

Fine-Grained AC^0 Lower Bounds for k -OV, k -XOR, and k -SUM via Colored Subgraph Isomorphism

Haoxing Lin *

National University of Singapore
haoxingl@comp.nus.edu.sg

Abstract

We prove lower bounds for k -OV, k -XOR, and k -SUM in nonuniform AC^0 , tracking how the circuit-size exponent scales with k and using no running-time hypothesis. Our framework gives depth-zero projections from colored subgraph isomorphism to the three targets at dimension, row count, or bit width $O(k \log n)$, without increasing depth or size, and preserving gate orientation. For every fixed depth and every sufficiently large fixed k , we obtain unconditional bounds $n^{\Omega(k)}$ for k -OV and $(n/k)^{\Omega(k)}$ for k -XOR and k -SUM, with an absolute exponent-rate constant independent of both k and the depth, while the onset threshold may depend on (d, k) . For growing $k = n^{o(1)}$, we obtain, for every fixed depth d , the unconditional floor $n^{\Omega_d(\min\{\sqrt{k}, \log n\})}$. This strengthens to $n^{\Omega(k)}$ at depth two for both top-gate orientations, i.e., top conjunction and top disjunction, and at depth three for top-disjunction circuits $\text{OR} \circ \text{AND} \circ \text{OR}$, with no restriction on bottom fan-in or literal polarity. The depth-three argument rests on a minterm bound for a single CNF: a fixed CNF is very unlikely to become true for the first time exactly when a randomly planted copy is completed, with no circuit-correctness hypothesis. Assuming a pattern-uniform strengthening of the Li–Razborov–Rossman source lower bound, the same projections complete the $k = n^{o(1)}$ frontier with $n^{\Omega_d(k)}$ for the missing top-conjunction depth-three orientation and for every fixed depth $d \geq 4$. The interface of our projection framework does not depend on which source lower bound is used, so new lower bounds for suitable source families or improved source exponents pass directly to all three targets. All direct k -XOR bounds stated above concern odd k ; a black-box odd-to-even lift transfers any such lower bound through a supplied admissible parameter decomposition. The k -SUM projection works for both parities. At the bit width $m = \Theta(k \log(en/k))$ used by our projection, a block-carry Σ_3 upper bound of size $(n/k)^{O(k)}$ matches the top-disjunction depth-three lower bound $(n/k)^{\Omega(k)}$ up to constants in the exponent. The remaining upper-versus-lower-bound gaps concern depth two, top-conjunction depth three, and other width regimes.

*This work was supported by the National Research Foundation, Singapore, under its NRF Fellowship programme, award no. NRF-NRFF14-2022-0010.

Contents

1	Introduction	4
1.1	Problems, fine-grained complexity, and constant-depth circuits	4
1.2	Main results and the depth-versus- k frontier	6
1.3	Proof architecture and technical ideas	9
1.4	Related work	11
1.5	Organization	14
2	Colored Subgraph Isomorphism and the Source Menu	14
2.1	The structured source and the circuit model	14
2.2	Fixed-pattern hardness	15
2.3	The source menu	18
3	The Projection Framework	18
3.1	Targets and depth-zero projections	18
3.2	The common universe skeleton	19
3.3	Master projections and host substitutions	21
4	Three Instantiations	23
4.1	k -Orthogonal Vectors	24
4.2	Odd k -XOR	26
4.3	k -SUM	29
5	Unconditional Lower Bounds for Fixed k	32
6	Unconditional Lower Bounds for Growing k	34
6.1	A clique fallback at every fixed depth	34
6.2	Depth two: an exponent linear in k	38
6.3	Depth three: an exponent linear in k with a top disjunction	40
6.4	The row-rich PARITY comparison	44
7	The Pattern-Uniform Conjecture and the Conditional Completion	45
8	Synthesis and Open Problems	47
8.1	Cross-target synthesis	47
8.2	Open problems	48
A	Source Formalities	53
A.1	Coordinate equivalence with the LRR structured problem	53
A.2	Citation assembly for the fixed-pattern lower bound	53
A.3	The switching-scale obstruction	55
B	Pattern Constructions and Constants	56
B.1	The bounded-degree expander family	56
B.2	The clique family	59

C	Full Projection Proofs	59
C.1	The k -OV projection	59
C.2	The odd- k k -XOR projection	61
C.3	The k -SUM projection	63
D	Parameter Completions	66
D.1	Extension beyond the constructed budget	66
D.2	General row counts for odd k -XOR	67
D.3	General bit widths for k -SUM	68
D.4	Two boundary cases for k -SUM	70
E	Supplementary Results for k-XOR	71
E.1	The even- K lifting projection	71
E.2	Consequences of the lift	73
E.3	The row-rich PARITY route	74
F	Elementary Upper Bounds and Quantitative Tightness	76
F.1	Brute-force circuits	76
F.2	The structural constant floor for the expander route	79
F.3	The depth-sensitive comparison for k -SUM	79
G	Depth-Three Probability Estimates	80
G.1	The single-CNF minterm bound	80
G.2	The planted-copy good event	82

1 Introduction

The problems studied in this paper have different algebraic forms, but they share one witness-selection task: choose k objects and verify one joint relation. Informally, their accepting witnesses have the following forms:

$$\begin{aligned}
k\text{-OV} : \quad & \exists (u_1, \dots, u_k) \in U_1 \times \dots \times U_k \quad \forall t \in [D], \quad \prod_{i=1}^k u_i(t) = 0, \\
k\text{-XOR} : \quad & \exists S \in \binom{[n]}{k} \quad \bigoplus_{i \in S} a_i = 0^m, \\
k\text{-SUM} : \quad & \exists S \in \binom{[n]}{k} \quad \sum_{i \in S} z_i \equiv 0 \pmod{2^m}.
\end{aligned}$$

For k -OV, the input consists of k arrays of n pairwise-distinct D -dimensional binary vectors, and a witness chooses one vector from each array; by the distinctness promise, each array is simply a set of n vectors listed in some order. For k -XOR and k -SUM, the input is a single indexed list of size n whose entries are m -dimensional binary vectors (k -XOR) or m -bit integers (k -SUM), and a witness chooses k distinct indices; different selected indices may carry equal values. The formal definitions and input conventions appear in Section 3.1.

Fine-grained complexity refines coarse running-time classifications by asking for the correct exponent of the running time and by using reductions that preserve the significance of an exponent improvement [IP01, IPZ01, VW18]. We ask the analogous quantitative question for nonuniform constant-depth computation: how does the exponent of the required circuit size scale with k and with the dimension, row count, or bit width of the instance? This is an analogy in granularity, not a transfer of a running-time hypothesis.

We answer this circuit analogue by projecting from colored subgraph isomorphism rather than by importing a time-complexity hypothesis. A common family of depth-zero projections yields fixed- k lower bounds for all three targets and an unconditional depth-versus- k frontier when k grows. The exponent is linear in k at depth two and for top-disjunction depth-three circuits. A precise source-side conjecture would extend that exponent to the missing top-conjunction depth-three orientation and to every fixed depth $d \geq 4$.

1.1 Problems, fine-grained complexity, and constant-depth circuits

The three target problems. The targets share a selection skeleton but use three different verification algebras. Orthogonality asks that every coordinate avoid the all-one pattern among the selected vectors. XOR asks that independent linear constraints over $\mathbb{F}_2$ vanish. Modular SUM asks that an integer sum vanish in $\mathbb{Z}_{2^m}$, where carries couple the bit positions. Orthogonal Vectors is a canonical hub of fine-grained complexity, and its two-vector form has a substantial equivalence class under truly subquadratic reductions [VW18, CW19]; the k -set formulation studied here exposes the number of selected vectors as an additional quantitative parameter. The k -SUM and k -XOR problems have their own algorithmic, average-case, and generalized-birthday lineages, discussed in Section 1.4. We do not place all three targets under one algorithmic conjecture. They are grouped here because the same source witness can be realized through each of their three algebras.

Our sense of “fine-grained” in this paper. Here *fine-grained* refers to the dependence of the circuit-size exponent on k and to the preservation of that dependence under our projections.

Theorems A–E in Section 1.2 are unconditional. Theorem F assumes only Pattern-Uniform LRR, a circuit lower-bound conjecture for colored subgraph isomorphism; it is neither SETH, OVC, nor any running-time hypothesis for the target problems. The fine-grained literature therefore motivates the quantitative question and the choice of targets, but it supplies no hardness assumption used in our unconditional proofs.

Constant-depth circuits. The class AC^0 consists of nonuniform families of polynomial-size, constant-depth Boolean circuits with unbounded-fan-in AND and OR gates and negations at the inputs. To state quantitative lower bounds, we allow arbitrary size at each fixed depth and let $size_d(f)$ denote the minimum number of non-input gates in a depth-at-most- d circuit computing f , with negations normalized to input leaves. Throughout the paper, $\Omega(\cdot)$ hides an absolute positive constant, whereas $\Omega_d(\cdot)$ may hide a constant depending on the fixed depth d ; in general, a subscript lists every parameter on which the hidden constant may depend. Foundational work of Ajtai, Furst, Saxe, Sipser, Yao, and Håstad established strong limitations of this model, most famously through lower bounds for PARITY and related explicit functions [Ajt83, FSS84, Yao85, Hås86].

Those classical results usually express complexity as a function of the total input length. Our question is parameter-sensitive: for structured witness-selection functions, how precisely can the circuit-size exponent scale with the number k of selected objects and with the target’s constraint budget, namely dimension, row count, or bit width?

Time lower bounds do not answer the circuit question. SETH- and OVC-based lower bounds concern the running time of uniform algorithms [IP01, IPZ01, VW18, CW19]. They do not by themselves rule out nonuniform constant-depth circuits, whose circuit may be chosen separately for each parameter tuple. Our route therefore starts with unconditional constant-depth lower bounds for a source Boolean function and transfers them by depth-zero projections. Such a projection replaces every input of a target circuit by a constant, a source input bit, or its negation; composition preserves depth and introduces no new gates. The resulting bounds are direct worst-case circuit lower bounds for the target functions.

The resulting viewpoint. This function-level bridge yields a modular methodology rather than three unrelated reductions. Four proved source statements and one conjecture feed a common projection interface and then branch into orthogonality, addition over $\mathbb{F}_2$, and modular integer addition. Quantitatively, the outcome depends jointly on circuit depth, the witness parameter k , the target budget (dimension for k -OV, row count for k -XOR, or bit width for k -SUM), and the way the target universe packs the source candidates.

Structurally, the direct k -XOR qualification arises while enforcing exact selection over $\mathbb{F}_2$: linear equations over $\mathbb{F}_2$ count selected entries only modulo two, so they cannot by themselves exclude a spurious selection that takes an even number of entries from every group of the constructed list; such a selection has even size, so an odd k excludes it while an even k does not. Signed integer selectors give k -SUM no analogous parity restriction. Both list targets nevertheless remain PARITY-hard after a simple restriction that fixes the k selected indices. Quantitatively, for a fixed support $S \in \binom{[n]}{k}$, its k -XOR predicate has a depth-two verifier of size $m \cdot 2^{O(k)}$, whereas the direct k -SUM verifier costs $2^{\Theta(km)}$; at the bit-width scale $m = \Theta(k \log(en/k))$ supplied by our projection, a block-carry depth-three circuit removes the extra factor of m in the latter exponent. We next state the resulting depth-versus- k frontier before turning to the underlying proof architecture.

Regime and circuit class	Common displayed lower bound	Status
Sufficiently large fixed k , every fixed depth	$n^{\Omega(k)}$ for k -OV; $(n/k)^{\Omega(k)}$ for k -XOR and k -SUM	Unconditional; Theorem B
Growing $k = n^{o(1)}$, every fixed depth d	$n^{\Omega_d(\min\{\sqrt{k}, \log n\})}$	Unconditional; Theorem C
Growing $k = n^{o(1)}$, depth two, both orientations	$n^{\Omega(k)}$	Unconditional; Theorem D
Growing $k = n^{o(1)}$, depth three with top disjunction	$n^{\Omega(k)}$	Unconditional; Theorem E
Growing $k = n^{o(1)}$, top-conjunction depth three; every fixed depth $d \geq 4$	$n^{\Omega_d(k)}$	Conditional on Pattern-Uniform LRR; Theorem F

Table 1: The depth-versus- k frontier. In the growing- k rows the common base- n display uses $k = n^{o(1)}$; the exact natural base for k -XOR and k -SUM is n/k , as in Table 2. Here $\Omega(\cdot)$ hides an absolute constant and $\Omega_d(\cdot)$ a depth-dependent one. The direct k -XOR statements are for odd k .

1.2 Main results and the depth-versus- k frontier

We refer to the principal results as Theorems A–F. As a cue to the notation used below, γ denotes a source-host exponent rate and β or c a target-size rate; Section 2.1 gives the full convention and its exceptions. Throughout this subsection, the direct k -XOR statements concern odd k ; the even- k lift is summarized after Theorem F.

Theorem A: one projection interface for three targets. Let P be a simple loopless pattern with exactly $k - 1$ edges. In colored P -subgraph isomorphism at host size N , written $P\text{-SUB}_N$ and formalized in Definition 2.1, the host has one color class $\{v\} \times [N]$ of N vertices for every $v \in V(P)$, and the input specifies the possible host edges between the color classes corresponding to edges of P . The source accepts when one can choose one vertex from each color class so that every pattern edge is present; the coloring fixes which host vertex plays each pattern role. Under the projection, an accepting target witness consists of one candidate object for each of the $k - 1$ selected host edges together with one anchor object, and therefore has size k . For k -OV the anchor merely supplies the k th array, whereas for the two list targets it is also the reference against which the constraints force exactly one candidate per pattern edge. Writing $\leq_{\text{proj}}$ for a depth-zero projection, we obtain $P\text{-SUB}_N \leq_{\text{proj}} k\text{-OV}$, $P\text{-SUB}_N \leq_{\text{proj}} k\text{-XOR}$ for odd k , and $P\text{-SUB}_N \leq_{\text{proj}} k\text{-SUM}$ for both parities, in each case without increasing circuit depth or size and while preserving gate orientation. The exact host capacities and $O(k \log n)$ -scale budgets appear in Table 2. The k -OV construction stores each edge group in a separate array and therefore permits $N = \Theta(\sqrt{n})$, whereas the two list targets k -XOR and k -SUM store all $k - 1$ groups together and permit only $N = \Theta(\sqrt{n/k})$; this packing difference produces the two natural lower-bound bases, n for k -OV and n/k for the list targets. Consequently, any lower bound for a suitable colored-subgraph source transfers through the same three target constructions, and source hardness and target realization can be selected independently. In particular, a new source family or an improved source exponent rate upgrades all three target bounds automatically, the source rate carrying to the target base with a factor-two loss, because each target base is of order N^2 .

Target	Source host size	Target budget	Natural base and qualification
k -OV	$N = \lfloor \sqrt{n} \rfloor$	$D \geq C_{\text{OV}} k \log n$	base n ; both parities
k -XOR	$N = \left\lfloor \sqrt{\frac{n-1}{k-1}} \right\rfloor$	$m \geq C_{\oplus} k \log(en/k)$	base n/k ; direct projection for odd k
k -SUM	$N = \left\lfloor \sqrt{\frac{n-1}{k-1}} \right\rfloor$	$m \geq C_{\Sigma} k \log(en/k)$	base n/k ; both parities; top-disjunction depth-three exponent-scale tight when $m = \Theta(k \log(en/k))$

Table 2: Target packing and principal constraint budgets. We call the displayed dimension, row-count, and bit-width thresholds the *principal budgets*: they are the $O(k \log n)$ -scale budgets supplied directly by the projections of Theorem 3.4. The universal constants $C_{\text{OV}}, C_{\oplus}, C_{\Sigma} > 0$ are fixed in Theorem 3.4 and Lemma 3.5, and these budgets are used by Theorems 5.1 and 6.4, Theorem 6.8, and Theorem 7.3. Theorem 6.2 instead uses the more demanding k -SUM width $m \geq C_{\Sigma} k \log n$, because $O(k \log(kN))$ is not uniformly $O(k \log(n/k))$ throughout its range $k \leq n^{1-\varepsilon}$.

Table 1 summarizes the resulting circuit frontier. Its growing- k rows use the common simplification $k = n^{o(1)}$, under which $n/k = n^{1-o(1)}$ and the list-target lower bounds can also be displayed in base n . The exact target bases and constraint budgets appear in Table 2 and in the formal theorem statements.

The depth-two theorem is proved in the larger ranges $k \leq \xi_D n^{1/4}$ for k -OV and $k \leq \xi_D n^{1/5}$ for k -XOR and k -SUM, where $\xi_D > 0$ is a universal range constant. The depth-three unconditional row concerns only $\Sigma_3 = \text{OR} \circ \text{AND} \circ \text{OR}$ circuits; it gives no lower bound for the top-conjunction orientation Π_3 and hence no lower bound for unrestricted depth-three circuits, where either orientation is allowed.

Theorem B: a universal fixed- k rate. Let $k_0 \in \mathbb{N}$ be the universal threshold in Lemma 2.5 and for each $k \geq k_0$, let P_k^{exp} denote the padded expander pattern supplied there. Let $C_{\text{OV}}, C_{\oplus}, C_{\Sigma}$ be the universal construction constants fixed by Theorem 3.4. There exists a universal target-side exponent rate $\beta > 0$, independent of k , the depth d , and the instance size n . For every fixed depth d and every fixed $k \geq k_0$, once n exceeds a threshold depending on (d, k) , every depth- d circuit for k -OV has size at least $n^{\beta(k-1)}$, while every depth- d circuit for k -XOR or k -SUM has size at least $(n/k)^{\beta(k-1)}$, at the budgets in Table 2. The exponent-rate constant is therefore uniform in both d and k ; only the onset threshold is not. At the principal budgets, the top-disjunction depth-three specializations for all three targets match elementary upper bounds at the exponent scale. This comparison makes no tightness claim at depth two or for the top-conjunction depth-three orientation.

Theorem C: an unconditional growing- k floor at every fixed depth. When k grows, the fixed-pattern source theorem provides no uniform control of the onset threshold, so we replace the padded expander source by a clique core padded with a matching, whose source hardness is Beame’s small-clique lower bound [Bea90] and remains uniform as the clique order grows. For every fixed depth d and every subpolynomial $k = n^{o(1)}$ tending to infinity, the resulting lower bound for each target has the common scale

$$n^{\Omega_d(\min\{\sqrt{k}, \log n\})}.$$

More generally, the theorem holds throughout $k \leq n^{1-\varepsilon}$ with the exact base- n and base- (n/k) expressions in Theorem 6.2; its k -SUM line uses the larger width $m \geq C_{\Sigma} k \log n$ recorded in

Table 2. The exponent first grows as $\sqrt{k}$ because a clique using at most $k - 1$ pattern edges has only $\Theta(\sqrt{k})$ vertices. Beame’s bound, however, covers only cliques of order at most $\log N$, so beyond $k = \Theta(\log^2 n)$ the usable clique order stops growing with k and the exponent levels off at $\Omega_d(\log n)$. Where this transition occurs does not depend on d ; only the level at which the exponent settles may depend on it.

Theorem D: both depth-two orientations. At depth two, direct counting restores an exponent linear in k . The DNF branch is the standard minimal-positive-input count: every colored copy is a minimal positive input of the source and therefore forces a distinct accepting term. The source-specific CNF branch constructs an explicit exponentially large family of minimal transversals of the colored-copy hypergraph, each of which forces a clause; a single pattern vertex of degree at least two already generates the required family. Transferring the resulting source bound yields $n^{\Omega(k)}$ for k -OV when $k \leq \xi_D n^{1/4}$ and $(n/k)^{\Omega(k)}$ for the two list targets when $k \leq \xi_D n^{1/5}$. Thus both the top-conjunction and top-disjunction orientations are covered.

Theorem E: top-disjunction depth three. The same padded bipartite expander family also supports an unconditional Σ_3 lower bound with exponent linear in growing k , without any restriction on bottom fan-in or literal polarity. For every constant c with $0 < c < 1/16$, choose any γ with $2c < \gamma < 1/8$. Then there is a polynomial-threshold exponent $C'(\gamma, c) > 1$ such that, for all sufficiently large k and every $n \geq k^{C'(\gamma, c)}$, every Σ_3 circuit computing k -OV has size at least n^{ck} , while every such circuit computing k -XOR or k -SUM has size at least $(n/k)^{ck}$. The main new ingredient is one support maximum, fixed before the relevant planted labels are revealed, that controls every state of an adaptive commitment process and avoids a factor- k loss from union-bounding over clause histories. The top disjunction selects one accepting middle CNF; the argument provides no corresponding Π_3 or size₃ lower bound.

Theorem F: a conditional completion. Pattern-Uniform LRR asserts that, for each fixed depth d , there are constants $\gamma_d, c'_d > 0$ such that the padded expander family $(P_k^{\text{exp}})_{k \geq k_0}$ satisfies $\text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma_d k}$ for every $k \geq k_0$ and every $N \geq k^{c'_d}$. Assuming this source-side statement, the unchanged projections of Theorem A give $n^{\Omega_d(k)}$ lower bounds throughout the sub-polynomial regime for all three targets. At depth three this supplies the missing top-conjunction orientation, and for every fixed depth $d \geq 4$ it supplies the full linear-in- k exponent. We use Pattern-Uniform LRR only as a sufficient condition and do not claim that it is necessary. It strengthens the fixed-pattern Li–Razborov–Rossman theorem by requiring one polynomial host threshold uniformly along the growing family, and it is not known to follow from that theorem.

Two target-specific qualifications complete the frontier. First, the direct k -XOR projection enforces exact selection through equations over $\mathbb{F}_2$ and is therefore clean for odd k . For an even target parameter K , Lemma E.2 gives a black-box lift from any supplied odd parameter r in an admissible decomposition $K = tr + s$; Corollary E.4, Corollary E.5, and Corollary E.6 record the corresponding quantitative consequences. Second, the k -SUM projection works for both parities because signed integer coordinates enforce exact selection counts. Although the direct fixed-support depth-two verifier costs $2^{\Theta(km)}$, Proposition F.3 gives a top-disjunction depth-three block-carry circuit of size $(n/k)^{O(k)}$ at the principal width $m = \Theta(k \log(en/k))$. At fixed k , this makes the top-disjunction depth-three specialization of Theorem 5.1 exponent-scale tight. Together with Theorem 6.8, the same conclusion holds for growing k in the stated polynomial-host regime. This does not settle the optimal size at depth two or for the opposite depth-three orientation. Nor does it make fixed-support verification polynomial-size: Proposition D.6 gives a two-point PARITY




Figure 1: Proof architecture. Solid boxes and arrows denote proved ingredients; the dashed box and arrow denote Pattern-Uniform LRR. The expander family carries B, D, E, and F, while the clique family carries C; see Table 3. All five paths use the same target interface in Theorem A.

restriction for k -SUM, and the same one-line restriction applies to k -XOR. Rather, block carries remove the extra factor of m in the exponent of the direct k -SUM verifier.

1.3 Proof architecture and technical ideas

Figure 1 records the proof dependencies. Five source statements feed two pattern families, each padded to exactly $k - 1$ edges, and every resulting colored subgraph source passes through the same target interface. In particular, the conditional path changes only the source statement: it introduces no new reduction to any target.

One witness skeleton, three verification algebras. A source witness selects one host edge for each edge of the pattern. It is valid exactly when every selected edge is present and all edges incident to the same pattern vertex assign that vertex the same host label. The target instances represent each proposed host edge by a candidate object, add one anchor, and pad each list or array to length n with dummies, while coordinates or rows enforce selection, presence, consistency, dummy exclusion, and, only for k -OV, pairwise distinctness within each array. The algebra changes but the semantics do not: a k -OV witness already chooses one vector from each array, k -XOR uses equations over $\mathbb{F}_2$ to enforce the group pattern, and k -SUM uses exact signed coordinates that are packed into one residue modulo a power of two. Because the target constructions are depth-zero projections, all circuit lower-bound work remains on the colored-subgraph side.

Why two pattern families appear. The strongest available source statement depends on the circuit regime. The common bounded-degree family has a bipartite 4-regular expander core. Its expansion gives $\kappa(P_k^{\text{exp}}) = \Omega(k)$ for the fixed-pattern Li–Razborov–Rossman theorem; its $\Theta(k)$ vertices, absence of isolated vertices, and degree-4 vertex support the depth-two counting argument; and one bipartition side supplies $(1 - o(1))k/4$ vertices with pairwise disjoint incident-edge blocks for the top-disjunction depth-three argument. A clique core instead supports Beame’s lower bound [Bea90] uniformly while the clique size grows. The edge allowance limits its order to $\Theta(\sqrt{k})$, while Beame’s depth-independent range requires $r \leq \log N$, yielding the source scale

$\Theta(\min\{\sqrt{k}, \log N\})$ used in [Theorem C](#). Substituting the source host size N in terms of the target size n , as in [Table 2](#), gives $\log N = \Theta(\log n)$ for k -OV and $\log N = \Theta(\log(n/k))$ for the list targets; both are $\Theta(\log n)$ in the subpolynomial- k regime. Disjoint matching edges pad both families to exactly $k - 1$ edges without weakening the relevant source lower bound.

Exact counting at depth two. After flattening adjacent gates with the same connective, a depth-two circuit is either a DNF, an OR of AND-terms, or a CNF, an AND of OR-clauses. Here an edge variable is the Boolean source input recording one possible colored host edge. For the DNF orientation, consider the minimal positive input consisting exactly of the edge variables of one source witness, or colored copy. Deleting any selected variable removes the only present edge from one pattern-edge group, so the input becomes rejecting. Any DNF term accepting this input must therefore contain every selected edge variable positively. Because the pattern has no isolated vertices and has $\Theta(k)$ vertices, distinct source maps give distinct colored copies and force $N^{\Omega(k)}$ distinct terms. For the CNF orientation, monotonicity lets us remove negative literals, after which each clause is a set of edge variables meeting every colored copy; call such a set a transversal. For each inclusion-minimal transversal, any CNF for the function must contain a clause included in it, and minimality forces equality. To generate many of them, choose a pattern vertex a of degree $\delta \geq 2$. Independently for each host label $i \in [N]$, choose one incident pattern edge $h(i)$ and include every variable on $h(i)$ that assigns label i to a . Every colored copy assigns some label i to a and therefore meets this set. Conversely, for any included variable x , fixing the endpoint labels encoded by x and choosing all other labels arbitrarily produces a colored copy whose unique intersection with the set is x , proving minimality. The $\delta^N \geq 2^N$ choices of h give distinct minimal transversals, and $2^N \geq N^{\Omega(k)}$ in the theorem's range $N \geq k^2$. Thus both depth-two orientations require size $N^{\Omega(k)}$.

The single-CNF depth-three mechanism. A Σ_3 circuit is an OR of middle CNFs. A random experiment forms a sparse background G and adds one uniformly planted colored copy F^* . One bipartition side I of the common expander core has size $L = (1/4 - o(1))k$, and the planted edges incident to its vertices form pairwise disjoint blocks. With probability $1 - o(1)$, no planted edge was already present and F^* is the unique copy in $G \cup F^*$, so deleting any planted edge makes the input rejecting. Correctness and the top disjunction then select one middle CNF T that accepts $G \cup F^*$ but rejects every such deletion: T must turn from false to true exactly when the planted copy is completed. Call this the minterm event $\mathcal{M}(T)$.

To see why that is unlikely, add the blocks of I one vertex at a time, in a random order. At every stage before the last, some clause of T is still false, so the vertex added next must be one whose planted block satisfies that clause. Only a few vertices can, because each pairing of a vertex of I with a label it might receive carries its own chance p of already appearing in the background: a clause that many such pairs could satisfy would, with high probability, have been satisfied by the background alone. Accordingly, fix the planted labels outside I but leave G random, and let Z be the largest number of such pairs still available to satisfy a clause, in the worst case over all 2^L stages and all s clauses. Then $\Pr[Z \geq t] \leq 2^L s e^{-pt}$, and hence $\|Z\|_L := (\mathbb{E}[Z^L])^{1/L} = O((L + \log(es))/p)$. Only afterward are G and the labels inside I exposed. With q vertices left to add, at most Z of the qN available pairs satisfy the false clause, so the process advances with probability at most $Z/(qN)$; on $\mathcal{M}(T)$ some vertex must work at every stage, which happens with probability at least $1/q$. Multiplying over the L stages and canceling the two $L!$ factors yields

$$\Pr[\mathcal{M}(T)] \leq \left(\frac{C_0(L + \log(es))}{pN} \right)^L.$$

Two union bounds should be distinguished here. Because one Z serves every clause at every stage, the argument never records which clause is satisfied at which stage; union-bounding over those ordered histories instead would cost a factor k in the exponent. The only union bound taken over the circuit is the outer one, over its middle CNFs: by Equation (4), a size- S source circuit has at most $S + 2M_{\text{src}}$ of them, each with that many clauses, and in the polynomial host range $M_{\text{src}} \leq kN^2$ is absorbed into the contradicted size bound. Taking $pN = N^{1/2-o(1)}$ and union-bounding over at most $2N^{\gamma k}$ middle CNFs gives $o(1)$ whenever $\gamma < 1/8$, contradicting the good event's $1 - o(1)$. The quadratic host substitutions then convert the source coefficient into every target coefficient $c < 1/16$ in [Theorem E](#).

Where parity and carries enter. The two target effects arise at different logical stages. For k -XOR, the obstruction appears while enforcing selection: equations over $\mathbb{F}_2$ count the selected candidates in each group only modulo two, so they cannot by themselves exclude the branch in which the anchor is unselected and every group contributes an even number of candidates. The selected candidates are then even in number, so an odd k forces an odd number of the k chosen indices onto dummies, and a single anti-dummy row, which is 1 exactly on the dummies, rejects any such selection. An even k leaves the branch open. For k -SUM, signed integer coordinates enforce the desired support exactly for both parities, so the projection itself loses nothing. Once a support $S \in \binom{[n]}{k}$ is fixed, that is, once the k selected indices have been chosen, the remaining k -XOR or k -SUM predicate checks only whether those entries satisfy the corresponding zero-XOR or zero-sum relation. Under a two-point restriction, either fixed-support predicate becomes the complement of PARITY, and hence remains PARITY-hard at every fixed depth as k grows. The quantitative distinction is that k -XOR has a depth-two fixed-support verifier of size $m \cdot 2^{O(k)}$, whereas the direct k -SUM verifier costs $2^{\Theta(km)}$; the top-disjunction depth-three circuit of [Proposition F.3](#) guesses block carries and removes the extra factor of m in the exponent at the principal width $m = \Theta(k \log(en/k))$. In summary, the odd- k loss is a selection phenomenon specific to k -XOR, while the carry overhead is a verification-cost phenomenon specific to k -SUM. [Section 8](#) develops this distinction further after the formal proofs.

1.4 Related work

Direct lower bounds for Orthogonal Vectors. Kane and R. Ryan Williams prove lower bounds for Orthogonal Vectors in branching programs, Boolean formulas over bounded-fan-in complete bases, and formulas whose unbounded-fan-in gates compute arbitrary symmetric functions; in their main regimes these bounds are within $\tilde{O}(D)$ factors of the natural upper bounds [[KW19](#), Thms. 2, 3, and 5]. Choudhury, Limaye, Sreenivasaiah, and Srinivasan later sharpen the bounded-fan-in formula and branching-program bounds in the regime $D > c \log n$, where $c > 1$ is constant, obtaining $\Omega(n^2 D)$ for formulas and $\Omega(n^2 D / \log(nD))$ for branching programs [[CLSS26](#), Thms. 6 and 25]. Kane and Williams's lower-bound argument uses a non-product restriction. Separately, for each fixed bias $p \in (0, 1)$, they construct AC^0 formulas of size $n^{2-\varepsilon_p}$, for some $\varepsilon_p > 0$, that are correct on a $1 - o(1)$ fraction of instances under the p -biased product distribution [[KW19](#), Thm. 6]. They also formulate a depth-three AC^0 circuit lower bound for worst-case Orthogonal Vectors as a conjecture [[KW19](#), Conj. 10]. These formula and branching-program models differ from general unbounded-fan-in AC^0 circuits, so the results are precursors and motivation rather than statements subsumed by ours.

Choudhury and Sreenivasaiah study layered top-disjunction depth-three circuits through disjunctions of restricted bottom functions. They prove that, for $k \leq D$, a disjunction of t -CNFs computing the *unpromised* k -OV function, in which vectors may repeat within an array, has top

fan-in, and hence size under our convention, $\Omega((n/t)^k)$; more generally, the bottom gates may be unate functions of unbounded fan-in provided that at most t inputs have negative polarity [CS26, Thms. 1.1, 1.2, and 5.12]. They also prove a parameterized lower bound for the opposite depth-three orientation with no bottom-fan-in restriction, yielding a $2^{\Omega(n)}$ lower bound for the unpromised 2-OV function when $D = \Omega(n^2)$ [CS26, Thm. 1.3]. When $k = O(1)$ and $k \leq t = O(1)$, their main bound matches the n^k exponent of the brute-force construction; for general t , their paper records a remaining gap in the top-/bottom-fan-in tradeoff. The promise obstruction differs between the two arguments. In the top-disjunction argument, every maxterm has at most one vector other than $\vec{1}$ in each array, so no maxterm satisfies pairwise distinctness once $n \geq 3$ [CS26, Prop. 4.4 and proof of Thm. 5.12]. In the opposite-orientation argument, the image of the projection contains arrays with repeated vectors, so a circuit required to be correct only on promised inputs need not decide the entire projected source function. Thus neither proof by itself lower-bounds $\text{size}_3(k\text{-OV}_{n,D,k})$ under the distinctness promise of Definition 3.1. Our top-disjunction theorem instead treats ordinary Σ_3 circuits with no restriction on clause width or literal signs, allows k to grow, and transfers to the promised k -OV function, and hence also to its unpromised array variant, as well as to odd- k k -XOR and k -SUM. The two results concern complementary and incomparable circuit models and parameter regimes.

The Choudhury–Limaye–Sreenivasaiah–Srinivasan work cited above also proves quantitatively stronger lower bounds for unrestricted fixed-depth circuits in a different dimension regime. Its main constant-depth theorem shows that, for fixed constants $k, d \geq 2$ and $\varepsilon > 0$, every depth- d circuit computing the monotone complement of k -OV at dimension $D = n^\varepsilon$ has size at least $n^{k-\varepsilon}$; because the circuit model is nonmonotone, they observe that the same statement may be phrased directly for k -OV [CLSS26, Thm. 4 and Footnote 4]. They also prove the monotone-circuit analogue of OVC for the complement of the unpromised two-set Orthogonal Vectors function at logarithmic dimension; in the proof’s 0-distribution, every unchosen position is filled with 1^D , and Lemma 20 uses that filler directly [CLSS26, Thm. 5; construction of $\mathcal{D}_0$, Step 4; and Lem. 20]. Their 1-distribution is pairwise distinct with high probability, so the promise sensitivity enters specifically through the 0-distribution. Their transfer is projection-based: a split-and-list construction maps ℓ -hyperclique in a u -uniform hypergraph on v vertices to the target by listing all ℓ/k -subsets and paying dimension $O(v^u)$ [CLSS26, Thm. 11]. The listed vectors are pairwise distinct within each array, so this transfer is compatible with the promise in Definition 3.1. This gives a near-maximal exponent at every fixed depth for fixed k , while their discussion records a barrier at $D = \Omega((\log n)^2)$ for the underlying reduction framework [CLSS26, Rem. 16]. Our results instead work at $O(k \log n)$ -scale budgets, allow k to grow, and cover k -XOR and k -SUM through the same interface.

Broader Orthogonal Vectors–circuit interfaces. Abboud, Bringmann, Dell, and Nederlof show that falsifying the moderate-dimensional Orthogonal Vectors Conjecture would yield unexpected speedups for weighted hypergraph clique and sparse TC^1 -circuit satisfiability [ABDN18]. R. Ryan Williams develops a different win–win connection between OVC and nonuniform lower bounds [RW25]: one branch gives lower bounds for exact-threshold and equality-matrix representations, while the other yields nearly linear algorithms for logarithmic-dimensional Orthogonal Vectors together with further nonuniform consequences. These works establish important interfaces between fine-grained complexity and circuit complexity, but they do not give direct worst-case lower bounds for the present target functions in standard AC^0 . Our route instead transfers unconditional source-function lower bounds through depth-zero projections and does not assume OVC.

