Harmonic and Effective Lower Bounds for Sunflower-Free Families

Shiva Kintali*

September 6, 2026

Abstract

Let $f_r(k)$ be the maximum size of a family of distinct k-element sets containing no sunflower with r members. We give two complementary lower bounds. For each fixed $d \geq 2$ and $\delta > 0$,

$$f_r(d) \geq r^d - (H_d + \delta)r^{d-1}, \qquad H_d = \sum_{i=1}^{d} \frac{1}{i},$$

for all sufficiently large r. The coefficient H_d is optimal within the bounded-extension certificate class used in the construction. Independently, an effective construction gives

$$f_r(k) \geq \frac{1}{2}\left(r - 300\frac{\log\log r}{(\log r)^{1/4}}\right)^k \qquad \big(r \geq \exp(\exp(40)),\ k \geq 1\big).$$

Both statements include both parities of r. The constructions share a hierarchical sunflower-avoidance certificate but use different realization mechanisms. A fixed-motif decomposition argument gives the sharper fixed-rank coefficient with an unspecified threshold. A fixed-length fractional correction and one-sided geometric rounding give an explicit growing-rank threshold, at the cost of a larger average degree deficit. We also determine an exact extremum in a restricted triple-shadow class for sufficiently large even r.

1 Introduction

For integers $r \geq 3$ and $k \geq 0$, an r-*sunflower* is a collection of r distinct sets $A_1, \ldots, A_r$ with $A_i \cap A_j = C$ for all $i \neq j$. Its *core* C may be empty. Write $f_r(k)$ for the maximum size of an r-sunflower-free family of k-element sets, over all finite ground sets. Thus the forcing threshold is $f_r(k) + 1$, $f_r(0) = 1$, and $f_r(1) = r - 1$. The sunflower lemma guarantees finiteness. Throughout, $r \geq 3$ and ranks are integers; all logarithms are natural unless a base is displayed. Families are simple unless multiplicities are explicitly specified.

The sunflower conjecture asks for an upper bound C_r^k when r is fixed. This paper concerns lower bounds, with particular attention to the dependence on the petal count r. We distinguish two regimes: a fixed seed rank with a harmonic second-order estimate, sharp within a specified certificate class, and a growing seed rank with a fully specified construction threshold.

*Email: `shiva.kintali@gmail.com` Homepage: shivakintali.github.io

1.1 Motivation and the scale of the problem

The classical sunflower lemma of Erdős and Rado [9] gives

$$f_r(k) \le k!(r-1)^k.$$

They conjectured that, for every fixed r, the factorial dependence on k can be replaced by C_r^k. The conjecture is an upper-bound problem, whereas the purpose of this paper is to understand the opposite side: how large the exponential base of a sunflower-free family can be, especially when the number of petals grows.

The disjoint-support product inequality (Lemma 2.3)

$$f_r(a+b) \ge f_r(a)f_r(b)$$

and Fekete's lemma, applied to $\log f_r(k)$, imply that the extended exponential growth rate

$$\beta_r := \lim_{k\to\infty} f_r(k)^{1/k} = \sup_{k\ge 1} f_r(k)^{1/k}$$

exists, with the value $+\infty$ allowed. Thus the root of a good finite-rank seed gives a lower bound for β_r, while repeated products, supplemented by rank-one factors, give bounds at every uniformity. The elementary Erdős–Rado construction gives $\beta_r \ge r-1$. The central large-petal question considered here is whether this lower-bound base can be improved from $r-1$ to $r-o(1)$ as $r\to\infty$.

There are two reasons to keep fixed-rank and growing-rank estimates separate. At fixed rank one can seek a refined lower estimate

$$f_r(d) \ge r^d - c_d r^{d-1} + o_d(r^{d-1}),$$

where the coefficient c_d records detailed local losses. Within the product approach used here, an explicit statement about β_r with an additive loss tending to zero requires the seed rank to grow with r. Every existence threshold and rounding error must then remain uniform in that growing rank. The two main theorems address these complementary tasks.

1.2 Past work and related directions

Upper bounds. For three petals, Spencer [20] and later Kostochka [16] improved the Erdős–Rado estimate while remaining on a factorial scale; in particular, Kostochka proved an upper bound of the form

$$ck!\left(\frac{\log\log\log k}{\log\log k}\right)^k$$

for an absolute constant c and all sufficiently large k, in the form recorded by [4]. A major change came with Alweiss, Lovett, Wu and Zhang [4], who proved, in the present notation, an upper bound of the form

$$f_r(k)+1 \le \left(Cr^3 \log k \log\log k\right)^k$$

for an absolute constant C, after enlarging C to cover the finitely many small values of k. Related counting refinements appear in Frankston, Kahn, Narayanan and Park [11]; Rao's encoding argument [19] gives $f_r(k)+1 \le (Cr\log(rk))^k$. Bell, Chueluecha and Warnke [6] then sharpened the general dependence to

$$f_r(k)+1 \le (Cr\log k)^k \qquad (r,k\ge 2).$$

The Bell–Chueluecha–Warnke estimate is the strongest published general upper bound in the unrestricted uniform setting. A 2025 preprint of Fukuyama [13] claims the further estimate

$$f_r(k) + 1 \le \left(Cr^2 \frac{\log k}{\log\log k}\right)^k \qquad (r, k > 2).$$

We record this separately as a preprint claim; none of our arguments depends on it.

Lower constructions and products. The Erdős–Rado block construction gives $(r-1)^k$ sets. Abbott [1] formalized the disjoint-support product and constructed improved graph seeds. Abbott, Hanson and Sauer [3] determined the exact rank-two extremum and, by a different substitution construction, proved $f_3(k) \ge 10^{k/2-O(\log k)}$. Abbott and Exoo [2] developed further analytic and computational constructions, including the classical four- and six-petal seeds. Frankl and Wang [12] later introduced center-size-dependent petal caps and proved product and asymptotic results for the resulting asymmetric extremal function. Bennett and Priestley [7] studied a random greedy sunflower-free process on $\binom{[n]}{k}$, principally in the regime $k = n^\alpha$ with $1/2 < \alpha < 1$. Their result analyzes that process and is not asserted there to improve the strongest lower bounds for the uniformity-only function $f_r(k)$ considered here.

Fixed rank and growing petal count. Kostochka, Rödl and Talysheva [17] proved that, for every fixed d,

$$f_r(d) = (1 + o_d(1))r^d \qquad (r \to \infty).$$

Their result determines the leading term but not the coefficient of r^{d-1}. Frankl and Wang's diagonal specialization recovers this leading asymptotic while their full theorem treats unequal core-size caps [12]. Our fixed-rank result starts at the next order and therefore does not claim the leading term as new.

The recent matching-recurrence comparison. Axante, Budala, Chitic, Dumitru and Nacu [5] introduced matching-constrained recurrences which, in our notation, imply

$$\beta_r \ge \rho_r := \begin{cases} (r + \sqrt{r^2 - 2r})/2, & r \text{ even}, \\ (r + \sqrt{r^2 - 2r + 2})/2, & r \text{ odd}. \end{cases}$$

Thus $\rho_r = r - \frac{1}{2} + O(r^{-1})$, improving the graph-product benchmark in both parities. For completeness, we spell out the fixed-rank consequence against which we compare our coefficient. Write their petal parameter as r and their uniformity as d. The Chvátal–Hanson theorem [8] says that the maximum number of edges in a graph with maximum degree at most D and matching number at most q is

$$Dq + \left\lfloor \frac{D}{2} \right\rfloor \left\lfloor \frac{q}{\lceil D/2 \rceil} \right\rfloor.$$

At $(D, q) = (r-1, r-2)$ this is

$$A_r := (r-1)(r-2) + \left\lfloor \frac{r-1}{2} \right\rfloor = r^2 - \tfrac{5}{2}r + O(1).$$

For $r \ge 4$, Corollary 3.2, Lemma 3.3 and Proposition 3.4 of [5] permit the compatible initial values

$$a_3 = (r-1)A_r + \left\lfloor \frac{r-1}{2} \right\rfloor (r-1) = r^3 - 3r^2 + O(r), \qquad b_3 = a_3 + E_r = r^3 - 2r^2 + O(r).$$

Here $E_r = f_r(2)$ is given explicitly in (5); the second identity follows from their equality between the relevant intersecting three-uniform extremum and $f_r(2)$, together with their disjoint-support split. Because this choice also has $b_3 = a_3 + b_2$ with $b_2 = E_r$, their scalar recurrence, normally stated from rank five, is valid at rank four as well:

$$b_d = rb_{d-1} - \left\lceil \frac{r-1}{2} \right\rceil b_{d-2}.$$

Induction on d therefore gives the fixed-rank consequence

$$f_r(d) \geq r^d - \frac{d+1}{2} r^{d-1} + O_d(r^{d-2}) \qquad (d \geq 3).$$

Indeed, the displayed value of b_3 is the base case; at each subsequent step, multiplication of b_{d-1} by r preserves its second-order coefficient $d/2$, while the subtracted term contributes another $1/2$ because $\lceil (r-1)/2 \rceil = r/2 + O(1)$ and $b_{d-2} = r^{d-2} + O_d(r^{d-3})$. The comparison with our two main theorems is made explicit below. Their preprint also supplies new unrestricted upper bounds for triple families and several finite constructions; our paper does not contain competing general upper bounds.

Nearby formulations. The nonuniform problem in which the ground set $[n]$ is fixed has a different extremal function. Naslund and Sawin [18] proved an exponential upper bound for three-petal-free subfamilies of $2^{[n]}$ by the polynomial method. Results for that Erdős–Szemerédi capacity, as well as weak, multicolored or vector-space sunflower variants, are not directly comparable with the uniformity-only function $f_r(k)$ studied here.

1.3 Two complementary lower bounds

Write

$$H_d := \sum_{i=1}^{d} \frac{1}{i}$$

for the dth harmonic number.

Theorem 1.1 (Fixed-rank second-order bound)**.** *Fix $d \geq 2$ and $\delta > 0$. There is a threshold $R(d, \delta)$ such that, for every integer $r \geq R(d, \delta)$,*

$$f_r(d) \geq B_{d,\delta}(r) := r^d - (H_d + \delta) r^{d-1} > 0. \tag{1}$$

For every integer $k \geq 0$, write $k = ad + b$, where $a \geq 0$ and $0 \leq b < d$. Then

$$f_r(k) \geq B_{d,\delta}(r)^a (r-1)^b. \tag{2}$$

The construction for this theorem lies in a class defined by bounded one-point extension degrees. Proposition 2.2 identifies H_d as the second-order coefficient of that class's extremum. This is a restricted optimality statement, not an upper bound on unrestricted $f_r(d)$. The threshold $R(d, \delta)$ is not evaluated.

Theorem 1.2 (Explicit growing-rank bound)**.** *For every integer $r \geq \exp(\exp(40))$ and every integer $k \geq 1$,*

$$f_r(k) \geq \frac{1}{2} \left(r - 300 \frac{\log \log r}{(\log r)^{1/4}} \right)^k. \tag{3}$$

The displayed additive loss is positive, less than one throughout this range, and tends to zero as $r \to \infty$.

Theorem 1.2 uses a separate construction, not an inversion of the unspecified threshold in Theorem 1.1. It permits

$$d = \left\lfloor \sqrt{\frac{\log r}{2^{16}}} \right\rfloor, \qquad N = dr - 1 \geq 2^{50000d^2}.$$

At a layer of rank j, its average degree deficit is at most $35\sqrt{j}$. Consequently its proved fixed-rank lower bound uses the second-order coefficient

$$T_d = H_d + 35 \sum_{j=2}^{d} \frac{\sqrt{j}}{d-j+1} \leq (35\sqrt{d}+1)H_d,$$

rather than $H_d + \delta$. The two main theorems therefore retain different information: effectiveness does not recover the sharper coefficient, and the sharper fixed-rank theorem does not supply the explicit rate.

Corollary 1.3. *For every $\varepsilon > 0$ there is $R_\varepsilon > \max\{3, \varepsilon\}$ such that*

$$f_r(k) \geq \tfrac{1}{2}(r-\varepsilon)^k \qquad (r \geq R_\varepsilon,\ k \geq 0). \tag{4}$$

The threshold is independent of k.

Proof. Put $g(r) = 300(\log\log r)/(\log r)^{1/4}$. Since $g(r) \to 0$, choose $R_\varepsilon \geq \exp(\exp(40))$, enlarging it if necessary so that $R_\varepsilon > \max\{3, \varepsilon\}$ and $g(r) \leq \varepsilon$ for all $r \geq R_\varepsilon$. Theorem 1.2 then gives the claim for $k \geq 1$. For $k = 0$, its right-hand side is $1/2$, while $f_r(0) = 1$. □

1.4 Comparison with earlier bounds

The graph extremum of Abbott, Hanson and Sauer [3] is

$$E_r := f_r(2) = \begin{cases} r(r-1), & r \text{ odd}, \\ r(r-\frac{3}{2}), & r \text{ even}. \end{cases} \tag{5}$$

It is also stated in [12, Theorem 1.4], whose scalar parameter is $r-1$ in our notation. Products, as used by Abbott [1, Theorem 1], give

$$f_r(k) \geq E_r^{\lfloor k/2 \rfloor}(r-1)^{k \bmod 2}.$$

The resulting bases are

$$\sqrt{E_r} = \begin{cases} r - \frac{1}{2} + O(r^{-1}), & r \text{ odd}, \\ r - \frac{3}{4} + O(r^{-1}), & r \text{ even}. \end{cases}$$

In contrast, (3) gives an explicit base $r - o(1)$ for both parities, with a prefactor independent of k. The fixed additive losses in the graph bases are thus replaced, for sufficiently large r, by a loss tending to zero.

Against the stronger recent benchmark ρ_r from [5], the $r - o(1)$ base in (3) is eventually larger. This does not hold throughout the whole displayed cutoff range: at $\log\log r = 40$ our loss is approximately 0.544799, whereas

$$r - \rho_r = \begin{cases} \frac{1}{2} + \frac{1}{4r} + O(r^{-2}), & r \text{ even}, \\ \frac{1}{2} - \frac{1}{4r} + O(r^{-2}), & r \text{ odd}. \end{cases}$$

The simpler condition $r \geq \exp(\exp(41))$ is sufficient for our displayed base to exceed ρ_r. Indeed, the displayed loss is then less than 0.435: as a function of $t = \log\log r$, it is $300te^{-t/4}$, whose derivative is negative for $t > 4$, and its value at $t = 41$ is less than 0.435. Rationalizing the two definitions gives

$$r - \rho_r = \begin{cases} \dfrac{r}{r + \sqrt{r^2 - 2r}}, & r \text{ even}, \\ \dfrac{r-1}{r + \sqrt{r^2 - 2r + 2}}, & r \text{ odd}, \end{cases}$$

and hence $r - \rho_r > (r-1)/(2r) > 0.49$ in the stated range. At fixed rank, $H_d < (d+1)/2$ for $d \geq 3$: this starts from $H_3 = 11/6 < 2$ and follows by induction because $1/(d+1) < 1/2$. Therefore Theorem 1.1, with $0 < \delta < (d+1)/2 - H_d$, improves the direct fixed-rank consequence of their recurrence for all sufficiently large r. At finite parameters, [5] reports $f_4(3) \geq 39$ and $f_6(3) \geq 153$: our supplementary count 39 matches the former, whereas its count 152 is below the latter.

Already $d = 5$ and $\delta = 1/60$ in Theorem 1.1 yield

$$f_r(5) \geq r^5 - \frac{23}{10}r^4 \qquad (r \text{ sufficiently large}). \tag{6}$$

Here $H_5 = 137/60$, and the fifth root is $r - 23/50 + O(r^{-1})$. Its limiting advantages over the odd- and even-petal graph bases are respectively $1/25$ and $29/100$.

For fixed d, the leading asymptotic $f_r(d) = (1 + o_d(1))r^d$ was already established by Kostochka, Rödl and Talysheva [17]; see also [12, Theorem 1.6]. Theorem 1.1 concerns the lower-order term, not the leading term itself. Neither the base comparison nor the uniform factor 1/2 asserts pointwise improvement over every existing construction at every rank. For example, for odd r the graph product remains stronger than (1) at ranks two and four. The explicit threshold in Theorem 1.2 is enormous, and that theorem makes no assertion for a specified small petal count. These comparisons do not establish priority over every earlier construction; [15] surveys the construction literature.

1.5 What is novel and what is derived

What is novel. The first new assertion is the harmonic second-order construction in Theorem 1.1. The fixed-rank asymptotics cited above determine only the leading term [17, 12], while the direct graph and matching-recurrence benchmarks have second-order losses of order d. Our coefficient is $H_d \sim \log d$. The accompanying capacity calculation proves that H_d is optimal within the bounded-extension certificate class used here. This is a structural optimality statement about the method, not a determination of the unrestricted second-order term of $f_r(d)$.

The second new assertion is the effective, all-rank estimate in Theorem 1.2. Its content beyond the fixed-rank theorem is quantitative: the seed rank grows explicitly with r, the construction threshold is fully checked, and the final prefactor is uniform in k. The proof introduces a fixed 80-term fractional correction coupled to a one-sided vector-balancing and repair argument. This permits an average layer deficit of at most $35\sqrt{j}$ while retaining the maximum-deficit control needed for the downward hierarchy.

The third new assertion is Theorem 5.1: for sufficiently large even r, the value

$$U(r) = \left\lfloor r^3 - \frac{11}{6}r^2 + \frac{r}{3} + \frac{1}{3} \right\rfloor$$

is attained and is exact within the stated triple-shadow class. The restriction to that class is essential; this theorem does not give an unrestricted upper bound on $f_r(3)$.

What is derived. Several displayed conclusions are consequences of these new seeds rather than separate construction theorems. Equation (2) follows from Theorem 1.1 by Abbott's standard disjoint-support product [1], with rank-one factors supplying the remainder. Corollary 1.3 is obtained from the effective theorem by observing that its additive loss tends to zero. The rank-five estimate (6), the corresponding root expansions, and all comparisons between exponential bases are algebraic consequences of the stated bounds.

The qualitative conclusion $\beta_r \geq r - o(1)$ can also be derived non-effectively from Theorem 1.1. For fixed d and δ,

$$B_{d,\delta}(r)^{1/d} = r - \frac{H_d + \delta}{d} + O_{d,\delta}(r^{-1}).$$

Given $\varepsilon > 0$, first choose d and then $\delta > 0$ so that $(H_d + \delta)/d < \varepsilon$; this is possible because $H_d/d \to 0$. For all sufficiently large r, the theorem and the product inequality give $\beta_r \geq B_{d,\delta}(r)^{1/d} > r - \varepsilon$. Theorem 1.2 is not presented as a second proof of only that qualitative fact; its additional contribution is the explicit rate, explicit cutoff and uniform pointwise statement for every k. Finally, the finite certificates in the supplement instantiate the certificate mechanism at selected small parameters. They support reproducibility but are not used to prove either asymptotic theorem. >

1.6 Overview of the proof strategy

The common certificate. We work with a downward-closed complex $\mathcal{K}$ on $N = dr - 1$ vertices. For every i-face e with $1 \leq i < d$ we enforce the strict extension budget

$$|N_{\mathcal{K}}(e)| \leq (d - i)r - 1.$$

If r top-dimensional faces formed a sunflower with core e, their petals would require too many vertices outside e; the budgets also handle nonempty intermediate cores, while $N = dr - 1$ excludes an empty-core sunflower. This reduces the construction problem to building a large complex whose successive degrees nearly saturate all budgets. Multiplying the incidence inequalities produces the certificate-class upper bound

$$\frac{1}{d!} \prod_{i=1}^{d} (ir - 1) = r^d - H_d r^{d-1} + O_d(r^{d-2}),$$

which explains the harmonic coefficient before either realization argument begins.

