma_0\Delta$. The bound $u_1u_2SY^4/\Gamma_0 < 31/10$ then yields

$$v_n \leq 1 + \frac{B(\mathbf{1})}{4^5 n^4}(1+17\vartheta) + \Gamma_0(-\Delta + 8\sqrt{Q}\, n^{-5/2}\sqrt{\Delta}). \tag{3.7}$$

Here we bounded the factor $(1+\vartheta/3)^4(1+\vartheta)^2(1+3\vartheta)^4$ by $1+17\vartheta$. Indeed, $\log(1+x) \leq x$ bounds this product by $\exp(46\vartheta/3)$, and convexity together with $\exp(46/3000) < 1.017$ bounds the exponential by $1+17\vartheta$ on $[0, 10^{-3}]$. The coefficient 8 follows from $22599\sqrt{10}S^2/81920 < 8$ for $S < 3.001$.

By (3.3) and $P_{\mathrm{base}} \leq 1$, we have $P_{\mathrm{add}} \geq B(\mathbf{1})/(4^5n^4)$. The upper coordinate bounds imply

$$P_{\mathrm{add}} \leq u_1u_2\left(\frac{Y}{4}\right)^4\left(\frac{9S}{4n}\right)^4 B(\mathbf{1}).$$

Cancel $B(\mathbf{1}) > 0$ to get $u_1u_2 > 1/10000$, hence $\min(u_1,u_2) > 1/11000$ and $\Gamma_0 > 1/12000$. In (3.7), the extra term $17B(\mathbf{1})\vartheta/(4^5n^4)$ is at most $5Q^2n^{-8}$. Comparison with (3.3) therefore gives $\Delta \leq 8\sqrt{Q}\, n^{-5/2}\sqrt{\Delta} + 60000Q^2n^{-8}$, so

$$\Delta \leq 64Qn^{-5} + 120000Q^2n^{-8} \leq 65Qn^{-5}.$$

Here the last inequality follows from $n \geq 2Q$, $Q \geq 10^9$. Equation (3.6) now gives

$$\left|\frac{3n\alpha_i}{S} - 1\right|, \left|\frac{3n\beta_i}{S} - 1\right| \leq \sqrt{10n^3\Delta} < \frac{26\sqrt{Q}}{n} < \frac{1}{2000}.$$

Multiplying the 4 factors, each close to one, gives $|z_j/(S/(3n))^4 - 1| \leq 130\sqrt{Q}/n$, completing the proof of (ii) in the finite case. For the upper bound in (i), use

$$-\Delta + 8\sqrt{Q}\, n^{-5/2}\sqrt{\Delta} = -\left(\sqrt{\Delta} - 4\sqrt{Q}\, n^{-5/2}\right)^2 + 16Qn^{-5} \leq 16Qn^{-5}.$$

Substituting into (3.7) and using $\Gamma_0 \leq 1/2$ gives

$$v_n \leq 1 + \frac{B(\mathbf{1})}{4^5n^4} + 8Qn^{-5} + 5Q^2n^{-8} \leq 1 + \frac{B(\mathbf{1})}{4^5n^4} + 9Qn^{-5}.$$

Together with (3.3), this proves (i) in the finite case.

We now allow the common color set of the fixed component list to be countable. The claimed attainment concerns the objective obtained after optimizing over the vertex weights $\boldsymbol{y}$. We prove it by uniformly approximating the resulting outer objective with finite restrictions.

Restrict the fixed components to an increasing sequence of finite color sets that exhausts their common color set, starting with a set containing an edge of every component, and write $B_m(\boldsymbol{\theta})$ for the maximum of their weighted sum at total weight 4, with coefficients $\boldsymbol{\theta}$. Every restricted component still has maximum one at that total weight. Finite-support approximation gives $B_m(\boldsymbol{\theta}) \uparrow B(\boldsymbol{\theta})$ for nonnegative coefficients. For every m, the function B_m is 1-Lipschitz in the ℓ^1 norm, since each component takes values in $[0,1]$; the same holds for B. Pointwise convergence on a finite net therefore gives uniform convergence on every compact coefficient box.

For fixed outer coordinates the optimized finite objective is $P_{\mathrm{base}} + u_1 u_2 (Y/4)^4 B_m(z_1, \dots, z_Q)$. The outer variables $t_a, u_b, \alpha_i, \beta_i, d_*, d, Y$ are nonnegative and have total weight 10, so they range over a compact finite-dimensional simplex. The products z_j lie in a fixed compact coefficient box. Thus these objectives converge uniformly to the outer objective obtained by replacing B_m with B. Their maxima converge, and a subsequence of maximizers in this simplex converges to a maximizer of the limiting objective. The finite localization bounds pass to the limit along this subsequence. This proves both assertions for countable components.

□

3.2 A cubic optimization on trees

We construct a fixed list of 3-graph polynomials on colors indexed by a rooted tree. Each polynomial has maximum one over nonnegative weightings of total weight 3, and the maximum of their weighted sum is determined by a scalar recursion on the tree. Varying the sequence of integer leaf counts gives an interval of maximum values. Adjoining a fixed new vertex gives 4-graph components whose maximum at total weight 4 is also one; these supply the variable components of the construction from Section 3.1.

Fix the alphabet $\mathcal{A} := \{10, 11, \dots, 48\}$ and $s \in [999/100, 1001/100]$. A finite prefix $d_0, \dots, d_{h-1}$ gives a path $v_0 \dots v_{h-1}$, with d_i leaves at v_i. Retain 48 labeled leaf slots at each path vertex, treating unused slots as isolated indices. Add 3 distinct isolated reserve indices. The infinite version uses the ray $v_0 v_1 \dots$ and a sequence $\boldsymbol{\omega} = (d_0, d_1, \dots) \in \mathcal{A}^{\mathbb{N}}$. Its maximum degree is at most 50. Denote the tree component by T, rooted at v_0, and index its root by 0. All color-weight sums below also include the isolated indices.

Color path edges alternately with labels $1, 2$ and assign label $k+2$ to the edge joining a path vertex to its kth active leaf, for $1 \le k \le 48$. Thus the edges carrying any one of the 50 labels form a matching. Let f_ℓ exchange endpoints of the corresponding matching and fix other indices, for $1 \le \ell \le 50$, and put $f_0 := \mathrm{id}$. The reserve indices are fixed by every f_ℓ.

At index i use colors a_i, b_i, c_i and denote their nonnegative weights by the same symbols. Put $\boldsymbol{x}_i := (a_i, b_i, c_i)$. We use the identity and matching maps f_ℓ to compare the 3 types at individual indices and across tree edges. Put $A := \sum_i a_i$, $B := \sum_i b_i$, and $C := \sum_i c_i$. For $\ell = 0, \dots, 50$ define

$$H_{\ell,a} := \sum_i a_i (b_{f_\ell(i)} + C - c_{f_\ell(i)})(B - b_{f_\ell(i)} + c_{f_\ell(i)}), \tag{3.8}$$

and define $H_{\ell,b}, H_{\ell,c}$ by cyclic permutation of the 3 types. This gives 153 polynomials. Figure 4 shows their edge types for $H_{\ell,a}$.

Put $M_T := \sum_i a_i b_i c_i$, and define

$$T_3 := \sum_{ij \in E(T)} (a_i b_i c_j + a_i c_i b_j + b_i c_i a_j + a_j b_j c_i + a_j c_j b_i + b_j c_j a_i),$$
$$R_0 := a_0 b_0 (C - c_0) + a_0 c_0 (B - b_0) + b_0 c_0 (A - a_0).$$

The remaining components are $ABC - M_T, T_3, R_0$, which are 3-partite simple 3-graph polynomials. Their edges are shown in Figure 5.

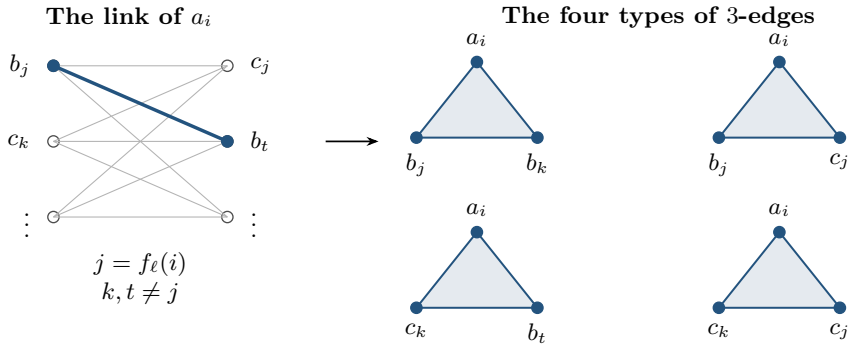


Figure 4: The control polynomial $H_{\ell,a}$. For fixed i and $j = f_\ell(i)$, the link at a_i is complete bipartite between b_j and the other c-colors on one side, and c_j and the other b-colors on the other. The four choices give the four triangle types on the right. The controls $H_{\ell,b}$ and $H_{\ell,c}$ cyclically permute the color types.

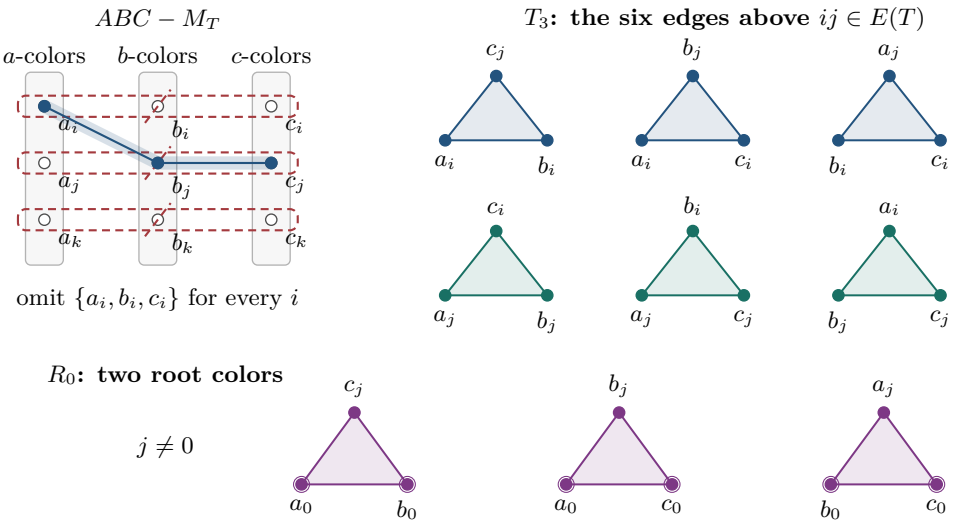


Figure 5: Cubic components associated with T. The left panel omits equal-index triples from the complete tripartite 3-graph and highlights one retained edge. The six triangles above show all contributions of $ij \in E(T)$ to T_3. The bottom row shows the three types in R_0, with root colors circled.

For a weighting of finite total weight with $a_i = b_i = c_i = x_i$, we have

$$sM_T - \frac{1}{3}T_3 = s\sum_i x_i^3 - \sum_{ij \in E(T)} (x_i^2 x_j + x_i x_j^2).$$

At a leaf with weight y and current cubic coefficient $a > 0$, attached to a parent of weight x, the terms involving y are $ay^3 - x^2y - xy^2$. Minimizing this expression gives the change in the parent's coefficient.

For $a > 0$ define

$$c(a) := \frac{1}{\sqrt{1+3a} - 1}, \qquad \psi(a) := \frac{2c(a) + c(a)^2}{3}.$$

Then $3ac(a)^2 - 2c(a) - 1 = 0$ and $2ac(a) - 1 > 0$. Hence, for $x, y \geq 0$,

$$ay^3 - x^2y - xy^2 + \psi(a)x^3 = (y - c(a)x)^2\big(ay + (2ac(a) - 1)x\big).$$

In particular the minimum over $y \geq 0$ of the first 3 terms is $-\psi(a)x^3$, attained at $y = c(a)x$. Differentiating the relation between a and $c(a)$ gives $\psi'(a) = -c(a)^3 < 0$.

For a finite tree T, successive minimization from the leaves to the root gives the effective coefficients

$$\xi_v := s - \sum_{w \text{ child of } v} \psi(\xi_w). \tag{3.9}$$

A leaf has coefficient s. On the path this is $\xi_i = s - d_i\psi(s) - \psi(\xi_{i+1})$, with the final term $\psi(\xi_{i+1})$ omitted at the terminal path vertex. The root coefficient is denoted $\xi_T(s)$. Figure 6 illustrates the tree and one elimination step.

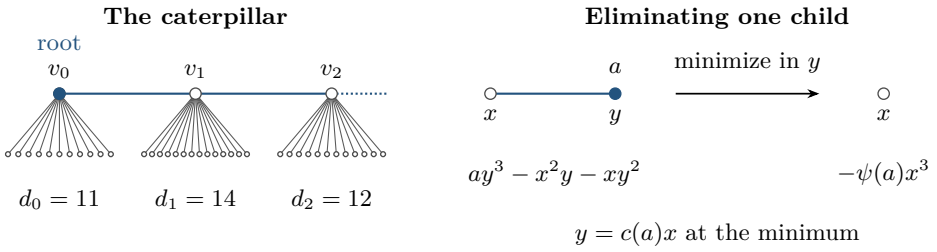


Figure 6: A caterpillar with initial leaf counts $11, 14, 12$. Eliminating a child of coefficient a subtracts $\psi(a)$ from its parent's cubic coefficient, as in Equation (3.9). Isolated leaf slots and reserve indices are omitted.

To pass to the infinite ray, rewrite the path recursion in terms of $\psi(\xi_i)$. For $d \in \mathcal{A}$ and $0 \le t \le 1$, define $\varphi_{d,s}(t) \coloneqq \psi(s - d\psi(s) - t)$. For an infinite sequence $\boldsymbol{\omega} = (d_0, d_1, \ldots) \in \mathcal{A}^{\mathbb{N}}$, write

$$t_{\boldsymbol{\omega}}(s) \coloneqq \lim_{h\to\infty} \varphi_{d_0,s} \circ \cdots \circ \varphi_{d_{h-1},s}(0).$$

For the infinite tree T determined by $\boldsymbol{\omega}$, define $\xi_T(s)$ by $\psi(\xi_T(s)) = t_{\boldsymbol{\omega}}(s)$. Define the coefficients on its tail trees in the same way. Claim 3.7 justifies these limits and the recursion (3.9).

Theorem 3.6. *Let $\mu_{\ell,\tau} > 0$ for $0 \le \ell \le 50$ and $\tau \in \{a, b, c\}$, and let $\gamma, \nu, \beta > 0$. Put $s \coloneqq \gamma/(3\nu) \in [9.99, 10.01]$ and suppose*

$$\min_{\ell,\tau} \mu_{\ell,\tau} \ge 2500(100\nu + \beta), \qquad \frac{1750\beta}{9\nu} < 1.$$

For any finite or infinite caterpillar constructed above, the supremum at total weight 3 of

$$\sum_{\ell,\tau} \mu_{\ell,\tau} H_{\ell,\tau} + \gamma(ABC - M_T) + \nu T_3 + \beta R_0$$

is attained and equals

$$\sum_{\ell,\tau} \mu_{\ell,\tau} + \gamma + \frac{4\beta^3}{9(\nu\xi_T(s) + \beta)^2}. \tag{3.10}$$

In the infinite case the supremum is attained by a summable weighting, possibly with infinite support. The supremum over finite-support weightings is the same number.

Proof.

Claim 3.7. *All these effective coefficients, finite or infinite, lie in $(1, 11)$. More precisely, for $d \in \mathcal{A}$ and $0 \le t \le 1$,*

$$\frac{8819}{8100} < s - d\psi(s) - t < 11, \qquad c(s) < \frac{2}{9}, \qquad c(s - d\psi(s) - t) < \frac{24}{25}.$$

The maps $\varphi_{d,s}$ map $[0, 1]$ into itself and are increasing contractions with common Lipschitz constant $(24/25)^3 < 1$. The compositions defining $t_{\boldsymbol{\omega}}(s)$ converge uniformly in s and $\boldsymbol{\omega}$, and the resulting effective coefficients satisfy (3.9) on every subtree rooted at a vertex of the ray.

Proof of the claim. Since $a = (2c + 1)/(3c^2)$ decreases strictly as a function of $c > 0$, direct rational comparisons at the endpoints of the s-interval give

$$\frac{7}{32} < c(s) < \frac{2}{9}, \qquad \frac{4}{25} < \psi(s) < \frac{40}{243} < \frac{1}{6}.$$

Consequently $s - 48\psi(s) - 1 > 999/100 - 48(40/243) - 1 = 8819/8100$. This last number is greater than $(2(24/25)+1)/(3(24/25)^2) = 1825/1728$, so the bound on c follows. Also $\psi(1) = 1$, which proves invariance of $[0, 1]$. Since $\psi'(a) = -c(a)^3$, we have $\varphi'_{d,s}(t) = c(s - d\psi(s) - t)^3 \le (24/25)^3$. Backward induction

proves the finite coefficient statement. The common contraction constant makes the compositions defining $t_{\boldsymbol{\omega}}(s)$ converge uniformly in both parameters. The inverse of ψ is continuous on $\psi([1,11])$, so the finite root coefficients converge to $\xi_T(s)$. Passing to the limit in the recursion on each tail of the digit sequence proves (3.9) there. □

To compare the tree recursion with the edge polynomials, we measure discrepancies between the 3 types. For nonnegative triples $\boldsymbol{x} = (x_1, x_2, x_3)$ and $\boldsymbol{y} = (y_1, y_2, y_3)$, define, with j, k denoting the 2 indices other than i in each summand,

$$D(\boldsymbol{x}) \coloneqq \sum_{i\in[3]} x_i(x_j - x_k)^2, \qquad D(\boldsymbol{x}, \boldsymbol{y}) \coloneqq \sum_{i\in[3]} x_i(y_j - y_k)^2.$$

Both quantities are nonnegative. A vector supported on a single coordinate has $D(\boldsymbol{x}) = 0$, which is why the mixed terms $D(\boldsymbol{x}, \boldsymbol{y})$ and $D(\boldsymbol{y}, \boldsymbol{x})$ are needed.

Let

$$\mathscr{D}_T \coloneqq \sum_i D(\boldsymbol{x}_i) + \sum_{ij\in E(T)} \{D(\boldsymbol{x}_i, \boldsymbol{x}_j) + D(\boldsymbol{x}_j, \boldsymbol{x}_i)\}.$$

Claim 3.8. *For either the finite or infinite tree T, with root 0, every nonnegative color weighting of finite total weight satisfies*

$$sM_T - \frac{1}{3}T_3 \geq \xi_T(s)a_0b_0c_0 - 32\mathscr{D}_T.$$

Proof of the claim. We express the tree polynomial as a sum of edge remainders and then bound each remainder. For $a > 0$ and nonnegative triples $\boldsymbol{x}, \boldsymbol{y}$, let

$$\mathcal{R}_a(\boldsymbol{x}, \boldsymbol{y}) \coloneqq ay_1y_2y_3 + \psi(a)x_1x_2x_3 - \frac{1}{3}\sum_{i\in[3]}(x_iy_jy_k + y_ix_jx_k).$$

Orient edges away from the root. The effective-coefficient recursion gives the exact decomposition

$$sM_T - \frac{1}{3}T_3 = \xi_T(s)a_0b_0c_0 + \sum_{v\to w} \mathcal{R}_{\xi_w}(\boldsymbol{x}_v, \boldsymbol{x}_w) + s\sum_{i \text{ isolated}} a_ib_ic_i.$$

We estimate each edge remainder by the deficits defined above.

With $\delta_i \coloneqq y_i - c(a)x_i$ there is an exact identity

$$\mathcal{R}_a(\boldsymbol{x}, \boldsymbol{y}) = \frac{1}{3}\sum_{i=1}^{3}\big(ay_i + (2ac(a) - 1)x_i\big)\delta_j\delta_k. \tag{3.11}$$

For $1 \leq a \leq 11$,

$$\mathcal{R}_a(\boldsymbol{x}, \boldsymbol{y}) \geq -\frac{1}{3}D(\boldsymbol{x}) - \frac{11}{3}D(\boldsymbol{y}) - \frac{4}{3}D(\boldsymbol{x}, \boldsymbol{y}) - \frac{1}{3}D(\boldsymbol{y}, \boldsymbol{x}). \tag{3.12}$$

Both sides of (3.11) vanish when $\boldsymbol{x} = (u, u, u)$ and $\boldsymbol{y} = c(a)(u, u, u)$.

To verify the identity and the lower bound, write $c \coloneqq c(a)$. Expanding the right side of (3.11) proves the identity, using $3ac^2 - 2c - 1 = 0$ and $(2ac - 1)c^2 = \psi(a)$. For the inequality use

$$\delta_j\delta_k \geq -\frac{1}{2}(\delta_j - \delta_k)^2, \qquad (\delta_j - \delta_k)^2 \leq 2(y_j - y_k)^2 + 2c^2(x_j - x_k)^2.$$

The resulting lower bound is

$$-\frac{1}{3}\{aD(\boldsymbol{y}) + \psi(a)D(\boldsymbol{x}) + (2ac - 1)D(\boldsymbol{x}, \boldsymbol{y}) + ac^2D(\boldsymbol{y}, \boldsymbol{x})\}.$$

For $a \in [1, 11]$, $c \leq 1$, $0 < 2ac - 1 = (c + 2)/(3c) < 4$, $ac^2 = (2c + 1)/3 \leq 1$, and $\psi(a) \leq 1$. Thus (3.12) holds.

Apply (3.12) to the edge remainders. Each nonroot vertex v is a child once and a parent at most 50 times. Thus the coefficient of $D(\boldsymbol{x}_v)$ in the sum of the loss terms is at most $(11 + 50)/3 < 32$. The root occurs only as a parent, so the coefficient of $D(\boldsymbol{x}_0)$ is at most $50/3 < 32$. The coefficients of $D(\boldsymbol{x}_v, \boldsymbol{x}_w)$

and $D(\boldsymbol{x}_w, \boldsymbol{x}_v)$ for $vw \in E(T)$ are at most $4/3 < 32$. Isolated products are nonnegative. This proves the bound. In the infinite case every series is absolutely convergent, since the tree has bounded degree, the total weight is finite, and the summands are cubic with uniformly bounded coefficients. Rearrangement of the displayed identity is therefore legitimate. □

For weightings of total weight 3, define the control deficit by

$$\mathscr{D} := \sum_{\ell=0}^{50} \sum_{\tau=a,b,c} (1 - H_{\ell,\tau}).$$

Each unit of $\mathscr{D}$ decreases the weighted sum of control polynomials by at least $\min_{\ell,\tau} \mu_{\ell,\tau}$. We will bound the excess of the remaining terms by $2500(100\nu + \beta)\mathscr{D}$; the hypothesis on the coefficients then makes their sum no larger than the claimed maximum.

For arbitrary class totals with $A + B + C = 3$, put $\delta_{\mathrm{cls}} := (A-1)^2 + (B-1)^2 + (C-1)^2$.

Claim 3.9. *The following statements hold.*

(i) *Each $H_{\ell,\tau}$ and each of $ABC - M_T, T_3, R_0$ is the edge polynomial of a simple 3-graph and has maximum 1 at total weight 3.*

(ii) *If $A = B = C = 1$, then*

$$H_{\ell,a} = 1 - \sum_i a_i (b_{f_\ell(i)} - c_{f_\ell(i)})^2.$$

Consequently $\mathscr{D} \geq \mathscr{D}_T$.

(iii) *If $A, B, C > 0$ and $\overline{\mathscr{D}}$ is $\mathscr{D}$ evaluated after normalizing each class to total weight one, then $\delta_{\mathrm{cls}} \leq 4\mathscr{D}$ and $ABC\,\overline{\mathscr{D}} \leq 2500\mathscr{D}$.*

Proof of the claim. For (i), first consider $H_{\ell,a}$. Fix i and write $k = f_\ell(i)$. The monomials of its summand are $a_i b_k b_j$ for $j \neq k$, $a_i b_k c_k$, $a_i b_j c_h$ for $j, h \neq k$, and $a_i c_k c_h$ for $h \neq k$. They are squarefree. The unique color of type a determines i. When 2 colors have the same type, exactly one has index $k = f_\ell(i)$, which determines its role in the monomial. The 2 mixed-type cases are disjoint. Hence every edge is counted once.

The two parenthesized sums in (3.8) add to $B + C$. Thus $H_{\ell,a} \leq A(B+C)^2/4 \leq 1$, since $A + B + C = 3$. The corresponding 3-graph contains an edge. Assigning weight one to the vertices of that edge and zero elsewhere attains value one.

In $ABC - M_T$, each triple is determined uniquely by its 3 types and indices. In T_3, the repeated index and singleton index determine the tree edge. In R_0, exactly 2 colors have index 0, so its 3 summands contain disjoint sets of monomials. The polynomial ABC is the edge polynomial of the complete 3-partite 3-graph on the color classes. Each component is the edge polynomial of a nonempty subhypergraph of that 3-graph and therefore also has maximum one at total weight 3.

For (ii), assume $A = B = C = 1$ and put $k = f_\ell(i)$. The 2 parenthesized factors in (3.8) are $1 + (b_k - c_k)$ and $1 - (b_k - c_k)$, proving the stated identity. The identity map contributes $\sum_i D(\boldsymbol{x}_i)$. Each matching involution contributes $D(\boldsymbol{x}_i, \boldsymbol{x}_j) + D(\boldsymbol{x}_j, \boldsymbol{x}_i)$ for every edge ij of its matching, together with $D(\boldsymbol{x}_i)$ at every fixed index i. All these terms are nonnegative, so summing gives $\mathscr{D} \geq \mathscr{D}_T$.

For (iii), we estimate the effect of class normalization. All terms in $\mathscr{D}$ are nonnegative. For $0 \leq A \leq 3$,

$$1 - \frac{A(3-A)^2}{4} = \frac{(A-1)^2(4-A)}{4} \geq \frac{(A-1)^2}{4}.$$

Apply this bound to $H_{0,a}, H_{0,b}, H_{0,c}$ and sum to obtain $\delta_{\mathrm{cls}} \leq 4\mathscr{D}$. If $A, B, C \geq 1/2$, completing the square in (3.8) gives

$$\sum_i a_i \left(b_{f_\ell(i)} - c_{f_\ell(i)} - \frac{B-C}{2} \right)^2 = \frac{A(B+C)^2}{4} - H_{\ell,a} \leq 1 - H_{\ell,a}.$$

To compare the normalized class weights, apply the following identity with $b = b_{f_\ell(i)}$ and $c = c_{f_\ell(i)}$,

$$\frac{b}{B} - \frac{c}{C} = \frac{2(b-c) - (B-C)}{2B} + \frac{B-C}{B}\left(\frac{1}{2} - \frac{c}{C}\right).$$

Square this identity, apply $(x+y)^2 \le 2x^2 + 2y^2$, multiply by a_i/A, and sum over i. Using the preceding deficit bound and $A, B \ge 1/2$ gives, for each ℓ,

$$\sum_i \frac{a_i}{A}\left(\frac{b_{f_\ell(i)}}{B} - \frac{c_{f_\ell(i)}}{C}\right)^2 \le 16(1 - H_{\ell,a}) + 2(B-C)^2.$$

Sum cyclically and over the 51 maps. Since $(A-B)^2 + (B-C)^2 + (C-A)^2 = 3\delta_{\mathrm{cls}}$, we obtain $\overline{\mathscr{D}} \le 16\mathscr{D} + 306\delta_{\mathrm{cls}} \le 1240\mathscr{D}$. Since $ABC \le 1$, this proves $ABC\overline{\mathscr{D}} \le 1240\mathscr{D}$ when $A, B, C \ge 1/2$. If some total is below $1/2$, then $\mathscr{D} \ge \delta_{\mathrm{cls}}/4 > 1/16$, while $\overline{\mathscr{D}} \le 153$ and $ABC \le 1$. Hence $ABC\overline{\mathscr{D}} < 2448\mathscr{D}$. Both cases give $ABC\overline{\mathscr{D}} \le 2500\mathscr{D}$. □

The root contribution is controlled by a scalar inequality. For $x, y, z \ge 0$ and $\beta, k > 0$,

$$\beta(xy + yz + zx) - kxyz \le \frac{4\beta^3}{k^2} + \frac{k}{9}D(x,y,z).$$

Equality holds at $x = y = z = 2\beta/k$.

Indeed, use $D(x,y,z) = (x+y+z)(xy+yz+zx) - 9xyz$. After subtracting $kD(x,y,z)/9$ from $\beta(xy+yz+zx) - kxyz$, the expression is $(\beta - k(x+y+z)/9)(xy+yz+zx)$. If its first factor is nonpositive there is nothing to prove. Otherwise use $xy+yz+zx \le (x+y+z)^2/3$ and maximize the resulting cubic in $x+y+z$; its maximum is $4\beta^3/k^2$, at $x+y+z = 6\beta/k$.

We now combine the tree estimate, the control bounds, and the scalar inequality to prove the claimed upper bound. First assume $A = B = C = 1$. Since $\gamma M_T - \nu T_3 = 3\nu(sM_T - T_3/3)$ and $\mathscr{D}_T \le \mathscr{D}$ by Claim 3.9(ii), Claim 3.8 gives, with $k \coloneqq 3\nu\xi_T(s) + 3\beta$,

$$\gamma(ABC - M_T) + \nu T_3 + \beta R_0 \le \gamma + \beta(a_0b_0 + b_0c_0 + c_0a_0) - ka_0b_0c_0 + 96\nu\mathscr{D}.$$

Apply the preceding scalar inequality, using $D(a_0, b_0, c_0) \le \mathscr{D}$ and $\xi_T < 11$, to obtain

$$\gamma(ABC - M_T) + \nu T_3 + \beta R_0 \le \gamma + \frac{4\beta^3}{k^2} + (100\nu + \beta)\mathscr{D}.$$

For general positive class totals, every monomial of $\gamma(ABC - M_T) + \nu T_3 + \beta R_0$ contains one variable from each color class. The value of this polynomial is therefore ABC times its value after normalizing each class to total weight one. The constant term $\gamma + 4\beta^3/k^2$ in the preceding bound is positive. Thus $ABC \le 1$ and Claim 3.9(iii) bound the noncontrol terms by $\gamma + 4\beta^3/k^2 + 2500(100\nu + \beta)\mathscr{D}$.

For the controls, $\sum_{\ell,\tau} \mu_{\ell,\tau} H_{\ell,\tau} \le \sum_{\ell,\tau} \mu_{\ell,\tau} - (\min_{\ell,\tau} \mu_{\ell,\tau})\mathscr{D}$, since each $1 - H_{\ell,\tau} \ge 0$. Adding this estimate to the preceding bound gives

$$\begin{aligned}&\sum_{\ell,\tau} \mu_{\ell,\tau} H_{\ell,\tau} + \gamma(ABC - M_T) + \nu T_3 + \beta R_0 \\ &\quad\le \sum_{\ell,\tau} \mu_{\ell,\tau} + \gamma + \frac{4\beta^3}{k^2} + \left(2500(100\nu + \beta) - \min_{\ell,\tau} \mu_{\ell,\tau}\right)\mathscr{D}.\end{aligned}$$

The coefficient of $\mathscr{D}$ is nonpositive by hypothesis, proving the upper bound (3.10). If a class has total weight zero, all 3 noncontrol terms vanish and the same bound holds.

To attain the bound, set

$$u_0 \coloneqq \frac{2\beta}{3\nu\xi_T(s) + 3\beta}, \qquad u_w \coloneqq c(\xi_w)u_v \quad \text{for every edge } v \to w.$$

Use $a_i = b_i = c_i = u_i$ on the nontrivial tree component. By Claim 3.7,

$$\sum_{i \text{ in the component}} u_i \le 25\left(1 + 48 \cdot \frac{2}{9}\right) u_0 = \frac{875}{3}u_0 < \frac{1750\beta}{9\nu} < 1.$$

Place each class's unused weight at a different reserve index, using only that type there. Every f_ℓ fixes these indices, so they contribute no control deficit. Assign zero to all other isolated slots. The class totals are now one, all controls equal one, and all remainders $\mathcal{R}_{\xi_w}(\boldsymbol{x}_v, \boldsymbol{x}_w)$ vanish on oriented tree edges $v \to w$. The root inequality is also an equality. Thus the upper bound is attained.

For the infinite tree, it remains to show that the supremum over finitely supported weightings has the same value. The weighting constructed above is summable. Truncate it to an increasing sequence of finite color sets that exhausts the color set, and rescale each truncation to total weight 3. The homogeneous cubic has nonnegative, uniformly bounded coefficients, so absolute convergence gives convergence of values to (3.10). □

We now turn this optimization into a fixed finite list whose coefficients may vary in a common neighborhood of one.

Put $Q \coloneqq (153 \cdot 10^6 + 31)2^{16} + 1 = 10027010031617$. Make a list with $10^6 \cdot 2^{16}$ copies of each of the 153 controls, $30 \cdot 2^{16}$ copies of $ABC - M_T$, 2^{16} copies of T_3, and one copy of R_0. If each entry is assigned a coefficient in $[1-10^{-4}, 1+10^{-4}]$, sum the coefficients of the repeated copies of each individual polynomial to obtain $\mu_{\ell,\tau}, \gamma, \nu, \beta$ in Theorem 3.6. Direct bounds give

$$9.99 < \frac{\gamma}{3\nu} < 10.01, \quad \min \mu_{\ell,\tau} > 2500(100\nu + \beta), \quad \frac{1750\beta}{9\nu} < 1.$$

For instance, $\nu \le (1 + 10^{-4})2^{16}$, $\beta \le 1 + 10^{-4}$, and $\min_{\ell,\tau} \mu_{\ell,\tau} \ge (1 - 10^{-4})10^6 \cdot 2^{16}$ give $\min_{\ell,\tau} \mu_{\ell,\tau} > 3 \cdot 2500(100\nu + \beta)$.

To pass from cubic to quartic components, adjoin the same new finite color z to every component edge and denote the resulting 4-graphs by $D_1(T), \dots, D_Q(T)$. For a fixed weight z, the remaining colors have total weight $4 - z$. By homogeneity, for fixed component coefficients, the maximum of the weighted sum of the quartic polynomials with this fixed z is $z((4-z)/3)^3$ times the maximum of the corresponding weighted cubic sum at total weight 3. This factor has maximum one on $0 \le z \le 4$, attained at $z = 1$, so each component has maximum one at total weight 4. For $\boldsymbol{\theta} \in [1 - 10^{-4}, 1 + 10^{-4}]^Q$, the maximum at total weight 4 of the weighted sum of these Q edge polynomials, with coefficient θ_j on $D_j(T)$, is given by (3.10), with the repeated coefficients aggregated as above. For the infinite tree determined by $\boldsymbol{\omega}$, denote this maximum by $\mathcal{B}_{\boldsymbol{\omega}}(\boldsymbol{\theta})$.

For a digit $d \in \mathcal{A}$, write d^∞ for the constant sequence $d, d, \dots$. For later use, at unit coefficients write $B_- \coloneqq \mathcal{B}_{10^\infty}(\mathbf{1})$ and $B_+ \coloneqq \mathcal{B}_{48^\infty}(\mathbf{1})$.

Restrict the components of each fixed infinite tree to an increasing sequence of finite color sets exhausting their common color set, starting with a set containing an edge of every component. Each restricted component still has maximum one at total weight 4. The finite-restriction argument in the proof of Proposition 3.4 gives uniform convergence of the weighted maxima to $\mathcal{B}_{\boldsymbol{\omega}}(\boldsymbol{\theta})$ on the coefficient box.

By Proposition 3.4(ii), for $n \ge 2Q$ there are maximizing outer coordinates for which the coefficient vector $\boldsymbol{\theta}$, defined by $\theta_j \coloneqq (3n/S)^4 z_j$ for $j \in [Q]$, lies in $[1-10^{-4}, 1+10^{-4}]^Q$. Indeed, $130\sqrt{Q}/n \le 65/\sqrt{Q} < 10^{-4}$ since $n \ge 2Q$ and $Q > 10^{13}$. Let Ω_n be the compact set of all nonnegative values of the outer coordinates $t_a, u_b, \alpha_i, \beta_i, d_*, d, Y$ satisfying $\sum_{a=1}^5 t_a + u_1 + u_2 + S + Y = 10$, $2 \le S \le 4$, $1/2 \le Y \le 4$, and this coefficient-box condition. Define

$$\Phi_n(\boldsymbol{\omega}) \coloneqq \max_{\Omega_n} \left\{ P_{\text{base}} + u_1 u_2 \left(\frac{SY}{12n} \right)^4 \mathcal{B}_{\boldsymbol{\omega}}(\boldsymbol{\theta}) \right\}. \tag{3.13}$$

Proposition 3.10. *For every $n \ge 2Q$ and $\boldsymbol{\omega} \in \mathcal{A}^{\mathbb{N}}$, let $H_{n,\boldsymbol{\omega}}$ be the countable 10-graph obtained from the component list above by (3.2). Then the following statements hold.*

(i) *The supremum of $\Lambda_{H_{n,\boldsymbol{\omega}}}$ over nonnegative weightings of total weight 10 is $\Phi_n(\boldsymbol{\omega})$. The maxima for increasing finite induced subgraphs exhausting $H_{n,\boldsymbol{\omega}}$ converge to this value.*

(ii) *There is $\varepsilon > 0$ such that*

$$(1, 1+\varepsilon] \subseteq \bigcup_{n \ge 2Q} \Phi_n(\mathcal{A}^{\mathbb{N}}).$$

Proof. For every point of Ω_n, sum the coefficients θ_j of repeated components as in Section 3.2 to obtain $\mu_{\ell,\tau}, \gamma, \nu, \beta$ for the corresponding weighted cubic sum. The coefficient-box condition ensures that Theorem 3.6 applies to this sum. Choose a maximizing weighting for it and adjoin z with weight 1. This gives a weighting of the common component vertex set of total weight 4 attaining $\mathcal{B}_{\boldsymbol{\omega}}(\boldsymbol{\theta})$. Rescale it to total weight Y. The resulting weighting of the 10-graph realizes the expression being maximized. Conversely, Proposition 3.4(ii) supplies a globally maximizing outer point in Ω_n. Restricting the outer optimization to Ω_n therefore preserves its value. Restricting a summable weighting to finite vertex sets and rescaling to total weight 10 approximates its value: the omitted monomials have sum tending to zero, since the edge coefficients are at most one and the total weight is finite. Thus finite induced subgraphs have maxima increasing to $\Phi_n(\boldsymbol{\omega})$, proving (i).

It remains to prove (ii). For a finite prefix u and $s \in [9.99, 10.01]$, the endpoint interval for sequences beginning with u is $[t_{u10^\infty}(s), t_{u48^\infty}(s)]$. Appending one digit gives its child intervals. We will show that the child intervals cover the parent interval and shrink uniformly as the prefix length increases.

Claim 3.11. *The following statements hold.*

(i) *For every $s \in [9.99, 10.01]$, the values $t_{\boldsymbol{\omega}}(s)$, as $\boldsymbol{\omega}$ ranges over $\mathcal{A}^{\mathbb{N}}$, form a nondegenerate closed interval $[t_-(s), t_+(s)]$, with $t_-(s) < 2/9$ and $t_+(s) > 4/9$. The endpoints are $t_{10^\infty}(s)$ and $t_{48^\infty}(s)$, respectively.*

(ii) *For every $s \in [9.99, 10.01]$, every finite prefix u, and $10 \le d < 48$, the endpoint intervals for adjacent one-digit extensions satisfy*

$$t_{u(d+1)10^\infty}(s) \le t_{ud48^\infty}(s).$$

The diameters of the endpoint intervals tend to zero uniformly in s and in prefixes of a given length as that length tends to infinity.

(iii) *For every fixed $n \ge 2Q$, the image of Φ_n is exactly $[\Phi_n(10^\infty), \Phi_n(48^\infty)]$.*

In particular $0 < B_- < B_+ \le Q$.

Proof of the claim. To prove (i) and (ii), first note that the maps $\varphi_{10,s}, \varphi_{48,s}$ are increasing contractions. The inequalities

$$s - 10\psi(s) - \frac{2}{9} > 8, \qquad \psi(8) = \frac{3}{16} < \frac{2}{9}, \qquad s - 48\psi(s) - \frac{4}{9} < 2, \qquad \psi(2) > \frac{4}{9}$$

locate the fixed points $t_-(s)$ of $\varphi_{10,s}$ and $t_+(s)$ of $\varphi_{48,s}$ as claimed. Since $t_+(s) - t_-(s) > 2/9 > \psi(s)$, we have $s - (d+1)\psi(s) - t_-(s) > s - d\psi(s) - t_+(s)$ for $10 \le d < 48$. The function ψ is decreasing, so

$$\varphi_{d+1,s}(t_-(s)) \le \varphi_{d,s}(t_+(s)).$$

Since $\varphi_{10,s}(t_-(s)) = t_-(s)$ and $\varphi_{48,s}(t_+(s)) = t_+(s)$, the 39 intervals $\varphi_{d,s}([t_-(s), t_+(s)])$, $d \in \mathcal{A}$, cover $[t_-(s), t_+(s)]$. Iterating the maps gives nested image intervals, and every point of a parent interval belongs to a child interval. Since the diameters tend uniformly to zero, successive choices of child intervals yield a digit sequence for every target value. Composing with the maps corresponding to a fixed prefix preserves the displayed inequalities.

If $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$ agree on their first h digits, the common contraction bound gives

$$\sup_{s \in [9.99, 10.01]} |t_{\boldsymbol{\omega}}(s) - t_{\boldsymbol{\omega}'}(s)| \le \left(\frac{24}{25}\right)^{3h}.$$

Formula (3.10) therefore makes their values of $\mathcal{B}_{\boldsymbol{\omega}}(\boldsymbol{\theta})$ uniformly close throughout the coefficient box. The only sequence dependence is in the final term of (3.10), which increases strictly with $t_{\boldsymbol{\omega}}(s)$, since ψ decreases strictly. Thus the intervals $[\mathcal{B}_{u10^\infty}(\boldsymbol{\theta}), \mathcal{B}_{u48^\infty}(\boldsymbol{\theta})]$ associated with finite prefixes u satisfy the same overlap inequalities pointwise in $\boldsymbol{\theta}$, and their diameters tend uniformly to zero. At unit coefficients this also proves $0 < B_- < B_+ \le Q$.

For (iii), fix $n \ge 2Q$. The preceding continuity and pointwise overlap estimates hold uniformly on the coefficient box, including the varying parameter $s = \gamma/(3\nu)$. In (3.13), the domain Ω_n is fixed, the coefficient of $\mathcal{B}_{\boldsymbol{\omega}}$ is nonnegative and bounded, and the remaining term is independent of the sequence. Taking maxima therefore makes Φ_n continuous and coordinatewise nondecreasing and preserves every overlap for adjacent one-digit extensions. Apply Lemma 2.4 with $D = \mathcal{A}$ to obtain (iii). □

Having obtained the full image interval for each fixed n, we compare its endpoints as n varies. By Proposition 3.4(i) and the definitions of B_- and B_+,

$$\Phi_n(10^\infty) = 1 + \frac{B_-}{4^5 n^4} + O_Q(n^{-5}), \qquad \Phi_n(48^\infty) = 1 + \frac{B_+}{4^5 n^4} + O_Q(n^{-5}).$$

Both endpoint values exceed one by the test weighting and tend to one as $n \to \infty$. Since $B_+ > B_-$, the endpoint expansions imply that consecutive image intervals overlap for all sufficiently large n. Indeed,

$$\Phi_{n+1}(48^\infty) - \Phi_n(10^\infty) = \frac{B_+ - B_-}{4^5 n^4} + O_Q(n^{-5}) > 0,$$

and the same endpoint expansions give $\Phi_n(48^\infty) > \Phi_{n+1}(10^\infty)$. Their union contains $(1, 1+\varepsilon]$ for some $\varepsilon > 0$. Figure 7 shows the two interval-covering steps.

□

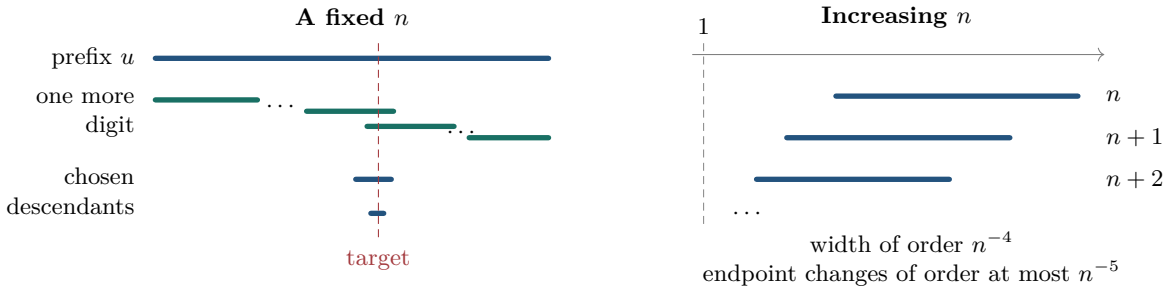


Figure 7: The two interval arguments in Proposition 3.10. At fixed n, overlapping child intervals cover each prefix interval and chosen descendants shrink to a target; some child intervals are omitted. The image intervals $\Phi_n(\mathcal{A}^{\mathbb{N}})$ then overlap for consecutive large n and approach 1. Positions and lengths are schematic.

3.3 The initial interval and the classification

Proof of Theorem 3.1. Fix $n \geq 2Q$ and $\boldsymbol{\omega} \in \mathcal{A}^{\mathbb{N}}$. Let $\varepsilon > 0$ be as in Proposition 3.10. By part (i), the normalized Lagrangians of finite induced subgraphs of $H_{n,\boldsymbol{\omega}}$ increase to $\alpha_{10}\Phi_n(\boldsymbol{\omega})$. Each finite Lagrangian is an ordinary Turán density, and $\Pi_\infty^{(10)}$ is closed. Hence this limit belongs to $\Pi_\infty^{(10)}$. Part (ii) of the proposition and the single-edge density give $[\alpha_{10}, (1+\varepsilon)\alpha_{10}] \subseteq \Pi_\infty^{(10)}$.

Peng [36] showed that if $c\alpha_r$ is a non-jump for r-graphs with $r \geq 3$, then $c\alpha_s$ is a non-jump for every $s \geq r$. Here the same relative scaling preserves an entire interval in $\Pi_\infty^{(r)}$. For $r > 10$, adjoin the same $r-10$ new vertices to every edge of a 10-graph. For a weighting of the resulting r-graph with total weight r, let $t \in [0, r]$ be the total weight on the original vertices. Homogeneity and the AM–GM inequality bound the new polynomial by the original maximum at total weight 10 multiplied by $(t/10)^{10}((r-t)/(r-10))^{r-10} \leq 1$. This bound is attained by taking an optimal weighting of total weight 10 on the original vertices and assigning weight 1 to every new vertex. The maxima at total weights 10 and r are therefore equal, so the operation preserves the relative Lagrangian. Applying it to approximating finite 10-graphs and taking limits proves Theorem 3.1. □

Products of the endpoint interval give an interval of fixed relative width in a sufficiently large uniformity. We then extend this interval to one by using a finite sequence of target values from the preceding uniformity.

Proof of Theorem 1.1(ii). We use the hypergraph product studied by Grosu [14, Section 4]. For an r-graph G and an s-graph F on disjoint vertex sets, let $G * F$ have edges $e \cup f$ for $e \in G$ and $f \in F$. For each fixed $t \in [0,1]$, the maximum of Λ_{G*F} over weightings assigning total weights t and $1-t$ to $V(G)$ and $V(F)$, respectively, is $t^r(1-t)^s\lambda(G)\lambda(F)/(r!s!)$. The factor $t^r(1-t)^s$ is maximized at $t = r/(r+s)$. Multiplying the resulting value by $(r+s)!$ gives $\lambda(G*F) = \frac{\alpha_{r+s}}{\alpha_r\alpha_s}\lambda(G)\lambda(F)$. By density of finite Lagrangians and closedness of $\Pi_\infty^{(r+s)}$, this gives $\alpha_{r+s}\frac{a}{\alpha_r}\frac{b}{\alpha_s} \in \Pi_\infty^{(r+s)}$ whenever $a \in \Pi_\infty^{(r)}$ and

$b \in \Pi_\infty^{(s)}$. Let $\varepsilon > 0$ be as in Theorem 3.1, and choose a positive integer m with $(1+\varepsilon)^m \geq 3$. Products of m values from the interval in uniformity 10 give $[\alpha_{10m}, 3\alpha_{10m}] \subseteq \Pi_\infty^{(10m)}$. Adding common vertices to every edge, as in the proof of Theorem 3.1, preserves this relative width. Thus $[\alpha_r, 3\alpha_r] \subseteq \Pi_\infty^{(r)}$ for every $r \geq 10m$. Figure 8 illustrates the product and the addition of common vertices.

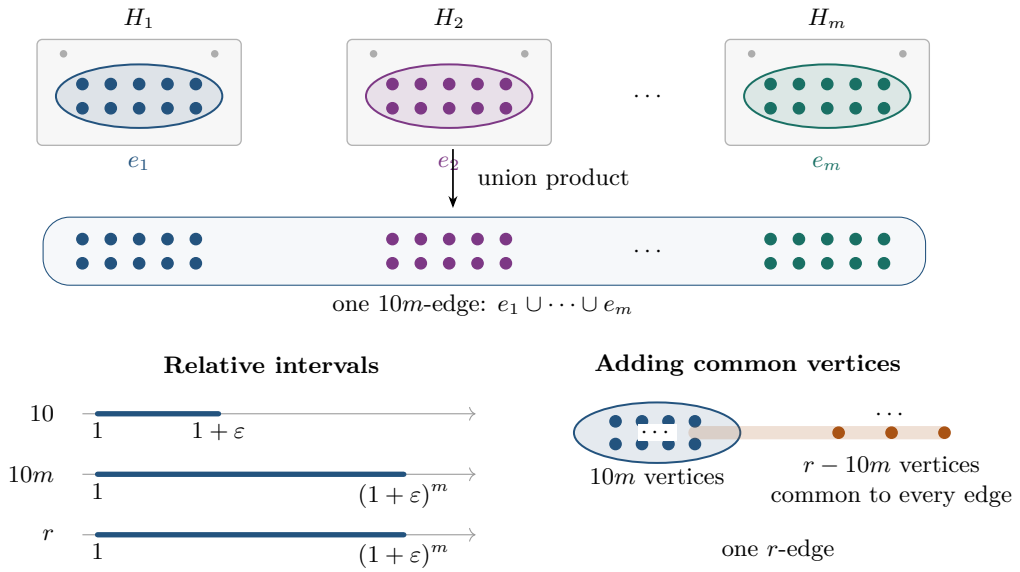


Figure 8: Products and lifting. Choose an edge e_i in each of m disjoint 10-graphs H_i; their union is an edge of the product. Adjoining the same $r - 10m$ new vertices to every product edge preserves the Lagrangian divided by α_r. The interval panel uses this relative normalization, with uniformities listed at the left; lengths are schematic.

We use Lemma 3.3 to extend the interval, choosing successive targets from the preceding uniformity.

Claim 3.12. *For every integer $s \geq 2$, there is a finite set $\mathcal{T}_s = \{c_0, \ldots, c_m\} \subseteq \Pi_\infty^{(s)}$ such that $\alpha_s = c_0 < \cdots < c_m = 1$ and $c_i \leq \frac{3}{2} c_{i-1}$ for every $1 \leq i \leq m$.*

Proof. Take $\mathcal{T}_2 := \{1/2, 2/3, 1\}$. The first two values are the normalized Lagrangians of K_2, K_3, respectively, by the Motzkin–Straus theorem [30], and the normalized Lagrangians of complete graphs tend to 1 as their order tends to infinity, giving $1 \in \Pi_\infty^{(2)}$ by closedness. For $s \geq 3$, define $\mathcal{T}_s := p_s \mathcal{T}_{s-1} \cup \mathcal{T}_{s-1}$. Membership in $\Pi_\infty^{(s)}$ follows from the inclusions $p_s \Pi_\infty^{(s-1)} \subseteq \Pi_\infty^{(s)}$ and $\Pi_\infty^{(s-1)} \subseteq \Pi_\infty^{(s)}$, and the minimum and maximum are α_s and 1.

Each of the two sets has consecutive ratios at most $3/2$. For $s = 3$, the intervals spanned by these sets are $[\alpha_3, 4/9]$ and $[1/2, 1]$, and their intervening endpoints satisfy $(1/2)/(4/9) = 9/8 < 3/2$. For $s \geq 4$, these intervals overlap because $p_s > e^{-1} > 1/3 > \alpha_{s-1}$. In this case, for an element x of the union below p_s, choose the least element of $p_s \mathcal{T}_{s-1}$ greater than x; for $p_s \leq x < 1$, choose the least element of $\mathcal{T}_{s-1}$ greater than x. Call the chosen element b and its predecessor in the selected set a. Then $a \leq x < b$. Thus $b \leq (3/2)a \leq (3/2)x$, and the successor of x in the union is at most b. Hence the union also has consecutive ratios at most $3/2$. □

Claim 3.13. *If $r \geq 3$ and $[\alpha_r, 3\alpha_r] \subseteq \Pi_\infty^{(r)}$, then $[\alpha_r, 1] \subseteq \Pi_\infty^{(r)}$.*

Proof. Since $3p_r > 1$, the hypothesis gives $[\alpha_r, \alpha_{r-1}] \subseteq \Pi_\infty^{(r)}$. Use the elements of $\mathcal{T}_{r-1}$ from Claim 3.12 as successive upper endpoints, starting at α_{r-1}. Suppose that v, c are consecutive elements and $[\alpha_r, v] \subseteq \Pi_\infty^{(r)}$. Then $v < c \leq 3v/2$ and $\alpha_r \leq p_r v$. Monotonicity, the explicit formula for Θ_r, and Bernoulli's inequality give

$$\Theta_r(\alpha_r, c) \leq \frac{3}{2} p_r v \left(1 - \frac{2p_r}{3r}\right)^{-(r-1)} \leq \frac{9p_r}{6 - 4p_r} v \leq \frac{18}{19} v < v.$$

Here Bernoulli's inequality gives $(1-2p_r/(3r))^{r-1} \geq 1-2(r-1)p_r/(3r) \geq 1-2p_r/3 > 0$. The fraction $(3/2)p_r/(1-(2/3)p_r)$ is increasing in $p_r \leq 4/9$, so it is at most $18/19$. Lemma 2.5 and closedness therefore extend $[\alpha_r, v]$ to $[\alpha_r, c]$. The target set is finite and ends at 1, proving the claim. □

For every $r \geq 10m$, the interval obtained from products satisfies the hypothesis of Claim 3.13. Hence $[\alpha_r, 1] \subseteq \Pi_\infty^{(r)}$. Forbidding a single edge gives $0 \in \Pi_\infty^{(r)}$, and the numerical gap from Erdős's theorem [5, Theorem 1] excludes every value in $(0, \alpha_r)$. This proves Theorem 1.1(ii). □

4 The ordinary three-uniform interval

We prove the interval theorem in Theorem 1.1(iii), starting with the three-uniform interval $[4/9, 1]$ and then extending it to higher uniformities. We use $\Pi_3 = \Pi_\infty^{(3)}$ as in the introduction. The normalized Lagrangian remains the one defined in Section 3; for a 3-graph it contains the factor $3! = 6$.

A number $a \in [0, 1)$ is a *jump* if $(a, a+\delta) \cap \Pi_3 = \varnothing$ for some $\delta > 0$; otherwise it is a *non-jump*. Frankl, Peng, Rödl, and Talbot [9, Theorem 1.5] proved that $5/9$ is a non-jump for 3-graphs. Yan and Peng [51] lowered this value to $12/25$ and also obtained infinitely many irrational non-jumps for every $r \geq 3$. Liu and Mubayi [26] proved that $4/9$ is a non-jump for 3-graphs. Their construction starts from the pattern consisting of all triples with one vertex in one part and two in another, whose optimal density is $4/9$. They conjectured that all numbers below $4/9$ are jumps and all numbers in $[4/9, 1)$ are non-jumps. The interval theorem establishes the second assertion by realizing every number in this interval as an ordinary Turán density.

For each $\alpha \in [4/9, 1]$ we construct finite 3-graphs H_j whose normalized Lagrangians satisfy $\lambda(H_j) \leq \alpha$ and $\lambda(H_j) \to \alpha$. For $\alpha \in [0, 1]$, put

$$\mathcal{F}_\alpha \coloneqq \{F \colon F \text{ is a finite 3-graph and } \lambda(F) > \alpha\}.$$

The proof has two stages. The local cone inequality and a sparse host give an interval beginning at $4/9$. Rounding internal substitution coefficients then enlarges this interval using stars and books, after which 3 explicit substitutions reach one.

4.1 Substitution and the local cone inequality

We shall keep track of approximation from below. Define

$$\mathcal{A} \coloneqq \{\alpha \in [0,1] \colon \text{for every } \delta > 0 \text{ there is a finite } H \text{ with } \alpha - \delta < \lambda(H) \leq \alpha\}.$$

The next lemma shows that it suffices to prove $[4/9, 1] \subseteq \mathcal{A}$.

Lemma 4.1. *For every $\alpha \in \mathcal{A}$, the family $\mathcal{F}_\alpha$ satisfies $\pi(\mathcal{F}_\alpha) = \alpha$. Moreover, if $\alpha_j \in \mathcal{A}$, $\alpha_j \leq \alpha$, and $\alpha_j \to \alpha$, then $\alpha \in \mathcal{A}$.*

Proof. A finite H is $\mathcal{F}_\alpha$-free exactly when $\lambda(H) \leq \alpha$. Uniform weights give the upper density bound. Conversely, blowups of any H with $\lambda(H) \leq \alpha$ are still $\mathcal{F}_\alpha$-free, since aggregating weights in each class preserves the Lagrangian; their densities approach $\lambda(H)$. Approximation from below proves equality. Choose a finite H_j with $\alpha_j - 1/j < \lambda(H_j) \leq \alpha_j$. Then $\lambda(H_j) \leq \alpha$ and $\lambda(H_j) \to \alpha$, so $\alpha \in \mathcal{A}$. □

We next record the exact effect of replacing vertices by internal 3-graphs.

Lemma 4.2. *Let B be a finite 3-graph on $[m]$. For $\boldsymbol{\beta} \in [0,1]^m$, define*

$$\Phi_B(\boldsymbol{\beta}) \coloneqq \max_{\boldsymbol{z} \in S([m])} \left(6 \sum_{ijk \in B} z_i z_j z_k + \sum_{i \in [m]} \beta_i z_i^3 \right).$$

If every β_i belongs to $\mathcal{A}$, then $\Phi_B(\boldsymbol{\beta}) \in \mathcal{A}$. The function Φ_B is nondecreasing in each coordinate and is 1-Lipschitz in the maximum norm.

Proof. First suppose that $\beta_i = \lambda(H_i)$ for finite 3-graphs H_i. Replace vertex i of B by a disjoint copy of H_i, keep all internal edges, and insert all transversal triples corresponding to edges of B. Once the total mass z_i of each copy is fixed, its internal weighting can be optimized independently. The Lagrangian of the resulting finite 3-graph is therefore exactly $\Phi_B(\boldsymbol{\beta})$.

Monotonicity follows directly from the defining polynomial. If two coefficient vectors differ by at most ε in maximum norm, their values at any $\boldsymbol{z} \in S([m])$ differ by at most $\varepsilon \sum_i z_i^3 \le \varepsilon$. Taking maxima proves the Lipschitz bound. For general coefficients in $\mathcal{A}$, approximate each β_i from below by finite Lagrangians and apply this bound to the finite substitution graphs. □

We will use a discrete set of coefficients approaching $4/9$. Let S_m be the complete 3-star with one apex and m leaves. At fixed apex mass a, equal leaf weights maximize its polynomial. Consequently,

$$\beta_m \coloneqq \lambda(S_m) = \max_{0\le a\le 1} 3\left(1-\frac{1}{m}\right) a(1-a)^2 = \frac{4}{9}\left(1-\frac{1}{m}\right). \tag{4.1}$$

The formula includes $m = 1$, when both sides are zero. In particular, $\beta_m \in \mathcal{A}$ for every $m \ge 1$, and Lemma 4.1 gives $4/9 \in \mathcal{A}$.

In subsequent maximization arguments, we restrict to the positive support before applying Lagrange multipliers. If a homogeneous cubic polynomial has value p at a maximum on a simplex, its first derivative in every positive coordinate is $3p$. The derivative in a zero coordinate is at most $3p$. These observations follow from stationarity and Euler's identity and also apply to the substitution polynomials above.

For a 3-graph Q, let $d_Q(i,j)$ be the number of edges containing $\{i,j\}$, and let $\Delta_2(Q) \coloneqq \max_{i\ne j} d_Q(i,j)$. We call Q *sparse* if $|Q[S]| \le |S| - 2$ for every $S \subseteq V(Q)$ with $|S| \ge 2$. Its *cone* is obtained by adding a new vertex, called the *apex*, and all triples consisting of the apex and two distinct vertices of Q.

We need a cone inequality that remains valid when internal 3-graphs are substituted at the bottom vertices. For Q on $[n]$ and nonnegative coefficients $\boldsymbol{\beta} = (\beta_i)_{i\in[n]}$, define

$$F_Q(\boldsymbol{\beta}) \coloneqq \max\left\{ 3a\left[\left(\sum_i y_i\right)^2 - \sum_i y_i^2\right] + 6\sum_{ijk\in Q} y_i y_j y_k + \sum_i \beta_i y_i^3 \right\}, \tag{4.2}$$

where the maximum is over nonnegative weights $a, y_1, \dots, y_n$ summing to one. Here a is the apex mass, and the coefficient β_i records the density substituted inside bottom part i. When all β_i vanish, the following result is the local cone estimate of Liu and Mubayi [26, Theorem 3.1]. The additional cubic terms are needed for the recursion in the proof of Proposition 4.7. Figure 9 shows the substitution and the three edge contributions.

Lemma 4.3. *If Q is sparse and $\Delta_2(Q) \le 2$, then*

$$F_Q(\boldsymbol{\beta}) \le \max\left\{\frac{4}{9}, \max_i \beta_i\right\}.$$

Proof. We first prove the assertion when every coefficient equals $4/9$. For $\boldsymbol{z} \in S([n])$, put $\rho \coloneqq \sum_i z_i^2$. If the total bottom mass is $b > 0$, write $y_i = bz_i$. The objective in Equation (4.2) becomes

$$3(1-b)b^2(1-\rho) + b^3\big(6\Lambda_Q(\boldsymbol{z}) + \frac{4}{9}\sum_i z_i^3\big), \qquad 0 \le b \le 1. \tag{4.3}$$

For $b = 0$, its value is zero. We will bound Equation (4.3) by $4/9$ for every $\boldsymbol{z}$ and b.

First, sparsity alone gives

$$\Lambda_Q(\boldsymbol{z}) \le \frac{1}{27}, \qquad 24\Lambda_Q(\boldsymbol{z}) + \sum_i z_i^3 \le 1. \tag{4.4}$$

For the first inequality, choose a maximizer of $\Lambda_Q(\boldsymbol{z})$ with minimum positive support S. If two positive coordinates belong to no common edge of $Q[S]$, keeping their sum fixed makes $\Lambda_Q(\boldsymbol{z})$ affine in either coordinate. Moving to an endpoint does not decrease the value and reduces the support. Thus every pair in S is covered. If $N \coloneqq |S| \ge 2$, pair coverage and sparsity imply

$$\binom{N}{2} \le 3|Q[S]| \le 3(N-2),$$

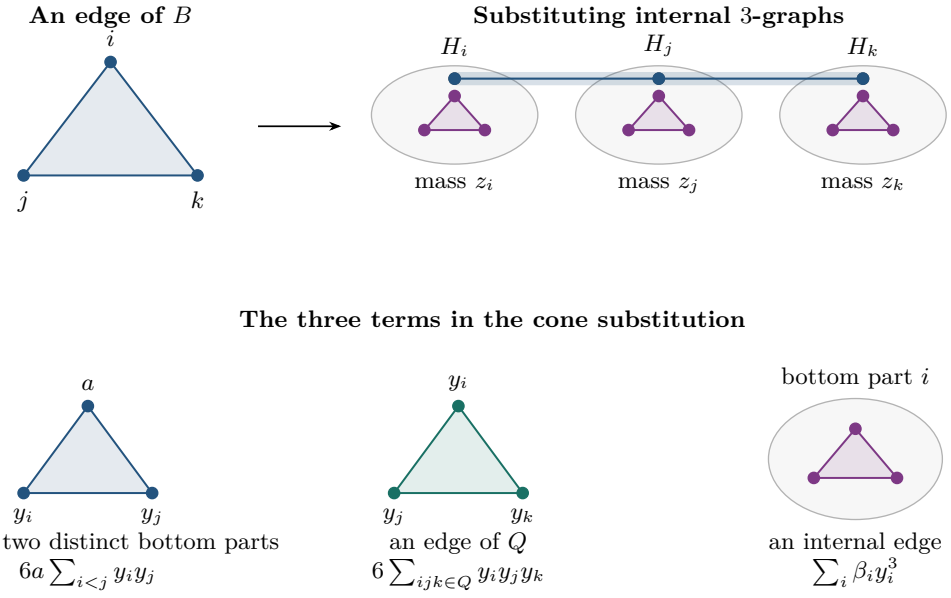


Figure 9: Substitution and the cone polynomial. A base edge supplies all transversal triples between the corresponding copies, while each copy retains its internal edges. The lower panels separate the three contributions to Equation (4.2); labels there denote masses. In the upper-right panel, only one transversal edge and one internal edge in each copy are shown.

so $N \in \{3, 4\}$. Two triples on four vertices cover at most five pairs, which excludes $N = 4$. Hence a nonzero maximum is attained on a single edge and is at most $1/27$.

