Electronic Colloquium on Computational Complexity, Report No. 158 (2026)

RANDOM 3-CNF FORMULAS ARE HARD FOR k-DNF RESOLUTION UP TO $k = O(\sqrt{\log n})$

GAIA CARENINI

ABSTRACT. We prove exponential lower bounds for k-DNF resolution on random 3-CNF formulas throughout the range $k = O(\sqrt{\log n})$ at every constant clause density above the elementary first-moment bound for unsatisfiability. For random 3-CNFs this improves Alekhnovich's range $k = O(\sqrt{\log n/\log\log n})$, and matches the $O(\sqrt{\log n})$ range obtained by Sofronova and Sokolov for random CNFs of sufficiently large constant width. We also obtain higher-density tradeoffs. In particular, random 3-CNFs with $n \log^h n$ clauses are exponentially hard for $k = O(\sqrt{\log n/\log\log n})$ for every fixed $h > 0$, while for every fixed K the same conclusion holds simultaneously for all $1 \le k \le K$ with $n^{1+\varepsilon}$ clauses whenever $\varepsilon < 1/(4K^2 + 2)$.

1. INTRODUCTION

Proof complexity asks how difficult it is to certify that a Boolean formula is unsatisfiable. The most basic proof system in the subject is Resolution: one starts from the clauses of a CNF formula and repeatedly resolves on variables until the empty clause is obtained. Its definition is elementary, but the known lower-bound arguments already involve several of the ideas that recur throughout proof complexity, among them expansion, restrictions, and the passage from local consistency to global contradiction; see, for example, [8, 15, 4, 11].

A natural strengthening of Resolution allows a proof line to be a disjunction of small conjunctions. In k-DNF resolution, denoted $\mathrm{Res}(k)$, every line is a disjunction of terms, each involving at most k literals. Thus $\mathrm{Res}(1)$ is essentially ordinary Resolution, while larger values of k allow increasingly complicated pieces of local information to be manipulated in a single inference. A basic question is therefore how rapidly Resolution lower bounds deteriorate as k grows.

Random CNF formulas are a particularly natural family on which to ask this question. They have long served as canonical hard instances for proof systems [5, 2, 4, 3, 6, 9, 14, 10, 7], and for several systems their hardness can be read off from suitable expansion properties. For $\mathrm{Res}(k)$, however, expansion is not by itself enough. A single proof line may contain a very large number of terms, and a restriction that simplifies one of them need not simplify the others. The main difficulty is therefore to find a restriction procedure that makes simultaneous progress on all the relevant terms while preserving enough local consistency to support a final lower-bound argument.

1.1. **Previous results and the width-three difficulty.**

The classical approach uses small random restrictions and switching arguments. Segerlind, Buss, and Impagliazzo [12] proved a switching lemma suited to this setting. Alekhnovich [1] adapted the method to random 3-CNFs and obtained exponential lower bounds in the range $k = O\left(\sqrt{\log n/\log\log n}\right)$.

More recently, Sofronova and Sokolov [13] developed a direct expansion-and-restriction argument. Their method reaches the larger range $k = O(\sqrt{\log n})$ at linear density, but requires the clause width to be a sufficiently large constant.

It is useful to explain the shape of their proof, since most of it will also be the shape of ours. Rather than trying to simplify an entire k-DNF in one step, one represents each proof line by a small decision tree whose leaves contain the residual DNFs. Different branches

Date: August 22, 2026.

ISSN 1433-8092

are allowed to query different variables, so information is exposed only where it is needed. A branch that is already locally inconsistent can be discarded. A branch containing many suitably independent terms is stored for later elimination by one global random restriction. Every remaining branch is queried further in such a way that a complexity measure decreases. Once all branches have become sufficiently simple, one moves backwards through the proof DAG and eventually contradicts local consistency of an initial clause.

The delicate part is the dichotomy used during the revision of a branch. Very roughly, Sofronova and Sokolov attach a small closure to each residual term. If many closures are needed to cover the branch, they extract many closure-independent terms and eliminate the branch probabilistically. If only a few closures are needed, they query their union and thereby make deterministic progress. Thus the mechanism is a dichotomy between many independent pieces and a small place on which to make progress.

At width three, the required independence statement is no longer available. The problem is not merely that two terms may share variables. Even terms with disjoint supports can be correlated under a locally satisfying restriction, because sparse parity constraints may create a hidden affine dependency between their supports. The obstruction is therefore to the independence lemma of [13, Lemma 15], rather than to the surrounding DNF-tree architecture.

The relevant parameter ranges are summarized below.

Result	Clause width	Clauses	Range of k
Segerlind–Buss–Impagliazzo [12]	$\Delta \gtrsim k^2$	$m \leq n^{7/6}$	$k = O(1)$
Alekhnovich [1]	$\Delta = 3$	$m = O(n)$	$k = O\Big(\sqrt{\log n/\log\log n}\Big)$
Sofronova–Sokolov [13]	$\Delta = O(1)$	$m = O(n)$	$k = O(\sqrt{\log n})$
This paper	$\Delta = 3$	$m = \alpha n$	$k = O(\sqrt{\log n})$
This paper	$\Delta = 3$	$m = n \log^h n$	$k = O\Big(\sqrt{\log n/\log\log n}\Big)$
This paper	$\Delta = 3$	$m = n^{1+\varepsilon}$	$1 \leq k \leq K$, $\varepsilon < 1/(4K^2+2)$

1.2. **Main results.** Let $\Phi_{n,m}$ be the random 3-CNF obtained by choosing m clauses independently and uniformly with replacement from all clauses on three distinct variables. Our main theorem is the following.

Theorem 1.1. *For every fixed $\alpha > \log 2/\log(8/7)$, there are constants $\delta, \gamma > 0$ such that, with probability tending to one as $n \to \infty$, the random formula $\Phi_{n,\lceil \alpha n\rceil}$ has no $\mathrm{Res}(k)$ refutation with fewer than 2^{n^δ} lines, simultaneously for every integer $1 \leq k \leq \gamma\sqrt{\log n}$.*

The density assumption is simply the elementary first-moment condition under which the expected number of satisfying assignments tends to zero. Thus, throughout the entire range where this elementary argument already proves unsatisfiability, we obtain the same $O(\sqrt{\log n})$ range previously known at sufficiently large constant width.

The same method gives the following tradeoffs when the density grows with n.

Theorem 1.2. *The following statements hold.*

(i) *For every fixed $h > 0$, there are constants $\delta, \gamma > 0$ such that, with probability tending to one, the random formula $\Phi_{n,\lceil n \log^h n\rceil}$ has no $\mathrm{Res}(k)$ refutation with fewer than 2^{n^δ} lines for every integer $1 \leq k \leq \gamma\sqrt{\log n/\log\log n}$.*

(ii) *Let $K \geq 1$ be fixed and let $0 < \varepsilon < 1/4K^2 + 2$. Then there is a constant $\delta > 0$ such that, with probability tending to one, the random formula $\Phi_{n,\lceil n^{1+\varepsilon}\rceil}$ has no $\mathrm{Res}(k)$ refutation with fewer than 2^{n^δ} lines, simultaneously for every integer $1 \leq k \leq K$.*

1.3. **The idea of the proof.** The proof has a deterministic part and a probabilistic part. Given a 3-CNF Φ, we associate to each clause the unique parity equation on the same three variables that is violated by the assignment falsifying that clause. This produces a system A_Φ of 3-variable equations over $\mathbb{F}_2$. We first prove an exponential $\mathrm{Res}(k)$ lower bound for

$\mathrm{CNF}(A_\Phi)$ from a suitable boundary-expansion hypothesis. We then show that the incidence graph of a random formula has the required expansion up to a scale determined by its density.