The fixed-rank realization. For fixed d, we choose a density sequence so that independent masks produce a complex with the required positive rooted extension counts. At each layer, an almost-perfect matching theorem of Frankl and Rödl [10] removes most of the deficit. The remaining divisible demand is encoded as a labelled fixed-family decomposition problem. Lattice membership, positive fractional decomposability and labelled extendability allow us to invoke the decomposition framework of Keevash [14]. A bounded residual correction completes the layer without destroying typicality. Iterating over the finitely many layers makes the mean deficits arbitrarily small and yields the coefficient $H_d + \delta$.

The effective growing-rank realization. The preceding decomposition proof has thresholds depending ineffectively on the fixed rank, so it cannot simply be run with $d = d(r)$. The effective branch therefore keeps only a closed collection of simplex and double-fan statistics. These counts give quantitative control of the relevant incidence Gram matrix. An explicit preconditioner and a fixed 80-term correction move a buffered fractional point close to the desired degree vector. We then use Banaszczyk's vector-balancing theorem in the Gram–Schmidt-walk formulation [21], followed by controlled deletion and repair inside the selected host. The resulting layer has average deficit at most $35\sqrt{j}$ and a separately bounded maximum deficit. Those two estimates propagate different requirements: the average controls the number of top faces, while the maximum preserves the fan counts required at later layers. Choosing

$$d = \left\lfloor \sqrt{\frac{\log r}{2^{16}}} \right\rfloor$$

and verifying $N \geq 2^{50000d^2}$ then gives the explicit base in Theorem 1.2; the common product argument extends the seed to every rank k.

The exact triple-shadow construction. For even r, the final branch constructs a pair-shadow with a prescribed nearly regular degree sequence, corrects its bounded parity leave, and realizes the remaining triangle demands through the same lattice and decomposition technology used in the fixed-rank branch. A redistribution step removes duplicate labelled demands. The matching upper bound inside the restricted class follows by combining the pair-codegree capacities with the unavoidable shadow-degree parity defect.

1.7 Organization and parameter conventions

Section 2 develops the common certificate, density sequence, and product and deletion principles. Section 3 proves Theorem 1.1, using rooted typicality, almost-perfect matchings and a labelled fixed-family decomposition theorem. Section 4 proves Theorem 1.2, using quantitative incidence-matrix estimates, a fixed-length correction and one-sided vector balancing. Section 5 states the exact even-petal triple result, separates the scopes of the two main constructions, and describes the finite certificates. Section A proves the triple-shadow theorem. Finite certificates and their verification data are supplied separately.

The hypotheses of the two branches are deliberately distinguished. The decomposition branch requires positive-face typicality for all patterns up to a fixed order, sufficient for labelled extendability. The effective branch tracks only specified simplex and double-fan patterns and their rank truncations. The latter condition is not used as a substitute for the former.

In Section 3 and Appendix A, the seed rank and auxiliary pattern orders and accuracies are fixed before taking a limit. The petal count and prescribed degrees grow with the ground-set size; constants implicit in $O_d(\cdot)$ may depend on d. In Section 4, the global rank d is allowed to grow and all required parameter inequalities are explicit. The notation $x = (1 \pm t)y$ means $(1-t)y \leq x \leq (1+t)y$.

2 Common framework

A *complex* on X is a family of subsets of X containing the empty set and every singleton and closed under taking subsets. Its rank is the maximum size of a face. Write $\mathcal{K}_i$ for the set of its i-faces. The rank need not be attained at every vertex. For $e \in \mathcal{K}$ define

$$N_{\mathcal{K}}(e) = \{x \in X \setminus e : e \cup \{x\} \in \mathcal{K}\}.$$

For a j-uniform family $\mathcal{A}$ and a $(j-1)$-set e, write $\deg_{\mathcal{A}}(e) = |\{T \in \mathcal{A} : e \subset T\}|$. Incidences in a multiset are counted with their multiplicities.

Lemma 2.1 (Sunflower certificate)**.** *Let $\mathcal{K}$ have rank at most d and $|X| = dr - 1$. If*

$$|N_{\mathcal{K}}(e)| \le (d-i)r - 1 \qquad (e \in \mathcal{K}_i,\ 1 \le i < d), \tag{7}$$

then $\mathcal{K}_d$ is r-sunflower-free.

Proof. Suppose $A_1, \dots, A_r \in \mathcal{K}_d$ form a sunflower with core e, and put $i = |e|$. If $i = 0$, their union has dr vertices, which exceeds $|X|$. If $i = d$, all the A_ℓ equal e, contrary to distinctness. For $1 \le i < d$, downward closure gives $e \in \mathcal{K}_i$ and $e \cup \{x\} \in \mathcal{K}$ for every $x \in \bigcup_\ell (A_\ell \setminus e)$. The latter union has $r(d-i)$ distinct vertices, contradicting (7). □

Define the certificate-class extremum explicitly by

$$C(r,d) := \max\{|\mathcal{K}_d| :\ |X| = dr-1,\ \mathcal{K} \text{ is a rank-at-most-}d \text{ complex on } X \text{ satisfying (7)}\},$$

and put

$$P_d(r) = \frac{1}{d!}\prod_{i=1}^{d}(ir-1).$$

The value is invariant under relabelling of X, and the maximum exists because there are only finitely many complexes on any fixed $(dr-1)$-element ground set. Allowing at most $dr-1$ vertices gives the same extremum: isolated vertices can be added to reach exactly $dr-1$ vertices without changing the top layer or any nontrivial one-point extension degree.

Proposition 2.2 (Certificate-class capacity)**.** *For all integers $r \ge 3$ and $d \ge 2$, $C(r,d) \le P_d(r)$. Moreover, Theorem 1.1 is realized within this class, and consequently*

$$C(r,d) = r^d - H_d r^{d-1} + o_d(r^{d-1}) \qquad (r \to \infty,\ d \text{ fixed}).$$

Proof. Set $M_i = |\mathcal{K}_i|$. Counting pairs (e,T) with $e \in \mathcal{K}_{i-1}$, $T \in \mathcal{K}_i$, and $e \subset T$ gives

$$iM_i = \sum_{e \in \mathcal{K}_{i-1}} |N_{\mathcal{K}}(e)| \le ((d-i+1)r-1)M_{i-1} \quad (2 \le i \le d).$$

Since $M_1 = dr-1$, multiplication gives the asserted upper bound. Indeed, the factors encountered are $((d-i+1)r-1)/i$ for $2 \le i \le d$, together with $M_1 = dr-1$, and their product is $P_d(r)$. Expanding the fixed-degree polynomial gives

$$P_d(r) = r^d \prod_{i=1}^{d}\left(1 - \frac{1}{ir}\right) = r^d - H_d r^{d-1} + O_d(r^{d-2}).$$

The construction in Section 3.4 gives, for every $\delta > 0$ and all sufficiently large r, a member of the same class with at least $r^d - (H_d+\delta)r^{d-1}$ top faces. Hence

$$\limsup_{r\to\infty} \frac{C(r,d) - (r^d - H_d r^{d-1})}{r^{d-1}} \le 0$$

by the upper bound, while for every $\delta > 0$,

$$\liminf_{r\to\infty} \frac{C(r,d) - (r^d - H_d r^{d-1})}{r^{d-1}} \ge -\delta.$$

Letting $\delta \downarrow 0$ proves the stated asymptotic equality. □

This proposition is not an upper bound on $f_r(d)$. Exact attainment of $P_d(r)$ can fail even for elementary divisibility reasons: when $d = 2$ and r is even, $P_2(r) = (2r-1)(r-1)/2$ is not an integer. The fixed-rank construction instead permits the pointwise deficit in (12) to be bounded and the average deficit in (13) to be arbitrarily small. The pointwise bound preserves rooted extension counts; the average bound determines the second-order coefficient.

2.1 Uniform products

Lemma 2.3 (Uniform products; Abbott [1])**.** *Let $s \geq 1$, let $X_1, \dots, X_s$ be pairwise disjoint, and let $\mathcal{F}_i \subseteq \binom{X_i}{k_i}$ be r-sunflower-free. Define their disjoint-support product by*

$$\mathcal{F}_1 * \cdots * \mathcal{F}_s := \{A_1 \cup \cdots \cup A_s : A_i \in \mathcal{F}_i \text{ for every } i\}.$$

This product is an r-sunflower-free $(\sum_{i=1}^{s} k_i)$-uniform family of cardinality $\prod_{i=1}^{s} |\mathcal{F}_i|$.

Proof. Disjointness of the X_i makes the representation of a product set unique, proving the uniformity and cardinality assertions. Suppose that r distinct product sets form a sunflower. Fix a block X_i and project the r sets onto it. The projected pairwise intersections are all equal. If two projections coincide, say in A, then their intersection is A; hence every projected set contains A. All projections have size k_i, so they must all equal A. If no two projections coincide, the r distinct projections form a forbidden sunflower in $\mathcal{F}_i$. The latter alternative is impossible, so the projections are constant on every block. The original product sets are therefore identical, a contradiction. □

We use the conventions that a rank-zero factor is $\{\varnothing\}$ and that the product of no factors is $\{\varnothing\}$. Consequently, for all nonnegative integers a, b, k,

$$f_r(a+b) \geq f_r(a) f_r(b) \qquad \text{and} \qquad f_r(k) \geq (r-1)^k.$$

The second inequality applies the product lemma to k disjoint copies of the $(r-1)$-element rank-one family; for $k = 0$ it uses the empty-product convention.

2.2 The density sequence

Fix an integer $d \geq 2$. Define $p_2, \dots, p_d$ recursively by

$$\prod_{t=2}^{j} p_t^{\binom{j-1}{t-1}} = \frac{d-j+1}{d} \qquad (2 \leq j \leq d). \tag{8}$$

At stage j, the exponent of p_j is one and all previously defined factors are positive; thus these equations determine a unique positive sequence.

Lemma 2.4 (Density bounds and the finite shift)**.** *For every $2 \leq j \leq d$, these numbers satisfy*

$$\frac{1}{d} \leq p_j \leq 1 - \frac{1}{2d^d}.$$

With the empty product interpreted as $Q_2 = 1$, write

$$Q_j = \prod_{t=2}^{j-1} p_t^{\binom{j-1}{t-1}}.$$

Then $1/d \leq Q_j \leq 1$. If $N = dr - 1$ and $\lambda_j = (d-j+1)r - 1$, then

$$\frac{\lambda_j}{Q_j N} = p_j - \frac{j-1}{d Q_j N}. \tag{9}$$

Proof. Let $a_m = \log(1 - m/d)$ for $0 \le m \le d-1$. Binomial inversion in (8) gives, for $n = j-1$,

$$a_n = \sum_{\ell=1}^{n} \binom{n}{\ell} \log p_{\ell+1}, \qquad a_0 = 0,$$

and therefore

$$\log p_j = \sum_{m=0}^{j-1} (-1)^{j-1-m} \binom{j-1}{m} a_m = \Delta^{j-1} a_0.$$

For the smooth function $a(x) = \log(1 - x/d)$,

$$a^{(j-1)}(x) = -\frac{(j-2)!}{(d-x)^{j-1}}.$$

The repeated fundamental theorem of calculus consequently gives

$$-\log p_j = (j-2)! \int_{[0,1]^{j-1}} \frac{d\mathbf{x}}{(d - x_1 - \cdots - x_{j-1})^{j-1}}.$$

Throughout the integration cube, $0 \le x_1 + \cdots + x_{j-1} \le j-1 \le d-1$, so the integrand is finite and positive. Hence $0 < p_j < 1$. Moreover, (8) reads

$$Q_j p_j = \frac{d-j+1}{d}.$$

Since all factors defining Q_j lie in $(0,1)$, we have $Q_j \le 1$ and thus $p_j \ge (d-j+1)/d \ge 1/d$; since also $p_j \le 1$, we have $Q_j \ge (d-j+1)/d \ge 1/d$.

The integral is at least $(j-2)!/d^{j-1} \ge d^{-d}$: the first inequality uses $d - x_1 - \cdots - x_{j-1} \le d$, and the second uses $j \le d$. It follows that

$$p_j \le e^{-d^{-d}} \le 1 - \frac{1}{2d^d},$$

where the last inequality follows from $e^{-x} \le 1 - x + x^2/2 \le 1 - x/2$ for $0 \le x \le 1$. Finally,

$$\lambda_j - Q_j N p_j = ((d-j+1)r - 1) - \frac{d-j+1}{d}(dr - 1) = -\frac{j-1}{d}.$$

Division by the positive number $Q_j N$ proves (9). $\square$

For example, at terminal rank $d = 5$,

$$p_2 = \frac{4}{5}, \qquad p_3 = \frac{15}{16}, \qquad p_4 = \frac{128}{135}, \qquad p_5 = \frac{3645}{4096}.$$

The displayed bounds hold simultaneously for every $2 \le j \le d$, with their dependence on d explicit. The decomposition branch needs only positivity at each fixed rank.

2.3 Rooted patterns and deterministic deletion

A j-set is *supported* by a lower complex if all of its proper subsets are present. At a rank-j incidence step, a *facet* means a $(j-1)$-face; it need not be a maximal face of the whole complex. A hypergraph is *linear* if distinct edges meet in at most one vertex.

A *positive rooted pattern* is a complex $\mathcal{P}$ on a finite vertex set W, together with a distinguished vertex $u \in W$. If $\mathcal{K}$ is a host complex on X, a *root embedding* is an injective map $\phi : W \setminus \{u\} \to X$ that sends every face of $\mathcal{P}$ not containing u to a face of $\mathcal{K}$. An *extension* of ϕ is a choice $x \in X \setminus \phi(W \setminus \{u\})$ such that the map ϕ_x, obtained by setting $\phi_x(u) = x$, sends every face of $\mathcal{P}$ to a face of $\mathcal{K}$. No nonface conditions are imposed. In particular, x must avoid the images of isolated root vertices as well. This convention is used in both branches.

Lemma 2.5 (Rooted deletion bound)**.** *Let $j \geq 2$. Let $\mathcal{K}$ be a complex of rank at most j, let $\mathcal{D} \subseteq \mathcal{K}_j$, and put $\mathcal{K}^- := \mathcal{K} \setminus \mathcal{D}$. Suppose that at most L members of $\mathcal{D}$ contain any one $(j-1)$-face. Fix a positive rooted pattern $(\mathcal{P}, u)$ of rank at most j and a root embedding ϕ that is valid in $\mathcal{K}^-$. If b_j is the number of j-faces of $\mathcal{P}$ containing u, then at most $b_j L$ extensions of ϕ in $\mathcal{K}$ fail to be extensions in $\mathcal{K}^-$.*

Proof. Because only j-faces are deleted, $\mathcal{K}^-$ is again a complex, and a root embedding valid in $\mathcal{K}^-$ is also valid in $\mathcal{K}$. Consider an extension that is valid in $\mathcal{K}$ but not in $\mathcal{K}^-$. A required face not containing the distinguished vertex belongs entirely to the root and remains valid by hypothesis. A required face of size less than j was not deleted. Thus the failed condition is a deleted j-face corresponding to some j-face of the pattern that contains the distinguished vertex. For each such pattern face, the images of its other $j-1$ vertices form a fixed $(j-1)$-face e of the host. At most L deleted j-faces contain e, and each determines at most one value of the extension vertex x. A union bound over the b_j possible pattern faces gives at most $b_j L$ destroyed extensions. No randomness of the deleted family is assumed. □

The distinction between average and maximum deletion is fundamental. Average degree loss controls the number of retained faces, whereas a maximum bound at each facet controls the rooted extension counts needed to construct the next layer. Neither kind of control is silently substituted for the other.

3 Sharp fixed-rank construction

In this section the seed rank and all local orders and accuracies are fixed before the ground-set size tends to infinity. The full positive rooted-pattern condition below is used to establish the extension hypotheses of the decomposition theorem. It is stronger than the restricted fan conditions used in the effective construction. Whenever $j = 2$ below, a displayed list $p_2, \ldots, p_{j-1}$ is empty and every product over that list is interpreted as one.

3.1 Rooted extension counts and random masks

Definition 3.1. Fix integers $j \geq 1$, $h \geq 1$, a number $\tau > 0$, and densities $p_2, \ldots, p_j \in (0, 1]$. A rank-at-most-$j$ complex $\mathcal{K}$ on a ground set X, $|X| = N$, is $(\tau, h; \mathbf{p})$-*typical* if the following condition holds. Let $\mathcal{P}$ be any complex of rank at most j on a vertex set W with $|W| \leq h$, choose $u \in W$, and fix an injective map $\phi : W \setminus \{u\} \to X$ such that $\phi(e) \in \mathcal{K}$ for every $e \in \mathcal{P}$ with $u \notin e$. For

$$b_t(\mathcal{P}, u) = |\{e \in \mathcal{P} : |e| = t,\ u \in e\}|,$$

the number of $x \in X \setminus \phi(W \setminus \{u\})$ for which $\phi \cup \{u \mapsto x\}$ preserves every face of $\mathcal{P}$ is

$$(1 \pm \tau) N \prod_{t=2}^{j} p_t^{b_t(\mathcal{P},u)}. \tag{10}$$

An empty product equals one. A root is required to preserve faces, but need not preserve nonfaces.

All pattern vertices, including isolated root vertices, count toward h. Thus (10) includes the requirement that the new vertex avoid the roots. At rank one its count is exactly $N-|W|+1$, so the complete singleton complex is (τ, h)-typical whenever N is sufficiently large. Successive use of (10) counts embeddings with several unembedded vertices. For the complete simplex on $1 \le i \le j$ vertices, each face of size t is counted when its last vertex is embedded. Multiplying the i extension estimates and dividing the number of ordered embeddings by $i!$ gives, if $h \ge j$ and $0<\tau<1$,

$$\frac{(1-\tau)^i N^i}{i!}\prod_{t=2}^{i} p_t^{\binom{i}{t}} \le |\mathcal{K}_i| \le \frac{(1+\tau)^i N^i}{i!}\prod_{t=2}^{i} p_t^{\binom{i}{t}}.$$

In particular, for $\tau \le 1/2$ these levels are nonempty and have size $\Theta(N^i)$, with constants depending only on the fixed densities and rank.