Constant-depth lower bounds for subgraph isomorphism. The source side belongs to the line of constant-depth lower bounds for detecting small subgraphs: Beame’s small-clique bound, Rossman’s k -Clique theorem, Amano’s lower bounds for uncolored general graph and hypergraph patterns (extending a weaker-exponent form of Rossman’s method), and the structural theory of colored subgraph isomorphism developed by Li, Razborov, and Rossman [Bea90, Ros08, Ama10, LRR17]. Li, Razborov, and Rossman identify the colored fixed-pattern exponent within a logarithmic factor of treewidth and prove minor-monotonicity of κ_{col} ; on the function side, their graph-minor reductions for structured colored subgraph isomorphism are linear-size monotone projections. Rosenthal subsequently closes their factor-two gap in the average-case exponent [Ros21, Thm. 1.3], proves that the embedding parameter $\text{emb}(G)$ underlying Marx’s ETH-hardness result is $O(\kappa(G))$ [Ros21, Thm. 1.5], and separates κ from treewidth: for the t -dimensional hypercube Q_t , one has $\kappa(Q_t) = \Theta\left(\text{tw}(Q_t)/\sqrt{\log \text{tw}(Q_t)}\right)$ [Ros21, Thm. 1.4 and the discussion following it]. Our source menu uses the fixed-pattern Li–Razborov–Rossman theorem for Theorem B and Beame’s clique lower bound for Theorem C, while the depth-two and top-disjunction depth-three source theorems are proved here. The fixed-pattern quantifiers are essential: Pattern-Uniform LRR strengthens the host-size uniformity across a growing pattern family and is not known to follow from the published theorems.

Fine-grained algorithms and Orthogonal Vectors. The ETH and SETH formulations of Impagliazzo and Paturi and the sparsification and reduction framework of Impagliazzo, Paturi, and Zane underlie the modern fine-grained program [IP01, IPZ01]. Vassilevska Williams surveys the resulting exponent-preserving methodology, and Chen and R. Ryan Williams identify a broad equivalence class around two-set Orthogonal Vectors under truly subquadratic reductions [VW18, CW19]. This literature motivates both the choice of Orthogonal Vectors and our insistence on preserving exponent dependence, but its conditional running-time lower bounds are logically separate from the nonuniform circuit lower bounds proved here.

Algorithmic and average-case work on k -SUM and k -XOR. Wagner’s generalized-birthday algorithm initiated a central algorithmic line for finding short zero-sum and XOR relations in dense groups [Wag02]. Dinur, Keller, and Klein establish essentially tight conditional bounds for dense average-case k -SUM and k -XOR for $k \leq 5$, and in restricted parameter ranges for larger k [DKK24], while Agrawal, Saha, Schwartzbach, Vanukuri, and Vasudevan study sparse and planted regimes, hardness amplification, reductions, and cryptographic applications [ASS⁺24]. Dalirrooyfard, Lincoln, Saha, and Vassilevska Williams establish conditional average-case hardness for parity-counting variants of k -OV, k -SUM, and k -XOR under explicit distributions [DLSV25]. This literature motivates the algebraic targets and their sensitivity to parameters, but it concerns uniform algorithms and distributional or counting problems with different success criteria and computational models. The present results are worst-case nonuniform AC^0 lower bounds and do not imply security for any cryptographic construction.

Comparison of regimes. Existing direct results give strong Orthogonal-Vectors-related lower bounds in four distinct settings: structurally restricted top-disjunction depth-three circuits when $k \leq D$, the opposite depth-three orientation without a bottom-fan-in restriction at dimension $D = \Omega(n^2)$, unrestricted fixed-depth circuits for fixed k at polynomial dimension, and monotone circuits for the unpromised two-set complement at logarithmic dimension. Our results occupy a complementary region: $O(k \log n)$ -scale constraint budgets, growing k , three target functions, and a common depth-zero transfer, together with a top-disjunction depth-three lower bound having

unrestricted clause width and literal signs. The resource units in these comparisons are not uniform: Kane–Williams use formula and branching-program size in some results [KW19, Thms. 2 and 3] and count wires in their symmetric-gate theorem and depth-three conjecture [KW19, Thm. 5; Conj. 10]; Choudhury–Limaye–Sreenivasaiah–Srinivasan use formula or branching-program size [CLSS26]; and Choudhury–Sreenivasaiah count gates or middle CNFs [CS26], whereas size_d here counts non-input gates. The comparisons above are therefore model-and-regime comparisons, not numerical identifications across complexity measures. The missing top-conjunction orientation and the higher-depth linear-in- k frontier are stated conditionally rather than folded into the unconditional claims.

1.5 Organization

Section 2 presents colored subgraph isomorphism, the source lower bounds, and the two pattern families. Section 3 isolates the depth-zero projection interface, and Section 4 realizes it for k -OV, odd- k k -XOR, and k -SUM. Section 5 derives the fixed- k consequences. Section 6 proves the unconditional growing- k frontier, including the depth-two and top-disjunction depth-three results, and records the complementary row-rich PARITY route together with its projection-saturation limit. Section 7 states Pattern-Uniform LRR and derives the conditional completion, and Section 8 synthesizes the parity and carry phenomena and records the open problems. Appendices A to G contain the source formalities, pattern constructions, full projection proofs, parameter completions, supplementary k -XOR results, elementary upper bounds, and depth-three probability estimates.

2 Colored Subgraph Isomorphism and the Source Menu

Throughout, k denotes the number of target objects selected by a witness, and every corresponding source pattern has exactly $k - 1$ edges. The source problem is a structured colored form of subgraph isomorphism in which the color classes are built into the input coordinates.

2.1 The structured source and the circuit model

The source function. Let P be a finite simple loopless undirected graph. Write $v(P) := |V(P)|$, $e(P) := |E(P)|$, and $\Delta(P)$ for its maximum degree, and fix an orientation $f = (a_f, b_f)$ of every edge $f \in E(P)$. The orientation is bookkeeping only and doesn't imply any directed structure.

Definition 2.1 (Colored P -subgraph isomorphism). For $N \geq 1$, the input to $P\text{-SUB}_N$ consists of variables

$$X_{f,u,w} \in \{0, 1\}, \quad f = (a_f, b_f) \in E(P), \quad u, w \in [N]. \quad (1)$$

The variable $X_{f,u,w}$ records whether the colored host contains the edge between (a_f, u) and (b_f, w) . The function accepts exactly when there is a map $\phi : V(P) \rightarrow [N]$ such that

$$X_{f,\phi(a_f),\phi(b_f)} = 1 \quad \text{for every } f \in E(P). \quad (2)$$

Equivalently, for each pattern edge $f = (a_f, b_f)$, one may select a host pair $(u_f, w_f) \in [N]^2$ with $X_{f,u_f,w_f} = 1$. These selections define an accepting map precisely when all edges incident to the same non-isolated pattern vertex assign that vertex the same host index; any isolated pattern vertices may then be assigned arbitrary indices.

The host vertex set is $V(P) \times [N]$, with one color class $\{a\} \times [N]$ for each $a \in V(P)$. The map ϕ need not be injective as a map into $[N]$: if $a \neq b$, then $(a, \phi(a))$ and $(b, \phi(b))$ are distinct host vertices even when $\phi(a) = \phi(b)$. Appendix A proves that this function is exactly the structured

problem $\text{SUBGRAPH}_{\text{col},N}(P)$ used by Li, Razborov, and Rossman [LRR17, Section 5, p. 960] under a bijection of input coordinates. Their host-size parameter is written n ; throughout this paper it is renamed N , while n is reserved for the target universe size.

If a pattern P is obtained from a core F by adjoining disjoint matching edges, we first fix an orientation of F , then fix an orientation of P whose restriction to F is the chosen orientation of F , and orient the padding edges arbitrarily.

Circuit conventions. Throughout the paper, $\log$ denotes the base-two logarithm. We normalize circuits so that negations occur only at input leaves, and we count only non-input gates. Thus $\text{size}_d(g)$ is the minimum number of non-input gates in an unbounded-fan-in AND/OR circuit of depth at most d computing g . At depth three, write $\Sigma_3 := \text{OR} \circ \text{AND} \circ \text{OR}$ and $\Pi_3 := \text{AND} \circ \text{OR} \circ \text{AND}$, with

$$\text{size}_3(g) = \min\{\text{size}_{\Sigma_3}(g), \text{size}_{\Pi_3}(g)\}. \quad (3)$$

For a CNF T , let $s(T)$ denote its number of clauses. After also deleting repeated literal inputs, a Σ_3 circuit of size at most S on M input variables can be written as $\bigvee_{i=1}^t T_i$, where

$$t \leq S + 2M \quad \text{and} \quad s(T_i) \leq S + 2M \quad \text{for every } i \in [t]. \quad (4)$$

The dual statement holds for Π_3 . Indeed, the non-input children are counted among the S gates, while at most $2M$ distinct literals can bypass a layer. For $P\text{-SUB}_N$, the source-variable count is $M_{\text{src}} = e(P)N^2$. These bounds will be used explicitly in Section 6.3.

Exponent rates, construction constants, and thresholds. When a scalar denotes an exponent rate, we generally use γ , possibly with a subscript, for a rate measured in the source host size, and β or c for a rate measured in the target size. The inherited LRR edge weighting $\beta : E(P) \rightarrow [0, 2]$ in Appendix A, the structural constant c_κ in Lemma 2.5, and Beame’s clique rate a_d in Fact 6.1, and the derived even- K rate a_E in Corollary E.5 keep their own letters. A capital C with one of the target subscripts OV , $\oplus$, or Σ denotes a universal construction-budget constant; a k with a subscript denotes an onset threshold on k ; and symbols such as n_1 , N_0 , $L_\oplus$, and L_Σ denote onset thresholds on the universe, host, or workable-width parameter. The symbols c'_d , $C_E(\gamma)$, $C_F(d)$, and $C'(\gamma, c)$ are exponents in polynomial host or universe thresholds such as $N, n \geq k^C$, not circuit-size rates.

2.2 Fixed-pattern hardness

The fixed-pattern source theorem and its expander corollary use the colored parameter $\kappa_{\text{col}}(P)$ of Li, Razborov, and Rossman [LRR17, Def. 2.12(iii)]. We abbreviate it to $\kappa(P)$, following Rosenthal [Ros21, Footnote 1]; the uncolored parameter denoted by the same letter in parts of [LRR17] is never used in this paper. For this paper, the operational role of $\kappa(P)$ is as the source-host exponent in the fixed-pattern lower bound below; Theorem 2.4 supplies the structural estimates used to lower-bound it.

Theorem 2.2 (Li–Razborov–Rossman fixed-pattern lower bound [LRR17]). *For every fixed simple loopless pattern P with at least one edge and no isolated vertices¹, every fixed depth d , and every*

¹We state the import only for patterns without isolated vertices. Li–Razborov–Rossman do permit a single vertex as a base element of a union sequence [LRR17, Def. 2.11], so this is a conservative draft-side scope rather than an empty-sequence convention; it matches every application below and avoids relying on an unused isolated-vertex edge case in the $\alpha \equiv 1$ citation assembly.

constant $\eta > 0$, there is a threshold $N_0(P, d, \eta)$ such that

$$\text{size}_d(P\text{-SUB}_N) + 2e(P)N^2 + 2 \geq N^{\kappa(P)-\eta} \quad (5)$$

for every $N \geq N_0(P, d, \eta)$.² If $\kappa(P) - \eta > 2$, then the threshold can be chosen so that the additive convention-bridge term is absorbed and $\text{size}_d(P\text{-SUB}_N) \geq N^{\kappa(P)-\eta}$.³

The theorem is pointwise in the pattern: both the loss absorbed by η and the onset threshold may depend on P . This distinction is harmless for fixed k , where the selected pattern is fixed as $N \rightarrow \infty$, but becomes the precise issue isolated by the pattern-uniform conjecture in Section 7. Appendix A gives the model, measure, notation, and citation assembly proving Theorem 2.2 for the structured function of Definition 2.1.

Definition 2.3 (Edge expansion). For a graph P with at least two vertices, let

$$h(P) := \min_{\substack{\emptyset \neq S \subseteq V(P) \\ |S| \leq v(P)/2}} \frac{|E_P(S, V(P) \setminus S)|}{|S|}. \quad (6)$$

Equivalently, Equation (6) is the minimum of $|E_P(S, V(P) \setminus S)| / \min\{|S|, |V(P) \setminus S|\}$ over all nonempty proper sets $S \subsetneq V(P)$, which is the normalization used in [LRR17, Thm. 4.9].

Theorem 2.4 (κ from expansion, treewidth, and minors). *For every simple loopless graph P with at least three vertices, at least one edge, and no isolated vertices,*

$$\kappa(P) \geq \frac{v(P)h(P)}{3\Delta(P)}. \quad (7)$$

Furthermore, $\kappa(P) \leq \text{tw}(P) + 1$. Finally, κ is minor-monotone: if F is a minor of P , then $\kappa(F) \leq \kappa(P)$.

Proof. The three assertions are [LRR17, Thm. 4.9], [LRR17, Prop. 4.3], and [LRR17, Thm. 5.1], respectively. The first citation states the expansion estimate without the condition $v(P) \geq 3$, but K_2 is a boundary counterexample under the cited definitions: the one-term union sequence consisting of its only edge has cost zero, so $\kappa(K_2) = 0$, whereas the right-hand side of Equation (7) is $2/3$. We therefore include the corrected boundary condition. The no-isolated-vertices condition also avoids the undefined degree ratios that arise for isolated vertices in the cited proof. None of the three results requires connectedness as a formal hypothesis, although $h(P) = 0$ makes the expansion estimate vacuous whenever P is disconnected. The stated scope is the only regime in which we invoke this theorem and is also the standing scope of Rosenthal’s corresponding restatements [Ros21, Thm. 3.8(ii),(iii)]; Footnote 2 of [Ros21] identifies those restatements with the corresponding Li–Razborov–Rossman results. In particular, every expander core and every ambient padded graph in the minor-monotonicity applications has at least three vertices and no isolated vertices. A bounded-degree family with $h(P) = \Omega(1)$ therefore has $\kappa(P) = \Omega(v(P))$ [Ros21, Cor. 3.9(i)]. $\square$

The theorem supplies all structural information about κ needed below. On a bounded-degree expander core F , Equation (7) gives $\kappa(F) = \Omega(v(F))$. Adjoining a disjoint matching need not

²The additive term in Equation (5) makes the import independent of whether the source’s phrase “number of gates” includes variable, constant, or input-negation nodes; no result derived from Corollary 2.6 carries it, while the quantifier display in Section 7 restates the theorem in full.

³No connectedness hypothesis is imposed: the two padded pattern families are cores together with disjoint matchings, and [LRR17, Rem. 2.7] computes the relevant parameters for a disconnected pattern.

preserve expansion, but F remains a minor of the resulting padded pattern, so minor-monotonicity transfers this lower bound. The treewidth upper bound in Theorem 2.4 situates κ below the familiar treewidth exponent; Remark B.5 records that the two parameters have the same linear scale on the expander cores used here. For $t \geq 0$, let J_t denote a matching of t pairwise disjoint edges on fresh vertices, with J_0 empty; both pattern families used in this paper are disjoint unions $F \dot{\cup} J_t$ of a core F with such a matching.

Lemma 2.5 (Padded bipartite expander patterns). *There are universal constants $k_0 \in \mathbb{N}$ and $c_\kappa > 0$ such that, for every integer $k \geq k_0$, there is a simple loopless graph*

$$P_k^{\text{exp}} = F_k^{\text{exp}} \dot{\cup} J_{t_k} \quad (8)$$

with exactly $k - 1$ edges, no isolated vertices, and maximum degree 4, where F_k^{exp} is a simple connected bipartite 4-regular expander and $0 \leq t_k < 4$. Fixing either bipartition side as I_k , one has

$$e(F_k^{\text{exp}}) > k - 5, \quad |I_k| = \left\lfloor \frac{k-1}{4} \right\rfloor > \frac{k-1}{4} - 1, \quad (9)$$

and

$$\kappa(P_k^{\text{exp}}) \geq c_\kappa(k-1). \quad (10)$$

One may take $c_\kappa = 1/147$; throughout, k_0 denotes a corresponding construction threshold enlarged, if necessary, so that $c_\kappa(k_0 - 1) > 4$.

Proof. Appendix B.1 fixes $\Delta = 4$ and $\varepsilon = 1/10$, applies Alon's exact-size transformation [Alo21] to the simple near-Ramanujan source graphs of Mohanty, O'Donnell, and Paredes [MOP20], and then takes the bipartite double cover of the resulting exact-size 4-regular graph. The cover preserves the required spectral gap, has two independent sides of the displayed size, and uses all but fewer than four of the available edges. The core and every padding edge have positive degree, so the padded graph has no isolated vertices. The discrete Cheeger inequality in Fact B.3, Theorem 2.4, and matching padding then give the stated κ bound and exact edge count. $\square$

Corollary 2.6 (Universal fixed- k source rate). *Fix any constant γ with $0 < \gamma < c_\kappa$. For every fixed depth d and every fixed integer $k \geq k_0$, there is a threshold $N_0(\gamma, d, k)$ such that*

$$\text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma(k-1)} \quad (11)$$

for every $N \geq N_0(\gamma, d, k)$.

Proof. The onset threshold k_0 fixed in Lemma 2.5 satisfies $c_\kappa(k_0 - 1) > 4$. Put $\bar{\gamma} := (c_\kappa + \gamma)/2$ and set $\eta_\gamma := (c_\kappa - \bar{\gamma})(k_0 - 1) > 0$. Fix d and $k \geq k_0$. By Lemma 2.5, the pattern P_k^{exp} is simple and loopless, has at least one edge, and has no isolated vertices, so it satisfies the hypotheses of Theorem 2.2. Moreover,

$$\kappa(P_k^{\text{exp}}) - \eta_\gamma \geq c_\kappa(k-1) - (c_\kappa - \bar{\gamma})(k_0 - 1) \geq \bar{\gamma}(k-1) \geq \frac{c_\kappa}{2}(k_0 - 1) > 2.$$

The pure non-input-gate consequence in Theorem 2.2 therefore gives, for every sufficiently large N ,

$$\text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\kappa(P_k^{\text{exp}}) - \eta_\gamma} \geq N^{\bar{\gamma}(k-1)} \geq N^{\gamma(k-1)}.$$

Absorbing all preceding thresholds into $N_0(\gamma, d, k)$ proves the claim. $\square$

The threshold $N_0(\gamma, d, k)$ in Corollary 2.6 is a threshold on the host size N . Once γ is fixed, the proof uses one loss parameter η_γ for the whole family, but the onset threshold imported from Theorem 2.2 may still depend arbitrarily on the fixed pattern P_k^{exp} , so $N_0(\gamma, d, k)$ retains a dependence on k ; Section 7 isolates this remaining nonuniformity. A later target-side threshold $n_1(d, k)$ is obtained only after substituting a host size $N = N(n, k)$ that tends to infinity with the target universe; Lemma 3.5 records the target-side substitutions and estimates used for this step.

2.3 The source menu

The paper uses two source families because the strongest available source lower bound depends on the circuit regime; these are the two pattern-family branches in Figure 1. Table 3 records where each lower-bound statement or proof and each pattern construction appears.

Theorem 6.4 is attached to the expander family because that family is already available with exactly $k - 1$ edges and the required elementary structure. Its DNF branch uses $v(P_k^{\text{exp}}) = \Theta(k)$ together with the absence of isolated vertices, whereas its CNF branch uses only a core vertex of degree 4; neither branch uses expansion or κ . The depth-three proof uses a different feature of the same core, namely one bipartition side whose incident edge blocks are pairwise disjoint; it likewise uses neither expansion nor κ . The clique family is not bounded-degree when r grows, so the two rows do not share a degree hypothesis.

Table 3: The source families used in the lower-bound transfers.

Source family	Property used	Detailed locations	Results carried
bipartite 4-regular expander core plus matching	$\kappa(P_k^{\text{exp}}) = \Omega(k)$ for the fixed-pattern; $v(P_k^{\text{exp}}) = \Theta(k)$ and no isolated vertices for the depth-two DNF branch; a degree-4 vertex for the depth-two CNF branch; an independent core side of size $(1 - o(1))k/4$ for the Σ_3 branch	Lemma 2.5, section 6.2, and Conjecture 7.1; construction in Appendix B.1; depth-three use in Section 6.3 and appendix G	Theorems 5.1, 6.4, 6.8 and 7.3
K_r plus matching	Beame’s small-clique lower bound [Bea90]	Fact 6.1; construction and projections in Appendix B.2	Theorem 6.2

3 The Projection Framework

The source function asks for one present host edge for each edge of the pattern, with agreement at shared endpoints. This section isolates that anatomy from the algebra used to enforce it. The resulting interface has one common universe skeleton, four semantic roles, and one domain-validity role; Section 4 realizes them separately through orthogonality, addition over $\mathbb{F}_2$, and addition over the integers.

3.1 Targets and depth-zero projections

We first fix the input conventions for the three targets, including the pairwise-distinctness promise for k -OV and the indexed-list semantics used by k -XOR and k -SUM.

Definition 3.1 (The three target problems). The target families are defined as follows.

- (i) k -OV: For $n \leq 2^D$, an input to k -OV $_{n,D,k}$ consists of k ordered arrays, each containing n vectors in $\{0,1\}^D$; this is a knD -bit encoding. Write $U_i = (u_{i,1}, \dots, u_{i,n})$, and impose the promise that the vectors within each array are pairwise distinct. The function accepts exactly when there are indices $j_1, \dots, j_k \in [n]$ such that $\prod_{i=1}^k u_{i,j_i}(t) = 0$ for every $t \in [D]$.
- (ii) k -XOR: An input to k -XOR $_{n,m,k}$ is an indexed list $(a_1, \dots, a_n) \in (\mathbb{F}_2^m)^n$. The function accepts exactly when there is a set $S \in \binom{[n]}{k}$ such that $\bigoplus_{i \in S} a_i = 0^m$.
- (iii) k -SUM: An input to k -SUM $_{n,m,k}$ is an indexed list $(z_1, \dots, z_n) \in (\mathbb{Z}_{2^m})^n$. The function accepts exactly when there is a set $S \in \binom{[n]}{k}$ such that $\sum_{i \in S} z_i \equiv 0 \pmod{2^m}$.

For the promised k -OV function, size_d denotes the minimum size of a circuit that is correct on every promised input.

The promise formulation of k -OV is equivalent to presenting each class as a set: the order of an array is irrelevant to the function value. A lower bound for this promised function automatically applies to the unrestricted array variant, because every circuit correct when repetitions are allowed is also correct on pairwise-distinct arrays. By contrast, the indices, rather than the values, are the objects selected in k -XOR and k -SUM, so distinct indices may carry equal vectors or equal residues. This asymmetry is essential: the two list reductions may reuse a single value at several padding indices, whereas the k -OV construction must land inside the pairwise-distinct promise.

Definition 3.2 (Projection). For possibly promised Boolean functions F and G , write $F \leq_{\text{proj}} G$ if there is a map R from the promise domain of F into the promise domain of G such that $F(x) = G(R(x))$ for every promised input x and every output bit of R is a constant, an input bit, or the negation of an input bit.

We call such a map a *depth-zero projection*: it substitutes constants and literals for target inputs and therefore adds neither gates nor circuit depth.

Lemma 3.3 (Projection monotonicity). *If $F \leq_{\text{proj}} G$, then $\text{size}_d(G) \geq \text{size}_d(F)$ for every depth d . The same implication holds for each fixed circuit orientation, such as Σ_3 or Π_3 .*

Proof. Substitute the constants and literals of the projection for the input leaves of a circuit computing G on its promise domain. Because the projection maps every promised input of F into that domain, the resulting circuit computes F on its promise domain. The substitution adds no gates and no depth, and it preserves the orientation of every internal gate. $\square$

3.2 The common universe skeleton

Fix an integer $k \geq 2$ and a simple loopless pattern P with $e(P) = k-1$, together with the orientation $f = (a_f, b_f)$ of each edge fixed in Section 2.1.

By Definition 2.1, a source witness chooses one present host edge for each of the $k-1$ pattern edges, and the choices must agree on the host label assigned to every shared endpoint. A target witness instead chooses k objects from lists or classes of prescribed length n . To capture and encode these accurately, we adopt a common universe skeleton that creates one *candidate* object for every possible edge choice, adds one *anchor* to guarantee valid selection as well as bridge the gap from $k-1$ edge choices to k selected objects, and uses *dummy* objects solely to pad each target list or class to length n . To enforce endpoint consistency, the skeleton picks one reference incident edge for every non-isolated pattern vertex and compares the label proposed by each non-reference incident edge with those by the reference edge.

For every $f = (a_f, b_f) \in E(P)$ and $(u, w) \in [N]^2$, introduce a candidate object $c_{f,u,w}$ representing the proposal that the pattern edge f is realized by the host edge from (a_f, u) to (b_f, w) . We also introduce one anchor object α . Thus the genuine object collection have the form

$$\{\alpha\} \dot{\cup} \bigcup_{f \in E(P)} \{c_{f,u,w} : (u, w) \in [N]^2\}. \quad (12)$$

For k -XOR and k -SUM, these objects occupy one indexed list and number $1 + e(P)N^2$; under $1 + e(P)N^2 \leq n$, dummies fill the list to length n . For k -OV, the anchor and the $e(P)$ edge groups become k separate classes: the anchor class contains one genuine object and each edge class contains N^2 genuine objects; under $N^2 \leq n$, dummies fill every class to length n .

The anchor compensates for the one-object gap between the $k - 1$ edge groups and a target witness of size k . In the additive targets it is also the pivot of the selector equations: comparing every group count with the anchor indicator forces one candidate per group rather than zero candidates per group. For k -OV, selection already chooses one vector from each class, so the anchor serves only as the additional class needed to reach k classes.

We next formalize how the skeleton determines the reference edges and how the comparison of proposed labels is encoded. For a pattern vertex a incident to an oriented edge $f = (a_f, b_f)$, define the label proposed by a candidate by

$$\text{lab}_a(f, u, w) := \begin{cases} u, & a = a_f, \\ w, & a = b_f. \end{cases} \quad (13)$$

Put $\ell_N := \lceil \log N \rceil$. Clearly, one can use the ℓ_N -bit binary encoding of $x - 1$ for each label $x \in [N]$. And we write $\text{bit}_t(x)$ for its t -th bit, where $t \in [\ell_N]$.

At a non-isolated pattern vertex a , comparing every incident edge with one reference edge is sufficient to ensure consistency because equality is transitive. Accordingly, fix one reference incident edge $\rho(a)$ for each such a , and let

$$\mathcal{I}(P) := \{(a, g) : a \in V(P), \deg_P(a) > 1, g \ni a, g \neq \rho(a)\}. \quad (14)$$

Thus $\mathcal{I}(P)$ indexes one comparison for every non-reference incidence, and

$$|\mathcal{I}(P)| = \sum_{a: \deg_P(a) \geq 1} (\deg_P(a) - 1) \leq \sum_{a \in V(P)} \deg_P(a) = 2e(P). \quad (15)$$

An implementation of the skeleton must enforce the following roles, summarized in Table 4.

- (R1) **Selection rigidity.** The witness selects the anchor and exactly one candidate from each edge group.
- (R2) **Guard.** A selected candidate $c_{f,u,w}$ is admissible only when $X_{f,u,w} = 1$.
- (R3) **Consistency.** For each pattern vertex a , all selected candidates from edges incident to a propose the same value under lab_a .
- (R4) **Purge.** No dummy belongs to an accepting witness.
- (R5) **Promise validity for k -OV.** Within each k -OV class, all n vectors are pairwise distinct, without changing which one-per-class selections are orthogonal.

Table 4: How the three target constructions realize the common witness roles.

Target	Selection rigidity (R1)	Guard and purge (R2, R4)	Consistency (R3)	Domain condition (R5)	Constructed budget
k -OV	built into the one-vector-per-class witness	one coordinate per class	two coordinates per compared label bit	inert identifiers enforce pairwise distinctness	$D_0 = O(k \log(2n))$
k -XOR	$e(P)$ selector rows, for odd k	compressed separating labels and one anti-dummy row	one row per compared label bit	total indexed-list domain	$m_{\oplus} = O(k \log(2N))$
k -SUM	exact signed group coordinates	one nonnegative guard and one dummy coordinate	one signed coordinate per comparison	total indexed-list domain	$m_{\Sigma} = O(k \log(2kN))$

The first four roles identify accepting witnesses with the present-and-consistent edge selections described in Definition 2.1. The fifth does not alter that witness semantics, but it is essential for the image of the k -OV map to lie in the promise domain of Definition 3.1. Obtaining it with only $O(\log n)$ additional coordinates is therefore a strength of the construction rather than a cosmetic representation choice.

The shared source and the three complete target instances in Figures 2 to 5 instantiate Table 4 row by row. The next section uses them as running examples, so the algebraic differences can be read directly from the coordinates that realize the same semantic role.

Three target-side effects determine the final parameter statements. First, each edge group contains N^2 candidates. Since k -OV stores different groups in different classes, it permits $N^2 \leq n$ and later yields a base- n bound; k -XOR and k -SUM store all groups in one list, require $1 + e(P)N^2 \leq n$, and later yield a base- (n/k) bound. Second, the $\mathbb{F}_2$ selector sees cardinalities only modulo two, which creates the odd- k restriction in the direct k -XOR projection. Third, the k -SUM construction packs $O(k)$ signed coordinates into one residue, producing the bit-width cost $O(k \log(kN))$. Constants such as c_{κ} and the onset threshold in the source lower bound are inherited from Section 2 and are logically separate from these three effects.

3.3 Master projections and host substitutions

The next theorem records exactly what the three implementations provide. Its bounds are construction budgets, not asymptotic lower bounds; source hardness enters only after a pattern family and a host size have been chosen.

Theorem 3.4 (A: master projections). *There are universal constants $C_{\text{OV}}, C_{\oplus}, C_{\Sigma} > 0$ such that the following holds. Let $k \geq 2$, let P be a simple loopless pattern with $e(P) = k - 1$, and fix an orientation of its edges. Then the following depth-zero projections exist.*

- (i) *If $N^2 \leq n$ and $D \geq D_0(P, N, n)$, then $P\text{-SUB}_N \leq_{\text{proj}} k\text{-OV}_{n,D,k}$, where $D_0(P, N, n) \leq C_{\text{OV}}k \log(2n)$. Moreover, $D_0(P, N, n) \geq \lceil \log n \rceil$, so the hypothesis on D automatically ensures $n \leq 2^D$.*

(ii) If k is odd, $1 + e(P)N^2 \leq n$, and $m \geq m_{\oplus}(P, N)$, then $P\text{-SUB}_N \leq_{\text{proj}} k\text{-XOR}_{n,m,k}$, where $m_{\oplus}(P, N) \leq C_{\oplus}k \log(2N)$. The compressed guard labels are fixed nonuniformly, once for the pair (P, N) and before the source input is read.

(iii) If $1 + e(P)N^2 \leq n$ and $m \geq m_{\Sigma}(P, N)$, then $P\text{-SUB}_N \leq_{\text{proj}} k\text{-SUM}_{n,m,k}$, where $m_{\Sigma}(P, N) \leq C_{\Sigma}k \log(2kN)$. This projection has no parity restriction.

Proof. Part (i) is Theorem 4.2, part (ii) is Theorem 4.4, and part (iii) is Theorem 4.5. The exact definitions of $D_0(P, N, n)$, $m_{\oplus}(P, N)$, and $m_{\Sigma}(P, N)$ appear in the corresponding subsections, and Appendix C contains the complete correctness and projection proofs. The absolute estimates follow from Equations (15), (24), (25), (33), (35) and (37), after fixing the three displayed constants sufficiently large. $\square$

The theorem deliberately leaves the host size N free. In particular, it does not impose the common substitution $N = \Theta(\sqrt{n/k})$, because that would discard the stronger packing available for $k\text{-OV}$.

Lemma 3.5 (Standard host substitutions). *Let $k \geq 2$ and $n \geq 16k$, and define the $k\text{-OV}$ host size and the common list-target host size by*

$$N_{\text{OV}} := \lfloor \sqrt{n} \rfloor, \quad N_{\oplus, \Sigma} := \left\lfloor \sqrt{\frac{n-1}{k-1}} \right\rfloor. \quad (16)$$

For every simple loopless pattern P with $e(P) = k - 1$, these choices satisfy

$$N_{\text{OV}}^2 \leq n, \quad 1 + (k-1)N_{\oplus, \Sigma}^2 \leq n, \quad (17)$$

and

$$N_{\text{OV}} \geq \frac{1}{2}\sqrt{n}, \quad \log N_{\text{OV}} \geq \frac{1}{4} \log n, \quad (18)$$

$$N_{\oplus, \Sigma} \geq \frac{1}{2}\sqrt{\frac{n}{k}}, \quad \log N_{\oplus, \Sigma} \geq \frac{1}{4} \log \frac{n}{k}. \quad (19)$$

The constants in Theorem 3.4 may be fixed so that, for the same substitutions,

$$D_0(P, N_{\text{OV}}, n) \leq C_{\text{OV}}k \log n, \quad (20)$$

$$m_{\oplus}(P, N_{\oplus, \Sigma}) \leq C_{\oplus}k \log(en/k), \quad (21)$$

$$m_{\Sigma}(P, N_{\oplus, \Sigma}) \leq C_{\Sigma}k \log n. \quad (22)$$

If in addition $n \geq k^3$, then the last estimate sharpens to

$$m_{\Sigma}(P, N_{\oplus, \Sigma}) \leq C_{\Sigma}k \log(en/k). \quad (23)$$

Proof. The capacity inequalities are immediate from the definitions. Because $n/k \geq 16$, both quantities inside the floors are at least 4, and $\lfloor x \rfloor \geq x/2$ for $x \geq 2$. Moreover, $(n-1)/(k-1) \geq n/k$ when $n \geq k$, which gives the two lower bounds on the host sizes. Taking logarithms yields $\log N_{\text{OV}} \geq \frac{1}{2} \log n - 1 \geq \frac{1}{4} \log n$ and $\log N_{\oplus, \Sigma} \geq \frac{1}{2} \log(n/k) - 1 \geq \frac{1}{4} \log(n/k)$.

For the budgets, Theorem 3.4 gives the native bounds in terms of $\log(2n)$, $\log(2N)$, and $\log(2kN)$. The first is $O(\log n)$. Also $N_{\oplus, \Sigma} \leq \sqrt{n/(k-1)} \leq \sqrt{2n/k}$, so $\log(2N_{\oplus, \Sigma}) = O(\log(en/k))$ and $\log(2kN_{\oplus, \Sigma}) = O(\log n)$ uniformly for $n \geq 16k$. If $n \geq k^3$, then $kN_{\oplus, \Sigma} \leq \sqrt{2kn} \leq \sqrt{2}n/k$, and hence $\log(2kN_{\oplus, \Sigma}) = O(\log(en/k))$. Enlarging the three universal constants once establishes all four displayed inequalities. $\square$

The lemma is purely target-side: it packages the repeated floor and budget arithmetic but inserts no source lower bound. Each later theorem remains responsible for checking that its chosen host size lies beyond the onset required by its own source statement.

4 Three Instantiations

We now implement the roles of Section 3.2. The constructions deliberately use the same candidate labels and the same reference-edge comparisons, so that their differences are confined to the algebra of the target primitive.