The new ingredient lies entirely in the deterministic part. Suppose that during the restriction process we have already exposed a small set of variables. The parity equations may then force additional relations. Ignoring these relations would create spurious independence, so we keep all of the locally visible affine information.

The basic object is a *closed state* (B, σ). Here B is a small seed set of variables. From B we generate a boundary closure by adjoining every small subsystem whose boundary is contained in B, together with all variables appearing in those equations. Informally, this closure contains every parity dependency that can already be detected from a small subsystem using only boundary variables in B. The assignment σ is defined on the entire variable closure and satisfies all equations in it. Boundary expansion keeps the closure small, makes it monotone in B, and, most importantly, guarantees that every locally satisfying assignment on a smaller closure extends to every larger closure.

Now let S be a set of variables outside the current closure. We define the *local affine space* $\mathcal{V}_{B,\sigma}(S)$ to be the set of assignments to S that can still be extended, together with σ, to the appropriate enlarged closure. This is a nonempty affine subspace of $\mathbb{F}_2^S$. Its dimension $d_{B,\sigma}(S) = \dim \mathcal{V}_{B,\sigma}(S)$ measures how many genuinely free affine degrees of freedom remain on S.

This dimension is exactly what is needed to repair the independence step. Two disjoint supports need not be independent, but if pairwise disjoint sets $S_1, \ldots, S_h$ satisfy $d_{B,\sigma}(\bigcup_{i=1}^h S_i) = \sum_{i=1}^h d_{B,\sigma}(S_i)$, then there is no hidden affine relation coupling the sets: the joint local space is the direct product of the individual local spaces.

For a compatible residual term t with support S, we therefore use the integer potential $\mu_{B,\sigma}(t) = |S| + d_{B,\sigma}(S)$. The first summand records the number of variables that remain unassigned. The second records the number of affine degrees of freedom that remain on those variables. This distinction matters: a query may simplify a term either by assigning one of its variables or by revealing a new affine relation among variables that remain unassigned. The potential notices both kinds of progress. Since both summands are at most the width, a nonconstant term of width at most k has potential between 1 and $2k$.

The central revision lemma is then a clean *product-or-progress* dichotomy. Given a family of compatible residual terms, we greedily select terms with disjoint supports while preserving additivity of affine dimension. If we can select h of them, they form a product family. Under a random locally satisfying restriction, each selected support is exposed with probability at least p^k, and conditional on being exposed the corresponding terms behave independently in the local affine distribution. This gives

$$\Pr[\text{the branch survives and no stored term becomes true}] \leq \exp\bigl(-h(p/2)^k\bigr).$$

If the greedy process stops before h selections, the union U of the selected supports has size less than kh. Querying the closure generated by U forces every compatible surviving term to lose either a variable or an affine degree of freedom. Hence its potential decreases strictly.

After at most $2k$ progress rounds, no compatible nonconstant term can remain on an active branch. Branches on which a large product family was found are all eliminated simultaneously by one global restriction. At that point we return to the Sofronova–Sokolov architecture: one searches backwards through the proof DAG, extending a single temporary locally satisfying assignment while traversing the DNF-trees of the premises. Soundness of the $\mathrm{Res}(k)$ rules moves the search to a premise. Eventually one reaches an initial clause, where the local parity equation forces the clause to be true while the chosen branch forces it to remain false.

1.4. **Organization.** In Section 2 we recall $\mathrm{Res}(k)$, boundary expansion, and the parity encoding, and state the deterministic theorem from which all the random lower bounds follow. In Section 3 we develop the boundary closure and the associated local affine spaces. The product-or-progress revision principle and the restriction estimate for product families are proved in Section 4. In Section 5 we combine these ingredients with the DNF-tree argument and the backward search through the proof DAG. Finally, Section 6 establishes the required boundary expansion for random 3-CNFs and derives the main theorems.

Acknowledgments. This problem was suggested to the author by Dmitry Sokolov during a visit to Susanna F. de Rezende's research group in the Department of Computer Science at Lund University in September 2025. The author is grateful to him for explaining the ideas underlying his joint work with Sofronova on expansion-based lower bounds for k-DNF resolution, and to Susanna F. de Rezende for many valuable discussions and comments on an earlier version of this manuscript. This work was supported by the CB European PhD Studentship funded by Trinity College, Cambridge.

Use of AI tools. AI-assisted tools were used solely to improve the wording, organization, and clarity of the exposition. They were not used to generate the mathematical arguments.

2. Preliminaries and the deterministic theorem

This section fixes the notation used later. There are only three ingredients: the proof system Res(k), boundary expansion for systems of parity equations, and the CNF encoding that connects the two.

2.1. k-DNF resolution.

A *partial assignment* is a map $\sigma : \mathrm{dom}(\sigma) \to \{0,1\}$, where $\mathrm{dom}(\sigma)$ is a set of variables. If $\mathrm{dom}(\sigma) \subseteq \mathrm{dom}(\sigma')$ and the two assignments agree on $\mathrm{dom}(\sigma)$, we say that σ' *extends* σ. Two partial assignments are *consistent* if they agree on the intersection of their domains; in this case their union is again a partial assignment.

A *term* is a conjunction of literals. We write $\mathrm{vars}(t)$ for the set of variables occurring in a term t, and call $|\mathrm{vars}(t)|$ its *width*. A k-DNF is a disjunction of terms of width at most k. Repeated literals are removed and contradictory terms are discarded throughout.

If t is a term and σ is a partial assignment, the restriction $t|_\sigma$ is defined in the usual way. If σ falsifies one of the literals of t, then $t|_\sigma = 0$. Otherwise we delete all literals satisfied by σ; if none remain, then $t|_\sigma = 1$, and otherwise $t|_\sigma$ is the resulting nonconstant term. For a DNF D, restriction is performed term by term. If some restricted term is 1, then $D|_\sigma = 1$; otherwise all zero terms are deleted, with $D|_\sigma = 0$ if no term remains. The nonconstant terms that remain will be called the *residual terms* of D under σ.

The system Res(k) has the following inference rules, where F and G are k-DNFs and $w \le k$:

$$\text{weakening:} \qquad \frac{F}{(F \vee \ell)},$$

$$\text{and-introduction:} \qquad \frac{(F \vee \ell_1), \dots, (F \vee \ell_w)}{(F \vee \bigwedge_{i=1}^{w} \ell_i)},$$

$$\text{and-elimination:} \qquad \frac{(F \vee \bigwedge_{i=1}^{w} \ell_i)}{(F \vee \ell_j)},$$

$$\text{cut:} \qquad \frac{(F \vee \bigwedge_{i=1}^{w} \ell_i), (G \vee \bigvee_{i=1}^{w} \neg \ell_i)}{(F \vee G)}.$$

A Res(k) refutation is a sequence of k-DNFs ending with the empty DNF, where each line is either an initial clause or follows from earlier lines by one of these rules. Its *length* is the number of lines. We refer to [11] for general background.

2.2. Boundary expansion and parity encodings.

Let A be a system of linear equations over $\mathbb{F}_2$, each involving exactly three distinct variables. Its incidence graph is the bipartite graph with an equation on the left, a variable on the right, and an edge for incidence. If Q is a set of equations, its *boundary* ∂Q is the set of variables that occur in exactly one equation of Q.