Lemma 3.2 (Mask and deletion estimates)**.** *Fix $j \ge 2$ and an order $h \ge j$. Fix lower densities $p_2, \dots, p_{j-1} \in (0,1)$ and a compact interval $I \subset (0,1)$. For every $\tau > 0$ there is $\tau_0 > 0$ with the following property. Let $\mathcal{K}$ be a rank-at-most-$(j-1)$ complex on X, $|X| = N$, which is $(\tau_0, h; \mathbf{p})$-typical. Select each j-set supported by $\mathcal{K}$ independently with probability $\theta \in I$, and call the selected family $\mathcal{A}$. With probability tending to one as $N \to \infty$, the complex $\mathcal{L} := \mathcal{K} \cup \mathcal{A}$ is $(\tau, h; p_2, \dots, p_{j-1}, \theta)$-typical. The conclusion is uniform for $\theta \in I$.*

If $\mathcal{F} \subseteq \mathcal{A}$ and

$$\deg_{\mathcal{A}\setminus\mathcal{F}}(e) \le \rho N \quad \textit{for all lower } (j-1)\textit{-faces } e,$$

then deleting $\mathcal{A} \setminus \mathcal{F}$ destroys at most $\binom{h-1}{j-1}\rho N$ extensions of any positive rooted pattern of order at most h whose root embedding is valid in $\mathcal{K} \cup \mathcal{F}$.

Proof. Fix a positive rooted pattern $(\mathcal{P}, u)$ and a root embedding that preserves all faces of rank at most $j-1$. Condition on the indicators of the required root-only j-faces. If any such face is absent, this root is not valid in $\mathcal{L}$ and need not be tested. Otherwise truncate $\mathcal{P}$ to rank $j-1$. Incoming typicality gives

$$S = (1 \pm \tau_0) N \prod_{t=2}^{j-1} p_t^{b_t}$$

candidate vertices satisfying its lower faces. For a candidate x, exactly $b_j = b_j(\mathcal{P}, u)$ distinct new indicators must equal one. For two different candidates these sets of indicators are disjoint: every relevant j-set contains its candidate and otherwise only root vertices. They are also disjoint from the conditioned root-only indicators. Hence, conditionally, the extension count has binomial distribution with parameters S and θ^{b_j}. Its expectation is bounded below by a fixed positive multiple of N. For any fixed relative error $\zeta > 0$, Chernoff's inequality bounds its failure probability by $2\exp(-c_\zeta N)$, uniformly over the roots and $\theta \in I$; when $b_j = 0$, the count is the deterministic value S. Choose τ_0, ζ so that $(1 \pm \tau_0)(1 \pm \zeta)$ is contained in $(1 \pm \tau)$. There are a fixed number of pattern types and $O_h(N^{h-1})$ root maps. The union bound proves the first assertion, including its conditional-root interpretation. All constants in the Chernoff bound depend on I only through $\min I > 0$; consequently the same argument proves that the supremum, over $\theta \in I$, of the failure probability tends to zero.

For the second assertion, apply Lemma 2.5 with $L = \rho N$. A pattern on at most h vertices has at most $\binom{h-1}{j-1}$ required new j-faces, giving the displayed bound for every root valid after the deletion. ∎

Because every density is bounded below and every order is fixed, the additive deletion error becomes an arbitrarily small relative error when ρ is sufficiently small. This estimate is deterministic and makes no assumption about how $\mathcal{F}$ is chosen.

3.2 Decomposition inputs and their framework

We state the particular external inputs and the definitions needed to apply them. Labelled-complex terminology below follows [14, Sections 2 and 7]; the ordinary rooted typicality of Definition 3.1 is our explicit sufficient condition.

3.2.1 Almost-perfect matchings

Theorem 3.3 (Frankl–Rödl, matching consequence)**.** *Fix $t \geq 2$, $a > 3$, and $\eta > 0$. There are $\gamma > 0$ and M_0 such that any t-uniform hypergraph on $M \geq M_0$ vertices satisfying*

$$(1-\gamma)D < \deg(v) < (1+\gamma)D, \qquad \deg(u,v) < D/(\log M)^a \quad (u \neq v)$$

has a matching covering at least $(1-\eta)M$ vertices.

This is the consequence of [10, Theorem 1.1 and the following remark] that we use. For clarity, its conversion from a cover is elementary. Apply their covering theorem with error $\zeta = \eta/t$. If a cover has $m \leq (1+\zeta)M/t$ edges, its redundant incidence $tm - M$ is at most ζM. While two edges intersect, delete one of them. More explicitly, for a collection $\mathcal{E}$ put $R(\mathcal{E}) = t|\mathcal{E}| - |\bigcup_{E\in\mathcal{E}} E|$. If a deleted edge meets another remaining edge, then $R(\mathcal{E})$ drops by at least one; it is always nonnegative. Thus at most ζM deletions leave a matching. Since every cover has at least M/t edges, this matching has at least $M/t - \zeta M$ edges and therefore covers at least $M - t\zeta M = (1-\eta)M$ vertices.

3.2.2 Labelled complexes and extension requirements

For a finite role set R, an *R-complex* Φ consists of sets $\Phi_B \subseteq \mathrm{Inj}(B,V)$ for $B \subseteq R$, closed under restriction of maps. Its vertex set $V(\Phi)$ is the union of the images of its singleton maps. Write Φ_t° for the unordered images of its maps with t-element domains.

Let Σ be a permutation group on R. The complex is *Σ-adapted* if precomposition by restrictions of elements of Σ preserves it. It is *exactly Σ-adapted* if, for every $\phi \in \Phi_B$ and every bijection $\alpha : B' \to B$,

$$\phi\alpha \in \Phi_{B'} \quad \Longleftrightarrow \quad \alpha = \sigma|_{B'} \text{ for some } \sigma \in \Sigma.$$

We call Φ *exactly adapted* if it is exactly Σ-adapted for some permutation group Σ on R.

Here is the full extension condition used below. For a positive integer s, let $R(s)$ be the R-complex of maps ψ into $R \times [s]$ with $\psi(i) = (i,x)$ on their domains. Choose an R-subcomplex $\mathcal{P} \subseteq R(s)$, a root set $F \subseteq V(\mathcal{P})$, and an injective map $\phi : F \to V(\Phi)$ which sends every root-only map of $\mathcal{P}$ to a map of Φ. An extension is an injection $\phi^+ : V(\mathcal{P}) \to V(\Phi)$ extending ϕ and satisfying $\phi^+\psi \in \Phi$ for every $\psi \in \mathcal{P}$. Put $v = |V(\mathcal{P}) \setminus F|$ and $n = |V(\Phi)|$. We call Φ *(ω,s)-extendable* if every such rooted problem has at least ωn^v solutions. There are at most $|R|s$ pattern vertices, not merely s.

For an unordered support $G \subseteq \Phi_t^\circ$, the condition that (Φ, G) is (ω,s)-extendable additionally permits any specified collection of t-domain maps of $\mathcal{P}$ not wholly in F to be required to have image in G. The same lower bound must hold for every such collection. Root-only maps are deliberately excluded from these additional support requirements. For an integer host vector $J \geq 0$, (Φ, J) means $(\Phi, \operatorname{supp} J)$.

3.2.3 The lattice and positive decomposition theorems

For a family $\mathcal{H}$ of t-uniform hypergraphs on the role set $[q]$, write $\mathcal{H}(\Phi)$ for their incidence vectors after embeddings in $\Phi_{[q]}$. The notation $\langle \mathcal{H}(\Phi) \rangle_{\mathbb{Z}}$ denotes their integer span, allowing negative coefficients. A decomposition, in contrast, is a sum with nonnegative integer coefficients equal to the host vector.

Theorem 3.4 (Clique-lattice input). *Fix $j \geq 2$. Let $H = \binom{[j]}{j-1}$ and let Ψ be an S_j-adapted $[j]$-complex on N vertices which is $(\omega, 3(j-1)^2)$-extendable. For a fixed $\omega > 0$ and sufficiently large N, a nonnegative integer vector W on Ψ°_{j-1} belongs to $\langle H(\Psi)\rangle_{\mathbb{Z}}$ if*

$$(j-i) \mid \sum_{e \supseteq u} W(e) \qquad (0 \leq i \leq j-2,\ u \in \tbinom{V(\Psi)}{i}).$$

This is [14, Lemma 5.20]. In that lemma the required lower bound $\omega > N^{-1/2}$ holds eventually for fixed ω. The divisors displayed here are the degrees of $\binom{[j]}{j-1}$ at its i-subsets. Faces u outside the complex have degree zero and cause no additional condition.

For $J \geq 0$, $(\mathcal{H}, c, \omega)$-*regularity in* Φ means that, for every $H \in \mathcal{H}$ and every $\phi \in \Phi_{[q]}$ for which $\phi(H) \leq J$ coordinatewise, there is a weight

$$\omega n^{t-q} \leq w_{H,\phi} \leq \omega^{-1} n^{t-q}, \qquad \sum_{H,\phi} w_{H,\phi}\phi(H) = (1 \pm c)J \tag{11}$$

coordinatewise. The weights are attached to ordered embeddings, not to isomorphism classes of copies.

Theorem 3.5 (Fixed-family decomposition input). *Fix $1 \leq t < q$ and a family $\mathcal{H}$ of t-uniform hypergraphs on $[q]$. There are constants $n_0(q)$ and $\omega_0(q) > 0$ with the following consequence. Put*

$$h_q = 2^{50q^3}, \qquad \delta_q = 2^{-1000q^5}, \qquad c = \omega^{h_q^{20}}.$$

Suppose $n > n_0(q)$, $n^{-\delta_q} < \omega < \omega_0(q)$, and Φ is an exactly adapted, (ω, h_q)-extendable $[q]$-complex on n vertices. If

$$J \in \langle \mathcal{H}(\Phi)\rangle_{\mathbb{Z}}, \qquad J \geq 0,$$

J satisfies (11), and (Φ, J) is (ω, h_q)-extendable, then J has a $\mathcal{H}$-decomposition using embeddings in $\Phi_{[q]}$.

This is the one-colour specialization ($D = 1$) of [14, Theorem 7.12]. Motifs may have isolated vertices. In particular, padding a motif to the full role set does not remove the weight requirement for embeddings of its isolated vertices. We verify that requirement explicitly.

3.2.4 From ordinary typicality to labelled extendability

Lemma 3.6 (Typed extension lemma). *Fix $j, s \geq 1$, $0 < a \leq b$, positive lower bounds for all densities, and $B \geq 0$. Let $\mathcal{L}$ be a rank-at-most-j complex on X, $|X| = N$, which is $(1/2, (j+1)s; \mathbf{p})$-typical. Let Y be disjoint from X, with $aN \leq |Y| \leq bN$. On roles $[j+1]$, require the first j roles to map into X and the last into Y, and require the X-image of every partial map to be a face of $\mathcal{L}$. This defines Φ. If $\mathcal{S} \subseteq \mathcal{L}_j$ satisfies $\deg_{\mathcal{S}}(e) \leq B$ for every $(j-1)$-face, put*

$$G = (\mathcal{L}_j \setminus \mathcal{S}) \ \cup\ \{e \cup \{y\} : e \in \mathcal{L}_{j-1},\ y \in Y\}.$$

For sufficiently large N, both Φ and (Φ, G) are (ω, s)-extendable for a fixed $\omega > 0$. Also, Φ is exactly adapted to the permutations of the first j roles fixing the last role.

The untyped $[j]$-complex of all injections with image a face of $\mathcal{L}$ is S_j-adapted and (ω', s)-extendable if $\mathcal{L}$ is $(1/2, js; \mathbf{p})$-typical, for some fixed $\omega' > 0$.

Proof. A partial bijection among the first j roles extends to a permutation of those roles fixing $j+1$; the same is true when both its domain and range contain $j+1$ and it fixes that role. Such precompositions preserve both the type constraints and the X-face condition. Conversely, if a bijection sends a first role to role $j+1$, or role $j+1$ to a first role, precomposition assigns a Y-vertex to an X-role, or an X-vertex to the Y-role. This is impossible because $X \cap Y = \varnothing$. These are all cases, proving exact adaptation to the group that permutes the first j roles and fixes the last.

Consider an arbitrary rooted problem in the preceding definition, including arbitrary additional positive-support requirements. It has at most $Q = (j+1)s$ vertices. Project the images of its maps onto their X-vertices and close under subsets. This gives an ordinary pattern complex of rank at most j on at most Q vertices. A valid labelled root induces a valid root for every already imposed face of this ordinary pattern.

Embed the unrooted vertices in any fixed order. At an X-step, use the induced ordinary pattern on the already embedded X vertices and the new vertex. Include previously used X vertices as isolated roots where needed. Typicality supplies at least κN choices, where $\kappa > 0$ is a fixed lower bound for one half of every product in (10) of order at most Q. To meet a positive all-X support requirement, reject a choice if the face completed at this step lies in $\mathcal{S}$. Its other $j-1$ vertices form a valid fixed facet, so at most B choices are rejected per requirement. There are at most 2^Q possible underlying face requirements. For large N, at least $\kappa N/2$ choices remain.

A positive mixed requirement has no restriction beyond having its X-facet in $\mathcal{L}_{j-1}$; this is already imposed by the partial-map face condition. At a Y-step there are at least $|Y| - Q \geq aN/2$ choices. These observations also cover a mixed map whose entire X-part was rooted originally. Additional root-only all-X requirements do not occur, by definition.

Write $n = N + |Y| \leq (1+b)N$ and choose

$$0 < \kappa_0 \leq \min\{1, \kappa/(2(1+b)), a/(2(1+b))\}.$$

If v vertices are unrooted, sequential multiplication gives at least $\kappa_0^v n^v \geq \kappa_0^Q n^v$ extensions. This is uniform over all patterns and roots. Take $\omega = \kappa_0^Q$. With no support requirements it also proves extendability of Φ. The untyped assertion follows by the same argument without Y or the deleted support. □

3.3 Construction of one layer

Lemma 3.7 (One-layer completion)**.** *Fix $j \geq 2$, densities $p_2, \dots, p_j \in (0,1)$, an output order $h_{\mathrm{out}} \geq j$, errors $\epsilon, \eta > 0$, and $C_0 > 0$. Put*

$$Q_j = \prod_{t=2}^{j-1} p_t^{\binom{j-1}{t-1}}, \qquad L = \mathrm{lcm}(1, \dots, j).$$

There are $h_{\mathrm{in}} \geq h_{\mathrm{out}}$, $\tau_{\mathrm{in}} > 0$ and N_0 such that the following holds for $N \geq N_0$. Let $\mathcal{K}$ be a rank-at-most-$(j-1)$, $(\tau_{\mathrm{in}}, h_{\mathrm{in}}; p_2, \dots, p_{j-1})$-typical complex on X, $|X| = N$. Let λ be an integer with $|\lambda - p_j Q_j N| \leq C_0$. There is a simple supported j-uniform family $\mathcal{F}$ for which

$$0 \leq \lambda - \deg_{\mathcal{F}}(e) \leq L - 1 \qquad (e \in \mathcal{K}_{j-1}), \tag{12}$$

$$\sum_{e \in \mathcal{K}_{j-1}} (\lambda - \deg_{\mathcal{F}}(e)) \leq \eta(L-1)|\mathcal{K}_{j-1}|, \tag{13}$$

and $\mathcal{K} \cup \mathcal{F}$ is $(\epsilon, h_{\mathrm{out}}; p_2, \dots, p_j)$-typical.

Proof. We may assume $\epsilon, \eta < 1$ by decreasing them. In the proof write $Q = Q_j$ and $\mathcal{B} = \mathcal{K}_{j-1}$. Choose an order

$$H \geq \max\{h_{\text{out}}, (j+1)h_{j+1}, 3j(j-1)^2\}.$$

Take H to be an integer and set $h_{\text{in}} = H$. Fix a compact interval $I \subset (0,1)$ having p_j in its interior. A positive margin ξ will be chosen below; set

$$\theta = \lambda/(QN) + \xi, \qquad D = \theta QN = \lambda + \xi QN.$$

Since $\lambda/(QN) = p_j + O(N^{-1})$, sufficiently small fixed ξ and sufficiently large N ensure that $\theta \in I$. Select a random mask $\mathcal{A}$ of supported j-faces with probability θ. By Lemma 3.2, making the input accuracy sufficiently small gives a mask such that $\mathcal{L} := \mathcal{K} \cup \mathcal{A}$ is typical to order H with any prescribed error $\tau > 0$. In particular,

$$\deg_{\mathcal{A}}(e) = (1 \pm \tau)D \qquad (e \in \mathcal{B}). \tag{14}$$

To see this last formula directly, the rooted pattern of a single supported j-face has $b_t = \binom{j-1}{t-1}$ for $2 \leq t \leq j$. We specify the order of all remaining parameter choices below.

The bounded residual. Write $\lambda = Lm + b_0$, where $0 \leq b_0 < L$. For large N, $m = \Theta(N) > 0$. Form the j-uniform auxiliary hypergraph on vertex set $\mathcal{B}$ with one edge $\binom{T}{j-1}$ for each $T \in \mathcal{A}$. Its degrees satisfy (14), and its pair-codegrees are at most one: two distinct facets determine at most one possible j-set. Also $M = |\mathcal{B}| = \Theta(N^{j-1})$, so $D/(\log M)^4 \to \infty$.

Let γ be the tolerance supplied by Theorem 3.3 with $t = j$, $a = 4$, and error η. Choose $\tau < \gamma/2$. For all sufficiently large N, the codegree bound $1 < D/(\log M)^4$ holds. The theorem therefore supplies a matching covering at least $(1-\eta)M$ vertices. Remove only the selected j-faces, retain *all* auxiliary vertices, and repeat for b_0 rounds. The degree at each auxiliary vertex falls by at most one per round. The bounded total change is at most $b_0 < L$; after enlarging N so that $L/D < \gamma/2$, all remaining degrees lie between $(1-\gamma)D$ and $(1+\gamma)D$. Pair-codegrees cannot increase. Thus the theorem continues to apply with the same nominal D in every round. Let $\mathcal{S}$ be the union of the selected faces; use $\mathcal{S} = \varnothing$ when $b_0 = 0$. The union is simple and

$$\deg_{\mathcal{S}}(e) \leq b_0, \qquad \sum_{e \in \mathcal{B}} (b_0 - \deg_{\mathcal{S}}(e)) \leq \eta b_0 M. \tag{15}$$

A signed lattice representation for the divisible part. Let Ψ be the untyped $[j]$-complex of injections with images in $\mathcal{L}$. Lemma 3.6, with $s = 3(j-1)^2$, verifies its required extendability. For every i-set u, $0 \leq i \leq j-2$, the degree of the vector $L\mathbf{1}_{\mathcal{B}}$ at u is $L|\{e \in \mathcal{B} : u \subseteq e\}|$, an integer multiple of L. Since $L = \operatorname{lcm}(1, \dots, j)$, it is divisible by $j - i$. Theorem 3.4 therefore gives integers z_T, $T \in \mathcal{A}$, such that

$$\sum_{T \in \mathcal{A}} z_T \mathbf{1}_{\binom{T}{j-1}} = L\mathbf{1}_{\mathcal{B}}. \tag{16}$$

Neither positivity nor boundedness of the z_T is needed.