A shared toy source (Figure 2). To make the common skeleton and the three target algebras directly comparable, all worked instances use $N = 2$, $k = 3$, and the oriented path P on vertices $\{p_1, p_2, p_3\}$ with $f_1 = (p_1, p_2)$ and $f_2 = (p_2, p_3)$. This choice is deliberate. A one-edge pattern would not activate the consistency role (R3). A triangle would have $e(P) = 3$ and hence $k = e(P) + 1 = 4$, outside the native odd- k k -XOR projection; the next compatible larger pattern has four edges and, already for $N = 2$, doubles the candidate columns from eight to sixteen. For the bitwise k -OV and k -XOR comparisons, encode $x \in [2]$ by $\text{bit}_1(x) = x - 1$; the k -SUM comparison instead uses the integer label x itself. Choose $\rho(p_1) = f_1$, $\rho(p_2) = f_1$, and $\rho(p_3) = f_2$. Then $\ell_N = 1$ and $\mathcal{I}(P) = \{(p_2, f_2)\}$, with $\text{lab}_{p_2}(f_1, u, w) = w$ and $\text{lab}_{p_2}(f_2, u, w) = u$. For compactness in the figures, write $X_{uw}^i := X_{f_i, u, w}$. The source assignment is shown in Figure 2; it has the unique accepting colored copy ϕ given by $(\phi(p_1), \phi(p_2), \phi(p_3)) = (2, 1, 2)$, while the additional present edge $X_{21}^2 = 1$ is inconsistent with the selected f_1 -edge at the shared vertex p_2 .




Figure 2: The shared source instance. Dashed gray host edges are absent, solid heavy edges form the unique colored copy, and the dash-dotted edge is present but inconsistent with that copy at pattern vertex p_2 . The labels inside the host vertices are indices in $[2]$, not their one-bit encodings. The right-hand table gives the complete assignment to the $e(P)N^2 = 8$ source bits X_{uw}^i .

The k -OV display in Figure 3 uses the minimal class capacity $n = N^2 = 4$, whereas the indexed-list displays in Figures 4 and 5 use $n = 12$: the nine genuine indices leave exactly three dummies, yielding a full all-dummy k -support that isolates the parity purge in k -XOR and the exact dummy counter in k -SUM. This deliberate difference reflects the capacities $N^2 \leq n$ and $1 + e(P)N^2 \leq n$. In all three displays, an entry $\overline{X}_{uw(b)}^i$ displays the projection literal $\overline{X}_{uw}^i$ together with its parenthesized value $b = 1 - X_{uw}^i$, as determined by the source-bit table in Figure 2; starred column headings mark the accepting target witness. Each subsection reads the starred witness and the diagnostic supports directly from its display. The role labels annotate how rows are used.

Role	coordinate	U_α				U_{f_1}				U_{f_2}			
		v_α^*	δ^1	δ^2	δ^3	v_1^{11}	v_1^{12}	v_1^{21*}	v_1^{22}	v_2^{11}	v_2^{12*}	v_2^{21}	v_2^{22}
R4	q_α	0	1	1	1	1	1	1	1	1	1	1	1
R2	q_{f_1}	1	1	1	1	$\overline{X}_{11(1)}^1$	$\overline{X}_{12(1)}^1$	$\overline{X}_{21(0)}^1$	$\overline{X}_{22(1)}^1$	1	1	1	1
R2	q_{f_2}	1	1	1	1	1	1	1	1	$\overline{X}_{11(1)}^2$	$\overline{X}_{12(0)}^2$	$\overline{X}_{21(0)}^2$	$\overline{X}_{22(1)}^2$
R3	$C_{p_2, f_2, 1}^{10}$	1	1	1	1	0	1	0	1	1	1	0	0
R3	$C_{p_2, f_2, 1}^{01}$	1	1	1	1	1	0	1	0	0	0	1	1
R5	id_1	0	0	0	0	0	0	1	1	0	0	1	1
R5	id_2	0	0	1	1	0	0	0	0	0	1	0	1
R5	id_3	0	1	0	1	0	1	0	1	0	0	0	0

Selected tuple	diagnostic	outcome
$\{v_\alpha, v_1^{21}, v_2^{12}\}$	every coordinate product is 0	accept
$\{v_\alpha, v_1^{11}, v_2^{12}\}$	product 1 on q_{f_1}	absent-edge rejection
$\{v_\alpha, v_1^{21}, v_2^{21}\}$	product 0 on $C_{p_2, f_2, 1}^{10}$ but 1 on $C_{p_2, f_2, 1}^{01}$	consistency rejection
$\{\delta^1, v_1^{21}, v_2^{12}\}$	product 1 on q_α	dummy rejection

Figure 3: The complete 8×12 matrix for the toy k -OV projection, where $v_i^{uw} := v_{f_i, u, w}$. Gray cells are fixed nonzero constants, and boxed cells display the literal $\overline{X}_{uw}^i$ together with the parenthesized value $1 - X_{uw}^i$ determined in Figure 2. The displayed consistency (R3) diagnostic realizes the disagreement $(0, 1)$, detected only by $C_{p_2, f_2, 1}^{01}$; the columns v_1^{12} and v_2^{12} realize the opposite disagreement on the consistency rows, although v_1^{12} also fails its guard (R2). The two columns v_1^{12} and v_1^{22} coincide on all five structural coordinates and are separated by id_1 , illustrating why the inert identifiers are needed even when pattern vertex p_1 participates in no comparison.

4.1 k -Orthogonal Vectors

An orthogonality coordinate rejects a selected tuple exactly when every chosen vector has entry 1 there. Selection rigidity is therefore automatic: a witness already chooses one vector from each class. For *consistency* (R3), A single coordinate cannot enforce equality of two compared label bits: after all inactive entries are fixed to 1, its product has the form $A(b_f)B(b_g)$ for functions $A, B: \{0, 1\} \rightarrow \{0, 1\}$. If this product were 1 on both disagreement patterns $(1, 0)$ and $(0, 1)$, then $A(0) = A(1) = B(0) = B(1) = 1$, so it would also be 1 on the agreement patterns $(0, 0)$ and $(1, 1)$. We therefore use one coordinate for each ordered disagreement pattern; the two consistency rows of Figure 3 make the asymmetry between those directions explicit.

Put $\ell_n := \lceil \log n \rceil$, and

$$s_{k,n} := \left\lceil \frac{k\ell_n}{k-1} \right\rceil, \quad D_0(P, N, n) := k + 2|\mathcal{I}(P)|\ell_N + s_{k,n}. \quad (24)$$

Here ℓ_N counts label bits, whereas ℓ_n counts identifier bits for *promise validity* (R5). Since $k \geq 2$, one has $s_{k,n} \leq 2\ell_n$; together with Equation (15), the inequality $N^2 \leq n$ implies

$$D_0(P, N, n) = O(k \log(2N) + \log(2n)) = O(k \log(2n)). \quad (25)$$

The k target classes, i.e., the k input arrays of Definition 3.1, are an anchor class U_α and one edge class U_f for every $f \in E(P)$. As explained in Section 3.2, the anchor is needed here only to supply the k th target class: the $e(P) = k - 1$ edge groups provide the other classes, while one-vector-per-class selection is already built into k -OV. The anchor class contains one genuine vector v_α , while U_f contains the N^2 genuine vectors $v_{f,u,w}$ indexed by $(u, w) \in [N]^2$; each class is padded to cardinality n with dummies. The three class blocks in Figure 3 make selection rigidity (R1) visible: the starred witness chooses exactly one column from each block.

Worked instance. Specialize first to Figure 2. Taking $n = 4$ gives $\ell_n = 2$, $s_{3,4} = 3$, and $D_0(P, 2, 4) = 8$. Since $n = N^2$, the edge classes U_{f_1} and U_{f_2} are already full, so only the anchor class U_α contains dummies. Let the three identifier coordinates be owned by $U_\alpha, U_{f_1}, U_{f_2}$, respectively, and assign the four two-bit strings on the two unpinned coordinates of each class. This is the tight balanced assignment from Lemma 4.1 below: every class owns one coordinate and is unpinned on exactly $\ell_n = 2$ coordinates.

Guards and dummy purge (R2, R4). There is one coordinate q_h for every class $h \in \{\alpha\} \cup E(P)$. On q_α , the genuine anchor has entry 0 and every other vector has entry 1. On q_f , a genuine vector $v_{f,u,w}$ has entry $\overline{X}_{f,u,w}$, while every vector in every other class and every dummy in U_f has entry 1. Thus a genuine edge vector satisfies its guard exactly when the proposed edge is present, and any dummy causes the guard of its own class to fail. In Figure 3, the starred present vectors v_1^{21} and v_2^{12} contribute 0 on their own edge guards, whereas replacing v_1^{21} by the absent candidate v_1^{11} leaves product 1 on q_{f_1} . Likewise, replacing the anchor by δ^1 leaves product 1 on q_α . Every dummy is set to 1 on every guard and consistency coordinate. Zero-vector padding would instead make every selected tuple containing a dummy automatically orthogonal, so we use the all-ones structural choice; Fact C.1 gives the one-line failure proof.

Consistency (R3). Fix $(a, g) \in \mathcal{I}(P)$, write $f := \rho(a)$, and let $t \in [\ell_N]$. On the coordinate $C_{a,g,t}^{10}$, a candidate in class f carries $\text{bit}_t(\text{lab}_a(f, u, w))$, a candidate in class g carries $1 - \text{bit}_t(\text{lab}_a(g, u, w))$, and every vector in every other class has entry 1. The coordinate $C_{a,g,t}^{01}$ complements the two active entries in the opposite way. For selected label bits $b_f, b_g \in \{0, 1\}$, the two coordinate products on the selected tuple are $b_f(1 - b_g)$ and $(1 - b_f)b_g$, so both vanish exactly when $b_f = b_g$. In Figure 3, the starred candidates both propose label 1 at p_2 , and both consistency products vanish. Replacing v_2^{12} by the present near-miss v_2^{21} changes the second proposal to 2: the product remains 0 on $C_{p_2,f_2,1}^{10}$ but becomes 1 on $C_{p_2,f_2,1}^{01}$. The entries of v_1^{12} and v_2^{12} show the reverse bit disagreement, which is detected by $C_{p_2,f_2,1}^{10}$. Ranging over all t and all $(a, g) \in \mathcal{I}(P)$ enforces full-label agreement at every non-isolated pattern vertex.

Promise-valid identifiers (R5). The structural coordinates need not distinguish all vectors within a class. In Figure 3, for example, v_1^{12} and v_1^{22} coincide on all five structural coordinates but are separated by id_1 . The ownership pattern is also visible: U_α is pinned to 0 on id_1 , U_{f_1} on id_2 , and U_{f_2} on id_3 , so every one-vector-per-class tuple has a zero factor on every identifier coordinate; in particular, the starred witness remains orthogonal. Simply appending unpinned binary identifiers

need not have this inertness property, as Fact C.2 records. We therefore use identifiers that are inert for every selected tuple.

Lemma 4.1 (Inert identifiers). *There are $s_{k,n}$ identifier coordinates such that every coordinate is pinned to 0 throughout one owner class, while every class is unpinned on at least ℓ_n coordinates. The n vectors of each class can consequently be assigned pairwise distinct identifiers, and the product over every selected tuple is identically zero on each identifier coordinate.*

The value $s_{k,n}$ is tight for this ownership scheme. With s identifier coordinates, each class can own at most $s - \ell_n$ of them if it is to remain unpinned on at least ℓ_n coordinates; assigning every coordinate an owner therefore requires $s \leq k(s - \ell_n)$, whose least integer solution is $s = s_{k,n}$. The complete proof and an explicit balanced ownership assignment appear in Appendix C.1. Combining the structural coordinates with these identifiers gives genuine promised inputs with n vectors in each class, without changing which selected tuples are orthogonal.

Theorem 4.2 (k -OV projection). *Let $k \geq 2$ and let P be a simple loopless pattern with $e(P) = k-1$. For every N, n with $N^2 \leq n$ and every $D \geq D_0(P, N, n)$, the construction above gives a depth-zero projection*

$$P\text{-SUB}_N \leq_{\text{proj}} k\text{-OV}_{n,D,k}. \quad (26)$$

Equivalently, the constructed k -OV instance has a one-vector-per-class orthogonal witness if and only if the source instance contains a colored copy of P . Since $s_{k,n} \geq \ell_n$, the hypothesis $D \geq D_0(P, N, n)$ also implies $n \leq 2^D$, so the projection lands in the promise domain of $k\text{-OV}_{n,D,k}$.

The proof in Appendix C.1 follows the five roles directly, in the same order displayed by Figure 3. Every output bit is a constant or a negated source variable. Dimensions above $D_0(P, N, n)$ are obtained by appending all-zero coordinates.

4.2 Odd k -XOR

Retain ℓ_N and the binary label encoding fixed in Section 3.2. For k -XOR, all objects lie in one indexed list, so selection rigidity must be enforced algebraically. The construction has one anchor column, one candidate column $c_{f,u,w}$ for each $f \in E(P)$ and $(u, w) \in [N]^2$, and dummy columns padding the list to length n . Repeated column values are permitted because the target selects indices; the three identical dummy columns in Figure 4 therefore require no analogue of role (R5).

We now define the general rows and use the displayed support sums to track what each role enforces.

Selection and dummy purge (R1, R4). For each $f \in E(P)$, an selector row (R1) is assigned to its edge group. As shown in Figure 4, each selector row contains two 1's on the starred support, while the anti-dummy row (R4) contains none, so all three sums vanish. The all-dummy support also vanishes on both selectors but has anti-dummy sum $3 \equiv 1 \pmod{2}$. In general, for a selected support S , let $\zeta \in \{0, 1\}$ indicate whether the anchor is selected, let c_f be the number of selected candidates from group f , and let d be the number of selected dummies. A selector row for f has entry 1 on the anchor and every group- f candidate and entry 0 elsewhere, so a zero sum of this row imposes

$$\zeta + c_f \equiv 0 \pmod{2}. \quad (27)$$

The anti-dummy row has entry 1 on every dummy and 0 on every genuine column. Together with the selector equations, the support-size identity is

$$\zeta + \sum_{f \in E(P)} c_f + d = k = e(P) + 1. \quad (28)$$

Role	row	α^*	c_1^{11}	c_1^{12}	c_1^{21*}	c_1^{22}	c_2^{11}	c_2^{12*}	c_2^{21}	c_2^{22}	δ_1	δ_2	δ_3
R1	S_{f_1}	1	1	1	1	1	0	0	0	0	0	0	0
R1	S_{f_2}	1	0	0	0	0	1	1	1	1	0	0	0
R2	H_1	0	0	0	0	0	$\overline{X}_{11(1)}^2$	$\overline{X}_{12(0)}^2$	$\overline{X}_{21(0)}^2$	$\overline{X}_{22(1)}^2$	0	0	0
R2	H_2	0	$\overline{X}_{11(1)}^1$	$\overline{X}_{12(1)}^1$	$\overline{X}_{21(0)}^1$	$\overline{X}_{22(1)}^1$	0	0	0	0	0	0	0
R2	H_3	0	0	0	$\overline{X}_{21(0)}^1$	$\overline{X}_{22(1)}^1$	0	0	$\overline{X}_{21(0)}^2$	$\overline{X}_{22(1)}^2$	0	0	0
R2	H_4	0	0	$\overline{X}_{12(1)}^1$	0	$\overline{X}_{22(1)}^1$	0	$\overline{X}_{12(0)}^2$	0	$\overline{X}_{22(1)}^2$	0	0	0
—	H_5	0	0	0	0	0	0	0	0	0	0	0	0
—	H_6	0	0	0	0	0	0	0	0	0	0	0	0
—	H_7	0	0	0	0	0	0	0	0	0	0	0	0
R3	$C_{p_2, f_2, 1}$	0	0	1	0	1	0	0	1	1	0	0	0
R4	A_{dum}	0	0	0	0	0	0	0	0	0	1	1	1

Selected support	unique nonzero row sum	reading
$\{\alpha, c_1^{21}, c_2^{12}\}$	none	accept
$\{\alpha, c_1^{11}, c_2^{12}\}$	$H_2 = 1$	one absent edge
$\{\alpha, c_1^{21}, c_2^{21}\}$	$C_{p_2, f_2, 1} = 1$	inconsistent shared label
$\{\delta_1, \delta_2, \delta_3\}$	$A_{\text{dum}} = 1$	odd dummy support

Figure 4: An worked instance specialized to the source assignment in Figure 2. The complete 11×12 matrix for the toy k -XOR projection with odd k , where $c_i^{uw} := c_{f_i, u, w}$. Gray cells are fixed nonzero constants, and boxed cells display the literal $\overline{X}_{uw}^i$ together with the parenthesized value $1 - X_{uw}^i$ determined in Figure 2. The starred column headings form the zero-sum witness. Each rejected diagnostic support has exactly the displayed nonzero row sum.

These equations prove support rigidity when k is odd. If $\zeta = 1$, every c_f is a positive odd integer by (27). Then (28) forces $c_f = 1$ for every f and $d = 0$. If $\zeta = 0$, every c_f is even and the zero-sum anti-dummy row makes d even, so the support size is even, contradicting that k is odd. Thus every zero-sum k -support contains the anchor, exactly one candidate from each edge group, and no dummies.

Compressed guards (R2). The guard must reject every support containing at least one absent edge. Assigning private guard rows to every columns would do so but would cost $e(P)N^2$ rows. Instead, a standard random-labelling argument with a union bound compresses them into

$$r_{\text{guard}}(P, N) := \lceil 2e(P) \log N + e(P) + 1 \rceil \quad (29)$$

rows.

Lemma 4.3 (Separating guard labels). *There are labels $\lambda_{f,u,w} \in \mathbb{F}_2^{r_{\text{guard}}(P,N)}$ such that, for every choice $(u_f, w_f) \in [N]^2$ for each $f \in E(P)$ and every nonempty $T \subseteq E(P)$,*

$$\bigoplus_{f \in T} \lambda_{f,u_f,w_f} \neq 0. \quad (30)$$

The labels depend only on (P, N) and may be fixed nonuniformly before the source input is read.

The probabilistic proof is in Appendix C.2. At the integer value in Equation (29), its union bound is

$$N^{2e(P)} 2^{e(P)} 2^{-r_{\text{guard}}(P,N)} \leq \frac{1}{2}. \quad (31)$$

Using these labels, write $r_g := r_{\text{guard}}(P, N)$. For each $h \in [r_g]$, the entry of candidate $c_{f,u,w}$ is $\bar{X}_{f,u,w} \lambda_{f,u,w}(h)$, and the anchor and all dummies have entry 0. Present candidates contribute 0, while the selected absent candidates contribute the XOR of their labels, which is nonzero by Equation (30).

In Figure 4, we use the following explicit seven-bit guard labels

$$\lambda_{f_1,u,w} := (0, 1, u-1, w-1, 0, 0, 0), \quad \lambda_{f_2,u,w} := (1, 0, u-1, w-1, 0, 0, 0). \quad (32)$$

Both of the two labels are nonzero, and the XOR of one label from each edge group begins with $(1, 1)$, so every nonempty subset of a selected two-group tuple has nonzero label XOR. For the active rows $H_1, \dots, H_4$ in Figure 4, the starred present candidates vanish throughout them, whereas the absent candidate c_1^{11} contributes 1 on H_2 , giving the displayed guard rejection. The rows H_5, H_6, H_7 are zero padding, allowing the example to retain the $r_g = 7$ general guard budget while displaying only four active guard rows.

Consistency (R3). For every $(a, g) \in \mathcal{I}(P)$, with $f := \rho(a)$, and every $t \in [\ell_N]$, one row places $\text{bit}_t(\text{lab}_a(\cdot))$ on every candidate in groups f and g and places 0 elsewhere. Once the selectors have chosen one candidate from each group, the row sum is the XOR of the two proposed label bits and vanishes exactly when they agree. In Figure 4, c_1^{21} and c_2^{12} both propose bit 0 at p_2 , whereas replacing c_2^{12} by the present candidate c_2^{21} changes that row sum to 1. Unlike orthogonality, one $\mathbb{F}_2$ equation forbids both disagreement patterns.

The exact constructed row budget is

$$m_{\oplus}(P, N) := e(P) + r_{\text{guard}}(P, N) + |\mathcal{I}(P)|\ell_N + 1 = O(k \log(2N)), \quad (33)$$

where the four terms are the selectors, compressed guards, consistency rows, and anti-dummy row, respectively.

Theorem 4.4 (k -XOR projection for odd k). *Let $k \geq 3$ be odd, and let P be a simple loopless pattern with $e(P) = k - 1$. For every N, n with $1 + e(P)N^2 \leq n$ and every $m \geq m_{\oplus}(P, N)$, the construction above gives a depth-zero projection*

$$P\text{-SUB}_N \leq_{\text{proj}} k\text{-XOR}_{n,m,k}. \quad (34)$$

The complete proof is in Appendix C.2. It shows that a zero-sum support uses only present candidates and is label-consistent, and it checks that every matrix entry is a constant or a negated source variable. Larger row counts are reached by appending all-zero rows.

The oddness condition is an artifact of exact selection over $\mathbb{F}_2$, not of colored subgraph isomorphism. The all-dummy diagnostic in Figure 4 pinpoints the issue. Whenever at least k dummy indices are available, the all-dummy support already satisfies every row. The black-box lifting projection is Lemma E.2; Corollary E.4 gives the quantitative consequence for any supplied admissible odd source parameter. The body-level bound is Equation (50).

4.3 k -SUM

The k -SUM construction works natively over the integers. A black-box additive, that is, group-homomorphic, encoding of the $\mathbb{F}_2$ construction into $\mathbb{Z}_{2^m}$ cannot preserve all zero relations; Proposition C.8 gives the formal obstruction. We instead encode selection, consistency, presence, and dummy rejection by exact signed coordinates and then pack those coordinates into one residue.

The indexed list again contains the anchor, all $e(P)N^2$ candidates, and repeated dummy residues. Because the target selects indices rather than distinct values, those repetitions are legal and no analogue of role (R5) is needed. Associate with every index j an integer vector $\text{vec}(j) \in \mathbb{Z}^{R(P)}$, where

$$R(P) := e(P) + |\mathcal{I}(P)| + 2 \leq 3e(P) + 2 = O(k). \quad (35)$$

For a selected support S , let $\zeta \in \{0, 1\}$ indicate whether the anchor is selected, let c_f count selected candidates from group f , and let d count selected dummies. We first display the resulting toy instance, then define the coordinates in general and read them against it.

Worked instance. Specialize again to the source assignment in Figure 2 and take $n = 12$. Unlike the bitwise comparisons in the first two instantiations, the consistency coordinate now uses the label values 1, 2 themselves, accounting for the entries $\pm 1, \pm 2$. Choose the coordinate order

$$(G_{f_1}, G_{f_2}, C_{p_2, f_2}, Q, Q_{\text{dum}}). \quad (36)$$

We now define the exact coordinates in general and read each one against the corresponding row of the figure.

Group coordinates (R1). We assign one group coordinate G_f for each $f \in E(P)$, where it has -1 on the anchor, $+1$ on every group- f candidate, and 0 elsewhere. On a selected support it equals $c_f - \zeta$, so vanishing enforces the exact equality $c_f = \zeta$ over $\mathbb{Z}$. In the first two rows of Figure 5, the starred anchor contributes -1 and the unique starred candidate from the corresponding group contributes $+1$. By contrast, the selection-violating support $\{c_1^{11}, c_1^{12}, \delta_1\}$ has G_{f_1} -sum 2 , so it is rejected even before the remaining nonzero coordinates are inspected.

Consistency coordinates (R3). For $(a, g) \in \mathcal{I}(P)$, write $f := \rho(a)$. The coordinate $C_{a, g}$ assigns $+\text{lab}_a(g, u, w)$ to a group- g candidate, $-\text{lab}_a(f, u, w)$ to a group- f candidate, and 0 elsewhere. Once one candidate is selected from each group, it is the signed difference of the two complete labels and vanishes exactly when they agree. In Figure 5, the starred columns contribute -1 from c_1^{21} and $+1$ from c_2^{12} , while the consistency near-miss c_2^{21} contributes $+2$ and leaves sum 1 . Thus one integer coordinate replaces the $\Theta(\log N)$ bitwise comparisons used by k -OV and k -XOR, although the carry-safe packing in Equation (37) later allocates $\Theta(\log(2kN))$ bits per integer coordinate.

Role	coordinate	α^*	c_1^{11}	c_1^{12}	c_1^{21*}	c_1^{22}	c_2^{11}	c_2^{12*}	c_2^{21}	c_2^{22}	δ_1	δ_2	δ_3
R1	G_{f_1}	-1	1	1	1	1	0	0	0	0	0	0	0
R1	G_{f_2}	-1	0	0	0	0	1	1	1	1	0	0	0
R3	C_{p_2, f_2}	0	-1	-2	-1	-2	1	1	2	2	0	0	0
R2	Q	0	$\overline{X}_{11(1)}^1$	$\overline{X}_{12(1)}^1$	$\overline{X}_{21(0)}^1$	$\overline{X}_{22(1)}^1$	$\overline{X}_{11(1)}^2$	$\overline{X}_{12(0)}^2$	$\overline{X}_{21(0)}^2$	$\overline{X}_{22(1)}^2$	0	0	0
R4	Q_{dum}	0	0	0	0	0	0	0	0	0	1	1	1

standard word, most to least significant	bits 19:16	bits 15:12	bits 11:8	bits 7:4	bits 3:0
power used by Φ	B^4	B^3	B^2	B^1	B^0
coordinate carrying that weight	Q_{dum}	Q	C_{p_2, f_2}	G_{f_2}	G_{f_1}

standard residue	α^*	c_1^{11}	c_1^{12}	c_1^{21*}	c_1^{22}	c_2^{11}	c_2^{12*}	c_2^{21}	c_2^{22}	δ_1	δ_2	δ_3
hexadecimal	FFFEF	00F01	00E01	FFF01	00E01	01110	00110	00210	01210	10000	10000	10000

Selected support	exact signed sum y_S	hexadecimal residue
$\{\alpha, c_1^{21}, c_2^{12}\}$	(0, 0, 0, 0, 0)	00000
$\{\alpha, c_1^{21}, c_2^{21}\}$	(0, 0, 1, 0, 0)	00100
$\{\alpha, c_1^{11}, c_2^{12}\}$	(0, 0, 0, 1, 0)	01000
$\{\delta_1, \delta_2, \delta_3\}$	(0, 0, 0, 0, 3)	30000
$\{c_1^{11}, c_1^{12}, \delta_1\}$	(2, 0, -3, 2, 1)	11D02

Figure 5: The full toy k -SUM construction. The first table is the primary signed-coordinate object; gray cells are fixed nonzero constants, while boxed cells display the literal $\overline{X}_{uw}^i$ together with the parenthesized value $1 - X_{uw}^i$ determined in Figure 2. The second table fixes how the signed coordinates are weighted, and the third lists all twelve target residues exactly as five hexadecimal digits, hence as twenty bits; these are standard residue digits, not the signed coordinate values themselves. The starred column headings form a witness. The diagnostic rows respectively show acceptance, a pure consistency violation, a pure guard violation, dummy rejection, and a selection violation witnessed already by $G_{f_1} = 2$. Different roles can fail simultaneously, so the diagnostics identify a decisive nonzero coordinate rather than claiming a partition of failures.

Guard and dummy coordinates (R2, R4). The guard coordinate Q assigns $\overline{X}_{f,u,w} \in \{0, 1\}$ to candidate $c_{f,u,w}$ and 0 to the anchor and dummies. Its selected sum is nonnegative and vanishes exactly when every selected edge is present. In Figure 5, replacing c_1^{21} by the absent candidate c_1^{11} leaves the pure guard sum $Q = 1$. The dummy coordinate Q_{dum} is 1 on dummies and 0 on genuine objects, so it counts selected dummies exactly; the all-dummy diagnostic has sum 3 in that row and zero in every lower signed coordinate.

These exact coordinates already force the desired support shape. If all coordinate sums vanish, then the dummy coordinate gives $d = 0$ and every group coordinate gives $c_f = \zeta$, so the support-size identity becomes $k = \zeta + \sum_{f \in E(P)} c_f + d = k\zeta$. Hence $\zeta = 1$, every $c_f = 1$, and no dummy is selected.

To pack these signed vectors, we use the standard large-base positional-encoding device introduced in Karp’s reduction from EXACT COVER to KNAPSACK [Kar72], where the base is chosen to exceed the largest attainable digit so that no carry propagates between coordinates. We use it in signed form, with digits in $[-kN, kN]$ and a power-of-two base, and correct for the modular ambient group in Equation (40). Set

$$q := \lceil \log(2kN + 1) \rceil, \quad B := 2^q, \quad m_\Sigma(P, N) := qR(P). \quad (37)$$

Fix an ordering of the $R(P)$ coordinates and define

$$\Phi(z_0, \dots, z_{R(P)-1}) := \sum_{i=0}^{R(P)-1} z_i B^i. \quad (38)$$

Every k -support sum $y_S := \sum_{j \in S} \text{vec}(j)$ satisfies $\|y_S\|_\infty \leq kN$, while $B > 2kN$. Consequently the balanced base- B representation is faithful:

$$\Phi(y_S) \equiv 0 \pmod{2^{m_\Sigma(P,N)}} \iff y_S = 0. \quad (39)$$

The proof is a digit induction given in Lemma C.11.

For each index j , let $\hat{a}_j := \Phi(\text{vec}(j)) \bmod 2^{m_\Sigma(P,N)}$ be the standard residue in $[0, 2^{m_\Sigma(P,N)})$. When the requested bit-width is $m \geq m_\Sigma(P, N)$, define

$$a_j := 2^{m-m_\Sigma(P,N)} \hat{a}_j \in [0, 2^m). \quad (40)$$

Simply viewing $\hat{a}_j$ as a low-bit m -bit integer is not sound: negative exact packed values wrap modulo $2^{m_\Sigma(P,N)}$, and a true witness can leave a nonzero multiple of that smaller modulus. The top-bit scaling in Equation (40) preserves divisibility exactly, because $2^m \mid 2^{m-m_\Sigma(P,N)}Z$ if and only if $2^{m_\Sigma(P,N)} \mid Z$.

In the worked instance shown in Figure 5, $R(P) = 5$, $q = \lceil \log 13 \rceil = 4$, $B = 16$, and $m_\Sigma(P, 2) = 20$. Thus G_{f_1} is the least significant signed coordinate and Q_{dum} is the most significant one. The lower panels of Figure 5 make the packing and modular wraparound explicit. In hexadecimal notation, the starred residues satisfy

$$\text{FFFEF}_{16} + \text{FFF01}_{16} + \text{00110}_{16} = \text{200000}_{16} = 2 \cdot 2^{20} \equiv 0 \pmod{2^{20}}. \quad (41)$$

Thus a valid exact zero need not sum to the integer 0 after each signed column is replaced by its standard residue; it sums to a multiple of the target modulus, as required.

The candidate c_1^{11} also shows why the signed table must not be identified with the standard digit blocks. Its signed vector is $(1, 0, -1, 1, 0)$, but

$$\Phi(1, 0, -1, 1, 0) = 3841 = \text{00F01}_{16}, \quad (42)$$

whose base-16 digits from least to most significant are $(1, 0, 15, 0, 0)$: the negative consistency coordinate borrows from the guard coordinate. Finally, for a requested width $m > 20$, Equation (40) multiplies every displayed residue by 2^{m-20} ; merely prefixing the twenty-bit words with zeros is the unsound low-bit extension ruled out by Remark C.13.

Theorem 4.5 (k -SUM projection). *Let $k \geq 2$ and let P be a simple loopless pattern with $e(P) = k - 1$. For every N, n with $1 + e(P)N^2 \leq n$ and every $m \geq m_\Sigma(P, N)$, the construction above gives a depth-zero projection*

$$P\text{-SUB}_N \leq_{\text{proj}} k\text{-SUM}_{n,m,k}. \quad (43)$$

Moreover, $m_\Sigma(P, N) = O(k \log(2kN))$, and the projection holds for both parities of k .

The complete proof is in Appendix C.3. The exact support argument above leaves one candidate from every group and no dummies; consistency and the nonnegative guard then recover a colored embedding. Although the packing uses ordinary integer carries, every packed candidate depends on only the single bit $X_{f,u,w}$, so each output bit is one of two constants as that bit varies and is therefore a constant, $X_{f,u,w}$, or $\overline{X}_{f,u,w}$.

5 Unconditional Lower Bounds for Fixed k

For fixed k , the strongest target consequences use the standard host substitutions of Lemma 3.5: $N_{\text{OV}} = \lfloor \sqrt{n} \rfloor$ for k -OV and $N_{\oplus, \Sigma} = \lfloor \sqrt{(n-1)/(k-1)} \rfloor$ for the two list targets. The first is available because every edge group occupies its own k -OV class, whereas the second is forced because all $k-1$ edge groups share one indexed list in k -XOR and k -SUM. Thus the base- n versus base- (n/k) distinction below is a target-packing effect, not a difference in source hardness.

Although the depth-two result Theorem 6.4 already applies to sufficiently large fixed k , and the top-disjunction depth-three result Theorem 6.8 does so as well, Theorem 5.1 is not subsumed by these later results. Its distinctive content is an exponent linear in k at every fixed depth, with one target-side rate uniform in both d and k .

Theorem 5.1 (B: unconditional lower bounds for fixed k). *Let $k_0 \in \mathbb{N}$ be the universal threshold of Lemma 2.5, and let $C_{\text{OV}}, C_{\oplus}, C_{\Sigma}$ be the universal construction constants fixed in Theorem 3.4 and Lemma 3.5. There is a universal constant $\beta > 0$, independent of k , d , and n , such that for every fixed depth d and every fixed integer $k \geq k_0$ there is a threshold $n_1(d, k)$ for which the following hold for every $n \geq n_1(d, k)$:*

$$\text{size}_d(k\text{-OV}_{n,D,k}) \geq n^{\beta(k-1)} \quad \text{for every } D \geq C_{\text{OV}}k \log n, \quad (44)$$

$$\text{size}_d(k\text{-XOR}_{n,m,k}) \geq (n/k)^{\beta(k-1)} \quad \text{for odd } k \text{ and every } m \geq C_{\oplus}k \log(en/k), \quad (45)$$

$$\text{size}_d(k\text{-SUM}_{n,m,k}) \geq (n/k)^{\beta(k-1)} \quad \text{for every } m \geq C_{\Sigma}k \log(en/k). \quad (46)$$

The same rate β works simultaneously for every fixed d and every fixed $k \geq k_0$; only the onset threshold depends on (d, k) .

Proof. Let c_{κ} be the constant of Lemma 2.5, and choose universal constants

$$0 < 2\beta < \gamma < c_{\kappa}. \quad (47)$$

Apply Corollary 2.6 with the source rate γ , so that for every fixed (d, k) and all $N \geq N_0(\gamma, d, k)$,

$$\text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma(k-1)}. \quad (48)$$

For k -OV, put $N = N_{\text{OV}}$. Enlarge $n_1(d, k)$ so that $n \geq 16k$, $N \geq N_0(\gamma, d, k)$, and $n^{\gamma/2-\beta} \geq 4^{\gamma/2}$. By Equations (17), (18) and (20), the k -OV projection is admissible at every $D \geq C_{\text{OV}}k \log n$, and Lemma 3.3 and eq. (48) give

$$\text{size}_d(k\text{-OV}_{n,D,k}) \geq N^{\gamma(k-1)} \geq (n/4)^{\gamma(k-1)/2} \geq n^{\beta(k-1)}.$$

For k -XOR and k -SUM, put $N = N_{\oplus, \Sigma}$. Enlarge $n_1(d, k)$ further so that $n \geq \max\{16k, k^3\}$, $N \geq N_0(\gamma, d, k)$, and $(n/k)^{\gamma/2-\beta} \geq 4^{\gamma/2}$. Then Equations (17), (19), (21) and (23) make the respective projections admissible at the displayed budgets. The odd- k projection for k -XOR and the parity-free projection for k -SUM, followed by Equations (19) and (48), yield

$$\text{size}_d(Q_{n,m,k}) \geq N^{\gamma(k-1)} \geq \left(\frac{n}{4k}\right)^{\gamma(k-1)/2} \geq (n/k)^{\beta(k-1)}, \quad Q \in \{k\text{-XOR}, k\text{-SUM}\},$$

with the oddness qualification in the k -XOR case. Enlarging the single threshold $n_1(d, k)$ to satisfy all preceding requirements proves the theorem. $\square$

Table 5: Same-depth exponent-scale comparisons at the principal fixed- k budgets. Each lower bound is specialized to top-disjunction depth three, and all elementary upper bounds are Σ_3 circuits; the k -XOR row uses the direct odd- k statement.

