Counting Clique-Containing Graphs by Transitive Block Weighting

Brian Klemm
Independent researcher
brian@brianklemm.com

August 2026

Abstract

Let $\mathcal{C}(n,c,k)$ be the set of labeled graphs on n vertices with exactly k edges containing at least one c-clique. We give an exact evaluation of $|\mathcal{C}(n,c,k)|$ that enumerates a single *block* — the k-edge graphs in which one fixed c-clique is present — rather than the full space of $\binom{\binom{n}{2}}{k}$ graphs. Writing h_r for the number of block graphs containing exactly r cliques in total, $|\mathcal{C}| = \binom{n}{c}\sum_{r\geq 1} h_r/r$.

The underlying identity is not new: it is the union-of-sets coverage identity of Karp and Luby, which expresses the size of a union as a sum of reciprocal-multiplicity weights. Our contribution is to isolate the structural hypothesis that makes it exactly evaluable. We prove an orbit-collapse proposition — if a group acts on the index set of the family compatibly with the sets themselves, the outer sum reduces to one term per orbit, weighted by orbit size — and observe that $\mathrm{Sym}(n)$ acting on the c-subsets of $[n]$ gives a single orbit. An identity normally estimated by Monte Carlo therefore becomes computable exactly by finite enumeration. Two further consequences follow at no cost: a single block determines not just $|\mathcal{C}|$ but the entire distribution of clique counts, via $H_r = \binom{n}{c}h_r/r$; and a second application, to monochromatic cliques in two-colourings, needs the orbit form essentially, having two orbits whose contributions differ.

We give the double-counting proof and an accounting of cost that we have tried to keep free of favourable comparisons. The saving over exhaustive enumeration is a factor of $\prod_{i<\binom{c}{2}}(N-i)/(k-i)$ with $N=\binom{n}{2}$, not the $\binom{n}{c}$ one might expect; the method is not a complexity-class improvement; sampling the same identity yields an approximation scheme polynomial in n and is preferable wherever exactness is not required; and a third method — enumerating the K_c-free graphs and subtracting — is cheaper than block weighting in an identifiable regime, which we measure and report rather than omit.

The case for exact evaluation is made computationally. We report 156 triples (n,c,k) on which the method agrees with independent exhaustive enumeration and 4 not attempted, with no disagreements; agreement with two further independent implementations; complete integer tables of $|\mathcal{C}(n,c,k)|$ for $c\in\{3,4\}$ and $n\leq 8$, for which we have found no published predecessor; and two external checks, in which the saturation threshold recovered by the counter reproduces the Turán number $\mathrm{ex}(n,K_c)$, and the two-colour count reproduces $R(3,3)=6$. All arithmetic is performed over $\mathbb{Q}$; no floating-point operation occurs in the counting path.

Finally we remove the enumeration in the regime where that is possible. For fixed c and k, decomposition by non-isolated vertex set makes $H_r(n)$ a polynomial in n, evaluable in microseconds at any n; this is an elementary observation rather than a theorem, and we say so. Regrading the same decomposition by *deficiency* $2k-v$, in the manner of Wright, restores a finite table at every order once k grows with n, and yields an exact expansion

whose truncation at order D is a certified lower bound, with relative error measured to fall as $(k^2/n)^{D+1}$. We give the exact leading ratio $2k(k-1)/(n-2k+1)$, which identifies the expansion parameter as k^2/n and so confines convergence to $k = o(\sqrt{n})$ — narrower than the $o(n)$ regime of the Chen–Stein approximation of DeSalvo and Rombach, and inside it. We report that limitation as the principal negative result of the paper rather than as future work.

1 Introduction

Counting graphs that contain a prescribed substructure is a basic problem of enumerative and extremal combinatorics. The complementary quantity treated here — the number of K_c-free labeled graphs — has an extensive asymptotic literature beginning with Erdős, Kleitman and Rothschild [7] and continuing through [6, 15, 22, 20, 2, 26, 3, 10]. That literature is asymptotic throughout; no exact formula is known, and exact values at fixed (n, c, k) are obtained, when they are obtained at all, by enumeration.

Direct enumeration becomes impractical quickly. Deciding membership in $\mathcal{C}(n, c, k)$ for each of the $\binom{\binom{n}{2}}{k}$ candidate graphs means classifying about 9×10^9 graphs in the largest layer at $n = 9$ — on the order of CPU-days at the throughput we measure in Section 6, and worse than exponentially so beyond. Inclusion–exclusion over the $\binom{n}{c}$ cliques replaces this with an alternating sum over subsets of cliques whose terms depend on the edge-set union of each subset, and the number of distinct overlap patterns grows quickly with c.

This paper takes a third route. Fix one c-clique Q and consider only the graphs in which Q is present — the *block* $\mathcal{B}_Q$. A graph containing r cliques lies in exactly r blocks, so summing block cardinalities overcounts each graph by a known factor; weighting each graph by $1/r$ removes the overcount exactly, in one pass, with no alternating signs. Because $\mathrm{Sym}(n)$ is transitive on c-subsets, every block yields the same answer, and one block suffices.

What is and is not new here

The weighting identity itself is due to Karp and Luby [12], refined in [13]; it is standard in approximate counting, where it underlies the coverage algorithm for DNF counting and network reliability. Stated in generality: for finite sets $A_1, \ldots, A_m$ and $d(x) = |\{i : x \in A_i\}|$,

$$\left|\bigcup_{i=1}^{m} A_i\right| = \sum_{i=1}^{m} \sum_{x \in A_i} \frac{1}{d(x)}. \tag{1}$$

We claim no priority for (1). What we contribute is the following.

1. **The orbit collapse, formalised.** Proposition 5 gives the hypothesis under which the outer sum of (1) reduces to one inner sum per orbit: a group acting on the index set compatibly with the family. Under transitivity that is a single sum, and the identity, normally estimated by Monte Carlo, becomes evaluable exactly by finite enumeration. Clique-containing graphs are the transitive case (Lemma 7, Theorem 8); monochromatic cliques in two-colourings are a two-orbit case where the transitive form would give the wrong answer (Section 7).

2. **The whole distribution, not just the count.** A single block determines H_r, the number of k-edge graphs with exactly r cliques, for every r — Corollary 10 — of which $|\mathcal{C}|$ is only the $r \geq 1$ marginal. The same corollary read backwards is a congruence constraint on h_r.

3. **Exact integer tables.** We give complete values of $|\mathcal{C}(n,c,k)|$ for $c \in \{3,4\}$ and all $n \le 8$ (Section 6), verified against exhaustive enumeration and against two further independent implementations. We have not found these tabulated in the literature. We stress that this is a claim about our search and not about difficulty: the tables are within reach of any of several methods, including one we describe in Section 4.5 that is cheaper than ours in part of the range.

4. **A finite grading where the obvious one fails.** Once k grows with n the vertex-set decomposition of Section 8 has no finite table. Grading by deficiency $2k - v$ instead does (Section 9), giving an exact expansion with certified truncation error. The grading is Wright's [29, 30] and we claim no priority for it; what is new is its use on H_r, and the exact ratio that identifies k^2/n as the expansion parameter — which also shows the window is narrower than the approximation literature's and lies inside it.

5. **An accounting without a straw man.** We give the saving against exhaustive enumeration, note that it is *not* $\binom{n}{c}$, state that the method is not a complexity-class improvement, and benchmark against the two competitors that actually matter: sampling the same identity (Section 4.4) and enumerating the complement (Section 4.5). Both beat block weighting in identifiable regimes, and we report where.

2 Preliminaries

Throughout, $n \ge c \ge 2$ are integers, $N = \binom{n}{2}$, $e_c = \binom{c}{2}$, and $m = \binom{n}{c}$. Graphs are simple, undirected and labeled on the vertex set $[n] = \{1, \dots, n\}$; we identify a graph with its edge set, a subset of $E(K_n)$. The edge count k is fixed throughout Sections 3, with $0 \le k \le N$.

Definition 1. $\mathcal{G}(n,k) = \{\, g \subseteq E(K_n) : |g| = k \,\}$, so $|\mathcal{G}(n,k)| = \binom{N}{k}$. For a c-subset $Q \subseteq [n]$ write K_Q for the set of $\binom{c}{2}$ edges spanned by Q. Say Q is *active* in g if $K_Q \subseteq g$.

Definition 2. $\mathcal{C}(n,c,k) = \{\, g \in \mathcal{G}(n,k) : \text{some } c\text{-subset is active in } g \,\}$, and for $g \in \mathcal{G}(n,k)$,

$$r(g) \;=\; \left|\{\, Q \in \tbinom{[n]}{c} : K_Q \subseteq g \,\}\right|.$$

Thus $\mathcal{C}(n,c,k) = \{g : r(g) \ge 1\}$.

Definition 3 (Block)**.** For a c-subset Q and the fixed k, the *block* of Q is

$$\mathcal{B}_Q \;=\; \{\, g \in \mathcal{G}(n,k) : K_Q \subseteq g \,\}.$$

Each $g \in \mathcal{B}_Q$ is determined by its remaining $k - e_c$ edges, chosen freely from the $N - e_c$ edges outside K_Q, so

$$|\mathcal{B}_Q| \;=\; \binom{N - e_c}{k - e_c} \tag{2}$$

for $k \ge e_c$, and $\mathcal{B}_Q = \emptyset$ otherwise.

Remark 4. Equation (2) counts the graphs containing *one fixed* clique. It is not $|\mathcal{C}(n,c,k)|$ and is in general much smaller: at $n = 6$, $c = 3$, $k = 9$ the block holds $\binom{12}{6} = 924$ graphs while $|\mathcal{C}| = 4995$. The two are easy to conflate.

3 The block-weighted identity

We first isolate the structure that makes (1) collapse, then apply it.