For the second inequality, maximize $24\Lambda_Q(\boldsymbol{z}) + \sum_i z_i^3$. At an uncovered positive pair, the contribution of $\Lambda_Q(\boldsymbol{z})$ is affine under the same transfer, while the sum of the two cubes is strictly convex. A maximizer therefore has support one or three, by the preceding argument. Support one gives value one. On three coordinates $x + y + z = 1$, the required inequality follows from

$$1 - (x^3 + y^3 + z^3 + 24xyz) = 3\left(\sum_{\mathrm{sym}} x^2y - 6xyz\right) \geq 0.$$

This proves Equation (4.4). In particular,

$$\Lambda_Q(\boldsymbol{z}) + \frac{2}{27}\sum_i z_i^3 \leq \frac{1}{24} + \frac{7}{216}\sum_i z_i^3 \leq \frac{2}{27}, \qquad 6\Lambda_Q(\boldsymbol{z}) + \frac{4}{9}\sum_i z_i^3 \leq \frac{4}{9}. \tag{4.5}$$

For the remaining weighted estimate, define

$$\tau(\rho) := \frac{1 - \rho - (1-\rho)^{3/2}}{2}, \qquad \Phi(\boldsymbol{z}) := \Lambda_Q(\boldsymbol{z}) + \frac{2}{27}\sum_i z_i^3 - \tau(\rho).$$

We claim that

$$\Phi(\boldsymbol{z}) < 0 \qquad \text{whenever } \rho \leq \frac{5}{9}. \tag{4.6}$$

To see why this suffices, first suppose $\rho \geq 5/9$. Then $1-\rho \leq 4/9$, and Equation (4.5) bounds Equation (4.3) by $\frac{4}{9}(3b^2 - 2b^3) \leq 4/9$. If $\rho \leq 5/9$, then Equation (4.6) gives $6\Lambda_Q(\boldsymbol{z}) + \frac{4}{9}\sum_i z_i^3 \leq 3(1-\rho)(1-\sqrt{1-\rho})$. With $x := b\sqrt{1-\rho} \in [0,1]$, the expression in Equation (4.3) is at most $3x^2(1-x) \leq 4/9$.

For $1/4 \leq \rho \leq 5/9$, use $\sum_i z_i^3 \leq \rho^{3/2}$ and Equation (4.4). It is enough to show that

$$H(\rho) := \tau(\rho) - \frac{1 + 2\rho^{3/2}}{27} > 0.$$

We have $H''(\rho) = -3/(8\sqrt{1-\rho}) - 1/(18\sqrt{\rho}) < 0$, while

$$H\left(\frac{1}{4}\right) = \frac{142 - 81\sqrt{3}}{432} > 0, \qquad H\left(\frac{5}{9}\right) = \frac{27 - 10\sqrt{5}}{729} > 0.$$

Concavity proves positivity throughout this interval.

It remains to consider $\rho < 1/4$. Suppose that Φ is nonnegative somewhere in this range, and maximize it over the compact subset $\rho \le 1/4$ of the simplex. The preceding paragraph excludes the boundary $\rho = 1/4$. Restrict to the positive support S of a maximizer and put $N \coloneqq |S|$. Since $\rho \ge 1/N$, we have $N \ge 5$. The constraint $\rho \le 1/4$ is inactive, so Lagrange multipliers apply on the relative interior of this face.

Set

$$s \coloneqq \sqrt{1-\rho}, \qquad B \coloneqq \tau'(\rho) = \frac{3s-2}{4}, \qquad A \coloneqq 3\tau(\rho) - 2\rho B = \frac{(1-s)(2-s)}{2},$$

and define $d_i \coloneqq \sum_{jk\colon ijk \in Q[S]} z_j z_k$ and $D \coloneqq \sum_{i \in S} d_i$. Stationarity gives

$$d_i + \frac{2}{9} z_i^2 - 2B z_i = \mu \qquad \text{for every } i \in S.$$

Multiplying by z_i and summing yields $\mu = 3\Phi + A \ge A$. Summing without weights therefore gives

$$D = N\mu + 2B - \frac{2}{9}\rho \ge NA + 2B - \frac{2}{9}\rho. \tag{4.7}$$

To bound D from above, form a graph whose vertices are the edges of $Q[S]$, joining two when they share a pair. This graph is a forest. Indeed, if $E_1, \dots, E_\ell, E_1$ is a cycle, each of $E_2, \dots, E_{\ell-1}$ adds at most one vertex to the union of its predecessors. The two pairs $E_\ell \cap E_{\ell-1}$ and $E_\ell \cap E_1$ are distinct, since otherwise one pair belongs to three hyperedges. Their union is E_ℓ, so the final edge adds no vertex. The ℓ edges then use at most $\ell + 1$ vertices, contrary to sparsity.

Let m be the number of hyperedges and m_2 the number of pairs of codegree two. Such pairs are in bijection with the edges of the forest. Thus $m_2 \le m - 1$ when $m > 0$, and

$$\sum_{j<k} d_{Q[S]}(j,k)^2 = 3m + 2m_2 \le 5m - 2 \le 5N - 12. \tag{4.8}$$

The same bound holds for $m = 0$.

By Equation (4.8) and Cauchy–Schwarz,

$$D^2 = \left(\sum_{j<k} d_{Q[S]}(j,k) z_j z_k \right)^2 \le \frac{5N-12}{2} \left(\rho^2 - \sum_i z_i^4 \right).$$

Hölder's inequality implies $\rho \le (\sum_i z_i^4)^{1/3} (\sum_i z_i)^{2/3}$, and thus $\sum_i z_i^4 \ge \rho^3$. With $Y \coloneqq \sqrt{(5N-12)/2}$, we obtain $D \le \rho s Y$.

Since $N = (2Y^2 + 12)/5$, the lower bound in Equation (4.7) minus this upper bound equals

$$NA + 2B - \frac{2}{9}\rho - \rho s Y = \frac{2A}{5} \left(Y - \frac{5\rho s}{4A} \right)^2 + \frac{12A}{5} + 2B - \frac{2}{9}\rho - \frac{5\rho^2 s^2}{8A}. \tag{4.9}$$

Substituting $\rho = 1 - s^2$ and the definitions of A, B, the last four terms are $q(s)/(180(2-s))$, where

$$q(s) \coloneqq 225s^5 + 225s^4 - 481s^3 + 665s^2 - 968s + 424.$$

For $\sqrt{3}/2 \le s \le 1$, the identity

$$q'(s) = 1125s^4 + 900s^3 - 1081 + (1-s)(1443s + 113)$$

gives $q'(s) \ge 2025 \cdot 9/16 - 1081 = 929/16 > 0$. Also,

$$q\left(\frac{\sqrt{3}}{2}\right) = \frac{33578 - 19235\sqrt{3}}{32} > 0,$$

because $33578^2 - 3 \cdot 19235^2 = 17526409 > 0$. Thus Equation (4.9) is positive, contradicting the two bounds on D. This proves Equation (4.6), and hence the assertion when all coefficients equal $4/9$.

Finally, let $M \coloneqq \max\{4/9, \max_i \beta_i\}$. At any feasible weighting, the objective with coefficients $\boldsymbol{\beta}$ is at most the objective with all coefficients equal to $4/9$ plus $(M - 4/9) \sum_i y_i^3 \le M - 4/9$. Its maximum is therefore at most M, as required. □

4.2 Sparse hosts and the initial interval

To make the width of the initial interval explicit, we use a quantitative matching lemma. All logarithms are natural. A hypergraph is *linear* if any two distinct edges intersect in at most one vertex.

Lemma 4.4. *Let H be a linear 3-graph on $N \geq 3$ vertices. Suppose that $D \geq e^{1000}\log(2N)$ and*

$$(1-e^{-100})D \leq d_H(v) \leq (1+e^{-100})D \qquad \text{for every } v \in V(H).$$

Then H has a matching covering at least 99% of its vertices.

The proof in Section 4.3 uses independent edge activations and a bounded-differences estimate.

The local inequality in Lemma 4.3 will exclude maximizing weight vectors with small support. We also need a bound on the number of edges in sets of intermediate size. The construction below obtains both properties by deleting a small number of edges from the union of two partial Steiner triple systems.

Applied to the auxiliary hypergraph whose vertices are pairs and whose edges are the 3 pairs of a triple, Lemma 4.4 finds a partial Steiner triple system in a 3-graph with nearly equal pair degrees.

Lemma 4.5. *For every integer n with $\log n \geq 10^8$, put $L_n := e^{2000}(\log n)^2$. There is a 3-graph Q_n on $[n]$ such that*

$$\Delta_2(Q_n) \leq 2, \qquad |Q_n| = \left\lfloor \frac{3n^2}{10} \right\rfloor, \tag{4.10}$$

$$|Q_n[S]| \leq |S| - 2 \qquad \text{for } 2 \leq |S| \leq 200, \tag{4.11}$$

$$|Q_n[S]| \leq \frac{|S|^2}{13} \qquad \text{for } 200 < |S| \leq \frac{n}{100L_n}. \tag{4.12}$$

The codegree and sparsity bounds are inherited by every subgraph of Q_n.

Proof. Write $T := \log n$, $L := L_n$, and $\delta := e^{-100}$. We shall use

$$\log L = 2000 + 2\log T \leq \frac{T}{1000} \qquad \text{for } T \geq 10^8. \tag{4.13}$$

This holds at $T = 10^8$, and the derivative of the left side is $2/T < 1/1000$. In particular, $0 < L/n < 1$.

Take independent random 3-graphs H_1, H_2, each including every triple of $[n]$ independently with probability L/n. We verify four properties, each with failure probability less than $1/8$.

For a fixed pair in H_i, its codegree is binomial with mean $L(1-2/n)$. Since $n \geq 8/\delta$, the two-sided Chernoff bound gives

$$\mathbb{P}\big[d_{H_i}(x,y) \notin [(1-\delta)L, (1+\delta)L]\big] \leq 2\exp\left(-\frac{\delta^2 L}{24}\right).$$

Indeed, the mean differs from L by at most $\delta L/4$, and deviation outside the stated interval requires deviation greater than $\delta L/2$ from the mean. A union bound over both hypergraphs and all pairs is at most

$$2n^2 \exp\left(-\frac{\delta^2 L}{24}\right) = 2\exp\left(2T - \frac{e^{1800}T^2}{24}\right) < \frac{1}{8}.$$

When all these degree bounds hold, the auxiliary 3-graph for each H_i is linear and has all degrees in $[(1-\delta)L, (1+\delta)L]$. Its order is $\binom{n}{2}$, and $L \geq e^{1000}\log\binom{n}{2}$. By Lemma 4.4, H_i contains a partial Steiner triple system M_i; we retain the weaker estimate

$$|M_i| \geq \frac{24}{25}\frac{n(n-1)}{6} \qquad \text{for } i \in \{1,2\}.$$

Next, independence gives $\mathbb{E}|H_1 \cap H_2| \leq nL^2/6$. Markov's inequality and Equation (4.13) imply

$$\mathbb{P}\left[|H_1 \cap H_2| > \frac{n^2}{1000}\right] \leq \frac{1000}{6}\exp(2\log L - T) \leq \frac{1000}{6}e^{-499T/500} < \frac{1}{8}.$$

Put $H \coloneqq H_1 \cup H_2$. Its distinct edge indicators are independent and each has success probability at most $2L/n$. Let X count pairs (S, A) with $2 \le s = |S| \le 200$ and A a set of $s-1$ edges of $H[S]$. Then

$$\mathbb{E}[X] \le \sum_{2 \le s \le 200} \binom{n}{s} \binom{\binom{s}{3}}{s-1} \left(\frac{2L}{n}\right)^{s-1} \le 200n(16000000L)^{199}.$$

By Equation (4.13),

$$\log\left(\frac{\mathbb{E}[X]}{n^2}\right) \le 4000 + 199 \log L - T \le -\frac{4T}{5}.$$

Thus $\mathbb{P}[X > n^2/1000] \le 1000e^{-4T/5} < 1/8$. Choose one edge from each counted configuration. Deleting these chosen edges from any subgraph of H removes at most X edges and leaves at most $s-2$ edges on every s-vertex set with $2 \le s \le 200$.

For the last property, fix $201 \le s \le n/(100L)$ and set $t_s \coloneqq \lfloor s^2/13 \rfloor + 1$. For a fixed s-set S, the mean of $|H[S]|$ is at most $Ls^3/(3n)$. The binomial upper-tail bound and a union bound over S give

$$\mathbb{P}\left[\exists S\colon |S| = s,\ |H[S]| > \frac{s^2}{13}\right] \le \exp(sF(s)), \tag{4.14}$$

where

$$F(s) \coloneqq 1 + \log\frac{n}{s} - \frac{s}{13}\log\frac{3n}{13eLs}.$$

Here $e\mathbb{E}|H[S]|/t_s \le 13eLs/(3n) \le 13e/300 < 1$, which justifies replacing the tail exponent t_s by $s^2/13$. On the stated range, $\log(3n/(13eLs)) > 2$, so

$$F'(s) = -\frac{1}{s} - \frac{1}{13}\left(\log\frac{3n}{13eLs} - 1\right) < 0.$$

Moreover, Equation (4.13) gives

$$F(201) = 1 - \frac{188}{13}(T - \log 201) + \frac{201}{13}\log\left(\frac{13eL}{3}\right) \le 135 - \frac{187799}{13000}T < -10T.$$

Summing Equation (4.14) over all admissible s gives a failure probability at most $\sum_{s \ge 201} e^{-10Ts} < 1/8$.

Fix H_1, H_2 for which all four properties hold, and choose M_1, M_2 as above. Their union has maximum codegree two. Accounting for the overlap and deleting the chosen small configurations leaves at least

$$\frac{24}{25}\frac{n(n-1)}{3} - \frac{2n^2}{1000} = \frac{159}{500}n^2 - \frac{8}{25}n > \frac{3}{10}n^2$$

edges. The last inequality holds for $n \ge 18$. Delete further edges until exactly $\lfloor 3n^2/10 \rfloor$ remain. All the required inequalities survive these deletions. □

We now apply the local inequality to the cones of subgraphs of Q_n. The first lemma controls every maximizing weighting of value greater than $4/9$, uniformly over all internal coefficients in $[0, 4/9]$. This uniformity is important: during the interpolation, both the coefficients and the maximizing weighting may change.

Lemma 4.6. *Let $n \ge \exp(10^8)$, let Q_n be given by Lemma 4.5, and put $K \coloneqq 400e^{2000}(\log n)^2$. Suppose $Q \subseteq Q_n$, $0 \le \beta_i \le 4/9$, and $F_Q(\boldsymbol{\beta}) > 4/9$. At every maximizing weighting in Equation (4.2), the apex has positive mass and*

$$\sum_i y_i > \frac{2}{3}, \qquad \max_i y_i < \frac{K}{n}. \tag{4.15}$$

Proof. Let $p \coloneqq F_Q(\boldsymbol{\beta}) > 4/9$, and let $(a, \boldsymbol{y})$ be a maximizer. Put $S \coloneqq \{i\colon y_i > 0\}$, $N \coloneqq |S|$, and $b \coloneqq \sum_i y_i$. If $N \le 200$, then Equation (4.11) makes $Q[S]$ sparse, and Lemma 4.3 gives $p \le 4/9$. Hence $N > 200$.

If $a = 0$, then $b = 1$. The stationarity equations on the positive bottom support are

$$p = 2 \sum_{jk\colon\, ijk \in Q[S]} y_j y_k + \beta_i y_i^2 \qquad \text{for every } i \in S.$$

Writing $\rho \coloneqq \sum_i y_i^2$ and summing these equations, the codegree bound gives $Np \le 2(1-\rho) + \frac{4}{9}\rho < 2$. Since $p > 4/9$, this forces $N \le 4$, a contradiction. Thus $a > 0$.

Normalize the bottom weights by $z_i \coloneqq y_i/b$ for $i \in S$, and put

$$\rho \coloneqq \sum_{i\in S} z_i^2, \qquad d_i \coloneqq \sum_{jk:\ ijk\in Q[S]} z_j z_k, \qquad h \coloneqq 1 - \frac{2a}{b}.$$

The apex and bottom stationarity equations, respectively, give

$$p = b^2(1-\rho), \tag{4.16}$$

$$(1-h)(1-z_i) + 2d_i + \beta_i z_i^2 = 1-\rho \qquad \text{for every } i \in S. \tag{4.17}$$

Since each pair has codegree at most 2, we have $\sum_i d_i \le 1-\rho$. Summing Equation (4.17), using $1-h = 2(1-b)/b$ and $\sum_i \beta_i z_i^2 \le \frac{4}{9}\rho$, gives

$$2\frac{1-b}{b}(N-1) + 2\sum_i d_i + \sum_i \beta_i z_i^2 = N(1-\rho), \qquad \frac{2(N-1)}{b} \ge 3N - 4 - \left(N - \frac{14}{9}\right)\rho.$$

Together with Equation (4.16) and $p > 4/9$, this implies

$$3N - 4 - \left(N - \frac{14}{9}\right)\rho < 3(N-1)\sqrt{1-\rho} \le 3(N-1)\left(1 - \frac{\rho}{2}\right).$$

Consequently,

$$\left(N + \frac{1}{9}\right)\rho < 2. \tag{4.18}$$

In particular, $1 \le N\rho < 2$ and $\rho < 1/100$.

Each link graph has maximum degree at most 2. Applying $z_j z_k \le (z_j^2 + z_k^2)/2$ to its edges gives $d_i \le \rho - z_i^2$. Multiplying Equation (4.17) by z_i and summing therefore yields

$$h(1-\rho) = 2\sum_i z_i d_i + \sum_i \beta_i z_i^3 \le 2\rho - \frac{14}{9}\sum_i z_i^3 \le 2\rho.$$

Thus

$$\frac{2\rho}{1-\rho} \ge h > 3(1-\sqrt{1-\rho}) \ge \frac{3}{2}\rho. \tag{4.19}$$

The strict lower bound follows from $b = 2/(3-h)$ and Equation (4.16). Rearranging Equation (4.17),

$$(1-h)z_i = 2d_i + \beta_i z_i^2 - h + \rho \le 3\rho - h - (2-\beta_i)z_i^2 < \frac{3}{2}\rho.$$

By Equation (4.19), $1 - h > 97/99 > 3/4$, and hence

$$\max_i z_i < 2\rho. \tag{4.20}$$

We now rule out $200 < N \le n/(100L_n)$ using Equation (4.12). Put $D_* \coloneqq \sum_i d_i$. Summing Equation (4.17) in terms of h gives

$$D_* = \frac{N(h-\rho) + 1 - h - \sum_i \beta_i z_i^2}{2}.$$

The upper bound in Equation (4.19) and $\sum_i \beta_i z_i^2 \le \frac{4}{9}\rho$ give $h + \sum_i \beta_i z_i^2 < 1/25$. The lower bound in Equation (4.19) therefore implies

$$D_* > \frac{N\rho}{4} + \frac{12}{25} > \frac{49}{100}N\rho,$$

where the second inequality uses $N\rho < 2$. By Cauchy–Schwarz and $d_{Q[S]}(j,k) \le 2$,

$$D_*^2 = \left(\sum_{\{j,k\}\in\binom{S}{2}} d_{Q[S]}(j,k) z_j z_k\right)^2 \le \left(\sum_{j<k} d_{Q[S]}(j,k)^2\right)\left(\sum_{j<k} z_j^2 z_k^2\right) \le 3|Q[S]|\rho^2.$$

Thus $|Q[S]| > 2401N^2/30000 > N^2/13$. Since $N > 200$, Equation (4.12) forces $N > n/(100L_n)$. Combining this with Equations (4.18) and (4.20), we get

$$\max_i y_i \le \max_i z_i < 2\rho < \frac{4}{N} < \frac{400L_n}{n} = \frac{K}{n}.$$

Finally, Equation (4.16) and $p > 4/9$ give $b > 2/3$. □

The preceding lemma lets us round real substitution coefficients without losing all weight on the coordinates used for the next recursive step.

Proposition 4.7. $[4/9, 4/9 + e^{-10^9}] \subseteq \mathcal{A}$.

Proof. Fix $p = 4/9 + \varepsilon$ with $0 < \varepsilon \le e^{-10^9}$. We first construct a finite substitution function φ and a number u such that

$$p = \varphi(u), \qquad \frac{4}{9} \le u \le \frac{4}{9} + \frac{\varepsilon}{2}, \tag{4.21}$$

and all fixed coefficients of φ are finite-star Lagrangians. We will then iterate this representation.

Put $n := \lfloor 1/(40\varepsilon) \rfloor$ and $K := 400e^{2000}(\log n)^2$. Then $\log n > 10^8$ and $1/80 \le n\varepsilon \le 1/40$.

At $a = 1/3$ and $y_i = 2/(3n)$, Equation (4.10) gives

$$F_{Q_n}\left(\frac{4}{9}, \dots, \frac{4}{9}\right) \ge \frac{4}{9} + \frac{4}{45n} - \frac{16}{9n^3} \ge \frac{4}{9} + \frac{1}{20n} > p.$$

Delete edges successively until the maximum with all coefficients equal to $4/9$ first becomes at most p. This happens because the edgeless bottom graph has maximum $4/9$, by Lemma 4.3 and a weighting concentrated at one bottom vertex. Let Q be the graph immediately before this crossing. At a maximizing weighting for Q, deleting the final edge changes the polynomial by at most $6K^3/n^3$. Hence

$$0 < F_Q\left(\frac{4}{9}, \dots, \frac{4}{9}\right) - p \le \frac{6K^3}{n^3}. \tag{4.22}$$

Lowering every coefficient of F_Q by the same amount changes its value at a fixed weighting at a rate of $\sum_i y_i^3$. At every maximizing weighting covered by Lemma 4.6, Jensen's inequality and Equation (4.15) give

$$\frac{8}{27n^2} < \sum_i y_i^3 \le \frac{K^2}{n^2}. \tag{4.23}$$

By Equation (2.3), these bounds also control the change in the maximum whenever both optimized values exceed $4/9$. Also $n > (729/4)K^3$. Indeed, for $T = \log n \ge 10^8$, we have $\log((729/4)K^3) < 6100 + 6\log T < T$; the last inequality holds at 10^8 and the derivative of $6100 + 6\log T$ is $6/T < 1$.

By continuity of $s \mapsto F_Q(s, \dots, s)$ and Equations (4.22) and (4.23), there is a coefficient $\beta_* = 4/9 - \delta$ with

$$F_Q(\beta_*, \dots, \beta_*) = p, \qquad 0 < \delta \le \frac{81K^3}{4n} < \frac{1}{9}. \tag{4.24}$$

Indeed, if lowering the common coefficient by $81K^3/(4n)$ left the maximum above p, comparison at a maximizer of the lowered polynomial would force a decrease greater than $6K^3/n^3$, a contradiction.

Choose a maximizing weighting $\boldsymbol{x}$ for $F_Q(\beta_*, \dots, \beta_*)$. More than $n/(3K-1)$ bottom coordinates satisfy $x_i \ge 1/(3n)$. Otherwise, if their number were $m \le n/(3K-1)$, then

$$\sum_i x_i \le \frac{mK}{n} + \frac{n-m}{3n} \le \frac{2}{3},$$

contrary to Lemma 4.6. Put $k := \lceil 10^{10}K^{10} \rceil$. To select k of these coordinates, we verify the stronger bound

$$8kK^3 < n. \tag{4.25}$$

With $T = \log n \ge 10^8$ and $K \ge 1$,

$$\log(8kK^3) \le \log(16 \cdot 10^{10}K^{13}) < 30 + 13\log 400 + 13(2000 + 2\log T) < T.$$

The last inequality holds at $T = 10^8$, where its left side is less than 26602, and its derivative is $26/T < 1$. Thus $k < n/(3K-1)$. Choose k coordinates with $x_i \geq 1/(3n)$ and let R be their index set. Then

$$\sum_{i \in R} x_i^3 \geq \frac{k}{27n^3}.$$

Set $C \coloneqq [n] \setminus R$ and let $g(s)$ be the value of F_Q with coefficient $4/9$ on R and coefficient s on C. Varying s changes the polynomial at a rate of $\sum_{i \in C} y_i^3$. Whenever $0 \leq s \leq 4/9$ and $g(s) > 4/9$, Lemma 4.6 and Equation (4.25) give $\sum_{i \in R} y_i < kK/n < 1/6$. The total bottom weight exceeds $2/3$, so the weight on C exceeds $1/2$. Jensen's inequality and the coordinate bound therefore give

$$\frac{1}{8n^2} < \sum_{i \in C} y_i^3 \leq \frac{K^2}{n^2}. \tag{4.26}$$

For $0 \leq s < t \leq 4/9$ with $g(s) > 4/9$, Equation (2.3) therefore gives

$$\frac{t-s}{8n^2} < g(t) - g(s) \leq \frac{K^2}{n^2}(t-s).$$

Evaluating first at $\boldsymbol{x}$, and then at a maximizer for $g(\beta_*)$, gives

$$\frac{\delta k}{27n^3} \leq g(\beta_*) - p \leq \frac{\delta k K^3}{n^3}. \tag{4.27}$$

The upper bound uses $g(\beta_*) > p > 4/9$, so Lemma 4.6 applies at that maximizer. If $g(\beta_* - 8\delta k K^3/n) > p$, the lower bound for the increment of g gives

$$g(\beta_*) - p > g(\beta_*) - g\left(\beta_* - \frac{8\delta k K^3}{n}\right) > \frac{\delta k K^3}{n^3},$$

contrary to Equation (4.27). By continuity, there is $s_* \coloneqq \beta_* - t$ with $g(s_*) = p$ and $0 < t \leq 8\delta k K^3/n$. The upper bound $g(\beta_*) - p \leq K^2 t/n^2$, together with Equations (4.25) and (4.27), gives

$$\frac{\delta k}{27K^2 n} \leq t \leq \frac{8\delta k K^3}{n} \leq \delta. \tag{4.28}$$

All coefficients remain at least $4/9 - 2\delta > 2/9$. In particular,

$$\delta \leq \frac{4}{9} - s_* = \delta + t \leq 2\delta.$$

Choose m with $\beta_m \leq s_* < \beta_{m+1}$, where β_m is given by Equation (4.1). Since $4/(9(m+1)) < 4/9 - s_* \leq 4/(9m)$ and $4/(9m) \leq 2(4/9 - s_*)$, we have

$$0 \leq s_* - \beta_m \leq \frac{4}{9m(m+1)} \leq 9\left(\frac{4}{9} - s_*\right)^2 \leq 36\delta^2.$$

Let $v \coloneqq g(\beta_m)$. Evaluate a maximizer for $g(s_*) = p$ at the rounded-down coefficient. The bound $n > 8kK^3 \geq 8 \cdot 10^{10} K^{13}$ implies $n^3 > 2361960 K^8$. Together with Equations (4.24) and (4.26) and $n\varepsilon \geq 1/80$, this gives

$$0 \leq p - v \leq \frac{36K^2\delta^2}{n^2} \leq \frac{59049}{4}\frac{K^8}{n^4} < \frac{1}{160n} \leq \frac{\varepsilon}{2}. \tag{4.29}$$

Thus $v > 4/9$.

Take a maximizing weighting $\boldsymbol{y}$ for v. Evaluate the same weighting with every coefficient equal to β_*. Since that maximum equals p,

$$v \leq p + \delta \sum_{i \in R} y_i^3 - (\beta_* - \beta_m) \sum_{i \in C} y_i^3.$$

We have $\beta_* - \beta_m \geq t$ and $k \geq 10^{10} K^{10} \geq 314928 K^7$. Thus Equations (4.26), (4.28) and (4.29) imply

$$\sum_{i \in R} y_i^3 \geq \frac{t - 8n^2(p-v)}{8\delta n^2} \geq \frac{1}{n^3}\left(\frac{k}{216K^2} - 729K^5\right) \geq \frac{k}{432K^2 n^3} > 0. \tag{4.30}$$

Thus the coordinates selected before rounding retain positive total cubic weight at every maximizer after rounding.

Finally, let $\varphi(u)$ be the substitution function with coefficient u on R and coefficient β_m on C. We have $\varphi(4/9) = v \le p$. Since $n\varepsilon \ge 1/80$ and $12754584 \cdot 80 < 10^{10}$, the choice of k gives $k \ge 12754584K^{10}/(n\varepsilon)$. By Equations (4.29) and (4.30), increasing the coefficient on R by $\varepsilon/2$ therefore recovers at least the rounding loss $p - v$. Evaluating at $\boldsymbol{y}$ gives

$$\varphi\left(\frac{4}{9} + \frac{\varepsilon}{2}\right) \ge v + \frac{\varepsilon}{2}\sum_{i\in R} y_i^3 \ge p.$$

Continuity proves Equation (4.21).

Starting with $p_0 = p$, repeat Equation (4.21) while the new parameter exceeds $4/9$. This gives functions φ_j and numbers p_j satisfying

$$p_j = \varphi_j(p_{j+1}), \qquad 0 \le p_j - \frac{4}{9} \le 2^{-j}\varepsilon.$$

Every φ_j is nondecreasing and 1-Lipschitz, by Lemma 4.2. At depth j, replace the remaining parameter by the finite-star value $\beta_M = \frac{4}{9}(1 - 1/M)$. All coefficients then arise from finite 3-graphs, and repeated substitution gives a finite 3-graph $H_{j,M}$ with Lagrangian

$$\lambda(H_{j,M}) = (\varphi_0 \circ \cdots \circ \varphi_{j-1})(\beta_M).$$

Monotonicity and the Lipschitz bounds yield

$$0 \le p_0 - \lambda(H_{j,M}) \le p_j - \beta_M \le 2^{-j}\varepsilon + \frac{4}{9M}.$$

If some $p_j = 4/9$, stop there and let $M \to \infty$; otherwise let both j and M tend to infinity. This proves $p_0 \in \mathcal{A}$. The endpoint $4/9$ already belongs to $\mathcal{A}$ by Equation (4.1). □

4.3 The quantitative matching estimate

Proof of Lemma 4.4. For completeness, we give the quantitative random selection argument. We repeatedly select isolated activated edges, controlling both the remaining degrees and the number of vertices removed without being matched.

We first record the concentration estimate used in each round. If X is a function of independent Bernoulli variables with success probabilities p_j, changing coordinate j changes X by at most c_j, and $V = \sum_j p_j(1-p_j)c_j^2$, $C = \max_j c_j$, then

$$\mathbb{P}\big[|X - \mathbb{E}[X]| \ge t\big] \le 2\exp\left(-\frac{3t^2}{2(3V + Ct)}\right). \tag{4.31}$$

Each one-sided estimate omits the factor 2. If $V = 0$, then X is almost surely constant. For $V > 0$, the Doob difference has the form $(B_j - p_j)A_j$ with $|A_j| \le c_j$. For $0 < \lambda < 3/C$, the conditional exponential moment is at most $\exp(\lambda^2 p_j(1-p_j)c_j^2/[2(1-\lambda C/3)])$, by the exponential series and $k! \ge 2\cdot 3^{k-2}$. Iterating conditional expectations and applying Markov's inequality with $\lambda = t/(V + Ct/3)$ gives the upper-tail bound. Applying the same argument to $-X$ gives the lower-tail bound.

Fix $\eta := 10^{-6}$, $\tau := e^{-100}$, and $T := 5000000$. At round i, let the current induced hypergraph H_i have N_i vertices and degrees in $[(1-\delta_i)D_i, (1+\delta_i)D_i]$. Initially, $D_0 = D$, $\delta_0 = \tau$, and $N_0 = N$. Activate every current edge independently with probability $p = \eta/D_i$. Add to the matching every activated edge that meets no other activated edge. Remove every vertex incident with an activated edge. A removed vertex not covered by one of the chosen matching edges is called *wasted*. The next induced hypergraph is still linear.

Define the deterministic parameters by

$$D_{i+1} := D_i\left(1 - \frac{\eta}{D_i}\right)^{2D_i - 2}, \qquad \delta_{i+1} := (1 + 10\eta)\delta_i + 3\tau. \tag{4.32}$$

If no vertices remain, stop. We first derive the one-round estimates assuming $D_i \geq 1$ and $\delta_i \leq 1/1000$, and then verify that these assumptions hold throughout.

Fix a current vertex v. Ignore the activation variables for edges containing v, and let X_v count the edges $\{v, x, y\}$ for which no activated edge not containing v meets x or y. If v survives the actual round, its remaining degree is exactly X_v. The ignored variables are independent of X_v.

Linearity gives the exact expectation

$$\mathbb{E}[X_v] = \sum_{\{v,x,y\} \in H_i} (1-p)^{d_{H_i}(x)+d_{H_i}(y)-2}. \tag{4.33}$$

For a fixed edge $\{v, x, y\}$, the non-v edges through x and y are disjoint collections of sizes $d_{H_i}(x) - 1$ and $d_{H_i}(y) - 1$. An edge in both would share the pair $\{x, y\}$ with $\{v, x, y\}$, violating linearity.

The edges through v have $2d_{H_i}(v)$ distinct other vertices, so X_v depends on at most $2d_{H_i}(v)\Delta(H_i)$ non-v activation variables. Changing one such variable affects at most three terms of X_v, since each vertex of the activated edge lies in at most one edge through v. Its variance parameter in (4.31) is at most

$$9p \cdot 2d_{H_i}(v)\Delta(H_i) \leq 18\eta(1+\delta_i)^2 D_i < 20\eta D_i.$$

Taking $t = \tau D_i$ and $C = 3$, and using $\tau < \eta$, gives

$$\mathbb{P}\big[|X_v - \mathbb{E}[X_v]| > \tau D_i\big] \leq 2\exp\left(-\frac{\tau^2 D_i}{42\eta}\right). \tag{4.34}$$

Since $p = \eta/D_i$ and $D_i \geq 1$,

$$-\log(1-p) \leq \frac{\eta}{D_i - \eta} \leq \frac{2\eta}{D_i}.$$

Combining Equations (4.32) and (4.33),

$$(1-\delta_i)(1-p)^{2\delta_i D_i} \leq \frac{\mathbb{E}[X_v]}{D_{i+1}} \leq (1+\delta_i)(1-p)^{-2\delta_i D_i}. \tag{4.35}$$

The lower bound is at least $1 - (1+4\eta)\delta_i$, by $e^{-x} \geq 1 - x$. For the upper bound, $e^{4\eta\delta_i} \leq 1 + 5\eta\delta_i$ and $\delta_i \leq 1/1000$ give at most $1 + (1+6\eta)\delta_i$. Also,

$$-\log\frac{D_{i+1}}{D_i} = (2D_i - 2)(-\log(1-p)) \leq \frac{2\eta D_i}{D_i - \eta} \leq \frac{2\eta}{1-\eta} < 3\eta. \tag{4.36}$$

Thus $D_{i+1} \geq D_i e^{-3\eta} > D_i/2$. A deviation of at most τD_i adds at most 2τ to the relative error. Together with Equation (4.35), this gives the error bound δ_{i+1} at every surviving vertex.

Let U be the number of vertices incident with no activated edge. Then

$$\mathbb{E}[U] = \sum_v (1-p)^{d_{H_i}(v)} \leq N_i e^{-\eta(1-\delta_i)} \leq (1 - 0.998\eta)N_i.$$

The last inequality follows from $e^{-x} \leq 1 - x + x^2/2$, $1 - \delta_i \geq 0.999$, and $\eta = 10^{-6}$. Changing one activation changes U by at most three. Since $|H_i| \leq (1+\delta_i)D_i N_i/3$, its variance parameter is at most $3\eta(1+\delta_i)N_i < 4\eta N_i$. Applying the upper-tail bound with deviation $0.008\eta N_i$ gives

$$\mathbb{P}[U > (1 - 0.99\eta)N_i] \leq \exp(-7 \cdot 10^{-6}\eta N_i) = \exp(-7 \cdot 10^{-12} N_i). \tag{4.37}$$

Here $0.008^2/[2(4 + 0.008)] > 7 \cdot 10^{-6}$.

Let J count unordered intersecting pairs of activated edges. Linearity gives

$$\mathbb{E}[J] = p^2 \sum_v \binom{d_{H_i}(v)}{2} \leq \frac{1}{2}\eta^2(1+\delta_i)^2 N_i.$$

At most $2J$ activated edges are nonisolated, so at most $6J$ vertices are wasted. If W denotes their number, then $\mathbb{E}W < 4\eta^2 N_i$, and Markov's inequality gives

$$\mathbb{P}[W > 400\eta^2 N_i] < \frac{1}{100}. \tag{4.38}$$

We verify that $D_i \geq 1$ and $\delta_i \leq 1/1000$ for every round. By Equation (4.36), for $0 \leq i \leq T$,

$$D_i \geq De^{-3\eta i} \geq e^{985}\log(2N) > 1.$$

This also verifies inductively the condition used to derive that decay estimate. Solving the other recurrence in Equation (4.32),

$$\delta_i = \tau(1+10\eta)^i + \frac{3\tau}{10\eta}\big((1+10\eta)^i - 1\big) \leq \left(1 + \frac{3}{10\eta}\right)\tau e^{10\eta i} \leq 300001e^{-50} < \frac{1}{1000}.$$

Thus all one-round estimates hold through round T.

The probability in Equation (4.34) is at most $2(2N)^{-100}$ for each current vertex. Summing over at most N vertices gives total degree failure probability at most $(2N)^{-99}$. If H_i is nonempty, linearity implies $N_i - 1 \geq 2(1-\delta_i)D_i$, and in particular $N_i \geq D_i$. The survivor failure probability in Equation (4.37) is then at most $(2N)^{-100}$. Together with Equation (4.38), these probabilities sum to less than one. In each round, we can therefore choose an outcome with the required new degree bounds and

$$N_{i+1} \leq (1 - 0.99\eta)N_i, \qquad W_i \leq 400\eta^2 N_i.$$

The selected edges from different rounds are disjoint because their vertices were removed when selected. After T rounds, $N_T \leq e^{-0.99\eta T}N = e^{-4.95}N$, and the total number of wasted vertices is at most

$$400\eta^2 \sum_{i<T} N_i \leq \frac{400\eta}{0.99}N.$$

Every uncovered vertex is either a final survivor or wasted. The uncovered proportion is therefore at most

$$e^{-4.95} + \frac{400 \cdot 10^{-6}}{0.99} < 0.01.$$

For a rational check, the Taylor sum $\sum_{0\leq j\leq 7}(99/20)^j/j!$ is greater than 110, so the last expression is less than $1/110 + 0.000405 < 0.01$. This completes the proof. □

4.4 Rounding and star substitutions

The interval in Proposition 4.7 is very short. To enlarge it, we use fixed finite templates whose maximizing weights remain comparable as their internal coefficients vary. A coefficient can then be rounded to a finite-star Lagrangian, and the resulting change in the maximum can be corrected by adjusting another coefficient.

For the interval extensions, put

$$\kappa := \left(\frac{10}{9}\right)^3 = \frac{1000}{729}, \qquad \beta_m = \frac{4}{9}\left(1 - \frac{1}{m}\right).$$

We first quantify the error in this rounding.

Lemma 4.8. *For $x \in (0,1]$, let $E_+(x) := x - 1/\lceil 1/x\rceil$, and put $E_+(0) := 0$. If $0 < B \leq 1$ and $M := \lceil 1/B\rceil$, then*

$$\sup_{0\leq x\leq B} E_+(x) = \max\left\{B - \frac{1}{M}, \frac{1}{M(M+1)}\right\}. \tag{4.39}$$

For $0 < x < 1$, rounding $\frac{4}{9}(1-x)$ upward or downward to the adjacent coefficient β_m incurs an error of at most $4x^2/(9(1-x))$.

Proof. Upward rounding uses $m = \lceil 1/x\rceil$ and has normalized error $E_+(x)$. On $[1/M, B]$, its maximum is $B - 1/M$. On $[1/m, 1/(m-1))$ with $m \geq M+1$, its supremum is $1/(m(m-1))$. These values decrease with m, proving Equation (4.39), including reciprocal endpoints.

Downward rounding uses $m = \lfloor 1/x\rfloor$. Since $m > 1/x - 1$, its normalized error $1/m - x$ is less than $x^2/(1-x)$. For upward rounding, $\lceil 1/x\rceil < 1/x + 1$ gives $E_+(x) < x^2/(1+x) \leq x^2/(1-x)$. At exact coefficients, the error is zero. □

Fix a target maximum $p > 4/9$. At each stage, one coefficient is chosen so that the substitution template has maximum p. We round this coefficient and adjust another to restore the same maximum. The *control coordinates* are used successively for these adjustments. One additional *reserved coordinate* will carry the final parameter.

Lemma 4.9. *Consider a finite substitution template with coefficients in $[0, 4/9]$. Suppose that, at every global maximum greater than $4/9$, the masses of any two designated control or reserved coordinates satisfy $y_i^3 \le \kappa y_j^3$. Suppose a control coefficient $4/9 - d$ gives maximum $p > 4/9$, while an unused control has coefficient $4/9$. If rounding the first coefficient upward increases it by Δ and $\kappa\Delta \le 4/9$, then lowering the unused control by some $d' \in [0, \kappa\Delta]$ restores the maximum to p.*

Proof. After the upward rounding, lower the unused coefficient by $\kappa\Delta$. If the modified maximum were greater than p, choose a maximizing weighting for that modified template. All coefficients are still in $[0, 4/9]$, so the ratio assumption applies there. The change from the template before rounding at this weighting is at most

$$\Delta y_{\text{old}}^3 - \kappa\Delta y_{\text{next}}^3 \le 0,$$

a contradiction. With no decrease in the unused coefficient, the maximum is at least p. Continuity gives the claimed d'. □

If the current coefficient equals 0, $4/9$, or some β_m, we retain it exactly and stop.

Lemma 4.10. *Assume the ratio bound in Lemma 4.9. Fix a target $p > 4/9$, and suppose there are q control coordinates, including the coordinate whose coefficient is chosen to give maximum p, and one reserved coordinate with coefficient $4/9$. Every other control initially has coefficient $4/9$, and all remaining fixed coefficients belong to $\{0, 4/9\} \cup \{\beta_m : m \ge 1\}$. Let*

$$(q, \varepsilon_{\min}) \in \{(7, 10^{-15}), (34, e^{-10^9})\}.$$

If $p = 4/9 + \varepsilon$ with $\varepsilon \ge \varepsilon_{\min}$, there is a finite substitution function φ with only the reserved coefficient variable, and all other coefficients in the same set, such that

$$\varphi\left(\frac{4}{9}\right) \le p \le \varphi\left(\frac{4}{9} + \frac{\varepsilon}{2}\right). \tag{4.40}$$

Consequently, $p = \varphi(u)$ for some $u \in [4/9, 4/9 + \varepsilon/2]$, and $u \in \mathcal{A}$ implies $p \in \mathcal{A}$.

Proof. Put $a := 4/9$ and write $a - d_j$ for the coefficient at the jth control that makes the maximum equal to p. If it is an allowed coefficient, retain it and stop. Otherwise round it upward in the first $q-1$ controls, compensating at the next control by Lemma 4.9. Let Δ_j be the rounding error and put $x_j := d_j/a$. By Lemma 4.8,

$$x_{j+1} \le \kappa E_+(x_j), \qquad E_+(x) \le \frac{x^2}{1+x}.$$

Since $x_1 \le 1$ and $E_+(x) \le 1/2$, we have $x_2 \le \kappa/2 < 11/16$. On $[0, 11/16]$, the sharper bound $E_+(x) \le 3/16$ follows from Equation (4.39) with $B = 11/16$ and $M = 2$. Thus $x_3 < 33/128$. With $u_j := \kappa x_j$, the next step gives

$$u_4 \le \kappa^2 \frac{x_3^2}{1+x_3} < \left(\frac{11}{8}\right)^2 \frac{33^2}{128 \cdot 161} = \frac{131769}{1318912} < \frac{1}{10}.$$

Every subsequent upward rounding satisfies $u_{j+1} \le u_j^2$, and hence $u_q < 10^{-2^{q-4}}$. All compensations are valid: the first has normalized size at most $\kappa/2 < 1$, and the subsequent bounds on x_j decrease.

Round the coefficient at the last control downward. Its error satisfies

$$2\kappa\Delta_q \le 2\kappa a \frac{x_q^2}{1-x_q} = \frac{2a}{\kappa - u_q} u_q^2 < \frac{6480}{9271} u_q^2 < 10^{-2^{q-3}}.$$

For $q = 7$ this is less than $10^{-16} < 10^{-15}$; for $q = 34$ it is less than $10^{-2^{31}} < e^{-10^9}$. In either case, $2\kappa\Delta_q < \varepsilon$.

Downward rounding gives $\varphi(4/9) \le p$. At a maximizer of the template before this last rounding, lowering its control by Δ_q and increasing the reserved coefficient by $\varepsilon/2$ changes the objective by at least

$$-\Delta_q y_{\mathrm{last}}^3 + \frac{\varepsilon}{2} y_{\mathrm{res}}^3 \ge \left(\frac{\varepsilon}{2} - \kappa\Delta_q\right) y_{\mathrm{res}}^3 \ge 0.$$

The ratio bound applies at this weighting because the template before rounding has all coefficients at most $4/9$. This proves Equation (4.40). If an exact coefficient is reached earlier, use the corresponding constant substitution function.

Finally, continuity yields $\varphi(u) = p$ for some $u \in [4/9, 4/9 + \varepsilon/2]$. All fixed coefficients belong to $\mathcal{A}$, by Equation (4.1), so the last assertion follows from Lemma 4.2. □

The star and book substitutions use the following bounds on leaf weights and stationary values.

Lemma 4.11. *Suppose that $D \ge 3$, $A, z > 0$, $0 \le \beta_i \le B$, and $3Bz < 2A$. Every maximizer of*

$$-A\sum_{i=1}^{D} y_i^2 + \sum_{i=1}^{D} \beta_i y_i^3 \qquad \text{over } y_i \ge 0 \text{ with } \sum_{i=1}^{D} y_i = z$$

has all $y_i > 0$. Writing $Y = \max_i y_i$ and $y = \min_i y_i$, it satisfies

$$Y < \frac{z}{D-1}, \qquad \frac{y}{Y} \ge 1 - \frac{3Bz}{2A(D-1)}, \qquad y \ge \frac{z}{D}\left(1 - \frac{3Bz}{2A(D-1)}\right).$$

Proof. At every positive coordinate the derivative $-2Ay_i + 3\beta_i y_i^2$ is negative. Moving a small amount of weight from such a coordinate to a zero coordinate would increase the value, so no coordinate vanishes. The stationary equations give $2Ay_i - 3\beta_i y_i^2$ independent of i. Put $K = 3B/(2A)$, so $Kz < 1$. Comparing a largest and a smallest coordinate yields $y \ge Y(1 - KY)$ and therefore $z \ge DY - (D-1)KY^2$.

If $Y \ge z/(D-1)$, the concave quadratic $Dt - (D-1)Kt^2$ has values greater than z at both endpoints of $[z/(D-1), z]$. The excesses are $z(1-Kz)/(D-1)$ and $(D-1)z(1-Kz)$, respectively. This contradicts the preceding inequality. Thus $Y < z/(D-1)$, and the other bounds follow from $y \ge Y(1-KY)$ and $Y \ge z/D$. □

Lemma 4.12. *Let f be a polynomial of degree at most 3 on a Euclidean space. Suppose that $h > 0$ and $\nabla^2 f(\boldsymbol{z}) \preceq -hI$. Every stationary point $\boldsymbol{x}$ satisfies*

$$f(\boldsymbol{x}) \le f(\boldsymbol{z}) + \frac{2}{3h}\|\nabla f(\boldsymbol{z})\|^2.$$

For a simplex, the derivatives and norms are taken in its affine span.

Proof. Put $\boldsymbol{w} = \boldsymbol{x} - \boldsymbol{z}$ and $\boldsymbol{g} = \nabla f(\boldsymbol{z})$. The exact cubic expansion has the form

$$f(\boldsymbol{z} + \boldsymbol{w}) = f(\boldsymbol{z}) + \boldsymbol{g}\cdot\boldsymbol{w} + \frac{1}{2}\boldsymbol{w}^{\mathsf{T}}\nabla^2 f(\boldsymbol{z})\boldsymbol{w} + T(\boldsymbol{w}),$$

where T is homogeneous of degree 3. Stationarity at $\boldsymbol{x}$ gives $\boldsymbol{g}\cdot\boldsymbol{w} + \boldsymbol{w}^{\mathsf{T}}\nabla^2 f(\boldsymbol{z})\boldsymbol{w} + 3T(\boldsymbol{w}) = 0$. Eliminating $T(\boldsymbol{w})$ yields

$$f(\boldsymbol{x}) - f(\boldsymbol{z}) = \frac{2}{3}\boldsymbol{g}\cdot\boldsymbol{w} + \frac{1}{6}\boldsymbol{w}^{\mathsf{T}}\nabla^2 f(\boldsymbol{z})\boldsymbol{w} \le \frac{2}{3}\|\boldsymbol{g}\|\|\boldsymbol{w}\| - \frac{h}{6}\|\boldsymbol{w}\|^2 \le \frac{2}{3h}\|\boldsymbol{g}\|^2.$$ □

The same identity with $\boldsymbol{g} = 0$ shows that a stationary reference point with nonpositive Hessian has value at least that of every other stationary point. Negative definiteness gives strict inequality away from the reference point.

We apply the rounding procedure to a complete 3-star with 37 leaves. Substitute a 3-graph of density $8/27$ at the apex, and give the leaves coefficients $\beta_i \in [0, 4/9]$. Its substitution polynomial is

$$P(a, \boldsymbol{y}) \coloneqq \frac{8}{27}a^3 + 3a\left[b^2 - \sum_{i\in[37]} y_i^2\right] + \sum_{i\in[37]} \beta_i y_i^3, \qquad b \coloneqq \sum_i y_i = 1 - a. \tag{4.41}$$

The three types of edges are illustrated in Figure 10.

The next lemma supplies the uniform ratio bound needed in Lemma 4.10.

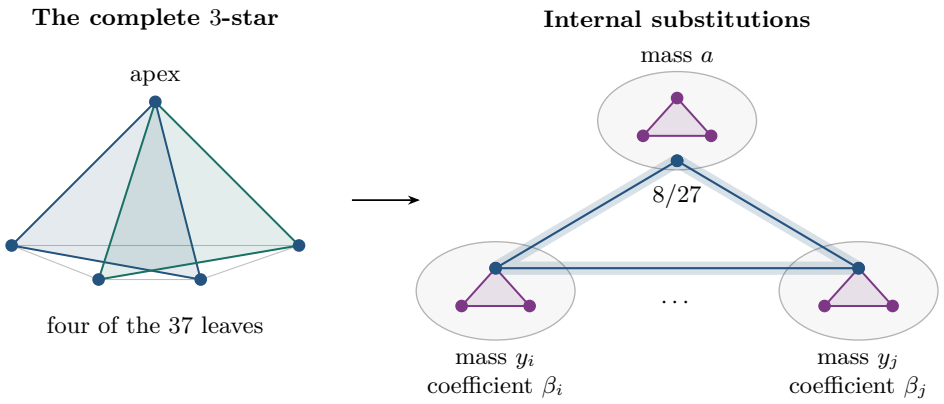


Figure 10: The 37-leaf star substitution in Equation (4.41). Every pair of leaves forms an edge with the apex; two such edges are highlighted on the left. The right panel shows internal edges and one transversal edge after substitution. The coefficients at the apex and leaves are $8/27$ and β_i, respectively.

Lemma 4.13. *At every global maximizer of Equation (4.41) of value greater than $4/9$, all leaves have positive weight and*

$$\frac{1}{4} < a < \frac{1}{2}, \qquad y_i > \ell_* := \frac{73}{5550}, \qquad \frac{y_i^3}{y_j^3} < \left(\frac{75}{73}\right)^3 < \kappa.$$

Proof. For upper bounds, first increase all leaf coefficients to $4/9$. At fixed a, b, choose a maximizer of $-3a\sum_i y_i^2 + \frac{4}{9}\sum_i y_i^3$ with minimum positive support. For two positive coordinates, stationarity gives $(y_i - y_j)(-6a + \frac{4}{3}(y_i + y_j)) = 0$. If $y_i \neq y_j$, the coefficient $6a - \frac{4}{3}(y_i + y_j)$ of $y_i y_j$ along their fixed-sum line is therefore zero. The polynomial is constant on that line, and merging the two coordinates preserves the value while reducing the support. Thus the positive coordinates are equal.

With k equal positive coordinates, the leaf contribution is $-3ab^2/k + 4b^3/(9k^2)$, a convex function of $1/k$. It is bounded by its values at $k = 1$ and $k = 37$. The first branch is $\frac{8}{27}a^3 + \frac{4}{9}b^3 \leq 4/9$. The second is

$$f(a) := \frac{8}{27}a^3 + \frac{108}{37}a(1-a)^2 + \frac{4}{9 \cdot 37^2}(1-a)^3. \tag{4.42}$$

For $a \leq 1/4$, monotonicity of $3a(1-a)^2$ gives $f(a) \leq 1/216 + 27/64 + 4/(9 \cdot 37^2) < 4/9$. For $a \geq 1/2$, use the identity

$$\frac{89}{216} - \frac{8}{27}a^3 - 3a(1-a)^2 = \frac{(2a-1)(-356a^2 + 470a - 89)}{216} \geq 0.$$

The quadratic factor is concave and positive at both endpoints of $[1/2, 1]$. Hence $f(a) \leq 89/216 + 4/(9 \cdot 37^2) < 4/9$ on that interval. Increasing coefficients only raises the objective, so every original maximizer of value greater than $4/9$ has $1/4 < a < 1/2$ and $1/2 < b < 3/4$.

At fixed a, b, apply Lemma 4.11 with $A = 3a$, $B = 4/9$, $D = 37$, and $z = b$. Since $2b/(9a) < 2/3$, every leaf is positive and

$$\frac{\min_i y_i}{\max_i y_i} > \frac{53}{54} > \frac{73}{75}, \qquad \min_i y_i > \frac{53}{3996} > \frac{73}{5550}.$$

These inequalities imply the stated bounds. □

Let A_* be the maximum when all leaf coefficients equal $4/9$. We will use the explicit bounds

$$\frac{4}{9} + 2 \cdot 10^{-6} < A_* < \frac{4}{9} + 2.5 \cdot 10^{-6}, \qquad A_* - \frac{4}{9} < \frac{4}{3}\ell_*^3. \tag{4.43}$$

At $a_0 := 3527805417/10^{10}$, differentiating Equation (4.42) gives

$$-f''(a_0) = \frac{6248829722647}{1283437500000} > 4.$$

Since $f(a_0) > 4/9$, Lemma 4.13 excludes the other branches and places every relevant maximum at an interior stationary point of f. Thus Lemma 4.12 gives $f(a_0) \le A_* \le f(a_0) + f'(a_0)^2/4$. Direct rational evaluation gives

$$f(a_0) - \frac{4}{9} = \frac{208915215870873322169027313}{855625000000000000000000000000000},$$
$$f(a_0) + \frac{f'(a_0)^2}{4} - \frac{4}{9} = \frac{643511093686260762320132402195448108240 1}{263553890625000000000000000000000000000000000000}.$$

These fractions verify Equation (4.43); for the last comparison, $\frac{4}{3}\ell_*^3 = 389017/128215406250$. The program `parts/ordinary/certificates/verify_star37.py` checks each inequality using exact arithmetic.

Proposition 4.14. *If $[4/9, 4/9 + e^{-10^9}] \subseteq \mathcal{A}$, then $[4/9, A_*] \subseteq \mathcal{A}$. In particular, $[4/9, 4/9 + 2 \cdot 10^{-6}] \subseteq \mathcal{A}$.*

Proof. Fix $p = 4/9 + \varepsilon \in (4/9, A_*]$ with $\varepsilon \ge e^{-10^9}$. Reserve one leaf and initially give every leaf coefficient $4/9$. Lower the other leaf coefficients toward zero successively until the maximum first equals p. If three coefficients were completely lowered without reaching p, comparison at a maximizer of the modified template and Lemma 4.13 would force a decrease greater than $\frac{4}{3}\ell_*^3 > A_* - 4/9$, which is impossible. Thus the crossing occurs by the third leaf.

At most two leaf coefficients have been fixed at zero. Excluding the reserved leaf, at least 34 controls remain, including the coordinate whose coefficient is chosen to give maximum p. Apply Lemma 4.10 with $q = 34$. Its hypotheses hold by Lemma 4.13, and all fixed coefficients, including the apex coefficient $8/27$, belong to $\mathcal{A}$. We obtain $p = \varphi(u)$ with $4/9 \le u \le 4/9 + \varepsilon/2$.

Repeat this representation whenever the new target exceeds $4/9 + e^{-10^9}$. Its excess over $4/9$ is at least halved at each step, so after finitely many steps it belongs to the assumed initial interval. Applying Lemma 4.2 in reverse order proves $p \in \mathcal{A}$. The numerical consequence follows from Equation (4.43). □

4.5 Book substitutions

For an integer $D \ge 10$, let a *book* be the 3-graph with two common vertices and D further vertices, with one edge consisting of the two common vertices and each further vertex. We use its cone. Denote the apex weight by a, the common-vertex weights by u, v, and the leaf weights by $y_1, \dots, y_D$. Put

$$b := 1 - a, \qquad s := u + v, \qquad z := \sum_{i \in [D]} y_i, \qquad a + s + z = 1.$$

Give the common vertices coefficients $q_1, q_2 \in [0, 4/9]$ and the leaves coefficients $\beta_i \in [0, 4/9]$. The substitution polynomial is

$$P := 3a\left[b^2 - u^2 - v^2 - \sum_i y_i^2\right] + 6uvz + q_1 u^3 + q_2 v^3 + \sum_i \beta_i y_i^3. \tag{4.44}$$

Figure 11 shows the book edges and the additional edges through the apex.

We first control all its maximizing weightings, before specializing the coefficients used in the finite computation.

Lemma 4.15. *At every global maximizer of Equation (4.44) of value $p > 4/9$, all coordinates are positive and*

$$\frac{1}{5} < a < \frac{1}{3}, \qquad s < \frac{2}{3}, \qquad z > \frac{1}{3}, \qquad |u - v| < \frac{3s}{20}, \qquad y_i > \frac{3}{10D}, \qquad \frac{y_i^3}{y_j^3} < \kappa.$$

Proof. If $a = 0$, increase all internal coefficients to $4/9$ and concentrate the leaf cubic sum in one leaf. The resulting value is at most $6uvz + \frac{4}{9}(u^3 + v^3 + z^3) \le 4/9$, since $u + v + z = 1$ and $24uvz + u^3 + v^3 + z^3 \le 1$. If a common vertex has weight zero, or at most one leaf has positive weight, the bottom graph on the support is empty or a single edge. In either case, Lemma 4.3 bounds the value by $4/9$. These cases are therefore excluded.

To show $a > 1/5$, increase all internal coefficients to $4/9$. At fixed a, u, v, z, the minimum-support argument from Lemma 4.13 reduces the leaf optimization to one positive leaf or D equal positive leaves.

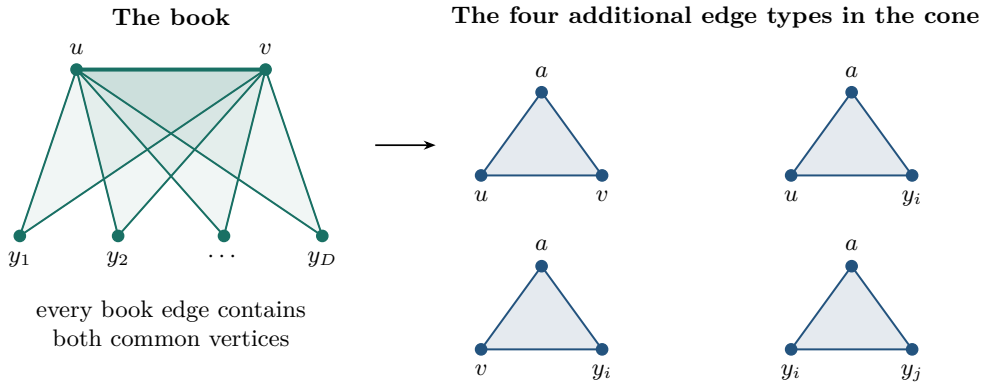


Figure 11: The book and its cone in Equation (4.44); labels denote vertex masses. In addition to every book edge, the cone contains the apex together with every pair of distinct bottom vertices, giving the four types on the right, with $i \neq j$ in the last type. The internal substitutions contribute q_1u^3, q_2v^3, and $\beta_i y_i^3$.

The first case has value at most $4/9$. In the second case, at fixed a, z, s, the coefficient of uv is $6(a+z)-\frac{4}{3}s$. The maximum is therefore attained at equal common-vertex weights or when one of them is zero; the latter case again has value at most $4/9$.

Write $u = v = x$. The resulting expression is convex in $1/D$, so it is bounded by the larger of the $D = 1$ value and the limiting value as $D \to \infty$. The latter expression is

$$3a(1-a)^2 + 6(1-2a)x^2 - \frac{100}{9}x^3.$$

For $0 \le a \le 1/5$, its maximum over $x \ge 0$ occurs at $x = 9(1-2a)/25$, which is feasible. Its value is

$$g(a) := 3a(1-a)^2 + \frac{162}{625}(1-2a)^3.$$

Since $625g'(a) = 903 - 3612a + 1737a^2 > 0$ on $[0, 1/5]$, we have $g(a) \le g(1/5) = 0.4399872 < 4/9$. This proves $a > 1/5$.

The apex equation is

$$p = b^2 - u^2 - v^2 - \sum_i y_i^2.$$

Hence $b > 2/3$ and $a < 1/3$. Also $s^2 \le 2(u^2+v^2) < 2((4/5)^2 - 4/9) = 88/225 < 4/9$, so $s < 2/3$.

At fixed a, u, v, z, apply Lemma 4.11 with $A = 3a$ and $B = 4/9$. Here $2z/(9a) < 8/9$, so every leaf is positive and $\max_i y_i < z/(D-1) < 4/45$. Their stationary equations are

$$2ay_i - \beta_i y_i^2 = 2ab + 2uv - p. \tag{4.45}$$

The difference of the two common-vertex equations gives $2(a+z)(u-v) = q_1u^2 - q_2v^2$. With $\delta := |u-v|/s \in [0,1)$, this yields

$$\delta \le \frac{s(1+\delta)^2}{18(1-s)} < \frac{(1+\delta)^2}{9}.$$

Solving the quadratic inequality gives $\delta < (7-\sqrt{45})/2 < 3/20$. Consequently,

$$4uv - (q_1u^2 + q_2v^2) \ge s^2\left(\frac{7}{9} - \frac{11}{9}\delta^2\right) > \frac{3}{4}s^2. \tag{4.46}$$

For clarity, the averages of the leaf equations and the two common-vertex equations are

$$p = 2ab - \frac{2az}{D} + 2uv + \frac{1}{D}\sum_i \beta_i y_i^2 = 2ab - as + sz + \frac{q_1u^2 + q_2v^2}{2}.$$

Subtracting the two expressions for p and multiplying by two gives

$$0 = 2s(a-z) + 4uv - (q_1u^2 + q_2v^2) - \frac{4az}{D} + \frac{2}{D}\sum_i \beta_i y_i^2. \tag{4.47}$$

Suppose $z \le 1/3$. Using $D \ge 10$ and Equation (4.46), the right side of Equation (4.47) is greater than

$$F(a,z) \coloneqq (1-a-z)\left(\frac{3}{4} + \frac{5}{4}a - \frac{11}{4}z\right) - \frac{2}{5}az.$$

On $1/5 \le a \le 1/3$ and $0 \le z \le 1/3$, its derivative in z is $-7/2 + 11a/10 + 11z/2 < 0$. But

$$F\left(a, \frac{1}{3}\right) = -\frac{5}{4}a^2 + \frac{13}{15}a - \frac{1}{9}$$

is concave in a, with endpoint values $11/900$ and $7/180$. It is therefore positive throughout the interval, a contradiction. Thus $z > 1/3$.

By Lemma 4.11,

$$\frac{\min_i y_i}{\max_i y_i} > 1 - \frac{8}{9(D-1)} \ge \frac{73}{81}, \qquad \min_i y_i > \frac{73}{243D} > \frac{3}{10D}.$$

Since $81/73 < 10/9$, the required ratio bound follows. □

We use these weight bounds to reduce targets below the book maximum toward $4/9$. Fix $q_1, q_2 \in \{\beta_m : m \ge 1\} \cup \{4/9\}$, and let $A = A(D, q_1, q_2) > 4/9$ be the maximum when every leaf coefficient equals $4/9$. Put $w_D \coloneqq 3(D-7)/(250D^3)$. For every $p = 4/9 + \varepsilon \in [A - w_D, A]$ with $\varepsilon \ge 10^{-15}$, there is a finite substitution function φ such that

$$p = \varphi(u), \qquad \frac{4}{9} \le u \le \frac{4}{9} + \frac{\varepsilon}{2}, \tag{4.48}$$

and all fixed coefficients belong to $\{0, 4/9\} \cup \{\beta_m : m \ge 1\}$.

To prove this, reserve one leaf and lower the other coefficients from $4/9$ toward zero successively. By Lemma 4.15, lowering $D-7$ entire coefficients without reaching p would decrease the maximum by more than w_D. Thus the crossing occurs by the $(D-7)$th leaf. At most $D-8$ leaves have been fixed at zero, leaving at least seven controls, including the crossing coordinate, besides the reserved leaf. Apply Lemma 4.10 with $q = 7$ to obtain Equation (4.48).

To apply the link operation with target $1/2$, the interval must reach $\Theta_3(4/9, 1/2) = 162/361$. We verify that the intervals certified by 132 choices of D, q_1, q_2 cover $[4/9, 162/361]$. Put $d \coloneqq u - v$.

When all leaf coefficients equal $4/9$, the proof of Lemma 4.15 gives $y_i < 4/45$. The function $t \mapsto 2at - \frac{4}{9}t^2$ is strictly increasing on this range, so Equation (4.45) forces $y_i = z/D$. Hence every relevant maximum is represented by the cubic polynomial

$$\begin{aligned} P_D(a,s,d) \coloneqq 3a\left[(1-a)^2 - \frac{s^2+d^2}{2} - \frac{z^2}{D}\right] + \frac{3}{2}(s^2-d^2)z \\ + q_1\left(\frac{s+d}{2}\right)^3 + q_2\left(\frac{s-d}{2}\right)^3 + \frac{4z^3}{9D^2}, \qquad z \coloneqq 1-a-s. \end{aligned} \tag{4.49}$$

For the cubic in Equation (4.49) and a feasible rational test vector $\boldsymbol{x}$, define

$$L \coloneqq P_D(\boldsymbol{x}), \qquad U \coloneqq P_D(\boldsymbol{x}) + \|\nabla P_D(\boldsymbol{x})\|^2. \tag{4.50}$$

Suppose $L > 4/9$ and $-\nabla^2 P_D(\boldsymbol{x}) \succeq (2/3)I$. Every global maximizer has positive apex, common-vertex, and leaf weights by Lemma 4.15. It therefore gives an interior stationary point of P_D. Applying Lemma 4.12 with $h = 2/3$ proves $L \le A \le U$. It follows that

$$[U - w_D, L] \subseteq [A - w_D, A]. \tag{4.51}$$

The 132 rational rows in `parts/ordinary/certificates/book_windows.json` specify D, $q_1 = \beta_m$, $q_2 = \beta_n$, and a weight vector $(a,u,v,z) = 10^{-8}(A_0, U_0, V_0, Z_0)$ of total weight one; the test vector in

Equation (4.50) is $\boldsymbol{x} = (a, u+v, u-v)$. The resulting interval bounds are recorded in `parts/ordinary/certificates/book_certificate.json`. The symbol ∞ denotes coefficient $4/9$. At all 132 test points, the 3 leading principal minors of $-\nabla^2 P_D(\boldsymbol{x}) - (2/3)I$ exceed 6, 2, and 2, respectively. These exact rational checks, performed by `parts/ordinary/certificates/verify_books.py`, give the required Hessian bound by Sylvester's criterion. The same program verifies feasibility, $L > 4/9$, and coverage of $[4/9, 162/361]$ by the intervals in Equation (4.51). More precisely, after sorting the left endpoints, every interval that extends the covered right endpoint overlaps the preceding union by more than $2 \cdot 10^{-7}$. The final right endpoint exceeds $162/361$ by more than 10^{-5}.