Target	Transferred lower bound	Elementary Σ_3 upper bound	Exponent-scale verdict
k -OV	$n^{\beta(k-1)}$	$n^{k+o(1)}$	same-orientation match at Σ_3 and scale $n^{\Theta(k)}$
k -XOR	$(n/k)^{\beta(k-1)}$	$(n/k)^{k+o(1)}$	same-orientation match at Σ_3 and scale $(n/k)^{\Theta(k)}$
k -SUM	$(n/k)^{\beta(k-1)}$	$(n/k)^{O(k)}$	same-orientation match at Σ_3 ; scale $(n/k)^{\Theta(k)}$, with only $O(k)$ on the upper side

The substantive uniformity assertion in Theorem 5.1 is its quantifier order. Suppressing only the constraint parameters and retaining the odd- k qualification for k -XOR, it is

$$\begin{aligned}
& \exists \beta > 0, k_0 \quad \forall d \quad \forall k \geq k_0 \quad \exists n_1(d, k) \quad \forall n \geq n_1(d, k) : \\
& \text{size}_d(k\text{-OV}) \geq n^{\beta(k-1)}, \quad \text{size}_d(k\text{-SUM}) \geq (n/k)^{\beta(k-1)}, \\
& k \text{ odd} \implies \text{size}_d(k\text{-XOR}) \geq (n/k)^{\beta(k-1)}.
\end{aligned} \tag{49}$$

In particular, the constant hidden by the customary notation $n^{\Omega(k)}$ is not allowed to deteriorate with either k or the depth. The proof obtains this target rate β from the universal source rate γ under the explicit slack condition $2\beta < \gamma$; Remark B.6 records the corresponding source-side quantifier form. The factor e inside $\log(en/k)$ merely keeps the displayed budget positive in small boundary regimes; it has no asymptotic effect once $n/k \rightarrow \infty$.

The direct k -XOR projection, and hence Equation (45), is restricted to odd k . The odd-to-even projection in Lemma E.2 applies whenever an even target parameter K is written as

$$K = tr + s, \quad t \geq 2, \quad r \geq k_0 \text{ odd}, \quad 0 \leq s < 2t, \quad s = 0 \text{ or } t \nmid s.$$

For every such supplied decomposition, Corollary E.4 gives a universal constant $C'_\oplus > 0$ such that every fixed depth d admits a threshold $n'_1(d, K, t, r, s)$ for which

$$\text{size}_d(k\text{-XOR}_{n', m', K}) \geq (n'/K)^{\beta(r-1)/2} \tag{50}$$

for all $n' \geq n'_1(d, K, t, r, s)$ and every $m' \geq C'_\oplus K \log n'$, where β is the universal rate in Theorem 5.1.

Theorems D.2 and D.3 give the corresponding lower bounds away from the principal $O(k \log n)$ -scale row and width budgets of Theorem 5.1. At the principal budgets, the elementary circuits of Propositions F.1 to F.3 give the lower-vs-upper comparison in Table 5.

Proposition F.4 shows that the specific κ -based regular-core-plus-matching route used for Theorem 5.1 remains separated from the brute-force leading exponent by a factor greater than the core degree; Remark F.5 delimits this as a limitation of that route rather than a general barrier.

For k -SUM, Proposition F.3 retains the depth-two DNF of size $2^{\Theta(km)}$ but also gives a top-disjunction depth-three block-carry circuit. At the principal width $m = \Theta(k \log(en/k))$ and in the asymptotic fixed- k regime, the block-carry circuit has size $(n/k)^{O(k)}$ and closes, for Σ_3 , the exponent-scale gap created by the direct $2^{\Theta(km)}$ verifier. It does not provide a same-depth upper bound at depth two or for Π_3 . The general-width estimates in Equations (203) and (204)

separate the residual support-enumeration gap at smaller widths from the carry overhead at very large widths. The parity and carry mechanisms underlying these comparisons are synthesized in Section 8.

6 Unconditional Lower Bounds for Growing k

The fixed-pattern theorem used in Section 5 does not control its onset threshold as the pattern grows. This section therefore changes the source lower bound while leaving every target projection unchanged. A padded clique pattern gives an unconditional floor at every fixed depth. On the padded expander family already used for fixed k , direct DNF-term and CNF-clause counting give an exponent linear in k at depth two, while a single-CNF minterm argument gives such an exponent for the top-disjunction orientation at depth three. The section closes by summarizing the unconditional frontier and comparing it with a complementary row-rich PARITY route.

6.1 A clique fallback at every fixed depth⁴

A clique packs many pattern edges into few pattern vertices. This vertex deficiency weakens the exponent from linear in k to order $\sqrt{k}$, but it also lets us invoke a source lower bound that is uniform while the clique size grows. The only other restriction is the small-clique range of Beame’s theorem, which produces the plateau at $k = \Theta(\log^2 n)$.

The imported PRAM statement and notation. Beame works with PRIORITY CRCW PRAMs⁵ on an N -node graph represented by $\binom{N}{2}$ Boolean input cells, one per unordered vertex pair; every non-input cell is initially 0, and the answer must occupy a designated cell C_1 after the final step. No computability condition is imposed on the transition functions, so the model is nonuniform. His function Clique_r^N is the total monotone function accepting exactly the N -node graphs containing an r -clique, hence it is the function $K_r\text{-CLIQUE}_N$ of Definition B.7. For a depth- d circuit with G non-input gates, the pointwise dictionary is

$$\begin{aligned} n_{\text{Bea90}} &= N, & m_{\text{Bea90}} &= \binom{N}{2}, & k_{\text{Bea90}} &= r, \\ T_{\text{Bea90}} &= T \leq d + 2, & c(n)_{\text{Bea90}} &= G + \binom{N}{2} + O(1), & \text{Clique}_k^n &= K_r\text{-CLIQUE}_N. \end{aligned}$$

Throughout [Bea90, Sec. 2], n is the number of graph vertices and $m = \binom{n}{2}$ is the number of input variables; the phrase “ n inputs” in Corollary 3.8 is therefore a notational slip. Theorem 3.1(b) of [Bea90] states the memory-cell bound $c(N) \geq N^{r/(43T^3)}$, even with infinitely many processors, for $r \leq \log N$, for r above one absolute threshold, and for all sufficiently large N ; its proof also uses $c(N) \geq N$. We quote this PRAM statement and, for the constant, the three terminal cases of its proof. We do not use [Bea90, Cor. 3.8]: that corollary does not define its circuit “size,” and its one-sentence simulation does not resolve the gate-versus-wire convention.

⁴The use of Beame’s small-clique lower bound [Bea90] as the source statement for this section was suggested by the AI language model GPT-5.5; the authors independently verified its applicability and developed the resulting gate-count consequence (Fact 6.1) and Theorem 6.2.

⁵In this model, many processors read from and write to a shared pool of memory cells concurrently; under the PRIORITY convention, when several processors write to the same cell in one step, the one of least index succeeds.

Fact 6.1 (Gate-count consequence of Beame’s small-clique lower bound). *For every fixed depth d , there are a constant $a_d > 0$ and integers $r_d, N_d \geq 1$ such that*

$$\text{size}_d(K_r\text{-CLIQUE}_N) \geq N^{a_d r} \quad (51)$$

where $K_r\text{-CLIQUE}_N$ is the total monotone function on $\binom{N}{2}$ edge variables that accepts exactly when the represented graph contains an r -clique; see Definition B.7. The bound holds whenever $N \geq N_d$ and $r_d \leq r \leq \log N$. In particular, the upper cap $r \leq \log N$ is depth-independent; only the exponent rate and the lower onset may depend on d . The proof permits $a_d := 1/(2000(d+2)^3)$ and $r_d \geq 3000(d+2)^3$, after enlarging r_d to absorb Beame’s absolute threshold for r .

Proof. The plan is to simulate a circuit for $K_r\text{-CLIQUE}_N$ by a PRIORITY CRCW PRAM whose memory cells are the circuit’s input variables and gates, so that Beame’s memory-cell bound becomes a gate-count bound. Let C be a depth- d circuit for $K_r\text{-CLIQUE}_N$ with G non-input gates. Propagate constant gates and flatten every edge joining gates of the same type, neither of which increases the depth or G . Give each of the $\binom{N}{2}$ input variables and each non-input gate one memory cell, together with one junk cell, and use one processor per circuit wire. The unrestricted processor count is legitimate because Beame’s theorem explicitly permits infinitely many processors. Every non-input cell starts at 0, as required by his model. An OR-gate cell stores the gate value, whereas an AND-gate cell stores its complement. Because adjacent gate types alternate, this convention makes every gate-to-gate wire act uniformly: when evaluating a layer, its processor is triggered writes 1 to its gate cell exactly when it reads a 0 in the predecessor cell, which is exactly when its predecessor makes the stored disjunction true. Indeed, an AND-gate cell holds 0 exactly when that gate is true, which makes a target disjunction true, and an OR-gate cell holds 0 exactly when that gate is false, which makes a target conjunction false and hence its stored complement true. Since the model is nonuniform, an input-wire processor can handle positive and negated literals by its local transition rule, so no cells are needed for negations. Beame’s model has every processor write somewhere at each step, so a nontriggered or idle processor writes 0 to the junk cell; writing 0 to the gate cell could overwrite a 1 written simultaneously by a triggered processor and would destroy the PRIORITY implementation of OR. Thus every processor writing to a gate cell writes the same value 1, so PRIORITY resolves the concurrent write correctly; if no processor writes there, the initial 0 is the correct stored value. Beame’s designated output cell C_1 is itself an input cell, so it starts at the value of the corresponding edge variable rather than at the 0 that the accumulation above requires; it therefore cannot serve as an accumulator, and the output gate must be evaluated in a cell of its own. Evaluating one circuit layer per PRAM step and then using one final step to copy the output into C_1 , with a polarity flip when the output gate is an AND gate, takes at most $T \leq d+2$ steps. Hence C yields a PRIORITY CRCW PRAM with $c(N) := G + \binom{N}{2} + O(1)$ memory cells. Moreover, $c(N) \geq N$ for $N \geq 3$, as used in Beame’s memory-cell argument.

Set $a := \log_N c(N)$. Beame states the exponent $r/(43T^3)$, but we use the weaker $r/(1000T^3)$ because it is what the three terminal cases in the proof of [Bea90, Thm. 3.1(b)] establish uniformly for all $T \geq 1$. If $a \geq r/6$, then $a \geq r/(1000T^3)$ because $T \geq 1$. If $T \geq (3/10)(r/a)^{1/3} - 2$, then

$$a \geq \frac{27r}{1000(T+2)^3} \geq \frac{r}{1000T^3};$$

where the second inequality uses $T \geq 1$. This branch reaches the published constant only once $T > 39$ (that is, $T \geq 40$ for integer T), whereas our $T = d+O(1)$ may be a small fixed integer. In the remaining terminal case, Beame’s displayed inequality $343T^3a/(8r) \geq 343/344$ gives $a \geq r/(43T^3)$. Consequently, the common conclusion of all three cases is

$$c(N) \geq N^{r/(1000T^3)} \geq N^{b_d r}, \quad b_d := \frac{1}{1000(d+2)^3},$$

throughout the range $r \leq \log N$, once the absolute clique-order threshold and the sufficiently-large- N condition are met.

Let $r_\star$ be the absolute clique-order threshold in Beame's theorem, set $a_d := b_d/2$, and choose

$$r_d := \max \left\{ r_\star, \left\lceil \frac{3}{b_d} \right\rceil \right\} = \Theta(d^3).$$

After enlarging N_d if necessary, for every $r \geq r_d$ the preceding memory bound and the definition of $c(N)$ give

$$G \geq N^{b_d r} - \binom{N}{2} - O(1) \geq \frac{1}{2} N^{b_d r} \geq N^{a_d r},$$

which proves Equation (51) under the paper's non-input-gate size convention. $\square$

The explicit choices in Fact 6.1 permit $a_d = \Theta(d^{-3})$ and $r_d = \Theta(d^3)$. Hence the interval $r_d \leq r \leq \log N$ can be nonempty only when $d = O((\log N)^{1/3})$; a convenient sufficient asymptotic regime is $d = o((\log N)^{1/3})$, subject also to $N \geq N_d$. All results below keep d fixed.

The upper range $r \leq \log N$ is the feature needed below. Rossman's stronger exponent is proved for every fixed clique order r , whereas a superconstant- r extension is presented only as a possibility [Ros08, Thm. 1.2; Sec. 6, final paragraph before "Open Questions"]. Rossman's theorem therefore supplies no uniform onset threshold for the choice $r = r(n)$ used here.

We combine Fact 6.1 with the padded clique pattern $P_{k,r}^{\text{clq}}$ of Definition B.7. By Lemma B.8, ordinary clique projects to colored clique because completeness forces the color-respecting map to be injective, and colored clique projects to the padded pattern by setting every padding-edge variable to 1.

Theorem 6.2 (C: growing k at every fixed depth). *Fix a constant $\varepsilon > 0$. For every fixed depth d , the following hold along every sequence $k = k(n)$ satisfying $k \rightarrow \infty$ and $k \leq n^{1-\varepsilon}$, for all sufficiently large n . Define*

$$\rho_{\text{OV}} := \min\{\sqrt{k-1}, \log n\}, \quad \rho := \min\{\sqrt{k-1}, \log(n/k)\}. \quad (52)$$

Then

$$\text{size}_d(k\text{-OV}_{n,D,k}) \geq \exp(\Omega_d(\rho_{\text{OV}} \log n)) \quad \text{for every } D \geq C_{\text{OV}} k \log n, \quad (53)$$

$$\text{size}_d(k\text{-XOR}_{n,m,k}) \geq \exp(\Omega_d(\rho \log(n/k))) \quad \text{for odd } k \text{ and every } m \geq C_{\oplus} k \log(en/k), \quad (54)$$

$$\text{size}_d(k\text{-SUM}_{n,m,k}) \geq \exp(\Omega_d(\rho \log(n/k))) \quad \text{for every } m \geq C_{\Sigma} k \log n. \quad (55)$$

The implied constants in the three lower bounds depend only on d .

Proof. Fix d , and let a_d, r_d, N_d be supplied by Fact 6.1. We use the same argument for the three targets, with base parameter

$$B := \begin{cases} n, & Q = k\text{-OV}, \\ n/k, & Q \in \{k\text{-XOR}, k\text{-SUM}\}, \end{cases}$$

and host size $N = N_{\text{OV}}$ or $N = N_{\oplus, \Sigma}$ from Equation (16), respectively. Since $k \leq n^{1-\varepsilon}$, one has $B \rightarrow \infty$ in every case, and eventually $n \geq 16k$. The host estimates in Equations (18) and (19) therefore give

$$\log N \geq \frac{1}{4} \log B. \quad (56)$$

Let ρ_Q denote ρ_{OV} for $Q = k\text{-OV}$ and ρ otherwise, and set

$$r := \left\lfloor \frac{\rho_Q}{4} \right\rfloor. \quad (57)$$

Because $k \rightarrow \infty$ and $B \rightarrow \infty$, we have $\rho_Q \rightarrow \infty$, so $r \geq r_d$, $N \geq N_d$, and $r \geq \rho_Q/8$ for all sufficiently large n . The pattern-edge allowance follows from $r \leq \sqrt{k-1}$:

$$\binom{r}{2} \leq \frac{r^2}{2} \leq k-1.$$

Beame's upper-range condition also holds, since $\rho_Q \leq \log B$ and therefore $r \leq \rho_Q/4 \leq \frac{1}{4} \log B \leq \log N$ by Equation (56). Thus Lemma B.8 and Fact 6.1 imply

$$\text{size}_d(P_{k,r}^{\text{clq}}\text{-SUB}_N) \geq \text{size}_d(K_r\text{-CLIQUE}_N) \geq N^{a_d r}. \quad (58)$$

The pattern $P_{k,r}^{\text{clq}}$ has exactly $k-1$ edges, so the appropriate projection in Theorem 3.4 applies. Its maximum degree is $r-1$, but this causes no budget loss: by Equation (15), the consistency count is governed by $|\mathcal{I}(P)| \leq 2e(P) = 2(k-1)$ rather than by $\Delta(P)$. With the universal constants fixed in Theorem 3.4 and Lemma 3.5, Equations (20) and (21) give the displayed $k\text{-OV}$ and $k\text{-XOR}$ budgets, while Equation (22) gives the $k\text{-SUM}$ width $m \geq C_\Sigma k \log n$. This is why Equation (55) requires the larger width $C_\Sigma k \log n$ rather than the sharper fixed- k budget.

Composing the target projection with Equation (58) and using $r \geq \rho_Q/8$ and Equation (56), we obtain

$$\log \text{size}_d(Q) \geq a_d r \log N \geq \frac{a_d}{32} \rho_Q \log B.$$

This is exactly Equations (53) to (55), with the oddness condition inherited only by $k\text{-XOR}$. $\square$

The restriction $k \leq n^{1-\varepsilon}$ ensures that the list-target host size grows and that $\log(n/k)$ is a meaningful asymptotic base. The larger $k\text{-SUM}$ width requirement in Equation (55) is also genuine: at $N = \Theta(\sqrt{n/k})$, the exact requirement $m_\Sigma(P, N) = O(k \log(kN))$ is not uniformly $O(k \log(n/k))$ when k approaches n . The sharper estimate in Equation (23) applies whenever $n \geq k^3$, including the parameter ranges used in Theorems 5.1 and 6.4, Theorem 6.8, and Theorem 7.3.

Corollary 6.3 (The growing- k frontier and its plateau). *For every fixed depth d and constant $\varepsilon > 0$, the budgets of Theorem 6.2 imply, along every sequence $k = k(n) \rightarrow \infty$ with $k \leq n^{1-\varepsilon}$,*

$$\log \text{size}_d \geq \Omega_{d,\varepsilon}(\min\{\sqrt{k}, \log n\} \cdot \log n) \quad (59)$$

for all three targets, with k odd for $k\text{-XOR}$. In the subpolynomial regime $k = n^{o(1)}$, this sharpens to the common form

$$\text{size}_d \geq n^{\Omega_d(\min\{\sqrt{k}, \log n\})}. \quad (60)$$

Consequently the exponent in base n grows as $\sqrt{k}$ until $k = \Theta(\log^2 n)$ and then plateaus at order $\log n$; at $k = \Theta(\log^2 n)$ the lower bound is $\exp(\Omega_d(\log^2 n))$.

Proof. Since $k \rightarrow \infty$, one has $\sqrt{k-1} = \Theta(\sqrt{k})$. Under $k \leq n^{1-\varepsilon}$, one has $\log(n/k) \geq \varepsilon \log n$, so the list-target expressions in Theorem 6.2 are bounded below by the right-hand side of Equation (59) up to a constant depending on ε ; the $k\text{-OV}$ expression already has that form. If $k = n^{o(1)}$, then $\log(n/k) = (1 - o(1)) \log n$, so the dependence on ε disappears and exponentiating gives Equation (60). The two terms in the minimum balance when $\sqrt{k} = \Theta(\log n)$. $\square$

The two branches have a simple source-side interpretation. The edge allowance permits a clique on only $\Theta(\sqrt{k})$ vertices, because K_r uses $\binom{r}{2}$ of the available $k-1$ edges; in contrast, the expander source has $\Theta(k)$ vertices. The source exponent tracks the clique's vertices and therefore rises only as $\sqrt{k}$. Once r reaches $\Theta(\log N)$, Beame's upper range becomes the binding cap, and the lower bound remains at $\exp(\Omega_d(\log^2 n))$ even as k continues to grow.

6.2 Depth two: an exponent linear in k

At depth two, no switching argument is needed. The DNF orientation uses the standard minimal-positive-input principle: each minimal positive input forces a distinct accepting term. The source-specific CNF orientation constructs exponentially many minimal transversals of the hypergraph of colored copies, each of which must appear as a distinct clause. We use the padded expander family P_k^{exp} of Lemma 2.5, but only because it has no isolated vertices, has $\Theta(k)$ vertices, and contains a vertex of degree at least two; expansion and κ play no role.

For a map $\phi : V(P) \rightarrow [N]$, let

$$E_\phi := \{X_{f, \phi(a_f), \phi(b_f)} : f = (a_f, b_f) \in E(P)\}. \quad (61)$$

The assignment with positive support exactly E_ϕ is a minimal positive input of $P\text{-SUB}_N$, and

$$P\text{-SUB}_N = \bigvee_{\phi: V(P) \rightarrow [N]} \bigwedge_{x \in E_\phi} x. \quad (62)$$

Theorem 6.4 (D: depth two, both orientations). *There are universal constants $k_D \geq k_0$, $\gamma_D > 0$, and $\xi_D > 0$ such that the following hold.*

(i) *For every $k \geq k_D$ and every $N \geq k^2$,*

$$\text{size}_2(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma_D k}. \quad (63)$$

(ii) *For every $k \geq k_D$, the following target bounds hold whenever the indicated range condition is satisfied:*

(a) *if $k \leq \xi_D n^{1/4}$, then*

$$\text{size}_2(k\text{-OV}_{n,D,k}) \geq n^{\Omega(k)} \quad (64)$$

for every $D \geq C_{\text{OV}} k \log n$;

(b) *if $k \leq \xi_D n^{1/5}$, then*

$$\text{size}_2(k\text{-XOR}_{n,m,k}) \geq (n/k)^{\Omega(k)} \quad \text{for odd } k \text{ and every } m \geq C_{\oplus} k \log(en/k), \quad (65)$$

$$\text{size}_2(k\text{-SUM}_{n,m,k}) \geq (n/k)^{\Omega(k)} \quad \text{for every } m \geq C_{\Sigma} k \log(en/k). \quad (66)$$

All implied constants are universal.

Proof. Fix $P = P_k^{\text{exp}}$. We first lower-bound both possible top-gate orientations of a depth-two source circuit.

DNF orientation. Let $\mathcal{D}$ be a DNF, with arbitrary input negations, computing $P\text{-SUB}_N$. For every ϕ , the minimal positive input E_ϕ satisfies some term T_ϕ of $\mathcal{D}$. Every positive literal of T_ϕ lies in E_ϕ . If the positive part omitted an edge variable $x \in E_\phi$, then the assignment setting exactly the positive literals of T_ϕ to 1 and every other variable to 0 would still satisfy the term, but its positive support would be contained in $E_\phi \setminus \{x\}$.

The set $E_\phi \setminus \{x\}$ contains no variable from the pattern-edge group to which x belongs, so it cannot support any colored copy. This contradicts correctness of $\mathcal{D}$, and therefore the positive part of T_ϕ is exactly E_ϕ . Distinct maps give distinct colored edge sets: if $\phi' \neq \phi$, choose a vertex a on which they differ; because P has no isolated vertices, an edge incident to a contributes different input coordinates to E_ϕ and $E_{\phi'}$. Hence the map $\phi \mapsto T_\phi$ is injective, so every such DNF has at least

$$N^{v(P)} = N^{\Omega(k)} \quad (67)$$

terms, where the maximum-degree bound gives $v(P) \geq 2e(P)/4 = (k-1)/2$.

CNF orientation. By Equation (62), the function $P\text{-SUB}_N$ is monotone. Let $\mathcal{C}$ be a CNF computing it. First discard every tautological clause, and then remove negative literals without increasing the number of clauses. Write a clause as $C = P_C \vee \neg N_C$, where P_C is its positive part. Every 1-input satisfies P_C : otherwise, raising all variables in N_C to 1 would preserve acceptance by monotonicity while falsifying C . Every 0-input still falsifies some positive part, because any original clause falsified on that input already has all its positive literals equal to 0. Thus replacing each clause by its positive part gives an equivalent monotone CNF.

In any monotone CNF computing the function, every clause is a transversal of the hypergraph $\{E_\phi : \phi : V(P) \rightarrow [N]\}$: a clause disjoint from some E_ϕ would be falsified by the 1-input with positive support E_ϕ . Moreover, every minimal transversal must occur as a clause. To see this, let W be minimal and set the variables in W to 0 and all others to 1. Since W meets every colored copy, this is a 0-input, so some clause C is falsified and hence satisfies $C \subseteq W$. The clause C is itself a transversal, and minimality of W forces $C = W$.

Choose a pattern vertex a of degree $\delta \geq 2$, and let $E_P(a)$ be its incident pattern edges. For every function $h : [N] \rightarrow E_P(a)$, define

$$W_h := \{X_{f,u,w} : \text{there is } i \in [N] \text{ with } f = h(i) \text{ and } \text{lab}_a(f, u, w) = i\}. \quad (68)$$

Every colored copy meets W_h : if $\phi(a) = i$, then its variable for the edge $h(i)$ belongs to W_h . The transversal is minimal: for a variable $x \in W_h$ belonging to edge $f = h(i) = \{a, b\}$ and assigning the opposite endpoint the label j , choose a map ϕ with $\phi(a) = i$, $\phi(b) = j$, and arbitrary labels elsewhere. Its copy meets W_h exactly in x , so deleting x destroys the transversal property. Distinct functions h give distinct transversals, and therefore every CNF has at least

$$\delta^N \geq 2^N \quad (69)$$

clauses. The function $x/\log x$ is increasing for all sufficiently large x , so for $N \geq k^2$ one has $N/\log N \geq k^2/\log(k^2)$. After choosing a universal $\gamma_D > 0$ and then increasing k_D , this implies $N \geq \gamma_D k \log N$, or equivalently $2^N \geq N^{\gamma_D k}$.

The DNF bound in Equation (67) also has the form $N^{\gamma_D k}$ after decreasing γ_D if necessary. Flattening every child gate that has the same connective as the output, treating literals as width-one terms or clauses, and deleting unsatisfiable terms or tautological clauses normalizes every depth-at-most-two circuit to a DNF or a CNF without increasing its size. Each counted term or clause has fan-in at least two once k_D is large, and therefore contributes a distinct non-input gate. This proves Equation (63).

Transfer to the targets. Use the host substitutions in Equation (16). If $k \leq \xi_D n^{1/4}$ and ξ_D is sufficiently small, then Equation (18) gives $N_{\text{OV}} \geq k^2$ after accounting for the floor. If $k \leq \xi_D n^{1/5}$, then Equation (19) similarly gives $N_{\oplus, \Sigma} \geq k^2$. Thus the source bound applies in each case.

After increasing k_D and decreasing ξ_D if necessary, the range hypotheses imply $n \geq 16k$, and the list-target range also implies $n \geq k^3$. Hence Equations (20), (21) and (23) give exactly the displayed projection budgets. The same host estimates give $N_{\text{OV}} \geq n^{1/4}$ and $N_{\oplus, \Sigma} \geq (n/k)^{1/4}$, so Equation (63) transfers to Equations (64) to (66). Every reduction in Theorem 3.4 is depth zero; the conclusion is genuinely at depth exactly two. $\square$

Remark 6.5 (Why a degree-2 vertex is necessary). The CNF argument fails for a matching pattern, and the hypothesis cannot be removed. If P is a matching, then

$$P\text{-SUB}_N = \bigwedge_{f \in E(P)} \bigvee_{u, w \in [N]} X_{f, u, w}$$

is an explicit polynomial-size monotone CNF. The padded expander family contains a degree-4 core vertex, so the necessary degree condition is automatic in the application.

Corollary 6.6 (Subpolynomial form of Theorem D). *Along every sequence $k = k(n) \rightarrow \infty$ with $k = n^{o(1)}$, the range conditions of Theorem 6.4 eventually hold and all three target lower bounds simplify to*

$$\text{size}_2 \geq n^{\Omega(k)} = \exp(\Omega(k \log n)), \quad (70)$$

with k odd for k -XOR.

Proof. One has $n/k = n^{1-o(1)}$, so the list-target base $(n/k)^{\Omega(k)}$ equals $n^{\Omega(k)}$. The polynomial range conditions follow from $k = n^{o(1)}$. $\square$

At $k = \Theta(\log^2 n)$, Equation (70) becomes $\exp(\Omega(\log^3 n))$. This is stronger than the all-depth clique fallback because depth two admits exact counting of the required DNF terms and CNF clauses; the argument does not use, and therefore does not advance, the higher-depth source machinery.

6.3 Depth three: an exponent linear in k with a top disjunction⁶

The depth-two counting argument does not survive one more layer: an accepting middle CNF need not expose a single full-copy witness whose positive literals can be counted directly. At the Σ_3 orientation, however, the top disjunction selects one middle CNF on every accepting input. We show that a fixed CNF is exceedingly unlikely to become true for the first time exactly when all edges of a planted copy are present, and then union-bound over the middle CNFs. This unrestricted-CNF statement is complementary to the Choudhury–Sreenivasiah top-fan-in bound [CS26] for disjunctions of bounded-width CNFs computing the unpromised k -OV function when $k \leq D$: their bottom layer is structurally restricted, whereas ours has unrestricted clause width and literal signs. The precise comparison appears in Section 1.4.

⁶While we were attempting to extend the depth-two counting argument, the AI language model GPT-5.6 Sol suggested combining two known ingredients: the adaptive commitment process used below is a conceptual descendant of the decision-tree querying in the algorithmic proofs of the switching lemma due to Razborov and to Beame, and its minterm-on-a-random-background formulation follows the planted-subgraph method of Rossman and of Li, Razborov, and Rossman. We verified the suggestion independently and then developed the depth-three results here for our padded expander family.

The same padded expander family now carries a third source argument. Let F_k^{exp} , P_k^{exp} , and one bipartition side I_k be supplied by Lemma 2.5. The core is bipartite and 4-regular, the padded pattern has exactly $k - 1$ edges, and, writing $L_k := |I_k|$,

$$e(F_k^{\text{exp}}) > k - 5, \quad L_k > \frac{k - 1}{4} - 1. \quad (71)$$

No expansion or κ estimate is used below; the proof sees only a regular core containing all but $O(1)$ of the edge allowance (Corollary B.4: $t_k < 4$) and an independent bipartition side whose vertices have pairwise disjoint incident-edge sets. The choice of degree 4 balances the two probabilistic constraints: the uniqueness estimate requires $\theta > 2/\Delta$ for the background density $p = N^{-\theta}$ of Equation (77), while one side of a bipartite Δ -regular core with $k - O(1)$ edges has $(1/\Delta - o(1))k$ vertices. Since Theorem 6.7 contributes one factor $1/(pN) = N^{-(1-\theta)}$ per such vertex, the resulting limiting coefficient $(1/\Delta)(1 - 2/\Delta)$ is maximized over integers $\Delta \geq 3$ at $\Delta = 4$, which is the cap $\gamma < 1/8$ in Theorem 6.8.

We first isolate the probabilistic statement for one CNF. Fix an oriented simple loopless graph F , a nonempty independent set $I \subseteq V(F)$ whose vertices all have positive degree, and write $L := |I|$. Identify an input to $F\text{-SUB}_N$ with the set of variables assigned value 1. Let G be a random background in which every variable is present independently with probability p , and choose planted labels $\phi^*(v) \in [N]$ independently and uniformly for all $v \in V(F)$. The planted copy is

$$F^* := \{X_{f, \phi^*(a_f), \phi^*(b_f)} : f = (a_f, b_f) \in E(F)\}. \quad (72)$$

For $v \in I$ and $a \in [N]$, define the block

$$B_{v,a} := \{X_{f,a, \phi^*(u)} : f = (v, u) \in E(F)\} \cup \{X_{f, \phi^*(u), a} : f = (u, v) \in E(F)\}. \quad (73)$$

Because I is independent, every neighbor of every $v \in I$ lies outside I . Consequently, once the labels outside I are fixed, all blocks are determined. They are pairwise disjoint for distinct pairs (v, a) : independence of I prevents one pattern edge from belonging to two vertices of I , and changing a changes the label at the v -endpoint.

Let F_I^* be the planted edges not incident to I . Since I is independent, every planted edge belongs either to F_I^* or to exactly one planted block. For $A \subseteq I$, set

$$H_A := G \cup F_I^* \cup \bigcup_{v \in A} B_{v, \phi^*(v)}, \quad (74)$$

so that $H_I = G \cup F^*$. For a CNF T , let $\mathcal{M}_I(T)$ be the event

$$T(H_I) = 1 \quad \text{and} \quad T(H_A) = 0 \quad \text{for every } A \subsetneq I. \quad (75)$$

Let $\mathcal{M}(T)$ be the full minterm event that $T(G \cup F^*) = 1$ while $T(G \cup R) = 0$ for every proper subset $R \subsetneq F^*$. Thus $G \cup F^*$ is a minimal 1-input of T within the subcube obtained by deleting planted edges. Because every vertex of I has positive degree and the planted blocks are disjoint, $\mathcal{M}(T)$ implies $\mathcal{M}_I(T)$.

Theorem 6.7 (Single-CNF minterm bound). *In the setting of Equations (72) to (75), with $L = |I|$, background density p , and $\Pr$ taken over the background G and the planted labels ϕ^* , there is an absolute constant $C_0 > 0$ such that, for every CNF T with $s := s(T) \geq 1$ clauses,*

$$\Pr[\mathcal{M}_I(T)] \leq \left(\frac{C_0(L + \log(es))}{pN} \right)^L. \quad (76)$$

Consequently, the same bound holds for $\Pr[\mathcal{M}(T)]$. The theorem requires neither an implicant hypothesis nor any circuit-correctness assumption. If $s = 0$, then $T \equiv 1$ and both minterm events are empty.

The proof, given in Appendix G.1, fixes the labels outside I and, with the background still random, maximizes the number of remaining vertex-label pairs whose blocks contain a positive literal of a given clause, over every clause and every state $A \subseteq I$ for which the canonical representatives of those pairs are absent from the background. Only after taking the L th moment of this single support maximum Z does it expose the labels in I through a randomized commitment process. The process has success probability at most $Z^L/(L!N^L)$ for a fixed Z , but at least $1/L!$ on $\mathcal{M}_I(T)$; the factorials cancel, and the moment bound for Z gives Equation (76). This single maximum, rather than a union bound over ordered clause histories, is what preserves an exponent linear in k .

We now specialize to $F = F_k^{\text{exp}}$ and set

$$p := N^{-\theta}, \quad \frac{1}{2} < \theta < 1. \quad (77)$$

By Lemma G.2, applied with $\Delta = 4$, there is a constant $C_{4,\theta}$ such that, uniformly for $N \geq k^{C_{4,\theta}}$, the following good event has probability $1 - o(1)$ as $k \rightarrow \infty$: no planted edge already belongs to G , and F^* is the unique copy of F_k^{exp} in $G \cup F^*$. On this event, every $G \cup R$ with $R \subsetneq F^*$ is a 0-input of F_k^{exp} -SUB $_N$.

Theorem 6.8 (E: depth three, top disjunction, exponent linear in k). *For every constant γ with $0 < \gamma < 1/8$, there are $k_E(\gamma) \geq k_0$ and $C_E(\gamma) > 0$ such that*

$$\text{size}_{\Sigma_3}(P_k^{\text{exp}}\text{-SUB}_N) \geq \text{size}_{\Sigma_3}(F_k^{\text{exp}}\text{-SUB}_N) > N^{\gamma k} \quad (78)$$

for every $k \geq k_E(\gamma)$ and every $N \geq k^{C_E(\gamma)}$. For every constant c with $0 < c < \gamma/2$, there are $k_1(\gamma, c) \in \mathbb{N}$ and $C'(\gamma, c) > 1$ such that, for every $k \geq k_1(\gamma, c)$ and every $n \geq k^{C'(\gamma, c)}$,

$$\text{size}_{\Sigma_3}(k\text{-OV}_{n,D,k}) \geq n^{ck} \quad \text{for every } D \geq C_{\text{OV}}k \log n, \quad (79)$$

$$\text{size}_{\Sigma_3}(k\text{-XOR}_{n,m,k}) \geq (n/k)^{ck} \quad \text{for odd } k \text{ and every } m \geq C_{\oplus}k \log(en/k), \quad (80)$$

$$\text{size}_{\Sigma_3}(k\text{-SUM}_{n,m,k}) \geq (n/k)^{ck} \quad \text{for every } m \geq C_{\Sigma}k \log(en/k). \quad (81)$$

Every fixed $c < 1/16$ is therefore attainable. The theorem concerns Σ_3 only and implies no lower bound on size_{Π_3} or on size_3 .