Proposition 5 (Orbit collapse)**.** *Let a finite group G act on a finite set X and on a finite index set I, and let $(A_i)_{i\in I}$ be a family of subsets of X satisfying the compatibility condition*

$$\sigma(A_i) \;=\; A_{\sigma(i)} \qquad \text{for all } \sigma \in G,\ i \in I. \tag{3}$$

Write $d(x) = |\{i \in I : x \in A_i\}|$, let $I = O_1 \sqcup \cdots \sqcup O_t$ be the decomposition of I into G-orbits, and pick a representative $i_j \in O_j$ for each. Then

$$\Big|\bigcup_{i\in I} A_i\Big| \;=\; \sum_{j=1}^{t} |O_j| \cdot \sum_{x\in A_{i_j}} \frac{1}{d(x)}.$$

In particular, if G acts transitively on I then $t = 1$ and the union is $|I|$ times a single inner sum.

Proof. d is G-invariant. Fix $\sigma \in G$ and $x \in X$. By (3), $x \in A_i$ if and only if $\sigma(x) \in \sigma(A_i) = A_{\sigma(i)}$. Since $i \mapsto \sigma(i)$ is a bijection of I, it restricts to a bijection $\{i : x \in A_i\} \to \{i : \sigma(x) \in A_i\}$, so $d(\sigma(x)) = d(x)$.

Inner sums agree within an orbit. Let $i, i' \in O_j$. Choose σ with $\sigma(i) = i'$. By (3), σ restricts to a bijection $A_i \to A_{i'}$, and d is preserved, so $\sum_{x\in A_i} 1/d(x) = \sum_{x\in A_{i'}} 1/d(x)$. Across orbits the inner sums need not agree, and in general do not; see Section 7.

Double counting. Let $\mathcal{I} = \{(x,i) : i \in I,\ x \in A_i\}$ and $S = \sum_{(x,i)\in\mathcal{I}} 1/d(x)$; every summand is defined because $(x,i) \in \mathcal{I}$ forces $d(x) \geq 1$. Grouping by x, each $x \in \bigcup_i A_i$ occurs in exactly $d(x)$ pairs, each contributing $1/d(x)$, and each $x \notin \bigcup_i A_i$ occurs in none; hence $S = |\bigcup_i A_i|$. Grouping by i, then by orbit, and using the previous paragraph, $S = \sum_j |O_j| \sum_{x\in A_{i_j}} 1/d(x)$. □

Remark 6. Condition (3) is what does the work and is easy to omit: without it, transitivity on I alone says nothing about the sets. The orbit form also guards against a tempting error. Faced with a family that is *almost* symmetric, one may assume transitivity, collapse to a single inner sum and scale by $|I|$. When $t > 1$ that is simply wrong, because the orbits carry different inner sums; Section 7 exhibits a natural family where $t = 2$ and the two sums differ by a factor of more than four.

Note also that Proposition 5 is not Burnside's lemma. Burnside counts *orbits* of a group acting on the objects being enumerated, via a sum of fixed-point counts. Here the group acts on the index set of the family, the objects counted are individual elements of X, and no fixed-point sum appears.

Lemma 7 (Blocks satisfy the hypothesis)**.** *Take $X = \mathcal{G}(n,k)$, $I = \binom{[n]}{c}$, $G = \mathrm{Sym}(n)$ acting on edges by $\sigma(\{u,v\}) = \{\sigma(u), \sigma(v)\}$, on graphs elementwise, and on I by direct image. Then $\sigma(\mathcal{B}_Q) = \mathcal{B}_{\sigma(Q)}$ for all σ and Q, the action of* $\mathrm{Sym}(n)$ *on I is transitive, and $d = r$.*

Proof. The action on edges is a bijection of $E(K_n)$, hence induces a bijection of $\mathcal{G}(n,k)$ preserving edge count. For any c-subset R we have $\sigma(K_R) = K_{\sigma(R)}$, since σ maps the pairs within R onto the pairs within $\sigma(R)$. Hence $K_Q \subseteq g$ iff $K_{\sigma(Q)} \subseteq \sigma(g)$, which gives both $\sigma(\mathcal{B}_Q) \subseteq \mathcal{B}_{\sigma(Q)}$ and, applying the same to σ^{-1}, equality. Transitivity on c-subsets is standard: any bijection between two c-subsets extends to a permutation of $[n]$. Finally $d(g) = |\{Q : g \in \mathcal{B}_Q\}| = |\{Q : K_Q \subseteq g\}| = r(g)$. □

Theorem 8 (Block weighting)**.** *Let $n \geq c \geq 2$ and $0 \leq k \leq N$. Then*

$$|\mathcal{C}(n,c,k)| \;=\; \binom{n}{c} \sum_{r \geq 1} \frac{h_r}{r},$$

where $h_r = |\{\, g \in \mathcal{B}_{Q_0} : r(g) = r \,\}|$ for any fixed c-subset Q_0; the value of h_r does not depend on the choice of Q_0.

Proof. For $k < e_c$ no graph with k edges can contain the e_c edges of a c-clique, so $\mathcal{C}(n,c,k) = \emptyset$, every $\mathcal{B}_Q$ is empty, every $h_r = 0$, and both sides are 0.

For $k \geq e_c$, apply Proposition 5 with the data of Lemma 7. Then $\bigcup_Q \mathcal{B}_Q$ is precisely the set of k-edge graphs in which some c-subset is active, that is $\mathcal{C}(n,c,k)$, and $|I| = \binom{n}{c}$, giving

$$|\mathcal{C}(n,c,k)| \;=\; \binom{n}{c} \sum_{g \in \mathcal{B}_{Q_0}} \frac{1}{r(g)}.$$

Grouping $\mathcal{B}_{Q_0}$ by the value of r turns the inner sum into $\sum_{r \geq 1} h_r/r$; the sum starts at $r = 1$ because Q_0 is itself active in every $g \in \mathcal{B}_{Q_0}$, so $h_0 = 0$. Independence of h_r from Q_0 follows from the bijection $\mathcal{B}_{Q_0} \to \mathcal{B}_Q$ of Proposition 5's second step, which preserves r. □

Corollary 9 (Integrality)**.** $\binom{n}{c} \sum_{r \geq 1} h_r/r \in \mathbb{Z}$.

Immediate, since the left side equals a cardinality. It is nonetheless a nontrivial arithmetic constraint — the inner sum is usually not an integer, for instance $\sum_r h_r/r = 689/4$ at $n = 6$, $c = 3$, $k = 6$ — and it makes a useful runtime assertion. Our implementation computes the sum in exact rational arithmetic and raises if the product fails to have denominator 1, so many implementation errors surface as an error rather than as a plausible rounded integer. Errors that happen to preserve integrality would of course pass.

3.1 The block determines the whole distribution

Theorem 8 extracts a single number from the histogram. The histogram carries strictly more, and at no additional cost.

Corollary 10. *For $r \geq 1$ let H_r be the number of graphs in $\mathcal{G}(n,k)$ containing exactly r c-cliques, counted over the whole space. Then*

$$H_r \;=\; \frac{\binom{n}{c}\, h_r}{r},$$

and consequently $|\mathcal{C}(n,c,k)| = \sum_{r \geq 1} H_r$ while $H_0 = \binom{N}{k} - |\mathcal{C}(n,c,k)|$.

Proof. Count the pairs (g, Q) with $r(g) = r$ and Q active in g in two ways. By g: there are H_r such graphs, each in exactly r pairs, giving rH_r. By Q: there are $\binom{n}{c}$ choices, and each block contains h_r graphs with $r(g) = r$ — the same h_r for every Q, by the block-independence established in Theorem 8 — giving $\binom{n}{c} h_r$. Hence $rH_r = \binom{n}{c} h_r$. □

So one block determines the entire distribution of clique counts, of which $|\mathcal{C}|$ is only the $r \geq 1$ marginal. We record h_r and H_r for every computed triple alongside the counts; a reader interested in the distribution need not recompute anything.

Corollary 11. $\operatorname{denom}\big(\sum_{r \geq 1} h_r/r\big)$ *divides* $\binom{n}{c}$.

Proof. $\binom{n}{c}\sum_r h_r/r = |\mathcal{C}| \in \mathbb{Z}$ by Theorem 8. □

Read forwards this is the runtime assertion of Corollary 9. Read backwards it is a congruence constraint on the histogram itself, and therefore a necessary condition on any closed form for h_r — which is the object of the open question in Section 11.

Remark 12 (The support need not be an interval)**.** The set of realisable values of r has interior gaps in 74 of the 160 triples we computed. The smallest instance is the clearest: on six vertices with six edges a graph may contain one, two or four triangles, but never exactly three, the jump from two to four being the arrival of a K_4. The full distribution there is $H_0 = 1560$, $H_1 = 2520$, $H_2 = 910$, $H_4 = 15$. Table 1 gives further cases.

Two distinct questions are nearby, and we claim neither. The *least* attainable clique count at prescribed order and size is the Erdős–Rademacher problem: Razborov [23] determined the asymptotic triangle-density profile, and Liu, Pikhurko and Staden [17] gave the exact minimum for all sufficiently large n. Which counts are attainable — the support, as opposed to its least element — is the question the gaps speak to, and we record them as an observation about computed data rather than as a result.

4 Algorithm and cost

4.1 The algorithm

Input: n, c, k
$m \leftarrow \binom{n}{c}$; $e_c \leftarrow \binom{c}{2}$; $h_r \leftarrow 0$ for all r; fix a base clique Q_0
$\mathcal{Q}' \leftarrow \{\, Q : |K_Q \setminus K_{Q_0}| \le k - e_c \,\}$ *(cliques that can possibly be active)*
for each $(k-e_c)$-subset F of $E(K_n) \setminus K_{Q_0}$:
  $g \leftarrow K_{Q_0} \cup F$
  $r \leftarrow |\{Q \in \mathcal{Q}' : K_Q \subseteq g\}|$
  $h_r \leftarrow h_r + 1$
return $m \cdot \sum_{r\ge 1} h_r/r$ (exact rational arithmetic)

The loop runs $\binom{N-e_c}{k-e_c}$ times; each iteration assembles g in $O(k)$ and tests $m' = |\mathcal{Q}'|$ cliques at $O(e_c)$ each, giving

$$O\left(\binom{N-e_c}{k-e_c} \cdot \big(k + m' e_c\big)\right), \qquad m' \le m.$$

The pruning to $\mathcal{Q}'$ is sound because every block graph has exactly $k - e_c$ edges outside K_{Q_0}, so a clique needing more than that many cannot be active. It is not merely cosmetic: at $n = 8$, $c = 4$ it removes most of the per-graph work, and it is why the measured ratio in Section 6 exceeds the structural prediction there.