Proposition 4.16. *If* $[4/9, 4/9 + 2 \cdot 10^{-6}] \subseteq \mathcal{A}$, *then* $[4/9, 162/361] \subseteq \mathcal{A}$.

Proof. Let $p \in (4/9 + 2 \cdot 10^{-6}, 162/361]$ and choose a certified interval containing p. Since $p - 4/9 > 10^{-15}$, Equation (4.48) writes $p = \varphi(u)$ with $0 \le u - 4/9 \le (p - 4/9)/2$ and all fixed coefficients in $\mathcal{A}$. Repeat whenever the new target exceeds $4/9 + 2 \cdot 10^{-6}$. After finitely many steps it belongs to the assumed initial interval. Applying Lemma 4.2 in reverse order proves $p \in \mathcal{A}$. □

4.6 Extension to one

Apply Lemma 3.3 with $r = 3$ and $F = K_k$, where $k \ge 2$ and $t = 1 - 1/k$ is the normalized Lagrangian of K_k. For every finite 3-graph H, this gives a finite 3-graph with normalized Lagrangian $\Theta_3(\lambda(H), t)$, where

$$\Theta_3(x,t) = \max_{0 \le b \le 1} \{3t(1-b)b^2 + xb^3\} = \begin{cases} \dfrac{4t^3}{(3t-x)^2}, & 0 \le x \le t, \\ x, & t \le x \le 1. \end{cases}$$

Preservation of approximation from below gives $\Theta_3(x,t) \in \mathcal{A}$ whenever $x \in \mathcal{A}$. This also holds for $t = 1$: let $k \to \infty$ and use continuity and monotonicity to approximate from below by the same finite constructions.

For $x < t$, an exact factorization gives

$$\Theta_3(x,t) - x = \frac{(t-x)^2(4t-x)}{(3t-x)^2} > 0. \tag{4.52}$$

Suppose $[4/9, b] \subseteq \mathcal{A}$ and $\Theta_3(4/9, t) \le b < t$. Since $\Theta_3(4/9, t) > 4/9$, we have $4/9 < b$. Apply Lemma 2.5 with $T(x) = \Theta_3(x,t)$ on $[4/9, t]$, using (4.52) and the preceding preservation property, to obtain $[4/9, t) \subseteq \mathcal{A}$. The endpoint follows by approximation from below and Lemma 4.1.

The three required starting values are

$$\Theta_3\left(\frac{4}{9}, \frac{1}{2}\right) = \frac{162}{361}, \qquad \Theta_3\left(\frac{4}{9}, \frac{2}{3}\right) = \frac{24}{49} < \frac{1}{2}, \qquad \Theta_3\left(\frac{4}{9}, 1\right) = \frac{324}{529} < \frac{2}{3}.$$

Therefore,

$$\left[\frac{4}{9}, \frac{162}{361}\right] \subseteq \mathcal{A} \implies \left[\frac{4}{9}, \frac{1}{2}\right] \subseteq \mathcal{A} \implies \left[\frac{4}{9}, \frac{2}{3}\right] \subseteq \mathcal{A} \implies \left[\frac{4}{9}, 1\right] \subseteq \mathcal{A}. \tag{4.53}$$

Proof of Theorem 1.1(iii). By Proposition 4.7, $[4/9, 4/9 + e^{-10^9}] \subseteq \mathcal{A}$. Proposition 4.14 extends this to $[4/9, 4/9 + 2 \cdot 10^{-6}]$, and Proposition 4.16 gives $[4/9, 162/361] \subseteq \mathcal{A}$. Applying Equation (4.53), we obtain $[4/9, 1] \subseteq \mathcal{A}$. Finally, Lemma 4.1 gives $\pi(\mathcal{F}_\alpha) = \alpha$ for every $\alpha \in [4/9, 1]$.

For the higher-uniformity assertion, recall from Lemma 3.3 that $p_r \Pi_\infty^{(r-1)} \cup \Pi_\infty^{(r-1)} \subseteq \Pi_\infty^{(r)}$. In particular, $[3/16, 27/64] \cup [4/9, 1] \subseteq \Pi_\infty^{(4)}$. Since $1/2 \in \Pi_\infty^{(3)}$, the map $x \mapsto \Theta_4(x, 1/2)$ preserves $\Pi_\infty^{(4)}$ and satisfies $\Theta_4(x, 1/2) > x$ for $0 \le x < 1/2$. Moreover, $\Theta_4(3/16, 1/2) = 6912/24389 < 27/64$. Applying Lemma 2.5 and using closedness gives $[3/16, 1/2] \subseteq \Pi_\infty^{(4)}$, and hence $[3/16, 1] \subseteq \Pi_\infty^{(4)}$.

For $r \ge 5$, suppose inductively that $[2\alpha_{r-1}, 1] \subseteq \Pi_\infty^{(r-1)}$. Since $2\alpha_{r-1} \le 3/16 < 1/3 < p_r$, the intervals $[2\alpha_r, p_r]$ and $[2\alpha_{r-1}, 1]$ overlap. Both intervals are contained in $\Pi_\infty^{(r)}$ by the preceding inclusions, and their union is $[2\alpha_r, 1]$. This proves the assertion for every $r \ge 3$. □

5 Minimum-degree densities and the three-uniform case

We prove Theorem 1.2 for $\ell \geq 3$ by cyclic links and for $(r,\ell) = (3,2)$ by symmetric vertex templates. Both constructions realize the minimum-degree and independent vertex-set densities simultaneously. The remaining cases $\ell = 2$ and $r \geq 4$ are treated in Sections 6 and 7.

For the simultaneous constructions, write $\rho_\ell^{(r)}(\mathcal{F}) \coloneqq \pi_{\{\{1\},\dots,\{\ell\}\}}(\mathcal{F})$. These tests restrict each of the first ℓ coordinates to an independently chosen vertex set. By Lemma 2.1 and the comparison of test systems in Section 2, $\gamma_\ell^{(r)} \leq \rho_\ell^{(r)} \leq \pi$. We work with monotone classes, which are precisely classes of the form $\mathrm{Forb}(\mathcal{F})$. For a partition into named parts, an edge has profile $X_1^{a_1}\cdots X_t^{a_t}$ if it has a_i vertices in X_i. Every construction below permits arbitrary subgraphs of the specified allowed edges.

We record the two elementary observations that allow the degree and local-density upper bounds to be proved together. Write $(x)_p \coloneqq x(x-1)\cdots(x-p+1)$. For a nonzero nonnegative vector $\boldsymbol{z}$ on $\mathbb{Z}_a$ and $1 \leq b < a$,

$$\min_{i\colon z_i>0} \sum_{j\in[b]} z_{i+j} \leq \frac{b}{a}\sum_i z_i. \tag{5.1}$$

Indeed, the average window has the right-hand value. If a minimizing index has weight zero, moving one step backwards changes its window by $-z_{i+b} \leq 0$. Continuing through zero coordinates reaches a supported index without increasing the window.

For $\ell = 2$, the independent vertex test with first two sets X, Y averages the normalized pair degrees over their distinct ordered pairs. If $\rho_2^{(r)}(\mathscr{P}) \leq x < 1$, then for every $\varepsilon > 0$ every sufficiently large $H \in \mathscr{P}$ has sets X, Y with $|X||Y| - |X\cap Y| \geq cn^2$, for some $c > 0$, and

$$\sum_{\substack{u\in X,\ v\in Y\\ u\neq v}} d_H(\{u,v\}) < (x+\varepsilon)(|X||Y| - |X\cap Y|)\binom{n-2}{r-2}.$$

To see this, choose a fixed error $\eta > 0$ excluding $(x+\varepsilon,\eta)$-dense hosts. A violating test has volume at least $\eta n^r/(x+\varepsilon)$ and gives the assertion after dividing by $(n-2)_{r-2}$.

Consequently, in a part of positive limiting size one may use a whole-part average, a pointwise bound on every pair of a fixed type, or a low-density rectangle in one common auxiliary hypergraph. In the last case, a uniform bound

$$\frac{d_H(uv)}{\binom{n-2}{r-2}} \leq A(\boldsymbol{z})\frac{d_{H_0}(uv)}{\binom{v(H_0)-2}{s-2}} + B(\boldsymbol{z}) + o(1), \qquad A, B \geq 0, \tag{5.2}$$

transfers the auxiliary upper bound x to $A(\boldsymbol{z})x + B(\boldsymbol{z})$. The same rectangle works for every outside contribution because the inequality is uniform in u, v. After passing to limiting part ratios $\boldsymbol{z}$, fix $d > A(\boldsymbol{z})x + B(\boldsymbol{z})$. Choosing the auxiliary error sufficiently small gives constants $\delta, c > 0$ such that, for all sufficiently large hosts, the resulting rectangle test K satisfies $|K| \geq cn^r$ and $e_H(K) \leq (d-\delta)|K|$. For $\eta < c\delta$, this violates the (d,η)-density inequality for the independent vertex-set tests. Sublinear parts contribute only $o(n^{r-2})$ completions.

Independent relabelings allow ordinary link densities to give uniform degree bounds.

Lemma 5.1. *Let $M \geq 1$ and $0 \leq a \leq s \leq m$ be integers, with $s \geq 1$, and let $H_1, \dots, H_M$ be s-graphs on $[m]$. Choose independent uniformly random permutations $\sigma_1, \dots, \sigma_M$ of $[m]$. For every $\varepsilon > 0$,*

$$\mathbb{P}\left[\max_{A\in\binom{[m]}{a}} \left|\frac{1}{M}\sum_{i=1}^{M}\frac{d_{\sigma_i H_i}(A)}{\binom{m-a}{s-a}} - \frac{1}{M}\sum_{i=1}^{M}\frac{|H_i|}{\binom{m}{s}}\right| > \varepsilon\right] \leq 2\binom{m}{a}e^{-2\varepsilon^2 M}.$$

Here $d_H(\varnothing) = |H|$. Moreover, if H is an s-graph on $[m]$, $0 \leq a \leq \ell \leq s$, and $A \in \binom{[m]}{a}$, then

$$\frac{d_H(A)}{\binom{m-a}{s-a}} \geq \frac{\delta_\ell(H)}{\binom{m-\ell}{s-\ell}}.$$

Proof. For fixed A, each normalized degree lies in $[0,1]$ and has expectation $|H_i|/\binom{m}{s}$, since $\sigma_i^{-1}(A)$ is uniform on $\binom{[m]}{a}$. Hoeffding's inequality [15] and a union bound give the first assertion. For the second,

double count pairs (L, E) with $A \subseteq L \subseteq E \in H$ and $|L| = \ell$ to obtain

$$\binom{s-a}{\ell-a} d_H(A) = \sum_{L \supseteq A,\ |L|=\ell} d_H(L) \geq \binom{m-a}{\ell-a} \delta_\ell(H).$$

Dividing by $\binom{s-a}{\ell-a}\binom{m-a}{s-a}$ proves the inequality. □

The following cyclic-link construction uses the terminal interval in Theorem 3.2 and the near-regular examples supplied by Lemma 2.2.

Fix $r > \ell \geq 2$, put $p := r - \ell$, and choose integers $p \leq b \leq a - \ell$. For a monotone class $\mathscr{Q}$ of ℓ-graphs, a member of $\mathscr{W}^{r,\ell}_{a,b}(\mathscr{Q})$ admits a cyclic partition $V_0 \sqcup \cdots \sqcup V_{a-1}$. Its link-type edges have ℓ vertices in V_i and one vertex in each of p distinct parts among $V_{i+1}, \ldots, V_{i+b}$; for each fixed outside transversal T, the link $L_{i,T} := \{S \in \binom{V_i}{\ell} : S \cup T \in H\}$ belongs to $\mathscr{Q}$. Also allow all edges meeting every part in at most $\ell - 1$ vertices. No other edges are allowed. A link-type edge has a unique source part, so these requirements are preserved by taking subgraphs, and empty links show that the class contains every edgeless hypergraph. This is the transversal form of the cyclic part construction of Mubayi and Zhao [32, Section 3.3]. Figure 12 gives a four-uniform example.

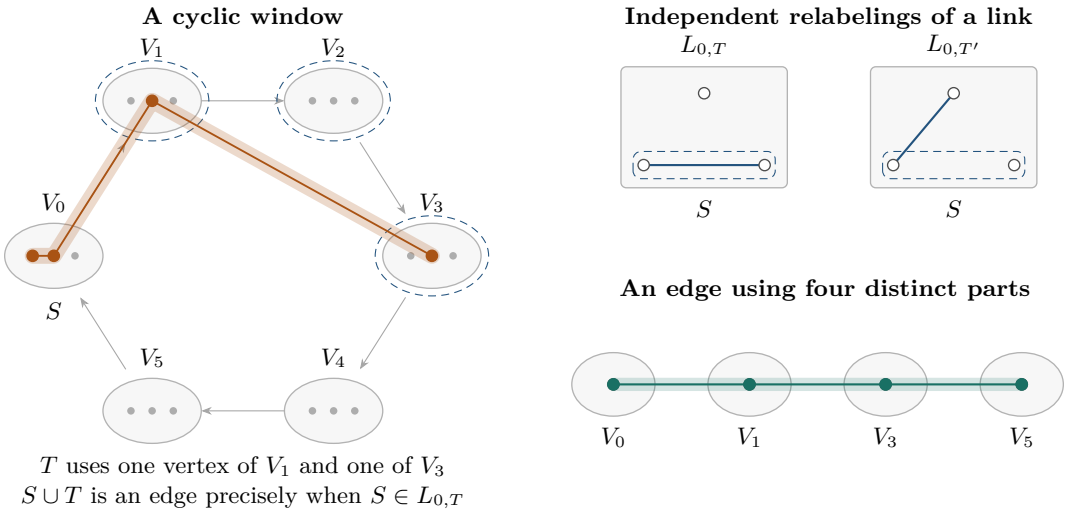


Figure 12: The cyclic-link construction for $r = 4$, $\ell = 2$, $a = 6$, and $b = 3$. For source part V_0, the outside vertices lie in distinct parts of the window V_1, V_2, V_3. The upper-right panels illustrate two relabelings of a one-edge link graph, chosen independently for different outside transversals. Edges meeting four distinct parts are also allowed, as shown below.

Proposition 5.2. *If $\pi(\mathscr{Q}) = \beta$, then*

$$\gamma^{(r)}_\ell\big(\mathscr{W}^{r,\ell}_{a,b}(\mathscr{Q})\big) = \rho^{(r)}_\ell\big(\mathscr{W}^{r,\ell}_{a,b}(\mathscr{Q})\big) = \beta \frac{(b)_p}{a^p}.$$

Proof. Use nearly equal parts for the minimum-degree lower bound. Choose an ℓ-graph of density $\beta - o(1)$ on each part and independently relabel a copy for every outside transversal. A fixed root inside a source part belongs to $\Theta(n^p)$ independent links. Lemma 5.1, applied with link uniformity and root size both equal to ℓ, gives normalized degree at least $\beta(b)_p/a^p - o(1)$ simultaneously. A root meeting at least 2 parts can be completed in p further distinct parts, giving normalized degree at least $(a-\ell)_p/a^p - o(1) \geq (b)_p/a^p - o(1)$.

For the local-density upper bound, fix $0 < \xi < 1/a$, put $w_i = |V_i|/n$, and set $z_i = w_i$ when $w_i \geq \xi$ and $z_i = 0$ otherwise. By Equation (5.1), some i has $|V_i| \geq \xi n$ and $\sum_{j \in [b]} |V_{i+j}| \leq (b/a + b\xi)n$. Let $\mathcal{T}_i$ be its outside transversals. Double counting and the ordinary-density bound give

$$\frac{1}{\binom{|V_i|}{\ell}} \sum_{S \in \binom{V_i}{\ell}} d_H(S) = \frac{1}{\binom{|V_i|}{\ell}} \sum_{T \in \mathcal{T}_i} |L_{i,T}| \leq (\beta + \varepsilon)|\mathcal{T}_i|, \qquad |\mathcal{T}_i| \leq \binom{b}{p} \left(\frac{1}{a} + \xi\right)^p n^p.$$

Restricting all ℓ tested vertex positions to V_i gives a test of volume $(|V_i|)_\ell(n-\ell)_p = \Omega_\xi(n^r)$ with density at most $(\beta+\varepsilon)(b)_p(1/a+\xi)^p + o(1)$. Let ε, ξ tend to zero, choosing the host-density error afterwards. This proves $\rho_\ell^{(r)} \le \beta(b)_p/a^p$, and the minimum-degree lower bound proves both equalities. □

For fixed p, the factors $(b)_p/a^p$ with $p \le b \le a-\ell$ are dense in $(0,1)$: take $b = \lfloor t^{1/p}a \rfloor$ for a prescribed $t \in (0,1)$ and let $a \to \infty$. For $\ell \ge 3$, Theorem 3.2 supplies every ordinary density $\beta \in [c_\ell, 1]$. For a target $\alpha \in (0,1)$, choose and fix an available factor λ with $\alpha < \lambda < \min\{\alpha/c_\ell, 1\}$, and take $\beta = \alpha/\lambda$. The preceding proposition then realizes α in both parameters. The edgeless and unrestricted classes give 0 and 1. This proves Theorem 1.2 for $\ell \ge 3$.

For $\ell = 2$, ordinary graph densities are discrete by the Erdős–Stone theorem [8], so the preceding argument does not give an initial interval. Write

$$\Sigma_r := \{x \in [0,1] : \text{some monotone class } \mathscr{P} \text{ has } \gamma_2^{(r)}(\mathscr{P}) = \rho_2^{(r)}(\mathscr{P}) = x\}.$$

In these constructions, a class *realizes* x when both parameters equal x, and its density formula denotes this common value.

5.1 Symmetric templates and cyclic substitution

The case $(r,\ell) = (3,2)$ uses symmetric relations on vertex colors. Their least weighted completion probabilities give both minimum-codegree and independent vertex-set densities. Compatible refinements realize limits of these probabilities. The construction below produces an interval adjacent to zero, and a cyclic substitution extends that interval to all of $[0,1]$.

A *symmetric vertex template* is a relation $P \subseteq C^r$ on a nonempty finite color set C, invariant under every permutation of the coordinates. Repeated colors are allowed. An r-graph is *P-colorable* if a map from its vertices to C sends every edge to an allowed color tuple. Symmetry makes this condition independent of the order of the vertices. Let $\mathcal{F}_P$ consist of the finite r-graphs that are not P-colorable, with one representative of each isomorphism class. Colorability is preserved by taking subhypergraphs, so an r-graph is $\mathcal{F}_P$-free if and only if it is P-colorable. Every edgeless r-graph is colorable.

A *homomorphism* $P \to P'$ is a color map preserving the allowed tuples. It makes every P-colorable r-graph P'-colorable. To obtain the density of these classes, we optimize the least probability of completing prescribed supported colors. For a finite relation $P \subseteq C^J$, write $\Lambda_P(\boldsymbol{w}) := \sum_{\boldsymbol{c}\in P} \prod_{j\in J} w_{c_j}$; its Lagrangian is $\lambda(P) = \max_{\boldsymbol{w}\in S(C)} \Lambda_P(\boldsymbol{w})$. The support condition requires the following compactness argument.

Lemma 5.3. *Let C, J be finite nonempty sets, let $P \subseteq C^J$, and let $\mathcal{I}$ be a nonempty family of subsets of J. The maximum*

$$\max_{\boldsymbol{w}\in S(C)} \min \left\{ \sum_{\substack{\boldsymbol{c}\in P \\ \boldsymbol{c}|_I = \boldsymbol{a}}} \prod_{j\in J\setminus I} w_{c_j} : I \in \mathcal{I},\ \boldsymbol{a} \in (\operatorname{supp} \boldsymbol{w})^I \right\}$$

is attained. A color map sending P into another finite relation on the same coordinate set cannot decrease this maximum. The value is algebraic; it is rational if every prescribed set leaves exactly one free coordinate.

This maximum is at most the Lagrangian $\lambda(P)$. Equality holds precisely when some maximizing weighting for Λ_P has every supported completion probability equal to $\lambda(P)$.

Proof. For $d \in [0,1]$, a weighting has all supported completion probabilities at least d precisely when

$$\left(\prod_{j\in I} w_{a_j}\right)\left(\sum_{\substack{\boldsymbol{c}\in P \\ \boldsymbol{c}|_I = \boldsymbol{a}}} \prod_{j\in J\setminus I} w_{c_j} - d\right) \ge 0$$

for every $I \in \mathcal{I}$ and $\boldsymbol{a} \in C^I$. These finitely many closed inequalities define a compact set of pairs $(\boldsymbol{w}, d)$. Maximizing d proves attainment.

For a color map θ, push the weighting forward by $w'_b := \sum_{\theta(c)=b} w_c$. Every supported assignment in the image has a supported preimage. Each permitted completion of that preimage maps to a permitted

completion in the image relation. Independent colors with distribution $\boldsymbol{w}$ have independent images with distribution $\boldsymbol{w}'$, so the completion probability cannot decrease. Taking the minimum and then the maximum proves monotonicity.

The feasible set is defined by rational polynomial inequalities. Quantifier elimination over real closed fields [2] therefore describes its maximum by polynomial sign conditions over $\mathbb{Q}$. Some nonzero polynomial in this description must vanish at the maximum, since otherwise the conditions would hold on a neighborhood rather than at a single value. This proves algebraicity. If only one coordinate is free, fix a nonempty candidate support and impose the corresponding linear completion constraints on probability weights there. Each resulting rational linear program has rational optimum. Maximizing over the finitely many candidate supports gives the capacity: the actual support of an optimal weighting is among them, and zero-weight colors in a feasible solution may be discarded.

For any $I \in \mathcal{I}$, average the completion probabilities over independent prescribed colors with law $\boldsymbol{w}$. The result is

$$\sum_{\boldsymbol{a}\in(\operatorname{supp}\boldsymbol{w})^I} \left(\prod_{j\in I} w_{a_j}\right) \sum_{\substack{\boldsymbol{c}\in P\\ \boldsymbol{c}|_I=\boldsymbol{a}}} \prod_{j\in J\setminus I} w_{c_j} = \Lambda_P(\boldsymbol{w}).$$

Their minimum is therefore at most $\Lambda_P(\boldsymbol{w}) \le \lambda(P)$. If the capacity equals $\lambda(P)$, take a weighting attaining it. Each supported completion probability is at least $\lambda(P)$, and their average is at most $\lambda(P)$. Hence they all equal $\lambda(P)$ and the weighting maximizes Λ_P. The converse follows by evaluating at the weighting in the statement. □

For d strictly above the maximum in Lemma 5.3, the same polynomial constraints give a uniform positive deficit,

$$\min_{\boldsymbol{w}\in S(C)} \max \left\{ d\prod_{i\in I} w_{a_i} - \sum_{\substack{\boldsymbol{c}\in P\\ \boldsymbol{c}|_I=\boldsymbol{a}}} \prod_{j\in J} w_{c_j} : I \in \mathcal{I},\ \boldsymbol{a} \in C^I \right\} > 0. \tag{5.3}$$

For every weighting, a supported assignment with completion probability below d makes the expression in braces positive. Its maximum is continuous in the weighting, so compactness makes the minimum positive as well.

The density formula now follows by counting completions in a blowup and testing suitable color classes.

Lemma 5.4. *For every finite symmetric vertex template $P \subseteq C^r$ and $1 \le \ell < r$,*

$$\gamma_\ell^{(r)}(\mathcal{F}_P) = \rho_\ell^{(r)}(\mathcal{F}_P) = \max_{\boldsymbol{w}\in S(C)} \min_{\boldsymbol{a}\in(\operatorname{supp}\boldsymbol{w})^\ell} \sum_{\substack{\boldsymbol{b}\in C^{r-\ell}\\ (\boldsymbol{a},\boldsymbol{b})\in P}} \prod_{i\in[r-\ell]} w_{b_i}.$$

Proof. Write d for the maximum on the right and put $p := r-\ell$. It is attained by Lemma 5.3, with $J=[r]$ and $\mathcal{I}=\{[\ell]\}$. Since $\gamma_\ell^{(r)} \le \rho_\ell^{(r)}$, it suffices to prove the minimum-degree lower bound and the independent vertex-set upper bound.

Choose a maximizing weighting $\boldsymbol{w}$. Partition n vertices into classes of sizes $nw_c+O(1)$, leaving unsupported classes empty, and include every allowed edge. For an ordered ℓ-tuple of distinct roots with color tuple $\boldsymbol{a}$, the number of ordered completions is

$$n^p \sum_{\substack{\boldsymbol{b}\in C^p\\ (\boldsymbol{a},\boldsymbol{b})\in P}} \prod_{i\in[p]} w_{b_i} + O_{r,|C|}(n^{p-1}).$$

Indeed, counting without the distinctness restrictions gives the leading term up to the rounding error; excluding repetitions and the roots removes $O_r(n^{p-1})$ tuples. This estimate is uniform over root color tuples, including repeated colors. Dividing by $(n-\ell)_p$ gives normalized minimum ℓ-degree at least $d - O_{r,|C|}(1/n)$.

For the upper bound, suppose $d<1$ and fix $0<\varepsilon<1-d$. Apply Equation (5.3) at $d+\varepsilon$ to obtain a deficit $\kappa>0$. In a witnessing coloring of a P-colorable n-vertex r-graph, let $w_c = |V_c|/n$. Choose a

supported assignment $\boldsymbol{a}$ realizing a deficit at least κ, and restrict the first ℓ positions to $V_{a_1}, \dots, V_{a_\ell}$. Leave the other positions unrestricted. The resulting injective test K satisfies

$$(d+\varepsilon)|K| - e_H(K) \geq \left(\kappa - O_r\left(\frac{1}{n}\right)\right) n^r.$$

Indeed, counting all tuples gives exactly the two polynomial terms in Equation (5.3), and excluding repeated vertices changes each count by $O_r(n^{r-1})$. For $\eta < \kappa/2$ and sufficiently large n, this contradicts $(d+\varepsilon, \eta, \{\{1\}, \dots, \{\ell\}\})$-density. Letting $\varepsilon \to 0$ proves the upper bound. If $d = 1$, the lower bound already gives equality. □

For $\ell = 1$, the maximum in Lemma 5.4 equals

$$\max_{\boldsymbol{w} \in S(C)} \sum_{\boldsymbol{c} \in P} \prod_{i \in [r]} w_{c_i}.$$

At a maximizing weighting, symmetry makes each supported one-color completion probability equal to $r^{-1}\partial\Lambda_P/\partial w_c$. The supported derivatives agree, and Euler's identity gives their common value $r\Lambda_P$. Thus every supported completion probability equals $\lambda(P)$, and the equality criterion in Lemma 5.3 applies. If P consists of all orderings of the edges of an r-graph H, this value is the normalized Lagrangian $\lambda(H)$.

For a finite k-coordinate palette P and a probability weighting $\boldsymbol{w}$, prescribe supported colors in all but one coordinate. Define the least completion probability by

$$\gamma_k(P, \boldsymbol{w}) := \min \left\{ \sum_{\substack{\boldsymbol{b} \in P \\ \boldsymbol{b}|_{[k]\setminus\{j\}} = \boldsymbol{a}}} w_{b_j} : j \in [k],\ \boldsymbol{a} \in (\operatorname{supp} \boldsymbol{w})^{[k]\setminus\{j\}} \right\}.$$

Put $\Gamma_k(P) := \max_{\boldsymbol{w}} \gamma_k(P, \boldsymbol{w})$. By Lemma 5.3, with $J = [k]$ and $\mathcal{I} = \binom{[k]}{k-1}$, the maximum is attained, is rational, and is nondecreasing under homomorphisms. A weighting is *d-feasible* if $\gamma_k(P, \boldsymbol{w}) \geq d$.

The cyclic construction of Mubayi and Zhao [32, Section 3.3] extends an interval adjacent to zero. We formulate the substitution for every $k \geq 3$.

Fix $1 \leq a < b$, with outer labels in $\mathbb{Z}_b$. For a finite k-coordinate palette P, define $\Psi^{(k)}_{a,b}(P)$ on $\mathbb{Z}_b \times C(P)$. If $k-1$ outer labels equal i and the remaining label is j, admit the tuple when $j - i \in [a]$ modulo b. Admit every other nonconstant outer pattern. When all outer labels agree, use P on the inner colors. The substitution preserves symmetry when P is symmetric.

Lemma 5.5. *For every $k \geq 3$ and $1 \leq a < b$, there is a nondecreasing, 1-Lipschitz function $f_{k;a,b} \colon [0,1] \to [0,1]$ such that*

$$\Gamma_k(\Psi^{(k)}_{a,b}(P)) = f_{k;a,b}(\Gamma_k(P)), \qquad f_{k;a,b}(0) = \frac{a}{b}, \qquad f_{k;a,b}(z) \geq \frac{a+z}{b}.$$

The substitution preserves homomorphisms. If the limits of Γ_k along projective sequences of finite palettes contain $[0,\eta]$ for some $\eta > 0$, then they contain $[0,1]$. The same assertion holds with all palettes required to be symmetric.

Proof. Fix outer weights $\boldsymbol{p} \in S(\mathbb{Z}_b)$, and write $c_i(\boldsymbol{p}) = \sum_{j \in [a]} p_{i+j}$. In a positive outer class i, choose a normalized inner weighting $\boldsymbol{w}^{(i)}$. Prescribing $k-1$ supported colors in this class gives minimum completion weight

$$c_i(\boldsymbol{p}) + p_i \gamma_k(P, \boldsymbol{w}^{(i)}).$$

For $k = 3$, a row with distinct outer labels i, j has completion weight

$$1 - p_i - p_j + p_i \mathbf{1}_{\{j-i \in [a]\}} + p_j \mathbf{1}_{\{i-j \in [a]\}}.$$

For $k \geq 4$, a mixed row with $k-2$ outer labels equal to i and its remaining label equal to j has completion weight $1 - p_i + p_i \mathbf{1}_{\{j-i \in [a]\}}$; every other mixed row has completion weight one. These expressions do not depend on the inner colors. Optimize the inner weighting separately in each positive block, take the

minimum over supported rows, and then take the supremum over $\boldsymbol{p}$. This defines $f_{k;a,b}$ and proves the capacity identity.

For fixed $\boldsymbol{p}$, the row minimum is a minimum of affine functions of z whose slopes belong to $[0,1]$. It is therefore nondecreasing and 1-Lipschitz, and taking the supremum preserves both properties. At $z=0$, a constant prescribed row and Equation (5.1) give the upper bound a/b.

Uniform outer weights give constant-row value $(a+z)/b$. For $k \geq 4$, every mixed row has weight at least $(b-1)/b$ when $a \leq b-2$, and weight one when $a = b-1$. For $k=3$, its weight is at least $(b-2)/b$ if $a \leq b-3$. If $a = b-2$, at least one of $j-i, i-j$ belongs to $[a]$, so the weight is at least $(b-1)/b$; if $a = b-1$, it is one. Each bound is at least $(a+z)/b$. This proves the lower bound and the equality at zero.

A homomorphism $\theta\colon P \to P'$ induces $(i,c) \mapsto (i,\theta(c))$ between the substitutions. Continuity therefore gives the same transformation on limits of capacities. In particular, the image of a projectively realized interval $[0,\eta]$, with $0 < \eta \leq 1$, contains $[a/b,(a+\eta)/b]$. The empty one-color palette gives a/b after substitution; empty and full one-color palettes give zero and one.

For an irrational $x \in (0,1)$, choose N with $1/N < \min\{x,\eta\}$. Two fractional parts among $0, x, \dots, Nx$ lie in the same one of N equal half-open intervals. Hence there are $q \geq 1$ and an integer p with $0 < |qx-p| < 1/N$. If $qx - p > 0$, take $b = q, a = p$. Otherwise put $s = p - qx$ and $m = \lfloor 1/s \rfloor$. Irrationality gives $0 < 1 - ms < s$, so $b = mq, a = mp - 1$ works. In either case $0 < bx - a < 1/N < x$ implies $1 \leq a < b$, and $x \in [a/b,(a+\eta)/b]$. Rational targets are already realized by constant chains. All constructions preserve symmetry, proving the last assertion as well. □

For the rest of this section, let $r=3$ and $\ell=2$. For $a,b \in C$, put $N_P(a,b) \coloneqq \{c \in C\colon (a,b,c) \in P\}$, and write $w(D) \coloneqq \sum_{c\in D} w_c$. Define

$$\gamma(P,\boldsymbol{w}) \coloneqq \min_{a,b\in\operatorname{supp}\boldsymbol{w}} w(N_P(a,b)), \qquad \Gamma(P) \coloneqq \max_{\boldsymbol{w}\in S(C)} \gamma(P,\boldsymbol{w}).$$

Thus $\Gamma(P) = \Gamma_3(P)$, and the minimum includes $a=b$. By Lemma 5.4, $\gamma_2^{(3)}(\mathcal{F}_P) = \rho_2^{(3)}(\mathcal{F}_P) = \Gamma(P)$.

Lemma 5.6. *Suppose that $P_0 \leftarrow P_1 \leftarrow P_2 \leftarrow \cdots$ is a sequence of finite symmetric ternary templates and homomorphisms, and put $\mathcal{F} \coloneqq \bigcup_{m\geq 0} \mathcal{F}_{P_m}$. Then*

$$\gamma_2^{(3)}(\mathcal{F}) = \rho_2^{(3)}(\mathcal{F}) = \lim_{m\to\infty} \Gamma(P_m).$$

Proof. The homomorphisms make $\mathcal{F}_{P_m}$ increasing. Apply Equation (2.2) and Lemmas 2.1 and 5.4 to the systems $\{\{1,2\}\}$ and $\{\{1\},\{2\}\}$. □

Let $\mathcal{R}$ denote the set of all limits of capacities along such projective sequences, including constant sequences. Thus $\mathcal{R} \subseteq \Gamma_2^{(3)}$, and each value in $\mathcal{R}$ has the simultaneous realization in Lemma 5.6.

For $k=3$, the substitution with an empty inner palette is symmetric and has capacity a/b. Empty and full one-color templates supply the endpoints. Conversely every finite symmetric ternary template has rational capacity by Lemma 5.3. Thus these templates realize exactly $\mathbb{Q} \cap [0,1]$ for both density parameters. The associated forbidden families need not be finite. By Lemma 5.5, it now suffices to construct a nondegenerate interval $[0,\eta] \subseteq \mathcal{R}$.

5.2 Recursive symmetric templates

We now construct capacities whose reciprocals are sums of independently chosen negative continued fractions. Each coarse color class is refined along a rooted tree. Some coarse classes will supply the prescribed pairs that impose the completion constraints; we call them *anchor classes*. The remaining *sink classes* have independently variable refinement sequences.

Let $\phi_m(t) \coloneqq 1/(m-t)$. An anchor tree has 29 initial children and 400 children at each refinement. A sink tree has 10 initial children, 800 children at its first refinement, and either 800 or 801 thereafter. At every stage child 1 is refined next, and child 2 is its guard. For $n \geq 0$, put

$$F_{\mathrm{a},n} \coloneqq \phi_{28} \circ \phi_{400}^{\circ n}, \qquad F_{z,0} \coloneqq \phi_9, \qquad F_{z,n} \coloneqq \phi_9 \circ \phi_{800} \circ \phi_{m_{z,1}} \circ \cdots \circ \phi_{m_{z,n-1}} \quad \text{for } n \geq 1,$$

where each sink has its own infinite word $(m_{z,j})_{j\geq 1} \in \{800,801\}^{\mathbb{N}}$.

Proposition 5.7. *For every integer $h \geq 1$ and every choice of $3h$ sink words there are finite symmetric ternary templates $P_n(h)$ with homomorphisms $P_{n+1}(h) \to P_n(h)$ and*

$$\Gamma(P_n(h))^{-1} = 180h\big(1 + F_{\mathrm{a},n}(0)\big) + \sum_{z\in[3h]} \big(1 + F_{z,n}(0)\big). \tag{5.4}$$

For $n \geq 1$, the template depends only on the first $n-1$ digits of each sink word.

For a coarse class D, fix branching numbers $N_D(0), \ldots, N_D(n)$. The initial children have indices $0, \ldots, N_D(0) - 1$. At step $j \geq 1$, replace child 1 of the preceding group by $N_D(j)$ new children. Identify D with the terminal leaves, which are its colors. Let $D_k^{(j)}$ be the descendant leaves of child k in group j. Thus

$$D = \dot{\bigcup_{0\leq k<N_D(0)}} D_k^{(0)}, \qquad D_1^{(j-1)} = \dot{\bigcup_{0\leq k<N_D(j)}} D_k^{(j)} \quad \text{for } 1 \leq j \leq n.$$

Every $D_k^{(j)}$ with $j < n, k \neq 1$, and every $D_k^{(n)}$, is a singleton. We call all the sets $D_k^{(j)}$ *blocks*, including the active nonsingleton blocks. An anchor has $N_D(0) = 29$ and $N_D(j) = 400$ for $j \geq 1$. A sink z has $N_D(0) = 10, N_D(1) = 800$, and $N_D(j) = m_{z,j-1}$ for $j \geq 2$.

There are 3 top classes indexed by $\mathbb{Z}_3$. Each contains h consecutive copies of a list of 60 anchor coarse classes and one sink coarse class. Write $K \coloneqq 61h$, and label the coarse classes in a top class by $a \in \mathbb{Z}_K$. Copy μ has labels $61\mu, \ldots, 61\mu + 60$, with the last a sink. A copy comprises these lists in all 3 top classes, hence has 180 anchors and 3 sinks. The color set is the disjoint union of all their leaf sets.

For each top class i and copy, assign a different unordered pair of its anchor classes to each pair (D, k), where D is a coarse class in top class $i+1$ of the same copy and $0 \leq k < N_D(0)$. There are enough pairs, since $\binom{60}{2} = 1770 \geq 60 \cdot 29 + 10 = 1750$. Fix these assignments independently of n and the sink words. The pair assigned to $(D, 2)$ will impose all subsequent recursive constraints on D. Denote its two anchor classes by U_D, V_D, ordered by their coarse labels. Finally define

$$g(i, a) \coloneqq (i + 1, a + \delta_i), \qquad \delta_0 = \delta_1 = 0, \quad \delta_2 = 1, \tag{5.5}$$

with coordinates read modulo 3 and K. Since $g^3(i, a) = (i, a + 1)$, this is a single cycle on all $3K$ coarse classes.

We specify allowed multisets of 3 colors and include every permutation. The initial rules are the following.

(T1) Every triple with 3 distinct top indices is allowed. No triple with all top indices equal is allowed.

(T2) When top index i occurs twice, the singleton top index must be $i + 1$.

(T3) If the repeated-top colors belong to the same coarse class (i, a), their completion set is exactly $g(i, a)$. This set is never modified.

(T4) For different repeated-top coarse classes, an unassigned pair has the whole next top class as completions. A pair assigned to (D, k) has the initial completion set $D \setminus D_k^{(0)}$.

Only the pair assigned to $(D, 2)$ is modified further. At refinement $1 \leq j \leq n$, choose an injection

$$\iota_{D,j} \colon \{0, \ldots, N_D(j) - 1\} \longrightarrow \{0, \ldots, N_{\mathrm{a}}(j - 1) - 1\}^2, \qquad \iota_{D,j}(2) = (1, 1),$$

where $N_{\mathrm{a}}(0) = 29$ and $N_{\mathrm{a}}(s) = 400$ for $s \geq 1$. These injections exist because $N_D(1) \leq 800 \leq 29^2$ and $N_D(j) \leq 801 \leq 400^2$ later. Choose them consistently for each possible pair of branching numbers. If $\iota_{D,j}(k) = (u, v)$, delete every triple with repeated-top colors in $U_{D,u}^{(j-1)}$ and $V_{D,v}^{(j-1)}$ and singleton color in $D_k^{(j)}$, together with all its permutations.

These rules give a well-defined symmetric relation. Every triple with exactly two equal top indices has a unique repeated top index and a unique repeated-top coarse pair. If assigned, that coarse pair has only one target D. Its two different coarse labels identify the roles in the injection. No deletion changes a completion set from (T3) or a triple with distinct top indices.

Map each newly created leaf to the active leaf it replaces, fixing other leaves. All old deletions pull back to the corresponding deletions at the new level, and the new level only deletes further triples. An allowed triple therefore projects to an allowed triple, giving the required homomorphisms.

We next derive the inequalities controlling all feasible weights, including weights with zero coordinates. For a target D, put $G_j^D := \bigcup_{0 \le s < j} D_2^{(s)}$. The guard blocks in this union are disjoint. Figure 13 shows the excluded blocks when a source pair leaves the active blocks at the second refinement.

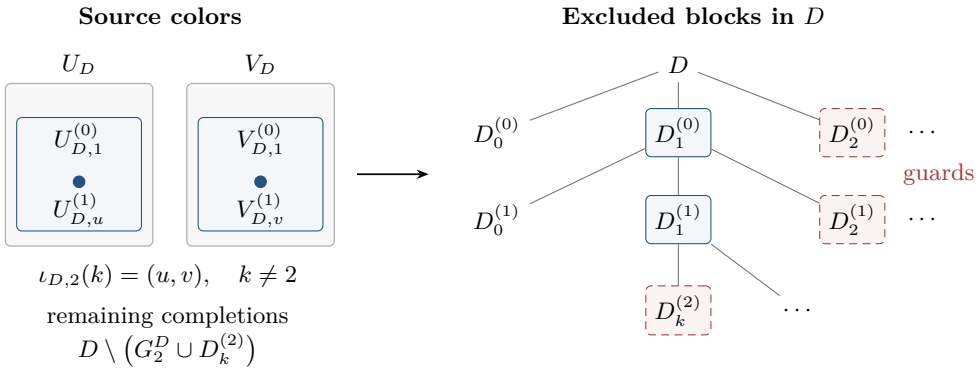


Figure 13: A source pair that leaves the active blocks at refinement 2. It excludes the initial guard $D_2^{(0)}$, the next guard $D_2^{(1)}$, and the child $D_k^{(2)}$ selected by the injection. Since $k \neq 2$, at least one source color leaves child 1, so no later deletion applies. Only child 1 is refined; omitted siblings are indicated by dots.

Lemma 5.8. *Let $\boldsymbol{w}$ be d-feasible with $d > 0$. The following statements hold.*

(i) *Every coarse class has weight at least d.*

(ii) *The bound*

$$w(D_k^{(j)}) \le w(D) - d - w(G_j^D) \tag{5.6}$$

holds for every initial child at $j = 0$. For $j \ge 1$, it holds whenever both source blocks specified by $\iota_{D,j}(k)$ have positive weight. In particular, all bounds at level j hold throughout a copy if its anchor blocks at level $j-1$ have positive weight. Conversely, if all bounds in Equation (5.6) hold, then every assigned pair of supported colors has completion weight at least d.

Proof. For (i), a supported color c in class (i, a) has $N_{P_n(h)}(c, c) = g(i, a)$, so that successor has weight at least d and contains a supported color. Since g cycles through all coarse classes in Equation (5.5), every coarse class has weight at least d.

For (ii), the anchor classes assigned to an initial child have positive weight by (i). A supported color from each gives a pair whose completion set is contained in $D \setminus D_k^{(0)}$, proving the bound at $j = 0$.

At level $j \ge 1$, choose supported colors in the specified source blocks. At every earlier refinement $s < j$ they belong to the active blocks of both source trees, so their block pair is $(1, 1)$. The injection therefore deletes the guard $D_2^{(s)}$. The initial guard was deleted by (T4), and the refinement at level j also deletes $D_k^{(j)}$. The completion set is contained in $D \setminus (G_j^D \cup D_k^{(j)})$, proving the assertion.

Conversely, a pair assigned to an initial child $k \neq 2$ excludes precisely that child. For a pair in U_D, V_D, each refinement matches at most one value of the injection. If its block pair differs from $(1, 1)$, at least one source color is outside the active block and no later deletion applies. If the pair is $(1, 1)$, the deleted child is a guard and further deletions may occur. Thus its excluded set consists of consecutive guards from level zero, followed by at most one further child. The corresponding bound in Equation (5.6) gives excluded weight at most $w(D) - d$. If the last deleted child is also a guard, apply the same bound to that guard. In each case the remaining completion weight is at least d. □

The next optimization turns the child bounds into an exact cost for each tree. Its last assertion measures the extra cost caused by a zero child.

Lemma 5.9. *Consider a tree with initial branching $R \ge 3$ and refinement word $(m_1, \dots, m_n)$, with every $m_i \ge 3$. Suppose nonnegative leaf weights of total weight p satisfy Equation (5.6) for $d > 0$. Put $F_n := \phi_{R-1} \circ \phi_{m_1} \circ \cdots \circ \phi_{m_n}$. Then $p/d \ge 1 + F_n(0)$, with equality attained by strictly positive leaf weights satisfying all the child bounds.*

There are $\lambda_n > 0$ and $0 < b_n < 1$ such that

$$F_n(t) = F_n(0) + \frac{\lambda_n t}{1 - b_n t} \qquad \text{for } 0 \le t \le 1,$$
$$b_0 = \frac{1}{R-1}, \quad \lambda_0 = \frac{1}{(R-1)^2}, \quad b_j = \frac{1}{m_j - b_{j-1}}, \quad \lambda_j = \frac{\lambda_{j-1}}{(m_j - b_{j-1})^2} \quad \text{for } 1 \le j \le n. \tag{5.7}$$

If a child in the last group has zero weight, then $p/d > 1 + F_n(0) + \lambda_n$.

Proof. Consider a group of N children, each of weight at most c, whose guard has weight g. Suppose the active child's descendants have total weight at most $E(c-g)$, where $E > 1$. The active child itself has weight at most c, so the total group weight is at most

$$(N-2)c + g + \min\{c, E(c-g)\} \le \left(N - \frac{1}{E}\right) c.$$

Indeed, $g + \min\{c, E(c-g)\}$ is maximized at $g = c - c/E$, where its value is $2c - c/E$.

For $n \ge 1$, apply this estimate upwards from the terminal group, with

$$E_n \coloneqq m_n, \qquad E_j \coloneqq m_j - \frac{1}{E_{j+1}} \quad \text{for } 1 \le j < n, \qquad E_0 \coloneqq R - \frac{1}{E_1}.$$

All these values exceed one. The child bound at level j is $c_j \coloneqq p - d - w(G_j)$, and subtracting the guard gives the next bound c_{j+1}. Hence $p \le E_0(p-d)$. Since $F_n(0) = 1/(E_0 - 1)$, this proves the stated lower bound. For $n = 0$, take $E_0 = R$ and use the initial child bounds.

For equality, set $p = (1 + F_n(0))d$ and $c_0 = p - d > 0$. In a nonterminal group at level j, give the active child and every ordinary child weight c_j, give the guard weight $c_j - c_j/E_{j+1}$, and continue inside the active child with $c_{j+1} = c_j/E_{j+1}$. At the terminal group give every child weight c_n. The group at level j has weight $E_j c_j$, so the assignments agree with their parent weights. Also $c_{j+1} = c_j - w(D_2^{(j)})$, which gives the required child bounds at every level. All leaf weights are positive.

The representation in Equation (5.7) is immediate at $n = 0$. Substituting $1/(m_j - t)$ into the preceding representation gives the recurrences and $0 < b_j < 1$.

If a terminal child is zero and $n \ge 1$, its group has weight at most $(m_n - 1)c_n$. Repeat the upper-bound calculation with terminal effective size $m_n - 1$. The resulting normalized cost is at least $1 + F_{n-1}(1/(m_n - 1))$, whose excess over $1 + F_n(0)$ is

$$F_{n-1}\left(\frac{1}{m_n - 1}\right) - F_{n-1}\left(\frac{1}{m_n}\right) = \frac{\lambda_{n-1}}{(m_n - 1 - b_{n-1})(m_n - b_{n-1})} > \lambda_n.$$

For $n = 0$, a zero child instead gives excess $1/(R-2) - 1/(R-1) > 1/(R-1)^2 = \lambda_0$. □

Fix one copy of the construction and write its proposed normalized weight as

$$C_n \coloneqq 180(1 + F_{\mathrm{a},n}(0)) + \sum_{z \in [3]} (1 + F_{z,n}(0)), \tag{5.8}$$

using its 3 sink words.

Lemma 5.10. *Every d-feasible weighting with $d > 0$ gives each copy total weight at least $C_n d$, with its own sink words in Equation (5.8).*

Proof. If every anchor block in the copy has positive weight, Lemma 5.8(ii) gives all the child bounds for every target. Summing the bounds in Lemma 5.9 over the 180 anchors and 3 sinks gives $w(\text{copy}) \ge C_n d$.

Otherwise, let j be the first level containing a zero-weight anchor block. Every earlier anchor block has positive weight, so the child bounds through level j hold for every target. Truncate every tree at level j, giving each terminal block the total weight of its descendants. This preserves coarse weights and the child bounds through that level. The chosen zero block is a terminal child of an anchor tree. Let $\lambda_{\mathrm{a},j}, b_{\mathrm{a},j}$ and $\lambda_{z,j}, b_{z,j}$ be the parameters from Equation (5.7). By the last assertion of Lemma 5.9 and the ordinary

bounds for the other trees, $w(\mathrm{copy})/d > C_j + \lambda_{\mathrm{a},j}$. It remains to show that the increase $C_n - C_j$ from the omitted levels is smaller than $\lambda_{\mathrm{a},j}$.

The initial sink-to-anchor coefficient ratio is $(28/9)^2 = 784/81 < 10$. Each anchor denominator in the recurrence is less than 400, whereas each sink denominator is greater than 799. Therefore

$$\frac{\lambda_{z,j}}{\lambda_{\mathrm{a},j}} \le \frac{784}{81}\left(\frac{400}{799}\right)^{2j} < 10. \tag{5.9}$$

A finite negative continued fraction with all digits at least $m \ge 3$ lies in $[0, 1/(m-1)]$, since each corresponding map preserves this interval. If $0 < b < 1$ and $0 \le t \le 1/(m-1)$, then $\lambda t/(1-bt) \le \lambda/(m-2)$. Apply this to the remaining digits after level j in every tree. Anchor tails have digits 400 and sink tails have digits at least 800. By Equation (5.9),

$$C_n - C_j \le \lambda_{\mathrm{a},j}\left(\frac{180}{398} + \frac{30}{798}\right) < \frac{1}{2}\lambda_{\mathrm{a},j} \qquad \text{for } 0 \le j \le n.$$

For the strict inequality, $180/398 < 23/50$ and $30/798 < 1/25$. Thus $w(\mathrm{copy})/d > C_n$. □

All source classes for a target lie in its own copy, so the bound on $C_n - C_j$ is independent of h.

Proof of Proposition 5.7. Summing Lemma 5.10 over all h copies gives

$$1 \ge d\left(180h(1 + F_{\mathrm{a},n}(0)) + \sum_{z\in[3h]}(1 + F_{z,n}(0))\right)$$

for every positive feasible value d. This is the upper bound in Equation (5.4).

For equality, let d be the reciprocal of the expression in parentheses. Give every tree the positive equality weighting from Lemma 5.9, with its prescribed total weight $(1 + F_{D,n}(0))d$. The weights sum to one and satisfy every child bound. Each coarse class has weight at least d.

For two colors in the same coarse class, the completion set is its whole successor class, whose weight is at least d. For colors in different coarse classes of the same top class, an unassigned pair has the entire next top class as completions, and an assigned pair has completion weight at least d by Lemma 5.8(ii). If the two prescribed colors have distinct top indices, all colors of the third top class complete them; that class also has weight at least d. These cases include repeated colors and exhaust the pairs, so the weighting is d-feasible. □

5.3 Continued fractions and an interval

The reciprocal capacities are sums of independently chosen parameters. The following elementary interval argument is sufficient to make these sums fill an interval.

Lemma 5.11. *Let E be the compact set obtained by repeatedly replacing each interval by 2 disjoint closed subintervals whose convex hull is the parent. Suppose each child has length at least c times the parent, where $0 < c < 1/2$. If $(k-1)c \ge 1$, then the sum of k copies of E is $[k \min E, k \max E]$.*

Proof. Fix a target in the latter interval and start with k copies of the convex hull. Repeatedly split a longest current interval, of length L, choosing a child so that the sum of the current intervals still contains the target. This choice is possible: maintain that every current interval has length at least c times the longest. The sum of the other $k-1$ lengths is then at least $(k-1)cL \ge L$, so the 2 possible sum intervals overlap. The invariant survives because both children have length at least cL. Each choice decreases the total length by at least cL, at least c/k times that total. The total tends to zero, so the nested intervals select k points of E summing to the target. □

For $m \ge 3$, put $r_m := (m - \sqrt{m^2-4})/2$, the fixed point of ϕ_m in $(0,1)$, and let $\alpha := r_{801}$, $\beta := r_{800}$, and $I := [\alpha, \beta]$. The maps ϕ_{800}, ϕ_{801} are increasing contractions of I, with disjoint images whose convex hull is I. Their derivative is at most $\beta^2 < 1$, and the gap is nonempty because $\beta - \alpha < 1$. Thus their iterated images determine a compact set K with hull I. Every $x \in K$ has a word satisfying

$$x = \lim_{n\to\infty} \phi_{m_1} \circ \cdots \circ \phi_{m_n}(0), \qquad m_i \in \{800, 801\}. \tag{5.10}$$

Indeed, the maps also preserve $[0,\beta]$, so replacing the terminal point of I by zero changes the nth value by at most $\beta^{2n}\beta$.

Put $f \coloneqq \phi_9 \circ \phi_{800}$, $C \coloneqq f(K)$, and $c_- \coloneqq f(\alpha) < c_+ \coloneqq f(\beta)$. To apply Lemma 5.11 to C, we verify a uniform child-length bound.

Every map $F = f \circ \phi_{m_1} \circ \cdots \circ \phi_{m_n}$ has the form $F(t) = F(0) + \lambda t/(1-bt)$ with $\lambda > 0$ and $0 < b < 1/799$. Initially $f(t) = (800-t)/(7199-9t)$, so $b = 9/7199 < 1/799$. Appending a digit m replaces b by $1/(m-b)$ and λ by $\lambda/(m-b)^2$, preserving the bound. Since $\beta < 1/799$,

$$\frac{\min_{t\in I} F'(t)}{\max_{t\in I} F'(t)} = \left(\frac{1-b\beta}{1-b\alpha}\right)^2 > \left(1-799^{-2}\right)^2 > \frac{1}{2}.$$

Also $|\phi_m(I)| > |I|/801^2$ for $m \in \{800, 801\}$. Integrating the derivative bounds shows that both children $F \circ \phi_m(I)$ have length at least $|F(I)|/(2\cdot 801^2)$. Their hull is $F(I)$. Thus Lemma 5.11 with $c = 1/(2\cdot 801^2)$ shows that the sum of k copies of C is $[kc_-, kc_+]$ for every integer $k \geq k_0 \coloneqq 1 + 2\cdot 801^2$.

Proposition 5.12. *There exists $\varepsilon_0 > 0$ such that every $z \in (0, \varepsilon_0]$ is the limit of $\Gamma(P_n(h))$ for a template sequence from Proposition 5.7. Its projection maps preserve the 3 top classes and the unmodified diagonal successor cycle.*

Proof. Let $a_* \coloneqq \phi_{28}(r_{400})$. For fixed $h \geq 1$ and independent $x_1, \ldots, x_{3h} \in K$, choose their infinite words from Equation (5.10) and use them as sink words in Proposition 5.7. The anchor parameters converge to a_* because $\phi_{400}^{\circ n}(0) \uparrow r_{400}$, and each sink parameter converges to $f(x_j)$. There are finitely many sinks for this fixed h, so the exact capacity formula and Lemma 5.6 give

$$\left(180h(1+a_*) + 3h + \sum_{j\in[3h]} f(x_j)\right)^{-1} \in \mathcal{R}. \tag{5.11}$$

Define

$$D_- \coloneqq 180(1+a_*) + 3(1+c_-), \qquad D_+ \coloneqq 180(1+a_*) + 3(1+c_+),$$
$$H \coloneqq \max\left\{1, \left\lceil \frac{k_0}{3} \right\rceil, \left\lceil \frac{D_-}{D_+ - D_-} \right\rceil\right\}, \qquad \varepsilon_0 \coloneqq \frac{1}{HD_-}.$$

Then $0 < D_- < D_+$. For every $h \geq H$, the sum of $3h$ copies of C is $[3hc_-, 3hc_+]$, so the denominators in Equation (5.11) fill $[hD_-, hD_+]$. Consequently

$$\left[\frac{1}{hD_+}, \frac{1}{hD_-}\right] \subseteq \mathcal{R}.$$

Since $h(D_+ - D_-) \geq D_-$, consecutive intervals overlap. Their endpoints tend to zero, and their union for $h \geq H$ is $(0, \varepsilon_0]$. The empty-relation one-color template supplies zero. Thus $[0, \varepsilon_0] \subseteq \mathcal{R}$.

□

Proof of Theorem 1.2 for $(r, \ell) = (3, 2)$. The interval $[0, \varepsilon_0]$ obtained in Proposition 5.12 lies in $\mathcal{R}$. Applying Lemma 5.5 to this initial interval gives $\mathcal{R} = [0, 1]$. Each target is the capacity limit of a projective sequence, so Lemma 5.6 supplies one forbidden family with both densities equal to that target. Thus $\Sigma_3 = [0, 1]$. □

6 The four-uniform minimum 2-degree case

Theorem 6.1. $\Sigma_4 = [0, 1]$.

Ordinary three-uniform densities first give $[0, 1/8]$. Two-part constructions then extend this interval to $1/2$, after which a suspension operation reaches every density below one.

Proposition 6.2. *For every integer $h \geq 4$ and every monotone class $\mathscr{Q}$ of 3-graphs with $\pi(\mathscr{Q}) = \beta$, there is a monotone class of 4-graphs realizing $2\beta/h^2$ simultaneously for $\gamma_2^{(4)}$ and $\rho_2^{(4)}$. Consequently, $[0, 1/8] \subseteq \Sigma_4$.*

Proof. Define a class of 4-graphs by requiring its members G to admit a cyclic partition $V(G) = V_0 \dot\cup \cdots \dot\cup V_{h-1}$ with the following properties.

(L1) For every i and every $y \in V_{i+1}$, the 3-graph $L_{i,y} \coloneqq \{E \in \binom{V_i}{3} : E \cup \{y\} \in G\}$ belongs to $\mathscr{Q}$.

(L2) Every edge is either of the form $E \cup \{y\}$ with $E \in L_{i,y}$ for some i, or meets four distinct parts.

The source part of an edge of type (L1) is unique, so this class is monotone. It contains every edgeless hypergraph because every empty link belongs to $\mathscr{Q}$.

For the lower bound, split the vertices equally among the h parts. For each i and $y \in V_{i+1}$, choose an independent uniformly random relabeling of a member $H \in \mathscr{Q}$ with edge density $\beta - o(1)$. Apply Lemma 5.1 with $s = 3$, $a = 2$, and $M = |V_{i+1}|$. Taking $\varepsilon = n^{-1/4}$ gives, simultaneously for every pair $S \subseteq V_i$ and every i,

$$d_G(S) = \sum_{y \in V_{i+1}} d_{L_{i,y}}(S) \geq (\beta - o(1))(|V_i| - 2)|V_{i+1}|.$$

After division by $\binom{n-2}{2}$, the normalized degree is at least $2\beta/h^2 - o(1)$.

If a pair meets two distinct parts, then choosing one vertex from each of two further distinct parts produces an edge meeting four distinct parts. Its normalized degree is at least $(h-2)_2/h^2 - o(1) \geq 2/h^2 - o(1) \geq 2\beta/h^2 - o(1)$.

For the upper bound, consider any sequence of witnessing partitions with part sizes $w_0, \dots, w_{h-1}$. After taking a subsequence, assume that $w_i/n \to z_i$ for every i. For each i, double counting the incidences between pairs in V_i and their completions in $V_i \times V_{i+1}$ gives

$$\sum_{S \in \binom{V_i}{2}} d_G(S) = \sum_{y \in V_{i+1}} 3|L_{i,y}|.$$

If the support of (z_i) is proper, choose i with $z_i > 0$ and $z_{i+1} = 0$. The trivial bound $|L_{i,y}| \leq \binom{w_i}{3}$ and the displayed identity bound the whole-part average by $(w_i - 2)w_{i+1} = o(n^2)$. If the support is full, all w_i tend to infinity. For fixed $\varepsilon > 0$, the definition of $\pi(\mathscr{Q}) = \beta$ gives $|L_{i,y}| \leq (\beta + \varepsilon)\binom{w_i}{3}$ for all sufficiently large n. The same identity therefore bounds, for every i, the average degree of pairs in V_i by $(\beta + \varepsilon)(w_i - 2)w_{i+1}$. The normalized density on one whole-part test is at most

$$2(\beta+\varepsilon)\min_i z_i z_{i+1} + o(1) \leq 2(\beta+\varepsilon)\left(\prod_i z_i\right)^{2/h} + o(1) \leq \frac{2(\beta+\varepsilon)}{h^2} + o(1).$$

Every selected part has positive limiting size. Letting $\varepsilon \to 0$ therefore proves the upper bound for $\rho_2^{(4)}$, and hence the simultaneous formula.

By Theorem 1.1(iii), every value in $[4/9, 1]$ is an ordinary three-uniform Turán density. Hence $[8/(9h^2), 2/h^2] \subseteq \Sigma_4$ for every $h \geq 4$. Consecutive intervals overlap because $2/(h+1)^2 \geq 8/(9h^2)$ for $h \geq 2$. Their union is $(0, 1/8]$, and the edgeless class realizes zero. □

Lemma 6.3. *Let $p \in \{2, 3\}$ and $x \in \Sigma_{p+2}$. For the unique $t \geq 0$ satisfying $x + pt + \binom{p}{2}t^2 = 1$, put $T_p(x) = (1+t)^{-p}$. Then $T_p(x) \in \Sigma_{p+2}$. The function T_p is continuous and strictly increasing on $[0, 1]$, with $T_p(0) \leq 1/2$, $T_p(1) = 1$, and $T_p(x) > x$ for $x < 1$.*

Proof. Take a monotone class realizing $x \in \Sigma_{p+2}$. Partition the vertices into A, B, use a member of this class inside A, and allow every cross edge with 1 or 2 vertices in B. For a limiting part ratio $a = |A|/n$, a low-density rectangle in A and the whole-part test on B give the bounds

$$xa^p + pa^{p-1}(1-a) + \binom{p}{2}a^{p-2}(1-a)^2, \qquad a^p.$$

A pair meeting both parts has at least $a^p - o(1)$ normalized completions by taking all remaining vertices in A. A near-extremal member of the original class in A, together with all allowed cross edges, therefore

attains the minimum of the 2 displayed bounds up to $o(1)$. The boundary ratios give x and zero. Hence the resulting class has simultaneous density

$$\max_{t\ge 0} \frac{\min\{1, x+pt+\binom{p}{2}t^2\}}{(1+t)^p}, \qquad t=\frac{1-a}{a}.$$

For $x<1$, let $t_x>0$ solve the equation in the statement. Above t_x the objective is $(1+t)^{-p}$ and decreases. Below t_x, its derivative has numerator $p(1-x)-(p-2)\binom{p}{2}t^2$, which is at least $p^2t_x+2\binom{p}{2}t_x^2>0$. Thus the maximum is $(1+t_x)^{-p}$. The case $x=1$ gives value one.

The root t_x is continuous and strictly decreasing in x. Also $(1+t_0)^p \ge 1+pt_0+\binom{p}{2}t_0^2=2$, so $T_p(0)\le 1/2$. For $x<1$, strict convexity gives $(1+t_x)^{-p}>1-pt_x\ge x$. This proves the remaining assertions. □

Proof of Theorem 6.1. Starting with Proposition 6.2, we extend $[0,1/8]$ to $[0,1/3]$ by a symmetric A^2B^2 construction. Every construction below admits arbitrary subgraphs of its permitted profiles. Let s be a prime power, put $\gamma := 1-1/s$, and let $\mathscr{P}_x$ be a monotone class realizing $x\in\Sigma_4$. A member G of the resulting class admits a partition $A\dot\cup B$ such that $G[A], G[B]\in\mathscr{P}_x$ and every cross edge has profile A^2B^2. For each $S\in\binom{A}{2}$, require the graph on B with edge set $\{T\in\binom{B}{2}: S\cup T\in G\}$ to be s-colorable; symmetrically, for each $T\in\binom{B}{2}$, require the graph on A with edge set $\{S\in\binom{A}{2}: S\cup T\in G\}$ to be s-colorable. The coloring may depend on S or T, respectively; no common coloring of either part is required. We also allow $s=\infty$, meaning $\gamma=1$ and imposing no colorability restriction on the cross links.

Claim 6.4. *The resulting class has density* $\max\{x,(\gamma+x)/4\}$.

Proof. Let the part ratios be a,b, with $a+b=1$. For the complete profile A^2B^2, Equation (2.6) gives normalized pair-degree contributions b^2, $2ab$, and a^2 for pairs in A, crossing A,B, and in B, respectively. Low-density rectangles in $G[A]$ and $G[B]$ give the upper bounds $xa^2+\gamma b^2$ and γa^2+xb^2, respectively, for $\rho_2^{(4)}$: each external rooted graph link has at most a $\gamma+o(1)$ fraction of graph edges. Suppose $a\ge 1/2$ and $x\le\gamma$. The first expression is smaller. As a convex function of $a\in[1/2,1]$, it is at most the larger of its endpoint values $(x+\gamma)/4$ and x. The other half is symmetric. If $x>\gamma$, both expressions are at most $x(a^2+b^2)\le x$.