We say that the incidence graph is an $(r, 3, c)$-*boundary expander* if every equation has degree three and every nonempty set Q of at most r equations satisfies $|\partial Q| \geq c|Q|$. The point of using the boundary rather than the ordinary neighbourhood is that boundary expansion immediately gives linear independence. Indeed, every nonempty set of at most r rows has a variable that occurs in exactly one of them. That variable survives in their sum, so the rows are linearly independent. Consequently every subsystem of at most r equations is satisfiable, irrespective of the right-hand sides.

For an equation $x + y + z = b$, its standard CNF encoding consists of the four width-three clauses excluding the four assignments that violate the equation. We write $\mathrm{CNF}(A)$ for the conjunction of the encodings of all equations in A.

2.3. **The deterministic theorem.** All probabilistic applications will be deduced from the following statement. The somewhat ungainly parameters are chosen so that the proof later becomes transparent: p is the density of the global random seed, q is the probability scale at which a width-k term is both exposed and satisfied, and R is the growth factor between successive revision thresholds.

Theorem 2.1. *Fix $c > 0$. Let A be an unsatisfiable system of equations over $\mathbb{F}_2$ on n variables, each equation involving exactly three distinct variables, and suppose that its incidence graph is an $(r, 3, c)$-boundary expander. Let*

$$\lambda = 1 + \frac{3}{c}, \qquad p = \frac{cr}{32n}, \qquad q = \left(\frac{p}{2}\right)^k, \qquad R = \frac{64\lambda k}{q}.$$

If, for some fixed $\delta > 0$,

$$4\lambda k n^\delta R^{2k} = o(r/k), \tag{2.1}$$

then every $\mathrm{Res}(k)$ refutation of $\mathrm{CNF}(A)$ has length at least 2^{n^δ}.

Here and below, asymptotic notation refers to a family indexed by n, with c fixed and with $r = r(n)$ and $k = k(n)$. The proof of Theorem 2.1 occupies Sections 3 to 5. The only role of (2.1) is to guarantee that every closure we construct remains far below the expansion scale r, while the product witnesses are large enough to union bound over a proof of length 2^{n^δ}.

3. Boundary closure and local affine geometry

The restriction distribution used later is supported on assignments satisfying small parity subsystems. Thus it is not enough to know which variables have been exposed: we must also know which affine relations among the remaining variables have become forced. This section packages exactly that information.

We regard each equation of A as a row vector over $\mathbb{F}_2$, supported on the three variables it contains. If Q is a set of equations, the *sum* of Q is the equation obtained by adding their rows over $\mathbb{F}_2$. An assignment satisfying every equation of Q necessarily satisfies this sum equation, which we call a *linear consequence* of Q.

A set B of variables will be called *admissible* if $|B| < cr/4$. For an admissible B, let $\mathcal{C}(B)$ be the family of all sets Q of equations such that $|Q| \leq \frac{r}{2}$ and $\partial Q \subseteq B$. Thus $\mathcal{C}(B)$ contains the small subsystems whose only uniquely occurring variables lie in B. The next proposition says that these subsystems have a well-defined maximal union, and that the solution spaces obtained by enlarging B behave exactly as one would hope.

Proposition 3.1 (Closure and local projections)**.** *For every admissible B, the family $\mathcal{C}(B)$ has a largest member K_B, satisfying $|K_B| \leq |B|/c$, and if $B \subseteq C$ are admissible then $K_B \subseteq K_C$. Let $N(K_B)$ be the set of variables occurring in some equation of K_B, and let $X_B = B \cup N(K_B)$. We call K_B the* equation closure *of B and X_B its* variable closure*. Then $|X_B| \leq \lambda|B|$, where $\lambda = 1 + 3/c$, and the following hold.*

(i) *If $B \subseteq C$ are admissible, every assignment on X_B satisfying K_B extends to an assignment on X_C satisfying K_C.*

(ii) *Fix an assignment σ on X_B satisfying K_B. If U is disjoint from X_B and $B \cup U$ is admissible, let $\mathcal{V}_{B,\sigma}(U)$ be the set of assignments on U that extend, together with σ, to a solution of $K_{B\cup U}$. Then $\mathcal{V}_{B,\sigma}(U)$ is a nonempty affine subspace of $\mathbb{F}_2^U$. Moreover, whenever $U \subseteq W$ and $B \cup W$ is admissible, the coordinate projection of $\mathcal{V}_{B,\sigma}(W)$ onto U is exactly $\mathcal{V}_{B,\sigma}(U)$.*

(iii) *Let $d_{B,\sigma}(U) = \dim \mathcal{V}_{B,\sigma}(U)$. For fixed admissible $B \cup W$ with $W \cap X_B = \varnothing$, the function $U \mapsto d_{B,\sigma}(U)$ on subsets of W is monotone and submodular. Furthermore, if pairwise disjoint sets $S_1, \dots, S_h \subseteq W$ satisfy*

$$d_{B,\sigma}\Big(\bigcup_{i=1}^{h} S_i\Big) = \sum_{i=1}^{h} d_{B,\sigma}(S_i),$$

then

$$\mathcal{V}_{B,\sigma}\Big(\bigcup_{i=1}^{h} S_i\Big) = \prod_{i=1}^{h} \mathcal{V}_{B,\sigma}(S_i).$$

Proof. We prove the statements in the order in which they will be used.

The closure. Suppose $Q_1, Q_2 \in \mathcal{C}(B)$. If a variable x has degree one in $Q_1 \cup Q_2$, then the unique equation of $Q_1 \cup Q_2$ containing x belongs to at least one of Q_1, Q_2, and in that set x also has degree one. Hence $\partial(Q_1 \cup Q_2) \subseteq \partial Q_1 \cup \partial Q_2 \subseteq B$. Boundary expansion gives $|Q_i| \le |\partial Q_i|/c \le |B|/c < r/4$, so $|Q_1 \cup Q_2| < r/2$. Therefore $Q_1 \cup Q_2 \in \mathcal{C}(B)$. The family $\mathcal{C}(B)$ is finite, nonempty, and closed under unions, and consequently has a unique largest member, which we call K_B.

Applying boundary expansion once more, $c|K_B| \le |\partial K_B| \le |B|$, so $|K_B| \le |B|/c$. If $B \subseteq C$, then $\mathcal{C}(B) \subseteq \mathcal{C}(C)$, and therefore $K_B \subseteq K_C$. Finally, $|X_B| \le |B| + 3|K_B| \le (1 + 3/c)|B| = \lambda|B|$.

Extension of local solutions. Fix admissible $B \subseteq C$. The only possible obstruction to extending a solution on X_B to one on X_C would be a linear consequence of K_C supported entirely on X_B that is not already forced by K_B. We show that this cannot happen.

Suppose $Q \subseteq K_C$ and the sum of the equations in Q is supported inside X_B. Then $\partial Q \subseteq X_B$, since every degree-one variable of Q necessarily occurs in the sum. Put $K = K_B$. We claim that $\partial(K \cup Q) \subseteq B$. Indeed, a variable in $X_B \setminus B = N(K) \setminus B$ has degree at least two in K, because $\partial K \subseteq B$, and so cannot lie in the boundary of $K \cup Q$. A variable outside X_B does not occur in K; if it had degree one in $K \cup Q$, it would therefore have degree one in Q, contradicting $\partial Q \subseteq X_B$.