A typed host with capacity-one face coordinates. Introduce Y, disjoint from X, with $|Y| = m$, and let Φ be the typed complex in Lemma 3.6 built from $\mathcal{L}$. Its role set is $[j+1]$. Use the j-uniform motifs

$$H_1 = \binom{[j+1]}{j}, \qquad H_2 = \{[j]\}.$$

The $(j+1)$st vertex of H_2 is isolated. Define the nonnegative integer vector J on Φ_j° by

$$J(T) = \begin{cases} 1 & T \in \mathcal{A} \setminus \mathcal{S}, \\ 0 & T \in \mathcal{S}, \end{cases} \qquad J(e \cup \{y\}) = L \quad (e \in \mathcal{B},\ y \in Y).$$

These are all coordinate types in Φ_j°.

For each T, y, choose any one ordering of T on the first j roles, and denote its motif vectors by $H_i(T, y)$. Lift (16) at each label y. The resulting signed sum $\sum_{y,T} z_T H_1(T, y)$ has mixed coordinate L everywhere and all-X coordinate $m z_T$ at T. Fix one label $y_0 \in Y$ and add

$$\sum_{T \in \mathcal{A}} (\mathbf{1}_{\{T \notin \mathcal{S}\}} - m z_T) H_2(T, y_0).$$

This gives exactly J, proving $J \in \langle \{H_1, H_2\}(\Phi) \rangle_{\mathbb{Z}}$. In this signed calculation faces of $\mathcal{S}$ remain in the ambient complex, even though their host capacity is zero.

Every ordered positive embedding and its weight. For either motif, an embedding with $\phi(H_i) \le J$ is precisely an ordering of a face $T \in \mathcal{A} \setminus \mathcal{S}$ on the first j roles, together with a label $y \in Y$. There are exactly $j!$ such embeddings for each pair (T, y). Assign every one of the H_1 embeddings weight

$$w_1 = \frac{L}{j!D}, \qquad \alpha = \frac{Lm}{D} < 1,$$

and every one of the H_2 embeddings weight

$$w_2 = \frac{1-\alpha}{j!m}.$$

For a mixed coordinate $e \cup \{y\}$ the coverage is

$$j! w_1 \deg_{\mathcal{A} \setminus \mathcal{S}}(e) = L \left(1 \pm \left(\tau + \frac{b_0}{D}\right)\right).$$

For an available all-X coordinate T the coverage is exactly $mj!w_1 + mj!w_2 = \alpha + (1-\alpha) = 1$. Zero coordinates receive no positive embeddings. Moreover, $D - Lm = b_0 + \xi QN$, so both w_1 and w_2 are between fixed positive constants times $(N+m)^{-1}$. The constants can depend on ξ; ξ is fixed here.

Extension conditions and choice of parameters. The ratios m/N lie in a fixed positive compact interval for large N. By (15), $\deg_{\mathcal{S}}(e) \le b_0 \le L-1$, and the support of J is exactly the family G in Lemma 3.6. Apply that lemma with $s = h_{j+1}$ and $B = L-1$. It proves exact adaptation and both extendability conditions for Φ, J. Its extension constants can be chosen uniformly for $\tau \le 1/2$ and for the above compact interval of mask densities. Choose $\omega > 0$ below these constants, below the lower bounds on $(N+m)w_i$, below the reciprocals of their upper bounds, and below $\omega_0(j+1)$ in Theorem 3.5. Set $c = \omega^{h_{j+1}^{20}}$. Choose τ smaller than $c/2$ and the matching theorem's tolerance; enlarge N until $b_0/D < c/2$ and $(N+m)^{-\delta_{j+1}} < \omega$. The coverage calculations above now give $\sum_{i,\phi} w_{i,\phi} \phi(H_i) = (1 \pm c)J$, and the bounds on w_1, w_2 give the weight window in (11). Theorem 3.5 now gives a decomposition of J.

Retain the all-X face from each H_1 copy and adjoin $\mathcal{S}$; call the result $\mathcal{F}$. Each all-X coordinate has capacity one, so no retained face is repeated, and none lies in $\mathcal{S}$. For a fixed $e \in \mathcal{B}$ and label y, exactly L selected H_1 copies contain the mixed coordinate $e \cup \{y\}$. Consequently

$$\deg_{\mathcal{F}}(e) = Lm + \deg_{\mathcal{S}}(e) = \lambda - (b_0 - \deg_{\mathcal{S}}(e)). \tag{17}$$

This proves (12)–(13). The final faces lie on X; none of the auxiliary labels is retained.

Output typicality. Equations (14) and (17) give

$$\deg_{\mathcal{A}\setminus\mathcal{F}}(e) \le \xi QN + \tau D + b_0 =: \rho N.$$

Put

$$B_h = \binom{h_{\text{out}} - 1}{j-1}, \qquad q_* := \prod_{t=2}^{j} p_t^{\binom{h_{\text{out}}-1}{t-1}} > 0.$$

Fix a positive rooted pattern of order at most h_{out} and a root embedding valid in $\mathcal{K} \cup \mathcal{F}$. Write $E_{\mathcal{F}}$ for its number of extensions in $\mathcal{K} \cup \mathcal{F}$ and

$$\mu = N \prod_{t=2}^{j} p_t^{b_t}$$

for the target count. Its root is also valid in the mask complex $\mathcal{L}$. Lemma 3.2 and the mask typicality estimate show that

$$|E_{\mathcal{F}} - \mu| \le N(\tau + B_h|\theta - p_j| + B_h\rho).$$

Indeed, the deletion term is at most $B_h\rho N$, and $|\theta^{b_j} - p_j^{b_j}| \le b_j|\theta - p_j| \le B_h|\theta - p_j|$. Also $\mu \ge q_* N$, $|\theta - p_j| \le \xi + C_0/(QN)$, and $\rho = \xi Q + \tau D/N + b_0/N$. On the fixed compact interval for θ, these inequalities give

$$\frac{|E_{\mathcal{F}} - \mu|}{\mu} \le C_1(\xi + \tau + N^{-1}),$$

where C_1 depends only on the fixed orders, densities, and C_0, and may be chosen uniformly before ξ and τ.

The choices can now be made in a noncircular order. First fix H and I, and hence a valid uniform constant C_1. Then choose $\xi > 0$ small enough that $C_1\xi < \epsilon/3$ and that $p_j + 2\xi \in I$. Next choose ω, c, then choose τ small enough for the decomposition, matching, and output errors. Lemma 3.2 then determines an input accuracy $\tau_{\text{in}} \le \min\{\epsilon, 1/2\}$, which we may decrease further. Finally choose N_0 to meet all the finitely many size requirements, including $C_1/N_0 < \epsilon/3$ and $C_0/(QN_0) < \xi$; the latter also makes the earlier assertion $\theta \in I$ explicit. This proves the quantified lemma. □

3.4 Induction and product amplification

Proof of Theorem 1.1. Fix d, δ, and use the densities from Lemma 2.4. Put $N = dr - 1$. At stage $j = 2, \dots, d$ prescribe

$$\lambda_j = (d - j + 1)r - 1, \qquad L_j = \operatorname{lcm}(1, \dots, j).$$

If Q_j has the meaning in Lemma 3.7, then

$$p_j Q_j N = \frac{d-j+1}{d}(dr - 1), \qquad \lambda_j - p_j Q_j N = -\frac{j-1}{d}.$$

Thus the lemma's budget hypothesis holds with $C_0 = 1$.

Choose $0 < \eta_j < 1$ such that

$$\sum_{j=2}^{d} \frac{\eta_j(L_j - 1)}{d - j + 1} < \delta. \tag{18}$$

This is possible because the sum has only $d - 1$ fixed positive coefficients. To organize the other constants, start with any fixed output order at least d and accuracy less than $1/2$ for stage d. Apply the parameter assertion of Lemma 3.7 backwards: the required input order at stage j becomes the prescribed output order at stage $j-1$, while the preceding output accuracy is chosen no larger than the required input accuracy at stage j. There are only $d-1$ stages. The rank-one complex satisfies the resulting first input condition for sufficiently large N. Increase r so that all the finitely many thresholds hold. Then apply the layers forwards. Lower levels are never changed.

At stage $j=i+1$, the degree of every i-face is at most $\lambda_{i+1}=(d-i)r-1$. Later stages do not change this layer, so (7) holds in the final complex. Lemma 2.1 excludes every core size. Let M_j be the number of j-faces. By incidence counting and (13),

$$jM_j=\sum_{e\in\mathcal{K}_{j-1}}\deg_{\mathcal{K}_j}(e),\qquad M_1=dr-1,\qquad M_j\geq\frac{\lambda_j-\eta_j(L_j-1)}{j}M_{j-1}.$$

Hence

$$M_d\geq\frac{dr-1}{d!}\prod_{j=2}^{d}((d-j+1)r-1-\eta_j(L_j-1)).\tag{19}$$

For sufficiently large r each factor is positive. After factoring out the leading r^d, the factors have the form $1-x_i$, with $0\leq x_i\leq 1$, and their sum is obtained by taking

$$x_1=\frac{1}{dr},\qquad x_j=\frac{1+\eta_j(L_j-1)}{(d-j+1)r}\quad(2\leq j\leq d).$$

Thus

$$\sum_i x_i=\frac{1}{r}\left(H_d+\sum_{j=2}^{d}\frac{\eta_j(L_j-1)}{d-j+1}\right).$$

The inequality $\prod_i(1-x_i)\geq 1-\sum_i x_i$, proved by induction from $(1-x)(1-y)\geq 1-x-y$, and (18) give (1). For $k=ad+b$, take the product of a disjoint copies of the rank-d family just constructed and b disjoint copies of the family of $r-1$ singletons. Lemma 2.3 gives (2). □

4 An effective growing-rank construction

This section proves Theorem 1.2 without using the decomposition inputs from Section 3. The approximate inverse below supplies a fractional vector with a small mean residual. A separate rounding lemma repairs that residual inside a near-full selected host. The final hierarchy verifies all hypotheses explicitly, including the rooted counts needed at every later layer.

4.1 Fractional correction from restricted fan counts

All matrix infinity norms in this subsection are induced norms, that is, maximum absolute row sums. Euclidean operator norms are denoted by $\|\cdot\|_2$. Incidence is unsigned. The extra lower-layer degree cap in the next lemma is essential; it is not inferred from approximate regularity.

Lemma 4.1 (Density-specific approximate inverse)**.** *Fix integers $d\geq 2$ and $1\leq h\leq d$, and suppose $N\geq 2^{160d}$. For $1\leq i\leq h$, let A_i be the incidence matrix from i-faces to $(i-1)$-faces of a complex on N vertices: rows are indexed by $(i-1)$-faces, columns by i-faces, and entries record containment. Assume all singletons are present, and put $Q_i=A_iA_i^{\mathsf{T}}$. Thus $Q_1=N$ as a 1-by-1 matrix. Suppose there are scalars $L_1=N$ and, for $2\leq i\leq h$,*

$$\frac{N}{2d}\leq L_i\leq L_{i-1}\leq N,\qquad \alpha_i=\frac{L_i}{L_{i-1}},\qquad c_i=L_i-\alpha_i(i-1),$$

such that

$$Q_i = c_i I + \alpha_i A_{i-1}^{\mathsf{T}} A_{i-1} + E_i, \qquad \|E_i\| * 2 \le N^{19/20}. \tag{20}$$

Assume additionally that every row of A_{i-1} has degree at most L_{i-1}. This cap is required only for the already selected lower layers: the current top layer may be an oversampled mask. Then, for every $1 \le i \le h$ and every real $s \ge 0$, there is a symmetric matrix $P_i(s)$ such that, with $T_i(s) = I - P_i(s)(Q_i + sI)$,

$$\|P_i(s)\| * \infty \le \frac{2^{3d}}{N+s}, \qquad \|T_i(s)\|_\infty \le 2^{4d}, \qquad \|T_i(s)\|_2 \le N^{-1/40}. \tag{21}$$

The parameter d remains the global target rank even when $h < d$.

Proof. Write

$$a_i = 1 - \frac{i-1}{L_{i-1}}, \qquad a = 1 - \frac{2d^2}{N}.$$

Then $a_i \ge a > 0$. The host threshold implies $N \ge 4d^3$, so Bernoulli's inequality gives

$$a^d \ge 1 - \frac{2d^3}{N} \ge \frac{1}{2}. \tag{22}$$

We construct approximate inverses without correcting them at lower ranks. Define

$$P_1(s) = \frac{1}{N+s}.$$

For $i \ge 2$, put $B = A_{i-1}$, $c = s + c_i$, and set

$$P_i(s) = \frac{1}{c}\left(I - B^{\mathsf{T}} P_{i-1}(c/\alpha_i) B\right). \tag{23}$$

Every shift in this recursion is positive. If $w = s/L_i$, then

$$c = L_i(w + a_i), \qquad \frac{c/\alpha_i}{L_{i-1}} = w + a_i. \tag{24}$$

Thus successive normalized shifts increase by almost one. The identity includes the diagonal centering term and remains exact when the top L_i is a mask prediction.

Let $F_i(w) = L_i \|P_i(L_i w)\|_\infty$. The lower-layer cap gives

$$\|B\| * \infty \le L * i - 1, \qquad \|B^{\mathsf{T}}\| * \infty = i - 1.$$

Equations (23) and (24) therefore imply

$$F_1(w) = \frac{1}{1+w}, \qquad F_i(w) \le \frac{1 + (i-1)F * i - 1(w + a_i)}{w + a_i}. \tag{25}$$

Expand this bound completely. After t descents the numerator is the falling factorial $(i-1)_t$, and successive denominators are at least $w + a, w + 2a, \ldots$. The final rank-one denominator adds 1, which is also at least a. Consequently

$$F_i(w) \le \sum_{t=0}^{i-1} \frac{(i-1) * t}{\prod * q = 1^{t+1}(w + qa)}.$$

For $w \geq 0$, $w + qa \geq a(w+q)$. Using (22), and then $w + q \geq q$ for $q \geq 2$, yields

$$F_i(w) \leq \frac{2}{1+w} \sum_{t=0}^{i-1} \frac{(i-1) * t}{(t+1)!} = \frac{2(2^i - 1)}{i(1+w)}. \tag{26}$$

Here $(i-1)_t/(t+1)! = \binom{i}{t+1}/i$. This binomial sum is the main cancellation: bounding each shifted lower inverse separately would lose it. Since $L_i + s \geq (N+s)/(2d)$,

$$\|P_i(s)\| * \infty \leq \frac{2^{3d}}{N+s}. \tag{27}$$

For $i \geq 2$, it suffices to note the sharper bound $4d2^i/[i(N+s)] \leq d2^{d+1}/(N+s) \leq 2^{3d}/(N+s)$; rank one is immediate. Induction in (23) shows that every $P_i(s)$ is symmetric. Therefore

$$\|P_i(s)\| * 2 \leq \|P_i(s)\| * \infty. \tag{28}$$

We next estimate the residual directly. Set

$$T_i(s) = I - P_i(s)(Q_i + sI).$$

Multiplying (23), with no commutation assumption, gives the exact identity

$$T_i(s) = -\frac{\alpha_i}{c} B^\mathsf{T} T_{i-1}(c/\alpha_i) B - P_i(s) E_i. \tag{29}$$

Indeed, multiplication by $cI + \alpha_i B^\mathsf{T} B$ gives $I + (\alpha_i/c) B^\mathsf{T} T_{i-1} B$; the remaining product is $P_i E_i$.

The degree cap also gives

$$\|B\| * 2^2 \leq (i-1)L * i - 1, \qquad \frac{\alpha_i}{c} = \frac{1}{L_{i-1}(w + a_i)}.$$

Let $e_i(w) = \|T_i(L_i w)\|_2$. By (26), (28), and (20),

$$e_1(w) = 0, \qquad e_i(w) \leq \frac{i-1}{w + a_i} e_{i-1}(w + a_i) + \frac{4d2^i}{i(1+w)} N^{-1/20}. \tag{30}$$

For a loose bound, drop i from the denominator of the last term. Expanding this recurrence, and using that the normalized shift after t descents is at least $w + ta$, gives

$$\begin{aligned} e_i(w) &\leq 4dN^{-1/20} \sum_{t=0}^{i-2} \frac{(i-1) * t 2^{i-t}}{a^t t!(1+ta)} \\ &\leq 8d2^i N^{-1/20} \sum * t = 0^{i-2} \binom{i-1}{t} 2^{-t} \\ &\leq 16d3^{i-1} N^{-1/20} \leq 2^{4d} N^{-1/20}. \end{aligned}$$

For the second line, use $1 + ta \geq a(t+1)$ and $a^{-(t+1)} \leq 2$, and then discard the factor $1/(t+1)$. The final inequality follows from $16d3^{d-1} \leq d2^{2d+2} \leq 2^{4d}$ for $d \geq 2$.

Every Q_i has row sum at most iN, since each incident i-face contributes i to that sum and a fixed facet has at most N one-vertex extensions. Consequently $\|Q_i + sI\|_\infty \leq iN + s \leq d(N+s)$, and hence

$$\|T_i(s)\|_\infty \leq 1 + d2^{3d} \leq 2^{4d}, \qquad \|T_i(s)\|_2 \leq N^{-1/40}. \tag{31}$$

The second bound uses $N \geq 2^{160d}$.

□

4.1.1 Obtaining the Gram decomposition from restricted fans

We use the following elementary graph estimate.

Lemma 4.2 (Graph degree and codegree estimate)**.** *Let L be a graph on n vertices. Suppose $0 \le p \le 1$ and $0 \le \varepsilon \le 1/2$, and*

$$|\deg(v) - pn| \le \varepsilon n, \qquad |\operatorname{codeg}(u, v) - p^2 n| \le \varepsilon n \quad (u \ne v).$$

Then

$$\| \operatorname{Adj}(L) - p(J - I)\| * 2 \le \sqrt{n + 3\varepsilon n^2} + 1.$$

Proof. Put $Z = \operatorname{Adj}(L) - pJ$. For $u \ne v$,

$$(Z^2) * uv = \operatorname{codeg}(u, v) - p(\deg(u) + \deg(v)) + p^2 n,$$

so its absolute value is at most $3\varepsilon n$. On the diagonal, positive semidefiniteness of Z^2 and the degree hypothesis give

$$0 \le (Z^2)_{vv} = p(1 - p)n + (1 - 2p)(\deg(v) - pn) \le n/4 + \varepsilon n.$$

The absolute row sum of Z^2 is at most $n + 3\varepsilon n^2$. Since Z is symmetric, this bounds $\|Z\|_2^2$. Adding pI costs at most 1 in operator norm. ☐

Lemma 4.3 (Fan interface, with the selected-layer cap)**.** *Fix $d \ge 2$, let $N \ge 2^{160d}$ and $1 \le h \le d$. At every $2 \le i \le h$, suppose each $(i - 2)$-face link has a common scalar size prediction L_{i-1}, its vertex degrees have common prediction L_i, and its distinct-pair codegrees have common prediction L_i^2/L_{i-1}. Assume that all these counts have relative error at most $\eta \le N^{-1/8}$, and*

$$L_1 = N, \qquad N/(2d) \le L_i \le L_{i-1} \le N.$$

If every already selected lower layer A_{i-1} has row degrees at most L_{i-1}, then all hypotheses of Lemma 4.1 hold. No maximum-degree cap at L_h is required for the current top mask.