For the lower bound, placing all vertices in one part gives x. For the balanced construction with finite s, take asymptotically extremal members of $\mathscr{P}_x$ inside A and B, identify both parts with $\mathbb{F}_s^d$, and declare (aa',bb') to be a cross edge if and only if $\langle a-a', b-b'\rangle\neq 0$. For fixed aa', the map $b\mapsto\langle a-a',b\rangle$ is an exactly balanced s-coloring, and its link is the corresponding complete s-partite graph; the symmetric statement holds for fixed bb'. Moreover, each pair meeting both A and B belongs to $(s^d-1)(s^d-s^{d-1})=(\gamma-o(1))|A||B|$ cross edges. For $s=\infty$, use all cross edges. Thus pairs contained in one part have degree $(x+\gamma)/4-o(1)$, while pairs meeting both parts have degree $\gamma/2-o(1)$. If the value from the balanced construction exceeds x, then $x<\gamma/3$, so $\gamma/2>(x+\gamma)/4$. □

By Claim 6.4, the map $x\mapsto\max\{x,(\gamma+x)/4\}$ preserves Σ_4, fixes $\gamma/3$, and strictly increases every $x<\gamma/3$. If $[0,M]\subseteq\Sigma_4$ with $\gamma/4\le M<\gamma/3$, Lemma 2.5 gives $[0,\gamma/3)\subseteq\Sigma_4$. The endpoint $\gamma/3$ follows from Proposition 5.2 with $r=4$, $\ell=2$, $a=6$, $b=4$, and $\beta=\gamma$, since $\gamma(4)_2/6^2=\gamma/3$. Successively take $\gamma=1/2,2/3,3/4,1$, using the exact identities

$$\frac{1}{2}\frac{(4)_2}{6^2}=\frac{1}{6}, \qquad \frac{2}{3}\frac{(4)_2}{6^2}=\frac{2}{9}, \qquad \frac{1}{2}\frac{(9)_2}{12^2}=\frac{1}{4}, \qquad \frac{1}{2}\frac{(25)_2}{30^2}=\frac{1}{3}.$$

Starting from $[0,1/8]$, these transfers give $[0,1/3]\subseteq\Sigma_4$.

To continue toward $1/2$, we use ordinary 3-graph classes as links in the A^3B-profile. Fix a monotone class $\mathscr{Q}_q$ of 3-graphs with ordinary Turán density q and a monotone class $\mathscr{P}_t$ realizing $t\in\Sigma_4$. Use a partition $A\dot\cup B$ with the following allowed profiles:

(C1) $G[B]\in\mathscr{P}_t$ and $G[A]$ is empty;

(C2) all A^2B^2 edges;

(C3) for each $b\in B$, the 3-graph on A formed by the A^3B-edges containing b belongs to $\mathscr{Q}_q$;

(C4) no AB^3 edges.

Claim 6.5. *For $0 \le t \le 1/2$, put $\zeta_q(t) \coloneqq q + \sqrt{q^2 + 1 - t}$ and $\Theta_q(t) \coloneqq (\zeta_q(t)^2 + t)/(1 + \zeta_q(t))^2$. The resulting class has density $\max\{t, \Theta_q(t)\}$.*

Proof. Let $\zeta \coloneqq a/b$ be the ratio of the part sizes. For the upper bound, average all A^3B links over pairs in A and take a low-density rectangle in $G[B]$. After multiplication by $(1+\zeta)^2$, the two normalized internal-pair bounds are $1 + 2q\zeta$ and $\zeta^2 + t$, respectively. They meet at $\zeta = \zeta_q(t)$, and the common normalized value is $\Theta_q(t)$. The boundary case $A = \varnothing$ gives t, while $B = \varnothing$ gives zero.

To complete the upper bound, observe that the first normalized bound is decreasing in ζ. The second decreases up to $\zeta = t$ and then increases, while $\zeta_q(t) \ge \sqrt{1-t} \ge t$ for $t \le 1/2$. Before the intersection the second bound is smaller, and its maximum on $[0, \zeta_q(t)]$ is $\max\{t, \Theta_q(t)\}$. After the intersection the first bound is smaller and is at most $\Theta_q(t)$. Both bounds are positive-volume tests (a whole-part average in A and a low-density rectangle in $G[B]$), so this also proves the $\rho_2^{(4)}$ upper bound.

For the lower bound, take part sizes with limiting ratio $\zeta = \zeta_q(t)$, and place an extremal member of $\mathscr{P}_t$ inside B. Choose a member of $\mathscr{Q}_q$ on A supplied by Lemma 2.2, with edge density and normalized minimum vertex degree at least $q - o(1)$. For each $b \in B$, let L_b be an independent uniformly random relabeling of this 3-graph. For a fixed pair $\{u, v\} \in \binom{A}{2}$,

$$d_G(\{u,v\}) = \binom{|B|}{2} + \sum_{b \in B} d_{L_b}(\{u,v\}).$$

By Lemma 5.1 with $s = 3$, $a = 2$, and $M = |B|$, the sum is at least $(q - o(1))|B|(|A| - 2)$ simultaneously for all pairs in A. Their normalized degrees are therefore at least $(1 + 2q\zeta)/(1+\zeta)^2 - o(1)$. Pairs in B have normalized degree at least $(\zeta^2 + t)/(1+\zeta)^2 - o(1)$ from the minimum degree in $G[B]$ and the complete A^2B^2 profile.

An AB-pair has normalized degree at least $(2\zeta + q\zeta^2)/(1+\zeta)^2 - o(1)$, using the minimum vertex degree in its link. At the intersection of the two internal-pair bounds, the numerator of the difference between this limiting lower bound and $\Theta_q(t)$ is $q\zeta^2 + 2(1-q)\zeta - 1$. For fixed q, this expression is minimized at $t = 1/2$: the quantity $\zeta_q(t)$ decreases with t, whereas $q\zeta^2 + 2(1-q)\zeta - 1$ increases with ζ for $0 \le q \le 1$. At $t = 1/2$, one has $q = (\zeta^2 - 1/2)/(2\zeta)$, and four times the difference numerator is $2\zeta^3 - 4\zeta^2 + 7\zeta - 2$. Its derivative $6\zeta^2 - 8\zeta + 7$ is positive for all real ζ, and at $\zeta = 1/\sqrt{2}$ its value is $4\sqrt{2} - 4 > 0$. Thus every AB-pair has normalized degree at least $\Theta_q(t) - o(1)$. □

The ordinary density $8/27$ is realized by subgraphs of blowups of the 3-graph with edges $012, 013, 023$, since

$$6x_0(x_1x_2 + x_1x_3 + x_2x_3) \le 2x_0(1-x_0)^2 \le \frac{8}{27},$$

with equality at $x_0 = 1/3$ and $x_1 = x_2 = x_3 = 2/9$. Balanced complete 4-partite 3-graphs realize $3/8$. We use these 2 values of q to reach $3/8$.

For $q = 8/27$, the root $\zeta_q(0)$ is less than $27/20$, since $z^2 - 2qz - 1$ is positive at $z = 27/20$. Thus $\Theta_q(0) < 729/2209 < 1/3$. At $t = 1/3$, evaluating the 2 internal-pair bounds at $\zeta = 7/6$ gives

$$\max\{t, \Theta_q(t)\} \ge \min\left\{\frac{548}{1521}, \frac{61}{169}\right\} = \frac{548}{1521} > \frac{9}{25}.$$

The density map is continuous, so its image of $[0, 1/3]$ fills the gap up to $9/25$. For $q = 3/8$, one has $\zeta_q(0) = (3 + \sqrt{73})/8 < 3/2$, so $\Theta_q(0) < 9/25$. At $t = 9/25$ and $\zeta = 5/4$, the same calculation gives

$$\max\{t, \Theta_q(t)\} \ge \min\left\{\frac{31}{81}, \frac{769}{2025}\right\} = \frac{769}{2025} > \frac{3}{8}.$$

Continuity now gives $[0, 3/8] \subseteq \Sigma_4$.

Finally, fix $t = 0$ and let q vary through the ordinary density interval $[4/9, 1]$ from Theorem 1.1(iii). The function $\Theta_q(0)$ is continuous in q, with $\Theta_1(0) = 1/2$. Since $\sqrt{97} < 10$, we have $\zeta_{4/9}(0) < 14/9$ and hence $\Theta_{4/9}(0) < 196/529 < 3/8$. These values fill the gap to $1/2$.

Apply Lemma 6.3 with $p = 2$. The map $T_2(x) = 1/(2-x)$ and Lemma 2.5 extend $[0, 1/2]$ to $[0, 1)$. The empty forbidden family realizes one. □

7 Minimum 2-degrees in higher uniformities

The preceding section proves the four-uniform base case. We now lift this interval result to every larger uniformity.

Theorem 7.1. *For every integer $r \geq 5$, one has $\Sigma_r = [0,1]$.*

Throughout this section put $p \coloneqq r-2$, so $p \geq 3$. Thus a minimum 2-degree is normalized by $\binom{n-2}{p}$. The constructions below use the profile convention from Section 5; in every case arbitrary subgraphs of the allowed edge set are admitted, so the resulting classes are monotone.

7.1 General transfer operations

We begin with a two-part construction that transfers densities from the four-uniform case.

Let $\mathscr{P}_x$ be a monotone class of 4-graphs with $\gamma_2^{(4)}(\mathscr{P}_x) = \rho_2^{(4)}(\mathscr{P}_x) = x$. When $x = 1$, we take $\mathscr{P}_1$ to be the class of all 4-graphs. Define a class $\mathscr{L}_p(\mathscr{P}_x)$ of $(p+2)$-graphs. A member admits a partition $V = X \dot\cup Y$ and an auxiliary 4-graph $H \in \mathscr{P}_x$ on X. An edge is allowed if it satisfies one of the following conditions.

(F1) It has profile X^4Y^{p-2}, and its four vertices in X form an edge of H.

(F2) It has profile XY^{p+1}.

Define the relevant scaling constants by

$$\kappa_p \coloneqq \binom{p}{2}\left(\frac{2}{p}\right)^2\left(\frac{p-2}{p}\right)^{p-2} = \frac{2(p-1)}{p}\left(\frac{p-2}{p}\right)^{p-2}, \qquad x_p \coloneqq \frac{(p-2)^2}{2p(p-1)}.$$

Lemma 7.2. *Suppose that $p \geq 3$ and $0 \leq x \leq \min\{1, 3x_p\}$, and let $\mathscr{P}_x$ be as above. Then $\gamma_2^{(p+2)}(\mathscr{L}_p(\mathscr{P}_x)) = \rho_2^{(p+2)}(\mathscr{L}_p(\mathscr{P}_x)) = \kappa_p x$.*

Consequently, $[0, \kappa_p \min\{1, 3x_p\}] \subseteq \Sigma_{p+2}$. In particular, $[0, ((p-2)/p)^p] \subseteq \Sigma_{p+2}$ for every $p \geq 3$; when $3x_p \leq 1$, the larger interval $[0, 3((p-2)/p)^p]$ is also contained in Σ_{p+2}.

Proof of Lemma 7.2. Let $|X| = (a+o(1))n$ and $|Y| = (b+o(1))n$, with $a+b=1$. For every pair $u, v \in X$, its completions satisfy $d_G(uv) \leq d_H(uv)\binom{|Y|}{p-2}$. A low-density rectangle in the common 4-graph H therefore gives normalized average at most $x\binom{p}{2}a^2b^{p-2} + o(1) \leq \kappa_p x + o(1)$. If a part is sublinear, a whole positive part instead gives density zero. This proves the local-density upper bound, including boundary supports.

For the lower bound, set $a = 2/p$ and $b = (p-2)/p$, choose a near-extremal $H \in \mathscr{P}_x$ on X, and include every edge allowed by this witness. An XX-pair has normalized degree $\kappa_p x - o(1)$. For an XY-pair uv, the edges of type (F2) contribute b^p. Moreover, $3d_H(u) = \sum_{w \in X\setminus\{u\}} d_H(uw) \geq (|X|-1)\delta_2(H)$, so the edges of type (F1) contribute $x\binom{p}{3}a^3b^{p-3} = (2/3)\kappa_p x$ asymptotically. Therefore every XY-pair has normalized degree at least $b^p + (2/3)\kappa_p x - o(1)$. Since $b^p = \kappa_p x_p$, this is at least $\kappa_p x - o(1)$ whenever $x \leq 3x_p$.

For a YY-pair, the edges of type (F2) contribute $pab^{p-1} = 2b^{p-1}$. When $p \geq 4$, the edges of type (F1) add $x\binom{p}{4}a^4b^{p-4}$; for $p = 3$ this term is absent. At the extreme value $x = 3x_p$, direct simplification gives

$$pab^{p-1} + 3x_p\binom{p}{4}a^4b^{p-4} = 3\frac{p-1}{p}b^{p-1} > 3b^p = \kappa_p(3x_p)$$

for $p \geq 4$, while for $p = 3$ the contribution from type (F2) is $2/9 > 1/9 = \kappa_3(3x_3)$. The required lower bound for YY-pairs is immediate at $x = 0$, and the difference between their normalized degree and $\kappa_p x$ is affine in x. Hence this bound holds throughout the stated interval and, together with the preceding XX- and XY-pair estimates, proves the matching lower bound. □

We next construct overlapping intervals extending to one. Fix $a \geq p+2$ and a monotone class $\mathscr{P}_x$ realizing $x \in \Sigma_4$. Partition the vertices into a parts $V_0, \ldots, V_{a-1}$ and choose one auxiliary 4-graph $H_i \in \mathscr{P}_x$ on each V_i. Allow the following edge profiles and arbitrary subgraphs of them.

(i) Every transversal edge meeting $p+2$ distinct parts is allowed.

(ii) An edge may meet one part in 2 or 3 vertices and meet each remaining part in at most one vertex. Every edge of these profiles is allowed.

(iii) An edge may consist of an edge of H_i and one vertex in each of $p-2$ distinct other parts. The same H_i is used for every outside transversal with source part V_i.

No other profiles are allowed. Denote this monotone class by $\mathscr{T}_{a,p}(\mathscr{P}_x)$. For integers $q \geq p-1$, put

$$U_{q,p} \coloneqq \frac{(q)_p}{q^p}, \qquad W_{q,p} \coloneqq \binom{p}{2}\frac{(q-1)_{p-2}}{q^p}.$$

For a finite index set J and a vector $\boldsymbol{z} = (z_i)_{i\in J}$, let $e_d(\boldsymbol{z}) \coloneqq \sum_{I\in\binom{J}{d}} \prod_{i\in I} z_i$, with $e_0(\boldsymbol{z}) = 1$, and write $\boldsymbol{z}_I$ for its restriction to I.

Lemma 7.3. *Let $0 \leq x < 1$ and let $\mathscr{P}_x$ be as above. Then $\gamma_2^{(p+2)}\big(\mathscr{T}_{a,p}(\mathscr{P}_x)\big) = \rho_2^{(p+2)}\big(\mathscr{T}_{a,p}(\mathscr{P}_x)\big) = U_{a,p} + xW_{a,p}$. Consequently, these classes realize $[U_{a,p}, U_{a,p}+W_{a,p})$. Moreover, $U_{a,p}+W_{a,p} \geq U_{a+1,p}$.*

Proof. Pass to a subsequence along which $|V_i|/n \to z_i$ for every i, let J be the positive support of $\boldsymbol{z}$, and put $q = |J|$. Parts outside J contribute $o(n^p)$ completions. For a pair contained in V_i, the profiles with 2 or 3 vertices in that part have normalized contribution at most

$$p!\left(e_p(\boldsymbol{z}_{J\setminus\{i\}}) + z_i e_{p-1}(\boldsymbol{z}_{J\setminus\{i\}})\right) = p!e_p(\boldsymbol{z}_J).$$

On a low-density rectangle in H_i, the four-uniform profile adds at most $\frac{p!}{2}xz_i^2 e_{p-2}(\boldsymbol{z}_{J\setminus\{i\}}) + o(1)$. If $q \leq p-2$, both terms are zero. Otherwise choose $i \in J$ with $z_i \leq 1/q$ and take a low-density rectangle in the common auxiliary 4-graph H_i. Maclaurin's inequality applied to all q positive weights gives $p!e_p(\boldsymbol{z}_J) \leq U_{q,p}$. Applied to the other $q-1$ weights, it also gives

$$e_{p-2}(\boldsymbol{z}_{J\setminus\{i\}}) \leq \binom{q-1}{p-2}\left(\frac{1-z_i}{q-1}\right)^{p-2}.$$

The function $t^2(1-t)^{p-2}$ is increasing for $0 \leq t \leq 2/p$, and $1/q \leq 2/p$ when $q \geq p-1$. Hence the average over that rectangle is at most $U_{q,p} + xW_{q,p} + o(1)$.

It remains to compare different support sizes. For $q \geq p$, writing $y = 1/q$ gives

$$U_{q,p} + xW_{q,p} = \prod_{j=1}^{p-2}(1-jy)\left(1-(p-1)y + x\binom{p}{2}y^2\right).$$

Every factor in the product is nonincreasing and nonnegative on $y \in (0, 1/p]$. The derivative of the last factor is $(p-1)(xpy-1) \leq 0$, so the whole expression is nonincreasing in y and hence nondecreasing in q. The remaining comparison from $q = p-1$ to $q = p$ is linear in x. At $x = 0$ it is immediate. At $x = 1$ it is equivalent to $(p/(p-1))^p \leq p+1$, and $(p/(p-1))^p \leq 27/8 \leq p+1$ because $(1-1/p)^p$ is increasing for $p \geq 3$ and equals $8/27$ at $p = 3$. We conclude that $U_{q,p} + xW_{q,p} \leq U_{a,p} + xW_{a,p}$.

The common auxiliary rectangle makes this an upper bound for $\rho_2^{(p+2)}$ on every support. Equal parts and near-extremal 4-graphs $H_i \in \mathscr{P}_x$ give normalized degree $U_{a,p} + xW_{a,p} - o(1)$ for pairs contained in one part. For a pair meeting 2 parts, the complete transversal profiles and those with 2 or 3 vertices in one part contribute

$$\frac{(a-2)_p + \frac{p(p+3)}{2}(a-2)_{p-1} + \frac{p(p-1)(p+4)}{6}(a-2)_{p-2}}{a^p} + o(1).$$

Writing $d \coloneqq a - p \geq 2$, the main term exceeds $U_{a,p}$ by $\frac{p(p-1)(3d+p-2)}{6}\frac{(a-2)_{p-2}}{a^p}$. Dividing this surplus by $W_{a,p}$ gives $\frac{(d+1)(3d+p-2)}{3(p+d-1)} \geq 1$. Since $x < 1$, pairs meeting two parts therefore have normalized degree at least $U_{a,p} + xW_{a,p} - o(1)$.

For the overlap, factor $(a-1)_{p-2}/a^{p-2}$ from both sides. The inequality $U_{a,p} + W_{a,p} \geq U_{a+1,p}$ follows from

$$\left(\frac{a}{a+1}\right)^{p-1} \leq 1 - \frac{p-1}{a+1} + \frac{\binom{p-1}{2}}{(a+1)^2} \leq 1 - \frac{p-1}{a} + \frac{\binom{p}{2}}{a^2}.$$

The first inequality follows by truncating the alternating binomial expansion after its quadratic term: the omitted terms alternate with decreasing absolute values because $a \geq p+2$. The second reduces to $a/(a+1)+(p-2)(a/(a+1))^2/2 \leq p/2$. This holds because its left side is at most $1+(p-2)/2 = p/2$. □

Since $U_{a,p} + W_{a,p} \geq U_{a+1,p}$ and $U_{a,p} \uparrow 1$, these interpolation intervals cover $[U_{p+2,p}, 1)$. If consecutive endpoints coincide, the common point belongs to the next interval as its left endpoint.

7.2 Uniformities at least seven

The general transfer intervals overlap in this range.

Proof of Theorem 7.1 for $r \geq 7$. First suppose that $p = r-2 \geq 6$. The initial interval from Lemma 7.2 reaches $((p-2)/p)^p$, while the interpolation intervals start at $U_{p+2,p} = (p+1)!/(2(p+2)^{p-1})$. To compare these endpoints, set $R_p := 2(p-2)^p(p+2)^{p-1}/(p^p(p+1)!)$. A direct calculation gives $R_6 = 262144/229635 > 1$. Furthermore,

$$\frac{R_{p+1}}{R_p} = \frac{p-1}{p+1}\left(1+\frac{p^2+p+4}{(p-2)(p+1)(p+2)}\right)^p \geq \frac{p-1}{p+1}\left(1+\frac{1}{p}\right)^p \geq \frac{2(p-1)}{p+1} > 1.$$

The first inequality uses the lower bound $1/p$ for the fraction inside the parentheses, and the second uses $(1+1/p)^p \geq 1+p(1/p) = 2$. Thus $R_p \geq R_6 > 1$, so the initial interval overlaps the interpolation intervals. The interpolation intervals from Lemma 7.3 then extend the realized interval to 1.

It remains to consider $p = 5$. The improved part of Lemma 7.2 gives $[0, 3(3/5)^5] = [0, 729/3125] \subseteq \Sigma_7$. The interpolation intervals start at $U_{7,5} = 6!/(2 \cdot 7^4) = 360/2401$. Since $729/3125 - 360/2401 = 625329/7503125 > 0$, the initial interval reaches beyond this endpoint. The interpolation intervals again cover every remaining density below 1. In both cases the empty forbidden family realizes 1.

□

7.3 The exceptional uniformities five and six

For $r = 6$, an additional two-part construction connects the initial interval to the interpolation intervals. For $r = 5$, we first reach $1/2$ and then iterate a suspension operation.

For positive integers $a_1, \ldots, a_m$ summing to r, the class allowing only the profile $V_1^{a_1} \cdots V_m^{a_m}$ has ordinary density

$$\max_{\substack{z_i \geq 0 \\ \sum_i z_i = 1}} \frac{r!}{\prod_i a_i!} \prod_i z_i^{a_i} = \frac{r!}{r^r} \prod_i \frac{a_i^{a_i}}{a_i!}.$$

Weighted AM–GM gives the maximum at $z_i = a_i/r$. In particular, the four-uniform profiles A^2BC and A^2B^2 realize $3/16$ and $3/8$, respectively.

Proof of Theorem 7.1 for $r = 6$. Now $p = 4$. The improved part of Lemma 7.2 gives $[0, 3/16] \subseteq \Sigma_6$, while the interpolation intervals start at $U_{6,4} = 5/18$. We cover this gap by a two-part construction.

Let $\mathscr{P}_x$ be a monotone class of 4-graphs realizing $x \in \Sigma_4$, and let $\mathscr{Q}_\theta$ be a monotone class of 4-graphs with ordinary Turán density θ. Define a class as follows. A member admits a partition $V = X \dot\cup Y$, an auxiliary 4-graph $H \in \mathscr{P}_x$ on X, and a 4-graph $L_S \in \mathscr{Q}_\theta$ on Y for every $S \in \binom{X}{2}$. Its edges may have the following forms:

(i) an edge of H together with two vertices of Y;

(ii) $S \cup F$, where $S \in \binom{X}{2}$ and $F \in L_S$;

(iii) one vertex of X together with five vertices of Y.

For $\zeta \in (1/2, 1)$, which will be the limiting ratio $|X|/|Y|$, put $T_\theta(\zeta) := \theta/((1-\zeta)(1+\zeta)^3)$.

Claim 7.4. *Let $0 < \theta < 1$ and $\zeta \in (1/2, 1)$, and put $x := \theta/(3\zeta(1-\zeta))$. Suppose that $x < 1$, $(1-\zeta)(1+\zeta)^3 \le 1$, and*

$$8\zeta^2\theta - 9\zeta\theta + 3\zeta + 3\theta - 3 \le 0, \tag{7.1}$$

$$17\zeta^3\theta - 18\zeta^2\theta + 12\zeta^2 + 3\zeta\theta - 12\zeta + 3\theta \le 0. \tag{7.2}$$

Then the resulting class has density $T_\theta(\zeta)$.

Proof. Consider an arbitrary sequence of members of the resulting class. If $|X| = o(n)$, every pair in Y has degree at most $|X|\binom{|Y|-2}{3} + \binom{|X|}{2}\binom{|Y|-2}{2} + \binom{|X|}{4} = o(n^4)$. If $|Y| = o(n)$, every pair in X has degree at most $\binom{|X|-2}{2}\binom{|Y|}{2} + \binom{|Y|}{4} = o(n^4)$. Thus neither case in which one part has sublinear size gives a positive asymptotic minimum 2-degree.

Suppose now that both parts have linear size, and pass to a subsequence such that $|X|/|Y| \to s \in (0, \infty)$. Take a low-density rectangle in the common auxiliary 4-graph H. The external ordinary-link contribution is bounded uniformly for every root, so its average normalized degree is at most $(\theta + 6xs^2)/(1+s)^4 + o(1)$. The derivative of this rational function is $4(3xs(1-s) - \theta)/(1+s)^5$. By the choice of x, its stationary points are $1-\zeta$ and ζ; the first is a local minimum and the second a local maximum. Its boundary values are θ and 0, while its value at ζ is $T_\theta(\zeta)$. The assumption $(1-\zeta)(1+\zeta)^3 \le 1$ gives $T_\theta(\zeta) \ge \theta$, so $T_\theta(\zeta)$ is the global maximum and gives the required upper bound.

For the lower bound, choose arbitrarily large $H \in \mathscr{P}_x$ on h vertices with $\delta_2(H) \ge (x - o(1))\binom{h-2}{2}$, and take $m = (1/\zeta + o(1))h$ vertices in Y. Choose $L \in \mathscr{Q}_\theta$ on Y with $|L| \ge (\theta - o(1))\binom{m}{4}$, let the 4-graphs L_S be independent uniformly random relabelings of L, and include every allowed edge determined by H and the 4-graphs L_S.

For an XY-pair, apply Lemma 5.1 to the $h-1$ links indexed by $\{u, v\}$ with $v \ne u$, using 4-uniform links and $a = 1$. For a YY-pair, use all $\binom{h}{2}$ links with $a = 2$. Taking $\varepsilon = n^{-1/4}$ and a further union bound over $u \in X$ makes both conclusions hold simultaneously. The last inequality in the lemma bounds the normalized vertex degrees and edge density of H from below by its normalized minimum 2-degree. After multiplication by $(1+\zeta)^4$, the resulting lower bounds for XX-, XY-, and YY-pairs are, respectively, $\theta + 6x\zeta^2$, $1 + 4\theta\zeta + 4x\zeta^3$, and $4\zeta + 6\theta\zeta^2 + x\zeta^4$. The first normalized bound equals $T_\theta(\zeta)$. After substituting $x = \theta/(3\zeta(1-\zeta))$, the other two bounds are at least this value precisely when (7.1) and (7.2), respectively, hold. This proves the matching lower bound. □

Take $\theta = 3/16$, realized by the profile A^2BC. Let $\zeta_0 \in (1/2, 1)$ be the unique solution of $(1-\zeta_0)(1+\zeta_0)^3 = 1$. The left side decreases for $\zeta > 1/2$, and its values at $5/6$ and $6/7$ are $1 + 35/1296$ and $1 - 204/2401$, respectively. Thus $5/6 < \zeta_0 < 6/7$.

For $\zeta_0 \le \zeta \le 10/11$, one has $x = 1/(16\zeta(1-\zeta)) \le 121/160 < 1$. At $\theta = 3/16$, the derivatives of the left sides of Equations (7.1) and (7.2) are

$$\frac{3}{16}(16\zeta + 7) \qquad \text{and} \qquad \frac{3}{16}(51\zeta^2 + 92\zeta - 61).$$

Both are positive for $\zeta \ge 5/6$. At $\zeta = 10/11$, the original expressions equal $-9/1936$ and $-6651/21296$, so the claim applies throughout this interval. Since $T'_\theta(\zeta)/T_\theta(\zeta) = 2(2\zeta - 1)/(1 - \zeta^2) > 0$, its image is

$$\left[\frac{3}{16}, \frac{14641}{49392}\right], \qquad \frac{14641}{49392} - \frac{5}{18} = \frac{307}{16464} > 0.$$

This connects the initial interval to the interpolation intervals in Lemma 7.3. They cover every remaining density below one, and the empty forbidden family realizes one.

□

Proof of Theorem 7.1 for $r = 5$. The case $p = 3$ of Lemma 7.2 gives $[0, 1/9] \subseteq \Sigma_5$. We next use a two-part construction to reach $1/2$.

For any $x \in [0, 1)$, let $\mathscr{P}_x$ be a monotone class of 4-graphs realizing $x \in \Sigma_4$, which is available by Theorem 6.1. Let $\mathscr{Q}_\theta$ and $\mathscr{R}_g$ be monotone classes of 3-graphs and 4-graphs with ordinary Turán densities θ and g, respectively. Define a class of 5-graphs whose members admit a partition $V = X \dot\cup Y$ and the following witnesses.

(i) There is a 4-graph $H \in \mathscr{P}_x$ on X, and every edge of profile X^4Y has its four vertices in X forming an edge of H.

(ii) For every $S \in \binom{X}{2}$, the 3-graph $L_S \coloneqq \{T \in \binom{Y}{3} : S \cup T \in G\}$ belongs to $\mathscr{Q}_\theta$.

(iii) For every $u \in X$, the 4-graph $R_u \coloneqq \{Q \in \binom{Y}{4} : \{u\} \cup Q \in G\}$ belongs to $\mathscr{R}_g$.

(iv) No other profiles are allowed.

The three nonempty profiles are disjoint.

For the density, define the two normalized internal-pair degree functions, for $\zeta \geq 0$, by

$$\varphi_x(\zeta) \coloneqq \frac{\theta + 3x\zeta^2}{(1+\zeta)^3}, \qquad \text{and} \qquad \psi_{\theta,g}(\zeta) \coloneqq \frac{3g\zeta + 3\theta\zeta^2}{(1+\zeta)^3}.$$

Their minimum extends continuously to $\zeta \in [0,\infty]$, with value zero at both endpoints. Define $\Phi_{\theta,g}$ and $P_{\theta,g}$ by

$$\Phi_{\theta,g}(x) \coloneqq \max_{\zeta \geq 0} \min\{\varphi_x(\zeta), \psi_{\theta,g}(\zeta)\}, \qquad P_{\theta,g}(\zeta) \coloneqq 3\theta\zeta^3 + (3g - 9\theta)\zeta^2 + (8\theta - 9g)\zeta + 3g.$$

We show that the mixed-pair lower bound is at least $\min\{\varphi_x(\zeta), \psi_{\theta,g}(\zeta)\}$ whenever $P_{\theta,g}$ is nonnegative on $[0,3]$.

Claim 7.5. *If $P_{\theta,g}(\zeta) \geq 0$ for $0 \leq \zeta \leq 3$, then the resulting class has density $\Phi_{\theta,g}(x)$.*

Proof. We first prove the upper bound. Consider a sequence of members of the resulting class and witnessing partitions. Suppose that both parts have linear size, and pass to a subsequence such that $|X|/|Y| \to \zeta \in (0,\infty)$. Take a low-density rectangle in the common auxiliary 4-graph H. Every root S in this rectangle has contributions only from the profile X^4Y determined by H and its link L_S. The ordinary bound for $\mathscr{Q}_\theta$ is uniform in S, so Equation (5.2) gives normalized average at most $\varphi_x(\zeta) + o(1)$ on that rectangle.

For a pair in Y, average over all pairs in $\binom{Y}{2}$. For every $S \in \binom{X}{2}$ and $u \in X$, respectively, one has $\sum_{T \in \binom{Y}{2}} d_{L_S}(T) = 3|L_S|$ and $\sum_{T \in \binom{Y}{2}} d_{R_u}(T) = 6|R_u|$. The ordinary Turán density upper bounds for $\mathscr{Q}_\theta$ and $\mathscr{R}_g$ consequently bound the whole-part average in Y by $\psi_{\theta,g}(\zeta) + o(1)$. Both estimates hold on positive-volume tests, so they bound $\rho_2^{(5)}$ by $\min\{\varphi_x(\zeta), \psi_{\theta,g}(\zeta)\} + o(1)$.

The boundary cases cannot increase the upper bound obtained from positive limiting part ratios. If $|X| = o(n)$, the same average over pairs in Y, now using the trivial complete-link bounds, bounds the whole-part average by $O(|X|^2|Y| + |X||Y|^2) = o(n^3)$. If $|Y| = o(n)$, any pair in X has at most $O(|Y||X|^2 + |Y|^3) = o(n^3)$ completions. Maximizing over the limiting ratio proves the upper bound $\Phi_{\theta,g}(x)$.

For the lower bound, fix $\zeta \in (0,\infty)$. Take $|Y|$ along the unbounded sequence supplied by Lemma 2.2 for $\mathscr{R}_g$, and put $|X| = \lfloor \zeta|Y| \rfloor$. On this sequence choose $R \in \mathscr{R}_g$ with $|R| \geq (g - o(1))\binom{|Y|}{4}$ and $\delta_1(R) \geq (g - o(1))\binom{|Y|-1}{3}$. The defining limits of the density parameters also supply $L \in \mathscr{Q}_\theta$ on the same vertex set Y with $|L| \geq (\theta - o(1))\binom{|Y|}{3}$, and $H \in \mathscr{P}_x$ on X with $\delta_2(H) \geq (x - o(1))\binom{|X|-2}{2}$. Take independent uniformly random relabelings of L and R as the links L_S and R_u, respectively, and include every edge of profile X^4Y determined by H.

For an XX-pair, relabeling preserves $|L_S|$, so the lower bound $\varphi_x(\zeta) - o(1)$ is deterministic. For an XY-pair, apply Lemma 5.1 to its $|X| - 1$ independent 3-links with $a = 1$; its 4-link contribution uses the minimum vertex degree of R. For a YY-pair, apply the lemma to the $\binom{|X|}{2}$ independent 3-links and the $|X|$ independent 4-links, both with $a = 2$. Taking $\varepsilon = n^{-1/4}$ and a union bound over the vertices of X makes these estimates simultaneous. The lemma also bounds the normalized vertex degrees of H from below by its normalized minimum 2-degree. Thus the normalized lower bounds for XX-, XY-, and YY-pairs are, respectively, $\varphi_x(\zeta)$, $(g + 3\theta\zeta + x\zeta^3)/(1+\zeta)^3$, and $\psi_{\theta,g}(\zeta)$, up to $o(1)$.

It remains to check the XY-pair bound. For $0 \leq \zeta \leq 3$,

$$\min\{\theta + 3x\zeta^2, 3g\zeta + 3\theta\zeta^2\} \leq \frac{\zeta}{3}(\theta + 3x\zeta^2) + \left(1 - \frac{\zeta}{3}\right)(3g\zeta + 3\theta\zeta^2),$$

and three times the difference between the XY-pair numerator and the right side is $P_{\theta,g}(\zeta)$. For $\zeta \geq 3$, the difference between the XY-pair numerator and $\theta+3x\zeta^2$ is $g-\theta+3\theta\zeta+x\zeta^2(\zeta-3) \geq 0$. The hypothesis on $P_{\theta,g}$ therefore makes the degree of every XY-pair no smaller than the minimum of the two internal-pair bounds. Maximizing over ζ gives the matching lower bound for the density of the resulting class. □

For the endpoint comparisons, differentiation gives

$$\psi'_{\theta,g}(\zeta) = \frac{3\big(g+2(\theta-g)\zeta-\theta\zeta^2\big)}{(1+\zeta)^4}.$$

The quadratic numerator inside the parentheses is concave. Its positivity on any interval $[0,z]$ therefore follows from positivity at zero and at z.

We first use $(\theta,g)=(8/27,3/32)$ to reach $3/20$. The first density was realized in Section 6; the second is the density of the complete 4-partite 4-graph class. The cubic Bernstein coefficients of $P_{\theta,g}$ on $[0,3/2]$ and $[3/2,3]$ are, respectively,

$$\left(\frac{9}{32},\frac{1805}{1728},\frac{65}{3456},\frac{235}{1152}\right) \quad \text{and} \quad \left(\frac{235}{1152},\frac{1345}{3456},\frac{3085}{1728},\frac{2129}{288}\right).$$

By Equation (2.7), $P_{8/27,3/32} \geq 65/3456 > 0$ on $[0,3]$. At $x=0$, split the range of ζ at $2/5$. The function φ_0 decreases, and $\psi_{8/27,3/32}$ increases on $[0,2/5]$, since the quadratic in the derivative equals $g>0$ at zero and $4501/21600>0$ at $2/5$. Also $\psi_{8/27,3/32}(2/5) < \varphi_0(2/5)$. Hence

$$\Phi_{8/27,3/32}(0) \leq \frac{1000}{9261} < \frac{1}{9}, \qquad \Phi_{8/27,3/32}(1) \geq \min\{\varphi_1,\psi_{8/27,3/32}\}\left(\frac{3}{2}\right) = \frac{31}{200} > \frac{3}{20}.$$

The substitution $u=\zeta/(1+\zeta)$ makes the objective a continuous function on the compact square $0\leq x,u\leq 1$, so $\Phi_{\theta,g}$ is continuous in x. Some $x_+<1$ therefore still gives a value above $3/20$, and the image of $[0,x_+]$ contains $[1/9,3/20]$.

Next let $0\leq s\leq 1$, take $\theta_s = 4/9+5s/9$ from the ordinary interval in Theorem 1.1(iii), and set $x_s=3s/4$ and $g=3/32$. The identity

$$P_{\theta,3/32}(\zeta) - P_{8/27,3/32}(\zeta) = \left(\theta-\frac{8}{27}\right)\zeta\left[3\left(\zeta-\frac{3}{2}\right)^2+\frac{5}{4}\right]$$

shows that the claim applies to every $\theta\in[4/9,1]$. The density $\Phi_{\theta_s,3/32}(x_s)$ is continuous in s by the same compactness argument. At $s=0$, split at $\zeta=1/2$. The function $\psi_{4/9,3/32}$ is increasing up to that point, because the quadratic in its derivative equals $g>0$ at zero and $1/3$ at $1/2$. Thus

$$\Phi_{4/9,3/32}(0) \leq \max\left\{\frac{91}{648},\frac{32}{243}\right\} < \frac{3}{20}.$$

At $s=1$, evaluating at $\zeta=1$ gives $\Phi_{1,3/32}(3/4) \geq \min\{13/32,105/256\} = 13/32 > 3/8$. The image of this continuous family contains $[3/20,3/8]$.

Finally take $(\theta,g)=(1,3/8)$, using the profile A^2B^2 for g. For every $\zeta\geq 0$,

$$8P_{1,3/8}(\zeta) = 24\zeta\left(\zeta-\frac{3}{2}\right)^2 + 9\left(\zeta-\frac{17}{18}\right)^2 + \frac{35}{36} > 0.$$

Splitting the range at $\zeta=2/5$ as above, where the derivative quadratic has endpoint values $3/8$ and $143/200$, gives

$$\Phi_{1,3/8}(0) \leq \max\left\{\frac{465}{1372},\frac{125}{343}\right\} < \frac{3}{8}.$$

At $x=1$ and $\zeta=8/9$, both internal-pair bounds equal $2457/4913 = 1/2+1/9826$. Continuity gives an $x_+<1$ with $\Phi_{1,3/8}(x_+)>1/2$, so the image of $[0,x_+]$ contains $[3/8,1/2]$. Together with the preceding intervals this proves $[0,1/2]\subseteq\Sigma_5$.

Apply Lemma 6.3 with $p=3$. Since $T_3(0)\leq 1/2$, $T_3(1)=1$, and $T_3(x)>x$ for $x<1$, Lemma 2.5 extends $[0,1/2]$ to $[0,1)$. The empty forbidden family realizes one. □

The rational identities and Bernstein coefficients in the four-, five-, and six-uniform cases are checked by `verify_degree_intervals.py`.

Completion of the proof of Theorem 1.2. The cyclic-link argument proves the assertion for $\ell \geq 3$. For $\ell = 2$, Section 5 gives $\Sigma_3 = [0, 1]$, and the constructions in Sections 6 and 7 give $\Sigma_r = [0, 1]$ for every $r \geq 4$. Each class is monotone, so its finite nonmembers give the required forbidden family. □

Part II

Uniform Turán densities for complete layers

8 Palette tools and finite-family realization

Fix $r \geq 3$. At the highest nontrivial level $r-2$, the results separate into finite attainment below $4r^{-r}$ and compatible constructions above it. We first develop the common palette operations and finite-family realization. The compactness argument is in Section 9, and the upper interval is constructed in Section 10. Recall that $\Pi^{(r)}_{r-2,\infty}$ is the set of uniform Turán densities at level $r-2$ and $\mathcal{D}_r$ is the set of finite-palette Lagrangians.

8.1 Palette operations and characterization

We extend the palette language of the introduction to countable color sets. Thus a k-coordinate palette is a relation $P \subseteq C(P)^k$ on a nonempty finite or countable color set, with ordered tuples that may repeat colors. The probability simplex $S(C)$ now also allows countable C. The polynomial and Lagrangian are

$$\Lambda_P(\boldsymbol{x}) \coloneqq \sum_{(c_1,\dots,c_k)\in P} \prod_{a\in[k]} x_{c_a}, \qquad \lambda(P) \coloneqq \sup_{\boldsymbol{x}\in S(C(P))} \Lambda_P(\boldsymbol{x}).$$

There is no factorial factor. The maximum is attained when $C(P)$ is finite. For countable $C(P)$, the sum converges absolutely because it is at most 1. Write $P[D] = P \cap D^k$. If $\boldsymbol{x}(D) = 1-\delta > 0$, let $\boldsymbol{x}^D$ be the normalized restriction to D. Splitting the admissible tuples according to whether they meet $C(P)\setminus D$ gives

$$\Lambda_P(\boldsymbol{x}) = (1-\delta)^k \Lambda_{P[D]}(\boldsymbol{x}^D) + R, \qquad 0 \leq R \leq 1-(1-\delta)^k.$$

Consequently,

$$|\Lambda_P(\boldsymbol{x}) - \Lambda_{P[D]}(\boldsymbol{x}^D)| \leq 1-(1-\delta)^k \leq k\delta. \tag{8.1}$$

Increasing finite D to $C(P)$ therefore proves

$$\lambda(P) = \sup\{\lambda(P[D])\colon \varnothing \neq D \subseteq C(P),\ |D| < \infty\}.$$

A *palette homomorphism* $P \to Q$ is a map of their color sets that preserves every admissible tuple coordinatewise. A *subpalette* is obtained by deleting colors and admissible tuples. Pushing weights forward under a palette homomorphism $P \to Q$ proves $\lambda(P) \leq \lambda(Q)$. A sequence $P_0 \leftarrow P_1 \leftarrow \cdots$ with such maps is *projective*.

The *relative Lagrangian* of an r-palette P is $r^r\lambda(P)$. By homogeneity, it equals the supremum of Λ_P over nonnegative weightings of total weight r. A palette consisting of one tuple of r distinct colors has relative Lagrangian 1, the minimum positive value. For a simple r-graph H, its edge polynomial is $\sum_{e\in H}\prod_{v\in e} x_v$.

Using the facet-coloring convention fixed in the introduction, write

$$\mathrm{Col}(P) \coloneqq \{H\colon H \text{ is an unordered } P\text{-colorable } r\text{-graph}\}.$$

For the reverse palette $\mathrm{rev}(P)$, reverse every admissible tuple. Reversing the vertex order gives $\mathrm{Col}(P) = \mathrm{Col}(\mathrm{rev}(P))$, and the palette polynomials are equal. A homomorphism $P \to Q$ also gives $\mathrm{Col}(P) \subseteq \mathrm{Col}(Q)$.

The following characterization and separation theorem allows us to work with finite palettes.

Theorem 8.1 ([24, Theorems 1.4 and 1.6])**.** *The following statements hold.*

(i) *For any family $\mathcal{F}$ of r-graphs, $\pi^{(r)}_{r-2}(\mathcal{F}) = \sup\{\lambda(P)\colon P$ is a finite $\mathcal{F}$-deficient r-palette$\}$.*

(ii) *For finite r-palettes P, Q, there is an r-graph in $\mathrm{Col}(P)\setminus\mathrm{Col}(Q)$ if and only if neither $P \to Q$ nor $P \to \mathrm{rev}(Q)$.*

The cited characterization uses the uniform weighting on palette colors. To obtain the weighted formulation, first consider a rational probability weighting $x_c = n_c/N$ and replace each color c by n_c clones, declaring a

tuple admissible exactly when its projection belongs to P. The resulting palette has density $\Lambda_P(\boldsymbol{x})$ under the uniform weighting and maps to P, so it preserves every deficiency condition of P. Approximating arbitrary probability weightings by rational ones proves Theorem 8.1(i).

Fact 8.2. *We have* $\mathcal{D}_r \subseteq \Pi^{(r)}_{r-2,\infty} \subseteq \{0\} \cup [r^{-r}, 1]$.

Proof. For a finite palette P, let $\mathcal{F}_{r,r-1}(P)$ consist of all finite r-graphs outside $\mathrm{Col}(P)$. The palette P is $\mathcal{F}_{r,r-1}(P)$-deficient. If Q is also $\mathcal{F}_{r,r-1}(P)$-deficient, then $\mathrm{Col}(Q) \subseteq \mathrm{Col}(P)$, since any Q-colorable r-graph outside $\mathrm{Col}(P)$ would belong to $\mathcal{F}_{r,r-1}(P)$. Hence Theorem 8.1(ii) gives $Q \to P$ or $Q \to \mathrm{rev}(P)$, so $\lambda(Q) \le \lambda(P)$. Theorem 8.1(i) proves $\pi^{(r)}_{r-2}(\mathcal{F}_{r,r-1}(P)) = \lambda(P)$.

To prove the second inclusion, suppose that $P \neq \varnothing$ and choose a tuple whose distinct colors have multiplicities $a_1, \dots, a_t$. The polynomial Λ_P is at least the corresponding monomial. Maximizing that monomial gives

$$\lambda(P) \ge \prod_{i=1}^{t} \left(\frac{a_i}{r}\right)^{a_i} = r^{-r} \prod_{i=1}^{t} a_i^{a_i} \ge r^{-r}. \tag{8.2}$$

Since $P \subseteq C(P)^r$, we have $\Lambda_P(\boldsymbol{x}) \le (\sum_{c \in C(P)} x_c)^r = 1$ for every $\boldsymbol{x} \in S(C(P))$, so $\lambda(P) \le 1$. Thus every finite palette has Lagrangian either zero or in $[r^{-r}, 1]$. By Theorem 8.1(i), $\pi^{(r)}_{r-2}(\mathcal{F})$ is the supremum of the Lagrangians of finite $\mathcal{F}$-deficient palettes. If this supremum is positive, at least one such palette has positive Lagrangian, and hence $r^{-r} \le \pi^{(r)}_{r-2}(\mathcal{F}) \le 1$. □

As recalled in Section 3, the normalized Lagrangians $\lambda(H)$ of finite s-graphs belong to $\Pi^{(s)}_{\infty}$ and are dense in it, by Brown and Simonovits [4]; see also Pikhurko [40, Theorem 24].

For palettes P, Q on disjoint color sets, let $P \star Q$ consist of all concatenations of their admissible tuples. For $r \ge k \ge 3$, let $\mathrm{lift}_r(P)$ be obtained from a k-palette by appending $r-k$ new distinct fixed colors to every tuple.

Fact 8.3. *The following statements hold.*

(i) *If H is a finite simple k-graph and $1 \le m \le k!$, selecting exactly m permutations of each edge gives a palette $P_m(H)$ with $\lambda(P_m(H)) = (m/k!)\lambda(H)$.*

(ii) *If $r \ge k \ge 3$, then $\lambda(\mathrm{lift}_r(P)) = (k^k/r^r)\lambda(P)$.*

(iii) *If P, Q have k, s coordinates, respectively, then $\lambda(P \star Q) = \frac{k^k s^s}{(k+s)^{k+s}} \lambda(P)\lambda(Q)$.*

The last 2 identities hold for finite or countable palettes, and both operations preserve homomorphisms.

Proof. The polynomial in (i) is m times the edge polynomial of H. For (ii), write u for the total weight on the old colors. For fixed u and new-color weights x_{z_i}, the supremum is $\lambda(P)u^k \prod_i x_{z_i}$, where $u + \sum_i x_{z_i} = 1$. Weighted AM–GM gives $u = k/r$ and $x_{z_i} = 1/r$. For (iii), fixing total weights $u, 1-u$ on the two color sets gives $\lambda(P)\lambda(Q)u^k(1-u)^s$, maximized at $u = k/(k+s)$. Finite restrictions and renormalization prove the countable cases. Homomorphisms act separately on the old color sets and fix the appended colors. □

A palette is *rainbow* if every admissible tuple has pairwise distinct colors. For a rainbow r-palette P, let $H(P)$ be its underlying simple r-graph, and let $m_e \in \{1, \dots, r!\}$ count the admissible permutations of each edge $e \in H(P)$. Then

$$\Lambda_P(\boldsymbol{x}) = \sum_{e \in H(P)} m_e \prod_{v \in e} x_v. \tag{8.3}$$

A repeated-color tuple has Lagrangian at least $4r^{-r}$ by (8.2). Hence every palette of smaller Lagrangian is rainbow. Weighting the colors of one underlying edge equally also shows that such an edge has at most 3 admissible permutations.

Proposition 8.4. *Every value of $\Pi^{(r)}_{r-2,\infty}$ is the limit of Lagrangians of finite r-palettes P_n with homomorphisms $P_{n+1} \to P_n$.*

Its proof combines finite attainment in Section 9 with the compatible upper-interval constructions in Section 10. Conversely, finite-palette realization and Equation (2.2) turn every such chain into a forbidden family of the limiting density.

8.2 Realization by finite forbidden families

We extend the finite-family realization theorem of King, Piga, Sales, and Schülke [17, Theorem 1.1] from 3-palettes to r-palettes for every $r \geq 3$. Following their rigidity and counting argument [17, Sections 5–6], we combine testability of hereditary properties of directed hypergraphs with the palette characterization and separation in Theorem 8.1.

Write $\Pi^{(r)}_{r-2,\mathrm{fin}}$ for the set of values $\pi^{(r)}_{r-2}(\mathcal{F})$ attained by finite families $\mathcal{F}$ of r-graphs.

Theorem 8.5. *For every $r \geq 3$ and every finite r-palette P_0, there is a finite family $\mathcal{G}$ of r-graphs with $\pi^{(r)}_{r-2}(\mathcal{G}) = \lambda(P_0)$. Consequently, $\mathcal{D}_r \subseteq \Pi^{(r)}_{r-2,\mathrm{fin}}$.*

We prove that finitely many local homomorphism constraints bound the Lagrangian by that of a fixed target palette. Fix $r \geq 3$. For a finite r-palette Q, write $d(Q) \coloneqq |Q|/|C(Q)|^r$ for its density under the uniform weighting. A *complete blowup* S of P on a nonempty finite set V is specified by a map $\phi\colon V \to C(P)$, with admissible tuples

$$S \coloneqq \{(v_1, \dots, v_r) \in V^r \colon (\phi(v_1), \dots, \phi(v_r)) \in P\}.$$

The map ϕ need not be surjective. If $C(P) = [m]$ and $V_i \coloneqq \phi^{-1}(i)$, then $d(S) = \Lambda_P((|V_i|/|V|)_{i\in[m]}) \leq \lambda(P)$. All tuple counts in this section are ordered and include tuples with repeated colors. Each tuple is counted once.

Call P *Lagrangian-minimal* if $\lambda(P[J]) < \lambda(P)$ for every nonempty proper subset $J \subsetneq C(P)$. Every finite palette has a Lagrangian-minimal induced subpalette with the same Lagrangian.

For a finite r-palette P and a positive integer k, let $\mathcal{C}_k(P)$ be the class of finite r-palettes Q satisfying $Q[J] \to P$ or $Q[J] \to \mathrm{rev}(P)$ for every nonempty $J \subseteq C(Q)$ with $|J| \leq k$. The choice between P and $\mathrm{rev}(P)$ may depend on J. Moreover, this class is closed under restriction, deletion of tuples, and taking complete blowups. For the last assertion, project at most k colors of a blowup to their image in Q, and compose with a homomorphism for that image.

Proposition 8.6. *Let P be a Lagrangian-minimal r-palette with $0 < \lambda(P) < 1$. There are positive integers k, N and a constant $\delta > 0$, depending only on P, such that every $Q \in \mathcal{C}_k(P)$ on $n \geq N$ colors with $d(Q) \geq \lambda(P) - \delta$ admits a complete blowup S of P or $\mathrm{rev}(P)$ on $C(Q)$ satisfying*

$$d(Q) \leq d(S) \leq \lambda(P), \qquad |Q \triangle S| \leq \frac{5}{3}\big(d(S) - d(Q)\big)n^r \leq \frac{5}{3}\big(\lambda(P) - d(Q)\big)n^r.$$

Every $Q \in \mathcal{C}_k(P)$, regardless of its number of colors or the value of $d(Q)$, satisfies $\lambda(Q) \leq \lambda(P)$.

Proof. Label the colors of P by $[m]$.

Claim 8.7. *The following statements hold.*

(i) *Every endomorphism of P is an automorphism, and $P \to \mathrm{rev}(P)$ implies $P \cong \mathrm{rev}(P)$.*

(ii) *For each $i \neq j$, there is a tuple in P and a position occupied by i such that replacing that occurrence of i by j gives a tuple outside P.*

(iii) *There are $\beta, \varepsilon_0 > 0$ such that $y_i \geq \beta$ for every $i \in [m]$ whenever $\boldsymbol{y} \in S([m])$ and $\Lambda_P(\boldsymbol{y}) \geq \lambda(P) - \varepsilon_0$.*

Proof of the claim. For (i), a nonsurjective endomorphism has image $J \subsetneq [m]$. Monotonicity of the Lagrangian under homomorphisms would give $\lambda(P) \leq \lambda(P[J]) < \lambda(P)$, a contradiction. A bijective endomorphism permutes the finite admissible set and is therefore an automorphism. Since $\Lambda_{\mathrm{rev}(P)[J]} = \Lambda_{P[J]}$, the same argument makes any homomorphism $P \to \mathrm{rev}(P)$ surjective. The two palettes have the same number of admissible tuples, so it is an isomorphism.

For (ii), suppose that replacing any one occurrence of i by j always preserves admissibility. Replacing all occurrences in succession then shows that the map sending i to j and fixing all other colors is an endomorphism. This contradicts (i).

For (iii), $0 < \lambda(P) < 1$ implies $m \geq 2$. By minimality,

$$g \coloneqq \lambda(P) - \max_{c \in [m]} \lambda(P[C(P) \setminus \{c\}]) > 0.$$

Deleting color c in Equation (8.1) gives $\Lambda_P(\boldsymbol{y}) \leq \lambda(P) - g + ry_c$ when $y_c < 1$; the same inequality is automatic when $y_c = 1$. Thus $\Lambda_P(\boldsymbol{y}) \geq \lambda(P) - g/2$ implies $y_c \geq g/(2r)$ for every c. Take $\varepsilon_0 = g/2$ and $\beta = g/(2r)$. □

Take pairwise disjoint sets $U_1, \ldots, U_m$, each of size rm, and let R_* be the rainbow part of the complete P-blowup on these parts. We use R_* as the root of the extensions below and write $h \coloneqq |C(R_*)| = rm^2$. Thus a tuple of distinct colors belongs to R_* exactly when the tuple of its part indices belongs to P. For $u \in C(R_*)$, a *one-color extension at* u retains R_*, adds one new color z, and adds all tuples obtained by replacing u by z in an admissible tuple of R_*. Every added tuple contains z exactly once.

Claim 8.8. *Every homomorphism $\theta\colon R_* \to P$ is constant on each U_i, and these constant values define an automorphism of P. If $u \in U_i$, then in every homomorphism from the one-color extension at u to P, the new color has the same image as U_i. If $P \ncong \mathrm{rev}(P)$, then $R_* \nrightarrow \mathrm{rev}(P)$.*

Proof of the claim. For each i, at least r members of U_i have the same image a_i. Every admissible tuple of P, including one with repeated part indices, can be lifted to distinct representatives among these members. Consequently, $i \mapsto a_i$ is an endomorphism of P, and hence an automorphism by Claim 8.7(i). Postcomposing θ with the inverse of this automorphism, assume that $a_i = i$ for all i.

Suppose that $u \in U_i$ has image $j \neq i$. Take an arbitrary admissible tuple containing i, and lift one occurrence of i to u. Lift the other positions to distinct representatives from the subsets of U_ℓ whose colors map to ℓ. Their images show that replacing this occurrence of i by j preserves admissibility. Since the tuple and occurrence were arbitrary, this contradicts Claim 8.7(ii). Thus θ is constant on each U_i.

For a one-color extension at $u \in U_i$, postcompose the homomorphism with an automorphism of P so that every color in U_ℓ maps to ℓ for each $\ell \in [m]$. If the new color has image $j \neq i$, lift any admissible tuple containing i by using u at one chosen occurrence and distinct original colors at the other positions, avoiding u. Replace that occurrence of u by the new color. The resulting tuple belongs to the extension and again contradicts Claim 8.7(ii). There are enough original colors even when the target tuple repeats i, since $|U_i| \geq r$.

Finally, the same choice of representatives applied to a homomorphism $R_* \to \mathrm{rev}(P)$ gives $P \to \mathrm{rev}(P)$. Claim 8.7(i) then gives the last assertion. □

Claim 8.9. *For every $\eta > 0$, there are positive integers k_0, N_0 such that every $Q \in \mathcal{C}_{k_0}(P)$ on $n \geq N_0$ colors admits a complete blowup S of P or $\mathrm{rev}(P)$ on $C(Q)$ with $|Q \setminus S| \leq \eta n^r$.*

Proof of the claim. Delete all tuples of Q with a repeated color, at a cost of at most $\binom{r}{2} n^{r-1}$ tuples. Regard the remaining palette as a directed r-graph, with its relation defined on ordered tuples of distinct vertices. The property of admitting a homomorphism to P or to $\mathrm{rev}(P)$ is invariant under relabeling and preserved under taking induced subgraphs. Notice that the source tuple has distinct colors, but its image in P may repeat colors.

The testability theorem of Austin and Tao [1, Theorem 1.5 in the arXiv version] supplies k_0, N_0 such that, when every restriction to at most k_0 vertices has this property, the directed r-graph differs in at most $(\eta/2)n^r$ ordered tuples from one that has the property globally. Their edit distance counts unordered r-sets on which the induced structure changes; each such set supports at most $r!$ ordered tuples, so tolerance $\eta/2$ gives at most $(\eta/2)r!\binom{n}{r} \leq (\eta/2)n^r$ ordered-tuple changes.

The defining condition for $\mathcal{C}_{k_0}(P)$ supplies the required local property. Let R be the modified directed r-graph, and let S be the complete blowup defined by a homomorphism from R to P or $\mathrm{rev}(P)$. Since $R \subseteq S$, the rainbow tuples of Q lying outside S are among the edits. Increasing N_0 so that $\binom{r}{2} n^{r-1} \leq (\eta/2)n^r$ proves the claim. □

For a palette Q on n colors and a complete P-blowup S on the same color set, put $A \coloneqq S \setminus Q$ and $B \coloneqq Q \setminus S$. It suffices to prove $|A| \geq 4|B|$, since then $|S| - |Q| = |A| - |B| \geq 0$ and $|A| + |B| \leq \frac{5}{3}(|S| - |Q|)$. We count pairs $(a, b) \in A \times B$ for which a is rainbow and the tuples share exactly 1 color. If, for some

$c > 0$, every b has at least cn^{r-1} partners and every a has at most $(c/4)n^{r-1}$ partners, double counting gives the required inequality.

The root has at most h^r tuples, so its union with at most r one-color extensions has at most $(r+1)h^r$ tuples. Take β, ε_0 from Claim 8.7(iii), and put $c := (\beta/2)^{r-1}/((2(r+1)h^r)$ and $\gamma := c/(4r)$.

Put $D_0 := rm^{r-1}$ and $p_0 := (\gamma/(4D_0))^m$. Choose $\delta > 0$ so that $2\delta < \varepsilon_0$ and $2h^r\delta(2/\beta)^r < \min\{p_0/2, 1/2\}$. Apply Claim 8.9 with $\eta = \delta$, and take $k \geq \max\{k_0, h+r\}$. Finally, choose an integer $N \geq \max\{N_0, k\}$ such that

$$\beta N \geq 2(h+r), \qquad r\binom{r}{2} \leq \frac{\gamma N}{2}, \quad \text{and} \quad (r-1)h \leq \frac{\gamma N}{4D_0}. \tag{8.4}$$

Let $Q \in \mathcal{C}_k(P)$ have $n \geq N$ colors and satisfy $d(Q) \geq \lambda(P) - \delta$. Claim 8.9 supplies a complete blowup of P or $\mathrm{rev}(P)$. Replacing P by $\mathrm{rev}(P)$ if necessary, assume that the target of this blowup is P. We have $\Lambda_{\mathrm{rev}(P)} = \Lambda_P$, and the preceding constants are unchanged. Among all complete P-blowups on $C(Q)$, choose S minimizing $|Q \setminus S|$. Use the sets A, B defined above and write the parts of S as $V_1, \ldots, V_m$. Then $|B| \leq \delta n^r$ and $|A| = |B| + |S| - |Q| \leq 2\delta n^r$. Moreover, we have $d(S) \geq d(Q) - |B|/n^r \geq \lambda(P) - 2\delta$. Hence Claim 8.7(iii) gives $|V_i| \geq \beta n$ for every $i \in [m]$.

We will repeatedly choose a uniformly random injective map $\psi\colon C(R_*) \to C(Q)$ with $\psi(U_i) \subseteq V_i$, avoiding a prescribed set of at most r colors. For any fixed ordered tuple of $t \leq r$ distinct root colors, the probability of a specified image is at most $(\beta n/2)^{-t}$. Indeed, the denominator is a product of falling factorials, and each factor is at least $\beta n - h - r \geq \beta n/2$. Since $|R_*| \leq h^r$, the probability that some root tuple maps to A is at most

$$h^r|A|\left(\frac{\beta n}{2}\right)^{-r} \leq 2h^r\delta\left(\frac{2}{\beta}\right)^r. \tag{8.5}$$

For $x \in C(Q)$, define $d_B(x) := |\{(v_1, \ldots, v_r) \in B\colon x \in \{v_1, \ldots, v_r\}\}|$. Thus each tuple is counted once even if x occurs in several positions.

Claim 8.10. *For every $x \in C(Q)$, we have $d_B(x) \leq \gamma n^{r-1}$.*

Proof of the claim. Suppose that $d_B(x) > \gamma n^{r-1}$. For each $i \in [m]$, change only the part assigned to x, placing it in V_i. Let $B_i(x)$ be the tuples of Q containing x that lie outside the resulting blowup. The choice of S gives $|B_i(x)| \geq d_B(x) > \gamma n^{r-1}$. There are at most $r\binom{r}{2}n^{r-2}$ tuples containing x and a repeated color. This bound follows by specifying a position occupied by x and a pair of positions with equal entries. By (8.4), at least $(\gamma/2)n^{r-1}$ members of each $B_i(x)$ are rainbow.

For each $i \in [m]$, classify the rainbow tuples in $B_i(x)$ by the position of x and the part indices of the other $r-1$ entries. There are at most D_0 types. Hence, for every i, choose a set $\mathcal{B}_i \subseteq B_i(x)$ of one type with $|\mathcal{B}_i| \geq (\gamma/(2D_0))n^{r-1}$. The corresponding tuple of part indices lies outside P when x is assigned to part i. Reserve $r-1$ root colors in the corresponding parts U_j for this type, using disjoint reserved colors for different i. There are $(r-1)m$ reserved colors in total, while each U_j has rm colors.

Choose ψ as above, avoiding x, and expose these reserved groups in order. Suppose that the previously exposed images form a set W. Then $|W| \leq h$, and at most $(r-1)|W|n^{r-2}$ members of $\mathcal{B}_i$ meet W. By (8.4), at least $(\gamma/(4D_0))n^{r-1}$ candidates remain. Conditional on the previously exposed images, the current group is uniformly distributed over its permissible ordered images, of which there are at most n^{r-1}. Its conditional probability of giving a tuple in $\mathcal{B}_i$, together with x, is therefore at least $\gamma/(4D_0)$. This bound holds for every previous exposure, so all m groups succeed with probability at least p_0.

By (8.5) and the choice of δ, the probability that some admissible tuple of R_* is mapped outside Q is less than $p_0/2$. Thus some such injection maps every admissible tuple of R_* into Q, and the images of the m reserved groups give tuples in $\mathcal{B}_1, \ldots, \mathcal{B}_m$, respectively, when combined with x.

Consider the subpalette of Q on $\psi(C(R_*)) \cup \{x\}$ consisting of the image of R_* and the m selected tuples from the sets $\mathcal{B}_i$. It has $h+1 \leq k$ colors but admits no homomorphism to P. Indeed, Claim 8.8 fixes the root labels up to an automorphism. After postcomposing a homomorphism with the inverse of the induced automorphism of P, each possible image i of x is excluded by the selected tuple from $\mathcal{B}_i$. There is no homomorphism to $\mathrm{rev}(P)$ either. If the targets are not isomorphic, the root already excludes such a homomorphism; otherwise, compose it with an isomorphism to P. Any homomorphism from the induced palette on these colors would also be a homomorphism from this subpalette. This contradicts $Q \in \mathcal{C}_k(P)$. □

We next give a lower bound on the number of missing tuples associated with each $b \in B$. For $b \in B$, let X be the set of colors occurring in b, and let $A_1(b)$ consist of the rainbow tuples in A containing exactly 1 color from X.

Claim 8.11. *For every $b \in B$, we have $|A_1(b)| \geq cn^{r-1}$.*

Proof of the claim. Write $X = \{x_1, \ldots, x_\ell\}$, where $\ell \leq r$. For each j, choose an original root color in U_i where $x_j \in V_i$, and form its one-color extension with new color z_j. Take the union of these extensions with their common root R_*. Do not add other tuples containing 2 distinct new colors. Finally, add the tuple obtained from b by replacing every occurrence of x_j by z_j.

Call the resulting palette T_b. It has at most $h + r \leq k$ colors. By Claim 8.8, any homomorphism $T_b \to P$ can be postcomposed with an automorphism of P so that each U_i maps to i; it must then send z_j to the part label of x_j. The image of the tuple corresponding to b therefore lies outside P, since $b \notin S$. As in the preceding claim, there is no homomorphism to $\mathrm{rev}(P)$. Consequently, there is no injective homomorphism $T_b \to Q$. Indeed, its image J would have at most k colors, so $Q \in \mathcal{C}_k(P)$ would give $Q[J] \to P$ or $Q[J] \to \mathrm{rev}(P)$. Composing with $T_b \to Q[J]$ would contradict the preceding exclusions.