Since $K \cup Q \subseteq K_C$ and $|K_C| \le r/2$, we have $K \cup Q \in \mathcal{C}(B)$. By maximality of K_B, this forces $Q \subseteq K_B$. Thus every linear consequence of K_C supported on X_B is already a consequence of K_B.

Now let σ be a solution of K_B on X_B. If it had no extension to a solution of K_C on X_C, Gaussian elimination would produce a linear combination of equations of K_C supported on X_B and violated by σ. The previous paragraph says that this combination is already a consequence of K_B, a contradiction. This proves (i).

The local affine spaces. Fix B, σ, and U as in (ii). By part (i), σ extends to a solution of $K_{B\cup U}$, so $\mathcal{V}_{B,\sigma}(U)$ is nonempty. Moreover, the set

$$\Theta = \{\theta \in \mathbb{F}_2^{X_{B\cup U}} : \theta \text{ satisfies } K_{B\cup U},\ \theta|_{X_B} = \sigma\}$$

is a nonempty affine subspace. Its coordinate projection onto U is precisely $\mathcal{V}_{B,\sigma}(U)$, proving that the latter is affine.

Now let $U \subseteq W$. If $\beta \in \mathcal{V}_{B,\sigma}(W)$, any witnessing solution of $K_{B\cup W}$ restricts to a witnessing solution for $\beta|_U$, so $\pi_U(\mathcal{V}_{B,\sigma}(W)) \subseteq \mathcal{V}_{B,\sigma}(U)$. Conversely, if $\alpha \in \mathcal{V}_{B,\sigma}(U)$, choose a witnessing solution on $X_{B\cup U}$ and extend it to $X_{B\cup W}$ using part (i). Restricting the extension to W gives a point of $\mathcal{V}_{B,\sigma}(W)$ projecting to α. Hence

$$\pi_U(\mathcal{V}_{B,\sigma}(W)) = \mathcal{V}_{B,\sigma}(U). \tag{3.1}$$

This proves (ii).

Rank properties and the product case. Monotonicity follows immediately from (3.1): a coordinate projection cannot increase affine dimension.

For submodularity, let $U, U' \subseteq W$. Translate $\mathcal{V}_{B,\sigma}(U \cup U')$ to a linear subspace $L \subseteq \mathbb{F}_2^{U \cup U'}$, and write $\rho(V) = \dim \pi_V(L)$. By (3.1), $\rho(V) = d_{B,\sigma}(V)$ for every $V \subseteq U \cup U'$. Rank–nullity gives $\rho(U) = \dim L - \dim(L \cap \mathbb{F}_2^{U' \setminus U})$. If $L' = \pi_{U'}(L)$, then similarly $\rho(U \cap U') = \rho(U') - \dim(L' \cap \mathbb{F}_2^{U' \setminus U})$. The projection $\pi_{U'}$ acts as the identity on vectors supported in $U' \setminus U$, and hence $\dim(L \cap \mathbb{F}_2^{U' \setminus U}) \le \dim(L' \cap \mathbb{F}_2^{U' \setminus U})$. Combining these three identities and using $\dim L = \rho(U \cup U')$ yields $\rho(U) + \rho(U') \ge \rho(U \cup U') + \rho(U \cap U')$, which is the required submodularity.

Finally, suppose $S_1, \ldots, S_h$ are pairwise disjoint and the dimensions add. By the projection identity, $\mathcal{V}_{B,\sigma}(\bigcup_i S_i) \subseteq \prod_i \mathcal{V}_{B,\sigma}(S_i)$. The left-hand side has dimension $d_{B,\sigma}(\bigcup_i S_i)$, while the right-hand side has dimension $\sum_i d_{B,\sigma}(S_i)$. These dimensions are equal by assumption. Two affine subspaces, one contained in the other and of the same finite dimension, must coincide. This proves (iii). $\square$

We now package the information that will be carried along a branch of the restriction process.

A *closed state* is a pair (B, σ) in which B is admissible and σ is an assignment on X_B satisfying every equation of K_B. If (B, σ) and (B', σ') are closed states with $B \subseteq B'$ and σ' extending σ, then (B', σ') is a *closed extension* of (B, σ).

Let t be a nonconstant term with support $S = \mathrm{vars}(t)$ disjoint from X_B, and suppose $B \cup S$ is admissible. The literals of t specify one point of $\mathbb{F}_2^S$. We call t *compatible* with (B, σ) if that point lies in $\mathcal{V}_{B,\sigma}(S)$, and *incompatible* otherwise. Equivalently, compatibility means that t can be made true by a local solution extending σ.

If (B', σ') is a closed extension, we say that t *survives the extension* if $t|_{\sigma'}$ is nonconstant and compatible with (B', σ'). For a compatible nonconstant term we define $\mu_{B,\sigma}(t) = |S| + d_{B,\sigma}(S)$. If t has width at most k, then $1 \le \mu_{B,\sigma}(t) \le 2k$.

Compatible terms $t_1, \ldots, t_h$ form a *product family over* (B, σ) if their supports $S_1, \ldots, S_h$ are pairwise disjoint, the seed $B \cup \bigcup_i S_i$ is admissible, and $d_{B,\sigma}(\bigcup_{i=1}^h S_i) = \sum_{i=1}^h d_{B,\sigma}(S_i)$. By Proposition 3.1(iii), this is equivalent to saying that the joint local affine space is the direct product of the individual local spaces.

The next elementary lemma is the reason for including affine dimension in the potential. Once the state is enlarged, affine freedom can only be lost. We state explicitly the admissibility condition needed to compare the old support with the larger state; this condition will be guaranteed by the size bookkeeping in the deterministic proof.

Lemma 3.2 (Stability under closed extensions)**.** *Let (B', σ') be a closed extension of (B, σ), let t be a nonconstant term with support S disjoint from X_B, and assume that $B' \cup S$ is admissible. Put $t' = t|_{\sigma'}$.*

- (i) *If t is compatible with (B, σ) and t' is nonconstant, then $d_{B',\sigma'}(\mathrm{vars}(t')) \le d_{B,\sigma}(S)$. In particular, if t survives the extension, its affine-dimension contribution to the potential cannot increase.*
- (ii) *If t is incompatible with (B, σ), then t' is either zero or incompatible with (B', σ'). In particular, an incompatible term can neither become true nor survive such a closed extension.*

Proof. Let $S' = \mathrm{vars}(t') = S \setminus X_{B'}$. Take any $\alpha' \in \mathcal{V}_{B',\sigma'}(S')$. By definition, $\sigma' \cup \alpha'$ extends to a solution of $K_{B' \cup S'}$. Since $B' \cup S' \subseteq B' \cup S$ and the latter seed is admissible, Proposition 3.1(i) extends this solution to $K_{B' \cup S}$.

Restrict the resulting assignment to S. On S' it equals α', while on $S \setminus S'$ its values are fixed by σ'. Restricting further to the smaller closure associated with $B \cup S$ shows that this assignment on S belongs to $\mathcal{V}_{B,\sigma}(S)$. Thus adjoining the fixed coordinates in $S \setminus S'$ gives an injective affine map $\mathcal{V}_{B',\sigma'}(S') \hookrightarrow \mathcal{V}_{B,\sigma}(S)$, which proves (i).