4.2 The saving is not $\binom{n}{c}$

It is tempting to read Theorem 8 as buying a factor of $m = \binom{n}{c}$, since one block replaces m of them. That is the saving relative to the *naive symmetric* method — enumerate every block and deduplicate — which is a different and much weaker baseline. Against exhaustive enumeration the ratio of objects examined is

$$\frac{\binom{N}{k}}{\binom{N-e_c}{k-e_c}} \;=\; \prod_{i=0}^{e_c-1} \frac{N-i}{k-i} \;\approx\; \left(\frac{N}{k}\right)^{\binom{c}{2}}, \tag{4}$$

the approximation holding when $e_c \ll k$. The two expressions are not equal as functions of (n, c, k) — they coincide only in degenerate cases, such as $c = 2$, $k = 1$, where both equal N — and (4) behaves quite differently across the range of k:

- For $k = \Theta(1)$ the ratio is $\Theta(n^{2e_c})$ — very large.
- For $c = 3$ at the triangle threshold $k = \Theta(n^{3/2})$ it is $\Theta(n^{3/2})$, much smaller than $\binom{n}{3} = \Theta(n^3)$.
- For $k = \alpha N$ with α constant it degenerates to α^{-e_c}, a constant: in the dense regime the method buys essentially nothing against exhaustive enumeration — and, as Section 4.5 shows, is beaten outright by a different method.

4.3 This is not a complexity-class improvement

Since $m \cdot |\mathcal{B}_Q| \geq |\mathcal{C}|$, the work performed is at least $|\mathcal{C}|/m$, and $|\mathcal{C}|$ is exponential in k over the relevant range. The method is a polynomial-factor speedup on an exponential exact algorithm, and no more. We state this explicitly because (4) is large enough at small k to be mistaken for something stronger.

4.4 Competitor I: sampling the same identity

If an approximation suffices, one should not enumerate the block at all. Sample g uniformly from $\mathcal{B}_{Q_0}$ — choose $k - e_c$ of $N - e_c$ edge slots — set $X = 1/r(g)$, and average over independent samples; $m \cdot |\mathcal{B}_{Q_0}| \cdot \mathbb{E}[X]$ is exactly $|\mathcal{C}|$ by Theorem 8, so the empirical mean scaled the same way is unbiased.

For the variance, note $0 < X \leq 1$, so $\mathrm{Var}(X) \leq \mathbb{E}[X^2] \leq \mathbb{E}[X]$ and hence

$$\frac{\mathrm{Var}(X)}{\mathbb{E}[X]^2} \;\leq\; \frac{1}{\mathbb{E}[X]} \;=\; \frac{m\,|\mathcal{B}_{Q_0}|}{|\mathcal{C}|} \;\leq\; m,$$

the last step because $\mathcal{B}_{Q_0} \subseteq \mathcal{C}$. The argument of Karp, Luby and Madras [13] then gives an approximation scheme using $O(m\,\varepsilon^{-2} \log \delta^{-1})$ samples — polynomial in n for fixed c. (We say "polynomial in n" rather than "fully polynomial" because with (n, c, k) encoded in binary the output alone has $\Theta(n^2 \log n)$ bits.) Whenever k grows with n this is exponentially cheaper than enumerating the block; for $k = \Theta(1)$ the block is already of polynomial size and exactness is free.

We regard this as a central limitation of exhaustive block enumeration and address it in Section 5 rather than leaving it to the reader. For $c = 3$ specifically, approximate counting in the neighbouring regime is available by other means [11].

4.5 Competitor II: enumerating the complement

Because $|\mathcal{C}(n, c, k)| = \binom{N}{k} - |\{K_c\text{-free graphs with } k \text{ edges}\}|$, any method that enumerates the K_c-free graphs answers the same question exactly. One does this by a vertex-by-vertex search:

place vertices $1, \dots, n$ in order, choosing at each step a neighbourhood among the already-placed vertices, and reject a neighbourhood exactly when it contains a $(c-1)$-clique of the graph so far, since that and the new vertex would form a c-clique. Every leaf is a distinct K_c-free labeled graph, tallied by edge count, and the whole column — every k at once — comes out in one pass. This is the same object as "counting independent sets in the c-clique hypergraph on $\binom{n}{2}$ vertices", which is the natural framing of the complement.

The two exact methods have incomparable costs. Block weighting visits $\binom{N-e_c}{k-e_c}$ objects for a single k, hence $\sum_k \binom{N-e_c}{k-e_c} = 2^{N-e_c}$ for a whole column, and is completely insensitive to how many graphs actually contain a clique. Complementary enumeration visits exactly as many objects as there are K_c-free graphs. So the complement wins precisely when clique-containing graphs are abundant — large c-density, k near or above the Turán threshold — and loses badly when they are rare.

Table 2 measures both, and the crossover is real. At $n = 8$, $c = 3$ the complement enumerates 4,682,270 objects against $2^{25} = 33{,}554{,}432$ for block weighting and finishes in roughly a third of the time. At $c = 4$ the ordering reverses by a wide margin: at $n = 8$ there are 147,141,138 K_4-free graphs against $2^{22} = 4{,}194{,}304$ block objects, and block weighting is faster by two orders of magnitude. The governing quantity is simply how many graphs are K_c-free relative to 2^{N-e_c}, and that ratio moves sharply with c. Any claim that block weighting is "the" exact method for this problem would be false, and we do not make it.

4.6 Competitor III: inclusion–exclusion by overlap signature

For completeness we record what inclusion–exclusion gives, since avoiding it is one of the method's stated selling points. The terms group by a signature.

Lemma 13. *Let $Q_1 \neq Q_2$ be c-subsets of $[n]$ sharing exactly s vertices. Then $|K_{Q_1} \cup K_{Q_2}| = u(s)$ where*

$$u(s) \;=\; 2\binom{c}{2} - \binom{s}{2},$$

and the number of k-edge graphs containing both cliques is $\binom{N-u(s)}{k-u(s)}$. The number of unordered pairs with signature s is $P(s) = \binom{n}{c}\binom{c}{s}\binom{n-c}{c-s}/2$.

Proof. K_{Q_1} and K_{Q_2} each have $\binom{c}{2}$ edges, and an edge lies in both exactly when both its endpoints lie in $Q_1 \cap Q_2$, so they share $\binom{s}{2}$ edges; inclusion–exclusion on the two edge sets gives $u(s)$. Fixing those $u(s)$ edges leaves $k - u(s)$ to choose freely from the remaining $N - u(s)$. For $P(s)$, choose Q_1 in $\binom{n}{c}$ ways, the s shared vertices within it in $\binom{c}{s}$ ways, and the other $c - s$ vertices of Q_2 outside Q_1 in $\binom{n-c}{c-s}$ ways; each unordered pair arises twice. □

Because the terms depend only on s, the pairwise sum collapses to at most c terms regardless of n, and truncating after pairs gives a Bonferroni bracket

$$S_1 - S_2 \;\le\; |\mathcal{C}(n,c,k)| \;\le\; S_1, \qquad S_1 = m\binom{N-e_c}{k-e_c}, \quad S_2 = \sum_{s=0}^{c-1} P(s)\binom{N-u(s)}{k-u(s)},$$

computable with no enumeration at all. It is exact at $k = e_c$, where no two cliques fit, and degrades to vacuity once triple overlaps matter — at $n = 7$, $c = 3$ the lower bound is already 0 by $k = 10$. Carrying the expansion further needs the number of clique triples by union size, which we have not tabulated; that is the sense in which the question below is open.

The identity $u(0) = u(1)$ is worth flagging because it is a natural place to slip. One might expect the union size to decrease with each shared vertex; it does not, because sharing one vertex shares no edge. The correct statement is $\binom{s}{2}$, and $\binom{1}{2} = 0$.

5 Why compute exactly

Given Section 4.4, exact evaluation needs a reason. We offer three, all turning on the fact that an approximation with relative error ε cannot certify an integer. Note that these are arguments for exactness, not for block weighting specifically — complementary enumeration is equally exact, and in its regime one should use it.

Ground truth for asymptotic results. The asymptotic literature on K_c-free graphs [7, 2, 26, 10] gives estimates whose error terms are invisible at small n. Exact values are what one checks such results against in the range where the asymptotics are not yet informative. An approximation scheme's output cannot serve this purpose: it is correct to within a factor chosen in advance.

Integer sequences. The values of $|\mathcal{C}(n, c, k)|$ form a two-parameter integer array, which we give for $c \in \{3, 4\}$ and $n \leq 8$ in Section 6. When this work began, neither that array nor its complement — the K_c-free counts refined by edge number — appeared in the literature we surveyed or in the OEIS. They do now. The K_c-free arrays have been submitted and accepted as A396974 ($c = 3$) and A399108 ($c = 4$); their row sums are A213434, which already existed, and A395158, which did not and was added at an editor's suggestion. The ambient tables are A084546, the labeled $\binom{N}{k}$ array of which these are the K_c-free part, and A008406, its unlabeled analogue.

Two things emerged from that submission, and both bear on claims made elsewhere in this paper. The arrays were independently recomputed by OEIS editors before acceptance, and were then extended well past the range reachable by any method described here — A396974 to $n = 14$ and A399108 to $n = 12$ — by unlabeled generation with automorphism weighting, which we discuss in Section 11. Separately, the arrays turn out to be edge refinements of columns of A395695, the number of labeled graphs on n vertices with clique number at most k, which existed throughout, with A395684 as its exact-clique-number companion. Neither is refined by edge count, which is the respect in which these tables are new. But our search had been organised around triangle-free counts rather than around clique number, and a search organised the other way would have found them at once. We record that as a caution about how little a negative search establishes.

A table of this kind is exact by definition or it is not a table, and that is the argument for exactness here independently of priority.

