Sharp Threshold Results for Computational Complexity*

Lijie Chen
lijieche@mit.edu
MIT
Cambridge, MA, USA

Ce Jin
jinc16@mails.tsinghua.edu.cn
Tsinghua University
Beijing, China

R. Ryan Williams
rrw@mit.edu
MIT
Cambridge, MA, USA

ABSTRACT

We establish several “sharp threshold” results for computational complexity. For certain tasks, we can prove a resource lower bound of n^c for $c \geq 1$ (or obtain an efficient circuit-analysis algorithm for n^c size), there is strong intuition that a similar result can be proved for larger functions of n , yet we can also prove that replacing “ n^c ” with “ $n^{c+\varepsilon}$ ” in our results, for any $\varepsilon > 0$, would imply a breakthrough $n^{\omega(1)}$ lower bound.

We first establish such a result for **Hardness Magnification**. We prove (among other results) that for some c , the Minimum Circuit Size Problem for $(\log n)^c$ -size circuits on length- n truth tables (MCSP $[(\log n)^c]$) does not have $n^{2-o(1)}$ -size probabilistic formulas. We also prove that an $n^{2+\varepsilon}$ lower bound for MCSP $[(\log n)^c]$ (for any $\varepsilon > 0$ and $c \geq 1$) would imply major lower bound results, such as NP does not have n^k -size formulas for all k , and #SAT does not have log-depth circuits. Similar results hold for time-bounded Kolmogorov complexity. Note that *cubic* size lower bounds are known for probabilistic De Morgan formulas (for other functions).

Next we show a sharp threshold for **Quantified Derandomization** (QD) of probabilistic formulas: (a) For all $\alpha, \varepsilon > 0$, there is a deterministic polynomial-time algorithm that finds satisfying assignments to every probabilistic formula of $n^{2-2\alpha-\varepsilon}$ size with at most 2^{n^α} falsifying assignments. (b) If for some $\alpha, \varepsilon > 0$, there is such an algorithm for probabilistic formulas of $n^{2-\alpha+\varepsilon}$ -size and 2^{n^α} unsatisfying assignments, then a *full derandomization* of NC¹ follows: a deterministic poly-time algorithm additively approximating the acceptance probability of *any polynomial-size formula*. Consequently, NP does not have n^k -size formulas, for all k .

Finally we show a sharp threshold result for **Explicit Obstructions**, inspired by Mulmuley’s notion of explicit obstructions from GCT. An *explicit obstruction against $S(n)$ -size formulas* is a poly-time algorithm A such that $A(1^n)$ outputs a list $\{(x_i, f(x_i))\}_{i \in [\text{poly}(n)]} \subseteq \{0, 1\}^n \times \{0, 1\}$, and every $S(n)$ -size formula F is inconsistent with the (partially defined) function f . We prove that for all $\varepsilon > 0$, there is an explicit obstruction against $n^{2-\varepsilon}$ -size formulas, and prove that there is an explicit obstruction against $n^{2+\varepsilon}$ -size formulas for some $\varepsilon > 0$ if and only if there is an explicit obstruction against *all*

polynomial-size formulas. This in turn is equivalent to the statement that E does not have $2^{o(n)}$ -size formulas, a breakthrough in circuit complexity.

CCS CONCEPTS

• **Theory of computation** → **Circuit complexity**.

KEYWORDS

hardness magnification, quantified derandomization, de morgan formulas

ACM Reference Format:

Lijie Chen, Ce Jin, and R. Ryan Williams. 2020. Sharp Threshold Results for Computational Complexity. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC '20)*, June 22–26, 2020, Chicago, IL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3357713.3384283>

1 INTRODUCTION

How far are we from proving major lower bounds such as $P \neq NP$, polynomial-size formula lower bounds for NP, or $P^{\#P} \neq NC^1$? Well-known complexity barriers [1, 14, 56] have shown that, in the absence of radically new ideas, we appear quite far from resolving many fundamental lower bound problems, given the current proof methods.

Suppose we accept that prognosis, and resign ourselves to being far (for now). Can we quantify “how far” we are from various open lower bounds? Is there a *fine-grained theory* of complexity barriers?¹ Such a theory would ideally help us order lower bound problems by their “distance” from known lower bounds, and gather better intuition overall for the relative difficulty of proving lower bounds in the complexity landscape.

Our idea for addressing these questions starts with asking: *what would a lower bound we are “close to” look like?* We will use the following natural heuristic: If we can prove a resource lower bound of n^c , and we have no evidence against (possibly also evidence for) proving an $n^{c+\varepsilon}$ lower bound for arbitrarily small $\varepsilon > 0$, it seems reasonable to say we are “close to” an $n^{c+\varepsilon}$ lower bound. Viewed from this perspective, we would be “close to” a lower bound L if there is a task Π such that:

- an n^c lower bound can be proved for Π ,
- we *know* how to prove $n^{c+\varepsilon}$ lower bounds for other natural tasks,
- yet an $n^{c+\varepsilon}$ lower bound for Π , for *any* $\varepsilon > 0$, would imply lower bound L .

(Again, this notion only makes sense if we believe the current lower bound against Π is not optimal.) Note, saying we are “close to” a

¹See Section 1.4 for a brief comparison with the methodology of fine-grained complexity.

*L.C. and R.W. are supported by NSF CCF-1741615 and a Google Faculty Research Award. Portions of this work were completed while L.C. and R.W. were visiting the Simons Institute at UC Berkeley.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

STOC '20, June 22–26, 2020, Chicago, IL, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6979-4/20/06...\$15.00

<https://doi.org/10.1145/3357713.3384283>

lower bound obviously does not mean that it will be easy to prove! It just means that what is necessary does not look *quantitatively* far, in our relative estimation of lower bound difficulty.²

In this paper, we establish several counterintuitive “sharp threshold” phenomena for some major open problems in lower bounds and derandomization. For certain tasks Π , we are able to prove a lower bound of n^c resources and/or provide an efficient derandomization for n^c -size objects, we have good reasons to believe that a stronger result is provable, yet improving our result to $n^{c+\varepsilon}$ for some $\varepsilon > 0$ would unexpectedly imply a superpolynomially-strong lower bound. Whether or not this means we are “close” to a major breakthrough remains to be seen. Our key message is that our collective intuition about what is a “weak” lower bound (and what is a “weak” derandomization) is still surprisingly poor and incomplete.

1.1 Sharp Thresholds for MCSP Lower Bounds Against Probabilistic Formulas

A major object of study in this paper is the computational model of probabilistic formulas. This model is a natural choice because it is an interesting expressive model for which we know how to prove nearly cubic lower bounds. For simplicity, throughout this paper, “formulas” means “De Morgan formulas” by default.

Definition 1.1. A *probabilistic formula* is a distribution $\mathcal{F}$ over De Morgan formulas. We say $\mathcal{F}$ computes a function f , if for all x , $\Pr_{F \sim \mathcal{F}}[F(x) = f(x)] \geq 2/3$.

Remark 1.2. By a simple “Chernoff bound, and union bound” argument, we may assume the distribution is a uniform distribution over $O(n)$ De Morgan formulas. Thus a probabilistic formula can always be written as an “Approximate-Majority” of $O(n)$ De Morgan formulas. This is helpful when discussing algorithms that take probabilistic formulas as inputs, as in that case we need a succinct way to describe a (potentially exponential-sized) distribution over De Morgan formulas.

We present a sharp hardness magnification threshold for probabilistic formulas. We first show that slightly super-quadratic probabilistic formula lower bounds for computing basic compression problems would imply breakthrough complexity separations. The Minimum Circuit Size Problem for size parameter $s(n)$ ($\text{MCSP}[s(n)]$) asks whether a given length- n truth-table has an $s(n)$ -size circuit on $m = \log n$ input bits. This is a basic problem that has seen a flurry of research activity in recent years (e.g. [7, 8, 30–34, 39, 50, 63]). The Minimum Kt-complexity Problem ($\text{MKtP}[s(n)]$) asks whether a given length- n input string has Kt-complexity at most $s(n)$. See Section 2.2 for formal definitions.

THEOREM 1.3 (FROM “WEAK” LOWER BOUNDS TO SUPER-POLYNOMIAL LOWER BOUNDS, ADAPTING [19]).

- (1) If there is an $\varepsilon > 0$ such that for all sufficiently small $\alpha > 0$, $\text{MCSP}[n^\alpha]$ does not have $n^{2+\varepsilon}$ -size probabilistic formulas, then $\oplus P \not\subseteq NC^1$.

²For example, it is reasonable to believe that it will be easier to prove SAT does not have an $O(n)$ -time algorithm in a random access model, than it will be to prove SAT does not have an $O(n^2)$ -time algorithm. That does not mean it will be easy to prove SAT isn’t in linear time! Note that this example does not quite fit our mold, in that, modulo time hierarchies, we do not really know how to prove super-linear time lower bounds.

- (2) If there is an $\varepsilon > 0$ such that for all sufficiently small $\alpha > 0$, $\text{MKtP}[n^\alpha]$ does not have $n^{2+\varepsilon}$ -size probabilistic formulas, then $\text{EXP} \not\subseteq NC^1$.
- (3) If there is an $\varepsilon > 0$ and a family of languages $\{L_\beta\}$ (indexed over $\beta \in (0, 1)$) such that L_β is a 2^{n^β} -sparse NP language and for all β L_β does not have $n^{2+\varepsilon}$ -size probabilistic formulas, then NP does not have n^k -size formulas, for all k .

In fact, lower bounds which are “barely” above quadratic size would still imply a breakthrough.

THEOREM 1.4. For any unbounded function $f(n)$, the following hold.