For (ii), if t' were compatible with (B', σ'), the same extension-and-restriction argument would produce a point of $\mathcal{V}_{B,\sigma}(S)$ satisfying t. If $t' = 1$, the values supplied by σ' on S give the same conclusion. Both contradict the incompatibility of t with (B, σ). □

4. The affine-rank revision principle

We can now isolate the combinatorial step that replaces the closure-covering dichotomy of Sofronova and Sokolov. The statement says exactly what the preceding discussion suggests: either we can find many terms that are genuinely independent in the local affine geometry, or a small set of queries decreases the potential of every surviving term.

Proposition 4.1 (Affine-rank revision principle)**.** *Let (B, σ) be a closed state, let F be a collection of compatible nonconstant terms of width at most k, and let $h \geq 1$. Assume that $B \cup \bigcup_{t \in F'} \mathrm{vars}(t)$ is admissible for every subfamily $F' \subseteq F$ of size at most h. Then at least one of the following holds.*

(i) *The family F contains a product family of h terms over (B, σ).*
(ii) *There is a set U of fewer than kh variables such that, for every closed state $(B \cup U, \sigma')$ extending (B, σ) and every $t \in F$ that survives this extension, if $t' = t|_{\sigma'}$, then $\mu_{B \cup U, \sigma'}(t') < \mu_{B,\sigma}(t)$.*

Proof. Start with $U = \varnothing$. As long as possible, choose a term $t \in F$ with support S such that $S \cap U = \varnothing$ and

$$d_{B,\sigma}(U \cup S) = d_{B,\sigma}(U) + d_{B,\sigma}(S), \tag{4.1}$$

and then replace U by $U \cup S$.

If the process selects h terms, their supports are pairwise disjoint and the equalities in (4.1) telescope. Hence the dimension on their union is the sum of the individual dimensions, and we have a product family.

Suppose instead that the process stops after fewer than h selections. Then $|U| < kh$. Consider a closed extension $(B \cup U, \sigma')$ and a term $t \in F$ that survives it. Write $S = \mathrm{vars}(t)$ and $S' = \mathrm{vars}(t|_{\sigma'})$. If $S \cap X_{B \cup U} \neq \varnothing$, then $|S'| < |S|$. The seed $(B \cup U) \cup S$ is admissible by the hypothesis of the proposition, so Lemma 3.2(i) applies and the affine dimension cannot increase. Therefore the potential decreases strictly.

It remains to consider the case $S \cap X_{B \cup U} = \varnothing$. In particular $S \cap U = \varnothing$. Since the greedy process has stopped, additivity fails, and therefore $d_{B,\sigma}(U \cup S) < d_{B,\sigma}(U) + d_{B,\sigma}(S)$. The assignments on S that remain possible after conditioning on the values $\sigma'|_U$ form a subset of a fibre of the projection $\mathcal{V}_{B,\sigma}(U \cup S) \longrightarrow \mathcal{V}_{B,\sigma}(U)$. By Proposition 3.1(ii), every fibre of this affine projection has dimension $d_{B,\sigma}(U \cup S) - d_{B,\sigma}(U) < d_{B,\sigma}(S)$. No variable of S has been assigned in this case, so the width is unchanged while the affine dimension decreases. Again the potential decreases strictly. □

The first alternative becomes useful because a product family is independent under the local affine distribution. We now define the global restriction and quantify this statement.

Definition 4.2 (Random closed restriction and survival)**.** Choose a random seed J by including every variable independently with probability p, and let $\mathcal{G} = \{|J| \leq cr/16\}$. On the event $\mathcal{G}$, choose ρ uniformly from the assignments on X_J satisfying K_J. Outside $\mathcal{G}$ we may define ρ arbitrarily. Fix a closed state (B, σ) with $|B| \leq cr/16$. We say that (B, σ) *survives* ρ if there is a solution of $K_{B \cup J}$ extending both σ and ρ. Equivalently, the two closed states have a common closed extension with seed $B \cup J$.

Survival is stronger than ordinary consistency on the overlap of the two assignments. It also excludes a parity contradiction that becomes visible only when the two closures are combined.

Proposition 4.3 (Restriction estimate for a product family)**.** *Let (B,σ) be a closed state with $|B| \le cr/16$, and let $t_1,\dots,t_h$ be a product family over (B,σ), each term having width at most k. Then*

$$\Pr[\mathcal{G},\ (B,\sigma) \text{ survives } \rho,\ \text{and } t_i|_\rho \ne 1 \text{ for all } i] \le \exp\bigl(-h(p/2)^k\bigr). \tag{4.2}$$

Proof. Let $S_i = \mathrm{vars}(t_i)$. Fix a seed J for which $\mathcal{G}$ holds, and write $I(J) = \{i : S_i \subseteq J\}$. Let $\mathcal{Y}_J$ be the set of assignments on X_J that, together with σ, extend to a solution of $K_{B\cup J}$. Conditional on survival, ρ is uniform on the affine space $\mathcal{Y}_J$.

We first identify the distribution seen on the fully exposed supports. Since $S_i \subseteq J$ for $i \in I(J)$, the projection of $\mathcal{Y}_J$ onto $\bigcup_{i\in I(J)} S_i$ is exactly $\mathcal{V}_{B,\sigma}(\bigcup_{i\in I(J)} S_i)$. One inclusion is immediate. For the other, any point of the local space on the union extends to the closure with seed $B \cup J$ by Proposition 3.1(i), because the union of the exposed supports is contained in J.

The product-family condition and Proposition 3.1(iii) now show that this joint local space is the product of the individual local spaces. Thus, conditional on survival, the assignments on the exposed supports are independent. A compatible term on S_i specifies one point of an affine space of dimension at most $|S_i| \le k$, so it is satisfied with probability at least 2^{-k}.

For this fixed J, the conditional probability that none of the exposed terms is satisfied is therefore at most $(1-2^{-k})^{|I(J)|}$. Discarding the conditional probability of survival, which is at most one, and averaging over the Bernoulli seed gives

$$\begin{aligned}\Pr[\mathcal{G},\ \text{survival, and no } t_i \text{ is true}] &\le \mathbb{E}\bigl[(1-2^{-k})^{|I(J)|}\bigr]\\ &= \prod_{i=1}^{h}\bigl(1 - p^{|S_i|}2^{-k}\bigr)\\ &\le \exp\bigl(-h(p/2)^k\bigr).\end{aligned}$$

Here we used that the supports are disjoint, so the events $S_i \subseteq J$ are independent, and that $|S_i| \le k$. This proves (4.2). ☐

5. The deterministic lower bound

We now prove Theorem 2.1. Assume for contradiction that $\pi = (D_1,\dots,D_s)$ is a $\mathrm{Res}(k)$ refutation of $\mathrm{CNF}(A)$ with $s < 2^{n^\delta}$. The proof has two phases. First we replace every proof line by a small DNF-tree and revise all nonconstant branches. A revision either lowers the potential of every compatible surviving term or stores a large product family. One global restriction then eliminates every branch carrying such a witness. In the second phase we search backwards through the proof DAG, always preserving a locally satisfying common extension. The search must eventually reach an initial clause, where local consistency gives the contradiction.

5.1. **DNF-trees and revision rounds.**

A *DNF-tree* is a rooted binary decision tree. Internal vertices are labelled by variables, the two outgoing edges by the values 0 and 1, and no variable is queried twice on a root-to-leaf path. Each leaf is labelled by 0, by 1, or by a nonconstant DNF. We identify a branch with the partial assignment given by its edge labels.

A tree *strongly represents* a DNF D if the path assignment of every zero-leaf falsifies every term of D, while the path assignment of every one-leaf satisfies at least one term of D. At a nonconstant leaf we impose no additional semantic condition: its label is simply the residual DNF that remains to be dealt with.