Recovering exact extremal thresholds. We use the counter to locate the saturation threshold and compare it against a closed form obtained by entirely different means.

Lemma 14. *For every* $0 \leq k \leq \mathrm{ex}(n, K_c)$ *there is a* K_c*-free graph on* $[n]$ *with exactly* k *edges; for* $k > \mathrm{ex}(n, K_c)$ *there is none. Consequently* $\max\{\, k : \text{some } k\text{-edge graph on } [n] \text{ is } K_c\text{-free} \,\} = \mathrm{ex}(n, K_c)$.

Proof. The second claim is the definition of ex. For the first, the Turán graph $T(n, c-1)$ is K_c-free with exactly $\mathrm{ex}(n, K_c)$ edges [27]; for $c = 3$ this is Mantel's theorem [16]. Deleting an edge from a K_c-free graph leaves a K_c-free graph, since any subgraph of a K_c-free graph is K_c-free. Deleting edges of $T(n, c-1)$ one at a time therefore realises every edge count from $\mathrm{ex}(n, K_c)$ down to 0. □

Our counter reports the cliqueless count $\binom{N}{k} - |\mathcal{C}(n, c, k)|$ for each k; the saturation threshold is one less than the least k at which that count reaches zero, and by Lemma 14 it must equal $\mathrm{ex}(n, K_c)$. Because the count is exact, the threshold is exact — an estimate straddling zero determines nothing. The agreement in Table 5 is a genuine external check on the implementation: it tests the counter against a closed form derived independently of anything in this paper.

6 Computational results

All figures below were produced by the accompanying implementation and are reproduced in this document directly from its output files by a generator script, so no figure is hand-transcribed. The environment was: Intel(R) Xeon(R) Processor @ 2.10GHz, 2 cores, CPython 3.11.15, Linux-6.18.5-fc-v20-x86_64-with-glibc2.39. The exact invocations that produced each result file are recorded in the repository. Nothing in the method or the implementation is randomized; every count is a deterministic enumeration, no seed is involved, and repeated runs return identical values — the timing harness asserts this across repetitions.

6.1 Validation

Table 4 lists representative triples. Across the full sweep, 160 triples (n, c, k) were computed, of which 156 were checked against exhaustive enumeration; the remaining 4 were not attempted, exhaustive enumeration for those being a matter of hours to days of CPU rather than of impossibility. There were no disagreements.

Three independent implementations were used, agreeing everywhere they overlap:

- *Exhaustive enumeration*, classifying every k-edge graph. It shares the construction of the clique masks with the block counter — so a fault in that construction would be invisible to this comparison — but nothing else.
- *Complementary enumeration* (Section 4.5), which shares no edge indexing, no clique masks and no block structure with the counter, and which agrees with it on every k for all (n, c) in Table 2. This is the check that discharges the shared-code concern above.
- *A vectorised bitmask reimplementation* of the block method itself, agreeing with the reference path on 642 triples.

Three further checks were run and are reported in full in the accompanying test suite (539 tests across 6 files):

- *Block invariance.* For each of several (n, c, k), the count was computed from *every* one of the $\binom{n}{c}$ blocks, not a sample. All blocks agreed and their histograms were identical. This exercises the symmetry step of Proposition 5 — the step that licenses evaluating a single block — though we note it is computed by the same code in each case, so it tests the mask construction's symmetry rather than the proposition itself.

- *Published ground truth.* Every row of an independently produced brute-force dataset for $c = 3$, $n = 3, \dots, 7$ and all k was reproduced, including cliqueless counts and block cardinalities.
- *Exactness.* In every case the rational $\sum_r h_r/r$ multiplied by $\binom{n}{c}$ to an exact integer, as Corollary 9 requires, and its denominator divided $\binom{n}{c}$ as Corollary 11 requires.
- *The distribution.* $H_r = \binom{n}{c} h_r/r$ was checked against direct classification of the whole space on every triple where that was feasible, and $\sum_{r \geq 1} H_r = |\mathcal{C}|$ throughout.
- *The two-colour count.* Every per-k value of $|\mathcal{M}(n, c, k)|$ for $n \leq 6$ was checked against enumeration of each graph together with its complement, and the transitive-form shortcut was confirmed to give a different, wrong answer.

6.2 The Turán check

Table 5 compares the saturation threshold located by the counter with $\mathrm{ex}(n, K_c)$ evaluated in closed form from the Turán graph $T(n, c-1)$. The two are computed by entirely disjoint routes: one is an enumeration over $\binom{N}{k}$-sized spaces, the other an arithmetic expression in n and c. They agree in every case tested. We regard this as the strongest single piece of evidence that the implementation computes what Theorem 8 says it computes, because a systematic error in the counter would have to conspire with Lemma 14 to survive it.

6.3 Timing

Table 6 reports medians of five timed runs after one discarded warm-up, measured with a monotonic clock. The "block" and "exhaustive" columns are both unvectorised Python; the vectorised block figure is given separately and should not be compared against the unvectorised exhaustive column.

The measured ratio tracks (4) in order of magnitude but is not a pure test of it: the block path prunes cliques that cannot be active (Section 4) while the exhaustive reference has no analogue, which is why the measured ratio exceeds the prediction at $n = 8$, $c = 4$. At the smallest instances it falls short, fixed overhead dominating.

The largest instance in the sweep is $n = 9$, $c = 4$, $k = 14$, where the block contains 5,852,925 graphs and the vectorised implementation returns $|\mathcal{C}| = 641{,}852{,}190$ in a median of 3.56 seconds against an exhaustive space of 3,796,297,200 graphs. The largest single instance attempted anywhere is $n = 10$, $c = 4$, $k = 12$, with $|\mathcal{C}| = 667{,}768{,}395$ from a block of 3,262,623 graphs in 2.04 seconds, against an exhaustive space of about 2.9×10^{10}.

6.4 Exact counts

7 A second application: monochromatic cliques

Proposition 5 is stated for an arbitrary group action, and the single-colour theorem uses only its transitive case. Here is an application that needs the orbit form, and needs it essentially.

Two-colour the edges of K_n, say red and blue, with exactly k red. Since a colouring is a choice of red edge set, this is again $\mathcal{G}(n, k)$; the question is now how many colourings contain a *monochromatic* c-clique — equivalently, for how many $g \in \mathcal{G}(n, k)$ does g or its complement contain a c-clique. Write $\mathcal{M}(n, c, k)$ for that set.

The covering family is indexed by $I = \binom{[n]}{c} \times \{\text{red}, \text{blue}\}$, with

$$A_{(Q,\text{red})} = \{g : K_Q \subseteq g\}, \qquad A_{(Q,\text{blue})} = \{g : K_Q \cap g = \emptyset\},$$

and $d(g) =$ the total number of monochromatic c-cliques of g, red plus blue. Compatibility (3) holds for $G = \mathrm{Sym}(n)$ acting on c-subsets and trivially on colours.

Proposition 15. $\mathrm{Sym}(n)$ *has exactly two orbits on* I, *namely* $\binom{[n]}{c} \times \{\text{red}\}$ *and* $\binom{[n]}{c} \times \{\text{blue}\}$, *and*

$$|\mathcal{M}(n,c,k)| \;=\; \binom{n}{c} \left(\sum_{g \in \mathcal{B}^{\text{red}}} \frac{1}{d(g)} \;+\; \sum_{g \in \mathcal{B}^{\text{blue}}} \frac{1}{d(g)} \right),$$

where $\mathcal{B}^{\text{red}} = A_{(Q_0,\text{red})}$ *and* $\mathcal{B}^{\text{blue}} = A_{(Q_0,\text{blue})}$ *for any fixed* Q_0*.*

Proof. $\mathrm{Sym}(n)$ preserves the colour coordinate, so no element carries a red index to a blue one; within each colour it is transitive on c-subsets, giving exactly two orbits, each of size $\binom{n}{c}$. Apply Proposition 5. □

Remark 16 (The transitive form does not apply)**.** It is tempting to say that $\mathrm{Sym}(n)$ acts transitively here and reuse Theorem 8 directly. It does not, and the two orbits are not interchangeable:

$$|\mathcal{B}^{\text{red}}| = \binom{N - e_c}{k - e_c}, \qquad |\mathcal{B}^{\text{blue}}| = \binom{N - e_c}{k},$$

which differ — 220 against 924 at $n = 6$, $c = 3$, $k = 6$ — and the two inner sums differ correspondingly, 6271/140 against 7191/35. Collapsing to a single block and doubling gives the wrong answer. There *is* a colour-swapping symmetry, but it sends k to $N - k$: one finds $\sum_{\mathcal{B}^{\text{red}}} 1/d$ at k equal to $\sum_{\mathcal{B}^{\text{blue}}} 1/d$ at $N - k$. It relates different values of k rather than acting within one, which is exactly why it supplies no transitivity at fixed k.

7.1 Ramsey numbers as an external check

Summing the complement over all k counts the 2-colourings of K_n with no monochromatic c-clique. By the definition of the diagonal Ramsey number, that total is zero if and only if $n \geq R(c,c)$. So the method reproduces Ramsey numbers, and Table 9 is a second external check of the same character as the Turán check of Section 5: for $c = 3$ the count is 18 at $n = 4$, drops to 12 at $n = 5$, and vanishes at $n = 6$, recovering $R(3,3) = 6$. For $c = 4$ it remains positive through $n = 7$, consistent with $R(4,4) = 18$ [9].

We claim nothing about computing Ramsey numbers this way — the enumeration is hopeless well below $R(4,4) = 18$ [9], and the state of the art on diagonal Ramsey bounds is a different subject entirely. The value here is as verification: the vanishing at $n = 6$ is a sharp, externally-known fact that an implementation error would be unlikely to reproduce, and it requires exact counting, since an estimate straddling zero decides nothing.

8 Removing the enumeration

Section 3.1 showed that one block determines the whole distribution H_r. This section answers the question of whether h_r — or equivalently H_r — has a closed form. It does, for fixed (c, k), and the reason is elementary. We give the elementary argument first, because it is the right one; the machinery that follows is about *computing* the coefficients, not about establishing the form.