Fix $z_j \mapsto x_j$ and choose a uniformly random injection of the root mapping each U_i into V_i, avoiding X. The tuple corresponding to b maps to $b \in Q$. Every other tuple of T_b maps into S and is rainbow; its image meets X in 0 or 1 color. Since no such injection is a homomorphism, the image of at least one of these tuples lies outside Q. By (8.5), with probability greater than $1/2$ every admissible tuple of R_* is mapped into Q. Thus with probability at least $1/2$ some tuple of T_b maps to $A_1(b)$.

Besides the tuple corresponding to b, T_b has at most $(r+1)h^r$ tuples. Fix $a \in A_1(b)$ and a tuple of T_b containing one new color. Its new color has a prescribed image, and its other $r-1$ colors are distinct root colors. The probability that its image is a is at most $(\beta n/2)^{-(r-1)}$. A union bound gives $\frac{1}{2} \leq |A_1(b)|(r+1)h^r(\beta n/2)^{-(r-1)}$, which proves the claim. The construction also applies when b repeats colors, since it uses one new color for each member of X, with repeated occurrences using that same new color. All the other added tuples still contain exactly 1 new color. □

By Claims 8.10 and 8.11, double counting the pairs (a, b) with $b \in B$ and $a \in A_1(b)$ gives

$$cn^{r-1}|B| \leq \sum_{b \in B} |A_1(b)| \leq r\gamma n^{r-1}|A|.$$

For the upper bound, each tuple $a \in A$ has at most r colors, and at most γn^{r-1} tuples of B contain any one of them. Thus $|A| \geq 4|B|$, so $|S| - |Q| = |A| - |B| \geq 0$ and we have

$$|Q \triangle S| = |A| + |B| \leq \frac{5}{3}(|A| - |B|) = \frac{5}{3}\big(d(S) - d(Q)\big)n^r.$$

Since $d(S) \leq \lambda(P)$, this proves the stability assertion.

To prove $\lambda(Q) \leq \lambda(P)$ for every $Q \in \mathcal{C}_k(P)$, suppose instead that some such Q has $\lambda(Q) > \lambda(P)$. Choose a rational probability weighting $\boldsymbol{x} \in S(C(Q))$ with $\Lambda_Q(\boldsymbol{x}) > \lambda(P)$. Complete blowups of Q in these proportions, on arbitrarily many colors, belong to $\mathcal{C}_k(P)$ and have density $\Lambda_Q(\boldsymbol{x})$ under the uniform weighting. These blowups satisfy the hypotheses of the density assertion just proved, which bounds their densities by $\lambda(P)$, a contradiction. □

Proof of Theorem 8.5. If $\lambda(P_0) = 0$, forbid a single-edge r-graph. If $\lambda(P_0) = 1$, use the empty forbidden family. Otherwise choose a Lagrangian-minimal induced subpalette $P \subseteq P_0$ with $\lambda(P) = \lambda(P_0)$, and take k from Proposition 8.6.

Up to relabeling, there are only finitely many palettes R on at most k colors with $R \nrightarrow P$ and $R \nrightarrow \mathrm{rev}(P)$. For each one, Theorem 8.1(ii) gives an r-graph $F_R \in \mathrm{Col}(R) \setminus \mathrm{Col}(P)$. Let $\mathcal{G}$ consist of these witnesses. It is finite, and P is $\mathcal{G}$-deficient. Each F_R has an edge, since every edgeless r-graph is P-colorable.

If Q is a finite $\mathcal{G}$-deficient palette, every induced subpalette on at most k colors maps to P or to $\mathrm{rev}(P)$. Otherwise that subpalette colors its corresponding F_R, and hence so does Q. Thus $Q \in \mathcal{C}_k(P)$, and Proposition 8.6 gives $\lambda(Q) \leq \lambda(P)$. The characterization in Theorem 8.1(i), together with the deficiency of P, now gives $\pi^{(r)}_{r-2}(\mathcal{G}) = \lambda(P) = \lambda(P_0)$. □

In the case $0 < \lambda(P_0) < 1$, the realizing family can be chosen as a finite truncation of the canonical forbidden family of the minimal target P. Choose an integer M at least the number of vertices of every member of $\mathcal{G}$, and let $\mathcal{F}_M(P)$ consist of one representative of each isomorphism class of r-graphs F with $|V(F)| \le M$ and $F \notin \mathrm{Col}(P)$. This finite family contains $\mathcal{G}$, while P and its blowups remain deficient. Hence $\pi^{(r)}_{r-2}(\mathcal{F}_M(P)) = \lambda(P)$.

9 Finite attainment and nonclosedness

9.1 Ladders and finite certificates

We prove the finite-attainment assertions in Theorem 1.3(i) by excluding ladders of a fixed length from the deficient palettes and then taking a limit of optimal weightings.

Definition 9.1. An *m-ladder* in a simple r-graph H consists of distinct vertices $b_1, \dots, b_m$ and $(r-2)$-sets $A_1, \dots, A_{m-1}$ such that $A_i \cap \{b_i, b_{i+1}, \dots, b_m\} = \varnothing$ for every $i \in [m-1]$ and $A_i \cup \{b_i, b_j\} \in H$ for every $1 \le i < j \le m$. The sets A_i may intersect one another and may contain earlier vertices b_j with $j < i$. We call the vertices b_i the *spine vertices* and the sets A_i the *cores*. A rainbow palette contains an m-ladder if its underlying hypergraph does.

Figure 14 gives an example in which two cores coincide and another contains an earlier spine vertex.

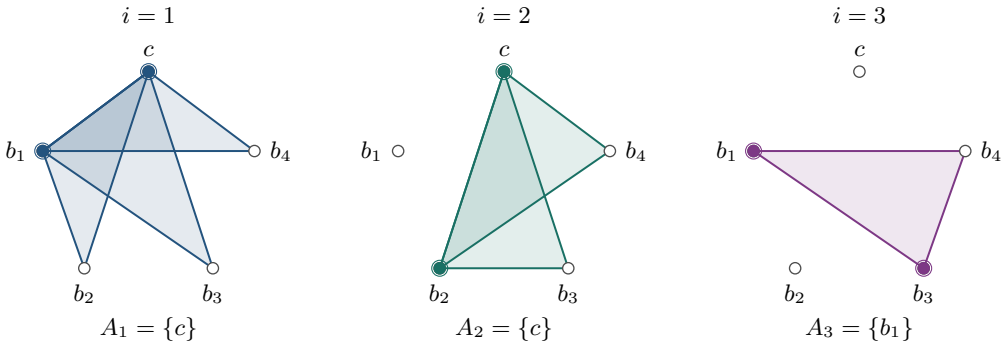


Figure 14: A 4-ladder in a 3-graph, drawn on the same five vertices in each panel. The panels show the required edges for $i = 1, 2, 3$. Each circled pair is $A_i \cup \{b_i\}$ and is completed by every later spine vertex.

Every pair of spine vertices b_i, b_j is contained in an edge. Consequently, if a rainbow palette containing an m-ladder maps to another rainbow palette, its image still contains an m-ladder. Indeed, the spine vertices remain distinct, and every required disjointness relation is witnessed inside a mapped edge.

The following compactness theorem is the main step. Its conclusion preserves every deficiency condition shared by the approximating palettes.

Theorem 9.2. *Fix $r \ge 3$ and an integer $M \ge 2$. Let P_n be finite rainbow r-palettes without M-ladders. If $\lim_{n\to\infty} \lambda(P_n) = t$, then there are a finite palette P and a subsequence (P_{n_j}) such that $\lambda(P) = t$ and $P \to P_{n_j}$ for every sufficiently large j.*

We first obtain the required ladder exclusion from finitely many forbidden hypergraphs. In the compactness proof, we use the Hessian to represent the colors by vectors on a sphere. Ladder exclusion forces the limiting distribution to be concentrated on countably many points of the sphere. We use these points to construct a countable palette attaining the limiting value, then find a finite subpalette with the same Lagrangian.

For distinct $p, q \in [r]$ and $m \ge 2$, let $S^{p,q}_m$ have colors $b_1, \dots, b_m$ and $c_{i,\ell}$ for $i \in [m-1]$ and $\ell \in [r] \setminus \{p, q\}$. For each $i < j$, include the tuple with b_i in position p, b_j in position q, and $c_{i,\ell}$ in position ℓ.

For $1 \le i < j \le m$, the tuples for $p < q$ and $q < p$ are, respectively,

$$(c_{i,1}, \dots, c_{i,p-1}, b_i, c_{i,p+1}, \dots, c_{i,q-1}, b_j, c_{i,q+1}, \dots, c_{i,r}),$$
$$(c_{i,1}, \dots, c_{i,q-1}, b_j, c_{i,q+1}, \dots, c_{i,p-1}, b_i, c_{i,p+1}, \dots, c_{i,r}).$$

Empty blocks of c-colors are omitted. Note that $S^{p,q}_m$ is an m-ladder palette with a fixed ordering of each edge.

The next lemma uses Ramsey's theorem [41, Theorem B] to obtain a homomorphism from some $S_m^{p,q}$ to any rainbow r-palette whose underlying r-graph contains a sufficiently long ladder.

Lemma 9.3. *For every $r \geq 3$ and $m \geq 2$, there is an integer $M \geq m$ such that every rainbow r-palette P whose underlying r-graph contains an M-ladder admits a homomorphism $S_m^{p,q} \to P$ for some distinct $p, q \in [r]$.*

Proof. Let $M \coloneqq R_2(m; r!)$, where $R_2(m; r!)$ is the $r!$-color Ramsey number for pairs. Fix an M-ladder in $H(P)$ with spine vertices $b_1, \dots, b_M$ and cores $A_1, \dots, A_{M-1}$, and order each core arbitrarily. For each $i < j$, choose an admissible permutation of $A_i \cup \{b_i, b_j\}$. Color the pair $\{i, j\}$ by the positions of b_i, b_j, and the ordered core vertices in this tuple. There are at most $r!$ colors, so Ramsey's theorem gives a homogeneous set of m indices.

On this set, the earlier and later spine vertices occupy fixed positions $p \neq q$, and each ordered core vertex also occupies a fixed position. Map the spine vertices of $S_m^{p,q}$ to the selected spine vertices in increasing order, and map each $c_{i,\ell}$ to the vertex in the corresponding core occupying position ℓ. Homogeneity makes this choice independent of the later spine vertex, so the map takes every tuple of $S_m^{p,q}$ to an admissible tuple of P. This gives the required homomorphism, which may identify auxiliary colors when different cores intersect. □

To obtain the required ladder exclusion, we compare the palettes $S_m^{p,q}$ with a single repeated-color tuple. Fix distinct positions $p, q \in [r]$, and let $R_{p,q}$ consist of a single tuple with one color a in positions p, q and a different color z_ℓ in every other position. For $p < q$, set

$$R_{p,q} \coloneqq \{(z_1, \dots, z_{p-1}, a, z_{p+1}, \dots, z_{q-1}, a, z_{q+1}, \dots, z_r)\}.$$

Empty blocks of z-colors are omitted. We have $R_{p,q} = R_{q,p}$ up to the names of its colors.

Lemma 9.4. *The following statements hold.*

(i) *We have $\lambda(R_{p,q}) = 4r^{-r}$.*

(ii) *For every $m \geq 2$, we have $\lambda(S_m^{p,q}) = r^{-r}$.*

(iii) *We have $\mathrm{Col}(R_{p,q}) = \bigcup_{m=2}^{\infty} \mathrm{Col}(S_m^{p,q})$. More precisely, for every $m \geq 2$, each $R_{p,q}$-colorable r-graph with at most m vertices is $S_m^{p,q}$-colorable.*

Proof. For (i), the polynomial of $R_{p,q}$ is the single monomial $x_a^2 \prod_{\ell \in [r] \setminus \{p,q\}} x_{z_\ell}$. Its maximum is attained at $x_a = 2/r$ and $x_{z_\ell} = 1/r$, giving $\lambda(R_{p,q}) = 4r^{-r}$.

For (ii), write x_i for the weight of b_i, $y_{i,\ell}$ for that of $c_{i,\ell}$, and put $B \coloneqq \sum_{i \in [m]} x_i$ and $Y_\ell \coloneqq \sum_{i \in [m-1]} y_{i,\ell}$. It follows from the definition of $S_m^{p,q}$ that

$$\Lambda_{S_m^{p,q}} = \sum_{i \in [m-1]} x_i \left(\sum_{j=i+1}^{m} x_j \right) \prod_{\ell \in [r] \setminus \{p,q\}} y_{i,\ell} \leq \frac{B^2}{4} \prod_{\ell \in [r] \setminus \{p,q\}} Y_\ell \leq r^{-r}.$$

For the first inequality, use $x_i \sum_{j=i+1}^{m} x_j \leq B^2/4$ and $\sum_{i \in [m-1]} \prod_{\ell \in [r] \setminus \{p,q\}} y_{i,\ell} \leq \prod_{\ell \in [r] \setminus \{p,q\}} Y_\ell$. The last inequality follows by applying AM–GM to $B/2, B/2$ and the $r-2$ numbers Y_ℓ. Assigning weight $1/r$ to each color of one admissible tuple gives equality.

For (iii), the map $b_i \mapsto a$, $c_{i,\ell} \mapsto z_\ell$ gives a homomorphism $S_m^{p,q} \to R_{p,q}$. To prove the reverse inclusion, fix an $R_{p,q}$-coloring of an r-graph with ordered vertices $1, \dots, N$. The conclusion is immediate for an edgeless r-graph, so assume $N \geq r$. For an $(r-1)$-face T, write $T[h]$ for its hth vertex. Define

$$h_\ell \coloneqq \begin{cases} r+1-q, & \text{if } \ell < q, \\ r-q, & \text{if } \ell > q, \end{cases} \quad \text{and} \quad \sigma(v) \coloneqq \begin{cases} v, & \text{if } p < q, \\ N+1-v, & \text{if } p > q. \end{cases}$$

Here h_ℓ is defined for every $\ell \in [r] \setminus \{q\}$ and belongs to $[r-1]$. Recolor each $(r-1)$-face contained in an edge according to $a \mapsto b_{\sigma(T[h_p])}$ and $z_\ell \mapsto c_{\sigma(T[h_\ell]),\ell}$. Consider an edge $v_1 < \cdots < v_r$. The faces in position p and in every position $\ell \neq p, q$ all select v_{r+1-q}. They therefore receive the same first index $i \coloneqq \sigma(v_{r+1-q})$. The face in position q selects v_{r+2-q} if $p < q$ and v_{r-q} if $p > q$. Its index j satisfies $i < j$ in both cases. The resulting tuple belongs to $S_N^{p,q}$. The color assigned to a face depends only on that

face, so the rule is consistent on shared faces. Unused faces can be colored arbitrarily. Finally $S_N^{p,q}$ maps to $S_m^{p,q}$ when $N \le m$. □

We now use Lemmas 9.3 and 9.4 to exclude long ladders from deficient palettes.

Lemma 9.5. *Suppose $\pi^{(r)}_{r-2}(\mathcal{F}) < 4r^{-r}$. There are a finite subfamily $\mathcal{F}_0 \subseteq \mathcal{F}$, with $|\mathcal{F}_0| \le \binom{r}{2}$, and an integer $M \ge 2$ such that every $\mathcal{F}_0$-deficient palette is rainbow and has no M-ladder in its underlying r-graph.*

Proof. For each unordered pair $\{p, q\}$, Lemma 9.4(i) gives $\lambda(R_{p,q}) = 4r^{-r}$, so $R_{p,q}$ is not $\mathcal{F}$-deficient by Theorem 8.1(i). Choose $F_{p,q} \in \mathcal{F} \cap \mathrm{Col}(R_{p,q})$, and set $\mathcal{F}_0 := \{F_{p,q} \colon 1 \le p < q \le r\}$. Write $F_{q,p} := F_{p,q}$ when the positions are listed in the opposite order. Put $m := \max\{2, |V(F_{p,q})| \colon p < q\}$ and let M be given by Lemma 9.3 for these r and m.

Let P be an $\mathcal{F}_0$-deficient palette. If an admissible tuple of P repeats a color in positions p, q, there is a homomorphism $R_{p,q} \to P$; the other colors of $R_{p,q}$ may also be identified. Thus $F_{p,q}$ would be P-colorable, contradicting the deficiency of P. This excludes all repeated-color tuples.

It remains to exclude an M-ladder. If P contains an M-ladder, then Lemma 9.3 gives $S_m^{p,q} \to P$ for some distinct $p, q \in [r]$. By Lemma 9.4(iii), $F_{p,q}$ is $S_m^{p,q}$-colorable and hence P-colorable, a contradiction. □

By Lemma 9.5, if a collection of rainbow palettes has underlying ladders of unbounded length, then any forbidden family for which every palette in the collection is deficient has density at least $4r^{-r}$. This conclusion is independent of the orientations of the underlying edges.

For example, fix distinct positions p, q. Let $\mathcal{F}$ consist of the r-graphs that are not $S_m^{p,q}$-colorable for any $m \ge 2$. By Lemma 9.4(iii), these are exactly the r-graphs outside $\mathrm{Col}(R_{p,q})$. The proof of Fact 8.2 gives $\pi^{(r)}_{r-2}(\mathcal{F}) = \lambda(R_{p,q}) = 4r^{-r}$, whereas $\sup_{m \ge 2} \lambda(S_m^{p,q}) = r^{-r}$.

9.2 Compactness below the threshold

Proof of Theorem 9.2. If $t = 0$, then $\lambda(P_n) < r^{-r}$ for all sufficiently large n. Since every nonempty finite r-palette has Lagrangian at least r^{-r} by Fact 8.2, we have $P_n = \varnothing$ for all such n, and a one-color palette with no admissible tuple suffices. Assume $t > 0$ and discard finitely many terms so that $\lambda(P_n) > 0$ for every n.

To justify the derivative calculations for countable palettes, we first establish convergence and continuity of the defining series. For a countable palette P, the tuple sum defining Λ_P extends to $\ell^1(C(P))$, the space of real weightings with $\|\boldsymbol{x}\|_1 := \sum_{c \in C(P)} |x_c| < \infty$. The series converges absolutely because $\sum_{(c_1,\dots,c_r) \in P} \prod_{i \in [r]} |x_{c_i}| \le \|\boldsymbol{x}\|_1^r$. Expanding a difference of products gives

$$|\Lambda_P(\boldsymbol{x}) - \Lambda_P(\boldsymbol{y})| \le r \max\{\|\boldsymbol{x}\|_1, \|\boldsymbol{y}\|_1\}^{r-1} \|\boldsymbol{x} - \boldsymbol{y}\|_1.$$

If P is rainbow, its first and second derivatives are obtained by differentiating the monomials in (8.3) term by term. For distinct colors u, v and every $\boldsymbol{x} \in \ell^1(C(P))$, the resulting series satisfy

$$|\partial_u \Lambda_P(\boldsymbol{x})| \le \sum_{\substack{e \in H(P) \\ u \in e}} m_e \prod_{w \in e \setminus \{u\}} |x_w| \le r \|\boldsymbol{x}\|_1^{r-1},$$

$$|\partial_u \partial_v \Lambda_P(\boldsymbol{x})| \le \sum_{\substack{e \in H(P) \\ u,v \in e}} m_e \prod_{w \in e \setminus \{u,v\}} |x_w| \le r(r-1) \|\boldsymbol{x}\|_1^{r-2},$$

and $\partial_u^2 \Lambda_P = 0$. These bounds follow from $m_e \le r!$ and $\sum_{A \in \binom{C(P)}{j}} \prod_{w \in A} |x_w| \le \|\boldsymbol{x}\|_1^j / j!$. Applying the same difference-of-products estimate to the derivative series shows that they depend continuously on $\boldsymbol{x}$ in the ℓ^1 norm, uniformly in their coordinate indices. In particular, the sums defining first and second variations along ℓ^1 directions converge absolutely.

The real separable *Hilbert spaces* used below can be viewed as finite-dimensional Euclidean spaces or as ℓ^2, the space of real sequences with $\|\boldsymbol{a}\|^2 := \sum_{i=1}^{\infty} a_i^2 < \infty$ and inner product $\langle \boldsymbol{a}, \boldsymbol{b} \rangle := \sum_{i=1}^{\infty} a_i b_i$. For finite palettes, finitely many dimensions suffice.

We first represent the entries of the Hessian at an optimal weighting as squared distances between vectors on a sphere. The passage from a quadratic-form inequality to a squared-distance representation follows Schoenberg [48]; the argument below also identifies the common radius and the weighted center.

Claim 9.6. *Let P be a finite or countable rainbow r-palette. Suppose a probability weighting $\boldsymbol{x}$ has all coordinates positive and attains $\lambda(P) > 0$. Let $\mathsf{H} := (\mathsf{H}_{uv}) \in \mathbb{R}^{C(P)\times C(P)}$ be the Hessian of Λ_P at $\boldsymbol{x}$. There are vectors $\boldsymbol{q}_v$ in a Hilbert space, indexed by $v \in C(P)$, with the following properties.*

(i) $\|\boldsymbol{q}_v\|^2 = r(r-1)\lambda(P)/2$ *for every* $v \in C(P)$.

(ii) $\mathsf{H}_{uv} = \|\boldsymbol{q}_u - \boldsymbol{q}_v\|^2$ *for every* $u, v \in C(P)$.

(iii) $\sum_{v\in C(P)} x_v \boldsymbol{q}_v = 0$.

(iv) *If P is finite and $\boldsymbol{x}$ has minimum support among its optimal weightings, then the $\boldsymbol{q}_v$ are pairwise distinct.*

Proof of the claim. Put $d := r(r-1)\lambda(P)$. Applying the method of Lagrange multipliers to any two coordinates u, v, with all other weights fixed, gives $\partial_u \Lambda_P(\boldsymbol{x}) = \partial_v \Lambda_P(\boldsymbol{x})$. Since Λ_P is homogeneous of degree r, Euler's identity gives $\sum_{v\in C(P)} x_v \partial_v \Lambda_P(\boldsymbol{x}) = r\Lambda_P(\boldsymbol{x}) = r\lambda(P)$. Since the partial derivatives are equal and $\sum_{v\in C(P)} x_v = 1$, the gradient vector satisfies $\nabla\Lambda_P(\boldsymbol{x}) = (\partial_v \Lambda_P(\boldsymbol{x}))_{v\in C(P)} = r\lambda(P)\mathbf{1}$.

For each color u, the first derivative $\partial_u \Lambda_P$ is homogeneous of degree $r-1$. Applying Euler's identity to this function gives $\sum_{v\in C(P)} \mathsf{H}_{uv} x_v = (r-1)\partial_u \Lambda_P(\boldsymbol{x}) = r(r-1)\lambda(P) = d$. Hence $\mathsf{H}\boldsymbol{x} = d\mathbf{1}$. For every finitely supported $\boldsymbol{z}$ with $\sum_{v\in C(P)} z_v = 0$, the weighting $\boldsymbol{x} + \varepsilon\boldsymbol{z}$ is feasible for all sufficiently small positive and negative ε, since $\boldsymbol{x}$ is positive on the support of $\boldsymbol{z}$. Optimality therefore gives $\frac{d^2}{d\varepsilon^2}\Lambda_P(\boldsymbol{x} + \varepsilon\boldsymbol{z})\Big|_{\varepsilon=0} = \boldsymbol{z}^{\mathsf{T}}\mathsf{H}\boldsymbol{z} \le 0$.

All Hessian entries are bounded by $r!/(r-2)! = r(r-1)$. Finitely supported zero-sum vectors are dense in the zero-sum subspace of ℓ^1, so the second-order inequality extends to every zero-sum ℓ^1 vector. The derivative identities above are absolutely convergent.

Since the polynomial is squarefree, we have $\mathsf{H}_{vv} = 0$. Set $\mathsf{G} := (d\mathbf{1}\mathbf{1}^{\mathsf{T}} - \mathsf{H})/2$. For a finitely supported vector $\boldsymbol{z}$, write $s := \sum_{v\in C(P)} z_v$ and $\boldsymbol{y} := \boldsymbol{z} - s\boldsymbol{x}$. The identity $\mathsf{H}\boldsymbol{x} = d\mathbf{1}$ gives

$$\boldsymbol{z}^{\mathsf{T}}\mathsf{G}\boldsymbol{z} = \frac{1}{2}\bigl(ds^2 - \boldsymbol{z}^{\mathsf{T}}\mathsf{H}\boldsymbol{z}\bigr) = -\frac{1}{2}\boldsymbol{y}^{\mathsf{T}}\mathsf{H}\boldsymbol{y} \ge 0.$$

Thus G is a positive semidefinite matrix, also when indexed by a countable set. Its Gram representation gives vectors $\boldsymbol{q}_v$ with $\langle \boldsymbol{q}_u, \boldsymbol{q}_v\rangle = \mathsf{G}_{uv}$. Since $\mathsf{G}_{vv} = d/2$, we have $\|\boldsymbol{q}_v\|^2 = d/2$ and $\|\boldsymbol{q}_u - \boldsymbol{q}_v\|^2 = \mathsf{G}_{uu} + \mathsf{G}_{vv} - 2\mathsf{G}_{uv} = \mathsf{H}_{uv}$, proving (i) and (ii). Moreover, the vector sum $\sum_{v\in C(P)} x_v \boldsymbol{q}_v$ converges absolutely and has squared norm $\boldsymbol{x}^{\mathsf{T}}\mathsf{G}\boldsymbol{x} = 0$, so it is zero, proving (iii).

For (iv), the weight-transfer argument of Motzkin and Straus [30] and Frankl and Rödl [10] shows that, at an optimal weighting of minimum support on a finite palette, any two colors of positive weight lie in a common underlying edge. Otherwise the polynomial is affine when weight is transferred between those colors, and optimality permits deleting one of them. A common edge, positive weights, and nonnegative coefficients give $\mathsf{H}_{uv} > 0$. The distance formula then implies $\boldsymbol{q}_u \ne \boldsymbol{q}_v$. □

For a Borel probability measure ν on a separable Hilbert space, a point $\boldsymbol{q}$ is an *atom* if $\nu(\{\boldsymbol{q}\}) > 0$. The set of atoms is at most countable. We call ν *atomic*, or *purely atomic*, if all its probability lies on these atoms, that is, $\sum_{\boldsymbol{q}:\, \nu(\{\boldsymbol{q}\})>0} \nu(\{\boldsymbol{q}\}) = 1$. The restriction of ν to the complement of its atoms is its *nonatomic part*.

To apply Claim 9.6 to each P_n, choose an optimal weighting of minimum support and restrict P_n to this support. Denote its positive weights by $x_v^{(n)}$. Put $c_n := r(r-1)\lambda(P_n)$ and $c := r(r-1)t$, so $\lim_{n\to\infty} c_n = c$. By Claim 9.6, there are pairwise distinct vectors $\boldsymbol{p}_v^{(n)}$ with $\|\boldsymbol{p}_v^{(n)}\|^2 = c_n/2$ and $\mathsf{H}_{uv}^{(n)} = c_n - 2\langle \boldsymbol{p}_u^{(n)}, \boldsymbol{p}_v^{(n)}\rangle$. Let $\mu_n := \sum_{v\in C(P_n)} x_v^{(n)} \delta_{\boldsymbol{p}_v^{(n)}}$ be the probability distribution of these vectors. Here $\delta_{\boldsymbol{q}}$ denotes the probability measure concentrated at $\boldsymbol{q}$.

To control uniformly in n the error from truncating inner products to finitely many coordinates, rotate each configuration so that its weighted second-moment matrix is diagonal. Its (i,j) entry is

$\sum_{v\in C(P_n)} x_v^{(n)} p_{v,i}^{(n)} p_{v,j}^{(n)}$. The matrix is real symmetric and positive semidefinite, so the spectral theorem provides an orthogonal change of coordinates that diagonalizes it and preserves all norms and inner products. Order the diagonal entries so that $\rho_{n,1} \geq \rho_{n,2} \geq \cdots \geq 0$; their sum is $c_n/2$.

Embed all these configurations into ℓ^2. They lie in a common closed ball centered at the origin. In the coordinatewise topology, this ball is a closed subset of a countable product of compact intervals, since its radius bound can be checked on finite sums of squared coordinates. The ball is therefore compact and metrizable. Passing to a subsequence, μ_n converges weakly on that ball to a probability measure μ, meaning that $\lim_{n\to\infty} \int \varphi \,\mathrm{d}\mu_n = \int \varphi \,\mathrm{d}\mu$ for every bounded continuous real-valued function φ on the ball with its coordinatewise topology.

The Skorokhod representation theorem [3] gives random vectors $\boldsymbol{p}_n$ and $\boldsymbol{p}$, with distributions μ_n and μ, on $([0,1], \mathrm{Leb})$ such that, for almost every s, we have $\lim_{n\to\infty} p_{n,j}(s) = p_j(s)$ for every $j \geq 1$. Here Leb denotes Lebesgue measure, and for measurable subsets of $[0,1]^k$ we write $|\cdot|$ for their Lebesgue measure. For almost every s and every positive integer L, coordinatewise convergence gives

$$\sum_{j=1}^{L} p_j(s)^2 = \lim_{n\to\infty} \sum_{j=1}^{L} p_{n,j}(s)^2 \leq \lim_{n\to\infty} \frac{c_n}{2} = \frac{c}{2}.$$

Letting L tend to infinity shows that $\|\boldsymbol{p}(s)\|^2 \leq c/2$ almost everywhere. Since the finite vectors are pairwise distinct, $\boldsymbol{p}_n(s)$ determines a unique color $\phi_n(s)$ of P_n, with distribution $(x_v^{(n)})_v$.

We next show that if μ has a nonzero nonatomic part or an atom of squared norm less than $c/2$, then one of the finite palettes P_n contains an M-ladder. For this purpose, we express the Hessian entries as symmetric kernels and prove their L^2 convergence. Define

$$K_n(s,t) \coloneqq c_n - 2\langle \boldsymbol{p}_n(s), \boldsymbol{p}_n(t)\rangle, \qquad K(s,t) \coloneqq c - 2\langle \boldsymbol{p}(s), \boldsymbol{p}(t)\rangle.$$

Claim 9.7. *The kernels satisfy* $\lim_{n\to\infty} \|K_n - K\|_{L^2([0,1]^2)} = 0$. *That is,* $\lim_{n\to\infty} \int_{[0,1]^2} |K_n(s,t) - K(s,t)|^2 \,\mathrm{d}s\,\mathrm{d}t = 0$.

Proof of the claim. Since $\lambda(P_n) \leq 1$, we have $c_n \leq r(r-1)$. For every $n \geq 1$ and positive integer L, the diagonal entries $\rho_{n,j}$ of the second-moment matrix are nonnegative and nonincreasing, with sum $c_n/2 \leq r(r-1)/2$, so $\rho_{n,L+1} \leq r(r-1)/(2(L+1))$. It follows that $\sum_{j=L+1}^{\infty} \rho_{n,j}^2 \leq \rho_{n,L+1} \sum_{j=L+1}^{\infty} \rho_{n,j} \leq r^2(r-1)^2/(4(L+1))$.

We now apply this estimate to the error in truncating the inner product. By the diagonalization, $\int_0^1 p_{n,i}(s) p_{n,j}(s) \,\mathrm{d}s = 0$ for $i \neq j$, and $\int_0^1 p_{n,j}(s)^2 \,\mathrm{d}s = \rho_{n,j}$. Expanding the square and integrating therefore gives

$$\iint_{[0,1]^2} \left| \langle \boldsymbol{p}_n(s), \boldsymbol{p}_n(t)\rangle - \sum_{j=1}^{L} p_{n,j}(s) p_{n,j}(t) \right|^2 \mathrm{d}s\,\mathrm{d}t = \sum_{j=L+1}^{\infty} \rho_{n,j}^2 \leq \frac{r^2(r-1)^2}{4(L+1)}.$$

For each fixed L, the truncated sums converge in L^2 by the almost-everywhere coordinatewise convergence of $\boldsymbol{p}_n$ to $\boldsymbol{p}$ and bounded convergence. To obtain the corresponding tail bound for $\boldsymbol{p}$, note that, for each j, bounded convergence gives $\rho_j \coloneqq \int_0^1 p_j(s)^2 \,\mathrm{d}s = \lim_{n\to\infty} \rho_{n,j}$, and, for $i \neq j$, $\int_0^1 p_i(s) p_j(s) \,\mathrm{d}s = \lim_{n\to\infty} \int_0^1 p_{n,i}(s) p_{n,j}(s) \,\mathrm{d}s = 0$. The sequence (ρ_j) is nonincreasing. Moreover, for every positive integer L, $\sum_{j=1}^{L} \rho_j = \lim_{n\to\infty} \sum_{j=1}^{L} \rho_{n,j} \leq r(r-1)/2$, so $\sum_{j=1}^{\infty} \rho_j \leq r(r-1)/2$. Thus $\sum_{j=L+1}^{\infty} \rho_j^2 \leq r^2(r-1)^2/(4(L+1))$, so the same tail estimate holds for $\boldsymbol{p}$. First choose L large, then let n tend to infinity, which proves the claim. □

In particular, convergence holds in L^1. The kernels K_n are nonnegative, being Hessian entries, and K is nonnegative because $\|\boldsymbol{p}(s)\|^2 \leq c/2$.

Let $m_e^{(n)}$ count the admissible permutations of $e \in H(P_n)$. Define a symmetric function $\omega_n \colon [0,1]^r \to [0, r!]$ by

$$\omega_n(s_1, \ldots, s_r) \coloneqq \begin{cases} m_e^{(n)}, & \text{if } e = \{\phi_n(s_1), \ldots, \phi_n(s_r)\} \in H(P_n), \\ 0, & \text{otherwise.} \end{cases}$$

Then $0 \le \omega_n \le r!$ and

$$\int_{[0,1]^{r-2}} \omega_n(s,t,u_3,\dots,u_r)\,\mathrm{d}u_3 \cdots \mathrm{d}u_r = (r-2)!K_n(s,t). \tag{9.1}$$

The factor $(r-2)!$ counts the orders of the remaining colors of an edge after the first two colors have been fixed; each Hessian monomial occurs once for every such order.

Claim 9.8. *The measure μ is purely atomic, and every atom has squared norm $c/2$.*

Proof of the claim. Let $U \subseteq [0,1]$ be the set of points for which $\boldsymbol{p}(s)$ is not an atom of μ having squared norm $c/2$. We have $K(s,t) > 0$ almost everywhere on $U \times U$. Indeed, $K(s,t) = 0$ requires $\boldsymbol{p}(s) = \boldsymbol{p}(t)$ and both vectors to have squared norm $c/2$. The event $K(s,t) = 0$ has positive product measure only at an atom with that squared norm, and all such atoms were removed in defining U.

Suppose $|U| > 0$, and fix $a \in (0, |U|]$. Choose $\gamma > 0$ such that $|\{(s,t) \in U^2 \colon K(s,t) < \gamma\}| < a^2/2$. For every measurable $S \subseteq U$ with $|S| \ge a$, at least half of S^2 lies where $K \ge \gamma$. By Claim 9.7, for all sufficiently large n,

$$\int_{S\times S} K_n \ge \frac{\gamma a^2}{2} - \|K_n - K\|_{L^1} \ge \frac{\gamma a^2}{4}.$$

This bound holds uniformly for all such S, including sets chosen after n.

Since $\omega_n \le r!$, equation (9.1) gives

$$\int_{S\times S\times [0,1]^{r-2}} \mathbf{1}_{\{\omega_n > 0\}} \ge \frac{1}{r(r-1)} \int_{S\times S} K_n.$$

If $|S| \ge a$, averaging over $(t,u_3,\dots,u_r) \in S \times [0,1]^{r-2}$ gives a choice with $\int_S \mathbf{1}_{\{\omega_n(s,t,u_3,\dots,u_r)>0\}}\,\mathrm{d}s \ge \gamma a^2/(4r(r-1))$, since the averaging domain has measure at most 1. Put $b \coloneqq \phi_n(t)$ and $A \coloneqq \{\phi_n(u_3),\dots,\phi_n(u_r)\}$. Then $S' \coloneqq \{s \in S \colon A \cup \{b, \phi_n(s)\} \in H(P_n)\}$ has measure at least $\gamma a^2/(4r(r-1))$. Positivity ensures that A has $r-2$ distinct colors and is disjoint from b.

Starting with $a_0 \coloneqq |U|$, make the following choices successively for $0 \le i < M-1$. Choose $\gamma_i > 0$ as above with $a = a_i$, and then choose $0 < a_{i+1} \le \min\{a_i/2, \gamma_i a_i^2/(8r(r-1))\}$. After these finitely many choices, choose a single n for which $\int_{S\times S} K_n \ge \gamma_i a_i^2/4$ holds for every $0 \le i < M-1$ and every measurable $S \subseteq U$ with $|S| \ge a_i$. Starting with $S_0 \coloneqq U$, repeat the preceding neighborhood selection $M-1$ times to obtain nested sets S_i of measure at least a_i, spine vertices b_i, and cores A_i. Choose $s \in S_{M-1}$ and put $b_M \coloneqq \phi_n(s)$. Every later spine vertex is the color of a point in each earlier selected neighborhood. By the definitions of these neighborhoods, it is distinct from the corresponding earlier spine vertex and from every color in its core. Thus, for every $1 \le i < j \le M$, we have $A_i \cap \{b_i,\dots,b_M\} = \varnothing$ and $A_i \cup \{b_i, b_j\} \in H(P_n)$. This is an M-ladder, a contradiction.

Thus $|U| = 0$, which proves the claim. □

By Claim 9.8, write the limit distribution as $\mu = \sum_{i\in I} w_i \delta_{\boldsymbol{p}_i}$, where I is finite or countably infinite, the $\boldsymbol{p}_i$ are distinct with $\|\boldsymbol{p}_i\|^2 = c/2$, and the weights satisfy $w_i > 0$ and $\sum_{i\in I} w_i = 1$. Put $\boldsymbol{w} = (w_i)_{i\in I}$.

We now use the atoms as colors while retaining the limiting Lagrangian. Sampling a point from the preimage of each atom under $\boldsymbol{p}$ and taking its color in P_n gives palettes with homomorphisms into P_n, and passing to a limit will recover the value t.

Let $B_i \coloneqq \{s \colon \boldsymbol{p}(s) = \boldsymbol{p}_i\}$, so $|B_i| = w_i$. Independently for $i \in I$, choose ξ_i according to normalized Lebesgue measure on B_i. Define a random palette Q_n on I by declaring $(i_1,\dots,i_r)$ admissible exactly when $(\phi_n(\xi_{i_1}),\dots,\phi_n(\xi_{i_r})) \in P_n$. It is therefore rainbow. For every finite $J \subseteq I$, the map $i \mapsto \phi_n(\xi_i)$ is a homomorphism $Q_n[J] \to P_n$. By the observation after Definition 9.1, homomorphisms between rainbow palettes preserve ladders. Thus an M-ladder in Q_n would give one in P_n, so every Q_n has no M-ladder.

Claim 9.9. *There is a rainbow palette Q on I with no M-ladder and with $\Lambda_Q(\boldsymbol{w}) = \lambda(Q) = t$. Along a subsequence, every finite $Q[J]$ admits a homomorphism to P_n for all sufficiently large n.*

Proof of the claim. Choose $s_1, \dots, s_r$ independently and uniformly from $[0,1]$. Since each $\phi_n(s_j)$ has distribution $(x_v^{(n)})_v$, the probability that $(\phi_n(s_1), \dots, \phi_n(s_r)) \in P_n$ is $\lambda(P_n)$. The expectation of $\Lambda_{Q_n}(\boldsymbol{w})$ is

$$\mathbb{E}[\Lambda_{Q_n}(\boldsymbol{w})] = \mathbb{P}\left[(\phi_n(s_1), \dots, \phi_n(s_r)) \in P_n \text{ and } \boldsymbol{p}(s_i) \neq \boldsymbol{p}(s_j) \text{ for every } 1 \le i < j \le r\right].$$

The second condition says that the points lie in pairwise different sets B_i. Indeed, for distinct indices $i_1, \dots, i_r$, the event $s_j \in B_{i_j}$ for every $j \in [r]$ has probability $w_{i_1} \cdots w_{i_r}$; conditional on this event, the points have the same joint law as $\xi_{i_1}, \dots, \xi_{i_r}$. Terms with repeated indices contribute zero because Q_n is rainbow. Thus $\lambda(P_n) - \mathbb{E}[\Lambda_{Q_n}(\boldsymbol{w})]$ is the probability that the color tuple is admissible and at least 2 sampled points lie in the same B_i.

Since $K = 0$ on each $B_i \times B_i$ and these sets are pairwise disjoint, we have

$$\sum_{i \in I} \int_{B_i \times B_i} K_n \le \|K_n - K\|_{L^1}.$$

By Claim 9.7, $\lim_{n\to\infty} \|K_n - K\|_{L^1} = 0$. Whenever the color tuple is admissible, $\omega_n(s_1, \dots, s_r) \ge 1$. A union bound over pairs of sampled points, followed by the symmetry of ω_n and (9.1), gives

$$0 \le \lambda(P_n) - \mathbb{E}[\Lambda_{Q_n}(\boldsymbol{w})] \le \binom{r}{2}(r-2)! \sum_{i \in I} \int_{B_i \times B_i} K_n.$$

Thus $\mathbb{E}[\Lambda_{Q_n}(\boldsymbol{w})] \to t$. For each n, choose a deterministic realization $\widehat{Q}_n$ whose value is at least this expectation minus $1/n$. The sampling map gives $\widehat{Q}_n \to P_n$, so

$$\mathbb{E}[\Lambda_{Q_n}(\boldsymbol{w})] - \frac{1}{n} \le \Lambda_{\widehat{Q}_n}(\boldsymbol{w}) \le \lambda(\widehat{Q}_n) \le \lambda(P_n).$$

Both outer bounds tend to t.

Since I^r is countable and every tuple indicator takes values in $\{0,1\}$, a diagonal subsequence makes each indicator eventually constant. Let Q be the resulting palette. The rainbow condition and the absence of an M-ladder pass to Q, since a violation is witnessed by finitely many tuples. For every finite $J \subseteq I$, we have $Q[J] = \widehat{Q}_n[J]$ for all sufficiently large n along the subsequence, and hence $Q[J] \to P_n$ and $\lambda(Q[J]) \le t$. Taking the supremum over finite J gives $\lambda(Q) \le t$.

For any palette R on I and any finite $J \subseteq I$, deleting tuples meeting $I \setminus J$ without renormalizing the remaining weights gives

$$0 \le \Lambda_R(\boldsymbol{w}) - \Lambda_{R[J]}((w_i)_{i \in J}) \le 1 - \left(\sum_{i \in J} w_i\right)^r \le r \sum_{i \notin J} w_i.$$

This bound is uniform in R. On J^r the indicators of $\widehat{Q}_n$ eventually agree with those of Q, so first taking $n \to \infty$ and then increasing J to I gives $\Lambda_Q(\boldsymbol{w}) = t$. Therefore $\Lambda_Q(\boldsymbol{w}) = \lambda(Q) = t$. □

It remains to show that a finite subpalette attains the same Lagrangian as the countable palette.

Claim 9.10. *Let P be a countable rainbow palette with no M-ladder for some integer $M \ge 2$. If a probability weighting $\boldsymbol{x}$ attains a positive value $\lambda(P)$, then a finite subpalette of P has Lagrangian $\lambda(P)$.*

Proof of the claim. Discard zero-weight colors and apply Claim 9.6 to P with its optimal weighting $\boldsymbol{x}$, writing $d \coloneqq r(r-1)\lambda(P)$ and $\boldsymbol{q}_v$ for the resulting vectors. Their common norm is $\sqrt{d/2}$. Suppose first that the vectors $\boldsymbol{q}_v$ take infinitely many distinct values. Choose colors $b_1, b_2, \dots$ with pairwise distinct vectors. Their boundedness allows passage to a subsequence along which $\boldsymbol{q}_{b_n}$ converges weakly to $\boldsymbol{q}_*$, meaning that $\lim_{n\to\infty} \langle \boldsymbol{q}_{b_n}, \boldsymbol{z} \rangle = \langle \boldsymbol{q}_*, \boldsymbol{z} \rangle$ for every vector $\boldsymbol{z}$ in the Hilbert space. For each $(r-1)$-set K of colors, set $m_{K \cup \{b_n\}} \coloneqq 0$ when $b_n \in K$. There are only countably many such sets K, and each coefficient takes values in $\{0, 1, \dots, r!\}$. A diagonal argument therefore gives a further subsequence along which, for every fixed K, the coefficient $m_{K \cup \{b_n\}}$ is constant for all sufficiently large n.

For a fixed color v, we have $\lim_{n\to\infty} \mathsf{H}_{b_n v} = d - 2\langle \boldsymbol{q}_*, \boldsymbol{q}_v \rangle$. The right side is nonnegative. It is zero only if $\|\boldsymbol{q}_*\| = \sqrt{d/2}$ and $\boldsymbol{q}_v = \boldsymbol{q}_*$. Hence it is positive for all but at most one of the selected colors b_i.

Expanding the Hessian entry gives

$$\mathsf{H}_{b_n v} = \sum_{\substack{A \subseteq C(P)\setminus\{v,b_n\} \\ |A|=r-2}} m_{A\cup\{v,b_n\}} \prod_{w\in A} x_w.$$

For all sufficiently large n, we have $b_n \neq v$. Using the zero-coefficient convention above, we may take the sum over the fixed family of $(r-2)$-sets $A \subseteq C(P)\setminus\{v\}$. For each such A, the preceding diagonal choice with $K = A\cup\{v\}$ makes $m_{A\cup\{v,b_n\}}$ eventually constant. Since $0 \le m_{A\cup\{v,b_n\}} \le r!$ and $\sum_{A\in\binom{C(P)}{r-2}} \prod_{w\in A} x_w \le 1/(r-2)!$, the dominated convergence theorem for series gives

$$\lim_{n\to\infty} \mathsf{H}_{b_n v} = \sum_{\substack{A \subseteq C(P)\setminus\{v\} \\ |A|=r-2}} \Big(\lim_{n\to\infty} m_{A\cup\{v,b_n\}}\Big) \prod_{w\in A} x_w.$$

If $\lim_{n\to\infty} \mathsf{H}_{b_n v} > 0$, at least one summand on the right is positive. Its core A therefore satisfies $\lim_{n\to\infty} m_{A\cup\{v,b_n\}} \ge 1$. Since this coefficient is eventually constant, we have $A\cup\{v,b_n\} \in H(P)$ for every sufficiently large n.

We now select indices $n_1 < \cdots < n_M$, avoiding the possible single color with zero Hessian limit. After choosing b_{n_i} for $i < M$, choose a core A_i by the preceding argument with $v = b_{n_i}$. Then choose $n_{i+1} > n_i$ so large that $b_{n_{i+1}} \notin \bigcup_{j=1}^{i} A_j$ and $A_j\cup\{b_{n_j}, b_{n_{i+1}}\} \in H(P)$ for every $j \in [i]$, again avoiding the exceptional color. These choices give $A_i \cap \{b_{n_i}, \dots, b_{n_M}\} = \varnothing$ for every $i \in [M-1]$, and every $A_i \cup \{b_{n_i}, b_{n_j}\}$ with $i < j$ is an edge. Thus the selected spine vertices and cores form an M-ladder, a contradiction.

The vectors $\boldsymbol{q}_v$ thus take only finitely many distinct values. Partition the colors according to those values. If u, v belong to one part, then $\mathsf{H}_{uv} = 0$. Since all weights are positive and all coefficients are nonnegative, no underlying edge contains both colors. The polynomial is therefore affine in the weights of any one part when the weights in the other parts are fixed. At the current optimal weighting, the coefficients of x_v in this affine function are the partial derivatives $\partial_v \Lambda_P(\boldsymbol{x})$, which are equal for all colors v in that part by the two-coordinate first-order condition. Concentrating the total weight of that part on one color leaves the value unchanged. The resulting weighting is still optimal, so the same argument applies successively to all the other parts. Since there are only finitely many parts, we obtain a finitely supported probability weighting attaining $\lambda(P)$. □

By Claims 9.9 and 9.10, choose a finite $J \subseteq I$ with $\lambda(Q[J]) = t$. Claim 9.9 gives $Q[J] \to P_n$ for all sufficiently large n along the subsequence. Composing with the inclusions into the original palettes, before restriction to optimal supports, proves the theorem. □

Proof of the finite-attainment assertions in Theorem 1.3(i). If $\pi^{(r)}_{r-2}(\mathcal{F}) = 0$, use a palette with no admissible tuples. Now suppose that $0 < \pi^{(r)}_{r-2}(\mathcal{F}) < 4r^{-r}$. By Theorem 8.1(i), choose finite $\mathcal{F}$-deficient palettes P_n with $\lim_{n\to\infty} \lambda(P_n) = \pi^{(r)}_{r-2}(\mathcal{F})$. By Lemma 9.5, the palettes P_n are rainbow and contain no M-ladder for some M independent of n. Theorem 9.2 gives a finite palette P with $\lambda(P) = \pi^{(r)}_{r-2}(\mathcal{F})$ and $P \to P_n$ along a subsequence. If a member of $\mathcal{F}$ were P-colorable, it would be P_n-colorable, a contradiction. Thus P is $\mathcal{F}$-deficient.

By Fact 8.2, $\mathcal{D}_r \subseteq \Pi^{(r)}_{r-2,\infty}$.

For a finite palette, take $\mathcal{I} = \{\varnothing\}$ in Lemma 5.3. Its capacity is the Lagrangian, so that lemma gives algebraicity. There are only countably many finite palettes, proving countability.

By Theorem 8.5, $\mathcal{D}_r \subseteq \Pi^{(r)}_{r-2,\mathrm{fin}} \subseteq \Pi^{(r)}_{r-2,\infty}$. Finite attainment makes these inclusions equalities below $4r^{-r}$. The palette with the single tuple $(a, a, c_3, \dots, c_r)$, where $a, c_3, \dots, c_r$ are pairwise distinct, has Lagrangian $(2/r)^2(1/r)^{r-2} = 4r^{-r}$. The same realization theorem puts this endpoint in all three sets. □

9.3 Nonclosedness

Proof of nonclosedness in Theorem 1.3(ii). By Theorem 1.1(iii), $[4/9, 1] \subseteq \Pi^{(3)}_{\infty}$. Normalized Lagrangians of finite 3-graphs are dense in this interval. Selecting one ordering of every edge and then applying Fact 8.3 shows that

$$\left[2r^{-r}, \frac{9}{2}r^{-r}\right] \subseteq \overline{\mathcal{D}_r} \qquad \text{for } r \ge 3.$$

Every transcendental point below $4r^{-r}$ is excluded from $\Pi^{(r)}_{r-2,\infty}$ by Theorem 1.3(i). In particular, all transcendental points of $[(54/25)r^{-r}, (9/4)r^{-r}]$ are missing limit points. This interval contains continuum many transcendental numbers, proving the claimed cardinality and nonclosedness. □

For the closure statement, we use the ordinary-density theorem in Theorem 1.1(ii). Thus we may fix an integer $s_0 \geq 5$ such that

$$\Pi^{(s)}_{\infty} = \{0\} \cup \left[\frac{s!}{s^s}, 1\right] \qquad \text{for every } s \geq s_0. \tag{9.2}$$

Proof of the closure assertion in Theorem 1.3(ii). We first show that $\mathcal{D}_s \cap [s^{-s}, 1]$ is dense in $[s^{-s}, 1]$ for every $s \geq s_0$. For $1 \leq m \leq s!$, Fact 8.3 gives $\lambda(P_m(H)) = (m/s!)\lambda(H)$ for each finite simple s-graph H. By (9.2) and the density of normalized Lagrangians of finite s-graphs, we have $[ms^{-s}, m/s!] \subseteq \overline{\mathcal{D}_s}$. Since $s!/s^s < 1/2 \leq m/(m+1)$, we have $m/s! \geq (m+1)s^{-s}$, so consecutive intervals overlap. The union of these intervals over $1 \leq m \leq s!$ is $[s^{-s}, 1]$. A single rainbow tuple gives s^{-s}, and a one-color palette containing its sole tuple gives 1.

Together with the inclusions $\mathcal{D}_r \subseteq \Pi^{(r)}_{r-2,\infty} \subseteq \{0\} \cup [r^{-r}, 1]$ from Fact 8.2, this proves the closure identity for every $r \geq s_0$. Theorem 1.3(i) shows that $\Pi^{(r)}_{r-2,\infty} \cap [0, 4r^{-r})$ is countable and consists of algebraic numbers. For $r \geq s_0$, the set $\mathcal{D}_r \cap [r^{-r}, 4r^{-r})$ is dense in $[r^{-r}, 4r^{-r})$. Countability also excludes every nondegenerate interval below $4r^{-r}$ for all $r \geq 3$. For $1 \leq C < 4$, any interval $[Cr^{-r}, 1]$ contains such a nondegenerate subinterval, so it cannot be a subset of $\Pi^{(r)}_{r-2,\infty}$. □

10 The upper interval at the highest nontrivial level

We now construct every value in $[4r^{-r}, 1]$ for $r \geq 3$, starting with a short interval at $4/27$ realized by compatible ternary palettes. Cyclic refinements preserve the bonding homomorphisms while controlling the global maximizing weights. The finite operations in Section 10.5 then extend the interval and increase the arity.

Put $\theta \coloneqq 19683/132845 = 4/27 + 61/3586815$. The central construction is the following.

Theorem 10.1. *For every $\beta \in [4/27, \theta]$, there is a projective sequence of finite ternary palettes whose Lagrangians decrease to β.*

Write $e_3(\boldsymbol{x}) \coloneqq \sum_{i<j<k} x_i x_j x_k$. Hessian inequalities on a simplex are understood on its tangent space $\mathbf{1}^{\perp}$.

Let $\mathcal{R}_k$ be the set of limits of Lagrangians along projective sequences of finite k-palettes. Constant sequences give $\mathcal{D}_k \subseteq \mathcal{R}_k$. If finite r-palettes satisfy $P_{j+1} \to P_j$ and $\lambda(P_j) \downarrow \alpha$, their canonical forbidden families increase. By Fact 8.2, each has $(r-2)$-uniform density $\lambda(P_j)$. Applying Equation (2.2) to their union gives density α. Hence $\mathcal{R}_r \subseteq \Pi^{(r)}_{r-2,\infty}$ for every $r \geq 3$.

10.1 Cubic estimates and directed cycle covers

The stationary-value estimate in Lemma 4.12 and the following boundary bound control the base palette.

Lemma 10.2 (Localization from a boundary bound)**.** *Let f have degree at most three, let $\boldsymbol{z}$ be interior, and assume $f \leq c$ on the boundary of its simplex. For $\boldsymbol{x} \neq \boldsymbol{z}$, write $\boldsymbol{x} = \boldsymbol{z} + \rho(\boldsymbol{v} - \boldsymbol{z})$, where $\boldsymbol{v}$ is on the boundary and $0 < \rho \leq 1$. If $\boldsymbol{h} = \boldsymbol{x} - \boldsymbol{z}$ and $\boldsymbol{g} = \operatorname{proj}_{\mathbf{1}^{\perp}} \nabla f(\boldsymbol{z})$, then*

$$f(\boldsymbol{x}) - c \leq (1-\rho)\left[(1+\rho+\rho^2)(f(\boldsymbol{z}) - c) + (1+\rho)\boldsymbol{g}\cdot\boldsymbol{h} + \frac{1}{2}\boldsymbol{h}^{\mathsf{T}}\nabla^2 f(\boldsymbol{z})\boldsymbol{h}\right]. \tag{10.1}$$

Proof. Since $\boldsymbol{x} = \boldsymbol{z} + \boldsymbol{h}$ and $\boldsymbol{v} = \boldsymbol{z} + \boldsymbol{h}/\rho$, the cubic terms cancel in the identity

$$f(\boldsymbol{x}) - \rho^3 f(\boldsymbol{v}) = (1-\rho^3)f(\boldsymbol{z}) + (1-\rho^2)\boldsymbol{g}\cdot\boldsymbol{h} + \frac{1-\rho}{2}\boldsymbol{h}^{\mathsf{T}}\nabla^2 f(\boldsymbol{z})\boldsymbol{h}.$$

Using $f(\boldsymbol{v}) \leq c$ and factoring out $1-\rho$ gives (10.1). □

10.1.1 Partial directed cycle covers

Let D be a directed graph with indegree and outdegree at most one, without loops or directed 2-cycles. Its components are directed paths, isolated vertices, or directed cycles of length at least three. On its n vertices define

$$F_D(\boldsymbol{x}) = e_3(\boldsymbol{x}) + \sum_{u \to v \in D} x_u^2 x_v, \qquad b_n = \frac{1}{6} - \frac{1}{2n} + \frac{4}{3n^2}.$$

Lemma 10.3. *For every $n \geq 3$,*

$$\max_{S([n])} F_D \leq \max\left\{\frac{4}{27}, b_n\right\}. \tag{10.2}$$

Equality holds if D is a directed cycle cover. In particular, every such cover on 25 vertices has boundary maximum $4/27$.

Proof. We induct on n, treating boundary supports by the inductive hypothesis. On one vertex the value is zero. On two vertices there is at most one arc, and $\max x^2(1-x) = 4/27$.

First consider $3 \leq n \leq 8$. If D is a single directed cycle, the uniform point $\boldsymbol{u}_n$ is stationary, has value b_n, and has tangent Hessian eigenvalues

$$-1 + \frac{4}{n}\left(1 + \cos\frac{2\pi j}{n}\right), \qquad 1 \leq j < n.$$

These are nonpositive for $n = 3, 4, 6, 7, 8$. The observation after Lemma 4.12 bounds all interior stationary values by b_n. For $n = 5$, use the tangent vectors $\boldsymbol{h}, \boldsymbol{k}$ with $h_i = \cos(2\pi i/5)$ and $k_i = \sin(2\pi i/5)$. At every point of the simplex, direct differentiation gives

$$\boldsymbol{h}^{\mathsf{T}} \nabla^2 F_D(\boldsymbol{x}) \boldsymbol{h} + \boldsymbol{k}^{\mathsf{T}} \nabla^2 F_D(\boldsymbol{x}) \boldsymbol{k} = -5 + 4 + 4\cos\left(\frac{2\pi}{5}\right) = \sqrt{5} - 2 > 0.$$

An interior local maximum is impossible. Boundary induction completes the single-cycle cases.

If D has no cycles, join its directed paths cyclically into one cycle. This adds only nonnegative terms. Otherwise D has a proper cycle component. Split the vertices into that component and its complement. Both parts have at most seven vertices, and by induction their maxima are at most $4/27$. If their masses are s and $1 - s$, then their quadratic elementary symmetric sums are at most $4s^2/9$ and $4(1-s)^2/9$. There are no repeated terms between the parts. Consequently

$$F_D(\boldsymbol{x}) \leq \frac{4}{27}\left(s^3 + (1-s)^3\right) + \frac{4}{9}s(1-s) = \frac{4}{27}.$$

This settles all cases with at most eight vertices.

Suppose now that $n \geq 9$, and put $k = n - |E(D)|$. At $\boldsymbol{u}_n$,

$$F_D(\boldsymbol{u}_n) = b_n - \frac{k}{n^3}, \qquad \nabla^2 F_D(\boldsymbol{u}_n)|_{\mathbf{1}^\perp} \preceq -\frac{n-8}{n} I.$$

Indeed, the tangent Hessian of e_3 is $(-1 + 2/n)I$ at $\boldsymbol{u}_n$, while the absolute row sums of the repeated-term Hessian are at most $6/n$. Let $d_i = 3 - 2d_D^+(i) - d_D^-(i)$. Then

$$0 \leq d_i \leq 3, \qquad \sum_i d_i = 3k, \qquad \sum_i d_i^2 \leq 9k.$$

The projected gradient at $\boldsymbol{u}_n$ therefore satisfies

$$\|\boldsymbol{g}\|^2 \leq \frac{9k(n-k)}{n^5}.$$

Lemma 4.12 bounds every interior stationary value by

$$b_n + \frac{1}{n^3}\left(-k + \frac{6k(n-k)}{n(n-8)}\right). \tag{10.3}$$

For $n \geq 14$, the parenthesis is nonpositive. For $9 \leq n \leq 13$, maximizing this quadratic in k bounds (10.3) by

$$b_n + \frac{(14-n)^2}{24(n-8)n^2} < \frac{4}{27}.$$

To verify the last inequality, use

$$\frac{4}{27} - b_n = \frac{(n-3)(24-n)}{54n^2}, \qquad 4(n-8)(n-3)(24-n) \geq 264 > 225 \geq 9(14-n)^2.$$

Finally $b_n \leq 4/27$ for $3 \leq n \leq 24$, and b_n increases for $n \geq 6$. The inductive bounds on proper supports are thus no larger than the right side of (10.2).

For a cycle cover, uniform weights attain b_n, and any arc with weights $2/3, 1/3$ attains $4/27$. This proves equality. On the boundary at $n = 25$, every support has size at most 24, so its value is at most $4/27$; an arc attains that value. □

10.1.2 An eight-cycle base palette

Use the color set $\mathbb{Z}_{25}$. Let the directed graph D consist of the following eight cycles, in the displayed order:

$$(15, 19, 20, 24), \quad (0, 4, 5), \quad (9, 10, 14), \quad \text{and} \quad (5j+1, 5j+2, 5j+3) \qquad \text{for } 0 \leq j < 5.$$

They partition all 25 colors. Let σ be the successor permutation and put

$$\mathcal{H} = \big\{\{5j, 5j+1, 5j+2\}, \{5j+2, 5j+3, 5j+4\} \colon 0 \leq j < 5\big\}.$$

The base palette P contains the increasing tuple on every rainbow support outside $\mathcal{H}$, together with all tuples $(i, \sigma(i), i)$. Its polynomial is

$$\Lambda_P(\boldsymbol{x}) = e_3(\boldsymbol{x}) + \sum_i x_i^2 x_{\sigma(i)} - \sum_{e \in \mathcal{H}} \prod_{i \in e} x_i. \tag{10.4}$$

For the ith cycle, of length d_i, put

$$R_i(\boldsymbol{A}) = \sum_{u \to w \text{ in cycle } i} A_u^2 A_w.$$

Thus $d_1 = 4$ and $d_2 = \cdots = d_8 = 3$.

We fix the following rational constants:

$$a = \frac{23}{1000}, \quad b = \frac{56}{1000}, \quad \ell = \frac{37}{1000}, \quad u = \frac{42}{1000}, \quad K = 10^7, \quad Y = \frac{183}{10^7}.$$

A pointwise bound on the excess.

Lemma 10.4. *The base palette satisfies*

$$\max_{\partial S([25])} \Lambda_P = \frac{4}{27}, \tag{10.5}$$

and, for every $\boldsymbol{A} \in S([25])$,

$$\Lambda_P(\boldsymbol{A}) - \frac{4}{27} \leq \frac{1}{2} \min_i A_i^2 A_{\sigma(i)}. \tag{10.6}$$

Proof. The upper bound on the boundary follows from $\Lambda_P \leq F_D$ and Lemma 10.3. Any retained repeated tuple, with its colors weighted $2/3, 1/3$, attains $4/27$.

The deletion degrees of $\mathcal{H}$ are two at the five colors congruent to 2 modulo 5 and one at the remaining colors. No two degree-two colors are adjacent in D. At the uniform point $\boldsymbol{z}_0 = (1/25, \ldots, 1/25)$, the tangent Hessian of F_D is at most $-17I/25$. For positive x, y, z,

$$\nabla^2(xyz) + \operatorname{diag}\left(\frac{yz}{x}, \frac{xz}{y}, \frac{xy}{z}\right) = xyz \begin{pmatrix} \frac{1}{x} \\ \frac{1}{y} \\ \frac{1}{z} \end{pmatrix} \begin{pmatrix} \frac{1}{x} & \frac{1}{y} & \frac{1}{z} \end{pmatrix} \succeq 0. \tag{10.7}$$

Hence deleting $\mathcal{H}$ increases this upper Hessian bound by at most $2I/25$.

Fix an arc $i \to \sigma(i)$ and define $J_i = \Lambda_P - \frac{1}{2}x_i^2 x_{\sigma(i)}$. At $\boldsymbol{z}_0$, the largest eigenvalue of $-\frac{1}{2}\nabla^2(x_i^2 x_{\sigma(i)})$ is $(\sqrt{5}-1)/50 < 1/25$. Thus

$$\nabla^2 J_i(\boldsymbol{z}_0)|_{\mathbf{1}^\perp} \preceq -\frac{14}{25}I.$$

Write $d_{\mathcal{H}}(j)$ for the degree of color j in $\mathcal{H}$ and put $D_j \coloneqq d_{\mathcal{H}}(j) + \mathbf{1}_{j=i} + \frac{1}{2}\mathbf{1}_{j=\sigma(i)}$. Then

$$\sum_j D_j = \frac{63}{2}, \qquad \sum_j D_j^2 = 40 + 2d_{\mathcal{H}}(i) + d_{\mathcal{H}}(\sigma(i)) + \frac{5}{4} \le \frac{185}{4}.$$

The projected gradient $\boldsymbol{g}_i$ of J_i at $\boldsymbol{z}_0$ has squared norm at most

$$\frac{1}{25^4}\left(\frac{185}{4} - \frac{63^2}{100}\right) = \frac{164}{9765625}.$$

Also $J_i(\boldsymbol{z}_0) = 4629/31250$. By Lemma 4.12, every interior stationary value of J_i is at most

$$\frac{4629}{31250} + \frac{25}{21}\frac{164}{9765625} = \frac{4}{27} - \frac{23}{147656250} < \frac{4}{27}.$$

On the boundary $J_i \le \Lambda_P \le 4/27$. Thus $J_i \le 4/27$ throughout the simplex. Applying this to every arc proves (10.6). □

A test weighting and localization.