Proof. Fix i, and let L_τ be the graph link of an $(i - 2)$-face τ. Its actual number of vertices n is at most N and equals $(1 \pm \eta)L_{i-1}$. Set $p = \alpha_i = L_i/L_{i-1}$. Comparing the given degree prediction $L_i = pL_{i-1}$ with pn, and the codegree prediction $L_i^2/L_{i-1} = p^2 L_{i-1}$ with $p^2 n$, gives errors at most $2\eta L_{i-1} \le 4\eta n$. Here $4\eta \le 1/2$ under the host threshold. Lemma 4.2 consequently yields

$$\| \operatorname{Adj}(L_\tau) - \alpha_i(J_\tau - I_\tau)\| * 2 \le \sqrt{N} + \sqrt{12\eta}\, N + 1.$$

Embed these matrices on the corresponding facet rows. Two distinct facets occur in a common i-face precisely when they intersect in one $(i - 2)$-face and are adjacent in its link. Thus

$$Q_i = \operatorname{diag}(D * \sigma) + \sum_\tau \operatorname{Adj}(L_\tau), \qquad A_{i-1}^\intercal A_{i-1} = \sum_\tau J_\tau.$$

Every facet lies in exactly $i - 1$ link row sets, so $\sum_\tau I_\tau = (i - 1)I$. Therefore

$$Q_i = [L_i - \alpha_i(i - 1)]I + \alpha_i A_{i-1}^\intercal A_{i-1} + E_i.$$

The centering convention here is consistently $\operatorname{Adj} - \alpha_i(J - I)$; centering at $\alpha_i J$ instead would require a different scalar center.

The diagonal error is at most ηN. To sum the embedded errors, apply their quadratic-form bounds and use that every facet lies in $i - 1$ link row sets. This gives

$$\begin{aligned}\|E_i\|_2 &\le \eta N + (i - 1)(\sqrt{N} + \sqrt{12\eta}\, N + 1) \\ &\le 5d\sqrt{\eta}\, N + 2d\sqrt{N} \le 7dN^{15/16} \le N^{19/20}.\end{aligned}$$

The final inequality is equivalent to $N^{1/80} \ge 7d$, which follows from $N \ge 2^{160d}$ and $2^{2d} \ge 7d$ for $d \ge 2$. All remaining hypotheses are exactly the stated density relations and selected-layer caps. ☐

4.1.2 Fixed-length correction

A fixed number of correction terms makes the remaining fractional error small in mean while retaining a separate maximum-error bound.

Lemma 4.4 (Fixed-length approximate right inverse)**.** *Under the hypotheses of Lemma 4.1, fix a rank $j \le h \le d$ and write $A = A_j$, $Q = AA^{\mathsf{T}}$. There is a matrix $\mathcal{R}$ such that*

$$\|\mathcal{R}\| * \infty \le \frac{2^{400d}}{N}.$$

For every vector b on the M facet rows with $\|b\|_\infty \le \epsilon N$, the residual $q = A\mathcal{R}b - b$ satisfies

$$\frac{1}{M}\sum_v |q_v| \le \frac{\epsilon}{N}, \qquad \|q\| * \infty \le 2^{320d}\epsilon N. \tag{32}$$

The construction uses exactly 80 *correction terms, independently of the global rank d and the facet count M.*

Proof. Take the symmetric approximate inverse $P = P_j(0)$ supplied by Lemma 4.1, and put $T = I - PQ$. Its bounds are

$$\|P\| * \infty \le 2^{3d}/N, \qquad \|T\|2 \le N^{-1/40}, \qquad \|T\|\infty \le 2^{4d}.$$

The transpose infinity bound also holds, but not because T is symmetric: indeed $T^{\mathsf{T}} = I - QP$ and $\|Q\|_\infty \le jN$, so

$$\|T^{\mathsf{T}}\| * \infty \le 1 + jN\|P\| * \infty \le 2^{4d}.$$

Set $m = 80$ and

$$P_m = (I + T + \cdots + T^{m-1})P, \qquad \mathcal{R} = A^{\mathsf{T}}P_m.$$

The identity $T^tP = P(I - QP)^t$ shows that each T^tP is symmetric. Thus P_m is symmetric, and both residual orientations are explicit:

$$P_mQ = I - T^m, \qquad QP_m = I - (T^{\mathsf{T}})^m.$$

Consequently $q = -(T^{\mathsf{T}})^m b$. The norm bound is

$$\|\mathcal{R}\| * \infty \le \|A^{\mathsf{T}}\| * \infty \sum * t = 0^{79}\|T\| * \infty^t \|P\| * \infty \le \frac{80j\,2^{319d}}{N} \le \frac{2^{400d}}{N} \qquad (d \ge 2).$$

The Euclidean bound $\|(T^{\mathsf{T}})^{80}\|_2 \le N^{-2}$, followed by Cauchy–Schwarz, gives

$$\frac{1}{M}\sum_v |q_v| \le \frac{\|q\|_2}{\sqrt{M}} \le N^{-2}\frac{\|b\|_2}{\sqrt{M}} \le \epsilon/N.$$

The transpose infinity bound separately gives $\|q\|_\infty \le 2^{320d}\epsilon N$. There is no independence assumption on the coordinates of q. $\square$

4.2 One-sided rounding with a small mean residual

We use Banaszczyk's vector-balancing theorem in the explicit form stated as Theorem 1.1 of [21]: vectors of Euclidean norm at most one have a signed sum in $5K$ whenever the convex body K has standard Gaussian measure at least one half. The body used below is symmetric. The argument is for existence, without a claim of computational efficiency.

Lemma 4.5 (One-sided rounding with a small mean residual)**.** *Let $\mathcal{H}$ be a simple linear j-uniform hypergraph on M vertices, where $M \geq 1$ and $j \geq 2$ are integers, and let $D > 0$. Let A be its unsigned incidence matrix. Define*

$$\begin{gathered} L = \left\lceil \sqrt{4\log(j+1)} \right\rceil, \qquad Q^+ = 2^{j-1}(j+1)^{10(j-1)}, \qquad \delta = \frac{1}{100Q^+}, \\ U = \sqrt{2\log(1000M)}, \qquad B_* = 10(U+1), \qquad R_* = 2LQ^+ + 2jU, \\ b = 80\log(j+1), \qquad P_* = 10\sqrt{j}\,R_* + 2b, \qquad a = \log(100M). \end{gathered} \tag{33}$$

Suppose the degrees and integer capacities λ_v satisfy

$$D \leq \deg_{\mathcal{H}}(v) \leq (1+\delta/2)D, \qquad 0 \leq \deg_{\mathcal{H}}(v) - \lambda_v \leq \delta D/2. \tag{34}$$

Assume

$$D \geq \max \left\{ \begin{array}{l} 2000Q^+ \log(1000ML), \\ 1000U \log(1000M), \\ 100P_* j^3 \log(100M), \\ 4000Q^+ \sqrt{j}\, B_* \end{array} \right\}. \tag{35}$$

Suppose $x \in [0,1]^{E(\mathcal{H})}$ satisfies

$$Ax = \lambda - \mathbf{1} + q, \qquad t = \|q\| * \infty \leq \frac{\delta D}{100j}, \qquad \bar{q} = \frac{1}{M}\sum_v |q_v| \leq \frac{1}{100j}. \tag{36}$$

Then $\mathcal{H}$ has a simple subhypergraph F such that

$$\deg_F(v) \leq \lambda_v \quad \textit{for every } v, \tag{37}$$

$$\frac{1}{M}\sum_v (\lambda_v - \deg_F(v)) \leq 35\sqrt{j}, \tag{38}$$

$$\lambda_v - \deg_F(v) \leq 25\sqrt{j}\,B** + 6b + 3a + 7 + (17j+1)t \quad \textit{for every } v. \tag{39}$$

Proof. We construct a Gaussian-large body with covers in every near-full subhost, round the given fractional vector into it, and then perform a prefix deletion and two repair rounds.

Elementary tail estimates. We record the bounded-variable inequality used below. If martingale differences Z_i have $|Z_i| \leq c$ and a deterministic upper bound V on the sum of their conditional variances, then

$$\mathbb{P}\left(\sum_i Z_i \geq z\right) \leq \exp\left[-\frac{z^2}{2(V+cz/3)}\right] \qquad (z > 0). \tag{40}$$

Indeed, the elementary series bound

$$e^w \leq 1 + w + \frac{w^2}{2(1-|w|/3)} \qquad (|w| < 3)$$

gives an exponential-moment upper bound $\exp(s^2V/(2(1-sc/3)))$ by conditional iteration, for $0<s<3/c$. Markov's inequality with $s=z/(V+cz/3)$ proves (40); zero variance is deterministic. This includes independent bounded variables after centering. For an independent Bernoulli sum T we also use the product bound

$$\mathbb{E}e^{sT}\le \exp((e^s-1)\mathbb{E}T).$$

It gives the two-sided multiplicative Chernoff estimate $2e^{-\mu\rho^2/3}$ at relative deviation $0<\rho\le 1$ and mean μ, and at $s=1$ gives the pointwise deletion estimates below.

Two separate certificates and a robust body. Let $\mathcal{C}$ be the finite collection of subhosts $\mathcal{H}'$ on the same vertices obtained by removing at most δD edges at each vertex. Each has minimum degree at least $(1-\delta)D$. For every such subhost require a vector $z\in\mathbb{R}^M$ to have separate nonnegative edge certificates y^+,y^-, supported on $\mathcal{H}'$, with

$$\begin{gathered} A_{\mathcal{H}'}y^+\ge z_+, \qquad A_{\mathcal{H}'}y^-\ge (-z)*+, \\ \sum_e y_e^\pm \le \frac{5M}{2j}, \qquad y_e^\pm\le R/D, \qquad \max_v (A_{\mathcal{H}'}y^\pm)v\le B*. \end{gathered} \tag{41}$$

Intersect these requirements over $\mathcal{C}$ and add the slab

$$\left|\sum_v z_v\right|\le 4\sqrt{M}. \tag{42}$$

Call the result K. Swapping the certificates proves symmetry. For a nonnegative certificate, $Ay\ge z_+$ is equivalent to $Ay\ge z$, so convex combinations of certificates prove convexity. Compact certificate boxes give closedness under limits. The two load bounds imply $|z_v|\le B_*$, giving boundedness. A sufficiently small constant edge weighting certifies a neighborhood of zero in each permitted subhost: its loads are bounded below by the positive minimum degree and above by the original maximum degree, while its mass and individual edge weight tend to zero with the constant. Finiteness of $\mathcal{C}$ gives a common neighborhood, which may be shrunk to satisfy the slab. Thus K is a symmetric convex body. We do not assert that a single certificate of mass $5M/(2j)$ covers $|z|$.

Gaussian levels, including the random zero level. Let G_v be independent standard Gaussians and put

$$\begin{gathered} u_v^+=(G_v)*+,\quad u_v^-=(-G_v)*+,\qquad S_m^+=\{v:G_v\ge m\},\quad S_m^-=\{v:G_v\le -m\}, \\ q_m=\mathbb{P}(G\ge m)=\tfrac12\mathbb{P}(|G|\ge m)\quad (0\le m<L),\qquad h_v^\pm=(u_v^\pm-L)_+. \end{gathered}$$

The zero level is random for both signs. Conditional on a vertex being active at a fixed sign and level, its incident edges have disjoint other vertices by linearity. Its induced degree is therefore binomial with success probability q_m^{j-1}.

We have $L\ge 3$ and $L^2\le 8\log(j+1)+2\le 10\log(j+1)$. Integrating the Gaussian density on $[L,L+1/L]$ and its negative counterpart proves $\mathbb{P}(|G|\ge L)\ge e^{-L^2}$: the logarithm of this integral lower bound divided by e^{-L^2} is

$$L^2/2-\log L+\log\sqrt{2/\pi}-1-1/(2L^2),$$

positive at $L=3$ and increasing thereafter. Hence

$$q_m^{j-1}\ge 2^{-(j-1)}e^{-(j-1)L^2}\ge 1/Q^+. \tag{43}$$

Chernoff at relative deviation 1/4 and a union bound over both signs and all vertex-level pairs have failure at most

$$4ML\exp[-D/(48Q^+)] < .001$$

by the first host bound. Outside that event the induced degrees lie between $.75\deg_{\mathcal{H}}(v)q_m^{j-1}$ and $1.25\deg_{\mathcal{H}}(v)q_m^{j-1}$. Every permitted subhost retains at least $.74\deg_{\mathcal{H}}(v)q_m^{j-1}$, since $\delta D \leq .01\deg_{\mathcal{H}}(v)q_m^{j-1}$.

On a fixed permitted subhost use level weights

$$y_e^{(m,\pm)} = \frac{3}{2Dq_m^{j-1}} \quad (e \subseteq S_m^{\pm}),$$

and zero on the other edges. Each active vertex receives load at least 1.11 and at most $(15/8)(1+\delta/2) < 1.9$. Summing levels covers $\min(u_v^{\pm}, L)$. The respective active-level counts are bounded by

$$Y_v^+ = (G_v) * ++ \mathbf{1} * \{G_v \geq 0\}, \qquad Y_v^- = (-G_v) * ++ \mathbf{1} * \{G_v \leq 0\}.$$

Writing $\mu = \mathbb{E}|G| < 5/6$, each has mean $(1+\mu)/2 < 11/12$ and variance $1-(\mu-1)^2/4 \leq 1$. Chebyshev shows that

$$\sum_v Y_v^+ \leq M, \qquad \sum_v Y_v^- \leq M \tag{44}$$

fail jointly with probability at most $288/M$. The first host bound implies $D > 10^6$, and linearity gives

$$M-1 \geq (j-1)D, \qquad M > 10^6, \qquad M \geq 500j, \qquad M \geq 1000(j-1). \tag{45}$$

Thus $288/M < .001$. On (44), each sign's level mass is at most $1.9M/j$, its maximum edge weight is at most $(3/2)LQ^+/D$, and its vertex loads are at most $1.9(u_v^{\pm}+1)$.

Gaussian tails and simultaneous robustness. For either sign,

$$\mathbb{E}h_v^{\pm} \leq \frac{1}{\sqrt{2\pi}}e^{-L^2/2} \leq \frac{1}{2(j+1)^2}.$$

Markov gives combined failure at most $2j/(j+1)^2 \leq 4/9$ for

$$\sum_v h_v^+ \leq \frac{M}{2j}, \qquad \sum_v h_v^- \leq \frac{M}{2j}. \tag{46}$$

The coordinate event $\max_v |G_v| \leq U$ fails with probability at most $2Me^{-U^2/2} = .002$.

Let $\mathcal{N}_v$ denote the vertices other than v sharing an original edge with it. Then $|\mathcal{N}_v| \leq 2(j-1)D$. For a fixed sign the independent clipped variables $\min(h_w^{\pm}, U)$ on this neighborhood have expectation sum at most $jD/(j+1)^2 \leq 2D/9 < D/2$ and variance sum at most $UD/2$. Equation (40), with increment bound U and deviation at least $D/2$, gives

$$\mathbb{P}\left(\sum_{w\in\mathcal{N}_v} \min(h_w^{\pm}, U) > D\right) \leq e^{-D/(6U)}.$$

The second host bound pays the union over both signs and all vertices with failure below .001. This clipped-variable estimate is unconditional. Intersecting its event with the coordinate event then bounds the actual neighboring tail sums by D.

On every permitted $\mathcal{H}'$ add separate tail certificates

$$y_e^{\mathrm{tail},\pm} = \sum_{v\in e} \frac{h_v^\pm}{\deg_{\mathcal{H}'}(v)}.$$

Their own contribution at v is $h_v^\pm$. Linearity bounds the other contributions by $D/((1-\delta)D) < 2$. Each sign has total edge mass exactly $\sum_v h_v^\pm \le M/(2j)$ and maximum edge weight at most $2jU/D$. Combining levels and tails gives each sign mass at most $2.4M/j < 5M/(2j)$, maximum edge weight at most R_*/D, and vertex loads at most $1.9(U+1)+U+2 < B_*$. These conclusions hold simultaneously for every permitted subhost on the same Gaussian events. No subhost union bound is used.

Finally $\sum_v G_v$ is Gaussian of variance M, so the slab (42) fails with probability at most $2e^{-8} < .001$. The full Gaussian failure sum is

$$.001 + .001 + 4/9 + .002 + .001 + .001 < 1/2.$$

The terms account for level degrees, level masses, tail masses, the coordinate maximum, neighboring tails, and the slab. Thus K has Gaussian measure greater than one half.

Rounding the arbitrary fractional vector. Normalize every incidence column by $\sqrt{j}$. The published vector-balancing theorem applies to each subset of columns, giving a signed sum in $5\sqrt{j}K$. For a vector with denominator 2^n, alter the odd-numerator coordinates by plus or minus 2^{-n} according to such a signed sum. Both choices stay in $[0,1]$, and all new numerators are even. Repeat at denominators $2^{n-1},\dots,2$. The accumulated error belongs to

$$5\sqrt{j}\left(\sum_{i=1}^{n} 2^{-i}\right) K \subseteq 5\sqrt{j}K$$

by convexity and $0 \in K$. Approximate a real vector x by dyadic vectors, pass to a subsequence with constant zero-one output, and use closedness. We obtain $Z \in \{0,1\}^{E(\mathcal{H})}$ such that

$$X = A(Z-x) \in 5\sqrt{j}K, \qquad |X_v| \le s := 5\sqrt{j}B_*, \qquad \left|\sum_v X_v\right| \le 20\sqrt{jM}. \tag{47}$$

By (45), the last absolute sum is less than M. On every permitted subhost there is a positive certificate satisfying

$$Ay \ge X_+, \qquad \sum_e y_e \le \frac{25M}{2\sqrt{j}}, \qquad y_e \le \frac{5\sqrt{j}R_*}{D}, \qquad \max_v (Ay) * v \le 5\sqrt{j}B**. \tag{48}$$

We next verify that a suitable selected host belongs to the permitted class before using its supported certificate.

A prefix at coordinates with residual greater than one. Let $\mathcal{H}_0$ be the host selected by Z. At each vertex with $q_v > 1$, independently choose a uniform subset of $k_v = \lceil q_v \rceil$ distinct incident edges of $\mathcal{H}_0$ and mark them. Put $k_v = 0$ at vertices with $q_v \le 1$. Delete the union of the marks to obtain $\mathcal{H}_1$.