8.1 Polynomiality, elementarily

Proposition 17. *Fix c and k. For $v \geq 0$ let $\eta_r(v)$ be the number of k-edge labeled graphs on $[v]$ with no isolated vertex and exactly r c-cliques. Then*

$$H_r(n) \;=\; \sum_v \binom{n}{v}\,\eta_r(v), \tag{5}$$

and η_r is supported on $v \leq 2k$. Consequently $n \mapsto H_r(n)$ is a polynomial in n, of degree exactly $\max\{v : \eta_r(v) \neq 0\}$.

Proof. Every graph in $\mathcal{G}(n,k)$ has a unique set of non-isolated vertices; classifying by that set and by its size gives (5). A k-edge graph has at most $2k$ non-isolated vertices, so the sum is finite. A binomial transform $\sum_v \binom{n}{v} a_v$ of a finitely supported sequence is a polynomial in n whose degree is the largest v with $a_v \neq 0$, since $\binom{n}{v}$ has degree v and the leading terms cannot cancel. □

Decomposition by non-isolated vertex set is standard practice in graph enumeration — it is how labeled counts are assembled from isomorph-free generation — so we claim Proposition 17 as an observation, not as a result. It does, however, pin the degree exactly, and the answer splits.

Corollary 18. *For fixed $c \geq 3$ and $k \geq \binom{c}{2}$,*

$$\deg H_r \;=\; \begin{cases} 2k, & r = 0, \\ c + 2\bigl(k - \binom{c}{2}\bigr), & r \geq 1 \ \textit{(when } H_r \not\equiv 0\textit{)}. \end{cases}$$

Proof. For $r = 0$, a perfect matching of k disjoint edges covers $2k$ vertices and contains no clique at all, so $\eta_0(2k) > 0$; no k-edge graph covers more. For $r \geq 1$ some K_c is present, consuming c vertices and $\binom{c}{2}$ edges, and each of the $k - \binom{c}{2}$ remaining edges introduces at most two new vertices; K_c together with that many disjoint edges attains the bound with $r = 1$. □

The two cases are genuinely different: $2k > c + 2(k - \binom{c}{2})$ for every $c \geq 3$, since $2\binom{c}{2} > c$. At $c = 3$, $k = 6$ we have $\eta_0(12) = 10395 = 11!!$, the perfect matchings on twelve labeled vertices, so $\deg H_0 = 12$ while $\deg H_1 = 9$. Folding the two cases together is a mistake we made before checking.

8.2 Computing the coefficients without isomorphism testing

Proposition 17 says the closed form exists; it does not say how to get η_r. Evaluating η_r directly means enumerating k-edge graphs with no isolated vertices, which is feasible only up to isomorphism — in practice with `geng` and automorphism-group orders to convert back to labeled counts. The following expansion computes the same numbers using only subset enumeration and integer arithmetic, with no isomorphism testing anywhere.

For a set S of c-cliques write $u(S) = |\bigcup_{Q \in S} K_Q|$. A k-edge graph contains every clique of S exactly when it contains those $u(S)$ edges, so the j-th binomial moment is

$$T_j \;=\; \sum_{|S|=j} \binom{N - u(S)}{k - u(S)} \;=\; \sum_{g \in \mathcal{G}(n,k)} \binom{r(g)}{j} \;=\; \sum_r H_r \binom{r}{j},$$

whence, expanding $x^r = (1 + (x-1))^r$,

$$\sum_r H_r x^r \;=\; \sum_{j\ge 0} T_j (x-1)^j, \qquad \text{so} \qquad H_r \;=\; \sum_{j\ge r} (-1)^{j-r}\binom{j}{r} T_j. \tag{6}$$

This much is the classical binomial-moment sieve; its graph instance, with $\binom{N-u}{k-u}$ as the probability that a family of copies is present, is textbook [8].

Proposition 19. *Let $A_j(v,u)$ be the number of j-element sets of c-cliques on a fixed v-element vertex set that cover all v vertices and whose edge union has exactly u edges. Then*

$$T_j \;=\; \sum_{v,u} \binom{n}{v} A_j(v,u) \binom{N-u}{k-u}, \tag{7}$$

and the sum is finite with range depending only on (c,k): every covered vertex has degree at least $c-1$ in the union, so $2u \ge v(c-1)$ and hence $v \le 2k/(c-1)$; and u is monotone under adding cliques, so terms with $u > k$ vanish and j is bounded by a function of (c,k).

Proof. Classify each S by its covered vertex set. A set covering a given v-subset contributes $\binom{N-u(S)}{k-u(S)}$, and the number of such S with union size u is $A_j(v,u)$ by definition, independent of which v-subset and of n. Sum over the $\binom{n}{v}$ subsets. The bounds are as stated. □

Grouping families of copies by the isomorphism type of their union, with configuration counts independent of n times a binomial vertex factor, is how moments of subgraph counts have been computed since at least Ruciński [24]; that literature is asymptotic in intent and we have not found (7) written down as an exact finite identity. We record it as a convenient normal form rather than as a theorem.

8.3 What it looks like

The tables are small. For $c = 3$, $k = 6$ the whole of A is six numbers:

j	v	u	$A_j(v,u)$
1	3	3	1
2	4	5	6
2	5	6	15
2	6	6	10
3	4	6	4
4	4	6	1

Feeding these through (7) and (6) and expanding gives, for every n,

$$\begin{aligned}
H_1(n) &= \tfrac{1}{288}\, n(n-1)(n-2)(n-3)(n-4)(n+3)\big(n^3+n^2-8n-92\big),\\
H_2(n) &= \tfrac{1}{72}\, n(n-1)(n-2)(n-3)(n-4)(10n+31),\\
H_3(n) &= 0,\\
H_4(n) &= \binom{n}{4},
\end{aligned}$$

with $H_0(n) = \binom{\binom{n}{2}}{6} - H_1 - H_2 - H_4$ of degree 12. The last two lines are checks one can make by hand: a 6-edge graph with four triangles is exactly a K_4, and three triangles are impossible because the third forces the fourth — which is the support gap of Remark 12. The factor $(n-4)$ in H_1 and H_2 records that at $n = 4$ the only 6-edge graph is K_4 itself.

8.4 Cost, and where this does and does not help

Once the table exists, evaluating the entire distribution at any n is a fixed number of binomial products: microseconds at $n = 1000$ or $n = 10^6$, where neither block enumeration nor any other enumeration is conceivable. That is the substantive gain, and it is a gain in n only.

Building the table is exponential in k. Table 10 reports the configurations visited. The comparison that matters is not against block enumeration — that would be comparing the method against our own weaker one — but against the isomorph-free route to η_r, whose cost is the number of isomorphism classes of k-edge graphs, OEIS A000664. On that comparison the type expansion *loses* at $c = 3$, by a factor of some 3000 at $k = 11$, and becomes competitive only for larger c, where the constraint $u \leq k$ prunes harder. Its advantage is qualitative rather than quantitative: it needs no graph-isomorphism machinery, only integers.

8.5 What remains open

Proposition 17 settles the closed-form question as it was posed — fixed c and k, growing n — and settles it cheaply. The regime it does not touch is k growing with n: the support bound $2k$ then grows too, so the table is no longer finite and the proposition degrades into a restatement. That is the regime of interest for the asymptotic questions of Section 5, and it is where DeSalvo and Rombach [5] work, proving a pre-asymptotic Chen–Stein bound, with constants explicit at all finite parameter values, on the total variation distance between a uniform graph with prescribed edge and triangle counts and an independent-process approximation to it.

The obstruction, however, is a property of the grading and not of the problem. Section 9 regrades the same decomposition so that every order is finite for all k at once, and quantifies exactly how far into this regime that buys us — which is less far than one might hope.

9 Regrading by deficiency

Proposition 17 grades the non-isolated-vertex decomposition by vertex count v, and is finite because η_r is supported on $v \leq 2k$. That bound is what fails once k grows with n. This section replaces the grading by one whose every order is finite for all k simultaneously, at the cost of an expansion rather than a closed form. We are explicit at the end about how narrow the resulting window is.

9.1 The grading

For a k-edge graph on v non-isolated vertices put

$$d \;=\; 2k - v, \tag{8}$$

the *deficiency*, which vanishes exactly on perfect matchings. Write the *excess* of a connected graph as $\varepsilon = 2e - v$. The bound that does the work is elementary: a connected graph has

$e \geq v-1$, so

$$\varepsilon \;=\; 2e - v \;\geq\; 2(v-1) - v \;=\; v-2, \qquad \text{that is} \qquad v \;\leq\; \varepsilon + 2. \tag{9}$$

Lemma 20. *Let g be a k-edge graph with no isolated vertex and deficiency d. Then g is a disjoint union of a* core *— those components of excess ≥ 1, of total excess exactly d, each spanning at most $\varepsilon_i + 2 \leq d+2$ vertices — and a perfect matching on the remaining vertices.*

Proof. Summing $\varepsilon_i = 2e_i - v_i$ over components gives $\sum_i \varepsilon_i = 2k - v = d$. A connected component with $\varepsilon_i = 0$ has $2e_i = v_i$ and $e_i \geq v_i - 1$, forcing $v_i \leq 2$, so it is a single edge; the components of excess 0 therefore form a perfect matching on the vertices they cover. Each remaining component has $\varepsilon_i \geq 1$ and, by (9), at most $\varepsilon_i + 2$ vertices. □

The point of (8) is the consequence: at fixed d the core ranges over a set of possibilities bounded by d alone, *independently of k and of n.* Where the vertex grading has one table that grows with k, this has infinitely many tables each of which is finite.

Proposition 21. *Let $A_d^{(c)}[V, r]$ be the number of cores on a distinguished V-element vertex set with total excess d and exactly r c-cliques, and let $(m-1)!!$ denote the number of perfect matchings on m labeled vertices. Then, for every n and k,*

$$\eta_r(2k-d) = \sum_V \binom{2k-d}{V} A_d^{(c)}[V, r] \left(2k - d - V - 1\right)!!, \tag{10}$$

$$H_r(n, k) = \sum_{d \geq 0} \binom{n}{2k-d} \eta_r(2k-d), \tag{11}$$

and each $A_d^{(c)}$ is a finite table depending on c and d only. The sum (11) taken to $d = 2k$ is exact; truncated at $d \leq D$ it is a lower bound, every omitted term being non-negative.