Proof. Fix $0 < \gamma < 1/8$, and choose

$$\frac{1}{2} < \theta < 1 - 4\gamma. \quad (82)$$

The choice of θ implies $\gamma < (1 - \theta)/4$. By continuity, choose $\eta > 0$ such that

$$\eta < \min \left\{ \frac{1}{4}, 1 - \theta \right\} \quad \text{and} \quad \gamma < \left(\frac{1}{4} - \eta \right) (1 - \theta - \eta). \quad (83)$$

By Equation (71), after increasing the threshold in k we have $L := |I_k| \geq (1/4 - \eta)k$.

Suppose that a Σ_3 circuit $\mathcal{C} = \bigvee_{i=1}^t T_i$ of size $S \leq N^{\gamma k}$ computes F_k^{exp} -SUB $_N$. Put $M_{\text{src}} := e(F_k^{\text{exp}})N^2$ and $\hat{S} := S + 2M_{\text{src}}$. By Equation (4), one has $t \leq \hat{S}$ and $s(T_i) \leq \hat{S}$ for every i . On the good event, correctness gives $\mathcal{C}(G \cup F^*) = 1$ and $\mathcal{C}(G \cup R) = 0$ for every proper $R \subsetneq F^*$. Some

T_i therefore accepts $G \cup F^*$, and the same T_i rejects every proper deletion, so $\mathcal{M}(T_i)$ occurs. By Theorem 6.7 and a union bound,

$$\Pr[\text{the good event holds and some } \mathcal{M}(T_i) \text{ occurs}] \leq \widehat{S} \left(\frac{C_0(L + \log(e\widehat{S}))}{pN} \right)^L. \quad (84)$$

Choose $C_E(\gamma)$ large enough to imply the good-event threshold $N \geq k^{C_{4,\theta}}$ in Lemma G.2, and to satisfy $C_E(\gamma) > 2/\eta$, so that $N \geq k^{C_E(\gamma)}$ forces $k \leq N^{\eta/2}$. Since $e(F_k^{\text{exp}}) \leq k - 1$, after increasing $k_E(\gamma) \geq k_0$ one has $\widehat{S} \leq 2N^{\gamma k}$ whenever $k \geq k_E(\gamma)$ and $N \geq k^{C_E(\gamma)}$. In the same range, $L = O(k)$, $\log(e\widehat{S}) = O(k \log N)$, and, after another increase of $k_E(\gamma)$, $C_0(L + \log(e\widehat{S})) \leq N^\eta$. Since $pN = N^{1-\theta}$, the right-hand side of Equation (84) is at most

$$2N^{\gamma k - (1-\theta-\eta)L} \leq 2N^{[\gamma - (1-\theta-\eta)(1/4-\eta)]k} = o(1). \quad (85)$$

The exponent is negative by Equation (83). This contradicts the fact that the good event has probability $1 - o(1)$. Hence $\text{size}_{\Sigma_3}(F_k^{\text{exp}}\text{-SUB}_N) > N^{\gamma k}$. The projection in Equation (118) and orientation-preserving projection monotonicity give the first inequality in Equation (78).

Fix $0 < c < \gamma/2$. Choose $C'(\gamma, c) > \max\{2C_E(\gamma) + 3, 3\}$, and then increase $k_1(\gamma, c)$ as needed. Use $N = N_{\text{OV}}$ for $k\text{-OV}$ and $N = N_{\oplus, \Sigma}$ for $k\text{-XOR}$ and $k\text{-SUM}$, as in Equation (16). The condition $n \geq k^{C'(\gamma, c)}$ makes both host sizes at least $k^{C_E(\gamma)}$ for all sufficiently large k and ensures $n \geq \max\{16k, k^3\}$. Thus Equations (20), (21) and (23) make the three projections admissible with the universal constants fixed in Theorem 3.4 and Lemma 3.5.

By Equations (18) and (19), after increasing $k_1(\gamma, c)$ so that $n^{\gamma/2-c} \geq 2^\gamma$ and $(n/k)^{\gamma/2-c} \geq 2^\gamma$ throughout the range, one has

$$N_{\text{OV}}^{\gamma k} \geq n^{ck}, \quad N_{\oplus, \Sigma}^{\gamma k} \geq (n/k)^{ck}.$$

The source lower bound, Theorem 3.4, and orientation-preserving projection monotonicity yield Equations (79) to (81). Given any fixed $c < 1/16$, choose γ with $2c < \gamma < 1/8$. $\square$

Corollary E.5 combines Theorem 6.8 with the even- k lift; because the lift is a projection between $k\text{-XOR}$ instances, it preserves the circuit orientation. In the subpolynomial regime $k = n^{o(1)}$, the polynomial relation in Theorem 6.8 eventually holds and $n/k = n^{1-o(1)}$, so the three bounds have the common form $n^{\Omega(k)} = \exp(\Omega(k \log n))$, with odd k in the principal $k\text{-XOR}$ statement.

At the principal budgets in the polynomial-host range of Theorem 6.8, the elementary circuits in Propositions F.1 to F.3 have sizes $n^{k+o(k)}$, $(n/k)^{k+o(k)}$, and $(n/k)^{O(k)}$, respectively, with odd k in the direct $k\text{-XOR}$ comparison. Thus the theorem matches the elementary upper bounds at the linear-in- k exponent scale in its own orientation.

The top disjunction is used only to select one accepting middle CNF. A top conjunction has a dual selection statement on a maximal negative input: one middle DNF rejects that input and must accept every prescribed restoration. What is missing is a hard distribution on such maximal negative inputs together with a one-DNF anti-concentration theorem. Accordingly, the present argument gives no Π_3 or size_3 lower bound.

Table 6: The unconditional growing- k frontier, including the full theorem ranges and the common simplification for $k = n^{o(1)}$. The direct k -XOR statements require odd k .

Circuit class	Range in which the theorem holds	$k = n^{o(1)}$ regime bound	Structural feature consumed
Every fixed depth d	$k \rightarrow \infty$ and $k \leq n^{1-\varepsilon}$ for fixed $\varepsilon > 0$	$n^{\Omega_d(\min\{\sqrt{k}, \log n\})}$	growing clique; edge allowance and Beame cap
Depth 2, both orientations	$k \geq k_D$; $k \leq \xi_D n^{1/4}$ for k -OV and $k \leq \xi_D n^{1/5}$ for the list targets	$n^{\Omega(k)}$	exact DNF-term and CNF-clause counting
Depth 3, top disjunction	$0 < \gamma < 1/8$, $0 < c < \gamma/2$, and $n \geq k^{C'(\gamma, c)}$ for $k \geq k_1(\gamma, c)$	$n^{\Omega(k)}$	one support maximum over all adaptive states

6.4 The row-rich PARITY comparison

The imported PARITY bound. Håstad's fixed-depth theorem counts AND/OR gates and places negations at the input leaves, so its size measure agrees with our non-input-gate convention; his depth parameter counts the AND/OR layers and therefore equals our d . More explicitly, there is an absolute constant $\ell_0 > 1$ such that, for every $d \geq 2$ and every $\ell > \ell_0^d$, every depth- d circuit for PARITY_ℓ has at least

$$2^{(1/10)^{d/(d-1)} \ell^{1/(d-1)}} = \exp(\Omega_d(\ell^{1/(d-1)}))$$

AND/OR gates [Hås86]. Our circuits are already normalized to input literals, so no model conversion is needed. The range condition is depth-dependent and is absorbed only because every application fixes d before sending the number ℓ of free variables to infinity; no growing-depth claim is made.

There is a complementary route for k -XOR through this bound. Write $\text{PARITY}_\ell(x) := \bigoplus_{i=1}^\ell x_i$. The route is row-rich: it operates at $m = \Theta(n)$ rows rather than at the principal $O(k \log n)$ -scale budget, but it is useful for comparing the depth dependence of two qualitatively different anchors.

Theorem 6.9 (Amplified row-rich PARITY projection). *Let $k \geq 3$ be odd and let $n \geq 4k$. Put*

$$b := \left\lfloor \log \frac{n}{2k} \right\rfloor, \quad B := (k-1)b + 1, \quad \ell := (k-1)B. \quad (86)$$

There is an integer $m_0 \leq 2n$ and a projection

$$\text{PARITY}_\ell \leq_{\text{proj}} k\text{-XOR}_{n, m_0, k}. \quad (87)$$

Consequently, for every fixed depth $d \geq 2$ and every $m \geq m_0$, once ℓ exceeds the depth-dependent onset in Håstad's theorem,

$$\text{size}_d(k\text{-XOR}_{n, m, k}) \geq \exp\left(\Omega_d\left((k^2 \log(n/k))^{1/(d-1)}\right)\right). \quad (88)$$

In particular, the bound holds in the row-rich regime $m = \Theta(n)$ whenever the implicit constant is at least 2.

Proposition 6.10 (Upper bound for projection-based PARITY reductions). *Let $n \geq k \geq 1$, and suppose that a projection $R : \{0, 1\}^\ell \rightarrow (\mathbb{F}_2^m)^n$ satisfies*

$$k\text{-XOR}_{n, m, k}(R(x)) = \text{PARITY}_\ell(x) \quad \text{for every } x \in \{0, 1\}^\ell. \quad (89)$$

Then

$$\ell \leq k^2 \log(en/k) + k. \quad (90)$$

Table 7: Depth-by-depth comparison of the PARITY and subgraph routes for k -XOR, with odd $k = \Theta(\log^2 n)$ and $m = \Theta(n)$.

Circuit class	PARITY route	Subgraph route	Comparison
Depth 2	$\exp(\Omega_2(\log^5 n))$	$\exp(\Omega(\log^3 n))$ by Theorem 6.4	PARITY stronger
Σ_3	$\exp(\Omega_3(\log^{5/2} n))$	$\exp(\Omega(\log^3 n))$ by Theorem 6.8	PARITY weaker
Π_3	$\exp(\Omega_3(\log^{5/2} n))$	$\exp(\Omega_3(\log^2 n))$ unconditionally by Theorem 6.2	PARITY stronger unconditionally
Fixed $d \geq 4$	$\exp(\Omega_d(\log^{5/(d-1)} n))$	$\exp(\Omega_d(\log^2 n))$ by Theorem 6.2	PARITY weaker

The construction in Theorem 6.9 matches the general projection upper bound of Proposition 6.10 up to constant factors and is therefore order-optimal among projections from PARITY. Thus, since the bound depends on the construction only through ℓ , the difference between the PARITY route and the subgraph route cannot be removed merely by packing a larger parity instance into the same k -XOR target.

The proofs of Theorem 6.9 and Proposition 6.10 appear in Appendix E.3. At $k = \Theta(\log^2 n)$, one has $\log(n/k) = \Theta(\log n)$, so Theorem 6.9 gives exponent $\Omega_d(\log^{5/(d-1)} n)$ in the logarithm of the circuit-size lower bound. Table 7 compares this with the subgraph route for the same circuit class at $m = \Theta(n)$.

The comparison is within one row-rich regime: the subgraph-route lower bounds hold for every row count above their displayed budgets and therefore remain valid at $m = \Theta(n)$. The PARITY route is complementary rather than superseded. Under Conjecture 7.1, the subgraph route gives $\exp(\Omega_d(\log^3 n))$ also for Π_3 and every fixed $d \geq 4$, and hence dominates the PARITY route in every unresolved row. The order-optimal embedded dimension $\ell = \Theta(k^2 \log(n/k))$ from Proposition 6.10 also explains why the PARITY route cannot reproduce the polynomial host parameter $N = \Theta(\sqrt{n/k})$ exploited by the subgraph projection.

7 The Pattern-Uniform Conjecture and the Conditional Completion

The fixed-pattern lower bound of Theorem 2.2 is pointwise in the pattern: once P and the depth are fixed, its exponent loss vanishes as the host size tends to infinity, but both that loss and the onset threshold may depend arbitrarily on P . When the pattern is P_k^{exp} and grows with k , the quantifier order (92) does not provide the polynomial host threshold needed by the target reductions. The following conjecture isolates exactly the uniform source statement used in this section.

Conjecture 7.1 (Pattern-Uniform LRR, depth-by-depth form). *For the padded expander family $(P_k^{\text{exp}})_{k \geq k_0}$, for every fixed depth d there are constants $\gamma_d > 0$ and $c'_d > 0$ such that*

$$\text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma_d k} \quad (91)$$

for every $k \geq k_0$ and every $N \geq k^{c'_d}$.

Equivalently, the conjecture asks for two forms of uniformity at once: a source exponent linear in k and an onset host size polynomial in k , with constants allowed to depend on the fixed depth.

The fixed-pattern theorem of Theorem 2.2 has the quantifier order

$$\begin{aligned} & \forall \text{ simple loopless } P \text{ with } e(P) \geq 1 \text{ and no isolated vertices } \forall d \forall \eta > 0 \\ & \exists N_0(P, d, \eta) \forall N \geq N_0(P, d, \eta) : \\ & \text{size}_d(P\text{-SUB}_N) + 2e(P)N^2 + 2 \geq N^{\kappa(P)-\eta}. \end{aligned} \quad (92)$$

whereas Conjecture 7.1 asks for

$$\forall d \exists \gamma_d, c'_d > 0 \forall k \geq k_0 \forall N \geq k^{c'_d} : \quad \text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma_d k}. \quad (93)$$

When $\kappa(P) - \eta > 2$, Theorem 2.2 also gives the pure non-input-gate form $\text{size}_d(P\text{-SUB}_N) \geq N^{\kappa(P)-\eta}$; this is the form used in Corollary 2.6. The threshold in Equation (92) may depend arbitrarily on P_k^{exp} , and hence on k ; the conjecture does not reorder these quantifiers but requires that dependence to be polynomially bounded.

Tracking the switching estimate by [LRR17] reveals the same obstruction. Even if the restriction parameter $\delta > 0$ of Equation (107) is fixed along the whole family and the desired source rate $\gamma > 0$ is fixed, Lemma A.3 shows that the estimate becomes useful only when $\log N = \Omega_{d,\delta,\gamma}(k)$, rather than at the polynomial scale $N = k^{O_d(1)}$ required here. The published analysis also permits $d = d(N) = o(\log N / \log \log N)$ for each fixed pattern [LRR17, p. 938]; this is a growing-depth allowance on a different uniformity axis and does not control the growing family P_k^{exp} . Remark A.4 records the scope of this limitation.

The depth-two instance is already proved, up to the harmless choice of the universal starting threshold: Theorem 6.4(i) is Equation (91) with $\gamma_2 = \gamma_D$ and $c'_2 = 2$ after replacing k_0 by $\max\{k_0, k_D\}$.

At depth three, Theorem 6.8 proves the polynomial-threshold Σ_3 lower bound for the same family P_k^{exp} , and therefore establishes one orientation of the $d = 3$ frontier in Conjecture 7.1. It does not establish the conjecture itself: by Equation (3), a lower bound on size_3 must cover both orientations. Thus the Π_3 orientation at $d = 3$, the resulting $d = 3$ case of Conjecture 7.1, and all cases $d \geq 4$ remain open.

Remark 7.2 (Universal-rate strengthening). A stronger form of Conjecture 7.1 asserts that the constants γ_d may all be bounded below by one universal constant $\gamma_0 > 0$, independent of the depth. The depth-by-depth form is the minimal hypothesis needed below and yields depth-dependent constants in the target exponents. The universal-rate form would replace every Ω_d in Theorem 7.3 by a universal Ω , paralleling the universal fixed- k rate of Theorem 5.1. Both formulations are sufficient source-side statements; neither is claimed to be necessary for the target lower bounds.

Theorem 7.3 (F: conditional all-depth completion). *Assume Conjecture 7.1. For every fixed depth d there are integers $k_F(d) \geq k_0$ and constants $C_F(d) > 1$ such that, for every $k \geq k_F(d)$ and every $n \geq k^{C_F(d)}$,*

$$\text{size}_d(k\text{-OV}_{n,D,k}) \geq \exp(\Omega_d(k \log n)) \quad \text{for every } D \geq C_{\text{OV}} k \log n, \quad (94)$$

$$\text{size}_d(k\text{-XOR}_{n,m,k}) \geq \exp(\Omega_d(k \log n)) \quad \text{for odd } k \text{ and every } m \geq C_{\oplus} k \log(en/k), \quad (95)$$

$$\text{size}_d(k\text{-SUM}_{n,m,k}) \geq \exp(\Omega_d(k \log n)) \quad \text{for every } m \geq C_{\Sigma} k \log(en/k). \quad (96)$$

Consequently, the conclusion applies along every sequence $k = k(n) \rightarrow \infty$ with $k = n^{o(1)}$. No new target reduction is used: the projections are exactly those of Theorem 3.4.

Proof. Fix the depth d , and let $\gamma_d, c'_d > 0$ be supplied by Conjecture 7.1. Choose $C_F(d)$ sufficiently large that

$$C_F(d) \geq \max\{2c'_d + 4, 4\}. \quad (97)$$

After increasing $k_F(d)$, every pair $k \geq k_F(d)$ and $n \geq k^{C_F(d)}$ satisfies $n \geq \max\{16k, k^3\}$, and both host sizes in Equation (16) are at least $k^{c'_d}$. Indeed, Equations (18) and (19) give

$$N_{\text{OV}} \geq \frac{1}{2}k^{C_F(d)/2} \geq k^{c'_d}, \quad N_{\oplus, \Sigma} \geq \frac{1}{2}k^{(C_F(d)-1)/2} \geq k^{c'_d} \quad (98)$$

for all sufficiently large k .

For k -OV, the projection in Theorem 3.4, projection monotonicity, and Conjecture 7.1 give

$$\text{size}_d(k\text{-OV}_{n,D,k}) \geq N_{\text{OV}}^{\gamma_d k} \quad (99)$$

whenever $D \geq C_{\text{OV}}k \log n$, where admissibility follows from Equation (20). Since Equation (18) gives $\log N_{\text{OV}} = \Omega(\log n)$, this proves Equation (94).

For k -XOR and k -SUM, Conjecture 7.1, Theorem 3.4, and Lemma 3.3 yield

$$\text{size}_d(k\text{-XOR}_{n,m,k}) \geq N_{\oplus, \Sigma}^{\gamma_d k} \quad \text{for odd } k \text{ and every } m \geq m_{\oplus}(P_k^{\text{exp}}, N_{\oplus, \Sigma}), \quad (100)$$

$$\text{size}_d(k\text{-SUM}_{n,m,k}) \geq N_{\oplus, \Sigma}^{\gamma_d k} \quad \text{for every } m \geq m_{\Sigma}(P_k^{\text{exp}}, N_{\oplus, \Sigma}). \quad (101)$$

Because $C_F(d) \geq 4$, one has $\log(n/k) \geq (3/4)\log n$. Together with Equation (19), this gives $\log N_{\oplus, \Sigma} = \Omega(\log n)$, so both lower bounds are $\exp(\Omega_d(k \log n))$. Finally, Equations (21) and (23) make the displayed hypotheses $m \geq C_{\oplus}k \log(en/k)$ and $m \geq C_{\Sigma}k \log(en/k)$ admissible, completing the proof. $\square$

Corollary E.6 combines Theorem 7.3 with the even- k lifting projection. For every fixed depth d and every supplied decomposition $K = tr + s$ with $t \geq 2$, odd $r \geq k_F^{\oplus}(d)$, $0 \leq s < 2t$, and either $s = 0$ or $t \nmid s$, it gives an exponent $\Omega_d(r \log n')$ under its own polynomial threshold in n' . In particular, any family of decompositions with $r = \Theta(K)$ preserves the linear-in- K exponent scale.

8 Synthesis and Open Problems

This section distills three cross-target lessons from the preceding proofs, i.e., how the targets enforce selection rigidity, when lower and upper bounds match at the exponent scale, and how verification cost depends on depth, and then separates the remaining source-side and target-side frontiers.

8.1 Cross-target synthesis

The three targets share the same source and the same present-and-consistent edge-selection interface, but they differ in how selection rigidity is represented. Over $\mathbb{F}_2$, selection counts are visible only modulo two. The selector equations in Equation (27), together with the support-size identity in Equation (28), yield the odd- k rigidity formalized in Lemma C.5; when k is even, the branch in which the anchor is unselected and every group contributes an even number of candidates survives, and it admits supports that satisfy every row without encoding a copy of P . This is an algebraic obstruction specific to parity-based selection, not a property of the source pattern.

The k -SUM construction instead uses signed integer coordinates that count the selected candidates in each group exactly; Proposition C.8 shows that this mechanism cannot arise from a

group-homomorphic encoding of the $\mathbb{F}_2$ primitive into $\mathbb{Z}_{2^m}$. The k -OV target builds one-vector-per-class selection into the witness semantics, as summarized in Table 4. These two selection mechanisms therefore work for both parities without an exponent loss.

For an even target parameter K supplied with a decomposition $K = tr + s$ satisfying the hypotheses of Lemma E.2, the black-box lift recovers the fixed- K , top-disjunction depth-three, and conditional lower bounds of Corollary E.4, Corollary E.5, and Corollary E.6. The inherited exponent is governed by the odd source parameter r . Being a depth-zero projection between k -XOR instances, the lift composes with the other odd- k k -XOR lower bounds in the same way.

At the principal fixed- k budgets, specializing the lower bounds to the top-disjunction depth-three orientation and comparing them with the elementary circuits of Propositions F.1 to F.3 shows that the upper and lower exponents have the same asymptotic scale for all three targets; the k -XOR comparison uses the direct odd- k statement. This makes no tightness claim at depth two or for top-conjunction depth three, and it does not identify the leading constant: Proposition F.4 and Remark F.5 show that the specific degree-4 regular-core route remains separated from the brute-force leading exponent constant by a factor greater than 4.

A separate comparison concerns the cost of checking a fixed choice of candidate indices. For k -OV and k -XOR, the fixed-support predicates have depth-two verifiers of size $O(D)$ and $m2^{O(k)}$, respectively. For k -SUM, the direct depth-two verifier costs $2^{\Theta(km)}$ at the principal width, whereas the top-disjunction depth-three block-carry construction reduces the additional exponent to the scale in Equation (203); Equation (204) records the remaining very-high-width overhead. The fixed-support restrictions in Equations (160) and (162) impose a size floor $\exp(\Omega_d(k^{1/(d-1)}))$ on both list targets for every fixed depth $d \geq 2$ as $k \rightarrow \infty$. The separation is therefore quantitative and depth-sensitive rather than the mere presence of a parity obstruction.

8.2 Open problems

- (1) **The unresolved depth frontier.** Conjecture 7.1 is proved at $d = 2$ by Theorem 6.4(i). At $d = 3$, Theorem 6.8 establishes the Σ_3 orientation for the conjecture’s own padded expander family, leaving the Π_3 orientation and hence the size₃ statement open; for every fixed $d \geq 4$, the full conjecture remains open. The target-side goal is weaker: without necessarily proving the source conjecture, obtain the missing lower bounds for top-conjunction depth three and every fixed depth $d \geq 4$. Lemma A.3 and Remark A.4 explain why the published switching estimate in [LRR17] does not reach the required polynomial host size.
- (2) **Even- k XOR without an exponent loss.** The projection in Lemma E.2 transfers an odd-parameter lower bound through any supplied admissible decomposition $K = tr + s$, but the resulting exponent is governed by the odd source parameter r rather than directly by K . Can one design a different linear lift that, for every even K , starts from an odd parameter $r = \Omega(K)$, or a non-parity selection mechanism that enforces exact cardinalities without this potential loss?
- (3) **The k -SUM depth-and-width frontier.** At fixed k , the block-carry circuit of Proposition F.3 closes the same-orientation exponent-scale gap at the principal width for top-disjunction depth three, but the optimal bounds at depth two and top-conjunction depth three remain open. Even for the top-disjunction depth-three upper bound, Equations (203) and (204) leave two width regimes unresolved: support enumeration can dominate when the lower bound is capped by the bit width m , and the square-root carry overhead can grow at very large width. Determine which of these depth and width overheads are necessary, or strengthen the lower bound by exploiting the coupled structure of modular addition.

Because the projection interface is fixed independently of the source, progress on the source side is inherited by all three targets through the same three constructions: a new colored-subgraph family, a better exponent rate, or an extension to further depths improves the target bounds without any change to Theorem 3.4. Planted and average-case variants, positioned in Section 1.4, require distributional lower bounds that do not follow from the present worst-case statements and lie outside the scope of this paper. A systematic theory of which fine-grained AC^0 reductions preserve the correct lower-bound exponent is likewise deferred to separate work.

Acknowledgments

I am deeply grateful to my advisor, Professor Prashant Nalini Vasudevan, who encourage me to study average-case planted k -XOR lower bounds in AC^0 and with whom I had many formative discussions throughout 2023. His guidance and support have been the foundation of this work. I also thank Professor Alexander Golovnev for a discussion in the summer of 2024 on possible routes to AC^0 lower bounds for the average-case planted problem, and for suggesting candidate source problems such as the Coin Problem; that exchange prompted the broader search that led to colored subgraph isomorphism.

AI Disclosure

Six specific technical contributions were aided by AI language models. We verified each of them independently before inclusion. GPT-5.5 pointed us to Beame’s small-clique lower bound for CRCW PRAMs [Bea90] as a candidate unconditional source statement for the growing- k , every-fixed-depth regime; we verified its applicability and developed the resulting argument (Fact 6.1, Theorem 6.2) in Section 6.1. While we were attempting to extend our depth-two counting argument, GPT-5.6 Sol suggested combining two known ingredients for the depth-three result in Section 6.3: an adaptive commitment process, which is a conceptual descendant of the decision-tree querying in the algorithmic proofs of the switching lemma due to Razborov and to Beame [Raz95, Bea94], and the formulation of the target as a minterm event on a random background, which follows the planted-subgraph method of [Ros08, LRR17]. We verified the suggestion independently and then developed the depth-three results for our padded expander family. GPT-5.5 suggested and derived the Erdős–Ginzburg–Ziv degeneracy argument of Proposition D.5 (Appendix D.4). While we were attempting to derive a direct projection for even k , GPT-5.5 suggested the odd-to-even lifting approach recorded as Lemma E.2 (Appendix E.1); we verified the construction and its admissibility conditions before developing the consequences. GPT-5.5 suggested and derived the constant-factor floor argument of Proposition F.4 (Appendix F.2). GPT-5.5 also suggested and derived the depth-sensitive exponent comparison at the principal bit-width for k -SUM in Appendix F.3, contrasting the depth-three block-carry construction against the direct depth-two verifier. In every case, the authors take full responsibility for the correctness of the final statements and proofs.

References

- [ABDN18] Amir Abboud, Karl Bringmann, Holger Dell, and Jesper Nederlof. More consequences of falsifying SETH and the orthogonal vectors conjecture. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*, pages 253–266. Association for Computing Machinery, 2018.

- [Ajt83] Miklós Ajtai. Σ_1^1 -formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983.
- [Alo21] Noga Alon. Explicit expanders of every degree and size. *Combinatorica*, 41(4):447–463, 2021.
- [Ama10] Kazuyuki Amano. k -subgraph isomorphism on AC^0 circuits. *Computational Complexity*, 19(2):183–210, 2010. Preliminary version in CCC’09.
- [ASS⁺24] Shweta Agrawal, Sagnik Saha, Nikolaj I. Schwartzbach, Akhil Vanukuri, and Prashant Nalini Vasudevan. k -SUM in the sparse regime: Complexity and applications. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology—CRYPTO 2024, Part II*, volume 14921 of *Lecture Notes in Computer Science*, pages 315–351. Springer, 2024.
- [Bea90] Paul Beame. Lower bounds for recognizing small cliques on CRCW PRAM’s. *Discrete Applied Mathematics*, 29(1):3–20, 1990.
- [Bea94] Paul Beame. A switching lemma primer. Technical Report UW-CSE-95-07-01, Department of Computer Science and Engineering, University of Washington, November 1994.
- [CLSS26] Tameem Choudhury, Nutan Limaye, Karteeek Sreenivasaiah, and Srikanth Srinivasan. New and improved concrete lower bounds for orthogonal vectors. Technical Report TR26-121, Electronic Colloquium on Computational Complexity, July 2026. Also available as arXiv:2607.23799.
- [CS26] Tameem Choudhury and Karteeek Sreenivasaiah. Depth-3 circuit lower bounds for k -OV. *ACM Transactions on Computation Theory*, 2026. Published online May 2026; preliminary version appeared in STACS 2024.
- [CW19] Lijie Chen and R. Ryan Williams. An equivalence class for orthogonal vectors. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*, pages 21–40. Society for Industrial and Applied Mathematics, 2019.
- [DKK24] Itai Dinur, Nathan Keller, and Ohad Klein. Fine-grained cryptanalysis: Tight conditional bounds for dense k -SUM and k -XOR. *Journal of the ACM*, 71(3):23:1–23:41, 2024.
- [DLSV25] Mina Dalirrooyfard, Andrea Lincoln, Barna Saha, and Virginia Vassilevska Williams. Average-case hardness of parity problems: Orthogonal vectors, k -SUM and more. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2025)*, pages 4613–4643. Society for Industrial and Applied Mathematics, 2025.
- [EGZ61] Paul Erdős, Abraham Ginzburg, and Abraham Ziv. Theorem in the additive number theory. *Bulletin of the Research Council of Israel, Section F: Mathematics and Physics*, 10F(1):41–43, 1961.
- [FSS84] Merrick Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17:13–27, 1984.

- [Hås86] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC 1986)*, pages 6–20, 1986.
- [IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.
- [Kar72] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103, New York, 1972. Plenum Press.
- [KW19] Daniel M. Kane and R. Ryan Williams. The orthogonal vectors conjecture for branching programs and formulas. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, volume 124 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 48:1–48:15. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019.
- [LRR17] Yuan Li, Alexander Razborov, and Benjamin Rossman. On the AC^0 complexity of subgraph isomorphism. *SIAM Journal on Computing*, 46(3):936–971, 2017.
- [MOP20] Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. Explicit near-Ramanujan graphs of every degree. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 510–523, New York, NY, USA, 2020. Association for Computing Machinery. Corrected full version: arXiv:1909.06988v3, 29 November 2022.
- [Raz95] Alexander A. Razborov. Bounded arithmetic and lower bounds in Boolean complexity. In Peter Clote and Jeffrey B. Remmel, editors, *Feasible Mathematics II*, volume 13 of *Progress in Computer Science and Applied Logic*, pages 344–386. Birkhäuser Boston, 1995.
- [Ros08] Benjamin Rossman. On the constant-depth complexity of k -clique. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, STOC 2008, pages 721–730, New York, NY, USA, 2008. Association for Computing Machinery.
- [Ros21] Gregory Rosenthal. Beating treewidth for average-case subgraph isomorphism. *Algorithmica*, 83(8):2521–2551, 2021.
- [RW25] R. Ryan Williams. The orthogonal vectors conjecture and nonuniform circuit lower bounds. *SIAM Journal on Computing*, pages FOCS24–1–FOCS24–28, 2025. Special Section FOCS 2024; published online 14 October 2025.
- [VW18] Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians (ICM 2018)*, volume 3, pages 3447–3487, 2018.
- [Wag02] David Wagner. A generalized birthday problem. In Moti Yung, editor, *Advances in Cryptology—CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 288–304. Springer, 2002.

- [Yao85] Andrew Chi-Chih Yao. Separating the polynomial-time hierarchy by oracles. In *Proceedings of the 26th Annual IEEE Symposium on Foundations of Computer Science (FOCS 1985)*, pages 1–10, 1985.

A Source Formalities

Throughout this appendix, N denotes the host-size parameter that Li, Razborov, and Rossman write as n ; our n is reserved for the target universe size and never denotes a source host parameter. This appendix verifies that the structured source of Definition 2.1 is exactly the colored function to which their lower bound applies, matches their circuit model and gate measure with ours, assembles the cited ingredients of Theorem 2.2, and finally records the switching-scale obstruction of Lemma A.3, which explains why the published estimate does not yield Conjecture 7.1 at the polynomial-in- k host size required by the target reductions. No reduction or target-specific construction is used here.

A.1 Coordinate equivalence with the LRR structured problem

Proposition A.1 (Coordinate-level equivalence). *For every simple loopless graph P with one fixed orientation $f = (a_f, b_f)$ of each edge and every $N \geq 1$, the function $P\text{-SUB}_N$ of Definition 2.1 coincides, after a bijection of Boolean input coordinates, with $\text{SUBGRAPH}_{\text{col},N}(P)$ of [LRR17, Sec. 5, p. 960].*

Proof. The structured problem introduced in the unnumbered paragraph of [LRR17, Sec. 5, p. 960] uses the host vertex set $V(P) \times [N]$ and first-coordinate coloring $\chi(a, u) = a$; this is the same coordinate space as the random colored graph of [LRR17, Def. 2.8(v)] at vertex exponent $\alpha \equiv 1$. Its Boolean input coordinates are the possible host edges

$$\{(a, u), (b, w)\} \quad \text{with } \{a, b\} \in E(P) \text{ and } u, w \in [N]. \quad (102)$$

Because P is simple and every undirected edge is oriented exactly once, the map

$$\{(a_f, u), (b_f, w)\} \mapsto X_{f,u,w} \quad (103)$$

is a bijection between the coordinates in Equation (102) and those in Equation (1).

A properly colored copy of P in the structured host selects one vertex from each color class $\{a\} \times [N]$, and therefore has the form $\{(a, \phi(a)) : a \in V(P)\}$ for a map $\phi : V(P) \rightarrow [N]$. Vertices of distinct colors are automatically distinct, even if their second coordinates agree. The selected vertices span every required pattern edge exactly when, for each $f = (a_f, b_f) \in E(P)$, the host edge $\{(a_f, \phi(a_f)), (b_f, \phi(b_f))\}$ is present. Under Equation (103), this condition is $X_{f,\phi(a_f),\phi(b_f)} = 1$, which is precisely the acceptance condition in Equation (2). Hence the two Boolean functions agree under the coordinate bijection. $\square$

A.2 Citation assembly for the fixed-pattern lower bound

Proof of Theorem 2.2. For a fixed d , let P be a fixed simple loopless pattern P with at least one edge and no isolated vertices. By Proposition A.1, it suffices to prove the lower bound for $\text{SUBGRAPH}_{\text{col},N}(P)$ on the structured host $V(P) \times [N]$. We use $\theta_{\text{col}}(P)$, $\kappa_{\alpha,\beta}(P)$, and $G_{\alpha,\beta}(N)$ only as the names of the objects defined in [LRR17, Defs. 2.8 and 2.12]; no property of them beyond the statements cited below is used.

By [LRR17, Cor. 4.2], there is a function $\beta : E(P) \rightarrow [0, 2]$ such that $(\mathbf{1}, \beta) \in \theta_{\text{col}}(P)$ and

$$\kappa_{\mathbf{1},\beta}(P) = \kappa_{\text{col}}(P) = \kappa(P), \quad (104)$$

where $\kappa_{\text{col}}(P)$ is from [LRR17, Def. 2.12(iii)]; the first equality is the conclusion of the cited corollary, and the second is the notation fixed in Section 2.2.

Consider the colored random graph $G_{\mathbf{1},\beta}(N)$ of [LRR17, Def. 2.8(v)]. The displays of [LRR17, Thms. 3.11 and 3.12] write $G_{\alpha,\beta}(P)$, but the definition and the surrounding proofs show that the argument is the host parameter; we therefore use the corrected notation $G_{\alpha,\beta}(N)$. Since the vertex exponent is identically one, its vertex set is

$$\{(a, u) : a \in V(P), 1 \leq u \leq \lfloor N^1 \rfloor\} = V(P) \times [N], \quad (105)$$

and its possible edges are exactly the coordinates of Equation (102), with the first-coordinate coloring. Thus every graph in the support of $G_{\mathbf{1},\beta}(N)$ is an input to the structured function of Proposition A.1. Any circuit computing $P\text{-SUB}_N$ on every input has zero error under this distribution and therefore satisfies the average-case correctness hypothesis of the Li–Razborov–Rossman theorem.