- (1) If there is a $d > 0$ such that $\text{MCSP}[(\log n)^d]$ does not have $n^2(\log n)^{f(n)}$ -size probabilistic formulas, then $\oplus P \not\subseteq NC^1$.
- (2) If there is a $d > 0$ such that, $\text{MKtP}[(\log n)^d]$ does not have $n^2(\log n)^{f(n)}$ -size probabilistic formulas, then $\text{EXP} \not\subseteq NC^1$.

Note it is widely believed that $\text{MCSP}[s(n)]$ and $\text{MKtP}[s(n)]$ are not in P_{poly} , for any $s(n)$ which is a sufficiently large $\text{poly}(\log n)$ [6, 39]. We show how to complement the hardness magnification theorem above in a sharp way, proving lower bounds for MCSP and MKtP that are quantitatively very close to the hypotheses in Theorem 1.3 and Theorem 1.4.

THEOREM 1.5 (SUB-QUADRATIC MKTP AND MCSP LOWER BOUNDS). There are $c > 0, d > 0, K > 0$ such that for all $t(n)$ satisfying $c \log n < t(n) \leq n/20$, and $s(n)$ satisfying $(\log n)^d < s(n) \leq n/200 \log n$,

- (1) $\text{MKtP}[t(n)]$ does not have probabilistic formulas of $n^{2-K/\log \log n / t(n)}$ size.
- (2) $\text{MCSP}[s(n)]$ does not have probabilistic formulas of $n^{2-K/\log \log n / s(n)}$ size.

More generally, the lower bound of Theorem 1.5 holds for “gap” versions of MCSP and MKtP (see Section 5.2 for details). The MKtP lower bound is particularly interesting because $\text{MKtP}[c \log n]$ is a *polynomially-sparse* language computable in P .³ If the lower bound of Theorem 1.5 can be improved to $n^{2+\varepsilon}$ for a subexponentially-sparse problem computable in $2^{n^{o(1)}}$ time, then $\text{EXP} \not\subseteq NC^1$.

When the circuit size parameter is close to maximal, Theorem 1.5 can be improved. We adapt recent MCSP lower bounds for larger circuit sizes [23] to achieve super-quadratic formula lower bounds.

THEOREM 1.6 (SUPER-QUADRATIC MCSP LOWER BOUNDS). For every $\alpha \in (0, 1)$ and $\varepsilon > 0$, $\text{MCSP}[n^\alpha]$ doesn’t have probabilistic $n^{2+\alpha-\varepsilon}$ -size formulas.

Thus, MCSP gets “harder” as the size parameter increases. Theorem 1.6 should be contrasted with Item (1) of Theorem 1.3. If we could prove a similar theorem for $\text{MCSP}[n^{o(1)}]$ (or any sparse-enough NP problem, by Item (3) of Theorem 1.3), then we could apply Theorem 1.3 and conclude a breakthrough complexity separation.

³To decide whether $\text{Kt}(x) \leq c \log n$, it suffices to enumerate all Turing machines of $c \log n$ -bit description length and run them for n^c steps, which takes $\text{poly}(n)$ time (see Definition 2.2).

Comparison With Prior Work. Several recent papers [19–21, 53–55] also establish hardness magnification results for problems such as MCSP and MKtP. Most recently, the reference [19] proves Theorem 1.3 for deterministic formula lower bounds of size $n^{3+\epsilon}$; we generalize their results for probabilistic formula lower bounds of size $n^{2+\epsilon}$. Our Theorem 1.5 generalizes $n^{2-\epsilon}$ -size formula lower bounds for MCSP [32, 54, 55] on larger circuit sizes. In [19, 54, 55] it is shown that an $n^{3+\epsilon}$ -size formula lower bound for MKtP $[n^{o(1)}]$ implies $\text{EXP} \not\subseteq \text{NC}^1$, and [32, 54] showed that MKtP[polylog(n)] does not have $n^{2-\epsilon}$ -size formulas. That is, the “gap” between the magnification threshold and known lower bounds was $n^{3+\epsilon}$ versus $n^{2-\epsilon}$. Theorems 1.3 and 1.5 improve both of these results, and show an *arbitrarily small* gap between the magnification threshold and known lower bounds for probabilistic formulas.

Similar sharp threshold results were known before, but apparently only in cases where lower bound techniques are few. For example, [46] show $\text{MCSP}[n^{o(1)}] \notin \text{SIZE}[n^{1+\epsilon}]$ implies $\text{NP} \not\subseteq \text{P}_{/\text{poly}}$, and it is not hard to establish an $\Omega(n)$ -size lower bound for MCSP $[n^{o(1)}]$. However, in the case of general circuits, researchers have been stuck for decades on proving even a $6n$ -size lower bound for De Morgan circuits,⁴ so an $n^{1+\epsilon}$ -size lower bound on MCSP seems *far*, relatively speaking. For another example, Allender and Koucky [9] (later refined by Chen and Tell [21]) show that if certain NC^1 -complete problems do not have $n^{1+\epsilon}$ size TC^0 -size circuits, then NC^1 does not have polynomial-size TC^0 circuits. However, there is still a gap between these results and the best known lower bounds for TC^0 [22, 36].

To compare with our case, note that any average-case lower bound for (deterministic) formulas immediately implies worst-case lower bounds for probabilistic formulas. Thus there are already multiple known methods for proving $n^{3-\epsilon}$ -size lower bounds against probabilistic formulas computing certain functions (see e.g., [23, 40]).

Additional Related Works on Hardness Magnification. Several other similar hardness magnification phenomena have been established: for $n^{1-\epsilon}$ approximation to CLIQUE [58], sublinear-depth circuit lower bounds for P [44], proof complexity [47], and lower bounds for non-commutative arithmetic circuits [15].

In [18], a barrier for proving strong circuit lower bounds via the hardness magnification techniques (termed as “The Locality Barrier”) was formulated, saying that:

- (a) almost all current lower bound proofs (random restrictions, polynomial approximation, etc.) satisfy a certain property, and,
- (b) lower bound proofs with such a property cannot be used to prove the required lower bounds for MCSP or MKtP to achieve hardness magnification.

It is discussed in [18] (HM Frontier D) that our magnification result for probabilistic formulas is also subject to this barrier.

⁴The state-of-the-art lower bound against De Morgan circuits for an explicit function is $5 \cdot n$ [38, 41], while the state-of-the-art lower bound against B_2 -circuits is the very recent $(3 + 1/86) \cdot n$ by [24].

1.2 Sharp Threshold for Quantified Derandomization

Next we discuss sharp thresholds for quantified derandomization of probabilistic formulas. We first provide the relevant definitions.

Definition 1.7 (Quantified Derandomization (QD) Problem, [28]). For a Boolean circuit class C and a function $B : \mathbb{N} \rightarrow \mathbb{N}$, the QD problem for C with B exceptional inputs is the following: *Given an n -input circuit $C \in C$ that evaluates to 0 on at most $B(n)$ inputs, deterministically output an n -bit string on which C evaluates to 1.*⁵

Note the standard derandomization problem has $B(n) = 2^n/3$, but studying the case of $B(n) = 2^{n^\alpha}$ for $\alpha < 1$ turns out to also be a very interesting and subtle problem [21, 28, 61, 62], with implications for general derandomization.

One way to design QD algorithms is by deterministically constructing hitting sets. A hitting set generator outputs a list L_n of inputs (independent of the given circuit), such that for all possible input circuits C with at most $B(n)$ exceptional inputs, at least one $x \in L_n$ makes C accept. A hitting set generator immediately implies a black-box QD algorithm, which only needs oracle access to C , rather than an explicit description of C . Conversely, from any black-box QD algorithm we obtain a hitting set generator: the hitting set consists of all input strings queried by the black-box algorithm on the all-zeroes circuit.

Generalized Probabilistic Formulas. By definition, a probabilistic formula has probability at least $2/3$ of accepting every YES instance, and at most $1/3$ of accepting every NO instance. We can also consider a generalization of probabilistic formulas which outputs ? if the input does not satisfy either of the two conditions. This allows us to consider arbitrary distributions of formulas. Specifically, a *generalized probabilistic formula* $\mathcal{F}$ defined to be (simply) a distribution over formulas that take n inputs with a different output criterion. On input $x \in \{0, 1\}^n$, define $p_x := \Pr_{F \sim \mathcal{F}}[F(x) = 1]$. We say the value of $\mathcal{F}$ on x is **1** if $p_x \geq 2/3$, **0** if $p_x \leq 1/3$, and **?** otherwise. We say $\mathcal{F}$ has at most $B(n)$ exceptional inputs, if the value of $\mathcal{F}$ is in $\{0, ?\}$ for at most $B(n)$ inputs. We can adapt the definition of hitting set accordingly.

Definition 1.8 (Hitting Sets for Generalized Probabilistic Formulas). $\mathcal{H} \subseteq \{0, 1\}^n$ is a *hitting set for generalized probabilistic formulas with at most $B(n)$ exceptions* if, for every generalized $\mathcal{F}$ with at most $B(n)$ exceptional inputs, there is an $x \in \mathcal{H}$ such that the value of $\mathcal{F}$ on x is **1** or **?**.

Sharp Threshold for QD of Generalized Probabilistic Formulas. We present a sharp threshold for quantified derandomization of generalized probabilistic formulas. First, we give a polynomial-time construction of hitting sets for QD of subquadratic-size generalized probabilistic formulas with subexponentially many exceptional inputs.