Proof. Lemma 20 classifies each graph by its core: choose the V core vertices from the $2k-d$ covered ones, choose the core ($A_d^{(c)}[V, r]$ ways, a count independent of n and k because the core spans exactly those V vertices), and match the rest, which is possible in $(2k-d-V-1)!!$ ways and in none if $2k-d-V$ is odd. That is (10). Equation (11) is the identity (5) of Proposition 17 reindexed by $d = 2k - v$, so it is exact when the reindexing is complete, i.e. summed to $d = 2k$. For finiteness of $A_d^{(c)}$, let the core have components of excess $\varepsilon_1, \ldots, \varepsilon_t \geq 1$ summing to d, so $t \leq d$; by (9),

$$V \;=\; \sum_{i=1}^{t} v_i \;\leq\; \sum_{i=1}^{t} (\varepsilon_i + 2) \;=\; d + 2t \;\leq\; 3d,$$

a bound attained at every order we computed (Table 11, column $\max V$). Its edge count is $E = (d+V)/2 \leq 2d$. The cores are therefore drawn from graphs on at most $3d$ vertices with at most $2d$ edges, a set determined by (c, d) alone. Non-negativity of the omitted terms is immediate, each being a product of counts. □

Table 11 gives the catalogue for $c = 3$. The top two orders have closed forms in k, both of which are hand-checkable and are asserted as tests:

$$\eta_0(2k) = (2k-1)!!, \qquad \eta_1(2k-3) = \binom{2k-3}{3} (2k-7)!! \quad (c = 3), \tag{12}$$

the first because the deficiency-0 graph is a perfect matching, the second because at $r = 1$ and top vertex count the graph is one triangle plus a perfect matching. Note $\eta_0(12) = 10395 = 11!!$, recovering the value that fixes $\deg H_0 = 2k$ in Corollary 18.

9.2 The expansion parameter

Finiteness at each order is worth little without knowing how fast the orders decay. They decay in k^2/n, and the first ratio is exact.

Proposition 22. *Write $t_d = \binom{n}{2k-d}\eta_0(2k-d)$ for the order-d term of* (11) *at $r = 0$. Then, for $2k \le n$,*

$$\frac{t_1}{t_0} = \frac{2k(k-1)}{n-2k+1}. \tag{13}$$

Proof. $\binom{n}{2k-1}/\binom{n}{2k} = 2k/(n-2k+1)$. For the other factor, the deficiency-1 core is a path on three vertices, of which there are 3 on a given triple, so (10) gives $\eta_0(2k-1) = \binom{2k-1}{3}\cdot 3\cdot(2k-5)!!$, and dividing by $\eta_0(2k) = (2k-1)!!$ leaves $k-1$. Multiplying the two factors gives (13). $\square$

So $t_1/t_0 \sim 2k^2/n$, and Table 12 shows the pattern continuing: truncating at order D costs a relative error empirically proportional to $(k^2/n)^{D+1}$, with fitted exponents 0.98, 2.03, 3.04, 4.20 for $D = 0, 1, 2, 3$ against the predicted $1, 2, 3, 4$, over the 162 measured points with $k^2/n < 0.4$. The reference values there are computed by the *vertex* grading of Section 8, which is exact at fixed (c, k) for every n; using it rather than the complete deficiency sum keeps the reference independent of the quantity being tested.

9.3 What this does and does not reach

The honest reading of (13) is a limitation, and it is the reason we do not claim to have answered the question of Section 8.5. Since the expansion parameter is k^2/n rather than k/n, truncation converges for

$$k = o(\sqrt{n}\,),$$

and not for $k = \Theta(n)$ or beyond. DeSalvo and Rombach [5] work with $k_S = o(n)$ and $k_T = o(n)$. The window opened here is therefore strictly narrower than theirs and sits inside it: it is a range in which their bound is already good. An exact expansion that converges only where a proven approximation is already tight is a modest thing, and we state it as such.

What the regrading does establish is smaller and, we think, still worth recording. First, the assertion that no finite table exists in this regime — which we made in an earlier draft of Section 8 — is false as stated: it holds for the vertex grading only, and (8) defeats it in two lines. Second, the expansion is exact when summed completely and a rigorous lower bound when truncated, so it yields certified brackets rather than estimates. Third, the closed forms (12) are in k, whereas everything else in this paper is closed in n at fixed k; they are the only statements here that survive k growing.

Two further remarks bound the claim from the other side. Building order d costs 13–20 times order $d-1$ over the range we built (Table 11), so exact evaluation via $D = 2k$ is confined to very small k and the vertex grading is the better route whenever it applies; the expansion earns its place only in the regime where no complete table is available at all. And we tested whether the sequences admit a P-recurrence — a linear recurrence in k with polynomial coefficients, which is the form in which an exact treatment of the growing-k regime would most plausibly arrive. Searching all orders ≤ 4 and degrees ≤ 5 against the exact K_3-free counts by edge number at $n = 8$ and $n = 9$, over every (order, degree) pair for which the fitting system is genuinely overdetermined, the coefficient system has *no* solution at all: the nullspace is trivial in each case. (Relaxing the overdetermination requirement does produce solutions, and they fail to predict held-out terms, which is the same conclusion arrived at less cleanly.) This is evidence and not proof — the search is bounded, and n is fixed with k ranging over one row rather than the two-parameter regime itself — but it is the reason we do not present the growing-k case as being one step away.

9.4 Attribution

Grading a graph decomposition by excess, with a finite catalogue of components at each order and the bound (9), is Wright's method [29, 30], developed to count connected graphs with $n + k$ edges and standard in the random-graph literature since; the components of excess ≥ 1 are what that literature calls complex components. We claim no priority for the grading. What is new here is only its application to the clique-count distribution H_r, the observation that it restores finiteness at every order where the vertex grading fails, and the exact ratio (13) identifying k^2/n as the parameter. The relationship is the same one this paper has with Karp and Luby: a known technique, applied to a statistic for which we have not found it used.

10 Related work

The identity underlying this method is a symmetric specialization of the union-of-sets estimator of Karp and Luby [12], refined by Karp, Luby and Madras [13], which expresses the cardinality of a union as a sum over elements weighted by the reciprocal of the number of sets containing each. Here the sets are the $\binom{n}{c}$ blocks and the weight is $1/r(g)$. The contribution of the present work is not the estimator but Proposition 5: the hypothesis under which the outer sum collapses to a single set, permitting exact evaluation by enumeration in place of approximation by sampling. As noted in Section 4.4, the sampled form of the same identity gives an approximation scheme polynomial in n; the exact evaluation given here is of interest only where an exact integer is required.

Reciprocal-multiplicity weighting is routine in the subgraph- and motif-counting literature, where sampled occurrences are reweighted by inverse sampling probability [14, 1], by automorphism or orbit factors [21, 4], or eliminated altogether by canonical-order symmetry breaking [28]; see [25] for a survey. That literature counts occurrences of a motif *within a fixed host graph*, whereas the problem here is to count how many graphs of a combinatorial family contain the motif at all. The two problems are distinct and share only their bookkeeping. Deduplication by canonical representative rather than by fractional weight underlies isomorph-free exhaustive generation [18], whose tooling is also the practical state of the art for the complementary enumeration of Section 4.5.

The complementary quantity, the number of K_c-free labeled graphs, was first studied asymptotically by Erdős, Kleitman and Rothschild [7] and extended in [6, 15, 22]. Results at fixed edge count are due to Osthus, Prömel and Taraz [20] and, via hypergraph containers, to Balogh, Morris and Samotij [2], Saxton and Thomason [26], and Balogh, Morris, Samotij and Warnke [3]; sharp asymptotics for triangle-free graphs with a prescribed number of edges were obtained by Jenssen, Perkins and Potukuchi [10], with polynomial-time approximate counting in [11]. All of these are asymptotic; none gives exact values.

11 Limitations and further work

Limitations. The method is exponential in k and offers no asymptotic improvement. Its advantage over exhaustive enumeration narrows to a constant as k approaches N, and in that same regime complementary enumeration beats it outright (Section 4.5). Where approximation is acceptable it is dominated by sampling (Section 4.4). The vectorised implementation packs one edge per bit of a 64-bit word and so requires $N \leq 64$, that is $n \leq 11$; beyond that the unvectorised path applies at a substantial constant-factor cost.

That word-size limit is not, however, what stops us in practice: time binds first, and by a wide margin. Table 13 measures complementary enumeration over all edge counts at $c = 3$, which is the cheaper of our two exact routes in that regime, in the Python implementation accompanying this paper. The cost grows by a factor that is itself growing — 29×, 43×, 64× for the last three steps — reaching 399 seconds at $n = 9$. Holding the last factor fixed at 64× puts $n = 10$ at about 7 CPU-hours and $n = 11$ at some 455; continuing the observed growth in the factor itself (a further 1.5× per step, giving 95×) puts $n = 10$ nearer 11 hours and $n = 11$ at about 62 days.

Those figures measure an implementation, not the problem, and the distinction is worth stating plainly. The search visits a number of nodes proportional to the number of K_c-free graphs on $[n]$, which is 4,682,270 at $n = 8$, 246,348,115 at $n = 9$ and 19,213,627,145 at $n = 10$ (A213434); everything else is constant factor, and constant factor here is a matter of language and data layout rather than of algorithm. A direct transcription of the same search into C, enumerating the same tree, completes $n = 9$ in under four seconds and $n = 10$ in under five minutes on one core.[1] The $n = 10$ row it produces sums to 19,213,627,145, matching A213434(10), and terminates at $k = 25 = \lfloor n^2/4 \rfloor$ as Mantel requires. At that rate $n = 11$ is some hours rather than some months, and the first genuinely prohibitive case is $n = 12$, where the count of graphs to be visited exceeds 3.6×10^{14}.