Li, Razborov, and Rossman call the relevant model type-I in [LRR17, Sec. 3.2, p. 948]: the standard definition of AC^0 . Their Theorem 3.11 is the type-II statement and is not used for this import. We also consider a conservative offset: if their depth convention counts the input layer or a leaf-negation layer separately, we apply their theorem at depth $d + 2$; because the statement holds at every fixed depth, this changes neither the conclusion nor any later quantifier.

Footnote 4 of [LRR17, p. 940] specifies that size is the number of gates, but the paper does not separately state whether variable nodes, constant nodes, or input NOT gates are included. Let $S := \text{size}_d(P\text{-SUB}_N)$, and recall that $M_{\text{src}} = e(P)N^2$. A normalized circuit with S non-input AND/OR gates can be presented in their model using at most

$$S + 2M_{\text{src}} + 2$$

gates under every reading of that omission: at most M_{src} variable nodes, at most one shared input-negation node per variable, plus at most two constant nodes. Thus the source’s gate lower bound implies a lower bound on this quantity, without assuming that its word “gate” already means our non-input gate. This also avoids relying on the proof-internal estimate $w = O(g^2)$ on p. 949 to infer which nodes are included in g .

To invoke their Thm. 3.12, suppose this additive quantity is smaller than $N^{\kappa(P)+1}$; otherwise the desired lower bound is immediate. Under this assumption the converted source-model circuit has this polynomial size in N , as required by the definition of type-I AC^0 used in [LRR17]. Fix any constant accuracy requirement in their average-case formulation. The circuit under consideration has zero error on the support of $G_{\mathbf{1},\beta}(N)$ and therefore satisfies that requirement. The type-I lower bound of [LRR17, Thm. 3.12] now gives

$$S + 2M_{\text{src}} + 2 \geq N^{\kappa_{\mathbf{1},\beta}(P) - o_{P,d}(1)} = N^{\kappa(P) - o_{P,d}(1)}, \quad (106)$$

where the equality uses Equation (104). The asymptotic is taken with P , $(\mathbf{1}, \beta)$, the adjusted fixed depth, and the accuracy requirement fixed; in particular it is not uniform in P or d , and β is determined by P through Equation (104). Choosing the $o_{P,d}(1)$ term to have magnitude at most η proves Equation (5). If $\kappa(P) - \eta > 2$, instead choose it to have magnitude at most $\eta/2$. Then $2e(P)N^2 + 2 = o(N^{\kappa(P) - \eta/2})$, and for all sufficiently large N ,

$$S \geq N^{\kappa(P) - \eta/2} - 2e(P)N^2 - 2 \geq N^{\kappa(P) - \eta},$$

which proves the pure non-input-gate consequence in Theorem 2.2. $\square$

Remark A.2 (Alternative citation route). One may instead use the equivalence between the structured and unstructured colored functions established in [LRR17, Sec. 5, p. 960], together with their worst-case corollary. The argument above is more direct for the present paper because it identifies the hard distribution with the exact fixed input space of Definition 2.1 before invoking the average-case theorem.

A.3 The switching-scale obstruction

The quantifier orders of the fixed-pattern theorem and the pattern-uniform conjecture are displayed side by side in Equations (92) and (93). In particular, the fixed-pattern threshold may depend arbitrarily on P_k^{exp} , and hence on k ; no bound of the form $N_0(P_k^{\text{exp}}, d, \eta) \leq k^{O_d(1)}$ follows from Theorem 2.2. The published switching estimate exhibits a separate obstruction at the level of direct parameter tracking.

The relevant bad-event bound in [LRR17, Appendix C] has the form

$$B_{d,\delta}(S, N) := S(5N^{-\delta/d}\delta \log N)^{\delta \log N}. \quad (107)$$

Here and below, logarithms follow the paper's base-two convention; replacing another fixed logarithm base by base two changes only fixed d - and δ -dependent constants in the scale estimates. The fixed-pattern analysis takes $S = N^{O(1)}$ and chooses the constant $\delta > 0$ after fixing the pattern. The next lemma grants the favorable assumption that one fixed positive δ is available along the whole family and shows that the same estimate nevertheless does not handle size $N^{\Theta(k)}$ at polynomial host size.

Lemma A.3 (Scale limitation of the published switching estimate). *Fix constants $d \geq 1$, $\gamma > 0$, and $\delta > 0$. Along every sequence with $k \rightarrow \infty$ and $N \rightarrow \infty$,*

$$\log B_{d,\delta}(N^{\gamma k}, N) = \gamma k \log N - \Theta_{d,\delta}((\log N)^2) + O_\delta(\log N \log \log N). \quad (108)$$

Consequently, for every fixed $C > 0$,

$$B_{d,\delta}(N^{\gamma k}, N) \longrightarrow \infty \quad \text{as } k \rightarrow \infty \text{ along } N = k^C. \quad (109)$$

Moreover, along every sequence with $k \rightarrow \infty$ and $N \rightarrow \infty$ for which $B_{d,\delta}(N^{\gamma k}, N) = o(1)$, one necessarily has

$$\log N = \Omega_{d,\delta,\gamma}(k). \quad (110)$$

Proof. Substituting $S = N^{\gamma k}$ into Equation (107) and taking logarithms gives

$$\begin{aligned} \log B_{d,\delta}(N^{\gamma k}, N) &= \gamma k \log N + \delta \log N \left(\log(5\delta) + \log \log N - \frac{\delta}{d} \log N \right) \\ &= \gamma k \log N - \frac{\delta^2}{d} (\log N)^2 + \delta \log N \log \log N + O_\delta(\log N), \end{aligned} \quad (111)$$

which proves Equation (108). If $N = k^C$, the first term in Equation (111) is $\Theta_{C,\gamma}(k \log k)$, whereas the negative quadratic term and the remaining positive terms have magnitude $O_{C,d,\delta}((\log k)^2)$ and $O_{C,\delta}(\log k \log \log k)$, respectively. The logarithm therefore tends to $+\infty$, proving Equation (109).

Now consider a sequence with $k \rightarrow \infty$, $N \rightarrow \infty$, and $B_{d,\delta}(N^{\gamma k}, N) = o(1)$. Writing $x := \log N$, Equation (111) and $\log B_{d,\delta}(N^{\gamma k}, N) < 0$ imply, after rearranging and dropping the nonpositive term $-\delta x \log x$, that for all sufficiently large terms of the sequence,

$$\gamma k x \leq \frac{\delta^2}{d} x^2 + O_\delta(x). \quad (112)$$

After dividing by $x > 0$, the right-hand side is $O_{d,\delta}(x)$ because $x \rightarrow \infty$. If $x = o(k)$, this is $o(k)$ and contradicts the left-hand side γk . Hence $x = \Omega_{d,\delta,\gamma}(k)$, proving Equation (110). $\square$

Remark A.4 (Scope of the switching-scale calculation). The lemma limits direct parameter tracking of the published estimate; it does not rule out a refinement of the Li–Razborov–Rossman method, a different restriction scheme, or another proof of Conjecture 7.1. At the host sizes used by the target reductions, the benchmark $k = \Theta(\log^2 n)$ gives $\log N = \Theta(\log n) = \Theta(\sqrt{k})$, whereas Equation (110) requires $\log N = \Omega_{d,\delta,\gamma}(k)$. Two additional pattern dependencies remain in the original argument: the auxiliary restriction parameter and the constants in the probabilistic estimates are chosen after the pattern is fixed and would also have to be controlled uniformly. By contrast, the hitting-set estimate $|\mathcal{H}| \leq 2^{|E(P)|}$ from the proof of [LRR17, Thm. 3.11] is not itself an obstruction at polynomial host size, because $2^{O(k)} = N^{O(k/\log N)} = N^{o(k)}$ whenever $\log N \rightarrow \infty$, and such a factor can be absorbed into an exponent $\gamma_d k$.

B Pattern Constructions and Constants

Recall from Section 2.1 that J_t denotes a matching of t pairwise disjoint edges on fresh vertices, with J_0 empty. Every padded pattern below is a disjoint union $F \dot{\cup} J_t$. All core and padding orientations follow the convention fixed there.

B.1 The bounded-degree expander family

The bounded-degree expander core supplies $\kappa = \Omega(k)$, while a disjoint matching realizes every sufficiently large edge count exactly.

Theorem B.1 (Explicit simple near-Ramanujan graphs of every degree and size). *For every fixed integer $\Delta \geq 3$ and every $\varepsilon > 0$, there is a finite threshold $n_A(\Delta, \varepsilon)$ such that, for every $v \geq n_A(\Delta, \varepsilon)$ with $v\Delta$ even, there is a simple Δ -regular graph on v vertices whose nontrivial adjacency eigenvalues, that is, all adjacency eigenvalues other than the Perron eigenvalue Δ , have absolute value at most $2\sqrt{\Delta-1} + \varepsilon$.*

Proof. Cited conclusion. Regularity, the exact vertex count, and the two-sided spectral bound are the conclusion of [Alo21, Thm. 1.3]. Simplicity is not part of Alon’s (n, d, λ) -graph convention, so it is verified separately below rather than attributed to that theorem.

Selected source graph. Alon’s exact-size proof invokes the near-Ramanujan, bicycle-free input of [Alo21, Thm. 3.3 and Sec. 3] for its intermediate graph H ; from that point onward, the transformation uses only the regularity, vertex-count, spectral, and local-cycle properties stated there. Alon’s restatement does not include simplicity, so we inspect the corrected full version underlying that input directly. Its lift-model branch constructs G_0 as a lift of the simple base graph $K_{\Delta+1}$ and then obtains $G_1, \dots, G_t$ by successive 2-lifts [MOP20, full version, Def. 2.8 and Sec. 5, Steps 1–2]. The resulting lift-model output has all the properties required of H and is therefore an admissible choice in Alon’s transformation. Here H is Alon’s notation; Mohanty–O’Donnell–Paredes use the same letter in their overview for the smaller starting graph, not for the final lift output. Their construction produces a dense sequence of sizes, while Alon’s transformation is the ingredient that reaches the prescribed exact size v .

Simplicity of the lift chain. Every lift of a simple base graph is simple. Indeed, a loop in the lift would project to a loop in the base. If two lifted edges had the same unordered pair of endpoints, simplicity of the base would force them to lie over the same base edge $\{u, v\}$. Orient that edge from u to v ; the two lifted edges share an endpoint (u, i) , and the permutation attached to $\{u, v\}$ sends i to a unique fiber index j , so both have other endpoint (v, j) and coincide. Thus G_0 is simple, and a 2-lift is a lift in the same sense, so induction makes every G_i , and hence H , simple.

Simplicity of Alon's exact-size modification. Set $r := \lceil 2/\varepsilon \rceil$. In the notation of [Alo21, Sec. 3], Alon deletes a set U chosen as in [Alo21, Lem. 3.1] and restores regularity by adding a perfect matching on the deficient set $W := N_H(U)$; Alon's H' is our $H - U$. For every $u \in U$, [Alo21, Lem. 3.1(2)] makes the $(r+1)$ -neighborhood of u cycle-free, so $N_H(u)$ is independent. For distinct $u, u' \in U$, [Alo21, Lem. 3.1(3)] gives $d_H(u, u') \geq 2r + 3$, so the two neighborhoods are disjoint; moreover, for $x \in N_H(u)$ and $x' \in N_H(u')$ one has $d_H(x, x') \geq 2r + 1 > 1$, and hence no edge joins them. The disjointness gives $|W| = \Delta|U|$ and makes every vertex of W lose exactly one incident edge when U is deleted, while vertices outside W lose none. Because H is Δ -regular and the final order $v = |V(H)| - |U|$ satisfies $v\Delta$ even, $|W| = |V(H)|\Delta - v\Delta$ is even; hence a perfect matching on W exists and restores regularity. Finally, W is independent in the simple graph $H - U$, so every perfect matching on W creates neither a loop nor an edge parallel to one already present. $\square$

Remark B.2 (The finite threshold). Alon's construction is deterministic polynomial-time for fixed Δ and ε , but we use only that $n_A(\Delta, \varepsilon)$ is finite; no theorem in the paper depends on a closed numerical value of this threshold. The later reductions are nonuniform: for each parameter choice, the pattern and all literal substitutions are fixed before the source input is read, so no strongly explicit construction is required. Accordingly, the universal threshold k_0 below is finite but intentionally unoptimized.

Fact B.3 (Discrete Cheeger inequality, adjacency form). *Let F be a Δ -regular graph. With $h(F)$ normalized as in Definition 2.3 and $\lambda_2(F)$ denoting the second-largest eigenvalue of the adjacency matrix, one has*

$$h(F) \geq \frac{\Delta - \lambda_2(F)}{2}. \quad (113)$$

Proof. This is the discrete Cheeger inequality in the form recorded in [Ros21, Fact 2.4]. Rosenthal's $h(F)$ uses the same denominator $\min\{|S|, |V(F) \setminus S|\}$ as the equivalent form of Definition 2.3, and his $\lambda_2(F)$ is the second-largest adjacency-matrix eigenvalue [Ros21, p. 2525]. $\square$

Corollary B.4 (The bipartite degree-4 expander family). *Set $\varepsilon := 1/10$ and $\lambda_0 := 2\sqrt{3} + \varepsilon$. There is a universal threshold $k_0 \geq 600$ such that, for every $k \geq k_0$, writing*

$$r_k := \left\lfloor \frac{k-1}{4} \right\rfloor, \quad t_k := k-1-4r_k \in \{0, 1, 2, 3\}, \quad (114)$$

there are a simple connected bipartite 4-regular graph F_k^{exp} with bipartition (I_k, I'_k) and a padded pattern

$$P_k^{\text{exp}} := F_k^{\text{exp}} \dot{\cup} J_{t_k} \quad (115)$$

such that

$$v(F_k^{\text{exp}}) = 2r_k, \quad e(F_k^{\text{exp}}) = 4r_k > k-5, \quad |I_k| = |I'_k| = r_k > \frac{k-1}{4} - 1, \quad (116)$$

and

$$\kappa(P_k^{\text{exp}}) \geq \frac{k-1}{147}. \quad (117)$$

Moreover, for every $N \geq 1$, setting every padding-edge variable to 1 gives a projection

$$F_k^{\text{exp}}\text{-SUB}_N \leq_{\text{proj}} P_k^{\text{exp}}\text{-SUB}_N. \quad (118)$$

Proof. Let $n_* := \max\{n_A(4, 1/10), 5\}$ and choose $k_0 \geq \max\{4n_* + 1, 600\}$. Fix $k \geq k_0$. By Theorem B.1, there is a simple 4-regular graph Q_k on exactly r_k vertices whose nontrivial adjacency eigenvalues have absolute value at most $\lambda_0 < 4$. Hence Q_k is connected, because otherwise the eigenvalue 4 would have multiplicity greater than one, and it is non-bipartite, because a bipartite 4-regular graph has the nontrivial eigenvalue -4 .

Let A_k be the adjacency matrix of Q_k , and let F_k^{exp} be its bipartite double cover. Thus the two sides are copies I_k and I'_k of $V(Q_k)$, and every edge $\{u, v\}$ of Q_k yields the two cross-edges joining the copy of u on either side to the copy of v on the other. Equivalently, F_k^{exp} is the 2-lift of Q_k in which every base edge carries the transposition. The lift-simplicity argument in the proof of Theorem B.1 therefore shows that the cover is simple, and 4-regularity is immediate from the construction. Its adjacency matrix is

$$\begin{pmatrix} 0 & A_k \\ A_k & 0 \end{pmatrix}. \quad (119)$$

Its spectrum is therefore $\{\pm\lambda : \lambda \in \text{spec}(A_k)\}$. Since Q_k is connected and non-bipartite, the eigenvalue 4 of the cover is simple, while every other eigenvalue except the minimum eigenvalue -4 is at most λ_0 in absolute value. In particular, F_k^{exp} is connected and $\lambda_2(F_k^{\text{exp}}) \leq \lambda_0$.

The vertex, edge, padding, and independent-set identities in Equations (114) and (116) are immediate. By Fact B.3,

$$h(F_k^{\text{exp}}) \geq \frac{4 - \lambda_0}{2}.$$

Both F_k^{exp} and P_k^{exp} have no isolated vertices, and F_k^{exp} is a minor of P_k^{exp} . Hence the minor-monotonicity and expansion assertions of Theorem 2.4 give

$$\kappa(P_k^{\text{exp}}) \geq \kappa(F_k^{\text{exp}}) \geq \frac{v(F_k^{\text{exp}})h(F_k^{\text{exp}})}{3 \cdot 4} \geq \frac{r_k(4 - \lambda_0)}{12} > \frac{r_k}{30}. \quad (120)$$

Here the last inequality uses $\sqrt{3} < 7/4$, and hence $4 - \lambda_0 > 2/5$. Since $r_k \geq (k - 1)/4 - 1$ and $k \geq 600$,

$$\frac{r_k}{30} \geq \frac{k - 1}{120} - \frac{1}{30} \geq \frac{k - 1}{147},$$

which proves Equation (117). Finally, preserve every core variable and set every padding-edge variable to 1. Any accepting core map extends over the fresh matching colors, and every accepting padded map restricts to an accepting core map, proving Equation (118). The orientation-extension convention from Section 2.1 makes the preserved coordinates unambiguous. In particular, Lemma 2.5 holds with $c_\kappa = 1/147$, and $c_\kappa(k_0 - 1) > 4$ follows from $k_0 \geq 600$. $\square$

Remark B.5 (κ versus treewidth). The source lower bound is stated in terms of κ , not treewidth. For the bounded-degree expander cores used here, however, Equation (7) gives $\kappa(F_k) = \Omega(v(F_k))$, while Theorem 2.4 gives $\kappa(F_k) \leq \text{tw}(F_k) + 1 \leq v(F_k)$, so $\kappa(F_k) = \Theta(\text{tw}(F_k)) = \Theta(v(F_k))$. The distinction therefore does not affect the exponent scale of the present construction.

Remark B.6 (Quantifier form of the fixed- k source bound). Combining Corollaries B.4 and 2.6, one may fix any universal γ with $0 < \gamma < c_\kappa = 1/147$, apply Corollary 2.6 with source rate γ , and read the source statement as

$$\exists \gamma > 0, k_0 \quad \forall d \quad \forall k \geq k_0 \quad \exists N_0(\gamma, d, k) \quad \forall N \geq N_0(\gamma, d, k) : \quad \text{size}_d(P_k^{\text{exp}}\text{-SUB}_N) \geq N^{\gamma(k-1)}. \quad (121)$$

The rate is universal in d and k , whereas, after γ is fixed, the host-size threshold is allowed to depend on the fixed pair (d, k) .

B.2 The clique family

The clique core trades pattern vertices for a source bound that remains usable when the clique order grows, while a matching fills the remaining edge allowance.

Definition B.7 (Ordinary clique and clique core plus matching). For integers $r \geq 2$ and $N \geq 1$, let $K_r\text{-CLIQUE}_N$ be the total monotone Boolean function on variables $y_{\{u,w\}}$, indexed by unordered pairs of distinct $u, w \in [N]$, that accepts exactly when the represented graph contains an r -clique.

For integers $k \geq 2$ and $r \geq 2$ with $\binom{r}{2} \leq k - 1$, define

$$P_{k,r}^{\text{clq}} := K_r \dot{\cup} J_{k-1-\binom{r}{2}}. \quad (122)$$

Then $e(P_{k,r}^{\text{clq}}) = k - 1$ and $\Delta(P_{k,r}^{\text{clq}}) = r - 1$.

Lemma B.8 (Clique hardness embeds into the padded pattern). *For every admissible k, r, N ,*

$$K_r\text{-CLIQUE}_N \leq_{\text{proj}} K_r\text{-SUB}_N \leq_{\text{proj}} P_{k,r}^{\text{clq}}\text{-SUB}_N. \quad (123)$$

Proof. Let the ordinary clique input have variables $y_{\{u,w\}}$ for distinct $u, w \in [N]$. For an oriented edge $f = (a_f, b_f)$ of the colored K_r pattern, project

$$X_{f,u,w} \mapsto \begin{cases} y_{\{u,w\}}, & u \neq w, \\ 0, & u = w. \end{cases} \quad (124)$$

If the projected colored instance accepts via $\phi : [r] \rightarrow [N]$, then every pair of distinct pattern vertices is joined by an edge of K_r , and Equation (124) forces their images to be distinct. Hence ϕ is injective and its image is an ordinary r -clique. Conversely, every ordinary r -clique gives an injective accepting map after assigning its vertices to the colors. This proves the first projection.

For the second projection, preserve every variable belonging to the K_r core and set every variable belonging to a padding edge to 1. Any accepting core map extends to the fresh matching colors by arbitrary host indices, and any accepting padded map restricts to an accepting core map. The orientation-extension convention recalled at the beginning of this appendix makes the preserved core coordinate unique, so the map is a projection. $\square$

C Full Projection Proofs

This appendix proves the three projection theorems of Section 4, together with their supporting lemmas, the two design facts motivating the k -OV dummy and identifier conventions, and the 2-torsion obstruction of Proposition C.8, which motivates the native integer encoding for k -SUM.. The proofs use only the common skeleton of Section 3.2; no property of the source pattern beyond simplicity, looplessness, the fixed orientation, and $e(P) = k - 1$ is required.

C.1 The k -OV projection

Call a choice of one vector from each of the k target classes a selected tuple. Relative to the roles in Section 3.2, this one-vector-per-class semantics supplies selection rigidity (R1), the guard coordinates certify presence (R2), Lemma C.4 establishes label consistency (R3), Lemma C.3 purges dummies (R4), and Lemma 4.1 enforces the pairwise-distinctness promise (R5).

We first record the two elementary failures that motivate the dummy and identifier designs.

Fact C.1 (Zero-vector dummies are unsound). *If a k -OV class contains the zero vector as a dummy, then every selected tuple using that dummy is orthogonal, independently of the remaining chosen vectors and independently of the source input.*

Proof. The selected zero vector contributes a zero factor on every coordinate, so the product of the selected entries is zero on every coordinate. $\square$

Fact C.2 (Unpinned shared identifiers need not be inert). *Appending classwise identifiers without pinning each coordinate to zero in some class can destroy a genuine orthogonal witness, even when the identifiers make every class pairwise distinct.*

Proof. Take any structural instance with an orthogonal witness. Append one shared coordinate that is 1 on the witness's selected vector in every class, and use further coordinates to assign pairwise distinct codes within each class. That witness then has product 1 on the shared coordinate and is no longer orthogonal. $\square$

Pinning every identifier coordinate to zero throughout one owner class makes that coordinate inert for every selected tuple.

We next prove the identifier lemma. Recall that $\ell_n = \lceil \log n \rceil$ and $s_{k,n} = \lceil k\ell_n/(k-1) \rceil$.

Proof of Lemma 4.1. Create $s_{k,n}$ identifier coordinates and assign each coordinate to one owner class. Balance these assignments as evenly as possible, so every class owns at most $\lceil s_{k,n}/k \rceil$ coordinates. The defining inequality $(k-1)s_{k,n} \geq k\ell_n$ is equivalent to $s_{k,n} - \ell_n \geq s_{k,n}/k$; since the left-hand side is an integer, it follows that

$$\left\lceil \frac{s_{k,n}}{k} \right\rceil \leq s_{k,n} - \ell_n. \quad (125)$$

Hence every class is unpinned on at least $s_{k,n} - \lceil s_{k,n}/k \rceil \geq \ell_n$ coordinates. Pin every vector of a coordinate's owner class to zero on that coordinate. For each class U_h , choose an injection from its n vectors into the bit strings on its unpinned coordinates; such an injection exists because there are at least ℓ_n unpinned coordinates and $2^{\ell_n} \geq n$. Assign those code bits on the unpinned coordinates and retain zero on the owned coordinates. The vectors within each class are then pairwise distinct. On every identifier coordinate, however, the owner class contributes zero to every selected tuple, so the product of the selected entries is identically zero. Finally,

$$s_{k,n} = \left\lceil \frac{k\ell_n}{k-1} \right\rceil \leq 2\ell_n \quad (126)$$

for $k \geq 2$. $\square$

For completeness, we spell out the structural entries. On the anchor guard q_α , set $v_\alpha(q_\alpha) = 0$ and set every other vector's entry to 1. On an edge guard q_f , set $v_{f,u,w}(q_f) = \bar{X}_{f,u,w}$ and set every other entry, including every dummy entry in U_f , to 1. For $(a, g) \in \mathcal{I}(P)$, with $f := \rho(a)$, and $t \in [\ell_N]$, define

$$v_{f,u,w}(C_{a,g,t}^{10}) := \text{bit}_t(\text{lab}_a(f, u, w)), \quad v_{g,u',w'}(C_{a,g,t}^{10}) := 1 - \text{bit}_t(\text{lab}_a(g, u', w')), \quad (127)$$

$$v_{f,u,w}(C_{a,g,t}^{01}) := 1 - \text{bit}_t(\text{lab}_a(f, u, w)), \quad v_{g,u',w'}(C_{a,g,t}^{01}) := \text{bit}_t(\text{lab}_a(g, u', w')), \quad (128)$$

and set every other entry on these coordinates to 1. Every dummy is therefore 1 on every structural coordinate; its identifier entries are supplied by Lemma 4.1.

Lemma C.3 (Dummy rejection for k -OV). *Every selected tuple containing a dummy vector is non-orthogonal.*

Proof. Suppose the selected vector in class $h \in \{\alpha\} \cup E(P)$ is a dummy. On the guard q_h , that dummy has entry 1, and every selected vector from every other class also has entry 1 by the guard definition. The product of the selected entries on q_h is therefore 1. $\square$

Lemma C.4 (The k -OV consistency test). *Suppose a selected tuple uses genuine candidates from the reference group $f = \rho(a)$ and a non-reference group g incident to a . For a fixed bit position t , the products on $C_{a,g,t}^{10}$ and $C_{a,g,t}^{01}$ both vanish if and only if the two proposed label bits agree.*

Proof. Let x and y be the two proposed bits. Every class other than U_f and U_g has entry 1 on both coordinates, so the products are $x(1-y)$ and $(1-x)y$. The first is one only on $(x, y) = (1, 0)$, and the second is one only on $(0, 1)$; both vanish exactly when $x = y$. $\square$

Proof of Theorem 4.2. We first take $D = D_0(P, N, n)$. The construction has one anchor class and one edge class for each of the $e(P) = k - 1$ pattern edges, hence exactly k classes. The condition $N^2 \leq n$ permits every edge class to contain all N^2 genuine candidates and enough dummies to reach cardinality n , while the anchor class contains its one genuine vector and $n - 1$ dummies. By Lemma 4.1, every class contains n pairwise distinct vectors. Moreover, $s_{k,n} \geq \ell_n$, so $D \geq D_0(P, N, n) \geq \ell_n$ and therefore $n \leq 2^{\ell_n} \leq 2^D$; the constructed instance lies in the promised target family of Definition 3.1.

Assume first that $P\text{-SUB}_N(X) = 1$, and let $\phi : V(P) \rightarrow [N]$ satisfy Equation (2). Select the genuine anchor and, for every oriented edge $f = (a_f, b_f)$, the genuine vector $v_{f, \phi(a_f), \phi(b_f)}$. The anchor guard has product zero because the anchor contributes zero. On the guard q_f , the selected group- f vector contributes $\bar{X}_{f, \phi(a_f), \phi(b_f)} = 0$, so that product also vanishes. Every pair of selected candidates incident to a vertex a proposes the common label $\phi(a)$, and Lemma C.4 therefore makes every consistency product zero. Finally, every identifier coordinate is zero in its owner class by Lemma 4.1. The selected tuple is an orthogonal witness.

Conversely, suppose the constructed instance has an orthogonal witness. By Lemma C.3, every selected vector is genuine, so the witness consists of v_α and one candidate v_{f, u_f, w_f} from every edge class. On q_f , all classes other than U_f contribute one, and orthogonality therefore forces $\bar{X}_{f, u_f, w_f} = 0$; hence every selected candidate represents a present edge. For every non-isolated vertex a , Lemma C.4 applied to every $(a, g) \in \mathcal{I}(P)$ and every bit position forces all incident candidates to propose the same label. Define $\phi(a)$ to be this common value; for an isolated vertex, choose any value in $[N]$. Then $X_{f, \phi(a_f), \phi(b_f)} = 1$ for every $f \in E(P)$, so $P\text{-SUB}_N(X) = 1$.

It remains to verify that the map is a projection. The only source-dependent entries are the edge-guard entries $\bar{X}_{f, u, w}$; every consistency bit, dummy bit, identifier bit, and inactive entry is a constant determined by (P, N, n) and the candidate index. Thus every output bit is a constant or a negated input literal. For $D > D_0(P, N, n)$, append $D - D_0(P, N, n)$ all-zero coordinates. They preserve pairwise distinctness and make the product of every selected tuple zero on the new coordinates, so correctness and the projection property remain unchanged. $\square$

C.2 The odd- k k -XOR projection

Here Lemma C.5 establishes selection rigidity (R1) and dummy purge (R4), Lemma C.6 certifies presence (R2), and Lemma C.7 establishes label consistency (R3). Because the target is an indexed list in which repeated vector values are permitted, no separate promise-validity role is needed.

We begin with the nonuniform guard compression. Write $r := r_{\text{guard}}(P, N)$.

Proof of Lemma 4.3. Choose every label $\lambda_{f,u,w} \in \mathbb{F}_2^r$ independently and uniformly. Fix one choice $(u_f, w_f) \in [N]^2$ for every $f \in E(P)$ and one nonempty set $T \subseteq E(P)$. Because the selected labels belong to distinct edge groups, they are independent uniform vectors, and their XOR is uniform on $\mathbb{F}_2^r$. Consequently,

$$\Pr \left[\bigoplus_{f \in T} \lambda_{f,u_f,w_f} = 0 \right] = 2^{-r}. \quad (129)$$

There are $N^{2e(P)}$ choices of the candidate tuple and fewer than $2^{e(P)}$ nonempty choices of T . Since $r = \lceil 2e(P) \log N + e(P) + 1 \rceil$, the union bound gives

$$\Pr[\text{some violation of (30)}] \leq N^{2e(P)} 2^{e(P)} 2^{-r} \leq N^{2e(P)} 2^{e(P)} 2^{-(2e(P) \log N + e(P) + 1)} = \frac{1}{2}. \quad (130)$$

Hence a deterministic labelling satisfying Equation (30) exists. Fix one such labelling for the parameter pair (P, N) ; it is independent of the source input X . $\square$

For the remaining proof, the matrix rows are exactly those described in Section 4.2: selector rows S_f , compressed guard rows H_h , consistency rows $C_{a,g,t}$, and one anti-dummy row A_{dum} . All unspecified entries are zero, and each dummy column is zero except for its entry one on A_{dum} .

Lemma C.5 (Support rigidity for odd k). *Assume $k = e(P) + 1$ is odd. If a k -support has zero sum on every selector row and on the anti-dummy row, then it contains the anchor, exactly one candidate from every edge group, and no dummies.*

Proof. Let $\zeta \in \{0, 1\}$ indicate whether the anchor is selected, let c_f count selected candidates from group f , and let d count selected dummies. The selector equations give $c_f \equiv \zeta \pmod{2}$ for every f , while the anti-dummy row gives $d \equiv 0 \pmod{2}$. If $\zeta = 1$, then every c_f is a positive odd integer, and the support-size identity

$$1 + \sum_{f \in E(P)} c_f + d = k = 1 + e(P)$$

forces $c_f = 1$ for every f and $d = 0$. If $\zeta = 0$, then every c_f and d is even, so the support size is even, contradicting that k is odd. $\square$

Lemma C.6 (Compressed-guard correctness). *Suppose a support contains exactly one candidate c_{f,u_f,w_f} from each edge group. Its vector of compressed-guard row sums is zero if and only if every selected edge is present.*

Proof. A present candidate has $\bar{X}_{f,u_f,w_f} = 0$ and contributes the zero label, while an absent candidate contributes λ_{f,u_f,w_f} . Hence the vector of guard sums is

$$\bigoplus_{f: X_{f,u_f,w_f}=0} \lambda_{f,u_f,w_f}. \quad (131)$$

If every selected edge is present, this XOR is empty and therefore zero. If at least one edge is absent, the indexing set in Equation (131) is a nonempty subset of $E(P)$, and the XOR is nonzero by Equation (30). $\square$

Lemma C.7 (k -XOR consistency). *Suppose a support contains exactly one candidate from each edge group. The consistency rows all sum to zero if and only if, at every non-isolated pattern vertex a , all incident selected candidates propose the same label under lab_a .*

Proof. Fix $(a, g) \in \mathcal{I}(P)$, write $f := \rho(a)$, and fix a bit position t . Only the selected candidates from groups f and g can contribute to $C_{a,g,t}$, and its sum is

$$\text{bit}_t(\text{lab}_a(f, u_f, w_f)) \oplus \text{bit}_t(\text{lab}_a(g, u_g, w_g)). \quad (132)$$

This sum is zero exactly when the two bits agree. Ranging over all $t \in [\ell_N]$ forces equality of the complete binary encodings, and ranging over all $(a, g) \in \mathcal{I}(P)$ ties every non-reference incident group to the reference group at a . $\square$

Proof of Theorem 4.4. First take $m = m_\oplus(P, N)$. If $P\text{-SUB}_N(X) = 1$ via an accepting map ϕ , select the anchor and the $e(P)$ candidates $c_{f, \phi(a_f), \phi(b_f)}$. Every selector row sees two ones, one from the anchor and one from its group, and therefore sums to zero. Every selected edge is present, so all compressed-guard contributions vanish. The selected candidates propose the common label $\phi(a)$ at every pattern vertex, so every consistency row sums to zero by Lemma C.7. No dummy is selected, so the anti-dummy row also sums to zero. These $k = e(P) + 1$ columns witness acceptance.

Conversely, let S be a zero-sum support of size k . By Lemma C.5, it contains the anchor, one candidate c_{f, u_f, w_f} from every group, and no dummies. By Lemma C.6, every selected edge is present. By Lemma C.7, all candidates incident to a non-isolated vertex a propose one common label; define this value to be $\phi(a)$, and assign an arbitrary value in $[N]$ to every isolated vertex. Then $X_{f, \phi(a_f), \phi(b_f)} = 1$ for every pattern edge, so $P\text{-SUB}_N(X) = 1$.

Every selector, consistency, anti-dummy, anchor, and dummy entry is constant. A compressed-guard entry is $\bar{X}_{f, u, w} \lambda_{f, u, w}(h)$, which is either the constant zero or the literal $\bar{X}_{f, u, w}$ because the label bit is fixed. Thus the map is a projection.

The column budget $1 + e(P)N^2 \leq n$ permits padding with the required number of dummy indices; repeated dummy columns are allowed by the indexed-list convention of Definition 3.1. For $m > m_\oplus(P, N)$, append all-zero rows. Finally, Equations (15) and (29) give

$$m_\oplus(P, N) = e(P) + r_{\text{guard}}(P, N) + |\mathcal{I}(P)|\ell_N + 1 = O(k \log(2N)) \quad (133)$$

for every $N \geq 1$. $\square$

C.3 The k -SUM projection

Here Lemma C.9 establishes selection rigidity (R1) and dummy purge (R4), while the guard and consistency coordinates certify presence (R2) and label consistency (R3) in the proof of Theorem 4.5. Lemma C.15 reduces the target modular equation to the exact integer-vector condition used by those arguments. Repeated list values are permitted, so no analogue of (R5) is needed.

We first justify why the construction is native to integer addition rather than imported from $\mathbb{F}_2$. The signed packing below is the large-base positional device introduced in Section 4.3; here we verify its range and modular faithfulness in full.