Whenever we query more variables below a leaf, we normalize the new leaf labels using the restriction convention from Section 2. A satisfied residual term produces a one-leaf, the disappearance of all residual terms produces a zero-leaf, and otherwise the remaining nonconstant terms label the leaf. This operation preserves strong representation.

Initially, for every proof line D_j, let T_j be the one-vertex tree whose only leaf is labelled by D_j. We revise all trees in $2k$ rounds, using the threshold $h_i = \lceil n^\delta R^i \rceil$ in round i, for $1 \le i \le 2k$.

Each branch β carries a seed B_β. Its path assignment σ_β has domain exactly X_{B_β}, and the leaf is labelled by the normalized restriction $D_j|_{\sigma_\beta}$. Initially $B_\beta = \varnothing$ and σ_β is empty. We use the following terminology.

(i) A branch is *dead* if σ_β fails to satisfy some equation of K_{B_β}. Such a branch does not define a closed state and is never revised again.

(ii) A nonconstant, nondead branch is declared *broken* if the first alternative of Proposition 4.1 produces a product family on it. We store that product family as a witness and never revise the branch again.

(iii) A nonconstant branch is *active* if it is neither dead nor broken. For an active branch, (B_β, σ_β) is a closed state.

Constant branches are left unchanged in every round.

Consider an active branch in round i, with closed state (B, σ). Among the residual terms at its leaf, let F_σ be the compatible ones. If F_σ is empty, we leave the branch unchanged. Otherwise apply Proposition 4.1 to F_σ with threshold h_i.

If the first alternative holds, the branch is declared broken and the resulting product family of h_i terms is stored. If the second alternative holds, let U be the set it provides. We query, in any fixed order, every variable in $X_{B\cup U} \setminus X_B$.

Every resulting branch has seed $B \cup U$ and carries the corresponding extension of σ to $X_{B\cup U}$. Its leaf is normalized. It is dead precisely when its path assignment fails to satisfy $K_{B\cup U}$; otherwise, if nonconstant, it is active.

We record three consequences of this construction.

Claim 1: size and admissibility bookkeeping. Let H_i be the maximum height of any tree after round i. Then

$$H_i \le 4\lambda k n^\delta R^i \qquad (1 \le i \le 2k), \tag{5.1}$$

and every seed, closure, and local affine space used during the revision process lies in the admissible range. In particular, $H := H_{2k} = o(r/k)$. Indeed, a progress step in round j adds fewer than kh_j variables to the seed. Thus before round i a seed has size less than $k\sum_{j<i} h_j$, and even after adjoining the supports of any h_i candidate terms its size is less than $k\sum_{j\le i} h_j$. Since $R \ge 2$, for all sufficiently large n we have $\sum_{j\le i} h_j \le 4n^\delta R^i$, and therefore $|B| + kh_i \le 4kn^\delta R^i \le 4kn^\delta R^{2k}$. By (2.1), the last quantity is $o(r)$, and hence is smaller than $cr/4$ for all sufficiently large n. This is exactly the admissibility required both in Proposition 4.1 and in every application of Lemma 3.2. Finally, $|X_B| \le \lambda |B|$ gives (5.1), and (2.1) gives $H = o(r/k)$.

Claim 2: compatible terms are exhausted. After $2k$ rounds, every residual term on an active branch is incompatible with the closed state of that branch.

Suppose otherwise, and take a compatible nonconstant term on an active branch after the last round. Trace this term backwards through the successive restrictions. By Lemma 3.2(ii), none of its ancestors could have been incompatible: an incompatible term cannot later survive as a compatible one. Thus in every round the branch containing the relevant ancestor had at least one compatible term and was revised. It could not have become broken, since broken branches are never extended. Consequently the second alternative of Proposition 4.1 applied in every round, and the positive integer potential of the surviving descendant decreased strictly each time. The initial potential was at most $2k$, so it cannot decrease strictly through all $2k$ rounds and remain positive. This contradiction proves the claim.

5.2. **One global restriction.** Let ρ be sampled as in Theorem 4.2, with seed J. A nondead branch β *survives* ρ if its closed state (B_β, σ_β) survives ρ. A dead branch never survives. We say that a broken branch is *eliminated by* ρ if either it does not survive or one of its stored witness terms is made true by the common extension of σ_β and ρ. In the first case the branch is incompatible with the global restriction; in the second its leaf DNF has already become true.

Claim 3: there is one restriction that eliminates every broken branch. There is a choice of J and ρ, with $|J| \le cr/16$, that eliminates all broken branches in all proof-line trees simultaneously.

Put $q = (p/2)^k$. A branch broken in round i stores h_i terms, so Proposition 4.3 bounds the probability that it survives without satisfying any stored term by e^{-qh_i} on the event $\mathcal{G}$. Before round i, the total number of branches in all trees is at most $s2^{H_{i-1}}$. A union bound therefore gives

$$\Pr[\mathcal{G} \text{ and some broken branch is not eliminated}] \le \sum_{i=1}^{2k} s2^{H_{i-1}}e^{-qh_i}. \tag{5.2}$$

Since $qR = 64\lambda k$ and $H_{i-1} \le 4\lambda kn^{\delta}R^{i-1}$, while $s < 2^{n^{\delta}}$, each summand in (5.2) is $\exp\bigl(-\Omega(\lambda kn^{\delta}R^{i-1})\bigr)$. The sum is therefore $o(1)$. On the other hand, $\mathbb{E}|J| = np = cr/32$, and a Chernoff bound gives $\Pr[\neg\mathcal{G}] \le e^{-cr/96}$. Hence a pair (J, ρ) with the desired property exists. Fix such a pair for the remainder of the proof.

5.3. **Backward search through the proof DAG.** The global restriction has disposed of the broken branches. We now show that the remaining branches cannot support a refutation.

Write $H = H_{2k}$. During one backward step we use a *temporary seed C*. It starts with the seed of the current branch together with J, and is enlarged while we traverse the premise trees of one inference. There are at most $k+1$ premise trees: one for weakening or and-elimination, at most k for and-introduction, and two for cut. Traversing one tree queries at most H variables, and at the reached leaf we may add at most another H variables from its seed. Thus a backward step adds at most $2(k+1)H$ variables, and we will occasionally reserve a further k variables for the support of one term.

By Claim 1, $H = o(r/k)$. Also (2.1) implies $k = o(r)$. Since $|J| \le cr/16$, it follows that, for all sufficiently large n, every temporary seed occurring in one backward step has size at most $cr/8$, and remains admissible even after adjoining any additional set of at most k variables. We will use this margin repeatedly without further comment.

At a proof line D_v we maintain a branch β_v with the following properties.

(a) β_v is either a zero-branch or an active branch;
(b) β_v survives ρ;
(c) no residual term at the leaf of β_v can become true in an admissible common closed extension: more precisely, if θ is a solution of $K_{B_{\beta_v}\cup J}$ extending σ_{β_v} and ρ, then there do not exist an admissible seed $C \supseteq B_{\beta_v} \cup J$, a solution θ' of K_C extending θ, and a residual term t of $D_v|_{\sigma_{\beta_v}}$ such that $C \cup \mathrm{vars}(t)$ is admissible and $t|_{\theta'} = 1$.

For a zero-branch, property (c) is vacuous because its normalized leaf has no residual term. For an active branch, Claim 2 says that every residual term is incompatible with the branch state, and Lemma 3.2(ii) gives (c). The final line of the refutation is the empty DNF, whose one-vertex zero-tree satisfies the invariant.