The practical ceiling for complete rows is therefore set by the constant factor of the implementation rather than by the $N \leq 64$ word-size limit: it is $n = 9$ for the Python code whose timings are reported here, and roughly two further values of n for a compiled one. Either way it is reached well before $N \leq 64$ becomes relevant, which is the point at issue. The complete tables in Section 6 stop at $n = 8$ because that is where both exact routes agree cheaply, and not for the word-size reason previously stated.

A method that beats all of these. Every route compared above enumerates labeled objects. There is a fourth that does not, and for complete rows it dominates them all: generate the K_c-free graphs up to isomorphism with `geng` [18, 19], then recover the labeled count by weighting each isomorphism class G by $n!/|\mathrm{Aut}(G)|$. The number of classes is smaller than the labeled count by a factor approaching $n!$, and that factor is what the method buys.

The size of the difference is best conveyed by what other people did with it. While the tables of Section 6 were being checked, A. Howroyd computed row 10 of the $c = 3$ array from precomputed collections of graphs, and S. A. Irvine computed rows to $n = 14$ at $c = 3$ and to $n = 12$ at $c = 4$ with a parallel implementation, in a couple of hours on twenty threads. Row 12 at $c = 4$ contains 4 920 462 917 148 212 818 labeled graphs; visiting them one at a time, at the rate measured for the compiled enumerator above, would take some two thousand CPU-years.

[1] Measured on a single core of a cloud container, compiled with `gcc -O3`; the C program is short enough to be given in full in the accompanying repository. These timings are not produced by the same measurement pipeline as Table 13 and are reported only to locate the ceiling, not to compare languages carefully.

We report this rather than omit it, on the same principle as Section 4.5. Two qualifications are worth stating and neither rescues the ceiling. The comparison is between methods for producing complete rows: at a single (n, c, k) with k small, block weighting examines $\binom{N-e_c}{k-e_c}$ objects and can still be the cheaper of the two. And the method assumes the isomorphism machinery that Section 8.2 is at pains to avoid, so the two are answering different questions about what may be taken as given. What the comparison does establish is that the limit reached in Table 13 is a property of the methods of this paper and not of the problem, and that the honest ceiling for exact rows of these arrays is currently set by `nauty` and not by us.

Open questions. We are not aware of a hardness result for exact computation of $|\mathcal{C}(n, c, k)|$ and do not assert one.

The closed-form question for h_r is answered in Section 8 for fixed (c, k), and answered by an elementary argument that makes the block enumeration unnecessary in that regime. What remains is k growing with n. Section 9 makes partial progress: regrading by deficiency restores a finite table at every order, and gives an exact expansion, with truncation a certified lower bound whose relative error is measured to fall as $(k^2/n)^{D+1}$. But the expansion parameter is k^2/n, so it converges only for $k = o(\sqrt{n}\,)$ — inside the range where the Chen–Stein bound of [5] is already tight — and the genuinely open case is $k = \Theta(n)$ and above. A bounded search for a P-recurrence in k found none — the coefficient system is unsolvable at every order and degree where it is overdetermined (Section 9) — which is weak evidence that the closed-form route is not merely unfinished. Section 4.6 is the partial result in the other direction: the pairwise inclusion–exclusion in closed form, with the obstruction to continuing being the number of clique triples by union size, which does not reduce to a single parameter as the pair case reduces to s.

A hybrid that selects between block weighting and complementary enumeration by comparing $\binom{N-e_c}{k-e_c}$ against a bound on the number of K_c-free graphs would dominate both and appears straightforward.

Finally, Proposition 5 needs only a compatible action, and we have given two instances. Others in the same shape — a symmetric family of constraints on a labeled structure, where one wants the count of structures satisfying at least one — should be findable, and the orbit form means the action need not be transitive for the method to apply.

Acknowledgements

I thank Andrew Howroyd and Sean A. Irvine for independently verifying Tables 7 and 8, for extending them well beyond the range reachable by the methods described here, and for correcting my understanding of the practical cost of these computations.

Portions of the implementation, the literature search and the preparation of this manuscript were carried out with the assistance of large language models. All mathematical claims were verified computationally against independent implementations, and all cited references were checked to exist.

References

[1] N. Alon, R. Yuster and U. Zwick. Color-coding. *Journal of the ACM* **42**(4), 844–856, 1995.

[2] J. Balogh, R. Morris and W. Samotij. Independent sets in hypergraphs. *Journal of the American Mathematical Society* **28**(3), 669–709, 2015.

[3] J. Balogh, R. Morris, W. Samotij and L. Warnke. The typical structure of sparse K_{r+1}-free graphs. *Transactions of the American Mathematical Society* **368**, 6439–6485, 2016.

[4] M. Bressan, F. Chierichetti, R. Kumar, S. Leucci and A. Panconesi. Motif counting beyond five nodes. *ACM Transactions on Knowledge Discovery from Data* **12**(4), Article 48, 2018.

[5] S. DeSalvo and M. P. Rombach. An independent process approximation to sparse random graphs with a prescribed number of edges and triangles. arXiv:1509.08585, 2015.

[6] P. Erdős, P. Frankl and V. Rödl. The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent. *Graphs and Combinatorics* **2**, 113–121, 1986.

[7] P. Erdős, D. J. Kleitman and B. L. Rothschild. Asymptotic enumeration of K_n-free graphs. In *Colloquio Internazionale sulle Teorie Combinatorie* (Rome, 1973), Tomo II, Atti dei Convegni Lincei 17, Accademia Nazionale dei Lincei, 19–27, 1976.

[8] A. Frieze and M. Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2016.

[9] R. E. Greenwood and A. M. Gleason. Combinatorial relations and chromatic graphs. *Canadian Journal of Mathematics* **7**, 1–7, 1955.

[10] M. Jenssen, W. Perkins and A. Potukuchi. On the evolution of structure in triangle-free graphs. *Advances in Mathematics* **480**, 110499, 2025.

[11] M. Jenssen, W. Perkins, A. Potukuchi and M. Simkin. Sampling, counting, and large deviations for triangle-free graphs near the critical density. In *Proceedings of the 65th IEEE Symposium on Foundations of Computer Science (FOCS)*, 151–165, 2024.

[12] R. M. Karp and M. Luby. Monte-Carlo algorithms for enumeration and reliability problems. In *Proceedings of the 24th IEEE Symposium on Foundations of Computer Science (FOCS)*, 56–64, 1983.

[13] R. M. Karp, M. Luby and N. Madras. Monte-Carlo approximation algorithms for enumeration problems. *Journal of Algorithms* **10**(3), 429–448, 1989.

[14] N. Kashtan, S. Itzkovitz, R. Milo and U. Alon. Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics* **20**(11), 1746–1758, 2004.

[15] Ph. G. Kolaitis, H. J. Prömel and B. L. Rothschild. K_{l+1}-free graphs: asymptotic structure and a 0-1 law. *Transactions of the American Mathematical Society* **303**(2), 637–671, 1987.

[16] W. Mantel. Problem 28. *Wiskundige Opgaven* **10**, 60–61, 1907.

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

[18] B. D. McKay. Isomorph-free exhaustive generation. *Journal of Algorithms* **26**(2), 306–324, 1998.

[19] B. D. McKay and A. Piperno. Practical graph isomorphism, II. *Journal of Symbolic Computation* **60**, 94–112, 2014.

[20] D. Osthus, H. J. Prömel and A. Taraz. For which densities are random triangle-free graphs almost surely bipartite? *Combinatorica* **23**(1), 105–150, 2003.

[21] A. Pinar, C. Seshadhri and V. Vishal. ESCAPE: efficiently counting all 5-vertex subgraphs. In *Proceedings of the 26th International World Wide Web Conference (WWW)*, 1431–1440, 2017.

[22] H. J. Prömel and A. Steger. The asymptotic number of graphs not containing a fixed color-critical subgraph. *Combinatorica* **12**(4), 463–473, 1992.

[23] A. A. Razborov. On the minimal density of triangles in graphs. *Combinatorics, Probability and Computing* **17**(4), 603–618, 2008.

[24] A. Ruciński. When are small subgraphs of a random graph normally distributed? *Probability Theory and Related Fields* **78**(1), 1–10, 1988.

[25] P. Ribeiro, P. Paredes, M. E. P. Silva, D. Aparício and F. Silva. A survey on subgraph counting: concepts, algorithms, and applications to network motifs and graphlets. *ACM Computing Surveys* **54**(2), Article 28, 2021.

[26] D. Saxton and A. Thomason. Hypergraph containers. *Inventiones Mathematicae* **201**(3), 925–992, 2015.

[27] P. Turán. On an extremal problem in graph theory (Hungarian). *Matematikai és Fizikai Lapok* **48**, 436–452, 1941.

[28] S. Wernicke. Efficient detection of network motifs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* **3**(4), 347–359, 2006.

[29] E. M. Wright. The number of connected sparsely edged graphs. *Journal of Graph Theory* **1**(4), 317–330, 1977.