Proposition C.8 (2-torsion obstruction). *Let $r \geq 2$, $m \geq 1$, and $k \geq 2$. There is no additive map $\psi : \mathbb{F}_2^r \rightarrow \mathbb{Z}_{2^m}$ such that, for every k -tuple $(v_1, \dots, v_k)$,*

$$\bigoplus_{i=1}^k v_i = 0 \iff \sum_{i=1}^k \psi(v_i) = 0 \text{ in } \mathbb{Z}_{2^m}. \quad (134)$$

Proof. Additivity gives $2\psi(v) = \psi(v \oplus v) = \psi(0) = 0$ for every v . The only elements of $\mathbb{Z}_{2^m}$ annihilated by two are 0 and 2^{m-1} , so the image of ψ has $\mathbb{F}_2$ -dimension at most one. Since $r \geq 2$, the kernel contains a nonzero vector v . The k -tuple $(v, 0, \dots, 0)$ has nonzero XOR but image sum zero, contradicting Equation (134). $\square$

We now analyze the signed vector construction. For a k -support S , write

$$y_S := \sum_{j \in S} \text{vec}(j) \in \mathbb{Z}^{R(P)}. \quad (135)$$

Lemma C.9 (Exact support structure). *If $y_S = 0$, then S contains the anchor, exactly one candidate from every edge group, and no dummies.*

Proof. Let $\zeta \in \{0, 1\}$ indicate whether the anchor is selected, let c_f count selected candidates from group f , and let d count selected dummies. Vanishing of the group coordinates gives $c_f = \zeta$ for every f , and vanishing of the dummy coordinate gives $d = 0$. Since $e(P) = k - 1$, the support-size identity becomes

$$k = \zeta + \sum_{f \in E(P)} c_f + d = k\zeta.$$

Hence $\zeta = 1$, every $c_f = 1$, and no dummy is selected. $\square$

Lemma C.10 (Coordinate range). *For every k -support S , one has $\|y_S\|_\infty \leq kN$.*

Proof. Every coordinate of every role vector has absolute value at most N : group, guard, and dummy entries lie in $\{-1, 0, 1\}$, while consistency entries lie in $[-N, N]$. Summing k vectors gives the claimed bound. $\square$

Recall that $B = 2^q > 2kN$, $R = R(P)$, $m_\Sigma(P, N) = qR$, and $2^{m_\Sigma(P, N)} = B^R$.

Lemma C.11 (Faithful balanced-digit packing). *For every $y \in \mathbb{Z}^R$ with $\|y\|_\infty \leq kN$,*

$$\Phi(y) \equiv 0 \pmod{B^R} \iff y = 0. \quad (136)$$

Proof. The reverse implication is immediate. For the forward implication, proceed by induction on R . When $R = 1$, divisibility by B gives $B \mid y_0$, while $|y_0| \leq kN < B/2$, so $y_0 = 0$.

Assume $R > 1$. Reducing $\Phi(y) = y_0 + B \sum_{i=1}^{R-1} y_i B^{i-1}$ modulo B shows that $B \mid y_0$, and the same balanced-range argument gives $y_0 = 0$. Dividing by B then gives

$$B^{R-1} \mid \sum_{i=1}^{R-1} y_i B^{i-1}. \quad (137)$$

The remaining digits satisfy the same ℓ_∞ bound, so the induction hypothesis gives $y_1 = \dots = y_{R-1} = 0$. $\square$

Lemma C.12 (Single-column range). *For every anchor, candidate, or dummy index j ,*

$$|\Phi(\text{vec}(j))| < B^R. \quad (138)$$

Proof. Every coordinate of $\text{vec}(j)$ has absolute value at most N , and $B > 2kN \geq 2N$. Hence

$$|\Phi(\text{vec}(j))| \leq N \sum_{i=0}^{R-1} B^i < \frac{B}{2} \cdot \frac{B^R - 1}{B - 1} < B^R, \quad (139)$$

where the last inequality uses $B > 2$. $\square$

Remark C.13 (Why low-bit modulus extension fails). Let S be the support arising from an accepting source map, so $y_S = 0$. Put $t_j := \Phi(\text{vec}(j))$ and let w be the number of selected indices with $t_j < 0$. The anchor has negative exact packed value, so $1 \leq w \leq k$. By Lemma C.12, reduction modulo B^R wraps each negative t_j exactly once and each nonnegative t_j not at all; therefore

$$\sum_{j \in S} \hat{a}_j = \sum_{j \in S} t_j + wB^R = wB^R. \quad (140)$$

If these residues are merely viewed as low-bit integers modulo 2^m and $2^{m-m_\Sigma(P,N)} > k$, then Equation (140) is nonzero modulo 2^m , so the true witness is rejected. This is the reason for the top-bit scaling.

Lemma C.14 (Top-bit scaling). *For every $m \geq m_\Sigma(P, N)$ and every integer Z ,*

$$2^m \mid 2^{m-m_\Sigma(P,N)} Z \iff 2^{m_\Sigma(P,N)} \mid Z. \quad (141)$$

Consequently, for every support S ,

$$\sum_{j \in S} a_j \equiv 0 \pmod{2^m} \iff \sum_{j \in S} \hat{a}_j \equiv 0 \pmod{2^{m_\Sigma(P,N)}}. \quad (142)$$

Proof. The first equivalence follows by cancelling the common power of two $2^{m-m_\Sigma(P,N)}$ from the divisibility relation. Apply it to $Z = \sum_{j \in S} \hat{a}_j$ to obtain the second equivalence. $\square$

Lemma C.15 (Packed supports are exact). *For every k -support S ,*

$$\sum_{j \in S} a_j \equiv 0 \pmod{2^m} \iff y_S = 0. \quad (143)$$

Proof. By Lemma C.14, the modular equation on the left is equivalent to $\sum_{j \in S} \hat{a}_j \equiv 0 \pmod{B^R}$. Since $\hat{a}_j \equiv \Phi(\text{vec}(j)) \pmod{B^R}$ and Φ is $\mathbb{Z}$ -linear, this is equivalent to

$$\Phi\left(\sum_{j \in S} \text{vec}(j)\right) = \Phi(y_S) \equiv 0 \pmod{B^R}. \quad (144)$$

By Lemma C.10, $\|y_S\|_\infty \leq kN$, so Lemma C.11 makes Equation (144) equivalent to $y_S = 0$. $\square$

Proof of Theorem 4.5. Assume first that $P\text{-SUB}_N(X) = 1$ via an accepting map ϕ , and select the support

$$S_\phi := \{\alpha\} \cup \{c_{f, \phi(a_f), \phi(b_f)} : f \in E(P)\}. \quad (145)$$

Every group coordinate of y_{S_ϕ} is $1 - 1 = 0$. Every consistency coordinate compares two occurrences of the common value $\phi(a)$ and is zero. Every selected edge is present, so every guard contribution is zero, and no dummy is selected. Hence $y_{S_\phi} = 0$, and Lemma C.15 gives a zero-sum k -support modulo 2^m .

Conversely, suppose the constructed k -SUM instance has a zero-sum support S of size k . By Lemma C.15, $y_S = 0$, and Lemma C.9 forces the anchor, exactly one candidate c_{f, u_f, w_f} from each group, and no dummies. The guard coordinate is the nonnegative sum $\sum_f \bar{X}_{f, u_f, w_f}$; since it is zero, every term is zero and every selected edge is present. Each consistency coordinate is the signed difference of the labels proposed by a non-reference group and the reference group at the same

pattern vertex, so its vanishing makes those labels equal. Define $\phi(a)$ to be the common proposed label at each non-isolated vertex and choose an arbitrary value for every isolated vertex. Then $X_{f,\phi(a_f),\phi(b_f)} = 1$ for every f , and the source instance accepts.

The column budget permits all genuine indices and enough dummies to reach length n , with repeated dummy residues allowed by Definition 3.1. The anchor and dummy residues are constants.

For a candidate $c_{f,u,w}$, every coordinate of $\text{vec}(c_{f,u,w})$ is constant except the guard coordinate $\bar{X}_{f,u,w}$, so the packed-and-scaled integer $a_{f,u,w}(X)$ takes one of two fixed m -bit values according to the single bit $X_{f,u,w}$. Thus every target input bit is a constant, $X_{f,u,w}$, or $\bar{X}_{f,u,w}$, and the map is a projection.

Finally, Equations (35) and (37) give

$$m_\Sigma(P, N) = R(P) \lceil \log(2kN + 1) \rceil = O(k \log(2kN)). \quad (146)$$

□

D Parameter Completions

This appendix completes the parameter picture left implicit in the main theorems. We first record how each projection extends beyond its constructed constraint budget and then optimize the source host size when the available row count or bit width is smaller than the principal $O(k \log n)$ -scale budget used in Theorem 5.1.

We do not pursue an analogous interpolation in the k -OV dimension D , i.e. the counterpart of Theorems D.2 and D.3 below. Unlike the two list targets, whose construction floors depend only on k , the k -OV floor $k + s_{k,n} \geq k + \ell_n$ grows with the target universe: the identifier block enforcing the promise in Definition 3.1 is sized by n , not by the source host N . For fixed k , this floor and the principal budget $C_{OV}k \log n$ are both $\Theta(\log n)$, so varying N trades one constant against another rather than opening a regime analogous to $m \ll k \log(n/k)$.

The final subsection records two boundary regimes outside this interpolation: a small-modulus degeneracy and a parity-hard restriction at the minimal universe $n = k$.

D.1 Extension beyond the constructed budget

Remark D.1 (Monotone extension of the constraint parameter). Let P be a pattern with $e(P) = k - 1$, and suppose one of the projections in Theorem 3.4 has been constructed at its native constraint parameter. Then the same source instance projects to every larger admissible target parameter: from $D_0 := D_0(P, N, n)$ to every $D \geq D_0$ for k -OV, from $m_0 := m_\oplus(P, N)$ to every $m \geq m_0$ for odd- k k -XOR, and from $m_0 := m_\Sigma(P, N)$ to every $m \geq m_0$ for k -SUM.

Proof. For k -OV, append $D - D_0$ coordinates that are 0 on every vector; this preserves pairwise distinctness within each array and makes the product of every one-vector-per-class selection identically zero on the new coordinates. For k -XOR, append $m - m_0$ all-zero rows, which do not change the sum of any support. For k -SUM, let $\hat{a}_j \in \mathbb{Z}_{2^{m_0}}$ be the residue produced at width m_0 and replace it at width m by $a_j := 2^{m-m_0}\hat{a}_j \in \mathbb{Z}_{2^m}$, as in Equation (40). For every support S ,

$$\sum_{j \in S} a_j \equiv 0 \pmod{2^m} \iff 2^m \mid 2^{m-m_0} \sum_{j \in S} \hat{a}_j \iff \sum_{j \in S} \hat{a}_j \equiv 0 \pmod{2^{m_0}}.$$

Every added target bit is still a Boolean function of at most the single source bit on which the original residue depended, and is therefore a constant, that bit, or its negation. Thus all three extensions remain depth-zero projections; this remark consolidates operations already built into Theorems 4.2, 4.4 and 4.5 rather than adding a separate reduction. □

Growing k for k -OV. Although the identifier floor leaves no qualitatively new fixed- k window, the explicit growing-pattern source bounds do yield a shifted interpolation above that floor. For every pattern P with $e(P) = k - 1$, one has $|\mathcal{I}(P)| \leq 2e(P)$ and therefore

$$D_0(P, N, n) \leq 4k\ell_N + k + 2\ell_n.$$

Define

$$w_{\text{OV}} := \left[\min \left\{ \frac{1}{2} \log n, \frac{D - k - 2\ell_n}{4k} \right\} \right]_+.$$

If $w_{\text{OV}} > 0$, put $N := 2^{\lfloor w_{\text{OV}} \rfloor}$. Then $\ell_N = \lceil \log N \rceil = \lfloor w \rfloor \leq w$, $N^2 \leq n$ and $D_0(P, N, n) \leq D$. Hence Theorem 4.2 applied to P_k^{exp} , followed by Theorems 4.2 and 6.4(i), gives

$$w_{\text{OV}} \geq 2 \log k + 1 \implies \log \text{size}_2(k\text{-OV}_{n,D,k}) \geq \gamma_D k \log N \geq \gamma_D k (w_{\text{OV}} - 1) \geq \frac{\gamma_D}{2} k w_{\text{OV}}$$

for every $k \geq k_D$. Likewise, for every fixed $0 < \gamma < 1/8$, applying the projection to P_k^{exp} and using Theorem 6.8 gives

$$w_{\text{OV}} \geq \max\{2, C_E(\gamma) \log k + 1\} \implies \log \text{size}_{\Sigma_3}(k\text{-OV}_{n,D,k}) \geq \frac{\gamma}{2} k w_{\text{OV}}$$

for every $k \geq k_E(\gamma)$. In particular, once $D \geq 2(k + 2\ell_n)$ and the corresponding onset condition holds, either conclusion has exponent $\Omega(\min\{k \log n, D\})$.

D.2 General row counts for odd k -XOR

The exact formula in Equation (33) yields a convenient absolute estimate. Since $e(P) = k - 1$, $|\mathcal{I}(P)| \leq 2e(P)$, and $r_{\text{guard}}(P, N) = \lceil 2e(P) \log N + e(P) + 1 \rceil$, for every $k \geq 2$ and $N \geq 1$,

$$m_{\oplus}(P, N) \leq 4(k - 1) \log(2N) + 2(k - 1) + 3 \leq 8k \log(2N). \quad (147)$$

The bound depends only on the edge count, so it applies uniformly to the expander pattern family used below.

Theorem D.2 (General row count for odd k -XOR). *There is a universal constant $c_{\oplus} > 0$ with the following property. For every fixed depth d and every fixed odd integer $k \geq k_0$, there is a threshold $L_{\oplus}(d, k)$ such that, whenever*

$$w_{\oplus} := \min \left\{ \log \frac{n}{k}, \frac{m}{k} \right\} \geq L_{\oplus}(d, k), \quad (148)$$

one has

$$\log \text{size}_d(k\text{-XOR}_{n,m,k}) \geq c_{\oplus} \min \left\{ k \log \frac{n}{k}, m \right\}. \quad (149)$$

Equivalently, $\text{size}_d(k\text{-XOR}_{n,m,k}) \geq \exp(\Omega(\min\{k \log(n/k), m\}))$, with a universal implied constant; only the onset threshold depends on (d, k) .

Proof. Fix a universal source rate $\gamma_{\star}$ with $0 < \gamma_{\star} < c_{\kappa}$, and let $N_0(\gamma_{\star}, d, k)$ be the threshold supplied by Corollary 2.6 for P_k^{exp} . Define

$$Y_{\oplus} := \min \left\{ 2^{m/(8k)-1}, \sqrt{\frac{n-1}{k-1}} \right\}, \quad N := \lfloor Y_{\oplus} \rfloor. \quad (150)$$

The second term ensures $1 + (k-1)N^2 \leq n$. The first gives $\log(2N) \leq m/(8k)$, so Equation (147) implies $m_\oplus(P_k^{\text{exp}}, N) \leq m$. Thus Theorem 4.4 gives a projection from P_k^{exp} -SUB $_N$ to k -XOR $_{n,m,k}$.

Set $x := w_\oplus$. Once x is above an absolute constant, one has $n \geq k$ and $Y_\oplus \geq 2$. Using $(n-1)/(k-1) \geq n/k$ and $\lfloor Y_\oplus \rfloor \geq Y_\oplus/2$ gives

$$\begin{aligned} \log N &\geq \min \left\{ \frac{m}{8k} - 1, \frac{1}{2} \log \frac{n-1}{k-1} \right\} - 1 \\ &\geq \min \left\{ \frac{x}{8} - 1, \frac{x}{2} \right\} - 1 \geq \frac{x}{16}, \end{aligned} \quad (151)$$

where the last inequality holds after increasing the absolute lower bound on x . Choose $L_\oplus(d, k)$ large enough that Equation (151) also forces $N \geq N_0(\gamma_\star, d, k)$. Projection monotonicity and Corollary 2.6 then yield

$$\begin{aligned} \log \text{size}_d(k\text{-XOR}_{n,m,k}) &\geq \gamma_\star(k-1) \log N \\ &\geq \frac{\gamma_\star}{32} k w_\oplus = \frac{\gamma_\star}{32} \min \left\{ k \log \frac{n}{k}, m \right\}, \end{aligned}$$

using $k-1 \geq k/2$. Taking $c_\oplus := \gamma_\star/32$ proves the theorem. $\square$

D.3 General bit widths for k -SUM

For this construction, the integer packing depends on kN , rather than on N alone. Indeed, Equations (35) and (37) imply, for every $k \geq 2$ and $N \geq 1$,

$$m_\Sigma(P, N) = R(P) \lceil \log(2kN + 1) \rceil \leq 9k \log(2kN), \quad (152)$$

because $R(P) \leq 3k$ and $\lceil \log(2kN + 1) \rceil \leq 3 \log(2kN)$. The factor k inside the logarithm produces the $-\log k$ term below.

Theorem D.3 (General bit width for k -SUM). *There is a universal constant $c_\Sigma > 0$ with the following property. For every fixed depth d and every fixed integer $k \geq k_0$, of either parity, there is a threshold $L_\Sigma(d, k)$ such that, with*

$$w_\Sigma := \left[\min \left\{ \log \frac{n}{k}, \frac{m}{9k} - \log(2k) \right\} \right]_+, \quad [x]_+ := \max\{x, 0\}, \quad (153)$$

the implication

$$w_\Sigma \geq L_\Sigma(d, k) \implies \log \text{size}_d(k\text{-SUM}_{n,m,k}) \geq c_\Sigma k w_\Sigma \quad (154)$$

holds. The constant c_Σ is universal; only the onset threshold depends on (d, k) .

Proof. Use the same universal $\gamma_\star$ as in the proof of Theorem D.2, and let $N_0(\gamma_\star, d, k)$ be the corresponding source threshold. Define

$$Y_\Sigma := \min \left\{ \frac{2^{m/(9k)}}{2k}, \sqrt{\frac{n-1}{k-1}} \right\}, \quad N := \lfloor Y_\Sigma \rfloor. \quad (155)$$

The column budget again holds by the second term. The first term gives $\log(2kN) \leq m/(9k)$, so Equation (152) implies $m_\Sigma(P_k^{\text{exp}}, N) \leq m$. Hence Theorem 4.5 gives a projection from P_k^{exp} -SUB $_N$ to k -SUM $_{n,m,k}$.

Write $x := w_\Sigma$. If $x > 0$, then both entries inside the minimum in Equation (153) are positive. Since $(n-1)/(k-1) \geq n/k$, one has

$$\log Y_\Sigma \geq \min \left\{ \frac{m}{9k} - \log(2k), \frac{1}{2} \log \frac{n}{k} \right\} \geq \frac{x}{2}.$$

Once $x \geq 4$, this gives $Y_\Sigma \geq 2$ and therefore

$$\log N \geq \log Y_\Sigma - 1 \geq \frac{x}{4}. \quad (156)$$

Choose $L_\Sigma(d, k)$ large enough that Equation (156) forces $N \geq N_0(\gamma_\star, d, k)$. By projection monotonicity and Corollary 2.6,

$$\begin{aligned} \log \text{size}_d(k\text{-SUM}_{n,m,k}) &\geq \gamma_\star(k-1) \log N \\ &\geq \frac{\gamma_\star}{8} k w_\Sigma, \end{aligned}$$

where $k-1 \geq k/2$ is used in the last step. Taking $c_\Sigma := \gamma_\star/8$ proves the theorem. $\square$

Growing k at depth two and three. The fixed- k qualification in Theorems D.2 and D.3 enters only through Corollary 2.6. Replacing that source bound by Theorem 6.4(i), the same host choices give, for every $k \geq k_D$, the implications

$$\begin{aligned} w_\oplus \geq 32 \log k &\implies \log \text{size}_2(k\text{-XOR}_{n,m,k}) \geq \frac{\gamma_D}{16} k w_\oplus && \text{for odd } k, \\ w_\Sigma \geq 8 \log k &\implies \log \text{size}_2(k\text{-SUM}_{n,m,k}) \geq \frac{\gamma_D}{4} k w_\Sigma. \end{aligned}$$

Indeed, Equations (151) and (156) then force $N \geq k^2$. Likewise, fix $0 < \gamma < 1/8$ and enlarge $C_E(\gamma)$ to at least 2. Replacing the source bound by Theorem 6.8, for every $k \geq k_E(\gamma)$ the same calculations give

$$\begin{aligned} w_\oplus \geq 16 C_E(\gamma) \log k &\implies \log \text{size}_{\Sigma_3}(k\text{-XOR}_{n,m,k}) \geq \frac{\gamma}{16} k w_\oplus && \text{for odd } k, \\ w_\Sigma \geq 4 C_E(\gamma) \log k &\implies \log \text{size}_{\Sigma_3}(k\text{-SUM}_{n,m,k}) \geq \frac{\gamma}{4} k w_\Sigma. \end{aligned}$$

Together with the preceding k -OV interpolation, the parameter completion is uniform for growing k at depth two and in the top-disjunction depth-three orientation once the corresponding logarithmic host size w_{OV} , $w_\oplus$, or w_Σ is $\Omega(\log k)$. No analogous low-row consequence follows from the even- K lift: its tag matrix and inert-column padding already cost $\Theta(K \log n')$ rows independently of the source host size.

The $[\cdot]_+$ and the subtraction of $\log(2k)$ are substantive. They keep the theorem silent whenever the available width does not permit a growing source host. Far above that boundary, the familiar minimum form is recovered.

Corollary D.4 (High-width minimum form for k -SUM). *For every fixed depth d and every fixed integer $k \geq k_0$, there is a threshold $L'_\Sigma(d, k)$ such that, whenever*

$$m \geq 18k \log(2k) \quad \text{and} \quad \min \left\{ \log \frac{n}{k}, \frac{m}{k} \right\} \geq L'_\Sigma(d, k), \quad (157)$$

one has

$$\text{size}_d(k\text{-SUM}_{n,m,k}) \geq \exp \left(\Omega \left(\min \left\{ k \log \frac{n}{k}, m \right\} \right) \right), \quad (158)$$

with a universal implied constant.

Proof. Put $u := \min\{\log(n/k), m/k\}$. The first hypothesis gives $m/(9k) - \log(2k) \geq m/(18k)$, and therefore

$$w_\Sigma \geq \min\left\{\log \frac{n}{k}, \frac{m}{18k}\right\} \geq \frac{u}{18}.$$

Choose $L'_\Sigma(d, k) := 18L_\Sigma(d, k)$. Then Theorem D.3 applies and gives

$$\log \text{size}_d \geq \Omega(ku) = \Omega\left(\min\left\{k \log \frac{n}{k}, m\right\}\right).$$

□

D.4 Two boundary cases for k -SUM

The next two propositions show why no uniform lower bound can cover the silent regime: one boundary family is constant, while another contains PARITY under a restriction.

Proposition D.5 (Erdős–Ginzburg–Ziv degeneracy). *If $k = 2^m$ and $n \geq 2k - 1$, then k -SUM $_{n,m,k}$ is identically one.*⁷

Proof. Write $G = \mathbb{Z}_{2^m}$, a cyclic group of order $2^m = k$. Fix an arbitrary input $x = (x_1, \dots, x_n) \in G^n$ and any $S \subseteq [n]$ with $|S| = 2k - 1$, which exists since $n \geq 2k - 1$. The Erdős–Ginzburg–Ziv theorem [EGZ61] asserts that every sequence of $2k - 1$ elements of G , repetitions permitted, admits a subsequence of length exactly k summing to 0;⁸ applied to $(x_i)_{i \in S}$ it yields $I \subseteq S$ with $|I| = k$ and $\sum_{i \in I} x_i = 0$ in G . Since I is a set of k distinct indices of $[n]$, it witnesses k -SUM $_{n,m,k}(x) = 1$; as x was arbitrary, the function is identically one. □

Proposition D.6 (The parity corner at every bit width). *Under the convention that PARITY $_k(x) = 1$ on odd Hamming weight, for every $m \geq 1$,*

$$k\text{-SUM}_{k,m,k}(2^{m-1}x_1, \dots, 2^{m-1}x_k) = 1 - \text{PARITY}_k(x). \quad (159)$$

Consequently, for every fixed depth $d \geq 2$ and every $m \geq 1$, once k is sufficiently large as a function of d ,

$$\text{size}_d(k\text{-SUM}_{k,m,k}) \geq \exp(\Omega_d(k^{1/(d-1)})). \quad (160)$$

Proof. When $n = k$, the only size- k support is all of $[k]$. On the displayed restriction, its sum is $2^{m-1}|x|$, which vanishes modulo 2^m exactly when $|x|$ is even; this proves Equation (159). For Equation (160), first observe that $\text{size}_d(f) = \text{size}_d(\neg f)$ exactly in our model because only non-input AND/OR gates are counted and complementing a leaf literal introduces no gate. Given a depth- d circuit for f , swap every $\wedge$ with $\vee$, complement every input literal, and interchange constant leaves 0 and 1 if they occur. Induction on gate height shows that the transformed gate computes the complement of the original gate, so the transformed circuit has the same size and depth and computes $\neg f$; applying the same involution to a circuit for $\neg f$ gives the reverse inequality. The substitution in Equation (159) realizes $1 - \text{PARITY}_k$ as a depth-zero restriction of the total indexed-list function $k\text{-SUM}_{k,m,k}$. Therefore Lemma 3.3, the dualization equality, and the fixed-depth PARITY bound recorded in Section 6.4 imply Equation (160) once k exceeds the depth-dependent onset. □

⁷This proposition and its proof were suggested by the AI language model GPT-5.5; the authors independently verified the argument before inclusion.

⁸[EGZ61] phrases this for a set of $2k - 1$ integers and a subset of k of them. The sequence form is what the proof there establishes, and the two are equivalent: given $g_1, \dots, g_{2k-1} \in \mathbb{Z}_k$ with representatives $\tilde{g}_i \in [0, k)$, the integers $z_i := \tilde{g}_i + ki$ are pairwise distinct, since $z_i \in [ki, (k+1)i)$ and these intervals are disjoint, and satisfy $z_i \equiv g_i \pmod{k}$.

The same fixed-support floor is shared by k -XOR. For every $m \geq 1$, set rows $2, \dots, m$ to zero and leave the first-row bits $x_1, \dots, x_k$ free. Since $n = k$, the only candidate support is $[k]$, and hence

$$k\text{-XOR}_{k,m,k}((x_1, 0^{m-1}), \dots, (x_k, 0^{m-1})) = 1 - \text{PARITY}_k(x). \quad (161)$$

This is a genuine depth-zero restriction of the total indexed-list function in Definition 3.1, not a promise-sensitive map. All k first-row bits remain free, and the substitution uses only literals and constants. By Lemma 3.3, the preceding dualization equality, and the fixed-depth PARITY bound recorded in Section 6.4, for every fixed depth $d \geq 2$ and every $m \geq 1$, once k is sufficiently large as a function of d ,

$$\text{size}_d(k\text{-XOR}_{k,m,k}) \geq \exp(\Omega_d(k^{1/(d-1)})). \quad (162)$$

Consequently, no fixed depth yields a polynomial-size per-support verifier uniformly in k for either list target. As Propositions F.2 and F.3 make explicit, the later comparison instead distinguishes the $m 2^{k-1}$ -clause CNF for a fixed k -XOR support from the $2^{(k-1)m}$ -term direct DNF for a fixed k -SUM support, and then from the block-carry Σ_3 verifier.

The two propositions concern different boundary phenomena. At $k = 2^m$ and $n \geq 2k - 1$, the function is trivial by additive combinatorics, whereas at $n = k$ every bit width contains a parity-hard two-point restriction. Neither contradicts Theorem D.3: in both cases the host-size parameter w_Σ is below the theorem's onset threshold. The hypothesis $n \geq 2k - 1$ is optimal: at $n = 2k - 2$ the input consisting of $k - 1$ zeros and $k - 1$ ones has no zero-sum witness, since any k of its coordinates contain j ones with $1 \leq j \leq k - 1$, and $j \not\equiv 0 \pmod{k}$.

E Supplementary Results for k -XOR

This appendix contains two black-box extensions of the principal odd- k k -XOR results. The first is a projection from an odd source parameter to an even target parameter, parameterized by a supplied admissible decomposition of the target parameter. The second is the amplified PARITY projection stated in Theorem 6.9, together with a matching limitation for every projection-based PARITY reduction.

E.1 The even- K lifting projection

The threshold k_0 fixed in Lemma 2.5 already satisfies $k_0 \geq 3$. The lift uses a nonuniformly fixed tag matrix to force every target witness of weight at most K to lie in the image of a repetition code. The t -fold repetition makes an odd-weight source witness of weight r occupy tr positions, while the s auxiliary copies of the source vector's parity bit raise that weight to $tr + s = K$; the image of an even-weight source vector has weight divisible by t , so the condition $t \nmid s$ eliminates that branch, with the case $s = 0$ handled by the opposite parities of r and an even weight.⁹

The next lemma is the standard probabilistic parity-check-matrix argument; we include it to fix the nonuniform tag construction and its exact row count.

Lemma E.1 (Separating tag matrix). *Let $n' \geq 2$ and $K \geq 1$, and let $U \subseteq \mathbb{F}_2^{n'}$ be a linear subspace. Put*

$$q := (K + 1)\lceil \log n' \rceil + 1. \quad (163)$$

⁹While we were attempting to derive a direct projection for even k , the AI language model GPT-5.5 suggested this odd-to-even lifting approach, namely the repetition code with auxiliary parity copies together with a separating tag matrix confining low-weight target witnesses to its image; the authors independently verified the construction before recording it here.

There is a matrix $\Pi \in \mathbb{F}_2^{q \times n'}$ such that $U \subseteq \ker \Pi$ and every $y \in \ker \Pi$ with $|y| \leq K$ belongs to U .

Proof. Let $\pi : \mathbb{F}_2^{n'} \rightarrow \mathbb{F}_2^{n'}/U$ be the quotient map. Choose q independent uniformly random linear functionals $\varphi_1, \dots, \varphi_q$ on $\mathbb{F}_2^{n'}/U$, and let Π be the matrix representing $(\varphi_1 \circ \pi, \dots, \varphi_q \circ \pi)$. Then $U \subseteq \ker \Pi$ deterministically. For each fixed $y \notin U$, the class $\pi(y)$ is nonzero, so every φ_h annihilates it with probability $1/2$, independently, and therefore

$$\Pr[\Pi y = 0^q] = 2^{-q}. \quad (164)$$

The number of vectors of weight at most K is at most

$$\sum_{j=0}^K \binom{n'}{j} \leq \sum_{j=0}^K (n')^j \leq 2(n')^K. \quad (165)$$

By Equation (163), the expected number of vectors $y \notin U$ with $|y| \leq K$ and $\Pi y = 0^q$ is at most

$$2(n')^K 2^{-q} \leq (n')^{-1} < 1.$$

Hence some fixed choice of the functionals leaves no such vector, and the corresponding matrix has the required property. $\square$

Lemma E.2 (Odd-to-even lifting projection). *Let $r \geq 1$ be odd, let $t \geq 2$ and $s \geq 0$ be integers, put $K := tr + s$, and assume that K is even and that either $s = 0$ or $t \nmid s$. Then, for all $n, m \geq 1$, there is a projection*

$$k\text{-XOR}_{n,m,r} \leq_{\text{proj}} k\text{-XOR}_{tn+s,m+q,K}, \quad q = (K+1)\lceil \log(tn+s) \rceil + 1. \quad (166)$$

Proof. For $\xi \in \mathbb{F}_2^n$, write $p(\xi) := \bigoplus_{i=1}^n \xi_i$ and define the linear repetition map $\text{rep} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{tn+s}$ by

$$\text{rep}(\xi) := (\underbrace{\xi_1, \dots, \xi_1}_t, \dots, \underbrace{\xi_n, \dots, \xi_n}_t, \underbrace{p(\xi), \dots, p(\xi)}_s). \quad (167)$$

Put $U := \text{im rep}$. The map is injective because ξ_i is recovered from the first copy of the i th block, and its Hamming weight satisfies

$$|\text{rep}(\xi)| = t|\xi| + s(|\xi| \bmod 2). \quad (168)$$

Let $n'' := tn + s$, and let $\Pi \in \mathbb{F}_2^{q \times n''}$ be supplied by Lemma E.1 for U and K , with columns $\Pi_1, \dots, \Pi_{n''}$. The target universe contains t copies of every source index followed by s auxiliary indices. For each target position j corresponding to the first copy of source index i , the target column is the concatenation of the source column $a_i \in \mathbb{F}_2^m$ (its payload) with that position's tag Π_j ; every remaining copy and every auxiliary position j has payload 0^m and tag Π_j . Every target bit is therefore a constant or a source bit, so this map is a projection.

Suppose first that ξ is the indicator of a source witness. Then $|\xi| = r$ and the selected source payloads sum to zero. Since r is odd, Equation (168) gives $|\text{rep}(\xi)| = tr + s = K$; moreover, $\text{rep}(\xi) \in U \subseteq \ker \Pi$, and the selected target payloads are exactly the selected source payloads. Thus $\text{rep}(\xi)$ is a target witness.

Conversely, let $y \in \mathbb{F}_2^{n''}$ indicate a target witness. Its tag sum is zero and $|y| = K$, so Lemma E.1 gives $y \in U$. Write $y = \text{rep}(\xi)$ for the unique $\xi \in \mathbb{F}_2^n$. If $|\xi|$ is odd, then Equation (168) and $|y| = K = tr + s$ imply $|\xi| = r$. If $|\xi|$ is even, then the same equations give $t|\xi| = tr + s$, so $t \mid s$. When $s > 0$ this contradicts the hypothesis, and when $s = 0$ it gives $|\xi| = r$, contradicting the opposite parities of $|\xi|$ and r . Hence $|\xi| = r$. Finally, the target payload sum is the XOR of the source columns indexed by $\text{supp}(\xi)$, so its vanishing makes ξ a source witness. $\square$

Remark E.3 (Nonuniformity and arbitrary target universes). The matrix Π is chosen by the probabilistic method and fixed once for each parameter tuple before the input is read, which is legitimate for the nonuniform circuit lower bounds considered here. To reach an arbitrary target universe size $n' \geq t + s$, let $n := \lfloor (n' - s)/t \rfloor$ and $\rho := n' - (tn + s) < t$. Apply Lemma E.2 to the first $tn + s$ columns and append ρ columns that are zero on every existing row, giving each one a private guard row equal to 1 on that column and 0 elsewhere. No zero-sum support uses any appended column, and the additional row count is less than t . Appending all-zero rows then reaches every larger target row count without changing the function.

E.2 Consequences of the lift

The lift is independent of the source of the odd- k lower bound. We first recover the fixed- k statement quoted in Section 5, then apply it to the proved top-disjunction depth-three theorem and to the conditional all-depth theorem.

Call a tuple (t, r, s) admissible for an even integer K when $t \geq 2$, $r \geq 1$ is odd, $0 \leq s < 2t$, $K = tr + s$, and either $s = 0$ or $t \nmid s$. By Lemma E.2, every admissible tuple gives an odd-to-even projection. We use the same elementary estimates in the consequences below. Let (t, r, s) be admissible for K , assume $r \geq 3$, and for a target universe size n' put $n := \lfloor (n' - s)/t \rfloor$. Since $tn + s \leq n' < t(n + 1) + s < t(n + 3)$ and $tr \leq K < t(r + 2)$, whenever $n \geq 3$ and $n'/K \geq 4$ one has

$$\frac{n}{r} \geq \frac{n'}{2K} \geq \left(\frac{n'}{K}\right)^{1/2}. \quad (169)$$

Moreover, the tag matrix and the inert-column padding add $O(K \log n')$ rows in total. More explicitly, let q be the tag-row count in Equation (166) and let $\rho < t$ be the number of inert columns in Remark E.3. There is a universal constant $C'_\oplus > 0$ such that, given a target budget $m' \geq C'_\oplus K \log n'$, the quantity $\tilde{m} := m' - q - \rho$ exceeds each source row budget used below. Applying the source theorem at row count $\tilde{m}$, then adding the q tag rows and the ρ private guard rows, fits within m' , and all-zero padding reaches the budget exactly.