The choices are well defined. The fourth host bound gives $s \le \delta D/8$. The first gives $\delta D \ge 20\log(1000ML)$, hence $a+1 \le \delta D/10$. Also $t \le \delta D/200$, so

$$\deg_{\mathcal{H}_0}(v) = \lambda_v - 1 + q_v + X_v \ge D - \delta D/2 - 1 - t - s \ge D/2.$$

Every nonzero quota obeys $k_v \le 2q_v \le 2t < D/2$. For fixed u, each other origin shares at most one edge with u by linearity. The contributions are independent Bernoulli variables, of individual probability at most $4t/D$ and mean sum at most $8(j-1)t$. The exponential-moment bound and $e-1<2$ give a prefix choice, with failure probability at most $Me^{-a} = .01$, whose incident losses C_v satisfy

$$C_v \le 2t + 16(j-1)t + a \le 17jt + a, \qquad \sum_v C_v \le j \sum_v k_v \le 2j \sum_v |q_v|. \tag{49}$$

The total bound holds for every choice; overlapping marks only reduce actual losses. Fix a choice with the pointwise bound. Charging a ceiling only when $q_v > 1$ is essential for this total bound.

The number of original edges absent at each vertex after the prefix is at most

$$\delta D/2 + 1 + t + s + 17jt + a \le \delta D \left(\frac{1}{2} + \frac{1}{200} + \frac{1}{8} + \frac{17}{100} + \frac{1}{10} \right) = \frac{9}{10} \delta D < \delta D.$$

Thus $\mathcal{H}_1 \in \mathcal{C}$. Let $h_v = (\deg_{\mathcal{H}_1}(v) - \lambda_v)_+$ be its integer overflow. When $q_v \le 1$, the overflow before prefix deletion was at most $(X_v)_+$. When $q_v > 1$, all its own marked edges disappear, and

$$\deg_{\mathcal{H}_1}(v) - \lambda_v \le X_v + q_v - 1 - \lceil q_v \rceil \le X_v.$$

Thus $h_v \le (X_v)_+$. Choose the certificate (48) for this now fixed $\mathcal{H}_1$. It covers h and is supported on edges actually present. Only its positive sign is used in the repair.

Independent deletion and the exact residual moment. Set $\beta = 4\log(j+1) \le b$. Independently delete each edge of $\mathcal{H}_1$ with probability

$$p_e = 2y_e + \frac{\beta}{(1-\delta)D}. \tag{50}$$

Then $p_e \le P_*/D < 1$ by (35). Let T_v be the deletion count and $\mu_v = \mathbb{E}T_v$. The current host has minimum degree $(1-\delta)D$ and maximum degree at most $2D$, giving

$$2h_v + \beta \le \mu_v \le 10\sqrt{j}B_* + 3\beta. \tag{51}$$

For the residual integer overflow $h'_v = (h_v - T_v)_+$,

$$\begin{aligned}
\mathbb{E}2^{-T_v} &= \prod_{e \ni v} (1 - p_e/2) \le e^{-\mu_v/2}, \\
\mathbb{E}h' * v &= \sum * m = 1^{h_v} \mathbb{P}(T_v \le h_v - m) \\
&\le e^{-(1-\log 2)h_v - \beta/2} \sum_{m=1}^{h_v} 2^{-m} \le e^{-\beta/2} = (j+1)^{-2}.
\end{aligned}$$

The conclusion is immediate if $h_v = 0$. Markov shows that

$$j \sum_v h' * v \le M \tag{52}$$

fails with probability at most $j/(j+1)^2 \le 2/9$. No independence of the residual variables is asserted. The Bernoulli exponential-moment bound also gives, with simultaneous failure at most $Me^{-a} = .01$,

$$T_v \le 2\mu_v + a \le 20\sqrt{j}B ** + 6b + a \quad \text{for every } v. \tag{53}$$

Local residual sums by bounded conditional influence. For each u put $W_u = \sum_{v\in\mathcal{N}_u} h'_v$. It depends on at most $4jD^2$ independent edge-deletion bits: there are at most $2(j-1)D$ relevant vertices, each with at most $2D$ incident original edges. Changing one bit alters at most j summands, each by at most one. Expose the bits in any order. Coupling the same future bits with the current bit fixed at zero and one shows that the difference Δ_i of conditional future expectations has absolute value at most j. The Doob increment is $(B_i - p_i)\Delta_i$, with absolute value at most j and conditional variance at most $p_i j^2 \le P_* j^2/D$. Its predictable total variance is at most $4P_* j^3 D$. Equation (40) consequently gives

$$\mathbb{P}(W_u > \mathbb{E}W_u + D) \le \exp[-D/(10P_* j^3)].$$

Moreover $\mathbb{E}W_u \le 2jD/(j+1)^2 \le 4D/9 < D$. The third host threshold and a vertex union bound thus yield

$$W_u \le 2D \quad \text{for every } u \tag{54}$$

with failure below .001. This argument uses the independent deletion bits, not independence of the overlapping residuals.

Total first-deletion incidence. Let V be the total incidence loss of this deletion round. Near-regularity and (48) give

$$\mathbb{E}V \le 2j\sum_e y_e + \frac{1+\delta/2}{1-\delta}\beta M \le \left(25\sqrt{j} + \frac{9}{2}\log(j+1)\right) M =: \mu_0,$$

where $(1+\delta/2)/(1-\delta) \le 9/8$. The independent summands of V are bounded by j and have variance sum at most $j\mu_0$. Applying (40) at additional tolerance $\sqrt{j}M$ gives

$$\mathbb{P}(V > \mu_0 + \sqrt{j}M) \le \exp[-M/(60\sqrt{j})] < 1/10.$$

The denominator estimate uses $2(25 + 4.5 + 1/3) < 60$ and $\log(j+1) \le \sqrt{j}$; the last inequality uses $M \ge 1000(j-1)$. For completeness, $\sqrt{z} - \log(1+z)$ is positive at $z = 1$ and has derivative $(\sqrt{z}-1)^2/(2\sqrt{z}(1+z)) \ge 0$, proving this logarithmic inequality for every $j \ge 1$.

The four first-deletion good events have total failure at most $.1 + 2/9 + .01 + .001 < 1$. Fix a simultaneous good realization of (52), (53), (54), and

$$V \le \left(26\sqrt{j} + \frac{9}{2}\log(j+1)\right) M. \tag{55}$$

All prefix choices and the supported certificate were fixed before the independent edge-deletion experiment.

Final uniform repair at overfull vertices. At each remaining overfull vertex v, independently choose a uniform h'_v-subset of its remaining incident edges and mark them. Delete the union of these marks to obtain F. The choices exist and enforce all capacities: an overfull degree exceeds $\lambda_v \ge (1-\delta/2)D > 0$, and all its own h'_v marked edges disappear. For fixed u, each other origin shares at most one edge with u by linearity. Its contribution is a Bernoulli variable, and the choices at distinct origins are independent. Their mean sum is at most

$$\frac{W_u}{(1-\delta/2)D} \le 3.$$

The Bernoulli exponential-moment bound and a vertex union bound give a final choice with at most $6 + a$ such other-origin marks at every vertex, with failure at most $Me^{-a} = .01$. Multiple marks on one edge only overcount actual deletions. For every final choice the total incidence cost is at most $j\sum_v h'_v \le M$. The first-deletion realization is already fixed, so this conditional selection changes none of its bounds.

Average and pointwise deficit bookkeeping. Initially $\lambda - \deg_{\mathcal{H}_0} = \mathbf{1} - q - X$. Equation (47) bounds its signed total deficit by $(2+\bar{q})M$. Add the prefix cost (49), the first-deletion cost (55), and the final cost at most M. The final average deficit is at most

$$26\sqrt{j} + \frac{9}{2}\log(j+1) + 3 + (2j+1)\bar{q} \le 30.5\sqrt{j} + \frac{121}{40} < 33\sqrt{j} < 35\sqrt{j}.$$

Here $(2j+1)\bar{q} \le 1/40$ and $121/40 < (5/2)\sqrt{2} \le (5/2)\sqrt{j}$.

For the maximum deficit, write Y_v for the final marks from other origins. Its own final marks cancel any remaining own overflow, so

$$\lambda_v - \deg_F(v) \le (\lambda_v - \deg_{\mathcal{H}_1}(v)) * + + T_v + Y_v.$$

The first term is at most $1 + t + |X_v| + C_v$. The bounds for X_v, C_v, T_v, Y_v now give

$$\lambda_v - \deg_F(v) \le 25\sqrt{j}B ** + 6b + 3a + 7 + (17j+1)t.$$

This proves all conclusions directly from the stated mean-error hypotheses. □

4.3 The quantitative hierarchy

The parameters below simultaneously guarantee the restricted fan counts, the fractional cube margins, and the hypotheses of the mean-error rounding lemma. The rank parameter d is global throughout the construction, including at all truncated lower-rank stages.

We prove Theorem 1.2 in this subsection. Write

$$g_{\mathrm{m}}(r) = 300\frac{\log\log r}{(\log r)^{1/4}}.$$

We prove the theorem with all numerical interfaces explicit. The matrix input is Lemma 4.4, with its hypotheses inherited from Lemma 4.1. In particular its rank parameter is a fixed *global* d even when the current complex is truncated at a smaller rank. The geometric input is Lemma 4.5, which allows a small mean fractional error. We denote its local level and robustness parameters by Q and δ_j below. No exact-feasibility claim is made about an uncorrected approximate fractional solution.

4.3.1 The closed collection of rooted positive patterns

For $1 \le t \le d$, let S_t be a complete simplex on t vertices, with one free vertex w and its other $t-1$ vertices rooted. For $2 \le t \le d$, let D_t^{fan} be the union of the simplices on

$$P = \sigma \cup \{u, w\}, \qquad Q = \sigma \cup \{v, w\}, \qquad |\sigma| = t-2.$$

Here u, v, w are distinct and outside σ, and w is free. The root is $\sigma \cup \{u, v\}$; it requires the two simplices on $\sigma \cup \{u\}$ and $\sigma \cup \{v\}$, not the t-face $\sigma \cup \{u, v\}$.

At stage a, retain the rank-a truncation of every one of these patterns. A root is valid if it is injective and every root-only required face is present. An extension is a choice of w outside the root satisfying all required positive faces. There is no nonface condition. Truncation followed by truncation is the earlier tracked pattern, so the collection is closed under all restrictions used in the induction. It contains at most $d+1$ vertices per pattern and at most $2dN^d$ labeled rooted tests at any one stage.

The numbers of required new s-faces are respectively

$$b_s(S_t) = \binom{t-1}{s-1}, \qquad b_s(D_t^{\mathrm{fan}}) = 2\binom{t-1}{s-1} - \binom{t-2}{s-1}.$$

The stage-a prediction of a pattern F is $\prod_{s=2}^{a} p_s^{b_s(F)}$. Consequently, by the exact density identities of Lemma 2.4,

$$P(S_t) = \frac{d-t+1}{d}, \qquad P(D_t^{\mathrm{fan}}) = \frac{(d-t+1)^2}{d(d-t+2)}. \tag{56}$$

The double-fan identity divides the square of the size-t simplex product by the size-$(t-1)$ simplex product. At $t=2$ the latter is an empty product. Each full prediction is at least $1/(2d)$. All truncated predictions are at least their full predictions, because $p_s \le 1$. There are at most $B = 2^d$ required new faces in any pattern, each containing the free vertex.

We say the selected rank-a complex has error at most τ_a if every valid rooted test has its number of extensions equal to its fixed p_s prediction times N, with relative error at most τ_a. This is a positive, non-induced notion of typicality.

4.3.2 Global suppressions and stage amplification

Fix $d \ge 2$ and put $N = dr - 1$. Define

$$\begin{gathered} B = 2^d, \qquad \gamma = \frac{1}{4d}, \qquad K = 2^{400d}, \qquad J = 2^{320d}, \\ Q_* = 2^{d-1}(d+1)^{10(d-1)}, \qquad A_0 = 10^6 d^2 B, \qquad G = 10^6 d^2 Q_* d^d, \\ R = 2^{500d}, \qquad H = R^{2d}. \end{gathered} \tag{57}$$

Impose the entirely numerical threshold

$$\boxed{N = dr - 1 \ge R^{100d} = 2^{50000d^2}.} \tag{58}$$

For $1 \le a \le d$ set

$$\tau_a = \frac{N^{-1/8} R^{a-d}}{100G}, \qquad \xi_a = \frac{\tau_a}{A_0}, \qquad \epsilon_a = \frac{\xi_a}{128 d^2 K}. \tag{59}$$

Only ϵ_a with $a \ge 2$ will be used. The global factor G is not part of the ratio of consecutive stage tolerances.

For all $d \ge 2$, elementary logarithmic estimates give

$$\begin{gathered} \log_2 G \le 20d^2, \qquad \log_2 A_0 \le 13d, \qquad Q_* \le G \le R^d \le H, \\ R \ge 12800 d^2 K A_0, \qquad R \ge 2048 d^2 K A_0, \qquad R > 10^6 d^8, \qquad \gamma \ge R^{-1}. \end{gathered} \tag{60}$$

For example, using $\log_2(d+1) \le d$,

$$\log_2 G \le 20 + 2d + (d-1) + 10d(d-1) + d^2 \le 20d^2.$$

Also $\log_2 A_0 \le 20 + 3d \le 13d$, and $\log_2(12800 d^2 K A_0) \le 14 + 415d < 500d$. It follows that

$$\begin{gathered} \tau_{a-1} \le \epsilon_a/16, \qquad \tau_1, \xi_a, \epsilon_a \ge N^{-1/8}/H \quad (a \ge 2), \\ \epsilon_a \le \xi_a \le \tau_a \le N^{-1/8}, \qquad \xi_a \le \frac{1}{4d^d}, \qquad \xi_a \le \frac{\delta_j}{8d} \quad (2 \le j \le d), \end{gathered} \tag{61}$$

where $\delta_j = [100 \cdot 2^{j-1}(j+1)^{10(j-1)}]^{-1}$. For the common lower bound, the smallest used ϵ_a is $\epsilon_2 \geq N^{-1/8}R^{1-d}/G$ by $R \geq 12800d^2KA_0$. Also $\tau_1 \geq N^{-1/8}R^{-d}/G$. Use $G \leq R^d$. The last two upper bounds use the factors d^d and Q_* in G.

The threshold also implies

$$\begin{gathered} N \geq 2^{160d}, \qquad \log N \geq \max\{d, 6\}, \qquad \log N \leq N^{1/3} \leq \sqrt{N}, \\ \sqrt{N}/H \geq R^{48d} > 160000d^7, \\ N^{3/8} \geq 6002dH, \qquad N^{7/8}/H \geq d+1, \qquad N \geq 100d, \qquad N > 4d^2. \end{gathered} \tag{62}$$

The first and last bounds follow directly from $R^{100d} = 2^{50000d^2}$. The middle power estimates follow by subtracting $2d$ from respectively $50d$, $37.5d$, and $87.5d$ in powers of R, and using $R > 10^6d^8$. For the logarithmic inequality put $z = \log N \geq 6$: $z - 3\log z$ is positive at 6 and increasing thereafter. These estimates also imply $(\log N)^{3/2} \leq \sqrt{N}$. In particular $N/(4d) > d$ and $N \geq 6d$; combining the latter with $\lambda_a \geq N/(2d)$ gives $\lambda_a - 1 \geq N/(3d)$ at every stage.

4.3.3 Masking and the full lower-rank matrix interface

Begin with all N singleton faces. Every rank-one test has N minus its root size extensions. Its relative error is at most $d/N \leq \tau_1$. Suppose the selected rank-$(a-1)$ complex has error at most τ_{a-1} and satisfies all earlier degree caps. Put

$$\lambda_a = (d-a+1)r - 1, \qquad Q_a = \prod_{s=2}^{a-1} p_s^{\binom{a-1}{s-1}}, \qquad \theta_a = \frac{\lambda_a}{Q_aN} + \xi_a, \qquad D_{\text{top}} = \lambda_a + \xi_aQ_aN.$$

By Lemma 2.4,

$$\frac{1}{d} \leq Q_a \leq 1, \qquad \frac{\lambda_a}{Q_aN} = p_a - \frac{a-1}{dQ_aN}, \qquad \lambda_a \geq \frac{N}{2d}. \tag{63}$$

The lower bound $\xi_aN \geq N^{7/8}/H \geq d$ gives $\theta_a \geq p_a$: indeed $(a-1)/(dQ_a) \leq a-1 < d \leq \xi_aN$. The upper bounds $p_a \leq 1 - 1/(2d^d)$ and $\xi_a \leq 1/(4d^d)$ give

$$\frac{1}{2d} \leq \theta_a \leq 1 - \frac{1}{4d^d} < 1.$$

Independently retain every supported a-face with probability θ_a.

Replacing p_a by θ_a in a test's prediction cannot decrease it. Thus the new prediction is at least $1/(2d)$. For a fixed root, condition on its root-only mask indicators. Distinct candidate free vertices use disjoint new mask indicators, so their extension count is a sum of independent indicators. If the truncation has no new a-faces, this count is deterministic; otherwise it is binomial. In either case its mean is at least γN, because $\tau_{a-1} < 1/2$ and the new prediction is at least $1/(2d)$. Its inherited relative error is at most $\tau_{a-1} \leq \epsilon_a/16$. A Chernoff bound with tolerance $\epsilon_a/2$ therefore bounds the probability of relative error greater than ϵ_a by

$$2\exp(-\epsilon_a^2\gamma N/12).$$

The conditional estimate bounds the unconditional probability of the conjunction that this root is valid and its test fails. No independence across different tests is asserted. The union bound succeeds because

$$\epsilon_a^2\gamma\sqrt{N} \geq N^{1/4}R^{-4d-1} \geq R^{21d-1} > 24d. \tag{64}$$

Indeed $\log N \le \sqrt{N}$ and $\log(4d) \le d\sqrt{N}$ imply $\epsilon_a^2\gamma N/12 > \log(4d) + d\log N$, which pays for the at most $2dN^d$ tests. Fix a mask satisfying all these tests.

Here is the precise check of the matrix input for this mask, with truncated rank $h = a$ and the same global d. Let A_i denote the incidence of i-faces and their $(i-1)$-facets. Use

$$L_1 = N, \qquad L_i = \frac{N(d-i+1)}{d}\ (i < a), \qquad L_a = D_{\text{top}}. \tag{65}$$

For $i < a$ the single and double fans of the selected lower complex give link size, vertex degree, and distinct-vertex codegree predictions

$$L_{i-1}, \qquad L_i, \qquad \frac{L_i^2}{L_{i-1}},$$

respectively, with relative error at most τ_{a-1}. For $i = a$ the mask fans give the same three predictions, now with relative error at most ϵ_a. The codegree denominator is correct because the intersection simplex of the two top faces has rank $a-1$: the top mask density occurs twice, whereas that intersection density is divided out. The root does not require the face $\sigma \cup \{u, v\}$.