Claim 4: the invariant moves to a premise. Suppose D_v is inferred from earlier lines $D_{u_1}, \dots, D_{u_a}$ and the invariant holds at D_v. Then it holds at one of the premises.

Start with $C = B_{\beta_v} \cup J$ and choose a solution θ of K_C extending both σ_{β_v} and ρ. We keep this same temporary assignment throughout the backward step, extending it whenever the seed grows.

Traverse the premise trees one after another. At a node querying a variable x, follow the value $\theta(x)$ if $x \in X_C$. Otherwise replace C by $C \cup \{x\}$, extend θ to the new closure using Proposition 3.1(i), and follow the newly assigned value. When a leaf branch β is reached, enlarge C to $C \cup B_\beta$ and extend once more. The resulting solution extends σ_β and ρ, and therefore witnesses that β survives. The size estimate above guarantees that all these extensions are admissible.

Suppose that for some premise the current temporary assignment satisfies no term of that premise. The reached branch is not dead, since the temporary assignment solves its closure. It is not broken either: a broken branch that survives is eliminated by our fixed global restriction, so one of its witness terms would be true under $\sigma_\beta \cup \rho$ and hence under the temporary extension. Finally it cannot be a one-branch, because strong representation would then give a satisfied term of the premise. Thus it is either a zero-branch or an active branch. In either case properties (a)–(c) follow exactly as above, and the invariant has moved to this premise.

We may therefore assume that the final temporary assignment satisfies a term of every premise. Notice that the premise trees were traversed successively but the temporary assignment was never resampled; it was only extended. Hence all these terms are simultaneously true under one common assignment.

The rules of $\mathrm{Res}(k)$ now imply that a term of the conclusion D_v is true under the same assignment. For weakening and and-elimination this is immediate. For and-introduction, either some premise is satisfied by a term of the common DNF F, in which case the same term appears in the conclusion, or every introduced literal ℓ_i is true, in which case their conjunction is true. For cut, if no term of $F \vee G$ were true, the first premise would force all ℓ_i to be true while the second would force at least one $\neg\ell_i$ to be true, an impossibility.

Let t be a term of D_v satisfied by the temporary assignment. Since the current branch is not a one-branch and the temporary assignment extends σ_{β_v}, the restriction $t|_{\sigma_{\beta_v}}$ is a nonconstant residual term. Our admissibility margin allows its support to be adjoined to the temporary seed. Thus the temporary solution violates property (c) of the invariant. This contradiction proves Claim 4.

Starting from the final line and applying Claim 4 repeatedly, we eventually reach an initial clause D in the standard CNF encoding of some equation e. Let W be the three variables of e. Starting from a solution witnessing survival of the current branch, adjoin W to the seed and extend the solution. Since $\partial\{e\} = W$, the equation e belongs to the enlarged closure. The extended assignment therefore satisfies e, and hence satisfies every clause in its standard CNF encoding, in particular D.

If the reached branch is a zero-branch, this is already impossible: strong representation says that its path assignment falsifies D, and the extended assignment agrees with that path assignment. If the branch is active, then $D|_{\sigma_\beta}$ has a residual literal that becomes true in the enlarged closed extension, contradicting property (c). This final contradiction proves Theorem 2.1.

6. Random 3-CNF formulas at variable density

It remains to verify that random 3-CNFs satisfy the hypotheses of the deterministic theorem. There are two probabilistic inputs, both elementary. First, a union bound gives boundary expansion up to a density-dependent scale. Second, the first moment of the number of satisfying assignments gives unsatisfiability in all the regimes considered here.

6.1. **Boundary expansion.** The next proposition records the expansion scale as a function of the clause density $D = m/n$.

Proposition 6.1 (Expansion at variable density)**.** *Fix $0 < \xi < 1$ and $0 < \tau < (1-\xi)/2$, and let $a = 2/(1-\xi)$. There is a constant $\eta = \eta(\xi) > 0$ with the following property. Let $D = D(n)$ satisfy $1 \le D \le n^{(1-\xi)/2-\tau}$, let $m = \lceil nD \rceil$, and let $r = \lfloor \eta n D^{-a} \rfloor$. Then, with probability tending to one, the clause–variable incidence graph of $\Phi_{n,m}$ is an $(r, 3, \xi)$-boundary expander.*

Proof. Let $\beta = (1-\xi)/2$ and $M = 2 - \beta = (3+\xi)/2$, so that $a = 1/\beta$. Call a set of clause positions *bad* if its boundary has size less than ξ times the number of selected positions.

Suppose a bad set has size s. Let $b < \xi s$ be the number of boundary variables and let v be the number of all other neighbouring variables. Every nonboundary neighbour occurs in at least two selected clauses, so counting incidences gives $3s \ge b + 2v$. If $z = b + v$ is the total number of neighbouring variables, then $z \le (3s+b)/2 < Ms$. Thus a bad set of s clause positions must have all its variables inside some set of fewer than Ms variables.

Fix s and z. Choose the s clause positions and a set of z variables containing their neighbourhood. Since $m \le 2nD$ for all sufficiently large n, a union bound gives

$$\begin{aligned}\Pr[\text{such a choice exists}] &\le \binom{m}{s}\binom{n}{z}\left(\frac{z}{n}\right)^{3s} \\ &\le \left(\frac{2eDn}{s}\right)^{s}\left(\frac{en}{z}\right)^{z}\left(\frac{z}{n}\right)^{3s}.\end{aligned}$$

Writing $z = \theta s$ with $0 < \theta < M$, the last expression becomes $[2e^{1+\theta}D\theta^{3-\theta}(s/n)^{2-\theta}]^s$. The function $e^{1+\theta}\theta^{3-\theta}$ is bounded on $0 < \theta \le M$, and $2-\theta > \beta$. Summing over the fewer than $Ms+1$ possible values of z and enlarging the constant if necessary, we obtain

$$\Pr[\text{there is a bad set of size } s] \le \left[C_\xi D\left(\frac{s}{n}\right)^{\beta}\right]^{s} \tag{6.1}$$

for a constant C_ξ depending only on ξ.

For $2 \le s \le \log n$, the assumption $D \le n^{\beta-\tau}$ gives $C_\xi D(s/n)^\beta \le C_\xi n^{-\tau}(\log n)^\beta = o(1)$. For $\log n < s \le r$, we instead have $C_\xi D(s/n)^\beta \le C_\xi D(r/n)^\beta \le C_\xi \eta^\beta$. Choose $\eta > 0$ so small that $C_\xi\eta^\beta < 1/2$. Summing (6.1) over the two ranges gives $o(1)$. The case $s = 1$ is automatic, since every clause contains three distinct variables and hence has boundary size three. □

6.2. **From random clauses to parity equations.** For each clause C of Φ, let a_C be the unique assignment to its three variables that falsifies C. Of the two parity equations on those variables, exactly one is violated by a_C; include that equation in a system A_Φ. By construction, C is one of the four clauses in the standard CNF encoding of its associated parity equation. Hence $\Phi \subseteq \mathrm{CNF}(A_\Phi)$. In particular, every solution of A_Φ satisfies Φ, so A_Φ is unsatisfiable whenever Φ is unsatisfiable. Conversely, any refutation using only the clauses of Φ is also a refutation from the larger axiom set $\mathrm{CNF}(A_\Phi)$. Therefore a lower bound for $\mathrm{CNF}(A_\Phi)$ applies a fortiori to Φ.