Corollary E.4 (Fixed even K from an admissible odd parameter). *Let K be even, and let (t, r, s) be admissible for K with $r \geq k_0$. For every fixed depth d there is a threshold $n'_1(d, K, t, r, s)$ such that, for every $n' \geq n'_1(d, K, t, r, s)$ and every $m' \geq C'_\oplus K \log n'$,*

$$\text{size}_d(k\text{-XOR}_{n', m', K}) \geq (n'/K)^{\beta(r-1)/2}, \quad (170)$$

where β is the universal target-rate constant of Theorem 5.1.

Proof. Put $n := \lfloor (n' - s)/t \rfloor$ and use the arbitrary-universe version of the lift from Remark E.3. For n' above a threshold depending on (d, K, t, r, s) , the source universe satisfies the onset condition in Theorem 5.1, the source theorem applies at row count $\tilde{m}$, and the common row estimate above fits the lifted instance within every $m' \geq C'_\oplus K \log n'$. Projection monotonicity therefore gives

$$\text{size}_d(k\text{-XOR}_{n', m', K}) \geq \text{size}_d(k\text{-XOR}_{n, \tilde{m}, r}) \geq (n/r)^{\beta(r-1)}. \quad (171)$$

After enlarging the threshold so that Equation (169) applies, substituting that estimate into Equation (171) proves Equation (170). $\square$

Corollary E.5 (Depth three for even K from an admissible odd parameter). *There are universal constants $a_E > 0$, $k_E^\oplus \in \mathbb{N}$, and $C_E^\oplus > 1$ such that the following holds. Let K be even, and let (t, r, s) be admissible for K with $r \geq k_E^\oplus$. If $n' \geq K^{C_E^\oplus}$, then, for every $m' \geq C'_\oplus K \log n'$,*

$$\text{size}_{\Sigma_3}(k\text{-XOR}_{n', m', K}) \geq (n'/K)^{a_E r}. \quad (172)$$

Proof. Fix $\gamma_E := 1/16$ and $c_E := 1/64$. Take $k_E^\oplus \geq \max\{k_1(\gamma_E, c_E), k_0\}$, and choose $C_E^\oplus$ sufficiently large that $n' \geq K^{C_E^\oplus}$ makes $n := \lfloor (n' - s)/t \rfloor$ satisfy both $n \geq r^{C'(\gamma_E, c_E)}$ and Equation (169). Apply the k -XOR line of Theorem 6.8 at parameter r and row count $\tilde{m}$, then compose with the lifting projection; the common row estimate above fits within $m' \geq C'_\oplus K \log n'$, and orientation preservation gives $\text{size}_{\Sigma_3}(k\text{-XOR}_{n', m', K}) \geq (n/r)^{c_E r} \geq (n'/K)^{c_E r/2}$. Taking $a_E := c_E/2$ proves the claim. $\square$

Corollary E.6 (Conditional all-depth bound for even K from an admissible odd parameter). *Assume Conjecture 7.1, and fix a depth d . There are $k_F^\oplus(d) \in \mathbb{N}$ and $C_F^\oplus(d) > 1$ such that the following holds. Let K be even, and let (t, r, s) be admissible for K with $r \geq k_F^\oplus(d)$. If $n' \geq K^{C_F^\oplus(d)}$ and $m' \geq C'_\oplus K \log n'$, then*

$$\text{size}_d(k\text{-XOR}_{n', m', K}) \geq \exp(\Omega_d(r \log n')). \quad (173)$$

Proof. Let $k_F^\oplus(d)$ be at least $k_F(d)$ from Theorem 7.3 and at least k_0 . Choose $C_F^\oplus(d) \geq \max\{C_F(d) + 3, 4\}$, and put $n := \lfloor (n' - s)/t \rfloor$. Since $r \leq K$ and $t \leq K$, the relation $n' \geq K^{C_F^\oplus(d)}$ makes $n \geq r^{C_F(d)}$ after increasing $k_F^\oplus(d)$ if necessary, so Theorem 7.3 applies to the odd- r source instance at row count $\tilde{m}$. The common row estimate above fits the lift within $C'_\oplus K \log n'$. Thus

$$\text{size}_d(k\text{-XOR}_{n', m', K}) \geq \exp(\Omega_d(r \log n)). \quad (174)$$

Moreover, $s < 2t$, $t \leq K$, and $n' \geq K^4$ imply

$$n = \left\lfloor \frac{n' - s}{t} \right\rfloor > \frac{n'}{K} - 3 \geq \sqrt{n'},$$

for every $K \geq 2$, and hence $\log n \geq \frac{1}{2} \log n'$. Substituting this estimate into Equation (174) proves Equation (173). $\square$

E.3 The row-rich PARITY route

We now prove the two statements from Section 6.4. The construction embeds a parity instance of length $\Theta(k^2 \log(n/k))$ while using only $O(n)$ target rows.

Proof of Theorem 6.9. Let b, B, ℓ be as in Equation (86), and put $q := 2^b$. The input variables are

$$x = (x_{r,t})_{r \in [B], t \in [k-1]} \in \{0, 1\}^\ell. \quad (175)$$

The target universe contains one anchor column $a_\star$ and, for every $t \in [k-1]$, a block of q columns (t, u) indexed by $u \in \{0, 1\}^b$. Since $q \leq n/(2k)$,

$$n_0 := 1 + (k-1)q < n, \quad (176)$$

and the remaining $n - n_0$ positions are dummies, each equipped with one private guard row. All dummy entries on the selector and data rows introduced below are set to 0. Thus no zero-sum support contains a dummy.

For each $t \in [k-1]$, add a selector row that is 1 on the anchor and on every column of block t , and 0 elsewhere. Let ζ indicate whether a support contains the anchor, and let c_t be the number of selected columns from block t . Vanishing of the selector rows gives $\zeta + c_t \equiv 0 \pmod{2}$ for every

t . If $\zeta = 0$, every c_t is even and the support size $\sum_t c_t = k$ is even, contradicting odd k . Hence $\zeta = 1$, every c_t is odd, and

$$1 + \sum_{t=1}^{k-1} c_t = k.$$

Because there are $k - 1$ positive odd integers in the sum, each $c_t = 1$. Every zero-sum support therefore has the form

$$\{a_\star\} \cup \{(t, u_t) : t \in [k - 1]\} \quad (177)$$

for a unique tuple $u = (u_1, \dots, u_{k-1}) \in (\{0, 1\}^b)^{k-1}$.

Define $y : (\{0, 1\}^b)^{k-1} \rightarrow \{0, 1\}^B$ by

$$y(u)_{(t-1)b+j} := u_{t,j} \quad \text{for } t \in [k - 1], j \in [b], \quad (178)$$

$$y(u)_B := 1 \oplus \bigoplus_{t \in [k-1]} \bigoplus_{j \in [b]} u_{t,j}. \quad (179)$$

The image of y is exactly the set of odd-parity strings in $\{0, 1\}^B$. We add B data rows. For a row $r = (t_r - 1)b + j_r \leq (k - 1)b$, put 0 on the anchor, and put the literal

$$x_{r,t} \oplus \mathbf{1}[t = t_r]u_{j_r} \quad (180)$$

on the column (t, u) . On the last row B , put 1 on the anchor and put

$$x_{B,t} \oplus \bigoplus_{j \in [b]} u_j \quad (181)$$

on the column (t, u) . Since u is part of the column index, every expression in Equations (180) and (181) is either an input variable or its negation, and the construction is a projection.

For the support in Equation (177), data row r has parity zero exactly when

$$\bigoplus_{t \in [k-1]} x_{r,t} = y(u)_r. \quad (182)$$

Such a tuple u exists if and only if the vector $(\bigoplus_t x_{r,t})_{r \in [B]}$ has odd parity, which is equivalent to

$$\bigoplus_{r \in [B]} \bigoplus_{t \in [k-1]} x_{r,t} = 1. \quad (183)$$

Thus the target accepts exactly the odd-parity inputs, proving Equation (87).

The number of rows before zero padding is

$$m_0 = (k - 1) + B + (n - n_0). \quad (184)$$

For $n \geq 4k$, one has $k - 1 \leq n/4$ and $B \leq k \log(n/k) + 1 \leq n/2 + 1$, where the last inequality uses $\log x \leq x/2$ for $x \geq 4$. Hence $m_0 \leq 2n$. Also, $k - 1 = \Theta(k)$ and $b = \Theta(\log(n/k))$, so

$$\ell = (k - 1)B = \Theta(k^2 \log(n/k)). \quad (185)$$

Appending zero rows extends the projection to every $m \geq m_0$. Every variable in Equation (175) remains free, every target bit is a constant or a possibly negated source literal, and k -XOR is a total function, so Equation (87) is an ordinary projection with no promise-side obligation. For fixed d , once ℓ exceeds the depth-dependent onset, the PARITY bound recorded in Section 6.4 gives size $\exp(\Omega_d(\ell^{1/(d-1)}))$, and Equation (185) yields Equation (88). $\square$

Proof of Proposition 6.10. Write $R(x) = (A_1(x), \dots, A_n(x))$. For each support $S \in \binom{[n]}{k}$, define

$$V_S := \left\{ x \in \{0, 1\}^\ell : \bigoplus_{j \in S} A_j(x) = 0^m \right\}. \quad (186)$$

Every output entry of a projection is a constant or one possibly negated input bit, hence an affine function over $\mathbb{F}_2$ involving at most one variable. Each row equation in Equation (186) is therefore affine and involves at most k input variables. Thus V_S is either empty or an affine subspace cut out by k -sparse equations.

Suppose V_S is nonempty. Choose a maximal linearly independent subset of the coefficient rows defining it, and write the resulting equivalent system as $M_S x = \eta_S$, where M_S has rank r_S and every row has Hamming weight at most k . For every $x \in V_S$, the support S witnesses acceptance of k -XOR($R(x)$), so Equation (89) implies that the all-ones linear functional has value 1 throughout V_S . Writing $V_S = x_0 + \ker M_S$, this means that the all-ones vector belongs to $(\ker M_S)^\perp$, which is the row space of M_S . It is therefore the XOR of at most r_S rows of M_S . The union of their supports has size at most $r_S k$, and consequently

$$\ell \leq r_S k. \quad (187)$$

It follows that

$$|V_S| = 2^{\ell - r_S} \leq 2^{\ell - \ell/k}. \quad (188)$$

Every odd-parity input is accepted and hence lies in V_S for at least one support S . The $2^{\ell-1}$ odd inputs are therefore covered by at most $\binom{n}{k}$ affine spaces satisfying Equation (188). Hence

$$2^{\ell-1} \leq \binom{n}{k} 2^{\ell - \ell/k}. \quad (189)$$

Cancelling the common factor, taking logarithms, and using $\binom{n}{k} \leq (en/k)^k$ gives

$$\frac{\ell}{k} - 1 \leq \log \binom{n}{k} \leq k \log(en/k),$$

which rearranges to Equation (90). □

F Elementary Upper Bounds and Quantitative Tightness

This appendix supplies the elementary circuits used to calibrate the lower bounds. Circuit size counts non-input gates, as fixed in Section 2.1; fan-in therefore affects the depth and the formula semantics but not the gate count directly.

F.1 Brute-force circuits

Proposition F.1 (Brute-force circuit for k -OV). *For all n, D, k ,*

$$\text{size}_{\Sigma_3}(k\text{-OV}_{n,D,k}) \leq 1 + n^k(D+1). \quad (190)$$

In particular, if $D = O(k \log n)$ and k is fixed, then

$$\text{size}_{\Sigma_3}(k\text{-OV}_{n,D,k}) \leq n^{k+o(1)}. \quad (191)$$

The fixed- k hypothesis is unnecessary for the sharper conclusion whenever $D = n^{o(1)}$: if $\log D = o(\log n)$, then

$$\text{size}_{\Sigma_3}(k\text{-OV}_{n,D,k}) \leq n^{k+o(1)}.$$

More generally, if $k \rightarrow \infty$ and $\log D = o(k \log n)$, then the same circuit has size $n^{k+o(k)}$.

Proof. For a tuple $(i_1, \dots, i_k) \in [n]^k$, orthogonality is the conjunction, over all coordinates $t \in [D]$, of the clause asserting that not all selected vectors have a 1 in coordinate t . Thus

$$k\text{-OV}_{n,D,k} = \bigvee_{(i_1, \dots, i_k) \in [n]^k} \bigwedge_{t \in [D]} \bigvee_{\ell \in [k]} \neg u_{\ell, i_\ell}(t). \quad (192)$$

The formula has one top OR gate, one middle AND gate for each of the n^k tuples, and one bottom OR gate for each pair of a tuple and a coordinate $t \in [D]$, proving Equation (190). If $D = O(k \log n)$ and k is fixed, then $1 + n^k(D + 1) = n^{k+o(1)}$. More generally,

$$\log \text{size}_{\Sigma_3}(k\text{-OV}_{n,D,k}) \leq k \log n + O(\log(D + 1)),$$

so $\log D = o(\log n)$ gives the $n^{k+o(1)}$ estimate, while $\log D = o(k \log n)$ gives $n^{k+o(k)}$ when $k \rightarrow \infty$. $\square$

Proposition F.2 (Brute-force circuit for k -XOR). *For all n, m, k ,*

$$\text{size}_{\Sigma_3}(k\text{-XOR}_{n,m,k}) \leq 1 + \binom{n}{k} (1 + m2^{k-1}) = O\left(\binom{n}{k} m2^k\right). \quad (193)$$

If k is fixed, $n/k \rightarrow \infty$, and $\log m = o(\log(n/k))$, then

$$\text{size}_{\Sigma_3}(k\text{-XOR}_{n,m,k}) \leq (n/k)^{k+o(1)}. \quad (194)$$

More generally, if $k \rightarrow \infty$, $\log(n/k) \rightarrow \infty$, and $\log m = o(k \log(n/k))$, then

$$\text{size}_{\Sigma_3}(k\text{-XOR}_{n,m,k}) \leq (n/k)^{k+o(k)}.$$

Proof. Fix a support $S \in \binom{[n]}{k}$ and a row $r \in [m]$. The condition $\bigoplus_{j \in S} A_{r,j} = 0$ is an even-parity predicate on k bits and has a CNF with exactly 2^{k-1} clauses, one excluding each odd-parity assignment. Conjoining these clauses over all m rows computes that S is a zero-sum support, and taking the OR over all supports computes k -XOR. This gives an $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit with one top OR gate, one middle AND gate per support, and $m2^{k-1}$ bottom clause gates per support, proving Equation (193). Finally, $\binom{n}{k} \leq (en/k)^k$, so

$$\log \text{size}_{\Sigma_3}(k\text{-XOR}_{n,m,k}) \leq k \log(n/k) + O(k) + \log m = (k + o(1)) \log(n/k)$$

under the stated fixed- k hypotheses. In particular, this condition holds at the principal $O(k \log n)$ -scale row budgets used in Theorem 5.1. Without fixing k , the same estimate has error term $O(k) + \log m$; the growing- k hypotheses in the proposition make this $o(k \log(n/k))$, proving that bound. $\square$

Proposition F.3 (Elementary circuits for k -SUM). *For all $n \geq k \geq 2$ and $m \geq 1$,*

$$\text{size}_2(k\text{-SUM}_{n,m,k}) \leq 1 + \binom{n}{k} 2^{(k-1)m}. \quad (195)$$

Consequently,

$$\log \text{size}_2(k\text{-SUM}_{n,m,k}) \leq (k-1)m + k \log(en/k) + O(1), \quad (196)$$

and the displayed depth-two circuit has size $2^{\Theta(km)}$ whenever $m \geq \log(en/k)$. There is also a top-disjunction depth-three block-carry circuit: for every integer $b \in [m]$, with $t := \lceil m/b \rceil$,

$$\text{size}_{\Sigma_3}(k\text{-SUM}_{n,m,k}) \leq 1 + \binom{n}{k} k^{t-1} (1 + t2^{kb}). \quad (197)$$

In particular,

$$\log \text{size}_{\Sigma_3}(k\text{-SUM}_{n,m,k}) \leq k \log(en/k) + O(\sqrt{km \log k} + k + \log m). \quad (198)$$

If $n \geq k^2$ and $m = \Theta(k \log(en/k))$, then

$$\text{size}_{\Sigma_3}(k\text{-SUM}_{n,m,k}) \leq (n/k)^{O(k)}. \quad (199)$$

Proof. Fix a support $S \in \binom{[n]}{k}$. Choosing arbitrary values for any $k-1$ of its residues uniquely determines the remaining residue that makes the sum zero modulo 2^m . Hence exactly $2^{(k-1)m}$ assignments to the km selected input bits satisfy the support equation. A conjunction of literals recognizes each such assignment, so the support predicate has a DNF with $2^{(k-1)m}$ terms. Taking the OR of these terms over every support gives one depth-two DNF with the gate count in Equation (195). The binomial estimate $\binom{n}{k} \leq (en/k)^k$ gives Equation (196). When $m \geq \log(en/k)$, the upper bound is $O(km)$ in the logarithm, while the construction contains $2^{(k-1)m}$ term gates for every fixed support, proving the stated $2^{\Theta(km)}$ size.

For the block-carry circuit, partition the m bit positions, from least to most significant, into $t = \lceil m/b \rceil$ consecutive blocks of width at most b . For a fixed support, the carry into the first block is zero, and every later carry lies in $\{0, \dots, k-1\}$: the sum of k b -bit chunks and an incoming carry at most $k-1$ is at most $k2^b - 1$. Guess the $t-1$ intermediate carries. For each nonfinal block, require that the sum of its k chunks and its incoming carry equal 2^b times the guessed outgoing carry; for the final, possibly shorter block, require only that this sum vanish modulo the corresponding power of two. Every block predicate is a Boolean function of at most kb input bits and therefore has a CNF with at most 2^{kb} clauses, one excluding each falsifying assignment. For a fixed support and carry guess, flatten the conjunctions of all block CNFs into one middle AND gate; then take the OR over all supports and carry guesses. This gives an $\text{OR} \circ \text{AND} \circ \text{OR}$ circuit. The true arithmetic carries give a satisfying branch whenever the modular sum is zero, and conversely any satisfying branch certifies zero in every block, so the circuit is correct. Its gate count is bounded by Equation (197).

Taking logarithms in Equation (197), using $\binom{n}{k} \leq (en/k)^k$, gives

$$\log \text{size}_{\Sigma_3}(k\text{-SUM}_{n,m,k}) \leq k \log(en/k) + t \log k + kb + O(\log m). \quad (200)$$

Choose $b := \lceil \sqrt{m \log k/k} \rceil$. Since $k \geq 2$ and $m \geq 1$, this choice lies in $[m]$, and $t \leq m/b + 1$. Hence both $t \log k$ and kb are $O(\sqrt{km \log k} + k)$, proving Equation (198). Finally, if $n \geq k^2$ and $m = \Theta(k \log(en/k))$, then $\log k = O(\log(n/k))$, so every term on the right of Equation (198) is $O(k \log(n/k))$, which proves Equation (199). $\square$

Along the growing- k polynomial-host regimes of Theorem 6.8, the principal budgets make the preceding k -OV and k -XOR estimates and Equation (199) read respectively as $n^{k+o(k)}$, $(n/k)^{k+o(k)}$, and $(n/k)^{O(k)}$. Thus the same-orientation top-disjunction depth-three upper and lower bounds match at the linear-in- k exponent scale for all three targets in those regimes as well as for fixed k .

F.2 The structural constant floor for the expander route

The preceding circuits show that the natural brute-force exponent is k for k -OV in base n and for k -XOR in base n/k . The lower bounds of Theorem 5.1 have the same linear scale, but the particular κ -based expander route cannot match the leading constant.

Proposition F.4 (A constant-factor floor for regular-core padding). *Let $P = F \dot{\cup} J_t$, where F is a nonempty Δ -regular core with $\Delta \geq 2$, J_t is a disjoint matching, and $e(P) = k - 1$. If one idealizes the LRR exponent $\kappa(P) - o(1)$ to $\kappa(P)$ and then applies a quadratic host substitution $N^2 = \Theta(n)$ or $N^2 = \Theta(n/k)$, the resulting target-base exponent is at most $(k - 1)/\Delta$. Consequently, compared with the brute-force exponent k , the ratio between the upper exponent and this ceiling for the κ -based route is strictly greater than Δ .¹⁰*

Proof. By Theorem 2.4, $\kappa(P) \leq \text{tw}(P) + 1$. Treewidth is the maximum over connected components; moreover, $\text{tw}(J_t) \leq 1$, while the minimum-degree condition on F forces a cycle, whose treewidth is 2, and treewidth is minor-monotone. Hence $\text{tw}(F) \geq 2$, while trivially $\text{tw}(F) \leq v(F) - 1$, and therefore

$$\kappa(P) \leq \text{tw}(P) + 1 = \max\{\text{tw}(F), 1\} + 1 \leq v(F). \quad (201)$$

Since F is Δ -regular,

$$v(F) = \frac{2e(F)}{\Delta} \leq \frac{2(k - 1)}{\Delta}. \quad (202)$$

The idealized LRR exponent $\kappa(P)$ in base N becomes $\kappa(P)/2$ in the corresponding target base after either quadratic host substitution. Combining Equations (201) and (202) gives the ceiling $(k - 1)/\Delta$. The brute-force exponent k divided by this ceiling is $\Delta k/(k - 1) > \Delta$. $\square$

Remark F.5 (Scope of the floor). The proposition is a limitation of this specific κ -based regular-core-plus-matching route, not an upper bound on what other source families or other lower-bound methods could prove. For the degree-4 family used in Theorem 5.1, it already forces a ratio greater than 4, even before accounting for slack in the expansion estimate or in the LRR exponent.

F.3 The depth-sensitive comparison for k -SUM¹¹

In the nondegenerate high-width regime of Corollary D.4, the transferred lower exponent has scale $\min\{k \log(n/k), m\}$. By Proposition F.3, a top-disjunction depth-three upper exponent is controlled by

$$U_\Sigma(n, m, k) := k \log(en/k) + \sqrt{km \log k} + k + \log m. \quad (203)$$

At the principal width $m = \Theta(k \log(en/k))$ and under $n \geq k^2$, which holds throughout the polynomial-host regimes of Theorem 6.4, Theorem 6.8, and Theorem 7.3, one has $U_\Sigma(n, m, k) = \Theta(k \log(n/k))$. Thus, after specializing the lower bound to top-disjunction depth three, the lower and upper bounds match at exponent scale there.

Away from the principal width, two different gaps may remain. When $m \leq k \log(n/k)$, the lower exponent can be width-limited while the elementary circuit still enumerates all supports,

¹⁰This comparison at the principal bit-width were suggested and derived by the AI language model GPT-5.5; the authors independently verified the argument before inclusion.

¹¹The depth-sensitive comparison in this section, including the block-carry width choice underlying the exponent bound $U_\Sigma(n, m, k)$ of Equation (203) and the resulting ratio estimate of Equation (204), was suggested and derived by the AI language model GPT-5.5; the authors independently verified the argument before inclusion.

contributing $k \log(en/k)$. When $m \geq k \log(n/k)$, the ratio between the displayed depth-three upper scale and the saturated lower scale satisfies

$$\frac{U_\Sigma(n, m, k)}{k \log(n/k)} = O\left(1 + \frac{\sqrt{m \log k/k}}{\log(n/k)} + \frac{k + \log m}{k \log(n/k)}\right). \quad (204)$$

The carry contribution in (204) becomes superconstant only beyond $m = \omega(k \log^2(n/k)/\log k)$. The resulting overhead is much smaller than the $\Theta(m/\log(n/k))$ gap suggested by the depth-two DNF, but it need not be constant at arbitrarily large width.

The block-carry improvement is depth-sensitive. The direct depth-two circuit resolves complete assignments to a support, whereas the top-disjunction depth-three circuit guesses intermediate block carries and verifies the blocks in parallel. No fixed depth yields a polynomial-size verifier for the fixed-support predicate uniformly in k : by Proposition D.6, restricting each selected residue to $\{0, 2^{m-1}\}$ turns that predicate into the complement of PARITY, so every fixed depth still requires size $\exp(\Omega_d(k^{1/(d-1)}))$. The matching restriction in Equations (161) and (162) shows that this parity floor is shared by k -XOR. Hence the PARITY obstruction itself does not distinguish the two list targets. The quantitative distinction is in the available upper bounds: the fixed-support k -XOR predicate has an $m 2^{k-1}$ -clause CNF, the direct fixed-support k -SUM predicate has a $2^{(k-1)m}$ -term DNF, and the block-carry Σ_3 construction replaces the latter dependence by the square-root carry overhead.

G Depth-Three Probability Estimates

This appendix proves the two probabilistic ingredients used in Section 6.3: the single-CNF minterm bound and the planted-copy good event. The common bipartite degree-4 expander core is constructed in Corollary B.4.

G.1 The single-CNF minterm bound

We use the notation of Theorem 6.7. Thus F has a nonempty independent set I of size L , every vertex of I has positive degree, T is a CNF with $s = s(T) \geq 1$ clauses, and the random experiment produces the background G , the planted labels ϕ^* , the blocks $B_{v,a}$, and the partial inputs H_A .

Fix the labels $\phi^*(u)$ for $u \notin I$. For a clause C of T and a state $A \subseteq I$, define its remaining positive support by

$$\begin{aligned} R_A(C) &:= \{(v, a) \in (I \setminus A) \times [N] : \\ &\quad C \text{ contains a positive literal whose variable lies in } B_{v,a}\}, \\ r_A(C) &:= |R_A(C)|. \end{aligned} \quad (205)$$

For each $(v, a) \in R_A(C)$, fix one canonical representative positive literal $x_{C,v,a} \in C \cap B_{v,a}$. The representatives are distinct because the blocks are pairwise disjoint. Call a pair (A, C) , with $A \subseteq I$ and C a clause of T , eligible when $x_{C,v,a} \notin G$ for every $(v, a) \in R_A(C)$, and define

$$Z := \max\{r_A(C) : (A, C) \text{ eligible}\}. \quad (206)$$

The eligible set is nonempty because $R_I(C) = \emptyset$ for every clause C . Once the outside labels are fixed, Z depends only on the product-random background G and is independent of the labels inside I .

Lemma G.1 (Support-maximum tail and moment). *With the outside labels fixed, let $\Pr$ and $\mathbb{E}$ be taken over G alone. For every real $t \geq 0$,*

$$\Pr[Z \geq t] \leq \min\{1, 2^L s e^{-pt}\}. \quad (207)$$

Consequently, for an absolute constant $C_0 > 0$,

$$\mathbb{E}[Z^L] \leq \left(\frac{C_0(L + \log(es))}{p} \right)^L. \quad (208)$$

Both estimates are uniform in the fixed outside labels.

Proof. Fix $A \subseteq I$ and a clause C with $r_A(C) = r$. Its r canonical representatives are distinct variables, each absent from G independently with probability $1 - p$, so the probability that all of them are absent is $(1 - p)^r \leq e^{-pr}$. A union bound over at most $2^L s$ pairs (A, C) proves Equation (207).

Put $\alpha := \ln(2^L s)$. Taking $t = (\alpha + \lambda)/p$ in Equation (207) gives

$$\Pr[pZ \geq \alpha + \lambda] \leq 2^L s e^{-(\alpha + \lambda)} = e^{-\lambda} \quad \text{for every } \lambda \geq 0. \quad (209)$$

Let Y be an exponential random variable of mean 1, so that $\Pr[\alpha + Y \geq \alpha + \lambda] = e^{-\lambda}$ for every $\lambda \geq 0$, while $\Pr[\alpha + Y \geq t] = 1$ for every $t < \alpha$. Comparing the first identity with Equation (209) and the second with the trivial bound $\Pr[pZ \geq t] \leq 1$ shows that pZ is stochastically dominated by $\alpha + Y$. Since $x \mapsto x^L$ is increasing on $[0, \infty)$, this domination gives $\mathbb{E}[(pZ)^L] \leq \mathbb{E}[(\alpha + Y)^L]$, and Minkowski's inequality for the norm $X \mapsto (\mathbb{E}[X^L])^{1/L}$, available because $L \geq 1$, gives

$$(\mathbb{E}[(pZ)^L])^{1/L} \leq \alpha + (\mathbb{E}[Y^L])^{1/L} = \alpha + (L!)^{1/L} \leq \alpha + L.$$

Since $\alpha + L = L(1 + \ln 2) + \ln s = O(L + \log(es))$, enlarging an absolute constant and dividing by p^L proves Equation (208). The argument is uniform in the outside labels. $\square$

Proof of Theorem 6.7. The order of conditioning is essential. First fix only the labels outside I and take the background moment in Equation (208). Then condition further on G , which fixes Z , and expose the mutually independent uniform labels $\{\phi^*(v) : v \in I\}$ adaptively. All probabilities in this proof are over the background, the planted labels, and the process's own uniform vertex choices, the last being independent of the other two; the conditioning in each display records which of them have been fixed.

Consider the following randomized commitment process. It starts at $A = \emptyset$. At a state $A \subseteq I$, it succeeds if $A = I$ and fails if $T(H_A) = 1$. In the remaining case, where $A \subsetneq I$ and $T(H_A) = 0$, call A proper and let C_A be the first clause, in a fixed ordering, that is false on H_A ; such a clause exists because $T(H_A) = 0$. The process chooses $v \in I \setminus A$ uniformly, reveals $\phi^*(v)$, and advances to $A \cup \{v\}$ exactly when C_A contains a positive literal whose variable lies in $B_{v, \phi^*(v)}$, that is, when $(v, \phi^*(v)) \in R_A(C_A)$; it fails otherwise.

Fix a proper state A and put $q := |I \setminus A|$. For each $v \in I \setminus A$, let $s_v := |\{a : (v, a) \in R_A(C_A)\}|$. Conditional on the preceding history, the next vertex is uniform among the q choices and its unrevealed label is uniform on $[N]$, so

$$\Pr[\text{advance from } A \mid \text{history}] = \frac{1}{q} \sum_{v \in I \setminus A} \frac{s_v}{N} = \frac{r_A(C_A)}{qN}. \quad (210)$$

The pair (A, C_A) is eligible for Equation (206). Indeed, C_A is false on H_A , so every positive literal of C_A has its variable outside H_A and hence outside $G \subseteq H_A$; in particular every canonical representative is absent from G . For a representative indexed by some (v, a) with $v \in I \setminus A$, the two absences are in fact equivalent: it lies in no planted block already added at A , the blocks being pairwise disjoint and $v \notin A$, and in no edge of F_I^* , which avoids I entirely, so its only possible source in H_A is G . Thus $r_A(C_A) \leq Z$. Since this bound holds uniformly over histories and Z does not depend on the labels inside I , multiplying the L conditional estimates $Z/(qN)$, one for each $q = L, L-1, \dots, 1$, yields

$$\Pr[\text{the process succeeds} \mid G, \phi^*|_{V(F) \setminus I}] \leq \frac{Z^L}{L! N^L}. \quad (211)$$

Now fix all planted labels and suppose that $\mathcal{M}_I(T)$ occurs. At every proper state A , the selected clause C_A is false on H_A but true on H_I . Passing from H_A to H_I only adds variables from the remaining planted blocks, so some remaining block contains a positive literal that repairs C_A . The uniform vertex choice therefore advances with probability at least $1/q$, and

$$\Pr[\text{the process succeeds} \mid G, \phi^*] \geq \frac{1}{L!} \quad \text{on } \mathcal{M}_I(T). \quad (212)$$

Since $\mathcal{M}_I(T)$ is determined by G and ϕ^* , Equation (212) says that $\Pr[\text{the process succeeds} \mid G, \phi^*] \geq \mathbf{1}[\mathcal{M}_I(T)]/L!$ pointwise; averaging over the labels inside I turns the right-hand side into $\Pr[\mathcal{M}_I(T) \mid G, \phi^*|_{V(F) \setminus I}]/L!$, and comparing with Equation (211) cancels the factorials and gives

$$\Pr[\mathcal{M}_I(T) \mid G, \phi^*|_{V(F) \setminus I}] \leq \frac{Z^L}{N^L}.$$

Averaging first over G and then over the outside labels, and applying Equation (208), proves Equation (76).

Finally, for every $A \subsetneq I$, the set

$$F_I^* \cup \bigcup_{v \in A} B_{v, \phi^*(v)}$$

is a proper subset of F^* because a missing vertex of I has positive degree and hence a nonempty planted block, and the planted blocks are disjoint. Hence $\mathcal{M}(T) \subseteq \mathcal{M}_I(T)$, which proves the full-minterm bound. $\square$

G.2 The planted-copy good event

The following statement is slightly more general than the 4-regular specialization used in Theorem 6.8. Its proof explains the threshold $\theta > 2/\Delta$ and the polynomial relation between N and k .

Lemma G.2 (Subthreshold planted input). *Fix an integer $\Delta \geq 3$ and a constant $\theta > 2/\Delta$. Let $F = F_k$ be a Δ -regular graph with $v(F) = O(k)$, and choose the background density*

$$p := N^{-\theta}. \quad (213)$$

There is a constant $C_{\Delta, \theta} > 0$ such that, uniformly for $N \geq k^{C_{\Delta, \theta}}$, the following event has probability $1 - o(1)$ as $k \rightarrow \infty$: no variable of the planted copy F^ already belongs to G , and F^* is the unique F -copy in $G \cup F^*$, where an F -copy means an accepting map $\psi : V(F) \rightarrow [N]$ in the sense of Definition 2.1. On this event, $G \cup R$ is a 0-input of $F\text{-SUB}_N$ for every proper subset $R \subsetneq F^*$.*

Proof. The planted copy is present in $G \cup F^*$ by construction. Since $e(F) = \Delta v(F)/2 = O(k)$,

$$\Pr[F^* \cap G \neq \emptyset] \leq e(F)p = o(1) \quad (214)$$

uniformly once $C_{\Delta, \theta}$ is sufficiently large.

It remains to rule out an alternative accepting map $\psi : V(F) \rightarrow [N]$. Let $J := \{v : \psi(v) \neq \phi^*(v)\}$ and $j := |J| \geq 1$. For each fixed J , there are at most N^j assignments to the labels on J ; label collisions are allowed and harmless. Every pattern edge meeting J is represented under ψ by a non-planted variable, these variables are distinct across pattern edges, and all must be supplied by G .

If $e_F(J)$ is the number of edges with both endpoints in J and $e_F(J, V(F) \setminus J)$ is the number with exactly one, then counting the Δj edge-endpoints at J gives

$$\Delta j = 2e_F(J) + e_F(J, V(F) \setminus J),$$

so the number $e_F(J) + e_F(J, V(F) \setminus J)$ of edges meeting J is at least $\Delta j/2$. Therefore

$$\begin{aligned} \mathbb{E}[\#\{\psi \neq \phi^* : \psi \text{ is a copy in } G \cup F^*\}] &\leq \sum_{j=1}^{v(F)} \binom{v(F)}{j} N^j p^{\Delta j/2} \\ &\leq \sum_{j \geq 1} (v(F) N^{-(\theta \Delta/2 - 1)})^j = o(1). \end{aligned} \quad (215)$$

The estimate is uniform for $N \geq k^{C_{\Delta, \theta}}$ after choosing $C_{\Delta, \theta}(\theta \Delta/2 - 1) > 2$: the geometric ratio $v(F)N^{-(\theta \Delta/2 - 1)}$ decreases in N and is $O(k^{1 - C_{\Delta, \theta}(\theta \Delta/2 - 1)}) = O(k^{-1})$ at $N = k^{C_{\Delta, \theta}}$; increasing the same constant so that $C_{\Delta, \theta}\theta > 2$ also makes Equation (214) uniform. Markov's inequality together with Equation (214) proves that the asserted good event has probability $1 - o(1)$.

Suppose the good event holds, and let $R \subsetneq F^*$ be arbitrary. If $G \cup R$ contained an F -copy, then, since $G \cup R \subseteq G \cup F^*$, the same copy would lie in $G \cup F^*$, and uniqueness would force it to be F^* ; in particular every variable of F^* would lie in $G \cup R$. But a variable of $F^* \setminus R$ lies outside R by choice and outside G because no planted variable belongs to G on the good event, a contradiction. $\square$