Lemma 10.5. *The base palette satisfies*

$$\theta < \lambda(P) < \frac{4}{27} + Y. \tag{10.8}$$

Every $\boldsymbol{A}$ with $\Lambda_P(\boldsymbol{A}) > 4/27$ satisfies

$$a < A_i < b \qquad \text{for every } i \in \mathbb{Z}_{25}. \tag{10.9}$$

Moreover, if $\boldsymbol{\alpha} \in [0, 1/40]^8$ and

$$F_{\boldsymbol{\alpha}} = \Lambda_P - \sum_{i=1}^{8} \alpha_i R_i,$$

then every global maximizer of $F_{\boldsymbol{\alpha}}$ of value greater than $4/27$ satisfies

$$\ell < A_i < u \qquad \text{for every } i \in \mathbb{Z}_{25}. \tag{10.10}$$

Proof. Take the rational weighting

$$z_i = \begin{cases} \dfrac{96}{2500}, & i \equiv 2 \pmod 5, \\ \dfrac{101}{2500}, & \text{otherwise.} \end{cases}$$

Call the five colors in the first case the centers. Each has two noncenter neighbors, and no noncenter has two center neighbors. Giving centers weight p and noncenters weight q yields

$$\Lambda_P = e_3 + 15q^3 - 5pq^2 + 5p^2q.$$

Substitute $p = 1/25 - 4t$, $q = 1/25 + t$, and $t = 1/2500$. We obtain

$$\Lambda_P(\boldsymbol{z}) = \frac{463019403}{3125000000}, \qquad \Lambda_P(\boldsymbol{z}) - \theta = \frac{87518307}{83028125000000} > 0, \qquad \Lambda_P(\boldsymbol{z}) - \frac{4}{27} = \frac{1523881}{84375000000}.$$

Let $\boldsymbol{g} = \operatorname{proj}_{\mathbf{1}^\perp} \nabla \Lambda_P(\boldsymbol{z})$. Differentiating the terms in (10.4) gives the following complete list of its coordinates:

Color type	Number of colors	$31250000\, g_i$
Center	5	1191
Predecessor of a center	5	-2854
Successor of a center	5	-2729
Neither	10	2196

Consequently

$$\|\boldsymbol{g}\|^2 = \frac{2665607}{19531250000000} < \left(\frac{37}{100000}\right)^2.$$

For example, before projection the derivative at color i is

$$e_2(\boldsymbol{z}|_{\mathbb{Z}_{25}\setminus\{i\}}) + 2z_i z_{\sigma(i)} + z_{\sigma^{-1}(i)}^2 - \sum_{\substack{e\in\mathcal{H}\\ i\in e}} \prod_{j\in e\setminus\{i\}} z_j.$$

This formula also verifies the four entries in the table.

For any positive weighting with smallest coordinate m and largest coordinate M, the tangent Hessian of e_3 is at most $(-1+2M)I$. The repeated terms have Hessian norm at most $6M$, and (10.7) bounds the deletion contribution by $2M^2 I/m$. Thus

$$\nabla^2 \Lambda_P|_{\mathbf{1}^\perp} \preceq \left(-1 + 8M + \frac{2M^2}{m}\right) I.$$

At $\boldsymbol{z}$ this is $-14203I/24000 \preceq -I/2$. Since $\Lambda_P(\boldsymbol{z}) > 4/27$ and (10.5) holds, every global maximizer is interior. Lemma 4.12 gives

$$0 < \lambda(P) - \frac{4}{27} \le \Lambda_P(\boldsymbol{z}) - \frac{4}{27} + \frac{4}{3}\|\boldsymbol{g}\|^2 = \frac{4810109051}{263671875000000} < Y.$$

This proves (10.8).

For (10.9), apply Lemma 10.2 with boundary bound $4/27$. If $\boldsymbol{h} = \boldsymbol{A} - \boldsymbol{z}$, the bracket in (10.1) is at most

$$3\left(\Lambda_P(\boldsymbol{z}) - \frac{4}{27}\right) + 2\|\boldsymbol{g}\|\|\boldsymbol{h}\| - \frac{14203}{48000}\|\boldsymbol{h}\|^2.$$

At $\|\boldsymbol{h}\| = 3/200$, replacing $\|\boldsymbol{g}\|$ by $37/100000$ bounds this by $-\frac{4658857}{3600000000000} < 0$. It decreases thereafter. Hence $\Lambda_P(\boldsymbol{A}) > 4/27$ forces $\|\boldsymbol{A} - \boldsymbol{z}\| < 3/200$, which implies $0.0234 < A_i < 0.0554$ and therefore (10.9).

Now let $\boldsymbol{A}$ be a global maximizer of $F_{\boldsymbol{\alpha}}$ above $4/27$. Since $\Lambda_P \ge F_{\boldsymbol{\alpha}}$, it is interior and satisfies $\|\boldsymbol{A} - \boldsymbol{z}\| < 3/200$. At $\boldsymbol{z}$,

$$-\nabla^2 F_{\boldsymbol{\alpha}}|_{\mathbf{1}^\perp} \succeq \left(\frac{14203}{24000} - \frac{6\cdot 101}{40\cdot 2500}\right) I = \frac{351439}{600000} I \succ \frac{7}{12} I.$$

Every coordinate of $\nabla(\sum_i \alpha_i R_i)$ lies between zero and $3(101/2500)^2/40$ at $\boldsymbol{z}$. A vector of length 25 whose entries lie in an interval of length w has centered squared norm at most $25w^2/4$. Thus the projected gradient $\boldsymbol{g}_{\boldsymbol{\alpha}}$ satisfies

$$\|\boldsymbol{g}_{\boldsymbol{\alpha}}\| \le \frac{37}{100000} + \frac{5}{2}\frac{3}{40}\left(\frac{101}{2500}\right)^2 < \frac{7}{10000}.$$

We next bound the Hessian variation along the tangent space. For $\boldsymbol{h} \perp \mathbf{1}$, the contribution of e_3 changes by $2\operatorname{diag}(\boldsymbol{h})$ on that space, with norm at most $2\|\boldsymbol{h}\|_\infty$. Each color has one predecessor and one successor, so the repeated terms, whose coefficients lie in $[39/40, 1]$, contribute at most $6\|\boldsymbol{h}\|_\infty$ by their absolute row sums. Each color belongs to at most 2 deleted triples, giving at most $4\|\boldsymbol{h}\|_\infty$ from the deletion terms. Consequently,

$$\left\| \left(\nabla^2 F_{\boldsymbol{\alpha}}(\boldsymbol{z}+\boldsymbol{h}) - \nabla^2 F_{\boldsymbol{\alpha}}(\boldsymbol{z})\right)\big|_{\mathbf{1}^\perp} \right\|_{\mathrm{op}} \le 12\|\boldsymbol{h}\|.$$

Put $t := \|\boldsymbol{A} - \boldsymbol{z}\| < 3/200$. Expanding the projected gradient at $\boldsymbol{z}$ and using stationarity at $\boldsymbol{A}$ gives

$$\frac{7}{12}t \le \|\boldsymbol{g}_{\boldsymbol{\alpha}}\| + 6t^2 < \frac{7}{10000} + 6t^2.$$

The function $7t/12 - 6t^2$ is increasing on $[0, 3/200]$, and

$$\frac{7}{12}\frac{1}{800} - 6\left(\frac{1}{800}\right)^2 - \frac{7}{10000} = \frac{19}{960000} > 0.$$

Hence $t < 1/800$, so $0.03715 < A_i < 0.04165$. This proves (10.10). All terminating decimals in these estimates denote exact rational numbers. □

10.2 Cyclic lifts and uniform control of the loss

10.2.1 A single lift and loss bounds

Suppose that a palette has tuples (u_i, u_{i+1}, u_i) along a directed cycle $u_0, \dots, u_{p-1}$, with indices modulo p. To lift this cycle by an integer factor $m \ge 2$, replace u_i by $u_{i,s}$, $s \in \mathbb{Z}_m$, and leave all other colors unchanged. First take the full blowup: a tuple is allowed exactly when its coarse-color image was allowed. Among the tuples above this cycle with exactly 2 distinct colors, retain only the arcs

$$u_{i,s} \to u_{i+1,s} \quad \text{for } 0 \le i < p-1, \qquad u_{p-1,s} \to u_{0,s+1}. \tag{10.11}$$

Thus the new repeated graph is a single cycle of length pm. All other tuples in the blowup remain unchanged.

The coarse-color projection is a palette homomorphism, since the operation only deletes tuples from the full blowup. The maps in (10.11) are bijections between successive color classes. Moreover, there is no directed cycle of length p in the new repeated graph. The lift with $p = 3$ and $m = 2$ is shown in Figure 15.

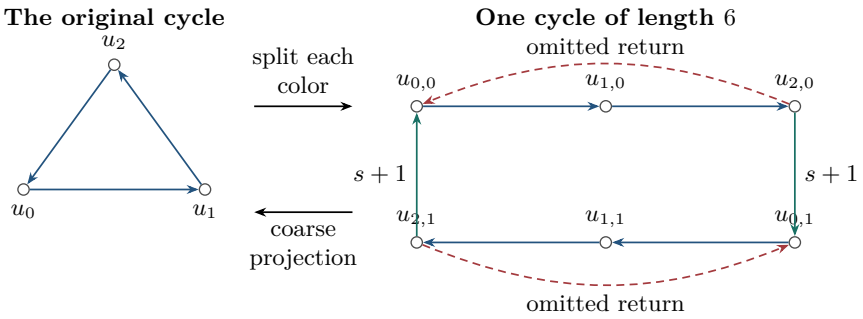


Figure 15: The cyclic lift in Equation (10.11), shown for $p = 3$ and $m = 2$. Each solid arc $v \to w$ represents the retained tuple (v, w, v). The layer changes only on returning from parent u_2 to parent u_0. Dashed arcs show the two omitted returns that would close 3-cycles.

If the parent masses are B_u, the loss at this step is exactly

$$L_0(\boldsymbol{x}) = \sum_u B_{\mathrm{succ}(u)} \sum_{v \text{ above } u} x_v^2 - R_{\mathrm{new}}(\boldsymbol{x}), \qquad R_{\mathrm{new}}(\boldsymbol{x}) = \sum_{v \to w} x_v^2 x_w.$$

Indeed, a parent repeat contributes $B_u^2 B_{\mathrm{succ}(u)}$ in the complete blowup. Removing its children with exactly 2 distinct colors subtracts $B_{\mathrm{succ}(u)} \sum_{v \text{ above } u} x_v^2$, and the retained arcs add R_{new} back. In particular,

$$L_0(\boldsymbol{x}) = \sum_{v \to w} x_v^2 \big(B_{\mathrm{type}(w)} - x_w\big) \ge 0. \tag{10.12}$$

The next two estimates bound individual weights when the loss is small.

Lemma 10.6 (A path inequality). *For $0 \le x, y \le 1$, let $\varphi(x) = x^2 - x^4/3$. Then*

$$x^2(1 - y) \ge \varphi(x) - \varphi(y). \tag{10.13}$$

Suppose a directed Hamilton cycle is partitioned into classes of size $m \geq 5$, and nonnegative weights y_v sum to one in each class. If

$$\sum_{v \to w} y_v^2(1 - y_w) < \frac{1}{2},$$

then $\max_v y_v < 4/5$.

Proof. The difference in (10.13) is $z^2/3 - yz + y^2 - y^4/3$, where $z = x^2 \in [0,1]$. For $y \leq 2/3$, its unrestricted minimum in z is $y^2/4 - y^4/3 \geq 0$. For $y \geq 2/3$, its minimum on $[0,1]$ occurs at $z = 1$ and equals

$$\frac{1}{3} - y + y^2 - \frac{y^4}{3} = \frac{(1-y)\big(y^3 + (1-y)^2\big)}{3} \geq 0.$$

This proves (10.13).

If a weight $M \geq 4/5$ occurs, another vertex in the same class has weight at most $(1-M)/(m-1) \leq 1/20$. Sum (10.13) along the simple directed path from the first vertex to the second. All omitted summands are nonnegative. Since φ is increasing on $[0,1]$, the total loss is at least

$$\varphi\left(\frac{4}{5}\right) - \varphi\left(\frac{1}{20}\right) = \frac{16031}{32000} > \frac{1}{2},$$

a contradiction. □

Lemma 10.7 (Small loss excludes a large atom)**.** *Consider a lift (10.11) of a parent cycle of length p. Suppose every parent mass is at least $a' > 0$. For $0 \leq s \leq 1$ put*

$$L_s(\boldsymbol{x}) = \sum_u B_{\mathrm{succ}(u)} \sum_{v \ \mathit{above}\ u} x_v^2 - (1-s) R_{\mathrm{new}}(\boldsymbol{x}).$$

If $L_s(\boldsymbol{x}) < (a')^3/32$, then $\max_v x_v < a'/4$.

Proof. Since $L_s \geq L_0$, it is enough to consider (10.12). Suppose a fine color v_0 has weight at least $a'/4$. Follow its successors through one round of the p parent classes. If its first successor has weight at most $a'/2$, the corresponding loss is at least $(a'/4)^2(a'/2) = (a')^3/32$. If a small successor first occurs later, its source has weight at least $a'/2$, and the loss is at least $(a')^3/8$.

Otherwise the next $p-1$ weights are all at least $a'/2$. The following successor is in the class of v_0 but is not v_0, because the lifted cycle has length $pm > p$. Its missing mass therefore includes x_{v_0}. The last loss is at least $(a'/2)^2(a'/4) = (a')^3/16$. All cases contradict the hypothesis. □

10.2.2 Successive lifts of an original cycle

Fix one original cycle of length $d \in \{3,4\}$ and prescribe its original masses $A_i \in [a,b]$. These masses are not required to sum to one; the other original cycles carry the remaining mass. Write

$$R_0(\boldsymbol{A}) = \sum_i A_i^2 A_{i+1}, \qquad r(\boldsymbol{A}) = \min_i A_i^2 A_{i+1}.$$

Lift the cycle successively by factors $m_1, m_2, \ldots$. Define

$$D_j = \prod_{h=1}^{j} m_h, \quad D_0 = 1, \qquad \alpha_j = \sum_{h=1}^{j} \frac{m_h - 1}{D_{h-1}^2 m_h^2}. \tag{10.14}$$

There are D_j descendants of each original color at level j. Let R_j be the repeated polynomial at that level, and let E_j be the cumulative loss relative to the full blowup of the original palette. If $\boldsymbol{B}$ is the pushforward of $\boldsymbol{x}$ to level $j-1$, then

$$E_j(\boldsymbol{x}) = E_{j-1}(\boldsymbol{B}) + \sum_u B_{\mathrm{succ}(u)} \sum_{v \text{ above } u} x_v^2 - R_j(\boldsymbol{x}), \qquad E_0 = 0.$$

Every added loss is nonnegative.

For $j \geq 0$, put $H_j(\boldsymbol{x}; s) \coloneqq E_j(\boldsymbol{x}) + sR_j(\boldsymbol{x})$. At level $j \geq 1$, dividing each parent weight equally among its m_j children gives $R_j = R_{j-1}/m_j^2$ and $E_j - E_{j-1} = (m_j - 1)R_{j-1}/m_j^2$, with the parent expressions evaluated at $\boldsymbol{B}$. Thus H_j reduces to $H_{j-1}(\boldsymbol{B}; T_{m_j}(s))$, where $T_m(s) \coloneqq (m - 1 + s)/m^2$. Controlling a range of s allows the induction hypothesis to be applied at this shifted parameter.

Let $\boldsymbol{z}_j$ give mass A_i/D_j to each descendant of original color i. Set

$$V_j(s) \coloneqq \left(\alpha_j + \frac{s}{D_j^2}\right) R_0(\boldsymbol{A}).$$

Direct substitution gives $R_j(\boldsymbol{z}_j) = R_0(\boldsymbol{A})/D_j^2$ and $H_j(\boldsymbol{z}_j; s) = V_j(s)$.

Theorem 10.8 (Uniform optimization under a loss bound)**.** *Suppose $m_1 \geq 40$ and $m_j \geq KD_{j-1}$ for every $j \geq 2$, where $K = 10^7$. For every $j \geq 1$ and every $0 \leq s \leq 1/(KD_j)$, any weighting with the prescribed original masses and*

$$H_j(\boldsymbol{x}; s) < \frac{r(\boldsymbol{A})}{2} \tag{10.15}$$

satisfies

$$H_j(\boldsymbol{x}; s) \geq V_j(s). \tag{10.16}$$

Equality implies $\boldsymbol{x} = \boldsymbol{z}_j$. If, in addition, $H_j(\boldsymbol{x}; s) \leq V_j(s) + a^3/(10^4 D_j^3)$, then

$$H_j(\boldsymbol{x}; s) - V_j(s) \geq \frac{a}{4D_j}\|\boldsymbol{x} - \boldsymbol{z}_j\|^2. \tag{10.17}$$

Proof. We prove the lower bound, its equality condition, and the quadratic estimate simultaneously by induction on j.

The first lift. Put $m = m_1$, $y_{i,t} = x_{i,t}/A_i$, and $w_{i,t} = y_{i,t} - 1/m$. Because $s \geq 0$ and $R_1 \geq 0$, the loss bound gives

$$E_1(\boldsymbol{x}) = \sum_i A_i^2 A_{i+1} \sum_t y_{i,t}^2 (1 - y_{\mathrm{succ}(i,t)}) < \frac{r(\boldsymbol{A})}{2}.$$

After division by $r(\boldsymbol{A})$, Lemma 10.6 yields $M := \max y_{i,t} < 4/5$. Define

$$W = \sum_i A_i^2 A_{i+1} \sum_t w_{i,t}^2.$$

Since every successor map is bijective and $\sum_t w_{i,t} = 0$, expansion at the uniform allocation gives

$$E_1(\boldsymbol{x}) - \frac{m-1}{m^2} R_0(\boldsymbol{A}) = \sum_i A_i^2 A_{i+1} \left[\left(1 - \frac{1}{m}\right) \sum_t w_{i,t}^2 - \frac{2}{m} \sum_t w_{i,t} w_{\mathrm{succ}(i,t)} - \sum_t w_{i,t}^2 w_{\mathrm{succ}(i,t)}\right].$$

The ratio of successive coefficients is at most $(b/a)^2$. Using $2pq \leq p^2 + q^2$ for the quadratic cross terms and $w_{\mathrm{succ}(i,t)} \leq M - 1/m$ for the cubic term, we obtain

$$\begin{aligned} E_1(\boldsymbol{x}) - \frac{m-1}{m^2} R_0(\boldsymbol{A}) &\geq \left[1 - M - \frac{1}{m}\left(1 + \frac{b^2}{a^2}\right)\right] W \\ &\geq \left[\frac{1}{5} - \frac{1}{40}\left(1 + \frac{56^2}{23^2}\right)\right] W = \frac{567}{21160} W \geq \frac{a}{50}\|\boldsymbol{x} - \boldsymbol{z}_1\|^2. \end{aligned} \tag{10.18}$$

The last inequality uses $W = \sum_i A_{i+1} \sum_t (x_{i,t} - A_i/m)^2 \geq a\|\boldsymbol{x} - \boldsymbol{z}_1\|^2$.

At $\boldsymbol{z}_1$, the partial derivatives of R_1 are constant within each original color class. Since $\boldsymbol{x} - \boldsymbol{z}_1$ sums to zero in each such class, the linear term in the expansion of R_1 vanishes. Its half-Hessian has operator norm at most $3b/m$. If $\boldsymbol{h} = \boldsymbol{x} - \boldsymbol{z}_1$, then $R_1(\boldsymbol{h}) \geq -(b/m)\|\boldsymbol{h}\|^2$, because $h_v \geq -b/m$ and every vertex has one outgoing arc. The exact cubic expansion therefore gives

$$R_1(\boldsymbol{x}) - R_1(\boldsymbol{z}_1) \geq -\frac{4b}{m}\|\boldsymbol{x} - \boldsymbol{z}_1\|^2.$$

Combining this with (10.18),

$$H_1(\boldsymbol{x}; s) - V_1(s) \geq \left(\frac{a}{50} - \frac{4sb}{m}\right) \|\boldsymbol{x} - \boldsymbol{z}_1\|^2 \geq \frac{a}{100}\|\boldsymbol{x} - \boldsymbol{z}_1\|^2 \geq \frac{a}{4m}\|\boldsymbol{x} - \boldsymbol{z}_1\|^2.$$

Here $s \leq 1/(Km)$, $m \geq 40$, and $b/a < 5/2$ verify both numerical inequalities. This proves all assertions at the first level.

The inductive step: localization of the parent. Let $j \geq 2$, put $D = D_{j-1}$, $m = m_j$, and let $\boldsymbol{B}$ be the parent pushforward. We initially restrict attention to an $\boldsymbol{x}$ satisfying both (10.15) and

$$H_j(\boldsymbol{x}; s) \leq V_j(s) + \frac{a^3}{10^4(Dm)^3}. \tag{10.19}$$

Any point violating (10.16) satisfies this additional inequality, so it suffices to prove the lower bound under this restriction. The quadratic assertion already assumes it.

The additional loss

$$L_s(\boldsymbol{x}) = \sum_u B_{\mathrm{succ}(u)} \sum_{v \text{ above } u} x_v^2 - (1-s)R_j(\boldsymbol{x})$$

is nonnegative. In particular, $E_{j-1}(\boldsymbol{B}) < r(\boldsymbol{A})/2$, so the inductive lower bound at parameter zero applies. Since $0 \leq s \leq 1$, we have $T_m(s) \leq 1/m$. Also $R_0(\boldsymbol{A}) \leq db^3$ and $m \geq KD$. Therefore

$$E_{j-1}(\boldsymbol{B}) \leq V_{j-1}(0) + \frac{1}{D^3}\left(\frac{db^3}{K} + \frac{a^3}{10^4 m^3}\right) < V_{j-1}(0) + \frac{a^3}{10^4 D^3}.$$

For the strict inequality, note that

$$K \geq 20000d\left(\frac{b}{a}\right)^3, \qquad \frac{db^3}{K} \leq \frac{a^3}{20000}, \qquad \frac{a^3}{10^4 m^3} \leq \frac{a^3}{80000}.$$

The bound on K follows from $d \leq 4$ and $b/a < 5/2$. The parent's quadratic estimate now yields

$$\|\boldsymbol{B} - \boldsymbol{z}_{j-1}\| \leq \frac{a}{50D}, \qquad \frac{a}{2D} \leq B_u \leq \frac{2b}{D}.$$

Here the quadratic estimate gives $\|\boldsymbol{B} - \boldsymbol{z}_{j-1}\|^2 \leq a^2/(2500D^2)$ after division by $a/(4D)$.

The inductive step: localization within fibers. Subtracting the inductive lower bound $E_{j-1}(\boldsymbol{B}) \geq V_{j-1}(0)$ from (10.19) gives

$$0 \leq L_s(\boldsymbol{x}) \leq V_j(s) - V_{j-1}(0) + \frac{a^3}{10^4(Dm)^3} = \frac{T_m(s)R_0(\boldsymbol{A})}{D^2} + \frac{a^3}{10^4(Dm)^3} < \frac{a^3}{10^4 D^3} < \frac{1}{32}\left(\frac{a}{2D}\right)^3.$$

The last two inequalities use the preceding bounds on $T_m(s)$, $R_0(\boldsymbol{A})$, and m. Apply Lemma 10.7 with $a' = a/(2D)$ to obtain

$$M_{\boldsymbol{x}} := \max_v x_v < \frac{a}{8D}.$$

Let $\boldsymbol{z}(\boldsymbol{B})$ be uniform within each parent class and put $\boldsymbol{h} = \boldsymbol{x} - \boldsymbol{z}(\boldsymbol{B})$. The fiber gradient of L_s at $\boldsymbol{z}(\boldsymbol{B})$ vanishes: all vertices above a fixed parent have equal source and predecessor contributions. The diagonal quadratic part has coefficient at least $a/(2D)$. The half-Hessian of R_j at $\boldsymbol{z}(\boldsymbol{B})$ has norm at most $3M_{\boldsymbol{z}(\boldsymbol{B})} \leq 6b/(Dm)$. Finally

$$R_j(\boldsymbol{h}) = \sum_{v \to w} h_v^2 h_w \leq M_{\boldsymbol{x}}\|\boldsymbol{h}\|^2,$$

since $h_w \leq x_w$ and the repeated graph has outdegree one. The exact cubic expansion gives

$$L_s(\boldsymbol{x}) - L_s(\boldsymbol{z}(\boldsymbol{B})) \geq \left(\frac{3a}{8D} - \frac{6b}{Dm}\right)\|\boldsymbol{h}\|^2 \geq \frac{a}{4D}\|\boldsymbol{h}\|^2. \tag{10.20}$$

We used $m \geq 48b/a$, which is implied by $m \geq KD$.

Verifying the parent hypothesis at the shifted parameter. At the uniform allocation, $L_s(\boldsymbol{z}(\boldsymbol{B})) = T_m(s)R_{j-1}(\boldsymbol{B})$, and

$$0 \le T_m(s) \le \frac{1}{m} \le \frac{1}{KD}.$$

Moreover, (10.20) proves the necessary loss bound at this parameter:

$$E_{j-1}(\boldsymbol{B}) + T_m(s)R_{j-1}(\boldsymbol{B}) \le H_j(\boldsymbol{x}; s) < \frac{r(\boldsymbol{A})}{2}.$$

Thus the induction hypothesis at $T_m(s)$ gives

$$H_j(\boldsymbol{x}; s) \ge V_{j-1}(T_m(s)) = V_j(s).$$

This proves the lower bound under (10.19), and hence for every point satisfying the loss bound.

For the quadratic assertion, the excess of the shifted parent above $V_{j-1}(T_m(s))$ is at most $a^3/(10^4(Dm)^3) \le a^3/(10^4D^3)$. Its inductive quadratic estimate, together with (10.20), yields

$$H_j(\boldsymbol{x}; s) - V_j(s) \ge \frac{a}{4D}\big(\|\boldsymbol{B} - \boldsymbol{z}_{j-1}\|^2 + \|\boldsymbol{h}\|^2\big).$$

The vector $\boldsymbol{h} = \boldsymbol{x} - \boldsymbol{z}(\boldsymbol{B})$ sums to zero in each fiber, whereas $\boldsymbol{z}(\boldsymbol{B}) - \boldsymbol{z}_j$ is constant on each fiber. Orthogonality therefore gives

$$\|\boldsymbol{x} - \boldsymbol{z}_j\|^2 = \frac{1}{m}\|\boldsymbol{B} - \boldsymbol{z}_{j-1}\|^2 + \|\boldsymbol{h}\|^2. \tag{10.21}$$

This proves (10.17). Equality in the lower bound forces both components to vanish. This completes the induction. □

10.3 The value of a finite refinement

We refine the eight original cycles independently. A cycle may be left unrefined. If it is used, its first factor is at least 40, and every later factor is at least KD, where D is its current cumulative scale. Each operation has a coarse-color homomorphism to its predecessor.

For cycle i, let D_i be its current scale and α_i its sum in (10.14); for an unrefined cycle set $D_i = 1$ and $\alpha_i = 0$. For a refined coordinate with factors $m_1, \dots, m_j$,

$$\alpha_i + \frac{s}{D_i^2} = T_{m_1} \circ \cdots \circ T_{m_j}(s).$$

Since every T_m with $m \ge 2$ maps $[0, 1/3]$ into itself, we have

$$0 \le \alpha_i + \frac{1}{3D_i^2} \le T_{m_1}\left(\frac{1}{3}\right) \le \frac{1}{40}. \tag{10.22}$$

For an unrefined coordinate, its next first step also lies in $[0, 1/40]$. In particular, every parameter vector obtained by these refinements belongs to $[0, 1/40]^8$.

Set

$$\Omega = S([25]) \cap [\ell, u]^{25}, \qquad G(\boldsymbol{\alpha}) = \max_{\boldsymbol{A} \in \Omega} \left\{ \Lambda_P(\boldsymbol{A}) - \sum_{i=1}^{8} \alpha_i R_i(\boldsymbol{A}) \right\}.$$

The function G is defined on the entire cube, including parameter vectors not obtained by these refinements.

Proposition 10.9. *Let Q be obtained by a finite sequence of refinements satisfying the factor bounds above, with parameter vector $\boldsymbol{\alpha}$. Then*

$$\lambda(Q) = \max\left\{ \frac{4}{27}, G(\boldsymbol{\alpha}) \right\}. \tag{10.23}$$

The function G is continuous, $G(0) = \lambda(P)$, and

$$d_i \ell^3 h \le G(\boldsymbol{\alpha}) - G(\boldsymbol{\alpha} + h e_i) \le d_i u^3 h \tag{10.24}$$

whenever both arguments are in the cube and $h \ge 0$.

Proof. Fix a weighting $\boldsymbol{x}$ of Q and let $\boldsymbol{A}$ be its pushforward to the original 25 colors. All original rainbow tuples have been copied in full. Losses from different original cycles involve disjoint fine variables once $\boldsymbol{A}$ is fixed. Hence the exact polynomial is

$$\Lambda_Q(\boldsymbol{x}) = \Lambda_P(\boldsymbol{A}) - \sum_{i=1}^{8} E_i(\boldsymbol{x}^{(i)}),$$

where $E_i = 0$ for an unrefined cycle and every loss is nonnegative.

Suppose $\Lambda_Q(\boldsymbol{x}) > 4/27$. Then $\Lambda_P(\boldsymbol{A}) > 4/27$, so Lemma 10.5 places all A_i in $[a, b]$. For every refined cycle i, the pointwise bound (10.6) implies

$$E_i(\boldsymbol{x}^{(i)}) \le \sum_k E_k(\boldsymbol{x}^{(k)}) < \Lambda_P(\boldsymbol{A}) - \frac{4}{27} \le \frac{1}{2} \min_{u\to w\in D} A_u^2 A_w \le \frac{1}{2} \min\left\{A_u^2 A_w \colon u \to w \text{ lies in cycle } i\right\}.$$

Thus Theorem 10.8, at parameter $s = 0$, applies in each refined direction and gives

$$\Lambda_Q(\boldsymbol{x}) \le \Lambda_P(\boldsymbol{A}) - \sum_i \alpha_i R_i(\boldsymbol{A}).$$

Conversely, giving equal mass to the descendants of each original color attains this coarse expression for every $\boldsymbol{A} \in S([25])$. A retained repeated tuple always attains $4/27$. Therefore

$$\lambda(Q) = \max\left\{\frac{4}{27}, \max_{\boldsymbol{A}\in S([25])} \left(\Lambda_P(\boldsymbol{A}) - \sum_i \alpha_i R_i(\boldsymbol{A})\right)\right\}.$$

If the coarse maximum exceeds $4/27$, all its maximizers belong to Ω by (10.10). If it does not, both the full and restricted coarse maxima are at most $4/27$. This proves (10.23) in all cases.

At every $\boldsymbol{A} \in \Omega$, one has $d_i \ell^3 \le R_i(\boldsymbol{A}) \le d_i u^3$. Apply Equation (2.3) while varying α_i, with coefficient $-R_i(\boldsymbol{A})$, to obtain (10.24). Summing these bounds over the 8 coordinates gives continuity on the cube. Since $\lambda(P) > 4/27$, a maximizer of Λ_P is in Ω, so $G(0) = \lambda(P)$. □

10.4 Refinements with prescribed limiting values

We choose successive refinements satisfying the factor bounds above so that the values of G decrease to a prescribed value.

In this subsection, abbreviate $L = 3\ell^3$ and $U = 4u^3$. The constants satisfy

$$U < 2L, \qquad \frac{39}{100}\ell^3 = \frac{1975467}{10^{11}} > Y = \frac{183}{10^7}.$$

The first inequality compares the maximal and minimal coordinate slopes in (10.24). The second provides the initial range, using one 4-cycle and four 3-cycles.

Lemma 10.10. *For every $\beta \in [\lambda(P) - Y, \lambda(P)]$, there is a sequence of finite refinements satisfying the factor bounds above, with parameter vectors $\boldsymbol{\alpha}^{(j)}$ such that*

$$G(\boldsymbol{\alpha}^{(j)}) \ge \beta, \qquad G(\boldsymbol{\alpha}^{(j)}) \downarrow \beta.$$

The sequence may terminate, in which case it is continued constantly.

Proof. The case $\beta = \lambda(P)$ uses no refinements. At any later stage let $r = G(\boldsymbol{\alpha}) - \beta$ be the residual, and stop if $r = 0$. For a chosen coordinate of current scale d, a new factor m changes its parameter by $a_m = (m-1)/(d^2m^2)$. Put $b_m = G(\boldsymbol{\alpha}) - G(\boldsymbol{\alpha} + a_m e_i)$. Every refinement satisfying these bounds gives a parameter vector in the cube by (10.22). The b_m are strictly decreasing to zero for $m \ge 3$.

Suppose the smallest allowed factor M satisfies $b_M > r$. Choose the least integer $m \ge M$ with $b_m \le r$. Then $m > M$ and $b_{m-1} > r$. By (10.24), the new residual and scale satisfy

$$0 \le r' < b_{m-1} - b_m \le \frac{U}{d^2m^2}, \qquad D' := dm \ge \frac{L}{2dr}. \tag{10.25}$$

For the first bound, use

$$\frac{m-2}{(m-1)^2}-\frac{m-1}{m^2}=\frac{m^2-3m+1}{m^2(m-1)^2}\le\frac{1}{m^2}.$$

For the second, use $b_m \ge La_m \ge L/(2d^2m)$ and $b_m \le r$.

Initial choices. Use previously unrefined coordinates in the order $1,2,\dots,8$. If factor 40 produces a decrease at most r, take it. Otherwise make the least-factor choice in (10.25), with $M=40$. There can be at most four factor-40 choices before the first least-factor choice. Indeed, if factor 40 were chosen for each of the first five coordinates, their total decrease would be at least

$$\frac{39}{100}\ell^3 > Y \ge \lambda(P)-\beta,$$

contrary to the rule of not crossing the target. Consequently, a least-factor choice occurs by coordinate five. Write its factor as $n_0 \ge 41$.

Unless the process terminates, use three more unrefined coordinates. They are available because the first least-factor choice occurred by coordinate five. Each of these three choices again uses (10.25). To see this, if the preceding factor was $n \ge 40$, its residual is at most $U/n^2 < 2L/40^2$, whereas an unused direction at factor 40 gives at least $39L/1600 > L/80$. If the three new factors are n_1, n_2, n_3, then

$$n_{j+1} > \frac{n_j^2}{4}, \qquad n_1 \ge 400, \quad n_2 \ge 40000, \quad n_3 \ge 4\cdot 10^8.$$

Here we used $U < 2L$ in the scale bound of (10.25).

Keep all coordinates except the last three fixed. Their sorted scales $d \le e \le N$ initially satisfy

$$e \ge \frac{d^2}{4}, \qquad N \ge \frac{e^2}{4}, \qquad N \ge 4\cdot 10^8, \qquad r \le \frac{U}{N^2}. \tag{10.26}$$

The first further refinement. Choose the direction of smallest scale d. Its least allowed factor is now $M = Kd$. We claim that $b_M > 2r$, so a least-factor choice is possible. The lower slope bound gives

$$b_M \ge \frac{L}{2Kd^3}.$$

From (10.26), $d \le (64N)^{1/4} < 3N^{1/4}$ and $d^3 < 24N^{3/4}$. Since $N \ge 4\cdot 10^8$,

$$N^{5/4} > 4\cdot 10^{10} > 192K, \qquad \text{and hence} \qquad N^2 > 8Kd^3.$$

Together with $U < 2L$, this implies $b_M > 2U/N^2 \ge 2r$. The least-factor choice therefore gives

$$D' > \frac{N^2}{4d} > \frac{N^{7/4}}{12} > 10^{13}, \qquad D' > N^{3/2}, \qquad r' \le \frac{U}{(D')^2}.$$

For example, $N^{1/4} > 100$ verifies the two growth comparisons. The old scales also satisfy $e \ge d^{3/2}$ and $N \ge e^{3/2}$, since the scales in (10.26) exceed 16. After this step the sorted scales are e, N, D'.

The infinite iteration. From now on maintain the invariant

$$d \le e \le N, \qquad e \ge d^{3/2}, \quad N \ge e^{3/2}, \quad N \ge 10^{13}, \quad r \le \frac{U}{N^2}. \tag{10.27}$$

Choose the direction of smallest scale. Then $d \le N^{4/9}$ and

$$\frac{N^2}{d^3} \ge N^{2/3} > 8K.$$

Thus $b_{Kd} > 2r$ by the preceding argument. Another least-factor choice yields

$$D' > \frac{N^2}{4d} \ge \frac{N^{14/9}}{4} \ge N^{3/2} > N.$$

The penultimate inequality follows from $N \ge 10^{13} > 4^{18}$. The new triple e, N, D' satisfies (10.27); the residual bound follows from (10.25).

Every chosen factor satisfies the required lower bound, and the values of G remain at least β. Either a residual is zero, or the maximum scale tends to infinity, forcing $0 \le r \le U/N^2 \to 0$. This proves the lemma. □

Theorem 10.11 (The critical interval)**.** *For the base palette P of Section 10.1.2,*

$$\left[\frac{4}{27}, \lambda(P)\right] \subseteq \mathcal{R}_3. \tag{10.28}$$

In particular, Theorem 10.1 holds.

Proof. Fix $\beta \in [4/27, \lambda(P)]$. Lemma 10.5 gives $0 \le \lambda(P) - \beta < Y$. Apply Lemma 10.10. At each stage of the refinement,

$$G(\boldsymbol{\alpha}) \ge \beta \ge \frac{4}{27}.$$

Proposition 10.9 identifies this value with the Lagrangian of the refined palette. Each update has the explicit coarse-color homomorphism of Section 10.2. Thus the constructed palettes form a projective sequence and their Lagrangians decrease to β. A finite termination is continued by a constant sequence. Finally $\lambda(P) > \theta$, so (10.28) contains the interval stated in Theorem 10.1. □

10.5 Interval extension and higher uniformities

We extend the ternary interval from Theorem 10.1 to $[4/27, 1]$ and then pass to higher uniformities.

A finite binary palette is a relation $B \subseteq D^2$; loops are allowed. Its polynomial and Lagrangian use the common definitions.

Lemma 10.12. *Fix a binary palette B and a nonempty set $J \subseteq \{1,2,3\}$ of size m. Put $q = m\lambda(B)$. There is a homomorphism-preserving operation on finite ternary palettes whose Lagrangian transformation is the function in (3.1), with the following formula*

$$\Theta_3\left(p, \frac{q}{3}\right) = \max_{0 \le s \le 1} \{ps^3 + qs^2(1-s)\} = \begin{cases} \dfrac{4q^3}{27(q-p)^2}, & p < \frac{q}{3}, \\ p, & p \ge \frac{q}{3}. \end{cases} \tag{10.29}$$

Consequently, $p \in \mathcal{R}_3$ implies $\Theta_3(p, q/3) \in \mathcal{R}_3$.

Proof. For a finite ternary palette P, use the color set $(C(P) \times D) \sqcup \{y\}$, where y is new. A tuple without y is allowed exactly when its old-color projection belongs to P. A tuple containing y once is allowed if the position of y is in J and the two remaining tags, in coordinate order, form a pair in B. Their old-color coordinates are unrestricted. No tuple contains y twice or three times. Distinct positions of y give disjoint sets of tuples, including when B has loops.

At total product weight s, the old-color and tag marginals bound the polynomial by $\lambda(P)s^3 + m\lambda(B)(1-s)s^2$. Product weights attain both bounds simultaneously, exactly as in Lemma 3.3. The variational formula is therefore $\Theta_3(\lambda(P), q/3)$, and its explicit value follows from that lemma.

A homomorphism $h : P \to P'$ extends by $(c, i) \mapsto (h(c), i)$ and $y \mapsto y$. It preserves both kinds of allowed tuples, even if h is noninjective. Applying the construction termwise to a projective sequence and using continuity of $p \mapsto \Theta_3(p, q/3)$ proves the last assertion. □

Proposition 10.13. *One has $[4/27, 1] \subseteq \mathcal{R}_3$.*

Proof. Let

$$A_k = \{(i,j) \colon 1 \le i < j \le k\}, \qquad B_k = \{(i,j) \colon i, j \in [k],\ i \ne j\},$$

and let $B_\circ = \{(1,1)\}$ on one color. The identity $\sum_{i<j} z_i z_j = (1 - \sum_i z_i^2)/2$ gives

$$\lambda(A_k) = \frac{k-1}{2k}, \qquad \lambda(B_k) = \frac{k-1}{k}, \qquad \lambda(B_\circ) = 1.$$

For each q supplied by Lemma 10.12, the map $p \mapsto \Theta_3(p, q/3)$ is continuous and increasing, fixes $q/3$, and satisfies $\Theta_3(p, q/3) > p$ for $p < q/3$, by (10.29). Thus Lemma 2.5 extends a realized interval $[c, b]$ to $[c, q/3)$ whenever $c < b < q/3$ and $\Theta_3(c, q/3) \le b$. If the initial interval is $[c, b)$, the last inequality must be strict. Each value in the extended interval is obtained by finitely many palette operations.

Starting with the closed interval $[4/27, \theta]$ from Theorem 10.11, apply this argument with the following choices. The left endpoint always remains $4/27$.

Relation	m	q	Old upper endpoint	$\Theta_3(4/27, q/3)$	New upper endpoint
A_{10}	1	9/20	θ	19683/132845	3/20
B_2	1	1/2	3/20	54/361	1/6
B_3	1	2/3	1/6	8/49	2/9
$B_\circ$	1	1	2/9	108/529	1/3
B_2	3	3/2	1/3	1458/5329	1/2
$B_\circ$	3	3	1/2	2916/5929	1

The first overlap is equality, which is permitted for the closed initial interval. The next five positive overlap margins are

$$\frac{3}{7220}, \quad \frac{1}{294}, \quad \frac{86}{4761}, \quad \frac{955}{15987}, \quad \frac{97}{11858}.$$

The resulting interval is $[4/27, 1)$. Every intermediate upper endpoint lies in a later interval. The value 1 is realized by the one-color palette containing its unique tuple, so the endpoint is included. □

To extend the interval to higher uniformities, recall $p_r = (1 - 1/r)^{r-1} = \alpha_r/\alpha_{r-1}$ from Part I.

Lemma 10.14. *There are three homomorphism-preserving operations from $(r-1)$-palettes to r-palettes with Lagrangian values*

$$\lambda(P), \qquad \frac{p_r}{r}\lambda(P), \qquad p_r\lambda(P).$$

If $0 < a \le 1/r$ and $[a, 1] \subseteq \mathcal{R}_{r-1}$, then

$$\left[\frac{p_r a}{r}, 1\right] \subseteq \mathcal{R}_r. \tag{10.30}$$

Proof. For the first operation, append an unrestricted old-color coordinate. The new polynomial is $\Lambda_P(\boldsymbol{x}) \sum_{c \in C(P)} x_c$, which agrees with Λ_P on the probability simplex.

For the other operations, insert a new color y in one fixed position or in all r positions. By Fact 8.3(ii), the first choice multiplies the Lagrangian by p_r/r. In the second, the position of the unique new color distinguishes the r copies, multiplying this value by r. The maps extend by fixing y. Applying the operations termwise to projective sequences gives three intervals in $\mathcal{R}_r$:

$$\left[\frac{p_r a}{r}, \frac{p_r}{r}\right], \qquad [p_r a, p_r], \qquad [a, 1].$$

The first two overlap because $ra \le 1$. The second and third overlap because

$$p_r = \left(1 - \frac{1}{r}\right)^{r-1} \ge \frac{1}{r} \ge a$$

by Bernoulli's inequality. Their union is (10.30). □

Completion of the proofs of Theorem 1.1(v) and Proposition 8.4. Proposition 10.13 gives $[4/27, 1] \subseteq \mathcal{R}_3$. Inductively suppose $[4/(r-1)^{r-1}, 1] \subseteq \mathcal{R}_{r-1}$. For $r \ge 4$,

$$(r-1)^{r-1} \ge 9(r-1) \ge 4r.$$

Thus $a = 4/(r-1)^{r-1} \le 1/r$, and Lemma 10.14 yields

$$\left[\frac{p_r a}{r}, 1\right] = \left[\frac{4}{r^r}, 1\right] \subseteq \mathcal{R}_r.$$

By Fact 8.2 and Proposition 2.3, these compatible palette sequences realize the upper interval in $\Pi^{(r)}_{r-2,\infty}$. The finite-attainment result proved in Section 9 gives the remaining part of Theorem 1.1(v). Every low value belongs to $\mathcal{D}_r \subseteq \mathcal{R}_r$, and every upper value has just been constructed in $\mathcal{R}_r$. This completes Theorem 1.1(v) and Proposition 8.4. □

11 Intermediate uniform densities: realization and construction

Let $r \ge 4$ and $1 \le j \le r-3$, and put $q \coloneqq j+1$ and $N \coloneqq \binom{r}{q}$. We use the conventional uniform density and notation from the introduction. Thus H is (d,η,j)-dense if

$$|H \cap \mathcal{K}_r(G)| \ge d|\mathcal{K}_r(G)| - \eta|V(H)|^r \qquad \text{for every } j\text{-graph } G \text{ on } V(H). \tag{11.1}$$

The following numerical gap supplies the lower bound in the classification.

Proposition 11.1. *For every $r \ge 4$ and $1 \le j \le r-3$, writing $N = \binom{r}{j+1}$, one has*

$$\Pi^{(r)}_{j,\infty} \subseteq \{0\} \cup [N^{-N}, 1].$$

Theorem 11.2. *There exists an absolute constant $\varepsilon > 0$ such that the following holds. Suppose that $r \ge 4$, $1 \le j \le r-3$, and $\mathcal{S} = \binom{[r]}{j}$, with $j \ge 2$ or $r \ge 6$. Then $[N^{-N}, (1+\varepsilon)N^{-N}] \subseteq \Pi^{(r)}_{\mathcal{S},\infty}$, where $N \coloneqq \binom{r}{j+1}$.*

The initial interval comes from the common 10-uniform construction with only 3 indexed color positions. The regularity proofs of the realization theorems are in Section 12.

11.1 Realization theorems

We use the common finite and countable palette definitions. Coordinates will now be assigned to q-faces rather than to facets. A countable palette Lagrangian is the supremum over finitely supported probability weightings, equivalently over all probability weightings by truncation and renormalization.

Fix $2 \le q \le r-1$, put $j \coloneqq q-1$ and $N \coloneqq \binom{r}{q}$, and index the coordinates of P by $\binom{[r]}{q}$. For palettes constructed with numbered coordinates, we specify their assignment to these faces. Given a coloring $\chi\colon \binom{V(H)}{q} \to C(P)$ of an ordered r-graph, an edge $e = \{v_1 < \cdots < v_r\}$ is *admissible* if

$$\big(\chi(\{v_i : i \in A\})\big)_{A \in \binom{[r]}{q}} \in P.$$

An unordered r-graph is *P-colorable* if some vertex order and some such coloring make every edge admissible. Let

$$\mathcal{F}_{r,q}(P) \coloneqq \{F : F \text{ is a finite } r\text{-graph that is not } P\text{-colorable}\},$$

taking one representative of each isomorphism class. Every edgeless r-graph is colorable. Since P-colorability is closed under taking subhypergraphs, an r-graph is $\mathcal{F}_{r,q}(P)$-free if and only if it is P-colorable. A palette homomorphism $P \to Q$ makes every P-colorable r-graph Q-colorable.

The counts in (11.1) concern unordered r-sets; palette coordinates are indexed by faces.

Suppose first that P is finite, and put $\mathcal{Q} \coloneqq \binom{[r]}{q}$. For a test system $\mathcal{S}$, a test can prescribe colors on those $A \in \mathcal{Q}$ contained in some $S \in \mathcal{S}$; different relative vertex orders permute these faces.

For $B \subseteq \mathcal{Q}$, let $\mathrm{Orb}(B) \coloneqq \{\sigma(B) : \sigma \in S_r\}$. Given a probability weighting $\boldsymbol{w}$ and a supported assignment $\boldsymbol{a}$ on $I \subseteq \mathcal{Q}$, define

$$E_{I,\boldsymbol{a}}(P,\boldsymbol{w}) \coloneqq \sum_{\substack{\boldsymbol{c} \in P \\ \boldsymbol{c}|_I = \boldsymbol{a}}} \prod_{A \in \mathcal{Q}\setminus I} w_{c_A}.$$

The *conditional capacity* is

$$m_B(P,\boldsymbol{w}) \coloneqq \min\left\{E_{I,\boldsymbol{a}}(P,\boldsymbol{w}) : I \in \mathrm{Orb}(B),\ \boldsymbol{a} \in (\operatorname{supp}\boldsymbol{w})^I\right\}, \qquad \lambda_B(P) \coloneqq \max_{\boldsymbol{w}} m_B(P,\boldsymbol{w}).$$

For $B = \varnothing$ this is $\lambda(P)$. By Lemma 5.3, the maximum is attained and $\lambda_B(P) \le \lambda_B(Q)$ whenever $P \to Q$. The orbit is induced by permutations of the vertex positions; arbitrary permutations of the N color coordinates need not preserve it.

Theorem 11.3. *Suppose that $r \ge 3$, $2 \le q \le r-1$, and $\mathcal{S}$ is a nonempty family of nonempty proper subsets of $[r]$ such that every $(q-1)$-subset of $[r]$ is contained in a member of $\mathcal{S}$. Put $\mathcal{Q} \coloneqq \binom{[r]}{q}$ and $B \coloneqq \{A \in \mathcal{Q} : A \subseteq S \text{ for some } S \in \mathcal{S}\}$. The following statements hold.*

(i) *Every finite palette P with coordinates indexed by $\mathcal{Q}$ satisfies*

$$\pi_{\mathcal{S}}(\mathcal{F}_{r,q}(P)) = \lambda_B(P).$$

(ii) *If finite palettes P_n with coordinates indexed by $\mathcal{Q}$ satisfy $P_{n+1} \to P_n$ and $\lambda_B(P_n) \to \alpha$, then*

$$\pi_{\mathcal{S}}\left(\bigcup_n \mathcal{F}_{r,q}(P_n)\right) = \alpha.$$

For the complete layer $\mathcal{S} = \binom{[r]}{q-1}$, we have $B = \varnothing$. The theorem and Table 1 give

$$\pi^{(r)}_{q-1}\big(\mathcal{F}_{r,q}(P)\big) = \lambda(P).$$

A palette consisting of one rainbow tuple has Lagrangian N^{-N} by AM–GM, so $N^{-N} \in \Pi^{(r)}_{q-1,\infty}$. If $P_{n+1} \to P_n$ and $\lambda(P_n) \to t$, part (ii) realizes t by the union of their canonical forbidden families.

The next two realization theorems use the same finite comparison rules. They differ in the positions at which indexed colors may interact. In both cases, the upper bound is uniform in the number of colors used by the host.

Fix $r \geq 3$ and $2 \leq q \leq r-1$, put $N := \binom{r}{q}$, and assign the palette coordinates to q-faces. Let P be a finite or countable palette with color set $C(P) = F_0 \sqcup (I \times A_0)$, where F_0, A_0 are finite and I is countably infinite. Fix a simple graph Γ on I of maximum degree at most Δ, a finite set of named indices $R \subseteq I$, and finitely many directed relations $R_1, \dots, R_b \subseteq I^2$, each supported on

$$\{(i,i) \colon i \in I\} \cup \{(i,k) \colon ik \in E(\Gamma)\}.$$

Membership in P is determined by the finite color or indexed type at each position, equality to named indices, and binary comparisons between distinct tuple positions. These comparisons may use equality and adjacency in Γ, the relations $R_1, \dots, R_b$, and arbitrary Boolean combinations of these conditions. No additional unary predicates on unnamed indices are permitted. For distinct nonadjacent unnamed indices, equality, adjacency, and all the R_t tests are false.

For the disjoint-face case, let J be a graph with $V(J) \subseteq \binom{[r]}{q}$. We impose the following conditions.

(i) Indexed colors appear only in positions in $V(J)$.

(ii) If $ab \in E(J)$, then $a \cap b = \varnothing$.

(iii) Binary comparisons are made only between indices at positions joined in J.

When A_0 is empty, every finite palette satisfies these conditions. The graph J specifies which indexed positions may interact; the comparison rules above specify what information about their indices may be used.

Theorem 11.4. *Every palette satisfying these conditions satisfies*

$$\pi^{(r)}_{q-1}(\mathcal{F}_{r,q}(P)) = \lambda(P).$$

More precisely, for every $\varepsilon > 0$ there are $c > 0, n_0$ such that every P-colorable r-graph H of order $n \geq n_0$ has a $(q-1)$-graph G on $V(H)$ with

$$|\mathcal{K}_r(G)| \geq cn^r, \qquad |H \cap \mathcal{K}_r(G)| \leq (\lambda(P) + \varepsilon)|\mathcal{K}_r(G)|.$$

At $q = 2$, the test is a vertex subset of size at least cn, after changing c. The constants depend on ε and the fixed data but not on how many indexed colors the host coloring uses.

We now allow the compared q-faces to share a common initial $(q-1)$-element core. The colors on these faces need not be independent. After fixing the core vertices, we compare the relevant counts with those obtained by choosing the private vertices independently. We select a positive proportion of cores and bound a finite palette polynomial separately on each one.

Theorem 11.5. *Let $3 \leq q \leq r-1$, put $N := \binom{r}{q}$, and assign the palette coordinates to $\binom{[r]}{q}$. Suppose P is a countable N-coordinate palette with the color set and finite comparison rules specified above. Indexed colors may occur only at coordinates assigned to faces $[q-1] \cup \{v\}$ with $q \leq v \leq r$; finite colors are also allowed there. Then*

$$\pi^{(r)}_{q-1}(\mathcal{F}_{r,q}(P)) = \lambda(P).$$

For each $\varepsilon > 0$ every sufficiently large P-colorable host H has a $(q-1)$-graph G on $V(H)$ such that

$$|\mathcal{K}_r(G)| \geq c|V(H)|^r, \qquad |H \cap \mathcal{K}_r(G)| \leq (\lambda(P) + \varepsilon)|\mathcal{K}_r(G)|.$$

The positive constant c and the required lower bound on $|V(H)|$ depend on the finite comparison data and ε, not on the number of indices used by the coloring.

11.2 The interval construction and the classification

Use the countable 10-graphs $H_{n,\boldsymbol{\omega}}$ from Proposition 3.10. Replace each unordered edge of the 10-graph defined in (3.2) by a single ordered tuple, choosing the orderings in the proof so that indexed colors occur only in positions $8, 9, 10$. In a base edge, put its 5 t-colors first, then u_b, then a fixed ordering of its 3 finite auxiliary colors, then its component color. This final color may be the finite color z or an indexed color. In an additional edge corresponding to the cubic component edge $\{w_1, w_2, w_3\}$, use

$$(u_1, u_2, \alpha_{2j-1}, \beta_{2j-1}, \alpha_{2j}, \beta_{2j}, z, w_1, w_2, w_3).$$

Theorem 11.6. *For every $n \geq 2Q$ and $\boldsymbol{\omega} \in \mathcal{A}^{\mathbb{N}}$, there is a countable 10-coordinate palette $L_{n,\boldsymbol{\omega}}$ with exactly one ordered tuple for each underlying edge of the 10-graph in (3.2), and*

$$\lambda(L_{n,\boldsymbol{\omega}}) = 10^{-10}\Phi_n(\boldsymbol{\omega}).$$

Each palette $L_{n,\boldsymbol{\omega}}$ has only 3 indexed positions, a finite auxiliary color set depending on n, and a finite collection of comparison rules. For each fixed n, the finite colors, indexed types, and the number and form of the comparison rules are independent of $\boldsymbol{\omega}$; the index relations are determined by its tree. Off the diagonal, its basic binary index relations are supported on a common graph of bounded degree.

Proof. We order the 3 colors w_1, w_2, w_3 of each component edge using only their types, the relations $k = f_\ell(i)$, equality, tree adjacency, and equality to the root index 0.

For $H_{\ell,a}$, put $k = f_\ell(i)$ and expand (3.8). Use the 4 orders (a_i, b_k, b_j) with $j \neq k$, (a_i, b_k, c_k), (a_i, b_j, c_h) with $j, h \neq k$, and (a_i, c_k, c_h) with $h \neq k$. In each repeated-type case, exactly one of the 2 colors of the repeated type is indexed by $f_\ell(i)$. The mixed-type cases are distinguished by whether both target indices equal $f_\ell(i)$ or neither does. Cyclically permute the color types to obtain the orderings for $H_{\ell,b}, H_{\ell,c}$.

For the 3-partite components $ABC - M_T, T_3, R_0$, order the colors by type a, b, c. In T_3 the repeated index and singleton index determine the tree edge, and in R_0 exactly 2 colors have index 0. Auxiliary blocks distinguish different list entries, and the t-colors distinguish base from additional edges.

These orders count each component edge once and use no numerical comparison of distinct index names. Each f_ℓ relates an index only to itself or to a tree neighbor. Their labeled relations are therefore supported on the diagonal and a graph of maximum degree 50. The condition on finite comparison rules of Theorem 11.5 is satisfied.

These choices give exactly one ordered tuple per edge of $H_{n,\boldsymbol{\omega}}$. Its edge polynomial is therefore the palette polynomial, so Proposition 3.10(i) gives $\lambda(L_{n,\boldsymbol{\omega}}) = 10^{-10}\Phi_n(\boldsymbol{\omega})$. □

We now place the 3 indexed coordinates on host faces to prove Theorem 11.2.

Proof of the initial-interval assertion in Theorem 11.2. Put $q = j + 1$ and $N = \binom{r}{q}$. Fix $n \geq 2Q$ and $\boldsymbol{\omega} \in \mathcal{A}^{\mathbb{N}}$. Throughout the stated range, $N \geq 10$; the smallest instance with $q \geq 3$ is $r = 5, q = 3$, and for $q = 2$ we have $r \geq 6$. Append the same $N - 10$ new distinct finite colors, in a fixed order, to every tuple of $L_{n,\boldsymbol{\omega}}$. After assigning its indexed coordinates $8, 9, 10$ to the faces below, assign all other coordinates bijectively to the remaining q-faces. Denote the resulting palette by P. The resulting palette still has finite comparison rules. Its index graph consists of the caterpillar together with isolated vertices and has maximum degree at most 50, and the finitely many involution relations are supported on its diagonal and edges.

If $3 \leq q \leq r - 2$, assign coordinates $8, 9, 10$ to the distinct faces $[q-1] \cup \{q\}$, $[q-1] \cup \{q+1\}$, and $[q-1] \cup \{q+2\}$. Only these 3 petals can carry indexed colors, so Theorem 11.5 applies. If $q = 2$ and $r \geq 6$, assign these coordinates to the disjoint pairs $12, 34, 56$ and apply Theorem 11.4.

In both cases, Fact 8.3, Theorem 11.6, and the appropriate realization theorem give

$$\pi^{(r)}_{q-1}(\mathcal{F}_{r,q}(P)) = \lambda(P) = N^{-N}\Phi_n(\boldsymbol{\omega}).$$

Together with Proposition 3.10(ii), this realizes $(N^{-N}, (1+\varepsilon)N^{-N}]$. A finite palette with one rainbow tuple has Lagrangian N^{-N}, and Theorem 11.3 realizes the endpoint. The same constant ε works in both cases. □

In particular, taking $q = r-2$ in Theorem 11.2 gives an interval at the last intermediate level for every $r \geq 5$. Here $N = \binom{r}{2}$, the common core has $r-3$ vertices, and exactly 3 private vertices remain. At $r = 5$ the interval is $[10^{-10}, (1+\varepsilon)10^{-10}] \subseteq \Pi^{(5)}_{2,\infty}$.

Proposition 11.7. *Let $r \geq 3$, $1 \leq j \leq r-2$, and $N := \binom{r}{j+1}$. Then*

$$[4N^{-N}, 1] \subseteq \Pi^{(r)}_{j,\infty}.$$

Proof. By the upper-interval construction in Section 10.5, every $x \in [4N^{-N}, 1]$ is a uniform Turán density at the highest nontrivial level for N-graphs. The representation in Proposition 8.4 gives a projective sequence of finite N-coordinate palettes with Lagrangians tending to x. Assign their coordinates bijectively to the $(j+1)$-faces of an ordered r-edge. Apply Theorem 11.3(ii) with $\mathcal{S} = \binom{[r]}{j}$ and $B = \varnothing$ to obtain density x for the union of their canonical forbidden families. □

Proof of the large-uniformity assertion in Theorem 1.1(iv). Let $\varepsilon > 0$ be as in Proposition 3.10, and choose an integer K with $(1+\varepsilon)^K > 4$. Fix $r \geq 6$ and $1 \leq j \leq r-3$, and put $q = j+1$, $N = \binom{r}{q}$, and $s = N - 10$. Assume $s \geq 10K$.

Finite restrictions of the palettes in Theorem 11.6 have relative values $10^{10}\lambda(P)$ dense in $[1, 1+\varepsilon]$, by Proposition 3.10; a single rainbow tuple gives the value one. Concatenate K such finite palettes on disjoint color sets. By Fact 8.3, their relative values multiply, so the resulting finite $10K$-coordinate palettes have relative values dense in $[1, (1+\varepsilon)^K]$. Appending $s - 10K$ fixed distinct colors preserves these values. We therefore obtain finite s-coordinate palettes Q for which $d := s^s\lambda(Q)$ is dense in an interval containing $[1, 4]$.

Fix such a finite Q, with colors disjoint from $L_{n,\boldsymbol{\omega}}$. The same product identity gives

$$N^N\lambda(L_{n,\boldsymbol{\omega}} \star Q) = d\Phi_n(\boldsymbol{\omega}).$$

Only the original 3 positions carry indexed colors. Assign them to the 3 one-sided petals when $q \geq 3$, or to $12, 34, 56$ when $q = 2$. All additional colors lie in a finite auxiliary set. Thus the same realization theorem used for the initial interval in Theorem 11.2 applies and gives every density in $(d, (1+\varepsilon)d]N^{-N}$. Replacing the critical palette by one rainbow 10-tuple gives the left endpoint by Theorem 11.3.

Consequently every available multiplier d gives $[d, (1+\varepsilon)d]N^{-N} \subseteq \Pi^{(r)}_{j,\infty}$. For $t \in (1+\varepsilon, 4]$, fix an available $d \in (t/(1+\varepsilon), t)$ and use the seed value $t/d \in (1, 1+\varepsilon)$. For $t \in [1, 1+\varepsilon]$, use $d = 1$. This realizes $[N^{-N}, 4N^{-N}]$. The upper interval is Proposition 11.7, and Proposition 11.1 excludes $(0, N^{-N})$. Forbidding one edge gives zero.

Finally, $\binom{r}{j+1} \geq \binom{r}{2}$ throughout $1 \leq j \leq r-3$. A single sufficiently large r_1 therefore makes $r \geq 6$ and $N \geq 10(K+1)$ simultaneously at every intermediate level. This proves the stated uniform choice of r_1.

For the quantitative bound $r_1 = 8$, one can instead obtain the finite palettes Q from ordinary densities. When $s \geq 18$, the bound $r_0 = 18$ recalled after Theorem 1.1 and the density of finite normalized Lagrangians give finite s-graphs H whose values $\lambda(H)/(s!/s^s)$ are dense in $[1, s^s/s!]$. Choosing one ordering of each edge gives a palette Q with $s^s\lambda(Q) = \lambda(H)/(s!/s^s)$ by Fact 8.3(i). Since $s^s/s! > 4$, these palettes supply all the multipliers needed above. For $r \geq 8$ and $1 \leq j \leq r-3$, we have $s = \binom{r}{j+1} - 10 \geq 18$, so the rest of the argument applies. □

12 Proofs of the palette realization theorems

We prove Theorems 11.3 to 11.5. The finite theorem uses conditional completion probabilities and a regularity–Ramsey extraction. For the countable palettes, the lower bound follows from random coloring, while the upper bounds retain a bounded set of indices and use disjointness or conditioning on the common core to control the discarded comparisons.

Lemma 12.1. *For finite or countable P,*

$$\pi^{(r)}_{q-1}(\mathcal{F}_{r,q}(P)) \geq \lambda(P).$$

For a fixed finitely supported probability weighting, the required lower bound on the number of vertices depends only on r, q and the density error, not on the palette or the number of its colors.

Proof. Fix a finitely supported probability weighting $\boldsymbol{x}$. Color every q-subset of $[n]$ independently according to $\boldsymbol{x}$, and keep exactly the admissible r-sets. For a fixed $(q-1)$-graph G, let Z_G be the number of kept edges in $\mathcal{K}_r(G)$. Then Z_G has mean $\Lambda_P(\boldsymbol{x})|\mathcal{K}_r(G)|$. Recoloring a single q-set can change the status of at most $\binom{n-q}{r-q} \leq n^{r-q}$ r-sets. The bounded-differences inequality [29] gives

$$\mathbb{P}\big[|Z_G - \mathbb{E}[Z_G]| > \eta n^r\big] \leq 2\exp\left(-\frac{2\eta^2 n^{2r}}{\binom{n}{q} n^{2r-2q}}\right) \leq 2\exp(-2\eta^2 n^q).$$

There are $2^{\binom{n}{q-1}}$ tests G, so their total failure probability is at most $2^{1+\binom{n}{q-1}}\exp(-2\eta^2 n^q) = o(1)$. Take the supremum over finitely supported probability weightings. □

12.1 Regularity and finite-palette realization

We use the layered partitions of Rödl and Schacht [45, Section 2] and specify the regularity and counting statements with their parameter dependencies.

For disjoint equal-sized clusters $V_1, \dots, V_t$, a layered partition through level $k-1$ recursively partitions the crossing ℓ-sets for $2 \leq \ell \leq k-1$. Once all proper subface cells have been fixed, exactly a_ℓ cells partition the supported ℓ-sets. Give them labels $1, \dots, a_\ell$. An *address* is a compatible assignment of labels to the faces of a fixed set of clusters. The actual cell at a face is determined by its label and its boundary address, not by its label alone. On a fixed set of ℓ clusters, the selected $(\ell-1)$-cells form a *polyad*; its ℓ-cliques are the ℓ-sets supported by that boundary address.

For an ℓ-partite $(\ell-1)$-uniform boundary B with $\mathcal{K}_\ell(B) \neq \varnothing$, a top relation $A \subseteq \mathcal{K}_\ell(B)$ is (δ, d, R)-*regular* if its relative density on every union of at most R clique sets of subboundaries differs from d by at most δ, whenever the tested union has at least a δ fraction of $\mathcal{K}_\ell(B)$. Here $0 < \delta \leq 1$, $d \in [0,1]$, and $R \geq 1$ is an integer. For $R = 1$ we omit R. Taking the whole boundary as a test gives $|A|/|\mathcal{K}_\ell(B)| = d \pm \delta$, so d may be a target density rather than the exact density.

For a (δ, d, R)-regular relation on B, the discrepancy against a product of $[-1,1]$-valued functions on proper subfaces is $O_\ell(\delta)$ under the uniform measure on $\mathcal{K}_\ell(B)$. To see this, write each $[0,1]$-valued factor as the integral of its level-set indicators, and regard each factor on a proper subface as a function on a containing $(\ell-1)$-face. The absolute discrepancy contributed by a test below the regularity threshold is bounded by its measure. Writing signed factors as differences costs only a factor depending on ℓ.

For a partition through level $k-1$, write $\boldsymbol{a} \coloneqq (t, a_2, \dots, a_{k-1})$ for its complexity vector, where t is the number of clusters. We use the strong regularity and counting theorems in the following form.