THEOREM 1.9 (QUANTIFIED DERANDOMIZATION OF GENERALIZED PROBABILISTIC FORMULAS). *Let $1 \leq s(n) \leq n/20$. There is a poly-time computable hitting set for generalized probabilistic formulas of*

⁵In this paper we consider the *one-sided* QD problem, as defined above. By convention, the majority output value is 1 for $B(n) < 2^n/2$. One can also study a stronger two-sided version, where the majority output value $b \in \{0, 1\}$ is unknown and needs to be decided deterministically.

size $n^{2-K/\log \log n}/s(n)^2$ with at most $B(n) = 2^{s(n)}$ exceptions, where $K > 0$ is a universal constant.

Complementing Theorem 1.9, we prove that if QD for generalized formulas of any slightly larger size is possible, then we obtain a full derandomization of all polynomial-size formulas, in polynomial time!

THEOREM 1.10 (FROM “WEAK” QD TO FULL DERANDOMIZATION). *Suppose there are $\varepsilon > 0$, $\alpha \in (0, 1)$, and a poly-time algorithm for QD of generalized probabilistic formulas of size $n^{2-\alpha+\varepsilon}$ with $B(n) = 2^{n^\alpha}$ exceptional inputs. Then, there is a poly-time algorithm finding a satisfying assignment to any poly-size formula with at most $2^n/3$ falsifying assignments. As a corollary, NP does not have n^k -size formulas for all k .*

Moreover, if the hypothesized algorithm is black-box (only calls the input formula as an oracle), then there is a poly-time algorithm that, for every polynomial-size formula, additively approximates its acceptance probability up to any desired constant error.

THEOREM 1.11. *Suppose there are $\varepsilon > 0$, $\alpha \in (0, 1)$, and a $2^{n^{o(1)}}$ -time algorithm for QD of generalized probabilistic formulas of size $n^{2-\alpha+\varepsilon}$, with $B(n) = 2^{n^\alpha}$ exceptional inputs. Then, there is a $2^{n^{o(1)}}$ -time algorithm that finds a satisfying assignment to every polynomial-size formula having at most $2^n/3$ falsifying assignments. As a corollary, $\text{NE} = \text{NTIME}[2^{O(n)}] \not\subseteq \text{NC}^1$.*

Moreover, if the above algorithm is black-box, then there is a $2^{n^{o(1)}}$ -time algorithm that, for every polynomial-size formula, additively approximates its acceptance probability up to any constant error.

Furthermore, when our probabilistic formulas satisfy the typical BPP-style promise (on every input, it outputs 1 with probability $\geq 2/3$, or 0 with probability $\geq 2/3$), we can obtain a slightly improved hitting set construction than that of Theorem 1.9; see Section 4 for details and more results.

Comparison With Prior Work. Several papers [21, 28, 61, 62] prove obtaining quantified derandomization “thresholds” for circuit classes such as AC^0 , TC^0 and $\text{AC}^0[\oplus]$. Most related to ours is [21, 61], showing if there is a QD algorithm for n^{1+1/ϕ^d} -wire TC_d^0 circuits with $B(n) = 2^{n^{o(1)}}$ exceptional inputs for any $\phi < 1.62$, then a full derandomization of TC^0 follows. They also provide a corresponding QD algorithm for n^{1+1/c^d} -wires TC_d^0 circuits for some constant $c \approx 30$ [61]. However, in their case, the gap is between two constants 1.62 and 30 in the exponent; in our case, the gap in the exponent can be made arbitrarily small.

1.3 Sharp Threshold for “Explicit Proofs” of Formula Lower Bounds

We can also prove a sharp threshold for proving lower bounds in a certain “explicit” sense, along the lines of Mulmuley [48] in his pioneering work on GCT. He suggests a concept of *explicit obstruction* where a polynomial-time algorithm produces a proof of a lower bound. It turns out that studying the problem of constructing “explicit proofs” against formulas and branching programs also reveals a sharp threshold phenomenon. We start by defining an *explicit obstruction* against a circuit class C . Our definition is more generic than that of Mulmuley (it has fewer constraints).

Definition 1.12. An *explicit C -obstruction* (or *explicit obstruction against C*) is a (deterministic) polynomial-time algorithm A such that for all large enough n , $A(1^n)$ outputs a list $L_n = \{(x_i, y_i)\}$ such that $x_i \neq x_j$ for all $i \neq j$, and for all n -input $C \in C$, there is an $(x_i, y_i) \in L_n$ such that $C(x_i) \neq y_i$.

First let us compare this obstruction notion with a similar one, that of *anti-checkers* [45, 54]. An *anti-checker for a function f against C* is a list $\{(x_i, f(x_i))\}$ such that every circuit in C fails to compute f on at least one x_i in the list. Observe that every anti-checker is also an obstruction, but an anti-checker is more constrained than an explicit obstruction (the input-output pairs must be from f). Mulmuley [48] argues philosophically that, in order to prove results such as $\text{NP} \not\subseteq \text{P}_{\text{poly}}$, we should strive to construct *explicit obstructions* against small circuits. His Geometric Complexity Theory program is rooted in the idea that this task may be more feasible to achieve for arithmetic circuits. It is clear that explicit obstructions against C imply lower bounds against C . For example, the following is easy to prove:

PROPOSITION 1.13. *If there is an explicit obstruction against C , then there is a function in P that is not in C .*

However, it still unclear what the extra advantage of this “explicit” approach might be. How could thinking about explicit obstructions be helpful, compared to other approaches to proving lower bounds? To understand this question, we ask: what *known* lower bound techniques can be used to achieve explicit obstructions in the Boolean world? First, we apply our technical results to construct an explicit obstruction against sub-quadratic size (De Morgan) formulas.

THEOREM 1.14 (EXPLICIT OBSTRUCTIONS AGAINST SUB-QUADRATIC DE MORGAN FORMULAS). *For a universal constant $K > 0$, there is an explicit obstruction against formulas of size $n^{2-K/\log \log n}$.*

Therefore certain lower bound techniques (based on random restrictions) can be made explicit. However, it is well-known that the best De Morgan formula lower bounds are not quadratic size; they are nearly cubic [29, 60]. So one would expect an improvement on Theorem 1.14.

We show that if any super-quadratic formula lower bound can be used to build an explicit obstruction, we would in fact prove *exponential-size* formula lower bounds for exponential time! In particular, we prove an *equivalence* between explicit obstructions for super-quadratic formulas, exponential-size formula lower bounds, explicit obstructions for *all* polynomial-size formulas, and the existence of optimal PRGs for polynomial-size formulas.

THEOREM 1.15 (BETTER EXPLICIT OBSTRUCTIONS ARE EQUIVALENT TO SUPER-POLYNOMIAL OBSTRUCTIONS). *The following are equivalent:*

- (1) *There is an $\alpha > 0$ and an explicit obstruction against formulas of size $n^{2+\alpha}$.*
- (2) *For all k , there is an explicit obstruction against formulas of size n^k .*
- (3) *There is an $\varepsilon > 0$ and a function in $\text{E} = \text{TIME}[2^{O(n)}]$ that does not have $2^{\varepsilon n}$ -size formulas, even infinitely often.*

As a corollary, any poly-time computable anti-checker for any f that does not have $n^{2+\varepsilon}$ -size De Morgan formulas would also imply

strong formula lower bounds sufficient for a full derandomization of poly(n)-size De Morgan formulas (any anti-checker is also an obstruction). In fact, extending our explicit obstruction even *slightly* beyond quadratic size would imply $n^{\omega(1)}$ size formula lower bounds:

THEOREM 1.16. *If there is an unbounded function $f(n)$ and an explicit obstruction against formulas of size $n^2 \cdot (\log n)^{f(n)}$, then $E \notin NC^1$.*

Explicit Obstructions for Other Models. One can obviously extend the notion of explicit obstruction to other well-studied models of computation. However, for two natural models just slightly more powerful than De Morgan formulas, branching programs and B_2 -formulas (formulas over AND, OR, and XOR), we observe that a “sharp threshold” already kicks in for explicit obstructions against super-linear size. This is very interesting, because $n^{2-o(1)}$ lower bounds for both branching programs and B_2 -formulas are well-known ([51], [65, p.253–255, p.422]). First, we observe:

PROPOSITION 1.17. *There are explicit obstructions against B_2 -formulas of size $o(n)$, and branching programs of size $o(n)$.*

However, we can show that any n^ϵ improvement in this simple result implies a breakthrough in formula or branching program complexity:

THEOREM 1.18. *The following are equivalent:*

- (1) *There is an $\alpha > 0$ and an explicit obstruction against B_2 -formulas of size $n^{1+\alpha}$.*
- (2) *For all k , there is an explicit obstruction against B_2 -formulas of size n^k .*
- (3) *There is an $\epsilon > 0$ such that E cannot be computed by $2^{\epsilon n}$ -size formulas, even infinitely often.*

The same equivalence also holds with “branching programs” in place of “ B_2 -formulas.”

Therefore, making lower bounds for branching programs and/or B_2 -formulas *explicit*, for any super-linear polynomial size bounds, would imply exponential size lower bounds for E . Such lower bounds, in turn, are equivalent to the existence of optimal pseudorandom generators for the corresponding model, by standard results [37, 52, 59]. Thus, explicit obstructions would also yield black-box derandomizations in this case.

1.4 Intuition

In this section we give a high-level description of our proof techniques.

The Key Technical Ingredient: An $O(\log n)$ -seed Pseudorandom Restriction Generator. Our primary technical contribution is a *pseudorandom restriction generator* for subquadratic-size formulas with only $O(\log n)$ seed length. Our starting point is the classical result that the “shrinkage exponent” of De Morgan formulas is 2, first proved by Håstad [29], with logarithmic factors later removed by [60].