[30] E. M. Wright. The number of connected sparsely edged graphs II: Smooth graphs and blocks. *Journal of Graph Theory* **2**(4), 299–305, 1978.

n	c	k	r	H_r	support gaps
6	3	6	1	2,520	$r = 3$
			2	910	
			4	15	
6	3	9	2	510	$r = 6$
			3	1,620	
			4	1,845	
			5	960	
			7	60	
7	3	9	1	57,330	$r = 6$
			2	115,080	
			3	81,165	
			4	20,090	
			5	7,385	
			7	210	
6	4	10	1	1,260	$r = 3, 4$
			2	300	
			5	6	
7	4	12	1	97,860	$r = 4$
			2	33,670	
			3	1,400	
			5	1,155	
8	4	12	1	4,238,920	$r = 4$
			2	462,175	
			3	5,600	
			5	8,568	
9	4	14	1	555,748,830	$r = 7, 8$
			2	80,545,500	
			3	3,596,040	
			4	79,380	
			5	1,775,340	
			6	105,840	
			9	1,260	

Table 1: The global clique-count distribution H_r recovered from a single block by Corollary 10, with the interior gaps in its support. Every row was checked against direct classification of the whole space wherever that was feasible.

		objects enumerated		time (s)		
n	c	block	complement	block	complement	cheaper
6	3	4,096	5,789	0.004	0.007	block
7	3	262,144	133,501	0.125	0.177	block
8	3	33,554,432	4,682,270	19.820	6.904	complementary
6	4	512	27,626	0.001	0.030	block
7	4	32,768	1,486,597	0.012	1.897	block
8	4	4,194,304	147,141,138	2.850	314.054	block

Table 2: Full-column cost at fixed (n, c): block weighting over all k against one pass of complementary enumeration. Both are exact and their outputs were checked to agree entry by entry. Neither dominates.

s	$c = 3$	$c = 4$	$c = 5$
0	6	12	20
1	6	12	20
2	5	11	19
3		9	17
4			14

Table 3: $u(s) = 2\binom{c}{2} - \binom{s}{2}$, the number of edges in the union of two c-cliques sharing s vertices. Note $u(0) = u(1)$: two cliques meeting in a single vertex share no *edge*.

n	c	k	$\binom{N}{k}$	$\lvert\mathcal{B}_Q\rvert$	$\lvert\mathcal{C}\rvert$	exhaustive check
6	3	6	5,005	220	3,445	agrees
6	3	9	5,005	924	4,995	agrees
7	3	9	293,930	18,564	281,260	agrees
8	3	9	6,906,900	177,100	5,743,360	agrees
6	4	9	5,005	84	1,200	agrees
6	4	10	3,003	126	1,566	agrees
7	4	12	293,930	5,005	134,085	agrees
8	4	12	30,421,755	74,613	4,715,263	agrees
9	4	12	1,251,677,700	593,775	71,295,315	*not attempted*
9	4	13	2,310,789,600	2,035,800	235,545,660	*not attempted*
9	4	14	3,796,297,200	5,852,925	641,852,190	*not attempted*

Table 4: Block weighting against exhaustive enumeration. $N = \binom{n}{2}$; $\lvert\mathcal{B}_Q\rvert$ is the block size (2); $\lvert\mathcal{C}\rvert$ is the number of k-edge graphs containing a c-clique.

n	c	saturation $k^* - 1$	$\mathrm{ex}(n, K_c)$	agreement
5	3	6	6	✓
6	3	9	9	✓
7	3	12	12	✓
8	3	16	16	✓
5	4	8	8	✓
6	4	12	12	✓
7	4	16	16	✓
8	4	21	21	✓

Table 5: The saturation threshold located by the counter against the Turán number $\mathrm{ex}(n, K_c)$, computed independently in closed form (Lemma 14).

			median time (s)			ratio	
n	c	k	block	block (vec.)	exhaustive	measured	predicted
6	3	6	0.00019	0.00015	0.00247	13	23
6	3	9	0.00079	0.00033	0.00181	2	5
7	3	7	0.00328	0.00096	0.08696	26	38
7	3	9	0.02345	0.00675	0.18431	8	16
6	4	9	0.00007	0.00012	0.00306	46	60
6	4	10	0.00008	0.00010	0.00182	22	24
7	4	12	0.00615	0.00184	0.29542	48	59
8	4	9	0.00105	0.00080	12.73139	12,084	4,485
8	4	10	0.00452	0.00203	24.07948	5,328	1,794
8	4	12	0.13731	0.02826	*not attempted*	–	408
9	4	12	*n/a*	0.29787	*not attempted*	–	2,108
9	4	13	*n/a*	1.03071	*not attempted*	–	1,135
9	4	14	*n/a*	3.56328	*not attempted*	–	649
10	4	12	*n/a*	2.03815	*not attempted*	–	8,815

Table 6: Median times over five runs after warm-up. "Measured" is the ratio of the two unvectorised columns; "predicted" is (4), which depends on no implementation detail.

k	$n = 4$	$n = 5$	$n = 6$	$n = 7$	$n = 8$
3	4	10	20	35	56
4	12	70	240	630	1,400
5	6	180	1,230	5,145	16,380
6	1	200	3,445	24,920	118,230
7		120	5,775	79,170	586,040
8		45	6,330	174,615	2,111,130
9		10	4,995	281,260	5,743,360
10		1	3,003	349,755	12,209,582
11			1,365	352,296	21,011,844
12			455	293,895	30,274,027
13			105	203,490	37,410,800
14			15	116,280	40,111,560
15			1	54,264	37,441,544
16				20,349	30,421,720
17				5,985	21,474,180
18				1,330	13,123,110
19				210	6,906,900
20				21	3,108,105
21				1	1,184,040
22					376,740
23					98,280
24					20,475
25					3,276
26					378
27					28
28					1

Table 7: $|\mathcal{C}(n, 3, k)|$: labeled graphs on n vertices with k edges containing at least one triangle. Rows with all-zero entries are omitted; values for $k < 3$ are zero.

k	$n=5$	$n=6$	$n=7$	$n=8$
6	5	15	35	70
7	20	135	525	1,540
8	30	540	3,675	16,170
9	10	1,200	15,715	107,240
10	1	1,566	45,381	501,746
11		1,125	92,316	1,752,408
12		440	134,085	4,715,263
13		105	138,285	9,944,116
14		15	100,485	16,567,320
15		1	52,262	21,847,308
16			20,244	22,797,180
17			5,985	18,879,756
18			1,330	12,550,580
19			210	6,829,340
20			21	3,101,343
21			1	1,183,760
22				376,740
23				98,280
24				20,475
25				3,276
26				378
27				28
28				1

Table 8: $|\mathcal{C}(n, 4, k)|$: labeled graphs on n vertices with k edges containing at least one 4-clique. Rows with all-zero entries are omitted; values for $k < 6$ are zero.

n	c	colourings 2^N	monochromatic	free	verdict
4	3	64	46	18	$R(3,3) > 4$
5	3	1,024	1,012	12	$R(3,3) > 5$
6	3	32,768	32,768	0	$n \geq R(3,3)$
7	3	2,097,152	2,097,152	0	$n \geq R(3,3)$
4	4	64	2	62	$R(4,4) > 4$
5	4	1,024	132	892	$R(4,4) > 5$
6	4	32,768	10,284	22,484	$R(4,4) > 6$
7	4	2,097,152	1,174,140	923,012	$R(4,4) > 7$

Table 9: Two-colourings of K_n containing a monochromatic c-clique, summed over all k, by Proposition 15. The vanishing of the free count at $n = 6$, $c = 3$ recovers $R(3, 3) = 6$. Every per-k value for $n \leq 6$ was checked against direct enumeration of both a graph and its complement.

c	k	table entries	configurations visited	A000664(k)
3	6	6	285	68
3	8	12	28,546	497
3	9	24	289,870	1,476
3	10	36	3,737,266	4,613
4	10	5	171	4,613
4	13	13	42,161	–
4	14	20	491,757	–
5	15	6	420	–

Table 10: Cost of building the n-independent table, against the number of isomorphism classes of k-edge graphs with no isolated vertices (OEIS A000664), which is what the isomorph-free route to η_r must enumerate. The type expansion is the more expensive of the two at $c = 3$.

order d	core types	configurations	max V	subsets visited	build (s)
0	1	1	0	0	0.000
1	1	3	3	3	0.000
2	2	106	6	23	0.000
3	4	9,366	9	354	0.001
4	7	1,581,131	12	4,737	0.010
5	9	437,264,616	15	85,587	0.199
6	13	179,890,797,606	18	1,749,658	4.512

Table 11: The deficiency catalogue at $c = 3$. Every order is a finite table, which is the property the vertex grading of Section 8 loses once k grows with n. The build enumerates edge subsets of K_v for $v \leq d + 2$, costing a factor of 13–20 more per order over the last three; $d = 6$ is the practical ceiling, and the implementation refuses higher orders rather than appearing to hang.

		relative error truncating at order D				
n	k^2/n	$D = 0$	$D = 1$	$D = 2$	$D = 3$	t_1/t_0
40	0.9000	8.1×10^{-1}	4.2×10^{-1}	1.2×10^{-1}	1.8×10^{-2}	60/29
80	0.4500	5.4×10^{-1}	1.5×10^{-1}	2.1×10^{-2}	1.5×10^{-3}	20/23
160	0.2250	3.2×10^{-1}	4.5×10^{-2}	3.1×10^{-3}	1.0×10^{-4}	60/149
320	0.1125	1.7×10^{-1}	1.2×10^{-2}	4.1×10^{-4}	6.9×10^{-6}	20/103
640	0.0563	9.0×10^{-2}	3.2×10^{-3}	5.4×10^{-5}	4.4×10^{-7}	60/629
1280	0.0281	4.6×10^{-2}	8.1×10^{-4}	6.8×10^{-6}	2.8×10^{-8}	20/423
2560	0.0141	2.3×10^{-2}	2.1×10^{-4}	8.6×10^{-7}	1.8×10^{-9}	60/2549
5120	0.0070	1.2×10^{-2}	5.2×10^{-5}	1.1×10^{-7}	1.1×10^{-10}	20/1703
10240	0.0035	5.8×10^{-3}	1.3×10^{-5}	1.4×10^{-8}	6.9×10^{-12}	60/10229

Table 12: Truncating the deficiency expansion at order D, for $c = 3$, $k = 6$, $r = 0$. Each column falls by a further factor of k^2/n, and the last column is the exact ratio (13). Reference values are from the vertex grading of Section 8, an independent route to the same numbers.

n	K_c-free graphs	time (s)	factor over $n-1$
3	7	¡0.001	–
4	41	¡0.001	–
5	388	¡0.001	–
6	5,789	0.005	–
7	133,501	0.145	29×
8	4,682,270	6.219	43×
9	246,348,115	398.592	64×

Table 13: Measured cost of exhaustive complementary enumeration at $c = 3$, over all edge counts, in the Python implementation accompanying this paper. The growth factor is itself increasing, so the practical ceiling is reached before the $N \leq 64$ word-size limit becomes relevant; a compiled implementation of the same search moves that ceiling by roughly two values of n without changing the conclusion.