Theorem 12.2. *Fix integers $k \geq 2$ and $M \geq k$, an initial ordered vertex equipartition, a top error $0 < \sigma \leq 1$, and a volume error $0 < \mu < 1$. The lower-level regularity tolerances may be prescribed by an arbitrarily small positive function of the final complexity vector $\boldsymbol{a}$, and the number R of allowed top tests by an arbitrarily large integer-valued function of $\boldsymbol{a}$. There are bounds T, n_0 such that every k-graph of order at least n_0 admits an equitable layered partition through $k-1$, refining the initial vertex partition, with the following properties.*

(i) *The number of clusters and every a_ℓ are at most T. Cluster sizes differ by at most one. A bounded deletion or a negligible exceptional set makes them equal when required.*

(ii) *Every compatible address on at most M clusters determines a lower complex with the prescribed regularity. On $u \le M$ clusters of common size m, an address through level a supports*

$$(1 \pm \mu) m^u \prod_{\ell=2}^{a} a_\ell^{-\binom{u}{\ell}} \tag{12.1}$$

transversal u-sets. Empty products are 1.

(iii) *All but a σ fraction of crossing k-sets belong to a top polyad on which the host k-graph is regular with top error σ and $R(\boldsymbol{a})$ tests.*

(iv) *If selected top polyads on the edges of a fixed k-graph on at most M vertices have densities at least $d_* > 0$, and all are regular with sufficiently small top error, then the selected top relations contain a copy of that k-graph, provided their lower addresses are compatible, the lower-level regularity tolerance is small enough as a function of $\boldsymbol{a}$ and d_*, and the clusters are sufficiently large.*

Proof. The partition comes from the strong regularity theorem of Rödl and Schacht [45, Theorem 2.3 and Lemma 4.1]. Reiher, Rödl, and Schacht [43, Theorem 3.5] and Lin, Sun, Wang, and Zhou [24, Theorem 3.7] also record the family formulation and the parameter order. Applying the dense counting lemma in Rödl and Schacht [44, Theorem 2.1] to the finitely many regular lower complexes determined by compatible addresses gives (12.1), with the lower-level regularity tolerance chosen sufficiently small for all $u \le M$. Positivity in that formula justifies recursive cell sampling for every formal compatible address. The embedding assertion follows from the counting lemma of Nagle, Rödl, and Schacht [33, Corollary 15], in the form stated by Lin, Sun, Wang, and Zhou [24, Theorem 3.8], which allows varying top densities above d_*. Divisibility conditions are met by deleting a bounded number of vertices after T is known; this changes all densities by $o(1)$. □

For later counting, we also use the relative-error counting lemma of Nagle, Rödl, and Schacht [33] for a fixed finite collection of top color classes. The lower-level regularity tolerance is chosen as a function of $\boldsymbol{a}$; the top error is a fixed constant chosen before the partition. With the top error fixed, relative-error counting cannot be applied uniformly to colors of arbitrarily small density.

The finite-palette upper bound uses a fixed number of colors. For countable palettes we need a complexity bound independent of that number, with simultaneous regularity of each fine color and each class of a prescribed coarse partition. Regularity of arbitrary unions of fine colors is not required.

12.1.1 A regular partition independent of the color count

Lemma 12.3. *Fix integers $k \ge 2$ and $u \ge k$, an initial ordered vertex equipartition, constants $0 < \sigma, \delta_k \le 1$ and $0 < \mu < 1$, and functions of the final complexity prescribing arbitrarily small lower-level regularity tolerances and arbitrarily large numbers of top tests. There are bounds T, n_0, independent of the number of fine colors, with the following property.*

For any coloring of all k-sets of a vertex set of size $n \ge n_0$ and any coarse labeling of these colors, there is an equitable layered partition through $k-1$, refining the initial vertex partition, with the following properties.

(i) *The number of clusters and every a_ℓ are at most T.*

(ii) *The volume conclusion (12.1) holds for at most u clusters.*

(iii) *Outside polyads supporting a total of at most σn^k k-sets, every fine-color indicator and every coarse-label indicator is simultaneously $(\delta_k, d, R(\boldsymbol{a}))$-regular, with its own actual density d.*

Proof. Let (Ω, μ) be a probability space, and let $(f_c)_{c \in C}$ be a finite or countable family of nonnegative functions with $\sum_c f_c = 1$ almost everywhere. For a finite measurable partition $\mathcal{A} = (A_i)$ put

$$I_f(\mathcal{A}) := \sum_{i:\ \mu(A_i) > 0} \frac{\sum_c (\int_{A_i} f_c \, \mathrm{d}\mu)^2}{\mu(A_i)}.$$

Since $\sum_c f_c = 1$, we have $0 \le I_f(\mathcal{A}) \le 1$. Applied to the fine-color and coarse-color indicators, the sum of the two indices is at most 2, independently of the number of colors. Regard $f = (f_c)_{c \in C}$ as an $\ell^2(C)$-valued function. Its norm is at most one. Completing the square on each atom gives

$$I_f(\mathcal{A}) = \sup_{(a_i)} \sum_i \int_{A_i} \left(2\langle f, a_i \rangle - \|a_i\|_2^2\right) \mathrm{d}\mu, \tag{12.2}$$

where each a_i is a probability vector; the maximum uses $a_i = \mathbb{E}[f \mid A_i]$ on nonnull atoms. If $\mathcal{C}$ refines $\mathcal{A}$, conditional orthogonality also gives

$$I_f(\mathcal{C}) - I_f(\mathcal{A}) = \int \|\mathbb{E}[f \mid \mathcal{C}] - \mathbb{E}[f \mid \mathcal{A}]\|_2^2 \, \mathrm{d}\mu. \tag{12.3}$$

Both identities hold for countably many colors because $\|f\|_2 \le 1$.

Claim 12.4. *Let $\mathcal{A}, \mathcal{B}$ be finite measurable partitions of Ω. The following 2 estimates hold.*

(i) *Align the labels of 2 partitions $\mathcal{A} = (A_i)$ and $\mathcal{B} = (B_i)$ by adding empty atoms if necessary. If their label functions disagree on a set of measure τ, then $|I_f(\mathcal{A}) - I_f(\mathcal{B})| \le 2\tau$. For the sum of 2 indices of this form the bound is 4τ, independently of both numbers of colors.*

(ii) *Suppose the union of atoms of $\mathcal{B}$ that are not contained in any atom of $\mathcal{A}$ has measure at most β. Then $I_f(\mathcal{B}) \ge I_f(\mathcal{A}) - \beta$. For the sum of 2 indices of this form the loss is at most 2β.*

Proof. For (i), the integrand in Equation (12.2) lies in $[-1, 1]$ for probability vectors f, a_i. Use the optimizing vector for each label of $\mathcal{A}$ also at that label of $\mathcal{B}$, choosing an arbitrary vector for an empty atom. The integrals differ only on the set of relabeled points, so $I_f(\mathcal{A}) \le I_f(\mathcal{B}) + 2\tau$. Interchanging the partitions proves the absolute bound. Adding the estimates for the fine and coarse indices gives 4τ.

For (ii), let $\mathcal{C} \coloneqq \mathcal{A} \vee \mathcal{B}$. By Equation (12.3), $I_f(\mathcal{C}) \ge I_f(\mathcal{A})$. On each atom of $\mathcal{B}$ contained in an atom of $\mathcal{A}$, the partitions $\mathcal{B}$ and $\mathcal{C}$ coincide. On the union of the atoms of $\mathcal{B}$ not contained in any atom of $\mathcal{A}$, the contribution to $I_f(\mathcal{C})$ is at most its measure, since $\sum_c f_c = 1$. Consequently $I_f(\mathcal{C}) - I_f(\mathcal{B}) \le \beta$. □

The hypotheses differ because a small relabeling error need not give an almost-refinement. To obtain the almost-refinement, we use the following slicing consequence for a partition with a bounded number of atoms.

Claim 12.5. *Fix an integer $\ell \ge 2$, a bound s on the number of labels, a positive number β, and a prescribed function for the lower-level regularity tolerances. Let $\mathcal{S}$ be a sufficiently regular family through level $\ell - 1$, and let $Y_1, \dots, Y_s$ partition the ℓ-sets and be sufficiently regular on every $(\ell-1)$-polyad of $\mathcal{S}$. Suppose the hierarchy formed by $\mathcal{S}$ and the partition (Y_i) refines the complete hierarchy of a fixed initial vertex partition. There is an integer b, bounded in terms of s, β and that initial hierarchy, and an equitable family extending $\mathcal{S}$ at level ℓ. On each polyad B of $\mathcal{S}$, this extension has b cells, each $(\varepsilon_{\mathrm{sl}}, 1/b)$-regular, where $\varepsilon_{\mathrm{sl}}$ is the prescribed tolerance at the extended complexity. Except for a union of new ℓ-cells of measure at most β in the normalized initial crossing space, every new ℓ-cell with supporting polyad B is contained in $Y_i \cap \mathcal{K}_\ell(B)$ for some i. The required regularity tolerances for $\mathcal{S}$ and the Y_i are given by an allowed function of the lower complexity. The clusters are assumed to be sufficiently large and to satisfy the divisibility conditions required for the equitable subdivision.*

Proof. This is the slicing step of Rödl and Schacht [45, Claim 6.10, properties (T1)–(T5), and its proof in Section 6.3], with $k = \ell$ and a complete initial hierarchy. Intersect each Y_i with the set of ℓ-sets supported by each polyad. There are at most s intersections in any polyad, and each is regular by assumption. Choose b sufficiently large compared with s/β, with the required divisibility. Their slicing lemma [45, Proposition 3.5] first cuts regular pieces with target density $1/b$ from each atom, leaving at most 1 remainder per intersection. These remainders have total relative measure $O(s/b)$, with rounding absorbed by the chosen input tolerance.

The union of these remainders consists of the supported ℓ-sets of the polyad that remain after removing the pieces already cut. Its indicator is the support indicator minus the sum of the indicators of those pieces. Choose the input tolerances so that the sum of the regularity errors of these pieces meets the requirements of the slicing lemma. A second application of the slicing lemma divides this union into the

remaining regular pieces with the same target density. Every cell of the resulting partition has actual relative density $1/b \pm \varepsilon_{\mathrm{sl}}$. The cells obtained by subdividing the union of the remainders are the only cells that may cross old atom boundaries, which gives the stated almost-refinement.

The fixed initial hierarchy is preserved by property (T4) in the cited claim. Its top density is 1 on each initial crossing tuple, so it imposes no additional regularity assumption. Choosing the tolerance function $\varepsilon_i(\boldsymbol{a})$ in the cited argument small enough gives the prescribed final tolerance at $(\boldsymbol{a}, b)$. The bound on the number of cells of the resulting partition depends on the atom partition only through its number s of labels. For any prescribed maximum pattern size u, the density errors change the product of target densities by a relative $O_u(b\varepsilon_{\mathrm{sl}})$. We choose this error and the counting error for the lower complex small enough to obtain (12.1) with its stated volume tolerance. □

We now return to the host coloring and construct the successive regular partitions.

Refine the prescribed initial vertex partition into a sufficiently fine ordered equipartition V^0. Its classes may be taken equal by a bounded initial deletion. Choose it so that the proportion of noncrossing k-sets is less than $\sigma/4$. Let Ω_k be the set of k-sets crossing V^0, with uniform measure, and compute all indices on this fixed probability space. Every subsequent vertex partition refines V^0, so a point of Ω_k remains crossing. The excluded noncrossing k-sets contribute less than $\sigma/4$ to the final exceptional proportion. We may enlarge the fixed initial partition further to meet any prescribed lower bound on the number of clusters.

At stage i, let $\boldsymbol{a}_i$ be the complexity vector of the current hierarchy and $\mathcal{A}_i$ its boundary-address partition on Ω_k. For the fine indicators f_c and coarse indicators g_b, define $\mathcal{E}(\mathcal{A}) \coloneqq I_f(\mathcal{A}) + I_g(\mathcal{A})$ on partitions of Ω_k. Each family of indicators sums to one, so $0 \le \mathcal{E}(\mathcal{A}) \le 2$, independently of their numbers of colors.

If the bad polyads in Ω_k have total measure more than $\sigma/2$, choose in each such polyad 1 fine or coarse indicator that fails the regularity condition and 1 witness. For a selected fine indicator f_c, let $B \subseteq \Omega_k$ be the k-set class supported by its bad polyad, and let $W \subseteq B$ be a witness. The set W is a union of at most $R(\boldsymbol{a}_i)$ subboundary clique sets, with $\mathbb{P}[W] \ge \delta_k \mathbb{P}[B]$ and $|\mathbb{E}[f_c \mid W] - \mathbb{E}[f_c \mid B]| > \delta_k$, where probabilities refer to the uniform distribution on Ω_k. By Equation (12.3) and Cauchy–Schwarz, refining the address partition so that W is a union of atoms increases the contribution of B to this indicator's index by at least

$$\mathbb{P}[W]\big(\mathbb{E}[f_c \mid W] - \mathbb{E}[f_c \mid B]\big)^2 \ge \mathbb{P}[B]\delta_k^3.$$

The same estimate holds for a selected coarse indicator. Refine the current partition of $(k-1)$-sets by membership in all facet relations defining the selected witness clique sets, retaining the old cell and initial vertex-class labels. Denote its classes by $X_1, \dots, X_s$. The induced address partition $\mathcal{A}_X$ makes every witness a union of atoms. Since the witnesses have disjoint ambient polyads, it satisfies

$$\mathcal{E}(\mathcal{A}_X) \ge \mathcal{E}(\mathcal{A}_i) + \gamma, \qquad \gamma \coloneqq \frac{\sigma}{2}\delta_k^3. \tag{12.4}$$

Complete this labeled partition outside the initial crossing sets by fixed labels. If the current number of cells and polyads is at most B_i, then $s \le B_i 2^{kR(\boldsymbol{a}_i)B_i} + B_i$. Only 1 witness per bad polyad enters this bound.

First suppose $k \ge 3$. Let Ω_{k-1} be the set of $(k-1)$-sets crossing V^0, with uniform measure, and put $\tau \coloneqq \gamma/(32k)$. The fixed initial partition gives a positive lower bound for $|\Omega_{k-1}|/n^{k-1}$ for all sufficiently large n. Since s is bounded, choose $\nu > 0$ so that $s\nu n^{k-1}/|\Omega_{k-1}| \le \tau$ uniformly for such n.

Apply the regular-approximation lemma of Rödl and Schacht [45, Lemma 4.2] at level $k-1$ over the fixed complete initial hierarchy $\mathcal{O}$ above V^0, with the labeled partition X as input. The partition X refines $\mathcal{O}^{(k-1)}$, as required by condition (RAL.c) in the cited proof. The hierarchy $\mathcal{O}$ has density 1 in its higher cells and satisfies every initial regularity requirement in (RAL.b).

The output consists of a family $\mathcal{S}$ through level $k-2$ and an aligned partition $Y_1, \dots, Y_s$ of $(k-1)$-sets. Each Y_a is regular on every polyad of $\mathcal{S}$, the initial hierarchy is preserved, and $|X_a \triangle Y_a| \le \nu n^{k-1}$ for each label a, by properties (RAL.Y1)–(RAL.Y3) in the cited proof. Summing these editing bounds over the s labels and dividing by $|\Omega_{k-1}|$ bounds the total label-disagreement probability on Ω_{k-1} by τ.

Let $\mathcal{A}_Y$ be the address partition on Ω_k induced by the edited labels. A top address changes only if at least 1 of its k boundary facets changes label. For equal initial classes, a uniformly chosen facet of a

uniform member of Ω_k is uniform on Ω_{k-1}. The sum of the k facet error probabilities is therefore at most $k\tau$. The union bound and Claim 12.4(i) give

$$\mathcal{E}(\mathcal{A}_Y) \geq \mathcal{E}(\mathcal{A}_X) - 4k\tau.$$

To obtain an equitable almost-refinement, intersect each Y_a with the supported $(k-1)$-sets of each polyad of $\mathcal{S}$. The resulting address partition refines $\mathcal{A}_Y$, so its index is no smaller. On each polyad there are at most s regular atoms. Apply Claim 12.5 with almost-refinement error allowance β/k, where $\beta := \gamma/16$. The resulting $(k-1)$-cells, together with $\mathcal{S}$, form the next equitable hierarchy. The bad lower cells are those $(k-1)$-cells not contained in any of these intersections. For a top address using no bad lower cell, all supported k-sets with that address lie in one atom of the address partition induced by the intersections. Applying the same union bound to the k boundary facets shows that the exceptional top atoms have total measure at most β on Ω_k. Thus Claim 12.4(ii) gives

$$\mathcal{E}(\mathcal{A}_{i+1}) \geq \mathcal{E}(\mathcal{A}_i) + \gamma - 4k\tau - 2\beta \geq \mathcal{E}(\mathcal{A}_i) + \frac{\gamma}{2}. \tag{12.5}$$

It remains to specify the parameter choices and verify termination. At a stage with bounded complexity B_i, the number s of labels is bounded as above. Prescribe the regular-approximation tolerance as a function of the lower complexity returned by regular approximation, sufficiently small for slicing, the final lower regularity, and all pattern counts of size at most u. Fix the number of slices in terms of s and the slicing error allowance β/k. Regular approximation then returns a bound B_{i+1} depending only on these quantities and the prescribed functions. Choose the bounds recursively for at most $\lceil 4/\gamma \rceil + 1$ stages, and take the original host large enough for all resulting size and divisibility requirements. This is the finite recursion of Rödl and Schacht [45, Sections 6.2–6.3], with the bounded number s of witness atoms as the number of colors in each lower-level application. The complete initial hierarchy remains fixed throughout.

For $k = 2$ the witness subsets refine the vertex classes directly, so no level-1 regular-approximation theorem is needed. Equitably divide the refined vertex classes, setting aside an arbitrarily small proportion of the vertices as exceptional. Choose this proportion so that the pairs meeting the exceptional vertices have measure less than $\gamma/4$ in Ω_2. Refinement does not decrease either index, and the contribution of these exceptional pairs to the sum of the 2 indices is at most twice their measure. Thus discarding them loses less than $\gamma/2$, while the witness refinements give the gain γ in (12.4). Each bad stage therefore increases the index by at least $\gamma/2$, and the same finite recursion applies. The resulting partition is regular for unions of at most $R(\boldsymbol{a})$ rectangle tests, as required.

There cannot be $\lceil 4/\gamma \rceil + 1$ bad stages, since each gains at least $\gamma/2$ and the index is at most 2. At termination the total measure of the bad polyads inside Ω_k is at most $\sigma/2$. The original noncrossing sets and the bounded deletions together cost less than the remaining $\sigma/2$. Every other polyad is regular for every fine color and every coarse label. The functions specifying the lower-level regularity tolerances ensure (12.1) with the stated volume tolerance. □

The regular partition allows us to select an ordered complex with all face polyads good simultaneously.

Lemma 12.6. *Given a positive integer h, a prescribed top regularity error, and a prescribed function for the number of top tests, the initial partition and exceptional allowance in Lemma 12.3 can be chosen so that one can select h equal-sized ordered clusters $U_1 < \cdots < U_h$ and a compatible lower complex through level $k-1$ with every k-face polyad good for every fine and coarse color. The cluster size is at least a positive constant times the original order. All constants are independent of the number of fine colors.*

Proof. It suffices to consider $h \geq k$, since a smaller selection follows by deleting clusters. Choose the initial number of consecutive vertex blocks sufficiently large compared with h^2, and the exceptional proportion sufficiently small compared with $\binom{h}{k}^{-1}$. Select h distinct clusters uniformly. Then, in increasing order of face size, choose a cell label uniformly for every face through level $k-1$. Summing over all selected k-subsets and using the equitability counts bounds the expected number of bad k-faces by $\binom{h}{k}$ times a fixed constant times the global exceptional proportion. The probability that 2 selected clusters have the same initial block tends to zero as the initial block count increases. A union bound therefore yields a selection with no bad face and no repeated initial block. Its clusters are ordered by those blocks. Trimming a size difference of at most 1 has negligible effect; alternatively use the divisible version of the partition. Bounded partition complexity gives the linear lower bound on the cluster size. □

12.1.2 Counting on regular complexes

We first compare arbitrary bounded functions on 2 disjoint faces, using only regularity below their dimension.

Lemma 12.7. *Fix $r \geq q \geq 2$ and $\zeta > 0$. Let $\mathcal{P}$ be a sufficiently regular $(r, q-1)$-complex on equal-sized ordered clusters, with positive lower densities $d_2, \dots, d_{q-1}$. Let $A, B \in \binom{[r]}{q}$ be disjoint. Let μ be uniform on the transversal r-sets supported by $\mathcal{P}$, and let μ_A, μ_B be uniform on the internally supported q-sets on A, B, respectively. Let X have law μ, and let X_A, X_B denote its restrictions to the clusters indexed by A, B, respectively. Then, for arbitrary functions f, g with values in $[-1, 1]$,*

$$|\mathbb{E}_\mu[f(X_A)g(X_B)] - \mathbb{E}_{\mu_A}[f]\,\mathbb{E}_{\mu_B}[g]| \leq \zeta. \tag{12.6}$$

The required lower-level regularity tolerances may depend on the density vector. At $q = 2$ the discrepancy is zero.

Proof. Work first with unnormalized expectations over independent uniform vertices of the r clusters. We prove the replacement step for arbitrary retained position sets $B_1, \dots, B_t$. Keep every lower-cell indicator whose position set is contained in some B_i, together with a $[-1, 1]$-valued function on each B_i. Replace the other lower-cell indicators by their densities times their boundary-support indicators, in decreasing dimension.

For a replaced cell on S, we have $S \not\subseteq B_i$ for every i. After fixing vertices outside S, each retained function therefore depends only on a proper subset of S. Every remaining cell indicator also has this property: larger cells not retained have already been replaced, and no retained cell contains S. Assign factors on smaller subsets to containing $(|S| - 1)$-faces. Lower regularity then bounds the replacement error; integration of level sets handles bounded weights, and expansion into positive and negative parts handles signs.

With lower errors at most α, the total absolute error is $O_{r,t}(\alpha)$. Boundary indicators introduced repeatedly equal their own powers. The resulting count is the weighted count on the retained support multiplied by the product of the replaced cell densities. Applying the same argument with all functions equal to one gives the corresponding support count. With no retained sets, it gives the full support probability as $\prod_{j=2}^{q-1} d_j^{\binom{r}{j}} + O_r(\alpha)$, which is at least half of this positive product for sufficiently small α.

For the lemma, take the retained sets to be A and B. They are disjoint, so their retained supports depend on disjoint groups of independent vertices. The weighted count factors into the two internal weighted counts. The absolute weighted count is at most the support count, so dividing the two replacement estimates gives an error at most $O_r(\alpha)/\prod_{j=2}^{q-1} d_j^{\binom{r}{j}}$. Choosing α so that this bound is at most ζ proves Equation (12.6). At $q = 2$, there are no lower cells and the discrepancy is zero. □

The product-test estimate also controls a sparse relation on countably many labels.

Lemma 12.8. *Let μ be a joint distribution of (X, Y), and let ν_X, ν_Y be reference distributions for X, Y, respectively. Assume the product-test estimate (12.6) in the form $|\mathbb{E}_\mu[f(X)g(Y)] - \mathbb{E}_{\nu_X}[f]\mathbb{E}_{\nu_Y}[g]| \leq \zeta$ for every pair of $[-1, 1]$-valued functions f, g. Let $\alpha(X), \beta(Y)$ take values in a countable set I. If $D \geq 1$ and a relation $R \subseteq I \times I$ has at most D entries in each row, then*

$$\left| \mathbb{P}_\mu[(\alpha(X), \beta(Y)) \in R] - \sum_{(i,k)\in R} \nu_X[\alpha = i]\nu_Y[\beta = k] \right| \leq D\zeta.$$

Proof. Assign independent uniform signs $\epsilon_k \in \{-1, 1\}$ to the column labels and put

$$f_\epsilon(i) = \frac{1}{D} \sum_{k:\ (i,k)\in R} \epsilon_k, \qquad g_\epsilon(k) = \epsilon_k.$$

Both functions lie in $[-1, 1]$, and independence gives

$$D\,\mathbb{E}_\epsilon[f_\epsilon(i)g_\epsilon(k)] = \mathbf{1}_{(i,k)\in R}.$$

Apply the product-test hypothesis to $f_\epsilon \circ \alpha$ and $g_\epsilon \circ \beta$, then average over the signs. Multiplication by D gives the stated error. Each row sum is finite and all functions are bounded, so Fubini's theorem also justifies the argument for countable labels. □

The next lemma shows that almost all supported tuples on a fixed set of positions have approximately the same number of extensions to the full complex.

Lemma 12.9. *Fix $u \ge q \ge 3$. Let $\mathcal{P}$ be a sufficiently regular layered $(u, q-1)$-complex on equal clusters of size m, with positive prescribed densities d_B. Cell constraints may be omitted provided the subfaces of every retained cell are retained. Set $d_B = 1$ for an omitted constraint. Fix a nonempty set $A \subseteq [u]$, and put*

$$D_A \coloneqq \prod_{B \subseteq A,\ 2 \le |B| < q} d_B, \qquad D \coloneqq \prod_{B \subseteq [u],\ 2 \le |B| < q} d_B.$$

For an internally supported tuple x on A, let $L(x)$ be the probability that independent uniform vertices in the other clusters extend x to the full complex. For every $0 < \kappa \le 1/10$, the lower-level regularity tolerances can be prescribed as a function of the density vector and m can be taken sufficiently large so that

$$\mathbb{E}[L] = (1 \pm \kappa)\frac{D}{D_A}, \qquad \mathbb{E}[L^2] = (1 \pm \kappa)\left(\frac{D}{D_A}\right)^2.$$

The expectation uses uniform measure on the internally supported A-tuples. In particular, $\mathbb{E}[L^2] \le 2\mathbb{E}[L]^2$. The marginal on A of a uniformly sampled full supported tuple can be made arbitrarily close in total variation to uniform measure on the internally supported A-tuples.

Proof. Consider the internal A-complex, 1 full copy, and 2 full copies identified on A. In the last pattern, the 2 variables supplied by each outside cluster are sampled independently from that cluster; no distinctness or additional cell constraints are imposed between the copies. Keep a single indicator for each shared lower face. The main terms for these 3 patterns are D_A, D, and $D_A(D/D_A)^2$, respectively. To verify these main terms, replace the lower-cell indicators in decreasing dimension by their prescribed densities times their boundary supports. In each replaced face, the variables still come from distinct original clusters. After fixing all other abstract variables, every remaining cell indicator depends on a proper subset of that face, since higher-dimensional indicators have already been removed and shared faces were imposed only once. Lower regularity therefore bounds each replacement error by a constant times the lower error. The argument also applies when 2 independent variables from the same outside cluster take the same actual vertex, since their abstract variable sets and the proper-subface tests are unchanged. Boundary indicators introduced during replacement duplicate indicators already present, whose powers equal themselves.

There are a bounded number of replacements on at most $2u - |A|$ abstract positions. Choose the lower error sufficiently small relative to all 3 positive main terms, so that each support probability differs from its main term by at most a relative error of $\kappa/10$. Divide the latter 2 probabilities by the internal support probability, which is $(1 \pm \kappa/10)D_A$. The resulting relative errors are at most $2(\kappa/10)/(1 - \kappa/10) < \kappa$, giving the 2 asserted moment estimates. Finally, $(1+\kappa)/(1-\kappa)^2 < 2$ for $\kappa \le 1/10$.

Let μ be uniform on the supported A-tuples, and let ν be the marginal of uniform sampling from the full support. The number of extensions of a root is proportional to L, so $\mathrm{d}\nu/\mathrm{d}\mu = L/\mathbb{E}_\mu L$. Consequently,

$$\|\nu - \mu\|_{\mathrm{TV}} = \frac{1}{2}\mathbb{E}_\mu\left[\left|\frac{L}{\mathbb{E}_\mu[L]} - 1\right|\right] \le \frac{1}{2}\left(\frac{\mathbb{E}_\mu[L^2]}{\big(\mathbb{E}_\mu[L]\big)^2} - 1\right)^{1/2}. \tag{12.7}$$

The moment estimates make this bound arbitrarily small as $\kappa \to 0$. It controls every root event, including one selected after the complex. □

We use the following estimate to convert an absolute test error into an error relative to the full support. On a probability space (Ω, μ), let $|f| \le 1$ and suppose $|\int_U f\,\mathrm{d}\mu| \le \delta$ for every union U of at most R basic tests, where $R \ge 1$ is an integer. Choose a measurable random basic test T and put $F(x) = \mathbb{P}[x \in T]$. If $C, \rho > 0$ and $\int F^2\,\mathrm{d}\mu \le C\rho^2$, then, for each integer $1 \le s \le R$,

$$\left|\int fF\,\mathrm{d}\mu\right| \le \frac{\delta}{s} + \frac{C}{2}(s-1)\rho^2. \tag{12.8}$$

Indeed, the union of s independent copies of T has membership probability $1 - (1-F)^s$. Average the test bound over these copies and use $0 \le sF - 1 + (1-F)^s \le \binom{s}{2}F^2$.

Suppose $C \geq 1$, $0 < \eta \leq 1$, $R \geq 1 + \rho^{-1}$, and $\delta \leq \eta^2/(8C)$. If $\rho \geq \eta/(4C)$, take $s = 1$. Otherwise take $s = \lceil \eta/(2C\rho) \rceil \leq R$. In the latter case the two terms in Equation (12.8) are each at most $\eta\rho/4$; in the former case $\delta \leq \eta\rho/2$. Thus $|\int fF \, d\mu| \leq \eta\rho$ in both cases.

Lemma 12.10. *Fix $u \geq q \geq 3$ and $\eta > 0$. Let $\mathcal{P}$ be a sufficiently regular $(u, q-1)$-complex on ordered equal-sized clusters. Fix $A \in \binom{[u]}{q}$ and a color c on its internally supported q-sets. Let $p_A(c)$ be the actual density of this color under uniform sampling of these q-sets. Suppose its indicator is $(\delta, p_A(c), R)$-regular, with a fixed error $\delta \leq \eta^2/16$ and sufficiently many tests R, where R may be an allowed function of the lower densities. In a count on the full complex, $\mathbf{1}_{\chi=c}$ at A may be replaced by $p_A(c)$ with error at most η times the full supported count, even with an additional product of arbitrary $[0,1]$-valued functions on other q-faces and on proper subfaces.*

Proof. For $\eta \geq 1$, the assertion follows from $|\mathbf{1}_{\chi=c} - p_A(c)| \leq 1$. Assume $0 < \eta < 1$. Let μ be uniform measure on the internally supported A-tuples, and let L be the extension probability of Lemma 12.9. Put $\rho := \mathbb{E}[L]$. Prescribe lower-level regularity tolerances for the internal pattern and the patterns with 1 or 2 extensions in that lemma. It gives $\mathbb{E}[L^2] \leq 2\rho^2$, with ρ bounded below by a positive function of the density vector. These choices are made before applying the regular partition theorem.

After fixing the outside vertices, every other face function depends only on vertices in a proper subset of A. Their product, together with the extension constraints, is a boundary test; factors on smaller subsets may be assigned to containing $(q-1)$-faces. Independent uniform $[0,1]$ thresholds turn the weighted factors into a random boundary test. Averaging the boundary-test indicator over the outside vertices and these thresholds gives a function F on the internally supported A-tuples, with $0 \leq F \leq L$. Top regularity bounds the absolute value of the integral of $\mathbf{1}_{\chi=c} - p_A(c)$ over any union of at most R tests by δ, including unions smaller than the regularity threshold. The positive lower bound for ρ depends only on the density vector, so the test-count function can be chosen to ensure $R \geq 1 + \rho^{-1}$.

Put $f := \mathbf{1}_{\chi=c} - p_A(c)$. Since $0 \leq F \leq L$, we have $\int F^2 \, d\mu \leq 2\rho^2$. Apply Equation (12.8) and its consequence with $C = 2$ and $\delta \leq \eta^2/16$ to obtain $|\int fF \, d\mu| \leq \eta\rho$. Multiplying by the number of internally supported A-tuples and by the number of choices of the outside vertices gives the asserted relative count error. □

Repeated replacements show that the joint color distribution is close to the product of the local color distributions.

Lemma 12.11. *Fix $r \geq q \geq 2$, a positive integer m_0, and an error $\varepsilon > 0$. On a sufficiently regular $(r, q-1)$-complex, let μ be uniform on the supported transversal r-sets. Suppose a q-face coloring uses at most m_0 colors and, for every $A \in \binom{[r]}{q}$ and every color c, the color class on A is top-regular on its boundary with actual density $p_A(c)$. Let $\boldsymbol{c}$ denote the face-color tuple in the chosen coordinate order. The top error can be chosen independently of the lower density vector; the lower errors and the number of top tests may be prescribed as functions of that vector. These parameters can be chosen so that*

$$\sup_E \left| \mathbb{P}_\mu[\boldsymbol{c} \in E] - \mathbb{P}_{\bigotimes_A p_A}[\boldsymbol{c} \in E] \right| \leq \varepsilon.$$

Proof. Put $N = \binom{r}{q}$ and $\eta = \varepsilon/(N m_0^N)$. For $q \geq 3$, choose the top error, lower regularity tolerances, and number of tests as in Lemma 12.10. Fix a color pattern $(c_A)_{A \in \binom{[r]}{q}}$. Replace its color indicators one at a time by their face densities. All remaining factors are $[0,1]$-valued functions on other faces, so each replacement has error at most η under the normalized supported measure. Thus

$$\left| \mathbb{P}_\mu[\boldsymbol{c} = (c_A)_A] - \prod_A p_A(c_A) \right| \leq N\eta.$$

Summing over the at most m_0^N patterns in any event gives the conclusion. At $q = 2$, after fixing the other vertices each replacement is a product test on the corresponding pair of clusters. Taking its error at most η gives the same bound. The top error depends only on r, q, m_0, ε; the lower errors and test counts may depend on the lower densities, as required. □

Lemma 12.12. *Fix integers $h \geq r > q \geq 2$ and $m \geq 1$, and let $\delta > 0$. Every coloring of the q-subsets of a sufficiently large ordered n-vertex set, using a color set C of size at most m, admits h equal-sized*

ordered clusters and a compatible lower complex with the following property. For each r-subset of clusters, let K_0 be its supported transversal tuples, indexed in increasing cluster order. Then $|K_0| \ge cn^r$, where $c > 0$ depends only on the fixed parameters, and the joint distribution ν of its q-face colors satisfies

$$\sup_{E \subseteq C^{\binom{[r]}{q}}} |\nu(E) - \boldsymbol{x}^{\otimes \binom{[r]}{q}}(E)| \le \delta$$

for one probability weighting $\boldsymbol{x}$ common to all these cluster subsets. Each K_0 is specified by cluster restrictions and relations on $(q-1)$-faces. For $q = 2$, it is the product of its r vertex clusters.

Proof. Apply Lemmas 12.3 and 12.6 with an initial interval partition to obtain sufficiently many ordered clusters and a compatible lower complex regular for every color. Assign each q-set of clusters a nearest point in a finite ξ-net of the probability simplex. Ramsey's theorem gives h clusters with a common net point $\boldsymbol{x}$, so their actual face marginals $\boldsymbol{p}_A$ differ from $\boldsymbol{x}$ by at most 2ξ in ℓ^1.

For each r-subset of clusters, Lemma 12.11 puts its joint law within $\delta/2$ of $\bigotimes_A \boldsymbol{p}_A$. Replacing the $N = \binom{r}{q}$ marginals costs at most $2N\xi$; choose $\xi \le \delta/(4N)$. Lower-complex counting gives $|K_0| \ge cn^r$. Assign each lower-cell restriction to a containing $(q-1)$-face. All errors are chosen for the finitely many r-subsets simultaneously, with lower-level tolerances depending on the density vector as in the cited lemmas. □

Let χ color the q-subsets of $[n]$ independently with a fixed distribution $\boldsymbol{w}$ on a finite color set. For every $\zeta > 0$, with probability tending to one,

$$\left| \sum_{u_1 < \cdots < u_q} \left(\mathbf{1}_{\chi(\{u_1,\dots,u_q\})=c} - w_c\right) \prod_{J \in \binom{[q]}{q-1}} f_J(\boldsymbol{u}|_J) \right| \le \zeta n^q \tag{12.9}$$

simultaneously for every color c and all functions $f_J \colon [n]^J \to [0,1]$ for $J \in \binom{[q]}{q-1}$. For indicator tests, Hoeffding's inequality [15] gives failure probability at most $2\exp(-2\zeta^2 n^q)$ per test, whereas there are at most $2^{qn^{q-1}}$ tests. A union bound followed by integration of level sets proves the assertion.

Proof of Theorem 11.3. We first prove (i). Put $m := |C(P)|$ and $N := \binom{r}{q}$. For the lower bound, fix a weighting $\boldsymbol{w}$ and $d = m_B(P, \boldsymbol{w})$. Color the q-sets independently and keep exactly the admissible edges. Suppose Equation (12.9) holds. Fix any $\mathcal{S}$-test and partition its tuples by relative vertex order. In one order class, the color coordinates contained in a test face form some $I \in \mathrm{Orb}(B)$. In each admissible-pattern expansion, replace the indicators on coordinates $A \notin I$ by their expectations, one at a time. With vertices outside A fixed, every other factor depends only on a proper subset of A: no test face contains A, and every other q-face intersects it properly. Group these factors by containing $(q-1)$-faces and apply Equation (12.9). Each replacement costs at most ζn^r.

After at most $r! N m^N$ replacements, the coefficient for the actual supported colors on I is $E_{I,\boldsymbol{a}}(P, \boldsymbol{w}) \ge d$. Taking $\zeta \le \eta/(r! N m^N)$ proves the density bound simultaneously for all tests, including those selected after the coloring. Optimize $\boldsymbol{w}$.

For the upper bound, the case $\lambda_B(P) = 1$ follows from the lower bound. Otherwise fix $\lambda_B(P) < d \le 1$. Let $\kappa > 0$ be the deficit in Equation (5.3) for this palette and $\mathcal{I} = \mathrm{Orb}(B)$, and choose $0 < \delta < \kappa/4$. Apply Lemma 12.12 with $h = r$. It gives a supported test set K_0 with $|K_0| \ge cn^r$, a common weighting $\boldsymbol{x}$, and a joint color distribution ν within δ of $\boldsymbol{x}^{\otimes \mathcal{Q}}$ on every event.

Choose $I \in \mathrm{Orb}(B)$ and a supported assignment $\boldsymbol{a}$ realizing the deficit, and put $A = \{\boldsymbol{c}|_I = \boldsymbol{a}\}$. Under the product law, $d\mathbb{P}[A] - \mathbb{P}[P \cap A] \ge \kappa$, so

$$d\nu(A) - \nu(P \cap A) \ge \kappa - (d+1)\delta \ge \frac{\kappa}{2}.$$

The set $K_0 \cap A$ is an allowed test. After permuting cluster positions to match I, each prescribed q-face and each lower $(q-1)$-face lies in an allowed test face. Assign its color, lower-cell, and cluster restrictions to such a containing face, multiplying restrictions assigned to the same face. Since every edge has an admissible color vector,

$$d|K_0 \cap A| - e_H(K_0 \cap A) \ge \frac{\kappa}{2}|K_0| \ge \frac{c\kappa}{2} n^r.$$

Taking $\eta < c\kappa/2$ excludes $(d, \eta, \mathcal{S})$-density. This proves (i).

For (ii), the canonical forbidden families increase under the homomorphisms $P_{n+1} \to P_n$. Their densities are $\lambda_B(P_n)$ by (i), so Equation (2.2) gives density α for their union. □

12.2 Disjoint faces

The proof retains a bounded set of frequently used indices. A canonical Ramsey selection gives a common retained set, and the product-test bound controls equal or adjacent indices outside it.

Proof of Theorem 11.4. The lower bound is Lemma 12.1. We prove the uniform local upper bound.

The following list selection will supply the common retained index set.

Claim 12.13. *For integers $q, L, h \geq 1$ with $h \geq 2q$, there is M such that the following holds. Assign to every q-subset A of an ordered M-set a set K_A of at most L elements of an arbitrary universe. There are an h-subset of the ordered M-set and a fixed set C_0 such that, for all q-subsets A, B of the selected h-subset,*

$$K_A \cap K_B = C_0 \qquad \text{whenever } A \cap B = \varnothing.$$

Moreover, $C_0 \subseteq K_A$ for every q-subset A of the selected h-subset, and $|C_0| \leq L$.

Proof. Enumerate each list and pad it to length L with symbols private to that list. For each list position a, let $f_a(A)$ denote the entry in position a of the list assigned to A. Finite canonical Ramsey theory [6] gives a large subset and $I_a \subseteq [q]$ such that $f_a(x_1, \ldots, x_q) = f_a(y_1, \ldots, y_q)$ if and only if $x_i = y_i$ for every $i \in I_a$. Apply this successively to all a. A further finite Ramsey extraction fixes, for each ordered type of a pair of q-sets and each a, b, the truth value of $f_a(A) = f_b(B)$. There are only finitely many such types, witnessed on at most $2q$ vertices.

Take an h-set with unused buffer points immediately after every selected point in the larger extracted set. Suppose A, B are disjoint and $f_a(A) = f_b(B)$. If I_a were nonempty, move a determining coordinate of A to a buffer point. This does not change the ordered pair type and hence preserves the equality to the unchanged $f_b(B)$, but canonicality changes $f_a(A)$, a contradiction. Thus $I_a = \varnothing$, and likewise $I_b = \varnothing$. Both entries are constant on the extracted set. Their common value therefore occurs in every list. Private padding symbols cannot have this property. Let C_0 consist of the values that occur in every list. The preceding argument proves both inclusions in the desired equality. □

For a finite set $C_0 \supseteq R$ of retained indices, keep $F_0 \cup (C_0 \times A_0)$ and introduce a symbol $*_a$ for each indexed type. When evaluating the defining rule, interpret every star occurrence as a fresh unnamed index, different from all other fresh indices and nonadjacent to them and to C_0. Call the resulting finite relation $\widehat{P}$. In particular, 2 occurrences of the same star are evaluated at distinct fresh indices; the stars specify a finite admissibility rule rather than actual original colors.

Choose the parameters in the following order.

(i) Choose $\tau, \zeta > 0$ so that $2(\Delta + 1)|E(J)|(\tau + \zeta) < \varepsilon/10$; if J has no edges, take any small positive choices. Put $L := \lceil (\Delta + 1)/\tau \rceil + |R|$. The compressed-color bound is $|F_0| + |A_0|(L + 1)$.

(ii) Choose a target number of clusters $h \geq \max\{r, 2q\}$, and at $q = 2$ additionally require $\binom{r}{2}/h < \varepsilon/10$. Choose a marginal error ξ with $N\xi < \varepsilon/10$. In each dimension up to the compressed-color bound, choose a finite simplex grid whose cells have ℓ^1 diameter at most ξ. Determine the Ramsey sizes for these grids and for Claim 12.13.

(iii) Choose a top regularity error δ_{comp} for the compressed colors as required by Lemma 12.11 with counting error $\varepsilon/10$. Choose the top regularity error δ_q for the original color classes small enough that every compressed-color indicator has this regularity; each is a sum or difference of at most $L + 1$ fine- or coarse-color indicators. Normalizing the absolute count discrepancy by the polyad's supported q-set count gives a bound of $2(L + 1)\delta_q$, so taking $\delta_q \ll \delta_{\mathrm{comp}}^2/(L + 1)$ is enough.

(iv) Prescribe functions specifying the lower-level regularity tolerances to ensure Lemma 12.7 with error ζ and all required counts.

All Ramsey parameters and top regularity errors are fixed before applying the regularity lemma.

Apply Lemmas 12.3 and 12.6 to a witnessing ordered q-face coloring of H. The coarse partition has a separate class for each finite color and one class for each indexed type. We obtain enough ordered clusters and a compatible $(q-1)$-complex, with every individual color class and every type class regular on every q-face boundary.

For a q-subset of clusters U, let p_i^U be the probability of index i under the uniform distribution on its internal boundary-supported q-sets; sum over all indexed types. Let K_U be the union of R and the closed neighborhoods in Γ of all indices i with $p_i^U > \tau$. Then $|K_U| \le L$. Apply Claim 12.13, retaining sufficiently many clusters for the subsequent simplex-grid Ramsey selection. Obtain a common finite set C_0, containing R, such that $K_U \cap K_V = C_0$ for disjoint q-sets U, V.

Use the relation $\widehat{P}$ defined above for this choice of C_0. We first prove that

$$\lambda(\widehat{P}) \le \lambda(P).$$

Given an optimal probability weighting of $\widehat{P}$ and an integer K, choose K unused indices for each type, all outside the closed neighborhood of $C_0 \cup R$, with no adjacency among any chosen indices. Bounded degree and infinitude permit the greedy choice. Replace each star weight by equal weights on its chosen original colors. A sampled tuple repeats a new index with probability at most $\binom{N}{2}/K$. Off this event, the new indices are distinct and nonadjacent, so equality, adjacency, and every R_t agree with the star interpretation. Their negations agree as well. The transferred weighting gives $\lambda(P) \ge \lambda(\widehat{P}) - \binom{N}{2}/K$. Letting $K \to \infty$ proves the desired inequality.

We now apply this compression to the host coloring, sending every other indexed color of type a to $*_a$. Original and compressed admissibility can differ only when 2 compared indices are equal or adjacent and at least 1 lies outside C_0.

Fix a compared pair of coordinates on disjoint faces U, V. In the product of the reference internal distributions, if both indices have probability greater than τ and are equal or adjacent, each lies in $K_U \cap K_V = C_0$. A related pair with at least one endpoint outside C_0 therefore has an endpoint of probability at most τ. The contribution from such indices on the first side is at most

$$\sum_{i:\, p_i^U \le \tau} p_i^U \sum_{k:\, k=i \text{ or } ik \in E(\Gamma)} p_k^V \le \tau \sum_k p_k^V (1 + \deg_\Gamma(k)) \le (\Delta+1)\tau.$$

Exchanging the sides gives a product-law bound of $2(\Delta+1)\tau$. The relation "equal or adjacent, with at least one endpoint outside C_0" has at most $\Delta+1$ entries per row and column. By Lemmas 12.7 and 12.8, its probability under the uniform distribution on the supported transversal r-sets is at most $2(\Delta+1)\tau + (\Delta+1)\zeta$. A union bound over $E(J)$ bounds the total probability of these exceptional comparison events. Type or position violations need no such estimate; the finite rule rejects them both before and after compression.

We next verify the regularity of the compressed color classes so that Lemma 12.11 applies. For the original face coloring χ, the star indicator of type a is

$$\mathbf{1}_{\chi \in (I \setminus C_0) \times \{a\}} = \mathbf{1}_{\chi \in I \times \{a\}} - \sum_{i \in C_0} \mathbf{1}_{\chi = (i,a)}.$$

The sum has at most $|C_0|$ terms. Hence each star indicator has the required top regularity by the choice of δ_q, for all the same subboundary tests.

The compressed color set has at most $|F_0| + |A_0|(L+1)$ elements, labeled consistently on all selected faces. A further q-uniform Ramsey extraction using the simplex grid in this dimension gives enough clusters so that every internal compressed color distribution is within ξ in ℓ^1 of a fixed distribution $\boldsymbol{x}$.

Fix any r of these clusters. By Lemma 12.11, the compressed admissibility probability on their supported transversal r-sets is at most its probability under the product of the reference face-color distributions plus $\varepsilon/10$. The marginal replacement estimate in the proof of Lemma 12.12 bounds the additional error by $N\xi$. Under the common product law $\boldsymbol{x}^{\otimes N}$, the admissibility probability is $\Lambda_{\widehat{P}}(\boldsymbol{x}) \le \lambda(P)$. The total error is at most $\varepsilon/10 + N\xi + |E(J)|[2(\Delta+1)\tau + (\Delta+1)\zeta] < 3\varepsilon/10 < \varepsilon/2$. Thus the original admissibility probability is at most $\lambda(P) + \varepsilon/2$.

For $q \geq 3$, take any r selected clusters and let G be the union of their selected $(q-1)$-cells. Since $q-1 \geq 2$, every r-clique of G is transversal. The family $\mathcal{K}_r(G)$ consists exactly of the r-sets supported by the selected complex and has size at least cn^r by (12.1). For $q = 2$, take the union of h selected clusters. The proportion of r-sets using some cluster twice is at most $\binom{r}{2}/h + o(1)$; counting all other sets by their class choices gives the same upper bound with final error ε.

If a P-colorable r-graph were $(\lambda(P) + 2\varepsilon, \eta, q-1)$-dense, its density inequality on this test would give $\varepsilon c \leq \eta$. Choosing $\eta < \varepsilon c$ proves the canonical-family upper bound. □

12.3 One-sided faces

Proof of Theorem 11.5. The lower bound is Lemma 12.1. If P is empty, every colorable host has no edges and the local upper bound is immediate. We therefore assume $P \neq \varnothing$ and prove the local upper bound. Fix a witnessing coloring of H and put $n := |V(H)|$.

Consider a layered $(q-1+h, q-1)$-complex on equal-sized ordered clusters, where $q \geq 3$ and $h \geq 1$. Here h is the number of private positions. The core positions are $C := [q-1]$, and the private positions are $v_t := q-1+t$ for $t \in [h]$. A *petal* is $C \cup \{v_t\}$. Retain all lower-cell indicators whose position sets are contained in a petal and discard all other lower-cell indicators. Call the resulting support the *petal product*.

Let μ_C be uniform measure on supported cores. For a supported core z, call a vertex in the tth private cluster usable if it satisfies all retained lower-cell constraints on the tth petal together with z. Let $\ell_t(z)$ be the proportion of usable vertices in that cluster and let p_z^t be their conditional color distribution; choose any probability distribution when $\ell_t(z) = 0$.

Every retained constraint involves the core and at most 1 private vertex. For a fixed supported core z, the permitted private tuples therefore form the Cartesian product of the usable-vertex sets. Under uniform sampling of the petal-product support, conditional on a core with positive support, the private vertices are consequently independent uniform choices from these sets. The number of such choices is proportional to $\prod_{t\in[h]} \ell_t(z)$, so the core marginal has density proportional to this product with respect to μ_C. Figure 16 shows the retained constraints and the conditional choices.

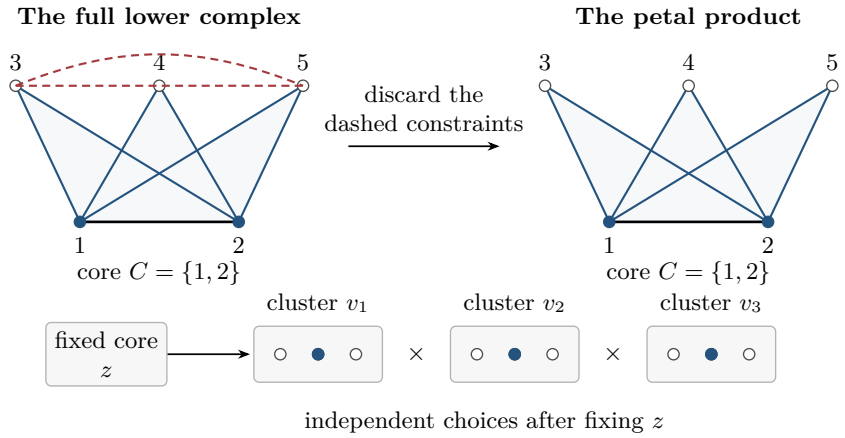


Figure 16: For $q = 3$ and 3 private positions, vertices in the upper panels denote positions and edges denote lower-cell constraints. The petal product retains precisely the constraints within $C \cup \{v_t\}$. Conditional on a supported core z, usable private vertices are chosen independently; their numbers weight the core marginal by $\prod_t \ell_t(z)$.

Claim 12.14. *The lower-level regularity tolerances may be chosen to satisfy the following statements with arbitrarily small counting and marginal errors. Counting errors are measured relative to the full supported count.*

(i) *For any product of a $[-1,1]$-valued core function and one such function on each petal, the weighted count on the full complex equals the weighted count on the petal product multiplied by the product of the discarded cell densities, up to the prescribed error.*

(ii) *If the indices of 2 indexed petal colors are tested for membership in a relation of bounded row degree, the comparison in (i) between counts on the full complex and on the petal product holds with error independent of the number of indices. The relation can also be restricted by conditions on an index and the common core.*

Proof. For (i), apply the replacement argument in the proof of Lemma 12.7 with retained sets $C \cup \{v_t\}$ for $t \in [h]$. Multiply the core function into one petal function. The retained lower-cell indicators define exactly the petal product. Choosing the lower errors small relative to the product of lower densities gives the required relative count error.

For (ii), let R_z be the sparse relation after imposing all conditions depending on the core z, and let D bound its row degrees. Use the sign representation in Lemma 12.8, now with

$$f_\epsilon(z,i) = \frac{1}{D} \sum_{k:\ (i,k)\in R_z} \epsilon_k, \qquad g_\epsilon(z,k) = \epsilon_k.$$

Set both factors to zero on nonindexed colors. Their product has sign expectation $D^{-1}\mathbf{1}_{(i,k)\in R_z}$. Each factor is a permitted petal function, so (i), followed by averaging, gives error at most D times the product-test error. For a retained set $J(z)$, take $R_z = \{(i,k) \in R\colon \{i,k\} \nsubseteq J(z)\}$. A further core restriction $z \in S$ is included by multiplying the first factor by $\mathbf{1}_S(z)$. □

We now return to the fixed host H and its uniformity r.

Write $h := r - q + 1$. Choose the parameters in the following order.

(i) Choose $0 < \tau, \xi < 1$ as the tail threshold and marginal tolerance, so that $2(\Delta+1)\binom{h}{2}\tau + N(2|F_0|+2)\xi < \varepsilon/2$. Put $K := \lceil (\Delta+1)/\tau \rceil + |R|$ and $B := |F_0| + |A_0|(K+1)$, where R is the set of named indices. Partition the probability simplex on B coordinates into g cells of ℓ^1 diameter at most ξ; vectors on fewer colors are padded with zeros.

(ii) Choose the number m of candidate private clusters large enough for Claim 12.13 with singleton index sets, list bound K, and target size hg. Put $\alpha := (2\binom{m}{h})^{-1}$, which will be a lower bound on the measure of cores selecting one fixed set of h private clusters.

(iii) Choose a relative replacement tolerance $\eta > 0$ so that $2NB^N\eta/\alpha < \varepsilon/4$. Choose a fixed top error $0 < \delta \le \eta^2/16$ small enough for the good-core selection, in terms of m, B, α, ξ.

(iv) Prescribe lower-level regularity tolerances and the number of allowed top tests as functions of the density vector to apply Lemma 12.10 and Claim 12.14 on the full selected r-complex. The required lower-level counts include the core, each petal, the full r-complex, and the petal product. Include 2 copies of each of the latter 3 configurations identified on the core, and 2 copies of the full complex identified on a q-face. Apply Lemma 12.9 so that each single-petal core marginal is within δ of μ_C, and the full-complex and petal-product core marginals are each within $\alpha/2$ of μ_C in total variation. Choose the lower-level tolerances simultaneously for this finite collection of patterns. Require the aggregate lower-level comparison error, including the errors for at most B^N color patterns, to contribute less than $\varepsilon/4$ after normalization by the restricted test volume.

All these parameters, including the required number of clusters, are fixed before applying the regularity lemma.

Use Lemmas 12.3 and 12.6 and a finite-grid q-uniform Ramsey selection to obtain ordered clusters, a compatible lower complex, and nearly equal marginal probabilities on the selected q-face boundaries for every finite color and every indexed type. Retain the first $q-1$ selected clusters as core clusters, with position set $C := [q-1]$, and the next m selected clusters as candidate private clusters, with positions $v_1 < \cdots < v_m$. For a q-face U of these clusters, let $p_U(f)$ be the proportion of internally supported q-sets on U colored f. In particular, there are numbers a_f such that

$$|p_U(f) - a_f| \le \xi \quad \text{for every } f \in F_0 \tag{12.10}$$

on every q-face boundary U under consideration. Top regularity is required for each fine color class and for each class of the coarse partition, which has a separate class for each color of F_0 and one class for each indexed type.

We next select cores with controlled finite-color probabilities. For a candidate private cluster t and supported core z, let p_z^t and $\ell_t(z)$ be as above. For a finite color f, consider separately the sets of cores where $p_z^t(f)$ is above or below its face marginal $p_{C\cup\{v_t\}}(f)$. For either set, restrict the core $(q-1)$-cell to that set and retain all other cells of the petal boundary. The resulting subboundary is a valid top-regularity test. Applying top regularity to these 2 tests gives

$$\int \ell_t(z)|p_z^t(f) - p_{C\cup\{v_t\}}(f)|\,\mathrm{d}\mu_C(z) \le 2\delta \int \ell_t(z)\,\mathrm{d}\mu_C(z). \tag{12.11}$$

For a core test below the regularity threshold, the discrepancy is bounded by the corresponding supported-face measure. Let ν_t be the core marginal of a uniformly sampled internally supported face on the tth petal. Dividing Equation (12.11) by $\int \ell_t\,\mathrm{d}\mu_C$ gives an integral bound 2δ under ν_t. By Equation (12.7), $\|\nu_t - \mu_C\|_{\mathrm{TV}} \le \delta$. Since the integrand lies in $[0,1]$,

$$\int |p_z^t(f) - p_{C\cup\{v_t\}}(f)|\,\mathrm{d}\mu_C(z) \le 3\delta.$$

Markov's inequality and a union bound over t show that the cores with $\sum_{f\in F_0} |p_z^t(f) - p_{C\cup\{v_t\}}(f)| > \xi$ for some t have measure at most $3m|F_0|\delta/\xi$. Choose δ so this is at most $1/2$. Combining with Equation (12.10), at least half the supported cores satisfy

$$\sum_{f\in F_0} |p_z^t(f) - a_f| \le (|F_0|+1)\xi \qquad \text{for every } 1 \le t \le m. \tag{12.12}$$

If F_0 is empty, all cores satisfy this condition.

For each good core z, let $K_t(z)$ be the union of R and the closed neighborhoods in Γ of all indices i with $p_z^t(\{i\}\times A_0) > \tau$. Each $K_t(z)$ has at most K elements. Apply Claim 12.13 with singleton index sets to choose hg lists whose pairwise intersections all equal a common set $J(z)$.

We next use the probability grid to select h petals with close compressed conditional color distributions. Retain $J(z)\times A_0$ and replace all other indexed colors of type a by $*_a$. The core of a petal face T is $T\setminus\{\max T\}$, so its compressed label is determined by T. Write $\widehat{p}_z^t$ for the resulting compressed conditional color distribution on the tth petal. Enumerate the elements of $J(z)$ by a fixed enumeration of I, using zero-weight padding so the compressed vectors have at most B labeled coordinates. Among these lists, h compressed conditional probability vectors lie in the same grid cell. Thus, for the selected private clusters, indexed by $t_1,\dots,t_h$,

$$K_{t_u}(z)\cap K_{t_v}(z) = J(z), \quad \|\widehat{p}_z^{t_u} - \widehat{p}_z^{t_v}\|_1 \le \xi \quad \text{for distinct } u,v\in[h]. \tag{12.13}$$

To build the density test on a fixed set of clusters, use a fixed tie-breaking rule and restrict to cores for which the same private clusters are chosen. There are only $\binom{m}{h}$ possible sets of selected private clusters. Some fixed choice occurs on a core set S with $\mu_C(S)\ge\alpha$.

Restrict the selected core $(q-1)$-cell to S, retaining all other selected $(q-1)$-cells. Their union is a legitimate $(q-1)$-graph G. As $q-1\ge 2$, an r-clique in this union is transversal to the r selected clusters. The r-cliques of G are exactly the r-sets with core in S that are supported by the full selected lower complex. Let W be the number of r-sets supported by the selected full complex without the restriction to S.

Let ν_{full} be the core marginal of uniform sampling from the full supported complex. The choice of lower tolerances and Equation (12.7) give $\|\nu_{\mathrm{full}} - \mu_C\|_{\mathrm{TV}} \le \alpha/2$. Since $\mu_C(S)\ge\alpha$,

$$|\mathcal{K}_r(G)| = W\nu_{\mathrm{full}}(S) \ge \frac{\alpha}{2}W \ge c\,n^r > 0. \tag{12.14}$$

The same marginal estimate for the petal product gives positive support after restricting its core to S.

We estimate the probability that 2 compared positions carry equal or adjacent indices, with at least 1 index not retained. Fix $z\in S$ and a compared pair of petals t_u, t_v, and temporarily work in the petal product. The private vertices are independent conditional on z. Consider the event that both petal colors are indexed, their indices i,k are equal or adjacent, and at least one lies outside $J(z)$. If both indices have conditional probability greater than τ on their respective petals, both belong to $K_{t_u}(z)\cap K_{t_v}(z) = J(z)$. Consequently, with probability taken over the original petal colors conditional on z,

$$\mathbb{P}[i=k \text{ or } ik\in E(\Gamma),\ \{i,k\}\not\subseteq J(z) \mid z] \le 2(\Delta+1)\tau. \tag{12.15}$$

The summation in the proof of Theorem 11.4, applied to the conditional index marginals at z, bounds the contribution of indices of probability at most τ on the first side by $\tau \sum_k p_k(1 + \deg_\Gamma(k)) \le (\Delta + 1)\tau$, where p_k is the index marginal on the other petal. Exchange the sides for the other contribution. The same bound covers all the finitely many labeled relations, since they are supported on the diagonal and edges of Γ. The conditional bound $2(\Delta + 1)\tau$ holds for every $z \in S$, so averaging over the core marginal of the petal product restricted to S gives the same bound in that product.

Apply Claim 12.14(ii) to the exceptional-event indicator multiplied by $\mathbf{1}_S$, and part (i) of that claim to $\mathbf{1}_S$. Dividing the resulting count estimates by $|\mathcal{K}_r(G)| \ge \alpha W/2$, as guaranteed by (12.14), gives the same probability bound under uniform sampling of $\mathcal{K}_r(G)$, up to the prescribed lower-level comparison error. Sum over the fixed number of compared position pairs.

The private clusters are fixed on S, while the retained index set $J(z)$ may still vary with z. We therefore compare a separate finite palette with P at each core before averaging. For each selected core $z \in S$, let $\widehat{P}_z$ be the palette on F_0, the retained labeled colors, and the stars. Evaluate each star occurrence as a distinct, unnamed index, nonadjacent to all retained and fresh indices. Then

$$\lambda(\widehat{P}_z) \le \lambda(P). \tag{12.16}$$

The argument proving $\lambda(\widehat{P}) \le \lambda(P)$ in Theorem 11.4 applies with retained index set $J(z)$, which contains all named indices.

We now use the palettes $\widehat{P}_z$ to bound the admissible-pattern count. Treat all selected petals as variable coordinates, including those where the rule allows only finite colors. Every edge of H has an admissible original pattern, and original and compressed admissibility can differ only on the exceptional comparison events bounded in (12.15). Expand the compressed admissibility indicator as a sum over at most B^N color patterns. Each term is the product of a $\{0,1\}$-valued function of the core, arbitrary compressed color indicators on the variable petals, and finite-color indicators at all other coordinates. The core function may depend on the entire finite relation table of $J(z)$.

The retained set $J(z)$ depends on the core, so the compressed color indicators on the petals need not be top-regular. At the other coordinates the indicators are original finite-color classes, which are top-regular. Apply Lemma 12.10 to these indicators. For a target q-face, every different q-face and the common core intersect it in a proper subset; the other factors therefore form admissible boundary tests after outside vertices are fixed. The absolute error is at most $NB^N\eta W$; by (12.14), it is at most $2NB^N\eta/\alpha$ times $|\mathcal{K}_r(G)|$.

After the nonvariable color indicators have been replaced by their densities, each pattern has only a core factor and petal factors, multiplied by constants. It therefore has the form required by Claim 12.14(i), so we discard the lower-cell constraints whose position sets are not contained in a petal, charging the comparison error in Claim 12.14(i) for each of the at most B^N patterns. In the petal product, conditional on z, the petal colors are independent with marginals $\widehat{p}_z^{\,t_u}$. Choose any one of these marginal distributions as $\boldsymbol{x}_z$. By (12.13), the marginal replacement error is at most $\sum_{u\in[h]} \|\widehat{p}_z^{\,t_u} - \boldsymbol{x}_z\|_1 \le N\xi$.

At every nonvariable coordinate admissibility requires a color in F_0, so only the finite-color part of its marginal matters. For every nonvariable face U, (12.10) and (12.12) give $\sum_{f\in F_0} |p_U(f) - (\boldsymbol{x}_z)_f| \le (2|F_0| + 1)\xi$. Replacing these finite-color subprobability vectors changes the admissibility probability by at most $\sum_{U \text{ nonvariable}} \sum_{f\in F_0} |p_U(f) - (\boldsymbol{x}_z)_f| \le N(2|F_0| + 1)\xi$. With the common marginal $\boldsymbol{x}_z$ at every coordinate, the admissibility probability is $\Lambda_{\widehat{P}_z}(\boldsymbol{x}_z)$. Combining the two marginal-replacement errors therefore bounds the probability before these replacements by $\Lambda_{\widehat{P}_z}(\boldsymbol{x}_z) + N(2|F_0| + 2)\xi$, which is at most $\lambda(P) + N(2|F_0| + 2)\xi$ by (12.16).

Averaging over the core marginal of the petal product restricted to S gives

$$\frac{\displaystyle\int_S \Lambda_{\widehat{P}_z}(\boldsymbol{x}_z) \prod_{u\in[h]} \ell_{t_u}(z)\,\mathrm{d}\mu_C(z)}{\displaystyle\int_S \prod_{u\in[h]} \ell_{t_u}(z)\,\mathrm{d}\mu_C(z)} \le \lambda(P).$$

The denominator is positive by the core-marginal estimate for the petal product.

Thus $|H \cap \mathcal{K}_r(G)|/|\mathcal{K}_r(G)| - \lambda(P)$ is at most

$$2(\Delta+1)\binom{h}{2}\tau + N(2|F_0|+2)\xi + \frac{2NB^N\eta}{\alpha},$$

plus the aggregate lower-level comparison error. By the chosen error budgets, the first 2 terms sum to less than $\varepsilon/2$, the third is less than $\varepsilon/4$, and the aggregate lower-level comparison error is less than $\varepsilon/4$. This proves the local upper bound and hence the theorem. $\square$

Part III

Turán densities for general test systems

13 Completion degrees and conditional densities

We use the completion capacity Γ_k and the cyclic substitution from Lemma 5.5. The recursive symmetric templates give a small interval of limits; an ordered lift preserves their completion constraints in every number of coordinates.