Roughly speaking, Håstad’s result says that the formula complexity of any function f is expected to “shrink” by a factor of about p^2 , when each variable is left unassigned with probability p , and is otherwise assigned a random bit. Such a variable restriction is called a *p-regular random restriction*. A uniform random

p -regular restriction requires $\Theta(pn \log(1/p))$ random bits to encode. Our pseudorandom restriction generator provides a *derandomized* version of the [29, 60] shrinkage theorem with an optimal $O(\log n)$ seed length, albeit with an $n^{o(1)}$ -size loss on the expected size.⁶ (A formal theorem statement can be found in Section 3.) In particular, it produces an explicit collection C of poly(n) partial assignments (random restrictions) $\rho : [n] \rightarrow \{0, 1, \star\}$ such that:

- (a) For every formula F , applying a random partial assignment ρ from C “shrinks” the formula complexity of F by about p^2 (with some $n^{o(1)}$ loss), and
- (b) A random ρ from C has decent probability of leaving at least $pn/2$ variables unset.

The collection C is highly explicit; we can generate any bit of its description with a poly($\log n$) size circuit.

Our construction improves generators implicit in [54] which have $O(\log^2 n)$ seed length. In more detail, our construction follows the iterated pseudorandom restriction paradigm of [32, 35, 54]. In each stage, a q -regular random restriction ϕ_i is applied to the remaining variables, where q is a non-zero constant. By composing r independent stages of random restrictions, one obtains a p -regular random restriction $\rho = \phi_r \circ \dots \circ \phi_1$ with $p = q^r$. The key observation is that, in each stage, ϕ_i only has to be $1/q^2$ -wise independent, which takes only $O(\log n)$ random bits if q is a non-zero constant. The total number of stages is $r = \log_q(p) \leq O(\log n)$, and the total seed length is $O(\log^2 n)$.

To obtain an $O(\log n)$ total seed length, our improved construction uses a different parameter setting: we set r to $O(\log n / \log \log n)$, and set q to be about $1/(\log n)^\alpha$ for small enough $\alpha > 0$. We then need to sample each restriction ϕ_i using only $O(\log \log n)$ random bits, while still ensuring the shrinkage property. To achieve this, our one-stage restriction construction uses the following two ideas:

(1) Let $S = \text{poly}(1/q)$. Fix a small set of restrictions $f_1, \dots, f_S \in \{0, 1, \star\}^{\text{poly}(1/q)}$, such that a random f_i ($i \in [S]$) has a shrinkage property on poly($1/q$)-variable $1/q^2$ -size formulas almost as good as q -regular random restrictions. Such $f_1, \dots, f_S$ exist, and can be found deterministically in $2^{\text{poly}(1/q)} < \text{poly}(n)$ time by brute force (for α small enough). After that, specifying an f_i only takes about $\log S$ bits, which is exponentially shorter than the seed-length of sampling a uniform $1/q^2$ -wise independent q -regular restriction.

(2) Construct n -variable restrictions, by composing a random poly($1/q$)-variable restriction f_i with an almost pairwise independent hash function $h : [n] \rightarrow [\text{poly}(1/q)]$, samplable using only $O(\log \text{poly}(1/q) + \log \log n) \leq O(\log \log n)$ bits. For a $1/q^2$ -size formula F , with at least $1 - \text{poly}(1/q)$ probability, the variables that F depends on are hashed into distinct variables. This allows us to extend the shrinkage property on poly($1/q$)-variable formulas to the shrinkage property on n -variable formulas.

Lower Bounds, QD Algorithms, and Explicit Obstructions From the Generator. Now we briefly discuss how our new optimal pseudorandom restriction generator implies some of our other results.

⁶Observe that $\Omega(\log n)$ random bits are required for pseudorandom restrictions for all $p \leq n^{-\Omega(1)}$. To see this, assume a pseudorandom restriction ρ has $o(\log n)$ seed-length. Find a restriction ϕ in its support with maximum $|\phi^{-1}(\star)| \geq pn$. Consider a formula F of maximum size $L(F) = 2^{\Omega(pn)}$ that only depends on $|\phi^{-1}(\star)|$. Then F does not shrink under ϕ . Hence the expected size of F under ρ is at least $n^{-o(1)} \cdot L(F) \gg p^2 L(F) + p\sqrt{L(F)}$.

Lower Bounds for MKtP $[c \log n]$. Let c be a sufficiently large constant. We first provide intuition for the proof that MKtP $[c \log n]$ cannot be computed by probabilistic formulas of size $n^{1.9}$. Suppose otherwise that MKtP $[c \log n]$ can be computed by a probabilistic formula $\mathcal{F}$ of size $n^{1.9}$. Applying our pseudorandom restriction generator with parameter $p = n^{-0.98}$, we show that there must be a restriction ρ with the following properties: (a) $L(F|_\rho) < 1$ with probability 0.9 for a random $F \sim \mathcal{F}$ and (b) $|\rho^{-1}(\star)| \geq pn/2 > n^{0.01}$. First we show that, if x^ρ is the input obtained by filling every $\star$ of ρ with 0, we show that the efficiency of our generator implies that MKtP $[c \log n](x^\rho) = 1$.

Item (a) implies that $\mathcal{F}$ must output the same value on every input consistent with ρ , but $\mathcal{F}$ must output 1 on every x^ρ . Therefore, $\mathcal{F}$ must output 1 on every input that is consistent with ρ . But there are at least $2^{n^{\Omega(1)}}$ many inputs consistent with ρ , and (by a counting argument) at least one of them must make MKtP $[c \log n]$ output 0. This is a contradiction.

Hitting Set for Quantified Derandomization. Now we sketch how to obtain a quantified derandomization algorithm for generalized probabilistic formulas. For simplicity we focus on the case of size $n^{1.9}$, and $B(n) = 2^{n^{0.01}}$. Apply the pseudorandom restriction generator again with $p = n^{-0.98}$, and fix a restriction ρ satisfying (a) and (b). Observe that for generalized probabilistic formulas, (a) implies for every input x consistent with ρ , $\mathcal{F}(x) \neq 0$. Therefore, we can simply let the hitting set be the collection of x^ρ (as defined above) for all ρ generated by our pseudorandom restriction generator.

Explicit Obstructions. We sketch how to obtain an explicit obstruction for $n^{1.9}$ -size formulas. As before, we use our pseudorandom restriction generator with $p = n^{-0.98}$. We show that for all formula F of size $n^{1.9}$, there is a restriction ρ from our generator such that $|\rho^{-1}(\star)| > 0$, and $L(F_\rho) < 1$. For each ρ and each $i \in \rho^{-1}(\star)$, we set $x^{\rho, i}$ be the input obtained by extending ρ so that the i -th variable is set to 1 and all other $\star$'s filled with 0's. We add both $(x^\rho, \text{PARITY}(x^\rho))$ and $(x^{\rho, i}, \text{PARITY}(x^{\rho, i}))$ to our list S .

We show this is an *anti-checker* for the PARITY function (hence it is an explicit obstruction). In particular, for every formula F , since $L(F_\rho) < 1$, F must output the same value on both x^ρ and $x^{\rho, i}$ for $i \in \rho^{-1}(\star)$, but their parities are different in S .

Better Explicit Obstructions Imply Breakthrough Lower Bounds. Finally, we sketch how an explicit obstruction for $n^{2+\varepsilon}$ formulas implies strong formula lower bounds. The idea is that, if such an obstruction exists and $2^{O(n)}$ time computations have $2^{o(n)}$ -size formulas, then there is a formula F of only $n^{o(1)}$ size whose truth table is the output of the poly(n)-time algorithm printing the obstruction. Using linear hashing tricks, we can use $O(n^2)$ -size PARITY formulas to construct a formula of size $n^{2+o(1)}$ which agrees with the obstruction on all of its input-output pairs, a contradiction. For branching programs and B_2 -formulas, PARITY can be implemented with only $O(n)$ extra size, improving the threshold for those two cases.

Comparison With Some “Barrier” Results in Fine-Grained Complexity. We note a relationship between our results and some in fine-grained complexity. For some central problems in fine-grained complexity (see e.g., [2–5, 13, 17]), an n^c time algorithm is known, and it is

known that $n^{c-\varepsilon}$ time (even $n^c/(\log n)^8$ in some cases [3]) would imply a complexity breakthrough, via an abnormally-faster SAT algorithm for some circuit class. This is generally viewed as an impossibility result, because the SAT algorithms may not exist.

In the settings of this paper, we strongly **believe** that $n^{c+\varepsilon}$ bounds will hold for our tasks, and it is not outlandish to think they can be also be proved. However, our results show that achieving an $n^{c+\varepsilon}$ bound suffice for proving super-polynomial bounds. Both types of results can be interpreted in a similar light. The fine-grained results say that we know an algorithm runs in n^c time, but improving (reducing) that exponent will probably require radically new techniques (if you believe it can be improved at all). In our case, we know a lower bound of n^c , and improving (increasing) that exponent may also require radically new techniques.

Organization. In Section 2 we introduce the necessary preliminaries. In Section 3, we present the main technical construction of our paper: an $O(\log n)$ -seed length pseudorandom restriction generator for formulas of sub-quadratic size. In Section 4, we discuss sharp thresholds for quantified derandomization of probabilistic formulas. In Section 5, we present sharp thresholds for lower bounds against probabilistic formulas. In Section 6, we establish sharp thresholds for explicit obstructions against (deterministic) formulas and branching programs.

2 PRELIMINARIES

2.1 Notation

We use $\tilde{O}(f)$ as shorthand for $O(f \cdot \text{polylog}(f))$ throughout the paper. All logarithms are base-2. We use n to denote the number of input bits. We say a language $L \subseteq \{0, 1\}^*$ is $f(n)$ -sparse if $|L_n| \leq f(n)$, where $L_n = L \cap \{0, 1\}^n$. We assume knowledge of basic complexity theory (see [12, 26]).