We have $N/(2d) \le L_i \le L_{i-1} \le N$. For the top monotonicity, $L_{a-1} - L_a$ is at least $N/d - \xi_a N > 0$; for earlier ranks it is immediate. Every selected lower incidence has the additional exact row cap required by the inverse:

$$\deg A_{i-1} \le \lambda_{i-1} = L_{i-1} - \frac{i-2}{d} \le L_{i-1} \qquad (3 \le i \le a). \tag{66}$$

For $i = 2$, the sole row of A_1 has degree exactly $N = L_1$. No such selected-row assumption is imposed on the top mask A_a. This condition follows from the integer caps of the induction, not merely from approximate regularity. All errors are at most $N^{-1/8}$ and $N \ge 2^{160d}$. The number of selected faces at every lower rank is at least N, as checked below. Lemma 4.3 therefore supplies the Gram decomposition and all hypotheses needed for Lemma 4.4.

4.3.4 A buffered fractional point and its residual

Aim at $\lambda_a - 1$, not at λ_a. Start with the constant edge vector

$$x_0 = \frac{\lambda_a - 1}{D_{\text{top}}}.$$

Its degree residual $b = (\lambda_a - 1)\mathbf{1} - A_a x_0$ has infinity norm at most $\epsilon_a N$. Apply Lemma 4.4, and put $x = x_0 + \mathcal{R}b$. Then

$$\begin{aligned} A_a x = (\lambda_a - 1)\mathbf{1} + q, &\qquad \|x - x_0\| * \infty \le K\epsilon_a = \frac{\xi_a}{128d^2}, \\ \overline{q} := \frac{1}{M}\sum_v |q_v| \le \frac{\epsilon_a}{N}, &\qquad t := \|q\| * \infty \le J\epsilon_a N = \frac{\xi_a N}{128d^2\, 2^{80d}}. \end{aligned} \tag{67}$$

There is no coordinate independence assumption on q.

The corrected vector really belongs to $[0,1]^E$. Since $\lambda_a - 1 \ge N/(3d)$ and $D_{\text{top}} \le N$, we have $x_0 \ge 1/(3d)$. Moreover

$$1 - x_0 = \frac{\xi_a Q_a N + 1}{D_{\text{top}}} \ge \frac{\xi_a}{d}.$$

The correction in (67) is smaller than both margins. This explicitly accounts for the unit buffer required by the prefix deletion.

4.3.5 All geometric thresholds and the pointwise residual

Let $\mathcal{H}$ be the a-uniform facet hypergraph of the mask: its vertices are selected $(a-1)$-faces, and each masked a-face gives its boundary as an edge. It is simple and linear, since two distinct $(a-1)$-faces can be together in only the a-face equal to their union. Write $j = a$, let M be its vertex count, and choose $D = \lambda_a$. Since $\epsilon_a \leq \xi_a/(2d)$,

$$D \leq \deg_{\mathcal{H}}(v) \leq D + 2\xi_a N \leq D + \frac{\delta_j D}{2} \leq 2D. \tag{68}$$

The lower bound uses $D_{\mathrm{top}} - \epsilon_a D_{\mathrm{top}} \geq D + \xi_a N/d - \epsilon_a N$. The upper near-full bound follows from $\xi_a \leq \delta_j/(8d)$ and $D \geq N/(2d)$.

For completeness, use the exact geometric notation

$$\begin{aligned} Q &= 2^{j-1}(j+1)^{10(j-1)}, & L &= \lceil \sqrt{4\log(j+1)} \rceil, \\ U &= \sqrt{2\log(1000M)}, & B_{\mathrm{geo}} &= 10(U+1), \\ R_{\mathrm{geo}} &= 2LQ + 2jU, & b_{\mathrm{geo}} &= 80\log(j+1), \\ P_0 &= 10\sqrt{j}R_{\mathrm{geo}} + 2b_{\mathrm{geo}}, & a_{\mathrm{geo}} &= \log(100M). \end{aligned}$$

Because $M \leq N^{j-1}$, $Q \leq Q_* \leq H$, and $\log N \geq \max\{d, 6\}$,

$$\begin{gathered} \log(1000M) \leq 4d\log N, \qquad \log(1000ML) \leq 5d\log N, \\ L \leq 3\sqrt{d}, \qquad U \leq 3\sqrt{d\log N}, \qquad B_{\mathrm{geo}} \leq 40\sqrt{d\log N}, \\ R_{\mathrm{geo}} \leq 6\sqrt{d}H + 6d^{3/2}\sqrt{\log N}, \qquad P_0 \leq 200d^2H\sqrt{\log N}. \end{gathered}$$

The four lower bounds on D in Lemma 4.5 are consequently at most

$$\begin{array}{ll} 2000Q\log(1000ML) & \leq 10000dH\sqrt{N}, \\ 1000U\log(1000M) & \leq 12000d^{3/2}\sqrt{N}, \\ 100P_0j^3\log(100M) & \leq 80000d^6H\sqrt{N}, \\ 4000Q\sqrt{j}B_{\mathrm{geo}} & \leq 160000dH\sqrt{N}. \end{array}$$

Here $(\log N)^{3/2} \leq \sqrt{N}$ was used in the third line. All four are at most $80000d^6H\sqrt{N}$ for $d \geq 2$. By (62), $\sqrt{N}/H > 160000d^7$, so $D \geq N/(2d)$ exceeds them. Notice that Q_* has only been bounded by H, not by R; the latter would fail uniformly as d increases.

The additional mean and maximum residual conditions also hold:

$$\bar{q} \leq \frac{1}{N} \leq \frac{1}{100j}, \qquad t \leq \frac{\xi_a N}{128d^2} \leq \frac{\delta_j N}{1024d^3} \leq \frac{\delta_j D}{100j}. \tag{69}$$

Apply Lemma 4.5. It supplies an actual selected a-face family with cap λ_a, average deficit at most $35\sqrt{a}$, and maximum deficit at most

$$E_a = 25\sqrt{a}B_{\mathrm{geo}} + 6b_{\mathrm{geo}} + 3a_{\mathrm{geo}} + 7 + (17a+1)t.$$

The preceding bounds imply

$$E_a \leq 3000d\log N + \frac{\xi_a N}{16}. \tag{70}$$

Indeed the residual part is at most $18d\,\xi_a N/(128d^2\,2^{80d}) \leq \xi_a N/16$. Also

$$\frac{3001d\log N}{N} \leq \frac{\xi_a}{2},$$

because it suffices that $N^{3/8} \geq 6002dH$, already included in (62). Therefore

$$\frac{E_a + d}{N} \leq \frac{9}{16}\xi_a < \xi_a. \tag{71}$$

The pointwise residual need not be $O(\log N)$: its larger $O(\xi_a N)$ term has been explicitly budgeted here.

These inequalities justify the face-count condition used above. Every selected facet degree is at least

$$\lambda_a - E_a \geq \frac{N}{2d} - \xi_a N \geq \frac{N}{4d} > d.$$

Hence incidence counting gives $M_a \geq M_{a-1}$, where M_i is the number of selected i-faces. Starting with $M_1 = N$, every lower face layer used as a matrix row set has size at least N. There is no reliance on a logarithmic-only deficit bound.

4.3.6 Propagation after arbitrary admissible rounding

Fix the selected output just constructed. At any selected $(a-1)$-face, the number of removed mask faces is at most

$$(\xi_a + \epsilon_a)N + E_a.$$

By Lemma 2.5, a positive rooted test loses at most B times this maximum facet deletion in its extension count. The change from mask prediction to the fixed density prediction is at most $B|\theta_a - p_a|N$, and

$$|\theta_a - p_a| \leq \xi_a + d/N.$$

Since every prediction is at least γ, the new relative error is bounded by

$$\frac{\epsilon_a + B\big(2\xi_a + \epsilon_a + (E_a + d)/N\big)}{\gamma} \leq \frac{5B\xi_a}{\gamma} = 20dB\xi_a \leq \tau_a. \tag{72}$$

Here (71) was used. A root valid in the output was already valid in the mask and in the previous truncated complex, so every count invoked is among the tests already controlled.

This closes the induction simultaneously for every fan truncation. Only maximum deletion counts, not a distributional assertion about rounded families, have been used. In particular the selected lower row caps needed at the next matrix step remain the exact integer caps supplied by the repair lemma.

4.3.7 Cardinality, sunflower avoidance, and all uniform ranks

The average conclusion of the repair lemma gives

$$M_1 = dr - 1, \qquad M_a \geq \frac{\lambda_a - 35\sqrt{a}}{a} M_{a-1} \quad (2 \leq a \leq d).$$

Let

$$T_d = H_d + 35\sum_{a=2}^{d} \frac{\sqrt{a}}{d-a+1} \leq (35\sqrt{d}+1)H_d \leq 36d^2,$$

where $H_d = \sum_{i=1}^{d} 1/i$. Factoring the successive main terms produces the losses

$$z_1 = \frac{1}{dr}, \qquad z_a = \frac{1+35\sqrt{a}}{(d-a+1)r} \quad (2 \leq a \leq d), \qquad \sum_{a=1}^{d} z_a = \frac{T_d}{r}.$$

The threshold estimate in the next display makes their sum less than $1/100$, so every z_a lies in $[0,1]$. Therefore $\prod_a(1-z_a) \geq 1 - \sum_a z_a$ gives

$$M_d \geq r^d(1 - T_d/r).$$

The threshold gives the stronger estimate

$$x := T_d/r \leq \frac{36d^3}{N+1} < \frac{1}{100}.$$

Indeed $N \geq 2^{50000d^2} > 3600d^3$; taking binary logarithms reduces the latter comparison to

$$\log_2(3600) + 3\log_2 d < 12 + 3d < 50000d^2.$$

For $0 \leq x < 1$, integrating the derivative of $(1-x)^{1/d}$ gives

$$1 - (1-x)^{1/d} = \frac{1}{d}\int_0^x (1-t)^{1/d-1}\,dt \leq \frac{x}{d(1-x)}.$$

Consequently, using $(100/99)\cdot 36 < 37$ and $T_d/d \leq 36(1+\log d)/\sqrt{d}$, we obtain

$$M_d^{1/d} \geq r - \frac{100T_d}{99d} \geq r - 37\frac{1+\log d}{\sqrt{d}}. \tag{73}$$

The resulting downward-closed complex has $N = dr - 1$ vertices, and every selected c-face, $1 \leq c < d$, has at most $(d-c)r - 1$ one-vertex extensions. Lemma 2.1 therefore makes its top faces an r-sunflower-free family.

Now suppose $\log\log r \geq 40$. Write $L_r = \log r$, $u = \log L_r$, and choose

$$d = \left\lfloor \sqrt{L_r/2^{16}} \right\rfloor. \tag{74}$$

Write $y = \sqrt{L_r}/256 = e^{u/2}/256$. Since $u \geq 40$ and $e > 2$, we have $y > 2^{12} > 81$. Thus $d = \lfloor y \rfloor \geq 2$ and $d \geq y - 1 \geq (80/81)y$. The threshold is permitted because

$$\log(R^{100d}) = 50000(\log 2)d^2 < 50000d^2 < L_r, \qquad N = dr - 1 \geq r.$$

The sharper floor estimate and the exact scale of y give

$$\frac{1}{\sqrt{d}} \leq 16\sqrt{\frac{81}{80}}\,L_r^{-1/4} < \frac{161}{10}L_r^{-1/4}, \qquad 1 + \log d \leq \frac{u}{2} + 1 - 8\log 2 < \frac{u}{2}.$$

For the strict numerical comparison, the two squared constants are $259.20 < 259.21$; the last inequality uses $\log 2 > 1/2$. Substitution in (73) gives

$$M_d^{1/d} \geq r - 297.85\,uL_r^{-1/4} \geq r - 300\frac{\log\log r}{(\log r)^{1/4}},$$

since $37 \cdot (161/10)/2 = 297.85 < 300$. The displayed loss is less than one on the whole range: $\log 300 + \log u - u/4$ is decreasing for $u \geq 40$ and is negative at 40, using $\log 300 < 6$ and $\log 40 < 4$.

By Lemma 2.3, products extend this one seed to every uniform rank with an absolute factor $1/2$. Put $b = r - g_{\mathrm{m}}(r)$. Then $M_d \geq b^d$ and $r - 1 < b < r$. For $k = qd + t$, $0 \leq t < d$, take q independent copies of the rank-d family and a product of t families of $r-1$ singletons. Then

$$f_r(k) \geq b^{qd}(r-1)^t \geq b^k(1-1/b)^{d-1} \geq \frac{1}{2}b^k,$$

because $b \geq r - 1 \geq 2(d-1)$. The estimate $r - 1 \geq 2(d-1)$ follows, for example, from $N > 4d^2$ and $r = (N+1)/d$, which were established in (62).

This completes Theorem 1.2.

5 Further consequences and scope

5.1 An exact even-petal triple-shadow extremum

The *pair-shadow* of a triple family is the graph of pairs contained in at least one triple. The *pair-codegree* of a pair is the number of triples containing that pair.

Theorem 5.1 (Triple-shadow extremum)**.** *For every sufficiently large even integer r,*

$$f_r(3) \geq U(r) := \left\lfloor r^3 - \frac{11}{6}r^2 + \frac{r}{3} + \frac{1}{3} \right\rfloor .$$

The value $U(r)$ is the exact maximum within the class of triple families on at most $3r-1$ vertices whose pair-shadow has maximum degree at most $2r-1$ and whose pair-codegrees are at most $r-1$.

We give the full proof in Appendix A. This restricted optimality statement is not an upper bound for unrestricted $f_r(3)$. Since

$$U(r)^{1/3} = r - \frac{11}{18} + O(r^{-1}), \qquad \sqrt{E_r} = r - \frac{3}{4} + O(r^{-1}) \quad (r \text{ even}),$$

its cube-root product base improves on the even-r graph-product base by $5/36 + O(r^{-1})$. It is separate from the all-parity construction in Theorem 1.1.

5.2 Scope of the constructions

Theorem 1.1 asymptotically attains the harmonic second-order coefficient within the certificate class of Proposition 2.2. Its decomposition inputs are applied with all required lattice, regularity, and extension conditions verified in Section 3. The unspecified existence thresholds of those inputs are not used quantitatively.

Theorem 1.2 has a different dependency chain. It uses a fixed 80-term approximate correction and repairs the remaining small mean error inside the selected host. The effective construction threshold is $N \geq 2^{50000d^2}$, and the final one-sided rounding lemma bounds the average deficit by $35\sqrt{j}$ while separately controlling its maximum. These estimates allow d of order $\sqrt{\log r}$ and yield the rate in (3). They do not assert exact fractional feasibility before the residual repair, or a better fixed-rank coefficient than Theorem 1.1.

The external inputs of the paper are correspondingly different. The fixed-rank argument uses the matching, lattice, and design decomposition theorems stated in Subsection 3.2; the exact-triple argument uses the lattice and design decomposition theorems; and the effective argument uses the vector-balancing theorem stated in Section 4. Restricted fan counts are not used to infer the stronger labelled extendability needed for designs.

All three stated results give lower bounds for the unrestricted function $f_r(k)$. The two optimality assertions concern only their specified certificate and triple-shadow classes. In particular, no matching upper bound for the unrestricted exponential base β_r is asserted. The constants and cutoff in (3) are conservative.

5.3 Finite certificates and reproducibility

The accompanying document `sfr_combined_finite_supplement.pdf` contains a structural four-petal certificate with 39 triples, additional verified certificates for $r = 6, 8, 10, 12$, and the associated data and exact verification instructions. These finite examples are independent of the unspecified asymptotic thresholds. They are not claims of unrestricted optimality, and no numerical search result is used to prove any theorem in this paper.

A Proof of the exact triple-shadow extremum

In this appendix $q = r - 1$ is odd and $N = 3q + 2 = 3r - 1$. A graph G on N vertices is *common-neighbourhood typical* at density p to order h with error τ if

$$\left|\bigcap_{v\in S} N_G(v)\right| = (1 \pm \tau)p^{|S|}N \qquad (S \subseteq V(G),\ |S| \le h).$$

The empty-set intersection is $V(G)$. When counting injective extensions, excluding a fixed number of other root vertices changes these counts by $O_h(1)$. The clique complex truncated at rank three then satisfies Definition 3.1 with $p_2 = p$, $p_3 = 1$, with a slightly larger error. A required triangle contributes no additional restriction once all its edges are present: when the new vertex is embedded, its required incident edges are precisely a common-neighbourhood condition on a fixed set of root vertices.

A.1 A typical shadow with the prescribed degree sequence

Lemma A.1. *For every fixed order and every positive accuracy, all sufficiently large odd q admit a graph G on $N = 3q + 2$ vertices and a matching P covering all but one vertex v_0 such that $G - P$ is $2q$-regular, $\deg_G(v_0) = 2q$, and all other degrees in G are $2q + 1$. The graph G is common-neighbourhood typical to the prescribed order and accuracy at density $p = 2/3$.*

Proof. Work on the cyclic group $\mathbb{Z}_N$. Among the $D_0 = (N - 1)/2$ distance classes $\{d, -d\}$ with $d \neq 0$, choose q classes uniformly and include every edge with a chosen distance. The resulting graph B is exactly $2q$-regular.

The common-neighbourhood count for $S = \varnothing$ is deterministic. Now fix S of size $s \ge 1$. Except for $O_s(1)$ vertices x, the distances from x to the elements of S are nonzero and belong to distinct classes. Indeed, equality of two such classes for distinct $a, b \in S$ forces $2x = a + b$; this equation has a unique solution since N is odd. For a nonexceptional x the probability of being a common neighbour is

$$\frac{(q)_s}{(D_0)_s} = \left(\frac{q}{D_0}\right)^s + O_s(N^{-1}),$$

where $(a)_s = a(a-1)\cdots(a-s+1)$. Thus the expected common-neighbourhood size is $(q/D_0)^sN + O_s(1)$.

Changing the selection by one swap of distance classes changes this size by at most $4s$: each of the two classes affects at most two possible neighbours per member of S. Expose the q selected classes in a random order. Conditional completions after two possible choices at any step can be coupled by swapping those two classes. The corresponding Doob martingale differences are therefore bounded by $4s$. The bounded-differences inequality gives a tail at most $2\exp(-c_st^2/N)$ for a deviation t. Apply the same argument to the unchosen classes to control the complement $\overline{B}$. Taking t to be a fixed positive multiple of N, the exponential tail dominates the polynomial number of sets S. A union bound over all sets of every size up to the prescribed fixed order therefore shows that B and $\overline{B}$ simultaneously have arbitrarily accurate common-neighbourhood estimates at limiting densities $2/3$ and $1/3$, respectively. Consequently a choice of B with both sets of estimates exists for all sufficiently large N.