Proposition 13.1. *For every $k \geq 3$ and $\alpha \in [0,1]$, there are finite k-coordinate palettes P_n and homomorphisms $P_{n+1} \to P_n$ with $\Gamma_k(P_n) \to \alpha$.*

We first transfer the recursive templates before their final cyclic substitutions. Let $P = P_n(h)$ be a template from Proposition 5.7, and let $\tau\colon C(P) \to \mathbb{Z}_3$ be its top-class label. The rules have 3 properties. Triples within one top class are forbidden; triples in 3 distinct top classes are allowed; and for a color c in coarse class D, the diagonal completion set $N_P(c,c)$ is its successor coarse class $g(D)$. The map g cycles through all coarse classes. The bonding maps preserve these labels.

For $k \geq 4$, define a k-coordinate palette $L_k(P)$ on the same colors. A tuple whose top labels all agree is forbidden. If exactly $k-1$ top labels agree, take the colors in the first 2 majority positions and the color in the singleton position; admit the tuple precisely when this triple belongs to P. Every remaining top-label pattern is admitted. The use of the first 2 positions makes this an ordered relation.

Lemma 13.2. *For every weighting $\boldsymbol{w}$ and every $d > 0$, the weighting is d-feasible for P if and only if it is d-feasible for $L_k(P)$. Thus $\Gamma_k(L_k(P)) = \Gamma(P)$, and the lifts preserve the bonding homomorphisms of the recursive templates.*

Proof. Suppose first that $\boldsymbol{w}$ is d-feasible for P. By Lemma 5.8(i), every coarse class, and hence every top class, has weight at least d. Prescribe $k-1$ supported colors. If their top labels agree, the permitted completions form one original same-top row $N_P(a,b)$, where a, b occupy the selected majority positions; its weight is at least d. If their labels do not agree, a completed tuple can have $k-1$ equal labels only by appending a label already occurring $k-2$ times. There is at most one such label because $k \geq 4$. All colors in any other top class are therefore completions, giving weight at least d.

Conversely, suppose $\boldsymbol{w}$ is d-feasible for $L_k(P)$. Prescribing one supported color c in all $k-1$ positions gives completion set exactly $N_P(c,c) = g(D)$, since both relations forbid a tuple within one top class. Iterating the successor cycle shows that every coarse class has weight at least d. For supported same-top colors a, b, leave the last coordinate free, put a, b in the first 2 positions, and fill the remaining prescribed positions with a. This reproduces exactly the row $N_P(a,b)$, so it has weight at least d. Two colors in distinct top classes have the entire third top class as completions in P. Hence every original row has weight at least d.

The bonding maps preserve top labels and preserve the selected original triple. They consequently preserve each of the 3 rules defining $L_k(P)$. □

Proof of Proposition 13.1. By Proposition 5.12, limits of the capacities of the recursive templates contain a fixed interval $[0, \varepsilon]$. Lemma 13.2 gives such an interval for every $k \geq 4$, while $k = 3$ follows directly from the same templates. Zero is supplied by an empty one-color relation.

Apply Lemma 5.5 to extend this interval to $[0,1]$. □

For three-uniform cherry density, the test system is $\{\{1,2\},\{2,3\}\}$. Its pair shadow consists of 2 coordinates out of 3, and its position orbit contains every such pair. Its conditional capacity is therefore Γ_3.

Proof of the cherry case of Theorem 1.1(i). Every singleton position is contained in one of the two test faces. Thus Theorem 11.3(ii) applies with $r = 3$, $q = 2$, and $\mathcal{S} = \{\{1,2\},\{2,3\}\}$ to the chains from Proposition 13.1. □

14 Three-uniform vertex-pair density

We prove the vertex-pair case of Theorem 1.1(i), with test system $\{\{1\},\{2,3\}\}$. The realization theorem follows from the common conditional-capacity theorem. As in the introduction, we denote the density parameter by π_{vp} and its set of possible values by $\Pi^{(3)}_{\mathrm{vp},\infty}$.

Number the pair coordinates $1,2,3$, corresponding to xy, xz, yz for $x<y<z$. For a finite ternary palette, put

$$g_i^P(c;\boldsymbol{w}) \coloneqq \sum_{\substack{\boldsymbol{a}\in P\\ a_i=c}} \prod_{j\neq i} w_{a_j}, \qquad m_i(P,\boldsymbol{w}) \coloneqq \min_{c\in\operatorname{supp}\boldsymbol{w}} g_i^P(c;\boldsymbol{w}), \qquad \Delta(P) \coloneqq \max_{\boldsymbol{w}} \min_{i\in[3]} m_i(P,\boldsymbol{w}).$$

By Lemma 5.3, with $J=[3]$ and $\mathcal{I}=\{\{1\},\{2\},\{3\}\}$, the maximum is attained and Δ is nondecreasing under homomorphisms. Moreover,

$$\sum_c w_c g_i^P(c;\boldsymbol{w}) = \Lambda_P(\boldsymbol{w}), \tag{14.1}$$

so $\Delta(P)\le\lambda(P)$.

To realize every value of Δ, we first construct an interval for the minimum degree in one coordinate. Mixing two palettes converts these directional values into the interval $[0,1/4]$ for Δ. Cyclic substitutions then extend this interval to $[0,4/9]$ and finally to $[0,1]$.

Proposition 14.1. *For every finite palette P, $\pi_{\mathrm{vp}}(\mathcal{F}_{3,2}(P))=\Delta(P)$. If $P_{n+1}\to P_n$, then*

$$\pi_{\mathrm{vp}}\left(\bigcup_n \mathcal{F}_{3,2}(P_n)\right) = \lim_n \Delta(P_n).$$

Proof. For $\{\{1\},\{2,3\}\}$, the only pair contained in a test face is $\{2,3\}$; its orbit contains all 3 pair coordinates. Every singleton position is contained in a test face, so Theorem 11.3 applies. □

Let $\mathcal{R}_{\mathrm{vp}}$ be the set of limits of $\Delta(P_n)$ along projective sequences of finite palettes. It suffices to prove $\mathcal{R}_{\mathrm{vp}}=[0,1]$. Constant sequences are allowed.

Lemma 14.2. *We have $\mathbb{Q}\cap[0,1]\subseteq\mathcal{R}_{\mathrm{vp}}$.*

Proof. For $0<a<b$, take colors $\mathbb{Z}_b$ and admit (c_1,c_2,c_3) when $c_2-c_1\in[a]$ modulo b. Uniform weights give $\Delta=a/b$. For any weighting, the first-coordinate minimum is at most a/b by Equation (5.1), giving the reverse bound. Empty and full one-color relations give the endpoints. □

14.1 A directional construction and its optimum

We begin with palettes for which the first two coordinates have a zero-degree supported color for every weighting, and vary the third-coordinate minimum. Define

$$\kappa_i(P) \coloneqq \max_{\boldsymbol{w}\in S(C(P))} m_i(P,\boldsymbol{w}) \qquad \text{for } i\in[3].$$

Apply Lemma 5.3 with $\mathcal{I}=\{\{i\}\}$ to obtain attainment and monotonicity under homomorphisms.

Proposition 14.3. *There are constants $0<a<b<1/2$ such that every $t\in[a,b]$ is realized by a sequence of finite palettes $Q_{n+1}\to Q_n$ with*

$$m_1(Q_n,\boldsymbol{w})=m_2(Q_n,\boldsymbol{w})=0 \quad \text{for every } \boldsymbol{w}\in S(C(Q_n)), \qquad \kappa_3(Q_n)\downarrow t. \tag{14.2}$$

Every $\kappa_3(Q_n)$ is attained at the uniform weighting, and $0<\kappa_3(Q_n)<1/2$.

The construction modifies the palette tower of Li, Liu, and Pikhurko [23]. We use their graph 2-lifts and their cubic graph estimate. The changes are that the first root color must be smaller than the second root color, and the optional deletions need only preserve regularity in the third coordinate.

Graph estimates and the root palette. For a simple graph D and nonnegative weights $\boldsymbol{x}$, define

$$E_D(\boldsymbol{x}) \coloneqq \sum_{uv \in E(D)} x_u x_v (x_u + x_v).$$

For a d-regular graph D_0 on N_0 vertices, identify $C_0 \coloneqq V(D_0)$ with $\mathbb{Z}_{N_0}$, ordered by the representatives $0, \dots, N_0 - 1$. Define Q_0 by allowing precisely the tuples (a, b, c) with $a < b$ for which either a, b, c are distinct, or $ab \in E(D_0)$ and $c \in \{a, b\}$. Its polynomial is

$$\Lambda_{Q_0}(\boldsymbol{x}) = 3e_3(\boldsymbol{x}) + E_{D_0}(\boldsymbol{x}), \qquad \Lambda_{Q_0}(\boldsymbol{u}_{N_0}) = \frac{1}{2} - \frac{3}{2N_0} + \frac{1+d}{N_0^2}.$$

Each color has exactly $\binom{N_0-1}{2} + d$ completions when fixed in coordinate 3.

To prove that this polynomial is maximized at uniform weights, we use the following estimate. By [23, Lemma 4.1], for each fixed d there are $c_d > 0$ and $N_{\mathrm{root}}(d)$ such that every d-regular graph D of order $N \geq N_{\mathrm{root}}(d)$ satisfies

$$\big(6e_3(\boldsymbol{u}_N) + 3E_D(\boldsymbol{u}_N)\big) - \big(6e_3(\boldsymbol{x}) + 3E_D(\boldsymbol{x})\big) \geq c_d \|\boldsymbol{x} - \boldsymbol{u}_N\|_2^2 \tag{14.3}$$

for every $\boldsymbol{x} \in S(V(D))$. The polynomial e_3 is maximized at the uniform weighting. Indeed, averaging two weights while keeping their sum fixed does not decrease e_3, since the coefficient of their product is the sum of the remaining nonnegative weights. If $N_0 \geq N_{\mathrm{root}}(d)$, the identity $3\Lambda_{Q_0} = (6e_3 + 3E_{D_0}) + 3e_3$ and Equation (14.3) give

$$3\big(\Lambda_{Q_0}(\boldsymbol{u}_{N_0}) - \Lambda_{Q_0}(\boldsymbol{x})\big) \geq c_d \|\boldsymbol{x} - \boldsymbol{u}_{N_0}\|_2^2. \tag{14.4}$$

A 2-*lift* replaces each vertex v by v^+, v^- and each edge uv by one of the 2 matchings $\{u^+v^+, u^-v^-\}$ and $\{u^+v^-, u^-v^+\}$. The *complementary lift* uses the other matching at each edge. The two lifts partition the edges of the complete blowup with 2 copies of each vertex, and both project to the original graph.

Choose D_0 from the regular 2-lift tower of Li, Liu, and Pikhurko [23, Proposition 2.7]. For sufficiently large d it provides d-regular graphs D_n with $|V(D_n)| = N_n = 2^n N_0$, where D_{n+1} is a 2-lift of D_n and D_{n+1}^- is its complementary lift. Their paired vertices are the fibers of the natural projections. All graphs are d-regular and all nontrivial adjacency eigenvalues of D_0, D_{n+1}, D_{n+1}^- have absolute value at most $d/100$. We fix such a $d \geq 480$. By taking preliminary lifts, choose the initial order so that

$$N_0 \geq \max\left\{N_{\mathrm{root}}(d),\ 7,\ \frac{19d}{5c_d},\ 10(d+1)\right\}. \tag{14.5}$$

Let $\mathrm{pr}_{n,j}$ be the natural projection from level n to level j.

Refinement and optional deletions. Write $C_n \coloneqq V(D_n)$. We will preserve the following property:

$$(a, b, c) \in Q_n \quad \Longrightarrow \quad \mathrm{pr}_{n,0}(a) < \mathrm{pr}_{n,0}(b). \tag{14.6}$$

The order in this condition is the fixed order on C_0, not a new order on the fine colors. For every weighting, a supported color with largest root label has zero degree in coordinate 1, and a supported color with smallest root label has zero degree in coordinate 2. Thus Equation (14.6) implies the two zero-degree assertions in Equation (14.2).

In addition, every edge uv of D_n, written so that $\mathrm{pr}_{n,0}(u) < \mathrm{pr}_{n,0}(v)$, will have the two tuples (u, v, u) and (u, v, v) in Q_n. Every edge of a lift projects to an edge of D_0, so its two root labels are distinct.

Given Q_n, first take its full pullback to C_{n+1}. For each edge uv of the complementary lift D_{n+1}^-, choose the endpoint order so that $\mathrm{pr}_{n+1,0}(u) < \mathrm{pr}_{n+1,0}(v)$ and delete

$$(u, v, u), \qquad (u, v, v). \tag{14.7}$$

These tuples were present in the pullback by the preceding edge property. If $\boldsymbol{y}$ is the pushforward of a child weighting $\boldsymbol{x}$, the resulting palette $\widetilde{Q}_{n+1}$ satisfies the exact identity

$$\Lambda_{\widetilde{Q}_{n+1}}(\boldsymbol{x}) = \Lambda_{Q_n}(\boldsymbol{y}) - E_{D_{n+1}^-}(\boldsymbol{x}). \tag{14.8}$$

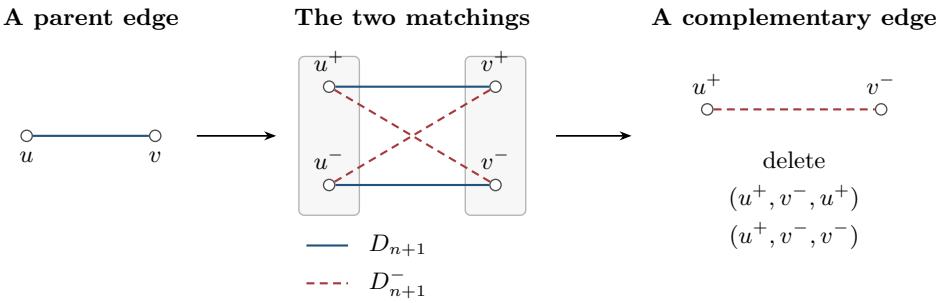


Figure 17: One choice of the two matchings above a parent edge uv, with the root label of u smaller than that of v. Each edge of D^-_{n+1} removes its two repeated-color tuples from the full pullback. The tuples belonging to edges of D_{n+1} remain available for the next refinement.

Figure 17 shows the two matchings above a parent edge and the tuples removed along a complementary edge.

To estimate the loss $E_{D^-_{n+1}}$, we use the cubic graph bound of Li, Liu, and Pikhurko [23, Lemma 4.4]. Suppose G is a d-regular graph on $2N$ vertices, paired as v^+, v^- for $v \in [N]$, and

$$\|A(G) - \frac{d}{2N} J\|_{\mathrm{op}} \le \rho \le \frac{d}{100}.$$

For a weighting $\boldsymbol{x}$ on $V(G)$, put $y_v := x_{v^+} + x_{v^-}$ and $z_v := x_{v^+} - x_{v^-}$. Then

$$3E_G(\boldsymbol{x}) - \frac{3d}{(2N)^2} \ge \frac{1}{2}\frac{d}{N}\|\boldsymbol{z}\|_2^2 - \frac{27}{20}\frac{d}{N}\|\boldsymbol{y} - \boldsymbol{u}_N\|_2^2. \tag{14.9}$$

This applies to $G = D^-_{n+1}$ with $N = N_n$, including weightings with zero coordinates.

The two corresponding tuples for every edge of D_{n+1} remain available for the next step. The projection to Q_n is a homomorphism, and Equation (14.6) is preserved. If the third-coordinate degree was constant on C_n, the pullback multiplies it by four and Equation (14.7) decreases it by d at each color. Hence this degree remains constant.

We next describe the optional deletions. On a color set C of size N, let $M \subseteq C^3$ have N tuples, with each color appearing exactly once in coordinate 3 and exactly three times in total over the three coordinates. Then

$$\Lambda_M(\boldsymbol{x}) - N^{-2} \ge -\frac{3}{N}\|\boldsymbol{x} - \boldsymbol{u}_N\|_2^2. \tag{14.10}$$

The proof of [23, Lemma 4.3] uses the scalar inequality

$$abc \ge a + b + c - 2 - (a-1)^2 - (b-1)^2 - (c-1)^2 \qquad \text{for } a, b, c \ge 0.$$

Apply it to the coordinates Nx_c of every tuple in M. Since each color occurs three times in total, summing gives

$$N^3\Lambda_M(\boldsymbol{x}) \ge N - 3\sum_{c\in C}(Nx_c - 1)^2,$$

which is Equation (14.10). Deleting M decreases every third-coordinate degree by exactly one.

We provide 3 disjoint choices at every level. For $j \in [3]$, start with

$$\{(t, t+1, t+j+1) \colon t \in \mathbb{Z}_{N_0}\},$$

and exchange the first two coordinates of each tuple when necessary to put them in increasing root order. Call the resulting set $M_{0,j}$. Its tuples have three distinct root colors and hence are allowed in Q_0. Every color occurs once in coordinate 3 and three times in total. The 3 sets are disjoint: the unordered first pair $\{t, t+1\}$ determines t, and then the third coordinate determines j.

Suppose that $M_{n,j}$ has been retained for further refinement. Its full pullback splits into four sets

$$M^{\alpha,\beta}_{n,j} := \{(a^\epsilon, b^{\alpha\epsilon}, c^{\beta\epsilon}) \colon (a,b,c) \in M_{n,j},\ \epsilon \in \{+1,-1\}\},$$

where (α, β) ranges over $\{+1, -1\}^2$. Each set has N_{n+1} tuples and again satisfies the two occurrence conditions used in Equation (14.10). Retain $M_{n,j}^{+,+}$ as $M_{n+1,j}$ and designate $M_{n,j}^{+,-}$ as an optional deletion at level $n+1$. These two sets are disjoint. The 3 designated sets, one for each j, are mutually disjoint as well. Their three root labels are distinct, so none is affected by Equation (14.7).

Choose a digit $k_{n+1} \in \{0, 1, 2, 3\}$, and delete the first k_{n+1} designated sets from $\widetilde{Q}_{n+1}$. This defines Q_{n+1}. The retained sets remain available regardless of the digit. Thus every infinite digit sequence defines a compatible palette sequence. The distinguished graph tuples and the third-coordinate regularity are preserved at each step.

We prove by induction on n that, for every weighting $\boldsymbol{x}$ on C_n and every choice of the deletion digits,

$$3\big(\Lambda_{Q_n}(\boldsymbol{u}_{N_n}) - \Lambda_{Q_n}(\boldsymbol{x})\big) \geq \frac{2d-27}{N_n}\|\boldsymbol{x} - \boldsymbol{u}_{N_n}\|_2^2. \tag{14.11}$$

The initial case follows from Equations (14.4) and (14.5), since $N_0 c_d \geq 19d/5 > 2d - 27$.

For the induction step put $N = N_n$. Given a weighting $\boldsymbol{x}$ on C_{n+1}, let $y_v = x_{v^+} + x_{v^-}$ and $z_v = x_{v^+} - x_{v^-}$. The orthogonal decomposition in Equation (10.21), applied to these paired fibers, gives

$$\|\boldsymbol{x} - \boldsymbol{u}_{2N}\|_2^2 = \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{u}_N\|_2^2 + \frac{1}{2}\|\boldsymbol{z}\|_2^2. \tag{14.12}$$

In Equation (14.8), the parent contributes the induction bound and the graph deletion contributes Equation (14.9). By Equation (14.10), the at most 3 optional deletions contribute at least $-27\|\boldsymbol{x} - \boldsymbol{u}_{2N}\|_2^2/(2N)$ after multiplication by 3. Thus

$$\begin{aligned} 3\big(\Lambda_{Q_{n+1}}(\boldsymbol{u}_{2N}) - \Lambda_{Q_{n+1}}(\boldsymbol{x})\big) &\geq \left(\frac{2d-27}{N} - \frac{27d}{20N}\right)\|\boldsymbol{y} - \boldsymbol{u}_N\|_2^2 + \frac{d}{2N}\|\boldsymbol{z}\|_2^2 - \frac{27}{2N}\|\boldsymbol{x} - \boldsymbol{u}_{2N}\|_2^2 \\ &= \frac{2d-27}{2N}\|\boldsymbol{x} - \boldsymbol{u}_{2N}\|_2^2 + \frac{3d-540}{20N}\|\boldsymbol{y} - \boldsymbol{u}_N\|_2^2 \geq \frac{2d-27}{2N}\|\boldsymbol{x} - \boldsymbol{u}_{2N}\|_2^2. \end{aligned}$$

The equality uses Equation (14.12), and the last inequality uses $d \geq 480$. This completes the induction. Since $2d - 27 > 0$, Equation (14.11) makes the uniform weighting the unique maximizer.

At the uniform maximizing weighting, all third-coordinate completion probabilities are equal. Their weighted average is $\Lambda_{Q_n}(\boldsymbol{u}_{N_n})$ by Equation (14.1). The equality criterion in Lemma 5.3, with $\mathcal{I} = \{\{3\}\}$, therefore gives

$$\kappa_3(Q_n) = \lambda(Q_n) = \Lambda_{Q_n}(\boldsymbol{u}_{N_n}). \tag{14.13}$$

Each refinement multiplies the common third-coordinate degree by four, subtracts d for the graph deletion, and subtracts k_i for the optional deletions. Hence

$$\kappa_3(Q_n) = \Lambda_{Q_0}(\boldsymbol{u}_{N_0}) - \sum_{i \in [n]} \frac{d + k_i}{N_i^2}. \tag{14.14}$$

In particular, all values are rational and nonincreasing along a branch.

Set

$$b \coloneqq \frac{1}{2} - \frac{3}{2N_0} + \frac{3+2d}{3N_0^2}, \qquad a \coloneqq b - \frac{1}{N_0^2}.$$

The choice in Equation (14.5) gives $a > 1/2 - 3/(2N_0) > 1/4$ and $b < \Lambda_{Q_0}(\boldsymbol{u}_{N_0}) < 1/2$. Since

$$\sum_{i \geq 1} \frac{d}{N_i^2} = \frac{d}{3N_0^2},$$

the limiting value in Equation (14.14) is

$$b - \frac{1}{N_0^2}\sum_{i \geq 1} k_i 4^{-i}.$$

By base-4 expansion, the series with digits $k_i \in \{0, 1, 2, 3\}$ fill $[0, 1]$. Thus the limiting directional values fill $[a, b]$.

All projections in the construction are palette homomorphisms. By Equation (14.6), the first two coordinate minima are zero at every level. Equations (14.13) and (14.14) show that uniform weights attain $\kappa_3(Q_n)$ and that these values converge to the prescribed target. This proves Proposition 14.3.

14.2 From directional densities to an interval

We now convert Proposition 14.3 into an interval for the minimum of all three coordinate degrees. One block of a new palette will control coordinates 2 and 3, while a second block will control coordinate 1. For this purpose, define

$$\mu(P) \coloneqq \max_{\boldsymbol{w} \in S(C(P))} \min\{m_2(P, \boldsymbol{w}), m_3(P, \boldsymbol{w})\}.$$

Both coordinates use the same weighting, so $\mu(P)$ need not equal $\min\{\kappa_2(P), \kappa_3(P)\}$. Attainment and monotonicity under homomorphisms follow from Lemma 5.3 with $\mathcal{I} = \{\{2\}, \{3\}\}$.

For disjoint palettes $P_1, \dots, P_m$, let a relation $R \subseteq [m]^3$ without constant tuples specify the allowed cross-block patterns. Give block i total weight p_i and normalized internal weighting $\boldsymbol{u}_i$. For $p_i > 0$, the least completion probability in coordinate j within this block is

$$\sum_{\substack{\boldsymbol{a} \in R \\ a_j = i}} \prod_{h \neq j} p_{a_h} + p_i^2 m_j(P_i, \boldsymbol{u}_i). \tag{14.15}$$

The first term counts cross-block completions and the second counts internal completions. All coordinates in one block use the same $\boldsymbol{u}_i$; zero-weight blocks are omitted from the minimum.

Let P, Q have disjoint color sets, denoted by A, B. Define $\mathsf{T}(P, Q)$ by retaining P and Q inside their respective blocks and allowing exactly the following cross-block patterns:

$$(A, A, B), \qquad (A, B, A), \qquad (A, B, B).$$

For each of these patterns, all choices of the internal colors are allowed. A homomorphism in each block induces a homomorphism of the resulting palettes.

Lemma 14.4. *Let $s \coloneqq \mu(P)$ and $t \coloneqq \kappa_1(Q)$. Then*

$$\Delta(\mathsf{T}(P, Q)) = \max\{\Delta(P), \Delta(Q), F(s, t)\}, \tag{14.16}$$

where

$$F(s, t) \coloneqq \max_{0 \le x \le 1} \min\{x - (1 - s)x^2,\ t(1 - x)^2\}. \tag{14.17}$$

The function F is continuous and nondecreasing in each variable. It is strictly increasing in each variable when both are positive and below one. Moreover,

$$F(s, 0) = 0, \qquad F(0, t) = \frac{t}{(1 + t)^2}, \qquad F\left(\frac{1}{2}, t\right) = \frac{t}{1 + 2t}.$$

Proof. Give the two blocks total masses x and $y = 1 - x$. If only one block is used, the largest possible minimum degree is the corresponding $\Delta(P)$ or $\Delta(Q)$. Suppose that both have positive mass. By Equation (14.15), the external contributions to the coordinate degrees are

block	coordinate 1	coordinate 2	coordinate 3
A	$2xy + y^2$	xy	xy
B	0	x	x

For any internal weightings, the common lower bound for coordinates $2, 3$ in A is at most $xy + x^2 s$, and the minimum in coordinate 1 in B is at most $y^2 t$. Conversely, choose internal weights attaining $\mu(P) = s$ and $\kappa_1(Q) = t$. The remaining coordinates are not smaller than the minimum of these two expressions: indeed, $xy + x^2 s \le x$ and $y^2 t \le y^2 \le 2xy + y^2$. Thus the optimum for mixed block weights is exactly Equation (14.17), proving Equation (14.16).

The objective in Equation (14.17) is jointly continuous on $[0, 1]^3$, so its maximum over $x \in [0, 1]$ is continuous in (s, t). Its two terms are nondecreasing in s and t, respectively, since their coefficients are x^2 and $(1 - x)^2$. Taking their minimum and then maximizing over x preserves this monotonicity.

For $t > 0$, the two expressions in Equation (14.17) meet first at a point in $(0, 1/2]$. Up to that point the first is increasing; from that point onward the second is decreasing and bounds the minimum. Hence the intersection gives the maximum. When $0 < s, t < 1$, increasing s moves this intersection to the left, where

the decreasing function $t(1-x)^2$ has a larger value. Increasing t moves it to the right, where $x-(1-s)x^2$ is strictly increasing. Thus F is strictly increasing in each parameter. When $s=0$, the intersection is $x=t/(1+t)$, which gives $t/(1+t)^2$. When $s=1/2$, write the first expression as $(1-(1-x)^2)/2$; the intersection gives $(1-x)^2=1/(1+2t)$ and hence $F(1/2,t)=t/(1+2t)$. □

We convert a lower bound in coordinate 3 into a common lower bound in coordinates 2 and 3. Suppose that $m_1(R,\boldsymbol{w})=m_2(R,\boldsymbol{w})=0$ for every weighting. Add one new color b, retain R on block $A=C(R)$, and allow the cross-block patterns (A,A,b) and (A,b,b). Do not allow (b,b,b). Denote the resulting palette by $\mathsf{C}(R)$.

The first-coordinate minimum of $\mathsf{C}(R)$ is always zero. If b is supported, it has no completion in that coordinate; otherwise use the first-coordinate degeneracy of R. Give A total weight x. By Equation (14.15), the minimum in coordinate 2 is $x(1-x)$, because the internal minimum of R in that coordinate is zero. The minimum in coordinate 3 is $\min\{x^2 m_3(R,\boldsymbol{w}),x\}$. Optimizing the internal weighting therefore gives

$$\mu(\mathsf{C}(R))=\max_{0\le x\le 1}\min\{x(1-x),x^2\kappa_3(R)\}=\frac{\kappa_3(R)}{(1+\kappa_3(R))^2}. \tag{14.18}$$

The maximum occurs at $x=1/(1+\kappa_3(R))$. The function on the right is strictly increasing for $0\le\kappa_3(R)<1$. Also, $R'\to R$ induces $\mathsf{C}(R')\to\mathsf{C}(R)$ by fixing b.

A second operation moves a two-coordinate interval toward $1/2$. Use $\mathsf{T}(P,P)$ on two disjoint copies. If the first-coordinate minimum of P is always zero, the same holds for $\mathsf{T}(P,P)$: in the last supported block, select an internal first-coordinate color of degree zero.

Lemma 14.5. *If $\mu(P)=s\le 1/2$, then*

$$\mu(\mathsf{T}(P,P))=f(s),\qquad f(s)\coloneqq\frac{1}{4(1-s)}. \tag{14.19}$$

If $\kappa_3(P)=t\le 1/2$, then $\kappa_3(\mathsf{T}(P,P))=f(t)$. These operations preserve palette homomorphisms.

Proof. For block masses x and $1-x$, the table in the proof of Lemma 14.4 shows that the best simultaneous bound in coordinates $2,3$ is

$$\min\{x(1-x)+sx^2,\ x+s(1-x)^2\}.$$

The second expression minus the first is $(x-s)^2+s(1-s)\ge 0$. The maximum is consequently that of $x-(1-s)x^2$, attained at $x=1/(2(1-s))$. This gives Equation (14.19). Using only coordinate 3 gives the identical calculation with s replaced by t. The one-block choices do not improve the answer, since $f(s)\ge s$ for $s\le 1/2$. Homomorphisms act independently on the two copies. □

For $s<1/2$, the map f is strictly increasing and satisfies $s<f(s)<1/2$. More precisely,

$$f\left(\frac12-\delta\right)=\frac12-\frac{\delta}{1+2\delta}\qquad\text{for }\delta>0.$$

Thus repeated applications move every compact subinterval of $(0,1/2)$ arbitrarily close to $1/2$, while preserving positive width.

Apply Equation (14.18) to the sequence in Proposition 14.3, and then apply a fixed number of self-mixing steps $P\mapsto\mathsf{T}(P,P)$. Thus there are nondegenerate intervals $[\ell,u]\subset(0,1/2)$ arbitrarily close to $1/2$ such that every $s\in[\ell,u]$ has a sequence $P_{n+1}\to P_n$ with $\Delta(P_n)=0$ and $\mu(P_n)\downarrow s$. All finite-level parameters remain below $1/2$, so Lemma 14.5 applies at every level.

To adjust the other block in Lemma 14.4, we need finite palettes with a dense set of directional values. The next lemma supplies all rational values below $1/2$, with the same weighting controlling coordinates 2 and 3.

Lemma 14.6. *For every $q\in\mathbb{Q}\cap[0,1/2)$, there is a finite palette P_q such that $m_1(P_q,\boldsymbol{w})=0$ for every weighting and $\mu(P_q)=\kappa_3(P_q)=q$. In particular, $\Delta(P_q)=0$.*

Proof. First suppose $q \in \mathbb{Q} \cap [0, 1/4]$, and write $4q = a/b$ with integers $0 \le a \le b$. Choose a directed graph D, allowing loops, so that

$$\max_{\boldsymbol{\nu} \in S(V(D))} \min_{v \in \operatorname{supp} \boldsymbol{\nu}} \sum_{u\colon uv \in E(D)} \nu_u = \frac{a}{b} = 4q.$$

For $0 < a < b$, take $V(D) = \mathbb{Z}_b$ and $uv \in E(D)$ when $v - u \in \{1, \dots, a\}$. The uniform weighting gives a/b. For the upper bound, let $s_v \coloneqq \nu_{v-1} + \cdots + \nu_{v-a}$. These quantities average to a/b. If a minimizing index has zero weight, then $s_{v+1} = s_v - \nu_{v-a} \le s_v$; move forward through zero-weight indices to reach a supported vertex. For $a = 0$ use the edgeless one-vertex graph, and for $a = b$ use the one-vertex loop.

Add a color $\star$ to $V(D)$ and define $Z(D)$ by allowing $(u, \star, v)$ for $uv \in E(D)$, $(u, v, \star)$ for $u, v \in V(D)$, and $(u, \star, \star)$ for $u \in V(D)$. If $\star$ is supported, it has zero first-coordinate degree. If it is unsupported, every allowed tuple has zero weight. Thus the first-coordinate minimum always vanishes.

Let x be the total mass of the ordinary colors and $\boldsymbol{\nu}$ their normalized weighting. If $0 < x < 1$, the minimum in coordinate 2 is $x(1-x)$: an ordinary fixed color gives this value, whereas fixing $\star$ adds the nonnegative term $x^2 \sum_{uv \in E(D)} \nu_u \nu_v$. In coordinate 3, an ordinary fixed color v gives

$$x(1-x) \sum_{u\colon uv \in E(D)} \nu_u,$$

and fixing $\star$ gives $x^2 + x(1-x) = x$. Consequently,

$$\mu(Z(D)) = \kappa_3(Z(D)) = \frac{a}{4b} = q.$$

The maximum uses $x = 1/2$. The boundary choices $x = 0, 1$ give zero and do not improve this value, so every rational $q \in [0, 1/4]$ is realized.

By Lemma 14.5, replacing P by $\mathsf{T}(P, P)$ when $\mu(P) = \kappa_3(P) = q \le 1/2$ changes both values to $f(q)$ and preserves the first-coordinate degeneracy. The iterates of f send $[0, 1/4]$ to successive adjoining intervals whose right endpoints increase to $1/2$. Explicitly,

$$f^{\circ j}(0) = \frac{j}{2(j+1)}, \qquad f^{\circ j}\left(\frac{1}{4}\right) = \frac{j+1}{2(j+2)}.$$

The inverse map on each such interval is $q \mapsto 1 - 1/(4q)$ and preserves rationality. Thus any rational $q < 1/2$ is obtained after finitely many operations from a rational in $[0, 1/4]$. This proves the lemma. □

We can now choose one block from a compatible interval and the other from a fixed finite palette. The continuous parameter gives an exact target value; the rational parameter is used only to place that target strictly between two endpoints.

Proposition 14.7. *We have* $[0, 1/4] \subseteq \mathcal{R}_{\mathrm{vp}}$.

Proof. Fix $\alpha \in (0, 1/4)$. Since $F(1/2, 1/2) = 1/4$, repeated applications of $P \mapsto \mathsf{T}(P, P)$ to the converted directional sequences give an interval $[\ell, u]$ sufficiently close to $1/2$ that $F(\ell, 1/2) > \alpha$. Choose $s_* \in (\ell, u)$. By Lemma 14.4, there is $q_* \in (0, 1/2)$ with $F(s_*, q_*) = \alpha$, and

$$F(\ell, q_*) < \alpha < F(u, q_*).$$

The continuity argument following Equation (2.4), with the dense set $\mathbb{Q} \cap (0, 1/2)$, gives a fixed rational q and $s \in (\ell, u)$ with $F(s, q) = \alpha$.

Take the corresponding projective sequence $P_{n+1} \to P_n$ for this s. Use Lemma 14.6 to obtain a fixed finite palette with $\Delta = 0$ and $\kappa_3 = q$, and exchange coordinates 1 and 3 to obtain Q_q with $\Delta(Q_q) = 0$ and $\kappa_1(Q_q) = q$. Set $T_n \coloneqq \mathsf{T}(P_n, Q_q)$. The bonding maps act on the first block and fix the second, so $T_{n+1} \to T_n$. By Equation (14.16),

$$\Delta(T_n) = F(\mu(P_n), q) \downarrow F(s, q) = \alpha.$$

Thus $\alpha \in \mathcal{R}_{\mathrm{vp}}$. The endpoints follow from Lemma 14.2. □

14.3 Cyclic substitutions and the full interval

The interval from Proposition 14.7 will now be enlarged by a finite collection of substitutions. The outer palettes are invariant under cyclic permutation of their three coordinates. This makes the external completion weight the same in all three directions, so the inner palettes enter the optimization only through their minimum degrees.

Let R be a palette on $[m]$ with no monochromatic tuple, satisfying $(i,j,k) \in R$ if and only if $(j,k,i) \in R$. For palettes $P_1, \dots, P_m$ on disjoint color sets, define $R[P_1, \dots, P_m]$ as follows. Inside one color block, use its original palette. For a tuple meeting more than one block, use the block-index tuple in R and impose no condition on the internal colors. Write $R[P]$ when every block is a disjoint copy of the same palette.

Proposition 14.8. *With $d_i := \Delta(P_i)$, we have*

$$\Delta(R[P_1, \dots, P_m]) = \max_{\boldsymbol{p} \in S([m])} \left\{ \Lambda_R(\boldsymbol{p}) + \sum_{i \in [m]} d_i p_i^3 \right\}. \tag{14.20}$$

Coordinatewise palette homomorphisms induce homomorphisms of the substituted palettes. In particular, if

$$G_R(t) := \max_{\boldsymbol{p} \in S([m])} \left\{ \Lambda_R(\boldsymbol{p}) + t \sum_{i \in [m]} p_i^3 \right\},$$

then $t \in \mathcal{R}_{\mathrm{vp}}$ implies $G_R(t) \in \mathcal{R}_{\mathrm{vp}}$. The function G_R is continuous, and for $0 \le t \le t' \le 1$,

$$\frac{t' - t}{m^2} \le G_R(t') - G_R(t) \le t' - t. \tag{14.21}$$

Proof. Give block i total mass p_i. Because R is cyclically invariant, fixing a color in block i contributes the same external completion weight in each coordinate; denote it by $E_i(\boldsymbol{p})$. For fixed block masses, each block can choose its own internal weighting. Thus the optimum is

$$\max_{\boldsymbol{p} \in S([m])} \min_{i:\ p_i > 0} \{E_i(\boldsymbol{p}) + p_i^2 d_i\}.$$

Put $F(\boldsymbol{p}) := \Lambda_R(\boldsymbol{p}) + \sum_i d_i p_i^3$. Cyclic invariance gives $\partial_i F = 3(E_i(\boldsymbol{p}) + p_i^2 d_i)$, and homogeneity gives

$$F(\boldsymbol{p}) = \sum_i p_i (E_i(\boldsymbol{p}) + p_i^2 d_i).$$

The minimum is therefore at most $\max F$. Conversely, at a maximizer of F, all partial derivatives indexed by its positive coordinates are equal. Euler's identity shows that $E_i(\boldsymbol{p}) + p_i^2 d_i = F(\boldsymbol{p})$ on that support. Choosing optimal internal weightings proves the reverse inequality in Equation (14.20).

Block-preserving maps give the homomorphism assertion. The maximum of continuous functions on the fixed simplex is continuous in t. More quantitatively, $m^{-2} \le \sum_i p_i^3 \le 1$ gives Equation (14.21) by evaluating at maximizers for t and t'. Applying the substitution to every level of a compatible sequence proves the assertion about $\mathcal{R}_{\mathrm{vp}}$. □

We next give a criterion that uses only an upper bound for the outer Lagrangian and one explicit block weighting.

Lemma 14.9. *Suppose that $[0,d] \subseteq \mathcal{R}_{\mathrm{vp}}$. Let R be cyclically invariant and have no monochromatic tuple, with $\lambda(R) \le d$. Suppose there are $e \in \mathbb{Q} \cap (d,1]$ and $\boldsymbol{p} \in S(C(R))$ such that*

$$A := \Lambda_R(\boldsymbol{p}), \qquad B := \sum_i p_i^3 < 1, \qquad A \ge (1-B)e.$$

Then $[0,e] \subseteq \mathcal{R}_{\mathrm{vp}}$.

Proof. We have $G_R(0) = \lambda(R) \le d$ and $G_R(t) \ge A + Bt$. Define $d_n := e - B^n(e-d)$. If $[0,d_n] \subseteq \mathcal{R}_{\mathrm{vp}}$, then Proposition 14.8 and continuity imply that $[G_R(0), G_R(d_n)] \subseteq \mathcal{R}_{\mathrm{vp}}$. This interval overlaps $[0,d_n]$, and

$$G_R(d_n) \ge A + Bd_n \ge d_{n+1}.$$

Starting from $d_0 = d$ proves $[0,d_n] \subseteq \mathcal{R}_{\mathrm{vp}}$ for every n. Every value smaller than e belongs to one of these intervals. The endpoint e belongs to $\mathcal{R}_{\mathrm{vp}}$ by Lemma 14.2. □

The required outer palettes have only two classes of colors. This permits all upper bounds to be reduced to one-variable cubic polynomials. Let the classes have sizes $a, b \geq 1$, and choose $\boldsymbol{k} = (k_0, k_1, k_2, k_3) \in \{0, 1, 2\}^4$. For each unordered triple of distinct colors of type AAA, AAB, ABB, BBB, respectively, allow k_0, k_1, k_2, k_3 of its two cyclic orbits. More explicitly, for colors $u < v < w$, the first orbit is $(u, v, w), (v, w, u), (w, u, v)$ and the second is $(u, w, v), (w, v, u), (v, u, w)$. No tuple with repeated colors is allowed. Denote this palette by $R(a, b; \boldsymbol{k})$.

If the first class has total mass x, averaging weights within each class maximizes every elementary symmetric polynomial that occurs. The coefficients of these polynomials are nonnegative. Hence

$$\lambda(R(a, b; \boldsymbol{k})) = \max_{0 \leq x \leq 1} f(x), \tag{14.22}$$

where

$$\begin{aligned} f(x) = {} & \frac{k_0(a-1)(a-2)}{2a^2} x^3 + \frac{3k_1(a-1)}{2a} x^2(1-x) \\ & + \frac{3k_2(b-1)}{2b} x(1-x)^2 + \frac{k_3(b-1)(b-2)}{2b^2} (1-x)^3. \end{aligned} \tag{14.23}$$

At the weighting that is uniform within each class, the diagonal sum is

$$B(x) \coloneqq \frac{x^3}{a^2} + \frac{(1-x)^3}{b^2}.$$

For the substituted palette at positive t, the weighting used in Equation (14.22) supplies the lower bound required by Lemma 14.9.

Proposition 14.10. *If* $[0, 1/4] \subseteq \mathcal{R}_{\mathrm{vp}}$, *then* $[0, 4/9] \subseteq \mathcal{R}_{\mathrm{vp}}$.

Proof. Set $d_0 \coloneqq 1/4$. For row i of Table 2, put $R_i \coloneqq R(a_i, b_i; \boldsymbol{k}_i)$, let f_i be its polynomial in Equation (14.23), and put $B_i(x) \coloneqq x^3/a_i^2 + (1-x)^3/b_i^2$. The row specifies a witness x_i and a rational endpoint d_i.

To bound f_i on $[0, 1]$, divide this interval into n_i equal subintervals and apply Equation (2.7) on each one. The largest cubic Bernstein coefficient among all these subintervals is the listed value U_i. The last column gives the exact difference $f_i(x_i) - (1 - B_i(x_i))d_i$. Thus the table verifies

$$\max_{x \in [0,1]} f_i(x) \leq U_i \leq d_{i-1}, \qquad f_i(x_i) \geq \big(1 - B_i(x_i)\big) d_i.$$

Every row has $0 < x_i < 1$ and hence $0 < B_i(x_i) < 1$. Applying Lemma 14.9 successively gives $[0, d_i] \subseteq \mathcal{R}_{\mathrm{vp}}$ for every $i \in [16]$. Since $d_{16} = 4/9$, the proposition follows. □

The verification script `parts/vertex_pair/scripts/verify_extension.py` also checks the data in Table 2, generating the cyclic tuple orbits and recomputing their polynomials and rational bounds.

Only one additional outer palette is needed after reaching $4/9$. Its relation is $\{AAB, ABA, BAA\}$, where all three tuples are retained.

Proposition 14.11. *If* $[0, 4/9] \subseteq \mathcal{R}_{\mathrm{vp}}$, *then* $\mathcal{R}_{\mathrm{vp}} = [0, 1]$.

Proof. For the two-color outer palette just described, Proposition 14.8 gives

$$G_*(t) = \max_{0 \leq x \leq 1} \{3x^2(1-x) + t(x^3 + (1-x)^3)\}.$$

We have $G_*(0) = 4/9$. If $t < 1$, choose $t < x < 1$. Then

$$3x^2(1-x) + t(x^3 + (1-x)^3) - t = 3x(1-x)(x-t) > 0.$$

Thus $G_*(t) > t$ for every $t < 1$.

Also $G_*(1) = 1$, and Equation (14.21) gives continuity and monotonicity. Apply Lemma 2.5 with $a = 0$, $b = 4/9$, $t = 1$, and $T = G_*$ to obtain $[0, 1) \subseteq \mathcal{R}_{\mathrm{vp}}$. The value 1 belongs to $\mathcal{R}_{\mathrm{vp}}$ by Lemma 14.2. □

Proof of the vertex-pair case of Theorem 1.1(i). Propositions 14.3 and 14.7 give $[0, 1/4] \subseteq \mathcal{R}_{\mathrm{vp}}$. By Propositions 14.10 and 14.11, we have $\mathcal{R}_{\mathrm{vp}} = [0, 1]$. Finally, Proposition 14.1 gives $\mathcal{R}_{\mathrm{vp}} \subseteq \Pi^{(3)}_{\mathrm{vp},\infty}$, while $\Pi^{(3)}_{\mathrm{vp},\infty} \subseteq [0, 1]$ follows from the definition. □

i	(a_i, b_i)	$\boldsymbol{k}_i$	x_i	d_i	n_i	U_i	Difference
1	$(1,3)$	$(0,0,1,2)$	$1/9$	$33/125$	4	$1/4$	$136/164025$
2	$(2,8)$	$(0,1,1,0)$	$9/20$	$67/250$	16	$8603/32768$	$25533/25600000$
3	$(2,10)$	$(0,1,1,0)$	$9/20$	$3/11$	16	$2743/10240$	$2727/2200000$
4	$(2,3)$	$(0,1,1,2)$	$1/4$	$17/60$	8	$139/512$	$3/5120$
5	$(3,4)$	$(1,1,1,0)$	$7/13$	$2869/10000$	8	$649/2304$	$17827/395460000$
6	$(3,7)$	$(0,1,1,0)$	$10/21$	$29/100$	32	$257/896$	$223231/204205050$
7	$(2,2)$	$(0,1,2,0)$	$2/5$	$309/1000$	16	$2369/8192$	$63/100000$
8	$(1,3)$	$(0,0,2,0)$	$7/20$	$159/500$	8	$115/384$	$3263/3000000$
9	$(2,3)$	$(0,2,1,0)$	$4/7$	$1/3$	8	$61/192$	0
10	$(1,3)$	$(0,0,2,1)$	$3/10$	$71/200$	16	$341/1024$	$203/900000$
11	$(2,4)$	$(0,2,1,1)$	$1/2$	$73/200$	8	$181/512$	$21/25600$
12	$(1,5)$	$(0,0,2,0)$	$4/11$	$3/8$	8	$23/64$	$21/66550$
13	$(1,3)$	$(0,0,2,2)$	$1/4$	$2/5$	4	$3/8$	0
14	$(1,4)$	$(0,0,2,1)$	$2/7$	$2077/5000$	8	$819/2048$	$381/5488000$
15	$(1,15)$	$(0,0,2,0)$	$11/30$	$271/625$	32	$17003/40960$	$68191/949218750$
16	$(3,4)$	$(1,2,1,1)$	$7/12$	$4/9$	32	$56829/131072$	$1081/2239488$

Table 2: Parameters and exact bounds for the 16 interval extensions.

15 Classification of general test systems

We now combine the constructions to prove Theorem 1.1(i). The excluded systems have the ordinary or complete-layer gaps. For the other systems, unused positions are handled by simultaneous realization, and covering systems are reduced to conditional capacities.

For $1 \le j \le r-2$ and $N = \binom{r}{j+1}$, the complete-layer results give $\Pi^{(r)}_{\binom{[r]}{j},\infty} \cap (0, N^{-N}) = \varnothing$. For $j = r-1$, the only possible density values are 0 and 1. At the highest nontrivial level, the gap follows from Reiher, Rödl, and Schacht [42, Theorems 1.2 and 6.1]; at intermediate levels it is Proposition 11.1.

Fix $\mathcal{Q} = \binom{[r]}{q}$ and a nonempty proper subset $\mathscr{B} \subsetneq \mathcal{Q}$. Define the simplicial complex

$$\Delta_{\mathscr{B}} := \{J \subseteq \mathcal{Q} \colon J \subseteq \sigma(\mathscr{B}) \text{ for some } \sigma \in S_r\}.$$

Each face of $\Delta_{\mathscr{B}}$ is a set of palette coordinates that can be prescribed together in one of the tests defining $\lambda_{\mathscr{B}}$. Every singleton is a face, while $\mathcal{Q}$ itself is not a face. A set J is a *minimal nonface* if $J \notin \Delta_{\mathscr{B}}$ and every proper subset of J belongs to $\Delta_{\mathscr{B}}$.

Restricting the palette rule to selected coordinates restricts the prescribed sets to their intersections with those coordinates.

Lemma 15.1. *Fix a nonempty set $J \subseteq \mathcal{Q}$ and a palette $R \subseteq C^J$. Let $\widetilde{R} \subseteq C^{\mathcal{Q}}$ consist of the tuples whose restriction to J belongs to R, and let $\mathcal{T}$ consist of the inclusion-maximal sets among $I \cap J$ for $I \in \mathrm{Orb}(\mathscr{B})$. For independent colors X_A with common distribution $\boldsymbol{w}$,*

$$m_{\mathscr{B}}(\widetilde{R}, \boldsymbol{w}) = \min\left\{\mathbb{P}[X_J \in R \mid X_T = \boldsymbol{a}] \colon T \in \mathcal{T},\ \boldsymbol{a} \in (\operatorname{supp} \boldsymbol{w})^T\right\}.$$

The construction preserves palette homomorphisms.

Proof. Prescribing colors on $I \setminus J$ does not affect the rule of R, so the conditional probability depends only on $I \cap J$. If $T \subseteq T'$, a probability conditional on supported colors on T is a weighted average of the probabilities obtained by prescribing the remaining colors on T'. Thus nonmaximal intersections do not lower the minimum. A color homomorphism acts on all coordinates and preserves the defining restriction. □

If J is a minimal nonface of size k, fix an ordering of its coordinates. Every $(k-1)$-subset of J is an intersection $I \cap J$, while no such intersection equals J. Hence $\mathcal{T} = \binom{J}{k-1}$, so $m_{\mathscr{B}}(\widetilde{R}, \boldsymbol{w}) = \gamma_k(R, \boldsymbol{w})$ and $\lambda_{\mathscr{B}}(\widetilde{R}) = \Gamma_k(R)$.

If all minimal nonfaces have size 2, the next lemma gives 3 q-faces no two of which belong to a common $\sigma(\mathscr{B})$.

Lemma 15.2. *Let $r \geq 3$, $1 \leq q \leq r-1$, and $\varnothing \neq \mathscr{B} \subsetneq \binom{[r]}{q}$. At least one of the following holds:*

(i) there is a three-element set $J \subseteq \binom{[r]}{q}$ with $|J \cap \sigma(\mathscr{B})| \leq 1$ for every $\sigma \in S_r$;

(ii) $\Delta_{\mathscr{B}}$ has a minimal nonface of size at least three.

Proof. Let G be the graph on $\mathcal{Q}$ in which two faces are adjacent precisely when they belong together to some $\sigma(\mathscr{B})$. Alternative (i) says $\alpha(G) \geq 3$. Suppose (ii) fails. Then all minimal nonfaces have size two, and a subset of $\mathcal{Q}$ belongs to $\Delta_{\mathscr{B}}$ if and only if it is a clique in G.

Taking complements of the q-sets is an S_r-equivariant bijection, so for this combinatorial proof we may assume $q \leq r/2$. If $q = 1$, an invariant graph is complete or empty. The complete graph would make $\mathcal{Q}$ a face of $\Delta_{\mathscr{B}}$, contrary to $\mathscr{B} \neq \mathcal{Q}$; the empty graph has an independent triple. Hence assume $q \geq 2$.

The graph G is not complete. Since adjacency depends only on intersection size, we may define $a \coloneqq \max\{|A \cap B| \colon A \neq B,\ AB \notin G\}$. If $a \geq q/2$, choose a common a-set and three pairwise disjoint $(q-a)$-sets outside it. They use $3q - 2a \leq 2q \leq r$ positions. The three unions are pairwise nonadjacent in G, giving (i). We may therefore assume $a < q/2$ and put $t \coloneqq a + 1$. Then $q \geq 2t - 1$ and $t < q$.

For a fixed t-set T, consider the complete star $\mathcal{A}_T \coloneqq \{A \in \mathcal{Q} \colon T \subseteq A\}$. It is a clique, since its pairwise intersections have size at least $t > a$. It is also maximal. To see this, take $D \notin \mathcal{A}_T$ and put $s = |D \cap T| \leq t - 1$. Choose $E \supseteq T$ by adding $t - 1 - s$ vertices from $D \setminus T$ and $q - 2t + 1 + s$ vertices outside $D \cup T$. The latter number is nonnegative because $q \geq 2t - 1$, and the required choices exist because $r \geq 2q$. Then $|D \cap E| = t - 1 = a$, so D is not adjacent to E.

Every clique belongs to $\Delta_{\mathscr{B}}$ and is consequently contained in some $\sigma(\mathscr{B})$. Maximality forces $\mathcal{A}_T = \sigma(\mathscr{B})$. Hence $\mathscr{B}$ is a complete t-star, and adjacency in G is exactly the condition $|A \cap B| \geq t$.

Choose a $(t+2)$-set W and consider the family $\{A \in \mathcal{Q} \colon |A \cap W| \geq t + 1\}$. Any two members intersect in at least t positions, so this is another clique. Its total intersection is empty: one can omit any specified position while selecting $t + 1$ positions of W and then filling to a q-set outside W; $r \geq 2q$ provides enough choices. Thus it is not contained in a complete t-star, contradicting the description of the faces of $\Delta_{\mathscr{B}}$. Alternative (i) must hold. □

We now apply these constructions to covering test systems.

Proposition 15.3. *Let $r \geq 3$, and let $\mathcal{S}$ be an antichain of nonempty proper subsets of $[r]$ with $\bigcup \mathcal{S} = [r]$. If $\mathcal{S} \neq \binom{[r]}{j}$ for every $1 \leq j < r$, then $\Pi^{(r)}_{\mathcal{S},\infty} = [0, 1]$.*

Proof. Let q be the least integer such that some q-subset of $[r]$ is contained in no member of $\mathcal{S}$. Covering gives $q \geq 2$. Put $B \coloneqq \{A \in \binom{[r]}{q} \colon A \subseteq S \text{ for some } S \in \mathcal{S}\}$. If B were empty, no member of $\mathcal{S}$ could have size at least q, while every $(q-1)$-set would have to belong to $\mathcal{S}$. The antichain would then equal the complete $(q-1)$-layer. Thus B is nonempty and proper. Also $q \leq r - 1$, since $q = r$ would force all $(r-1)$-sets to be members of $\mathcal{S}$.

Fix $\alpha \in [0, 1]$ and apply Lemma 15.2 to B. In alternative (i), let J be the independent triple. Each coordinate of J belongs to some $\sigma(B)$ because $B \neq \varnothing$ and S_r acts transitively on $\mathcal{Q}$. The maximal intersections $\sigma(B) \cap J$ are therefore exactly the 3 singletons. By Lemma 15.1, placing a ternary palette R on J gives $\lambda_B(\widetilde{R}) = \Delta(R)$. The compatible sequences from Section 14, extended in this way at every level, give a projective sequence of finite palettes $P_{n+1} \to P_n$ with $\lambda_B(P_n) \to \alpha$. In alternative (ii), choose a minimal nonface J of size $k \geq 3$. By Proposition 13.1, there is a k-coordinate chain with completion-degree limit α. Place its coordinates on J and apply Lemma 15.1 to obtain the same conditional-capacity limit and the required homomorphisms.

By the choice of q, every $(q-1)$-subset of $[r]$ is contained in a member of $\mathcal{S}$. Thus Theorem 11.3(ii) applies with this B and gives a forbidden family of density α. The trivial bounds give the reverse inclusion in $[0,1]$. □

The same chain realizes the same density for any test system whose members contain exactly the q-sets in B, provided every $(q-1)$-set is contained in a member of that system.

It remains to treat systems with unused positions. Let $\rho_\ell^{(r)}$ be the independent vertex-set parameter from Section 5. If $|\bigcup \mathcal{S}| = \ell < r$, then every monotone class satisfies

$$\gamma_\ell^{(r)} \le \pi_{\mathcal{S}} \le \rho_\ell^{(r)}.$$

The first inequality follows by summing minimum-degree bounds over the allowed roots; the test conditions depend only on their ℓ supported positions. For the second, each independent vertex-set restriction can be assigned to a containing test face. By Theorem 1.2, for $r \ge 3$ and $2 \le \ell < r$, each $\alpha \in [0,1]$ is therefore simultaneously realized by $\gamma_\ell^{(r)}$, $\rho_\ell^{(r)}$, and every $\pi_{\mathcal{S}}$ with $|\bigcup \mathcal{S}| = \ell$.

Proof of Theorem 1.1(i). Suppose that $|\bigcup \mathcal{S}| \ge 2$ and the system is not a complete layer. If $|\bigcup \mathcal{S}| < r$, the preceding simultaneous realization applies. If it covers $[r]$, apply Proposition 15.3.

For the converse, $|\bigcup \mathcal{S}| = 1$ gives the ordinary parameter by Table 1, and no ordinary Turán density lies in $(0, r!/r^r)$. Complete layers have the gaps recalled at the start of this section. These are precisely the excluded systems. □

By finite-subfamily approximation in Proposition 2.3, the densities defined by finite forbidden families are dense in $[0,1]$ for every test system under which arbitrary forbidden families realize all values in $[0,1]$.

References

[1] T. Austin and T. Tao. Testability and repair of hereditary hypergraph properties. *Random Structures Algorithms*, 36(4):373–463, 2010. doi:10.1002/rsa.20300. Preprint version: arXiv:0801.2179v2.

[2] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer, Berlin, second edition, 2006. doi:10.1007/3-540-33099-2.

[3] P. Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics. John Wiley & Sons, New York, second edition, 1999. doi:10.1002/9780470316962.

[4] W. G. Brown and M. Simonovits. Digraph extremal problems, hypergraph extremal problems, and the densities of graph structures. *Discrete Math.*, 48(2–3):147–162, 1984. doi:10.1016/0012-365X(84)90178-X.

[5] P. Erdős. On extremal problems of graphs and generalized graphs. *Israel J. Math.*, 2(3):183–190, 1964. doi:10.1007/BF02759942.

[6] P. Erdős and R. Rado. A combinatorial theorem. *J. London Math. Soc.*, 25(4):249–255, 1950. doi:10.1112/jlms/s1-25.4.249.

[7] P. Erdős and V. T. Sós. On Ramsey–Turán type theorems for hypergraphs. *Combinatorica*, 2(3):289–295, 1982. doi:10.1007/BF02579235.

[8] P. Erdős and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc.*, 52(12):1087–1091, 1946. doi:10.1090/S0002-9904-1946-08715-7.

[9] P. Frankl, Y. Peng, V. Rödl, and J. Talbot. A note on the jumping constant conjecture of Erdős. *J. Combin. Theory Ser. B*, 97(2):204–216, 2007. doi:10.1016/j.jctb.2006.05.004.

[10] P. Frankl and V. Rödl. Hypergraphs do not jump. *Combinatorica*, 4(2–3):149–159, 1984. doi:10.1007/BF02579215.

[11] P. Frankl and V. Rödl. Extremal problems on set systems. *Random Structures Algorithms*, 20(2):131–164, 2002. doi:10.1002/rsa.10017.

[12] J. Gao, O. Pikhurko, M. Rong, and S. Sun. Rational codegree Turán density of hypergraphs. *Bull. Lond. Math. Soc.*, to appear. arXiv:2601.00758.

[13] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math. (2)*, 166(3):897–946, 2007. doi:10.4007/annals.2007.166.897.

[14] C. Grosu. On the algebraic and topological structure of the set of Turán densities. *J. Combin. Theory Ser. B*, 118:137–185, 2016. doi:10.1016/j.jctb.2016.01.001. The cited Section 4 and Theorem 44 are in arXiv:1403.4653v3.

[15] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58(301):13–30, 1963. doi:10.1080/01621459.1963.10500830.

[16] G. O. H. Katona, T. Nemetz, and M. Simonovits. On a graph-problem of Turán. *Mat. Lapok*, 15:228–238, 1964. In Hungarian. Available online.

[17] D. King, S. Piga, M. Sales, and B. Schülke. On possible uniform Turán densities. Preprint, arXiv:2504.21220, 2025.

[18] A. Lamaison. Palettes determine uniform Turán density. Preprint, arXiv:2408.09643, 2024.

[19] J. Lenz and D. Mubayi. The poset of hypergraph quasirandomness. *Random Structures Algorithms*, 46(4):762–800, 2015. doi:10.1002/rsa.20521.

[20] J. Lenz and D. Mubayi. Perfect packings in quasirandom hypergraphs II. *Combin. Probab. Comput.*, 25(4):595–611, 2016. doi:10.1017/S0963548315000267.

[21] H. Li and X. Liu. On the structure of uniform Turán densities. Preprint, arXiv:2610.07049v1, 2026.

[22] H. Li, X. Liu, and D. Mubayi. The complete structure of possible Turán densities in large uniformities. Unpublished manuscript, 2026.

[23] H. Li, X. Liu, and O. Pikhurko. Intervals of uniform Turán densities. Preprint, arXiv:2608.04790v1, 2026.

[24] H. Lin, G. Sun, G. Wang, and W. Zhou. Uniform Turán densities of k-uniform hypergraphs. Preprint, arXiv:2605.15105v1, 2026.

[25] H. Lin and W. Zhou. The ℓ-degree Turán densities of hypergraphs: infinite families and their finite subfamilies. Preprint, arXiv:2610.17558, 2026.

[26] X. Liu and D. Mubayi. The number 4/9 is a non-jump for 3-graphs. Preprint, arXiv:2605.13567, 2026.

[27] X. Liu and O. Pikhurko. Intervals of hypergraph Turán densities. Preprint, arXiv:2605.25914, 2026.

[28] A. Lo and K. Markström. ℓ-degree Turán density. *SIAM J. Discrete Math.*, 28(3):1214–1225, 2014. doi:10.1137/120895974.

[29] C. McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics, 1989*, volume 141 of *London Math. Soc. Lecture Note Ser.*, pages 148–188. Cambridge Univ. Press, Cambridge, 1989. doi:10.1017/CBO9781107359949.008.

[30] T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canad. J. Math.*, 17:533–540, 1965. doi:10.4153/CJM-1965-053-6.

[31] D. Mubayi and V. Rödl. Uniform edge distribution in hypergraphs is hereditary. *Electron. J. Combin.*, 11(1):Research Paper 55, 2004. doi:10.37236/1808.

[32] D. Mubayi and Y. Zhao. Co-degree density of hypergraphs. *J. Combin. Theory Ser. A*, 114(6):1118–1132, 2007. doi:10.1016/j.jcta.2006.11.006.

[33] B. Nagle, V. Rödl, and M. Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures Algorithms*, 28(2):113–179, 2006. doi:10.1002/rsa.20117. The cited Corollary 15 is in the author version.

[34] Y. Peng. Using Lagrangians of hypergraphs to find non-jumping numbers. II. *Discrete Math.*, 307(14):1754–1766, 2007. doi:10.1016/j.disc.2006.09.024.

[35] Y. Peng. Using Lagrangians of hypergraphs to find non-jumping numbers. I. *Ann. Comb.*, 12(3):307–324, 2008. doi:10.1007/s00026-008-0353-2.

[36] Y. Peng. On jumping densities of hypergraphs. *Graphs Combin.*, 25(5):759–766, 2009. doi:10.1007/s00373-010-0874-4.

[37] Y. Peng. Hypergraph Lagrangian function. *Adv. Math. (China)*, 50(1):22–28, 2021. doi:10.11845/sxjz.2020010a.

[38] Y. Peng, V. Rödl, and J. Skokan. Counting small cliques in 3-uniform hypergraphs. *Combin. Probab. Comput.*, 14(3):371–413, 2005. doi:10.1017/S0963548304006546.

[39] Y. Peng and C. Zhao. Generating non-jumping numbers recursively. *Discrete Appl. Math.*, 156(10):1856–1864, 2008. doi:10.1016/j.dam.2007.09.003.

[40] O. Pikhurko. On possible Turán densities. *Israel J. Math.*, 201(1):415–454, 2014. doi:10.1007/s11856-014-0031-5.

[41] F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc. (2)*, 30(1):264–286, 1930. doi:10.1112/plms/s2-30.1.264.

[42] C. Reiher, V. Rödl, and M. Schacht. Hypergraphs with vanishing Turán density in uniformly dense hypergraphs. *J. Lond. Math. Soc. (2)*, 97(1):77–97, 2018. doi:10.1112/jlms.12095.

[43] C. Reiher, V. Rödl, and M. Schacht. On a generalisation of Mantel's theorem to uniformly dense hypergraphs. *Int. Math. Res. Not. IMRN*, 2018(16):4899–4941, 2018. doi:10.1093/imrn/rnx017.

[44] V. Rödl and M. Schacht. Regular partitions of hypergraphs: counting lemmas. *Combin. Probab. Comput.*, 16(6):887–901, 2007. doi:10.1017/S0963548307008565.

[45] V. Rödl and M. Schacht. Regular partitions of hypergraphs: regularity lemmas. *Combin. Probab. Comput.*, 16(6):833–885, 2007. doi:10.1017/S0963548307008553.

[46] V. Rödl and M. Schacht. Generalizations of the removal lemma. *Combinatorica*, 29(4):467–501, 2009. doi:10.1007/s00493-009-2320-x.

[47] V. Rödl and J. Skokan. Regularity lemma for k-uniform hypergraphs. *Random Structures Algorithms*, 25(1):1–42, 2004. doi:10.1002/rsa.20017.

[48] I. J. Schoenberg. Metric spaces and positive definite functions. *Trans. Amer. Math. Soc.*, 44(3):522–536, 1938. doi:10.1090/S0002-9947-1938-1501980-0.

[49] H. Towsner. σ-algebras for quasirandom hypergraphs. *Random Structures Algorithms*, 50(1):114–139, 2017. doi:10.1002/rsa.20641.

[50] P. Turán. On an extremal problem in graph theory. *Mat. Fiz. Lapok*, 48:436–452, 1941. In Hungarian.

[51] Z. Yan and Y. Peng. Non-jumping Turán densities of hypergraphs. *Discrete Math.*, 346(1):113195, 2023. doi:10.1016/j.disc.2022.113195.