We use U_ℓ to denote the uniform distribution over $\{0, 1\}^\ell$. The *statistical distance* between two random distributions X_1 and X_2 is

$$|X_1 - X_2| := \frac{1}{2} \sum_x |\Pr[X_1 = x] - \Pr[X_2 = x]|.$$

We consider De Morgan formulas (i.e. formulas with AND, OR, NOT gates). The *size* of a formula F , denoted by $L(F)$, is the number of leaves in F . For a Boolean function f we also use $L(f)$ to denote the size of the smallest formula that computing f . We say a formula is *read-once* if each input variable appears at most once as a leaf in the formula.

2.2 Definitions of MCSP and MKtP

The Minimum Circuit Size Problem (MCSP) [39] and the Minimum Kt Complexity Problem (MKtP, [6, 42]) are studied in this paper. We recall their definitions.

Definition 2.1 (MCSP). Let $s : \mathbb{N} \rightarrow \mathbb{N}$ satisfy $s(m) \geq m - 1$ for all m .

Problem: MCSP $[s(m)]$.

Input: A function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, presented as a truth table of $n = 2^m$ bits.

Decide: Does f have a (fan-in two) Boolean circuit C of size at most $s(m)$?

Recall that the Kt complexity (time-bounded Kolmogorov complexity) of string x is the smallest $c + \log(t)$, such that there is a Turing machine M of description length c that prints x in at most t steps.

Definition 2.2 (MKtP). Let $p : \mathbb{N} \rightarrow \mathbb{N}$ satisfy $p(n) \geq \log n$ for all n .

Problem: MKtP[$p(n)$].

Input: A string $x \in \{0, 1\}^n$.

Decide: Does x have Kt complexity at most $p(n)$?

In the Gap-MCSP[$f(n), g(n)$] problem (respectively, Gap-MKtP[$f(n), g(n)$]), we are given a string with circuit complexity (respectively, Kt complexity) which is either at least $g(n)$ or at most $f(n)$, and the goal is to distinguish between the two cases.

2.3 Useful Tools

We introduce some technical tools used in this paper. The most important tool in this paper is the explicit construction of (almost) k -wise independent distributions. We recall two explicit constructions of k -wise independent spaces.

THEOREM 2.3 (EXPLICIT k -WISE INDEPENDENT HASH FAMILY, [16]; SEE ALSO [64, COROLLARY 3.34]). For n, m, k , there is a family of k -wise independent functions $\mathcal{H} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ such that every function from $\mathcal{H}$ can be described in $k \cdot \max\{n, m\}$ random bits, and evaluating a function from $\mathcal{H}$ (given its description, and given an input $x \in \{0, 1\}^n$) takes time $\text{poly}(n, m, k)$.

THEOREM 2.4 (EXPLICIT k -WISE ϵ -DEPENDENT DISTRIBUTION, [10]). For $1 \leq k \leq n$ and $\epsilon > 0$, there is a distribution $\mathcal{R}$ on $\{0, 1\}^n$ which can be efficiently sampled using $O(k + \log \log n + \log(1/\epsilon))$ random bits, such that for every k positions $1 \leq i_1 < i_2 < \dots < i_k \leq n$, we have

$$|\mathcal{R}|_{i_1, i_2, \dots, i_k} - \mathbf{U}_k| \leq \epsilon,$$

where $\mathcal{R}|_{i_1, i_2, \dots, i_k}$ denotes the distribution of the k -bit string $x_{i_1}x_{i_2} \dots x_{i_k}$ for $x \sim \mathcal{R}$.

Moreover, the i -th coordinate of $x \sim \mathcal{R}$ can be computed in $\text{poly}(\log n, k, \log(1/\epsilon))$ time, given $i \in [n]$ and the seed as input.

We will also apply linear extractors in our paper. Here we recall the definition.

Definition 2.5 (Min entropy and extractors). The *min entropy* of a random variable X is the largest $k \in \mathbb{R}^+$ such that $\Pr[X = x] \leq 2^{-k}$ for every x in the range of X . A distribution over $\{0, 1\}^n$ with min entropy at least k is called an (n, k) -source.

A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) extractor, if for every (n, k) -source X ,

$$|\text{Ext}(X, \mathbf{U}_d) - \mathbf{U}_m| \leq \epsilon.$$

We say that Ext is a *linear extractor* if $\text{Ext}(\cdot, u)$ is a linear function over $\mathbb{F}_2$, for every $u \in \{0, 1\}^d$.

We also use random restriction methods in this paper; let us recall some notation for them. Fix $n \in \mathbb{N}$. Call a function $\rho : [n] \rightarrow \{0, 1, \star\}$ an n -variable restriction; a distribution $\mathcal{H}$ over n -variable restrictions is called a *random restriction*. We say $\mathcal{H}$ is p -regular if $\Pr_{\rho \sim \mathcal{H}}[\rho(i) = \star] = p$ and $\Pr_{\rho \sim \mathcal{H}}[\rho(i) = 0] = \Pr_{\rho \sim \mathcal{H}}[\rho(i) = 1] = (1-p)/2$ for every $i \in [n]$. We use $\mathcal{R}_p$ to denote the p -regular random restriction where all n coordinates are mutually independent.

Definition 2.6. We say $\mathcal{H}$ is p -regular k -wise δ -dependent, if for any k coordinates $i_1, \dots, i_k$ we have $|\mathcal{D}_{i_1, \dots, i_k} - \mathcal{R}_p| \leq \delta$, where $\mathcal{D}_{i_1, \dots, i_k}$ is the distribution of $(\rho(i_1), \dots, \rho(i_k))$ where $\rho \sim \mathcal{H}$. If $\delta = 0$, we say $\mathcal{H}$ is p -regular k -wise independent.

For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, we use $f|_\rho$ to denote the function $\{0, 1\}^{|\rho^{-1}(\star)|} \rightarrow \{0, 1\}$ obtained by restricting f according to ρ in the natural way. The composition $\rho = \rho_2 \circ \rho_1$ of two restrictions $\rho_1, \rho_2 : [n] \rightarrow \{0, 1, \star\}$ is defined as

$$\rho(i) := \begin{cases} \rho_1(i), & \text{if } \rho_1(i) \neq \star \\ \rho_2(i) & \text{otherwise} \end{cases}.$$

We observe the $\circ$ operator is associative, and $(f|_{\rho_1})|_{\rho_2} = f|_{\rho_2 \circ \rho_1}$.

3 PSEUDORANDOM RESTRICTION GENERATOR WITH LOGARITHMIC SEED LENGTH

In this section we construct our main technical component: a pseudorandom restriction generator for $O(n^{2-\epsilon})$ formulas with $O(\log n)$ seed length. We begin with the classical result that the shrinkage exponent of De Morgan formulas is 2.

THEOREM 3.1 (THE SHRINKAGE CONSTANT OF DE MORGAN FORMULAS IS 2 [29, 60]). Let f be a Boolean function. For every $p > 0$,

$$\mathbb{E}_{\rho \sim \mathcal{R}_p} [L(f|_\rho)] = O(p^2 L(f) + p\sqrt{L(f)}).$$

That is, when f is “hit” with a p -regular random restriction, we expect its formula complexity to shrink by about a p^2 factor. A uniform random p -regular restriction needs about $\Theta(pn \log(1/p))$ bits to encode. The main theorem of this section, stated below, provides a derandomized version of Theorem 3.1, with an optimal $O(\log n)$ seed length, albeit an $n^{o(1)}$ -size loss on the expected shrinkage size.

THEOREM 3.2. For all sufficiently large n , and probability parameters $p \in [20/n, 1]$, there is a family of pseudorandom restrictions $\rho : [n] \rightarrow \{0, 1, \star\}$ that is samplable using $O(\log n)$ random bits, such that the following properties hold for sufficiently large n and a universal constant $C > 0$:

(Shrinkage) For all n -variable formulas F ,

$$\mathbb{E}[L(F|_\rho)] \leq (p^2 L(F) + p\sqrt{L(F)}) \cdot n^{C/\log \log n}.$$

(“Close” to p -regular) $\Pr[|\rho^{-1}(\star)| \geq pn/2] > 2/3$.

(Constructive) There is a $\text{poly}(n)$ -time algorithm that outputs $\rho(1), \dots, \rho(n)$ given the seed s . There is also a $\text{polylog } n$ -size circuit that outputs $\rho(i)$, given index $i \in [n]$ and seed s as input.

Note that our random restriction ρ is not p -regular, but it is “close enough” on average: with $2/3$ probability, at least $p/2$ fraction of the coordinates are unrestricted. This turns out to be sufficient for us. Our proof of Theorem 3.2 will crucially exploit the following structure theorem for formulas, which decomposes a large formula into many small formulas.

LEMMA 3.3 (DECOMPOSITION OF FORMULAS [35, 60]). Let F be a formula over variables $X = \{x_1, \dots, x_n\}$, and $k \in \mathbb{N}$ be some parameter. There are $m \leq 1 + 36 \cdot L(F)/k$ formulas over X , denoted by $G_1, \dots, G_m$, each of size at most k , and a read-once formula F' of size m , such that $F'(G_1(x), \dots, G_m(x)) = F(x)$ for all $x \in \{0, 1\}^n$.

3.1 One-Stage Pseudorandom Restriction

We now describe one “stage” of our pseudorandom restriction generator. As stated in the introduction, our proof follows the iterated pseudorandom restriction paradigm [32, 35, 54]. We first need a small k -perfect hash family, as stated below.