For a fixed truth assignment, a uniformly random 3-clause is satisfied with probability 7/8. Thus

$$\mathbb{E}[\#\text{ satisfying assignments of } \Phi] = 2^n(7/8)^m. \tag{6.2}$$

This tends to zero whenever m/n is a fixed constant larger than $\log 2/\log(8/7)$, and therefore also whenever $D = m/n \to \infty$. By Markov's inequality, Φ and hence A_Φ are unsatisfiable with probability tending to one in every regime used below.

6.3. **Proofs of the main theorems.**

Proof of Theorem 1.1. Fix $\alpha > \log 2/\log(8/7)$ and apply Proposition 6.1 with $D = \alpha$, $\xi = 1/4$, and any fixed $0 < \tau < 3/8$. With high probability we obtain the deterministic theorem with $c = 1/4$ and $r = \Theta(n)$. Consequently $p = \Theta(1)$ and $R = O(kC^k)$ for a constant $C = C(\alpha)$, whence $R^{2k} = \exp(O(k^2) + O(k\log k))$. If $k \le \gamma\sqrt{\log n}$, then $4\lambda k n^\delta R^{2k} = n^{\delta+O(\gamma^2)+o(1)}$. On the other hand $r/k = n^{1-o(1)}$ uniformly in this range. Choose $\delta > 0$ and then $\gamma > 0$ sufficiently small that $\delta + O(\gamma^2) < 1$. Then (2.1) holds uniformly for every $1 \le k \le \gamma\sqrt{\log n}$. Combining Theorem 2.1 with the first-moment estimate (6.2) proves the theorem. □

Proof of Theorem 1.2. For part (i), take $D = \log^h n$ and again choose $\xi = 1/4$, say with $\tau = 1/4$. The hypothesis of Proposition 6.1 holds for all sufficiently large n. Since $a = 8/3$, with high probability we may take $r = \Theta(n/\log^{8h/3} n)$ and $p = \Theta(1/\log^{8h/3} n)$. It follows that $R = O(k(C\log^{8h/3} n)^k)$ for a constant $C = C(h)$, and hence

$$\log R^{2k} = \frac{16h}{3}k^2\log\log n + O(k^2 + k\log k). \tag{6.3}$$

If $k \leq \gamma\sqrt{\log n/\log\log n}$, then (6.3) gives $R^{2k} \leq n^{(16h/3)\gamma^2+o(1)}$. Since $r/k = n^{1-o(1)}$, we may choose $\delta > 0$ and then $\gamma > 0$ sufficiently small that $\delta + (16h/3)\gamma^2 < 1$. The condition (2.1) follows uniformly throughout the asserted range of k. The formula is unsatisfiable with high probability by (6.2), so Theorem 2.1 proves part (i).

For part (ii), the hypothesis on ε is equivalent to $2\varepsilon(2K^2+1) < 1$. Choose $\xi > 0$ sufficiently small that, with $a = 2/(1-\xi)$, we have

$$a\varepsilon(2K^2+1) < 1. \tag{6.4}$$

Set $D = n^{\varepsilon}$. Since $\varepsilon < 1/a$, Proposition 6.1 applies with any fixed $0 < \tau < 1/a - \varepsilon$ and gives, with high probability, $r = \Theta(n^{1-a\varepsilon})$ and $p = \Theta(n^{-a\varepsilon})$. For every fixed $1 \leq k \leq K$, we have $R = O_K(n^{a\varepsilon k})$ and $R^{2k} = O_K(n^{2a\varepsilon k^2})$. Choose $0 < \delta < 1 - a\varepsilon(2K^2+1)$, which is possible by (6.4). Uniformly for $1 \leq k \leq K$,

$$4\lambda k n^{\delta} R^{2k} = o(n^{1-a\varepsilon}/k) = o(r/k).$$

Thus (2.1) holds for all $1 \leq k \leq K$. A final application of Theorem 2.1, together with (6.2), completes the proof. ☐

References

[1] M. Alekhnovich, *Lower bounds for k-DNF resolution on random 3-CNFs*, Computational Complexity **20** (2011), 597–614. Preliminary version in Proceedings of the 37th annual ACM Symposium on Theory of Computing *(STOC 2005)*.

[2] P. Beame, R. Karp, T. Pitassi, and M. Saks, *The efficiency of resolution and Davis–Putnam procedures*, SIAM Journal of Computing **31** (2002), 1048–1075. Preliminary versions in Proceedings of the 37th Annual Symposium on Foundations of Computer Science *(FOCS 1996)* and the 30th ACM Symposium on Theory of Computing *(STOC 1998)*.

[3] E. Ben-Sasson and R. Impagliazzo, *Random CNFs are hard for the polynomial calculus*, Computational Complexity **19** (2010), 501–519. Preliminary version in Proceedings of the 40th Annual Symposium on Foundations of Computer Science *(FOCS 1999)*.

[4] E. Ben-Sasson and A. Wigderson, *Short proofs are narrow—resolution made simple*, Journal of the ACM **48** (2001), 149–169. Preliminary version in Proceedings of the 31st annual ACM Symposium on Theory of Computing *(STOC 1999)*.

[5] V. Chvátal and E. Szemerédi, *Many hard examples for resolution*, Journal of the ACM **35** (1988), 759–768.

[6] N. Fleming, D. Pankratov, T. Pitassi, and R. Robere, *Random $\Theta(\log n)$-CNFs are hard for cutting planes*, in Proceedings of the 58th Annual Symposium on Foundations of Computer Science *(FOCS 2017)*, 2017, 109–120.

[7] S. Gryaznov and N. Talebanfard, *Bounded-depth Frege lower bounds for random 3-CNFs via deterministic restrictions*, arXiv:2403.02275, 2024.

[8] A. Haken, *The intractability of resolution*, Theoretical Computer Science **39** (1985), 297–308.

[9] P. Hrubeš and P. Pudlák, *Random formulas, monotone circuits, and interpolation*, in Proceedings of the 58th Annual Symposium on Foundations of Computer Science *(FOCS 2017)*, 2017, 121–131.

[10] E. Khaniki, *On proof complexity of resolution over polynomial calculus*, ACM Transactions on Computational Logic **23** (2022), no. 3, Article 16, 24 pp.

[11] J. Krajíček, *Bounded arithmetic, propositional logic, and complexity theory*, Encyclopedia of Mathematics and its Applications, vol. 60, Cambridge University Press, 1995.

[12] N. Segerlind, S. R. Buss, and R. Impagliazzo, *A switching lemma for small restrictions and lower bounds for k-DNF resolution*, SIAM Journal of Computing **33** (2004), 1171–1200. Preliminary version in Proceedings of the 43rd Annual Symposium on Foundations of Computer Science *(FOCS 2002)*.

[13] A. Sofronova and D. Sokolov, *A lower bound for k-DNF resolution on random CNF formulas via expansion*, in *40th Computational Complexity Conference (CCC 2025)*, LIPIcs 339, Article 32, 2025.

[14] D. Sokolov, *Random $(\log n)$-CNFs are hard for cutting planes (again)*, in the Proceedings of the 56th Annual ACM Symposium on Theory of Computing *(STOC 2024)*, 2024, 2003–2016.

[15] A. Urquhart, *Hard examples for resolution*, Journal of the ACM **34** (1987), 209–219.

Trinity College Cambridge, Department of Pure Mathematics and Mathematical Statistics, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, United Kingdom.

Email address: gc645@cam.ac.uk

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il