We record why $\overline{B} - v_0$ has a perfect matching, for any predetermined v_0. More generally, suppose a graph L has degrees $(p' + O(\tau))N$ and codegrees $(p'^2 + O(\tau))N$, with $p' > 0$ fixed. For $S \subseteq V(L)$, put $d_S(v) = |N_L(v) \cap S|$. Counting degrees and ordered pairs in S gives

$$\sum_v d_S(v) = p'N|S| + O(\tau N|S|),$$

and

$$\sum_v d_S(v)^2 = \sum_{x\in S} \deg_L(x) + \sum_{\substack{x,y\in S\\ x\neq y}} |N_L(x)\cap N_L(y)| = p'^2 N|S|^2 + O(\tau N|S|^2 + N|S|).$$

Expanding the square therefore gives

$$\sum_v (d_S(v) - p'|S|)^2 \le C(\tau N|S|^2 + N|S|).$$

For any fixed $\alpha > 0$, sufficiently small τ and large N, there cannot be disjoint sets S, T of sizes at least αN with no edge between them: the left side would be at least $|T|p'^2|S|^2$, a contradiction.

Partition $V(L)\setminus\{v_0\}$ randomly into equal parts A_0, B_0. Hypergeometric concentration and a union bound give a partition in which every vertex of $A_0 \cup B_0$ has at least βN neighbours in the opposite part, for a fixed $0 < \beta < p'/4$: conditionally on the side containing a vertex, its expected opposite-part degree is $(p'/2 + O(\tau) + O(N^{-1}))N$. Use the preceding no-empty-pair property with $\alpha < \beta$. Hall's condition follows for $S \subseteq A_0$: for $0 < |S| \le \beta N$ it follows from one vertex's degree; for $\beta N < |S| \le |B_0| - \beta N$, a violation would leave a set $B_0 \setminus N(S)$ of size at least βN with no edge to S; and for $|S| > |B_0| - \beta N$, every vertex of B_0 has a neighbour in S. The empty set is harmless. Hall's theorem supplies the required perfect matching.

Apply this to $L = \overline{B}$ and let P be the resulting matching. Then $G = B \cup P$ has the required degree sequence. Adding a matching changes the common-neighbourhood count of an s-set by at most s. This preserves all the prescribed typicality estimates for large N. □

A.2 The upper bound and a bounded parity leave

For the upper bound in Theorem 5.1, pad the ground set of an admissible triple family to N vertices. A vertex link is a graph with at most $2q+1$ vertices and maximum degree at most q: its vertices are shadow neighbours, and the degree of a shadow neighbour is the codegree of the corresponding pair. Hence the link has at most

$$\mathcal{E}_q = \left\lfloor \frac{q(2q+1)}{2} \right\rfloor = q^2 + \frac{q-1}{2}$$

edges. Since N and $2q+1$ are odd, the handshaking lemma forbids every vertex of the pair-shadow from having degree $2q+1$. Some vertex therefore has shadow degree at most $2q$ and hence link size at most q^2. Summing link sizes yields

$$3|\mathcal{F}| \le (N-1)\mathcal{E}_q + q^2 = N\mathcal{E}_q - \frac{q-1}{2}. \tag{75}$$

Its floor after division by three is $U(r)$, as stated. Indeed, with $q = r-1$ and $N = 3q+2$ the unfloored quantity is

$$\frac{N\mathcal{E}_q - (q-1)/2}{3} = \frac{6q^3 + 7q^2 - 2q - 1}{6} = r^3 - \frac{11}{6}r^2 + \frac{r}{3} + \frac{1}{3}.$$

We next attain this bound. Choose G, P, v_0 from Lemma A.1, at an order and accuracy to be fixed below. Write

$$E = |G| = \frac{N(2q+1)-1}{2}, \qquad a = |P| = \frac{N-1}{2}.$$

Let $\ell \in \{0,1,2\}$ satisfy $\ell \equiv qE - a \pmod 3$. There is a triangle-free graph $D \subseteq G$ such that

$$|D| = a + \ell, \quad \Delta(D) \le 5, \quad \deg_D(v) \equiv \deg_G(v) \pmod 2 \quad \text{for every } v. \tag{76}$$

To construct it, choose a matched vertex u with mate u'. Among the neighbours of u other than u', v_0, at least $2q-1$ vertices occur in the other $(3q-1)/2$ matching pairs. At least $(2q-1)-(3q-1)/2 = (q-1)/2$ of those pairs have both endpoints adjacent to u. For large q, choose ℓ such pairs a_ib_i (with no choice needed when $\ell = 0$). In P replace each a_ib_i by a_iu, b_iu. The result is one star with $2\ell+1$ edges together with disjoint matching edges. Each replacement increases the edge count by one, preserves the degrees of a_i, b_i, and increases the degree of u by two. The parity vector of P already agrees with that of G: both degrees at v_0 are even, and both degrees at every other vertex are odd. The resulting graph is triangle-free and has maximum degree $2\ell + 1 \le 5$, so (76) holds.

A.3 A residual triangle multiset

Write $q = 6m + b$, where $b \in \{1,3,5\}$. Consider the pair-multiplicity vector

$$J_0 = b\mathbf{1}_{E(G)} - \mathbf{1}_{E(D)}.$$

It is nonnegative, its vertex degrees are even, and its total weight $bE - |D|$ is divisible by three. Indeed, b is odd and $\deg_D(v) \equiv \deg_G(v) \pmod 2$, while

$$bE - |D| \equiv qE - a - \ell \equiv 0 \pmod 3.$$

These are precisely the triangle divisibility conditions.

We apply the lattice and decomposition inputs from Subsection 3.2 to obtain a *multiset* of triangles $\mathcal{T}_0$ whose pair-incidence vector is J_0. Here are the required checks. Throughout put $p = 2/3$. If $b = 1$, use the clique complex of $G - D$ truncated at rank three. Deleting a graph of maximum degree five changes every common-neighbourhood count of a fixed s-set by at most $5s$, so this complex is extendable. The host is the simple graph $G - D$. Theorem 3.4 with $j = 3$ supplies its lattice membership, and weight $1/(6p^2N)$ on each ordered triangle covers every edge by $1 + O(\tau + N^{-1})$.

If $b = 3$ or 5, use the clique complex of G. The support of J_0 is all of G. Give an ordered triangle with underlying set T the weight

$$\frac{b - |E(T) \cap E(D)|}{6p^2N}. \tag{77}$$

The numerator is at least $b - 2 \ge 1$, since D is triangle-free. Each underlying triangle has six ordered embeddings. For an edge uv, summing over its common neighbours gives coverage

$$\frac{(b - \mathbf{1}_{uv\in D})|N_G(u) \cap N_G(v)| - \sum_{w \in N_G(u)\cap N_G(v)} (\mathbf{1}_{uw\in D} + \mathbf{1}_{vw \in D})}{p^2N}.$$

The subtracted sum is at most ten. This equals $J_0(uv)(1 + O(\tau + N^{-1}))$. Theorem 3.4 again supplies membership in the triangle lattice.

In both cases the ambient complex is exactly S_3-adapted, because every bijection between subsets of the three roles extends to a permutation in S_3. Its extensions follow by sequential common-neighbour choices. The positive graph support is the entire edge level of the corresponding ambient complex, so the support extension condition imposes nothing further. The weights are bounded above and below by fixed positive multiples of N^{-1}. Choose the extension and weight constant ω first, then the required regularity accuracy, then the initial graph accuracy and N.

Theorem 3.5, applied at uniformity two with the single triangle motif on three roles, gives $\mathcal{T}_0$. This role count is fixed and is unrelated to the petal parameter $q = r - 1$ used in this appendix. We do not assume that its unordered triangles are distinct.

Let $\mathcal{S}$ be the simple support of $\mathcal{T}_0$, and let $\mathcal{R} = \mathcal{T}_0 - \mathcal{S}$ be its multiset of additional copies. Since $\partial\mathcal{S} + \partial\mathcal{R} = J_0$ coordinatewise, every pair occurs at most b times in either $\mathcal{S}$ or $\mathcal{R}$. If $\deg_{\mathcal{R}}(v)$ counts triples with multiplicity, then

$$\deg_{\mathcal{R}}(v) \le \deg_{\mathcal{T}_0}(v) = \frac{b\deg_G(v) - \deg_D(v)}{2} \le \frac{b(2q+1)}{2} \le 2N.$$

The last inequality uses $b \le 5$ and $N = 3q + 2$.

A.4 Redistributing duplicate demands

Partition the multiset $\mathcal{R}$ into m classes $\mathcal{R}_y$, $y \in Y$, so that each class has maximum vertex degree at most $K = 1000$. A greedy assignment works: at an endpoint of a new triple at most $2N/K$ classes are full, so at most $6N/K$ classes are forbidden by its three endpoints. For sufficiently large q, $m \ge N/25 > 6N/K$. The first inequality already holds for $q \ge 20$, since $b \le 5$, and the second follows from $K = 1000$. Assign the triple to any remaining class.

Let $\delta_y(e)$ be the pair-incidence vector of $\mathcal{R}_y$. Then

$$\sum_{w \ne v} \delta_y(vw) \le 2K, \qquad 0 \le \sum_y \delta_y(e) \le b. \tag{78}$$

Indeed, the first left-hand side is $2\deg_{\mathcal{R}_y}(v)$, and the second is $\partial\mathcal{R}(e) \le J_0(e)$. Each δ_y is itself a sum of triangle vectors. The intended pair demand at label y will be $6\mathbf{1}_{E(G)} + \delta_y$.

Use four roles, the first three in $X = V(G)$ and the last in Y, and let Φ contain precisely the injective partial maps whose X-images are cliques in G. Take

$$H_1 = \{123, 124, 134, 234\}, \qquad H_2 = \{123\}.$$

Give an all-X triangle T host capacity one if $T \notin \mathcal{S}$ and zero if $T \in \mathcal{S}$; give the mixed coordinate $\{u, v, y\}$ capacity $6 + \delta_y(uv)$ for every edge $uv \in G$. These capacities are at most eleven because $\delta_y(uv) \le b \le 5$.

For each y, the vector $6\mathbf{1}_{E(G)} + \delta_y$ is triangle-divisible: its vertex degrees are even and its total weight is divisible by three, while δ_y is itself a sum of triangle vectors. Explicitly, its degree at v is $6\deg_G(v) + 2\deg_{\mathcal{R}_y}(v)$ and its total weight is $6E + 3|\mathcal{R}_y|$. It therefore lies in the triangle lattice by Theorem 3.4. Lift its signed representation to H_1 copies with label y, then correct all-X coordinates with signed H_2 copies, as in the explicit lattice construction in Lemma 3.7. This proves integer-span membership for the typed host. Again, $\mathcal{S}$ remains in the ambient complex for this signed calculation.

For every admissible ordered H_1 embedding with underlying triangle T and label y, use weight

$$w_{T,y} = \frac{6 + \sum_{e \in \binom{T}{2}} \delta_y(e)}{6p^2 N}.$$

There are six orderings for every (T, y). For a mixed coordinate $e = \{u, v\}$ at label y, the coverage is

$$\frac{(6 + \delta_y(e))|\{w \in N_G(u) \cap N_G(v) : \{u, v, w\} \notin \mathcal{S}\}|}{p^2 N} + O(K/N).$$

At most b common-neighbour choices are excluded by $\mathcal{S}$, and (78) bounds the remaining cross terms by $4K/(p^2N)$. The coverage is therefore $(6+\delta_y(e))(1+O(\tau+N^{-1}))$, uniformly. For an available all-X triangle T, the total H_1 coverage is exactly

$$\alpha_T = \frac{6m + \sum_{e\in\binom{T}{2}} \sum_y \delta_y(e)}{p^2N} = \frac{3}{4} + o(1)$$

uniformly in T, since the double sum lies between zero and $3b$. More explicitly,

$$p^2N = 8m + \frac{4b}{3} + \frac{8}{9}, \qquad 0 \le \sum_{e\in\binom{T}{2}} \sum_y \delta_y(e) \le 3b \le 15,$$

so $\alpha_T = 3/4 + O(N^{-1})$ uniformly and $1-\alpha_T$ is bounded away from zero for all sufficiently large N. Weight $(1-\alpha_T)/(6m)$ on every ordered H_2 embedding of T makes this coordinate's coverage exactly one. Both kinds of weights are uniformly positive of order $(N+m)^{-1}$ for large N.

The support is exactly the one in Lemma 3.6 with $j=3$, the clique complex of G, and the bounded pair-degree family $\mathcal{S}$. That lemma proves exact adaptation and both extension conditions. Its linear-size label hypothesis holds because $|Y| = m$ and $m/N \to 1/18$. Choose ω, then the required accuracy, then the initial graph accuracy and finally N, exactly as in Lemma 3.7. The fixed constants here, including K, do not depend on N. Theorem 3.5, now applied at uniformity three to the two motifs on four roles, decomposes the typed host. Again, the number of roles is fixed independently of the petal parameter.

Retain the all-X triangle from each selected H_1 and adjoin $\mathcal{S}$. Capacity one makes the result $\mathcal{F}$ simple. For each pair e and label y, exact filling of the mixed coordinate (e,y) counts precisely the retained H_1 triangles containing e. Its pair-incidence vector is therefore exactly

$$\begin{aligned}\partial\mathcal{F} &= 6m\mathbf{1}_{E(G)} + \sum_y \delta_y + \partial\mathcal{S}\\ &= 6m\mathbf{1}_{E(G)} + \partial\mathcal{R} + \partial\mathcal{S}\\ &= 6m\mathbf{1}_{E(G)} + J_0 = q\mathbf{1}_{E(G)} - \mathbf{1}_{E(D)}.\end{aligned}$$

In particular,

$$|\mathcal{F}| = \frac{qE-a-\ell}{3} = \left\lfloor \frac{qE-a}{3} \right\rfloor = \left\lfloor \frac{N\mathcal{E}_q - (q-1)/2}{3} \right\rfloor = U(r),$$

where $qE - a = N\mathcal{E}_q - (q-1)/2$ by direct substitution. The equality with $U(r)$ is the polynomial expansion displayed in the upper-bound calculation above. The ground set is X, not $X\cup Y$, and its size is $3r-1$. The pair-shadow is contained in G, so its maximum degree is at most $2q+1 = 2r-1$, while every pair-codegree is at most $q = r-1$. An empty core would require $3r$ vertices, a singleton core would require $2r$ shadow neighbours, and a pair core would require r triples on that pair. These three proper core sizes are excluded by the preceding strict bounds; a core of size three would force identical triples and is impossible because $\mathcal{F}$ is simple. Together with (75), this proves Theorem 5.1.

Use of AI: We prepared a list of proof strategies and techniques used in several papers related to sunflower conjectures and related problems in extremal combinatorics. We fed this to multiple AI tools (Claude, ChatGPT and two open-weight models with open-source harnesses) and prompted these AI tools to improve the lower bounds. When an AI tool was hallucinating, stuck, or going in the wrong direction, we stopped that AI tool, corrected the intermediate `.md` files, added more specific new ideas and switched to a different AI tool with more detailed fresh prompts. Equations written by one AI tool were critiqued, verified and expanded by a completely different AI tool in an iterative manner, till we were satisfied with their correctness and details. The author is solely responsible for the correctness of the proofs.

References

[1] H. L. Abbott, *Some remarks on a combinatorial theorem of Erdős and Rado*, Canadian Mathematical Bulletin **9** (1966), 155–160. doi:10.4153/CMB-1966-019-1.

[2] H. L. Abbott and G. Exoo, *On set systems not containing delta systems*, Graphs and Combinatorics **8** (1992), 1–9. doi:10.1007/BF01271703.

[3] H. L. Abbott, D. Hanson and N. Sauer, *Intersection theorems for systems of sets*, Journal of Combinatorial Theory, Series A **12** (1972), 381–389. doi:10.1016/0097-3165(72)90103-3.

[4] R. Alweiss, S. Lovett, K. Wu and J. Zhang, *Improved bounds for the sunflower lemma*, Annals of Mathematics **194** (2021), 795–815. doi:10.4007/annals.2021.194.3.5.

[5] E. Axante, C. Budala, D. Chitic, B. Dumitru and M. Nacu, *Sunflower-free uniform families: recursive constructions and explicit bounds*, arXiv:2609.06175v1 (2026). arXiv record.

[6] T. Bell, S. Chueluecha and L. Warnke, *Note on sunflowers*, Discrete Mathematics **344** (2021), article 112367. doi:10.1016/j.disc.2021.112367.

[7] P. Bennett and A. Priestley, *The sunflower-free process*, arXiv:2509.16355 (2025). arXiv record.

[8] V. Chvátal and D. Hanson, *Degrees and matchings*, Journal of Combinatorial Theory, Series B **20** (1976), 128–138. doi:10.1016/0095-8956(76)90004-6.

[9] P. Erdős and R. Rado, *Intersection theorems for systems of sets*, Journal of the London Mathematical Society **35** (1960), 85–90. doi:10.1112/jlms/s1-35.1.85.

[10] P. Frankl and V. Rödl, *Near perfect coverings in graphs and hypergraphs*, European Journal of Combinatorics **6** (1985), 317–326. doi:10.1016/S0195-6698(85)80045-7.

[11] K. Frankston, J. Kahn, B. Narayanan and J. Park, *Thresholds versus fractional expectation-thresholds*, Annals of Mathematics **194** (2021), 475–495. doi:10.4007/annals.2021.194.2.2.

[12] P. Frankl and J. Wang, *New bounds on families without large sunflowers*, Electronic Journal of Combinatorics **32**(2) (2025), P2.43. doi:10.37236/13277.

[13] J. Fukuyama, *Sunflower bound with a sub-logarithmic base*, arXiv:2510.19037v2 (2025). arXiv record.

[14] P. Keevash, *The existence of designs II*, arXiv:1802.05900 (2018). arXiv record; author-hosted version. The theorem numbering used here is that of the version dated 16 February 2018.

[15] A. V. Kostochka, *Extremal problems on Δ-systems*, survey, author-hosted version dated 25 October 2011. Author-hosted text.

[16] A. V. Kostochka, *A bound of the cardinality of families not containing* Δ*-systems*, in *The Mathematics of Paul Erdős II*, Algorithms and Combinatorics **14**, Springer, 1997, 229–235. doi:10.1007/978-3-642-60406-5_19.

[17] A. V. Kostochka, V. Rödl and L. A. Talysheva, *On systems of small sets with no large* Δ*-subsystems*, Combinatorics, Probability and Computing **8** (1999), 265–268. Publisher text.

[18] E. Naslund and W. Sawin, *Upper bounds for sunflower-free sets*, Forum of Mathematics, Sigma **5** (2017), e15. doi:10.1017/fms.2017.12.

[19] A. Rao, *Coding for sunflowers*, Discrete Analysis (2020), paper 2. doi:10.19086/da.11887.

[20] J. Spencer, *Intersection theorems for systems of sets*, Canadian Mathematical Bulletin **20** (1977), 249–254. doi:10.4153/CMB-1977-038-7.

[21] N. Bansal, D. Dadush, S. Garg, and S. Lovett, *The Gram–Schmidt Walk: A Cure for the Banaszczyk Blues*, Theory of Computing **15** (2019), article 21, 1–27. doi:10.4086/toc.2019.v015a021. Theorem 1.1 states the Banaszczyk vector-balancing theorem with constant 5.