LEMMA 3.4 (ADAPTED FROM [11, SECTION 4]). *For $1 \leq k \leq n$, there is a distribution $\mathcal{H}$ over functions $h : [n] \rightarrow [k^{20}]$ which can be efficiently sampled by $O(\log k + \log \log n)$ random bits, such that for every subset $S \subseteq [n]$ of size $|S| = k$,*

$$\Pr_{h \sim \mathcal{H}} [h(i) \text{ are distinct for all } i \in S] \geq 1 - 1/k^{18}.$$

PROOF. Denote $K := k^{20}$. Let $h' : [n \log K] \rightarrow \{0, 1\}$ be sampled from a $(2 \log K)$ -wise ε -dependent distribution (Theorem 2.4) with $\varepsilon := 1/K$, using $O(\log K + \log \log(n \log K) + \log(1/\varepsilon)) = O(\log k + \log \log n)$ random bits. The function $h : [n] \rightarrow [K]$ is derived from h' by grouping $\log K$ bits together. Then h is pairwise ε -dependent. For every $S \subseteq [n]$ of size k , by the union bound, the probability that there exist two different $i, j \in S$ with $h(i) = h(j)$ is at most $\binom{k}{2} \cdot (1/K + \varepsilon) \leq 1/k^{18}$. $\square$

Now we are ready to construct one stage of the pseudorandom restriction.

LEMMA 3.5 (ONE STAGE OF THE PSEUDORANDOM RESTRICTION). *Let $k \geq 2$ and $q := k^{-1/2}$. There is a pseudorandom restriction $\varphi : [n] \rightarrow \{0, 1, \star\}^n$, samplable using $O(\log k + \log \log n)$ random bits, with the properties:*

(Shrinkage) *For all n -variable formulas F ,*

$$\mathbb{E}[L(F|_{\varphi})] \leq O\left(q^2 L(F) + q\sqrt{L(F)}\right).$$

(Regular and Almost Pairwise Independent) *φ is q -regular pairwise $1/k^{18}$ -dependent.*

(Constructive) *There is a $(2^{O(k^8)} + \text{poly}(k, n))$ -time algorithm that outputs $\varphi(1), \dots, \varphi(n)$ given a seed s of length $O(\log k + \log \log n)$, and a $\text{poly}(k, \log n)$ -size circuit that outputs $\varphi(i)$, given an index $i \in [n]$ and seed s as input.*

PROOF. Let $K := k^{20}$, $q := k^{-1/2}$. Our proof works by first fixing an optimal pseudorandom restriction on K variables using non-uniform advice. We then show how to “amplify” that into the required pseudorandom restriction on n variables, by composing our small pseudorandom restriction with a random hash function from the k -perfect hash family of Lemma 3.4.

Construction of a “Micro” Pseudorandom Restriction on K Variables. Let $f : [K] \rightarrow [2/q]$ be a k -wise independent random function, samplable using $O(k \log k)$ random bits (Lemma 2.3), and let $g : [K] \rightarrow [2/q]$ be a pairwise independent random function, samplable using $O(\log k)$ random bits. Think of f and g as outputting $\log(2/q)$ -bit strings. Then $g \oplus f$ (where $\oplus$ denotes XOR on $\log(2/q)$ bits) is also a k -wise independent random function, from which we can obtain a q -regular k -wise independent random restriction $\rho : [K] \rightarrow \{0, 1, \star\}$. For example, given $g \sim g$, we can set $\rho(i) := \star$ if $g(i) \in \{0, 1\}$, and set $\rho(i) := g(i) \bmod 2$ if $g(i) \in [2/q] \setminus \{0, 1\}$. Let $F|_{g \oplus f}$ denote the restricted formula $F|_{\rho}$ where ρ is obtained from $g \oplus f$.

For any K -variable formula H of size $L(H) \leq k$, by the original shrinkage theorem (Theorem 3.1), we have

$$\begin{aligned} \mathbb{E}_f \mathbb{E}_g [L(H|_{g \oplus f})] &= O\left(q^2 L(H) + q\sqrt{L(H)}\right) \\ &\leq O\left(q\sqrt{L(H)}\right), \end{aligned}$$

where the last inequality follows from $q\sqrt{L(H)} \leq q\sqrt{k} = 1$.

From f , we draw S i.i.d. random samples $(f_1, \dots, f_S)$ for a parameter S to be determined later, and let $\mathcal{F}_S$ denote the uniform distribution over $(f_1, \dots, f_S)$. Then for any K -variable formula H of size $1 \leq L(H) \leq k$, the Chernoff-Hoeffding bound implies:

$$\begin{aligned} \Pr_{f_1, \dots, f_S} \left[\mathbb{E}_{f \sim \mathcal{F}_S} \mathbb{E}_g [L(H|_{g \oplus f})] \geq \mathbb{E}_f \mathbb{E}_g [L(H|_{g \oplus f})] + q\sqrt{L(H)} \right] \\ \leq \exp\left(-2S \cdot \left(\frac{q\sqrt{L(H)}}{k}\right)^2\right) \leq \exp\left(-\frac{2S}{k^3}\right). \end{aligned}$$

There are $K^{O(k)} \leq 2^{O(k \log k)}$ formulas H on K variables of size at most k . Set $S := O(k^5)$. Then a union bound implies that, with at least $2/3$ probability over the choice of $(f_1, \dots, f_S)$, for all K -variable formulas H of size at most k ,

$$\mathbb{E}_{f \sim \mathcal{F}_S} \mathbb{E}_g [L(H|_{g \oplus f})] = O\left(q\sqrt{L(H)}\right).$$

We observe that such an S -tuple $(f_1, \dots, f_S)$ can be found deterministically in $2^{O(k^8)}$ time. We simply enumerate all $2^{O(k \log k) \cdot S} \leq 2^{O(k^7)}$ possible $(f_1, \dots, f_S)$, and output the first tuple that satisfies the above property. To verify the property holds, we enumerate all $K^{O(k)} \leq k^{O(k)}$ relevant formulas H , and compute the formula size after shrinkage under each possible $g \oplus f$ by brute force. Recall that we define $\mathcal{F}_S$ to be the uniform distribution over the elements $f_1, \dots, f_S$ in this S -tuple.

Composing With the Hash Function. Let $\mathcal{H}$ be a family of perfect hash functions $h : [n] \rightarrow [K]$ as defined in Lemma 3.4. For a formula G of size $L(G) \leq k$ on the variable set $\{x_1, \dots, x_n\}$, we now consider its expected size under the random restriction defined by

$$\varphi(i) := \varphi_K(h(i))$$

for every $i \in [n]$, where $h \sim \mathcal{H}$ and $\varphi_K \sim g \oplus \mathcal{F}_S$. That is, we draw a function $g \sim g$ and a function $f \sim \mathcal{F}_S$, and let φ_K be the restriction obtained from the function $g \oplus f$ in the way we have defined above.

From the formula G , we obtain a new formula H on variables $\{y_1, \dots, y_K\}$ by replacing every leaf variable x_i of G with the variable $y_{h(i)}$. Since there are at most k leaves, by Lemma 3.4 says that with at least $1 - 1/k^{18}$ probability over the choice of $h \sim \mathcal{H}$, this replacement does not assign any distinct x_i 's to the same variable y_j . In this case, we have

$$\mathbb{E}_{\varphi_K \sim g \oplus \mathcal{F}_S} [L(G|_{\varphi_K \circ h})] = \mathbb{E}_{\varphi_K \sim g \oplus \mathcal{F}_S} [L(H|_{\varphi_K})] = O\left(q\sqrt{L(G)}\right).$$

Hence,

$$\begin{aligned} \mathbb{E}_{\varphi} [L(G|_{\varphi})] &= \mathbb{E}_{h \sim \mathcal{H}} \mathbb{E}_{\varphi_K \sim g \oplus \mathcal{F}_S} [L(G|_{\varphi_K \circ h})] \\ &\leq \left(1 - \frac{1}{k^{18}}\right) \cdot O\left(q\sqrt{L(G)}\right) + \frac{1}{k^{18}} \cdot L(G) = O\left(q\sqrt{L(G)}\right). \quad (1) \end{aligned}$$

For a formula F over $\{x_1, \dots, x_n\}$ of size $L(F) > k$, we apply the decomposition of Lemma 3.3. We obtain $m = O(L(F)/k)$ formulas $G_1, \dots, G_m$, each of size at most k , along with a read-once formula F' of size m such that $F'(G_1(x), \dots, G_m(x)) = F(x)$ for all $x \in \{0, 1\}^n$. By linearity of expectation, we have

$$\mathbb{E}[L(F|\varphi)] \leq \sum_{i=1}^m \mathbb{E}[L(G_i|\varphi)] \leq m \cdot O(q\sqrt{k}) = O(q^2 L(F)). \quad (2)$$

Combining (1) and (2) finishes the proof of the shrinkage property.

Note that $\varphi_K \sim \mathbf{g} \oplus \mathcal{F}_S$ can be specified using $O(\log k) + \log S = O(\log k)$ bits. Hence a restriction $\varphi = \varphi_K \circ h$ can be sampled by a seed s of $O(\log k + \log \log n)$ bits.

After computing $(f_1, \dots, f_S)$ in $2^{O(k^8)}$ time, a restriction $\varphi \in \{0, 1, \star\}^n$ can be computed in $\text{poly}(n, k)$ time given the seed s . After fixing a seed, we can also implement the function that maps indices i of $\log(n)$ bits to $\phi(i) \in \{0, 1, \star\}$ using a $\text{poly}(\log n, k)$ -size circuit (the tuple $(f_1, \dots, f_S)$ can be hardwired into the circuit).

Since $\mathbf{g}$ is pairwise independent, $\varphi_K \sim \mathbf{g} \oplus \mathcal{F}_S$ is also. (To see this, note that $\mathbf{g} \oplus f$ is pairwise independent for *any* fixed function f .) For two different $i_1, i_2 \in [n]$, with probability at least $1 - 1/k^{18}$ we have $h(i_1) \neq h(i_2)$, in which case $\varphi_K(h(i_1)), \varphi_K(h(i_2))$ are q -regular independent. Hence $\varphi \sim \varphi_K \circ h$ is q -regular pairwise $1/k^{18}$ -dependent. $\square$

3.2 Multi-stage Pseudorandom Restriction

Now we build our final pseudorandom restriction by composing the one-stage construction from Lemma 3.5.

Proof of Theorem 3.2. Assume $L(F) \geq 1$. Set $k := (\log n)^{1/10}$, $q := k^{-1/2}$, and $p := q^r$, where $r = \log(p^{-1})/\log(q^{-1}) \leq O(\log n / \log \log n)$. Let φ^i denote the distribution of the composed restriction

$$\varphi_i \circ \dots \circ \varphi_1,$$

where each φ_j ($1 \leq j \leq i$) is independently sampled from the one-stage restriction φ (Lemma 3.5), using $O(\log k + \log \log n) = O(\log \log n)$ bits. Our pseudorandom restriction ρ is defined as φ^r , samplable using $r \cdot O(\log \log n) = O(\log n)$ random bits. Since $2^{O(k^8)} = 2^{o(\log n)}$, each φ can be sampled in $\text{poly}(n)$ time. So ρ can also be sampled in $\text{poly}(n)$ time, or strongly explicitly by a $\text{polylog}(n)$ -size circuit (using the constructive property of the one-stage restriction).

Since $\mathbb{E}[L(F|\varphi)] \leq \frac{c}{2}(q^2 L(F) + q\sqrt{L(F)})$ for a constant $c > 1$, we have

$$\mathbb{E}[L(F|\varphi)] \leq cqL(F), \quad (3)$$

and

$$\mathbb{E}[L(F|\varphi)] \leq c(q^2 L(F) + 1). \quad (4)$$

By induction on (3),

$$\mathbb{E}[L(F|\varphi^w)] \leq c^w q^w L(F). \quad (5)$$

And by induction on (4),

$$\begin{aligned} \mathbb{E}[L(F|\varphi^u)] &\leq c^u q^{2u} L(F) + c(1 + cq^2 + c^2 q^4 + \dots + c^{u-1} q^{2u-2}) \\ &\leq c^u q^{2u} L(F) + 2c, \end{aligned} \quad (6)$$

for small enough q . Combining (5) and (6) with $w := r - u$, we have

$$\begin{aligned} \mathbb{E}[L(F|\rho)] &= \mathbb{E}_{\varphi^{r-u}} \mathbb{E}_{\varphi^u} [L(F|\varphi^{r-u} \circ \varphi^u)] \\ &\leq c^{r-u} q^{r-u} \cdot (c^u q^{2u} L(F) + 2c) \\ &\leq c^r \cdot q^r \cdot O(q^u L(F) + q^{-u}). \end{aligned}$$

Since $\min_{0 \leq u \leq r} \{q^u L(F) + q^{-u}\} \leq O(q^r L(F) + q^{-1} \sqrt{L(F)})$, we have

$$\begin{aligned} \mathbb{E}[L(F|\rho)] &\leq O\left(c^r q^{2r} L(F) + c^r q^{r-1} \sqrt{L(F)}\right) \\ &\leq (p^2 L(F) + p\sqrt{L(F)}) \cdot n^{O(1/\log \log n)}. \end{aligned}$$

Note that $\rho(i) = \star$ iff $\varphi_1(i) = \varphi_2(i) = \dots = \varphi_r(i) = \star$. Since $\varphi_1, \dots, \varphi_r$ are independent, and each of them is q -regular pairwise $1/k^{18}$ -dependent, for every $1 \leq i < j \leq n$, we have

$$\begin{aligned} \Pr[\rho(i) = \rho(j) = \star] &\leq (q^2 + 1/k^{18})^r \leq q^{2r} \exp(rq^{-2}/k^{18}) \\ &\leq p^2(1 + o(1)), \end{aligned}$$

and

$$\begin{aligned} \Pr[\rho(i) = \star] &\geq (q - 1/k^{18})^r \geq p(1 - o(1)), \\ \Pr[\rho(i) = \star] &\leq (q + 1/k^{18})^r \leq p(1 + o(1)). \end{aligned}$$

Hence,

$$np(1 - o(1)) \leq \mathbb{E}[|\rho^{-1}(\star)|] \leq np(1 + o(1)),$$

and

$$\begin{aligned} \text{Var}[|\rho^{-1}(\star)|] &= \mathbb{E}[|\rho^{-1}(\star)|^2] - \mathbb{E}[|\rho^{-1}(\star)|]^2 \\ &\leq np(1 - p + o(1)) + n^2 p^2 o(1). \end{aligned}$$

By Chebyshev's inequality,

$$\Pr[|\rho^{-1}(\star)| \geq np/2] \geq 1 - \frac{\text{Var}[|\rho^{-1}(\star)|]}{(np/2 - o(1))^2} > 2/3.$$

This completes the proof. $\square$

4 SHARP THRESHOLDS FOR QUANTIFIED DERANDOMIZATION

In this section we present sharp threshold results for quantified derandomization.

4.1 Threshold Theorems for Quantified Derandomization

To prove our threshold theorems for quantified derandomization, we need the following construction of linear extractors with a short seed.

THEOREM 4.1 (ADAPTATION OF [43, THEOREM 3.14]). *For all constant $\beta \in (0, 1)$ and sufficiently large n , there is an explicit (n^β, ε) linear extractor $\{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $\varepsilon = n^{-\beta}$, $d \leq (1 + O(\beta)) \log n$, and $m = n^{\beta/2}$.*

The theorem above is already implicit in [43]. We provide a proof in the full version of this paper.

We start with a lemma along the lines of Goldreich and Wigderson [28], showing that, given a poly-time algorithm for QD that finds any input on which the given probabilistic formula has non-zero probability of outputting 1, we can obtain a full derandomization.

LEMMA 4.2. *Suppose there is a constant $\alpha > 0$, and a deterministic poly-time algorithm that, given any generalized probabilistic formula $\mathcal{F}$ of size $n^{2+\alpha}$ with $B(n) = 2^{n^\alpha}$ exceptional inputs, finds an input $x \in \{0, 1\}^n$ on which $\mathcal{F}$ has non-zero probability of outputting 1. Then, there is a poly-time algorithm that finds an accepting assignment to any polynomial-size formula having at most $2^n/3$ falsifying assignments.*

PROOF. Our proof follows [28]; the idea is to apply a linear extractor appropriately. Let F be a (deterministic) polynomial size formula on m variables, which has at least $B(m) = 0.9 \cdot 2^m$ YES-inputs (we amplified the acceptance probability from $2/3$ to 0.9 by a standard argument). For simplicity, we assume the size of F is m , since otherwise we could add dummy variables so that m equals the size of F , without affecting the fraction of YES-inputs. Now we will deterministically find a YES-input to F , by reducing to a quantified derandomization instance.

Let β be a positive constant to be specified later. Let $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be the (n^β, ϵ) linear extractor from Theorem 4.1, where $m = n^{\beta/2}$, $\epsilon \leq 0.1$, and $d \leq (1 + O(\beta)) \log n$. We choose β small enough that $d \leq (1 + \alpha) \log n$, and $\beta \leq \alpha$.

Define the probabilistic formula $\mathcal{G} := F(E(x, U_d))$ on input $x \in \{0, 1\}^n$. Each output bit of $E(x, s)$ is the XOR of a subset of input bits, which can be computed by a n^2 -size formula. So for each $s \in \{0, 1\}^d$, the size of $F(E(x, s))$ is at most $n^2 \cdot m \leq n^{2+\alpha}$.

We claim that there are less than $2^{n^\beta} \leq 2^{n^\alpha}$ strings $x \in \{0, 1\}^n$ such that $\mathbb{E}[F(E(x, U_d))] \leq 2/3$. Assume not; then the uniform distribution X over all x such that $\mathbb{E}[F(E(x, U_d))] > 2/3$ has min-entropy at least n^β . Hence X is an (n, n^β) source, which implies that the statistical distance between $E(X, U_d)$ and U_m is at most 0.1 , since E is an $(n^\beta, 0.1)$ -extractor. Therefore, $\mathbb{E}[F(E(X, U_d))] \geq \mathbb{E}[F(U_m)] - 0.1 \geq 0.9 - 0.1 = 0.8$, contradicting $\mathbb{E}[F(E(x, U_d))] \leq 2/3$ for all x in the support of X .

Thus there are at most 2^{n^α} exceptional inputs on which the generalized formula $\mathcal{G}$ evaluates to **0** or **?**. By the hypothesis, on such a $\mathcal{G}$ we can deterministically find an input x where $\mathcal{G}$ has positive probability of outputting **1**. Then, by enumerating all $s \in \{0, 1\}^d$, we can find a YES-input $E(x, s)$ to F . $\square$

Reminder of Theorem 1.10. *Suppose there are $\epsilon > 0$, $\alpha \in (0, 1)$, and a poly-time algorithm for QD of generalized probabilistic formulas of size $n^{2-\alpha+\epsilon}$ with $B(n) = 2^{n^\alpha}$ exceptional inputs. Then, there is a poly-time algorithm finding a satisfying assignment to any poly-size formula with at most $2^n/3$ falsifying assignments. As a corollary, NP does not have n^k -size formulas for all k .*

Moreover, if the above algorithm is black-box, then there is a poly-time algorithm that, for every polynomial-size formula, additively approximates its acceptance probability up to any desired constant additive error.

PROOF. We show that, assuming the hypothesis of this theorem, we can achieve the hypothesis of Lemma 4.2, which in turn implies a (one-sided) derandomization of polynomial-size formulas with $2^n/3$ falsifying assignments. By [49, Theorem 1.1], this implies that NP does not have n^k -size formulas for all k .

Recall the requirement of Lemma 4.2: we need to show that there is a polynomial-time algorithm such that, on every generalized formula of size $n^{2+\alpha}$ with at most 2^{n^α} exceptions, the algorithm finds an input $x \in \{0, 1\}^n$ on which $\mathcal{F}$ has non-zero probability of outputting **1**.

Let $\beta > 0$ be a constant to be determined later, and let $\mathcal{F}(x) = \{(p_i, F_i(x))\}$ be a generalized probabilistic formula on m variables. In particular, each $F_i(x)$ is evaluated on input x with probability p_i . We assume each F_i has size $m^{2+\beta}$, and that there are at most 2^{m^β} exceptional inputs on which $\mathcal{F}$ evaluates to **0** or **?**.

Let $n := m^{1/t}$ for a constant $t \in (0, 1)$ to be determined. For simplicity, we assume m divides n . We construct a new probabilistic formula $\mathcal{G}(y)$ over n Boolean variables y as follows: denote $y = (x_1, x_2, \dots, x_{n/m})$ with $|x_j| = m$ for every $j \in [n/m]$, and define

$$\mathcal{G}(y) := \{(p_i, G_i(y))\}, \text{ where } G_i(y) := \bigvee_{j \in [n/m]} F_i(x_j).$$

Observe that, for every $j \in [n/m]$, $\mathbb{E}_{G_i \sim \mathcal{G}} G_i(y) \geq \mathbb{E}_{F_i \sim \mathcal{F}} F_i(x_j)$; that is, $\mathcal{G}$ is only more likely to output **1**. Thus there are at most $(2^{m^\beta})^{n/m} = 2^{n^{1-t+\beta t}}$ inputs on which $\mathcal{G}$ evaluates to **0** or **?**.

Now, $\mathcal{G}(y)$ is a probabilistic formula of size at most $n/m \cdot m^{2+\beta} = n^{1+t+\beta t}$. By assumption, we can find a satisfying assignment y to $\mathcal{G}$ in polynomial time, as long as $1+t+\beta t \leq 2-\alpha+\epsilon$ and $\alpha \geq 1-t+\beta t$. Setting $\beta := \alpha\epsilon/4$ and $t := (1-\alpha)/(1-\beta)$, both inequalities are satisfied. Hence by assumption, we can deterministically find a $y \in \{0, 1\}^n$ on which $\mathcal{G}$ has value **1** or **?**; that is, the probability $\mathcal{G}(y)$ outputs **1** is at least $1/3$. Recalling $y = (x_1, \dots, x_{n/m})$ for $x_i \in \{0, 1\}^m$, there is at least one x_i on which $\mathcal{F}$ outputs **1** with non-zero probability. The proof then follows from Lemma 4.2.

Moreover, if the QD algorithm is black-box, then the above algorithm immediately implies a hitting set generator for all polynomial-size formulas with at most $2^n/3$ falsifying assignments. By [28, Theorem 2.1] (which relies on [27]), this implies a *two-sided* derandomization algorithm for polynomial-size formulas, which allows us to estimate the acceptance probability of any polynomial-size formula up to any constant additive error. $\square$

Remark 4.3. *In fact, in the hypothesis of Theorem 1.10, we can assume that the generalized probabilistic formula has support size at most $n^{1-\alpha+\epsilon}$. To see this, we note that we can assume $\mathcal{F}$ has support size $m^{1+\beta}$ (which follows from the small seed length of the extractor used in proving Lemma 4.2). Then we observe that the support size of $\mathcal{G}$ is also $m^{1+\beta} = n^{t+\beta t} \leq n^{1-\alpha+\epsilon}$.*

PROOF SKETCH OF THEOREM 1.11. The proof follows precisely the same strategy as Theorem 1.10. The $\text{NE} \not\subseteq \text{NC}^1$ consequence follows from [66, 67]. $\square$

4.2 Quantified Derandomization of Probabilistic Formulas

Complementing the previous theorem, we can prove the following quantified derandomization for generalized probabilistic formulas.

Reminder of Theorem 1.9. *Let $1 \leq s(n) \leq n/20$. There is a polynomial time computable hitting set for generalized probabilistic*

formulas of size $n^{2-K/\log \log n}/s(n)^2$ with $B(n) = 2^{s(n)}$ exceptional inputs, where $K > 0$ is a universal constant.

PROOF. Let $\mathcal{F}$ be a generalized probabilistic formula with $B(n) = 2^{s(n)}$ exceptional inputs, where each formula F in its support has size $L(F) \leq n^{2-K/\log \log n}/s(n)^2$. Let ρ be the pseudorandom restriction guaranteed by Theorem 3.2 with parameter $p := 20s(n)/n$. Let $K := 4C$, where C is the universal constant from Theorem 3.2. Then by the shrinkage property of the pseudorandom restriction, for every formula F with $L(F) \leq n^{2-K/\log \log n}/s(n)^2$, we have

$$\mathbb{E}[L(F|_\rho)] \leq (p\sqrt{L(F)} + p^2L(F)) \cdot n^{C/\log \log n} \leq O(n^{-C/\log \log n}) \leq 1/18$$

for large enough n . Therefore

$$\mathbb{E}_{\rho \sim \mathcal{F}} \mathbb{E}[L(F|_\rho)] = \mathbb{E}_{F \sim \mathcal{F}} \mathbb{E}[L(F|_\rho)] \leq 1/18.$$

By Markov's inequality, with at least $2/3$ probability over the choice of ρ , we have $\mathbb{E}_{F \sim \mathcal{F}}[L(F|_\rho)] \leq 1/6$, which (again by Markov's inequality) implies $\Pr_{F \sim \mathcal{F}}[L(F|_\rho) < 1] \geq 5/6$. Note that when $L(F|_\rho) < 1$, $F|_\rho$ is simply a constant function.

By the “close to p -regular” property of the pseudorandom restriction, $\Pr_\rho[|\rho^{-1}(\star)| \geq np/2] \geq 2/3$. Hence there is a ρ from our pseudorandom restriction generator such that $|\rho^{-1}(\star)| \geq np/2 > 10s(n)$, and $\Pr_{F \sim \mathcal{F}}[F|_\rho \text{ is a constant}] \geq 5/6$. Let

$$p_0 := \Pr_{F \sim \mathcal{F}}[F|_\rho \text{ is the constant } 1].$$

Then for every $x \in \{0, 1\}^n$ consistent with ρ ,

$$\begin{aligned} \mathbb{E}_{F \sim \mathcal{F}}[F(x)] &= \mathbb{E}_{F \sim \mathcal{F}}[F(x)|F|_\rho \text{ is not constant}] \\ &\quad \cdot \Pr_{F \sim \mathcal{F}}[F|_\rho \text{ is not constant}] + 1 \cdot p_0 \\ &\in [p_0, p_0 + 1/6]. \end{aligned}$$

The number of such x is at least $2^{10s(n)}$, while the number of x that make $\mathcal{F}$ output $?$ and 0 is at most $2^{s(n)}$. Therefore $p_0 + 1/6 \geq 2/3$, and $\mathcal{F}$ on such an x has value either $?$ or 1 . $\square$

Finally, we observe that if the given probabilistic formula satisfies the standard BPP-style promise (for each input, it either outputs 1 with probability $\geq 2/3$, or outputs 0 with probability $\geq 2/3$), we can obtain a slightly better hitting set construction than that of Theorem 1.9.

THEOREM 4.4. *For $1 \leq b(n) \leq n/20$, there is a polynomial time computable hitting set for probabilistic formulas of size $n^{2-K/\log \log n}/b(n)$ with $B(n) = 2^{b(n)}$ exceptional inputs, where $K > 0$ is a universal constant.*

We provide a proof of this theorem in the full version of this paper.

5 SHARP THRESHOLDS FOR HARDNESS MAGNIFICATION

In this section we establish sharp threshold results for hardness magnification. We first prove our magnification theorems, and then establish corresponding lower bounds, which are just “an epsilon away” from the magnification results.

5.1 Magnification Theorems

We first establish the magnification theorems, which are an adaptation of the main results of [19]. We need the following theorem from [19].

THEOREM 5.1 ([19, THEOREM 3.6, LEMMA 5.1, LEMMA 5.2], ADAPTED). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a function of sparsity S_{sparse} , where $\log n \leq \log(S_{\text{sparse}}) \leq n^{1-\Omega(1)}$.*

There is a function $H : \{0, 1\}^{\Theta(\log S_{\text{sparse}})} \rightarrow \{0, 1\}$, computable in nondeterministic $O(n)$ time with one oracle query to f , and a randomized $O(n)$ -time algorithm O_{rand} computing f with error probability 0.01 using $O(\log n)$ random bits, $O(\log S_{\text{sparse}})$ bits of advice, and only $O(1)$ non-adaptive oracle queries to H . In particular, for each fixed choice of randomness, each bit of the queries of O_{rand} to H is a parity of a subset of the input x , and the algorithm O_{rand} simply returns the majority of the answers it receives from the oracles.

Moreover, if $f = \text{MKtP}[S_{\text{sparse}}]$, then we can assume $H \in E$, and we can assume $H \in (\Sigma_2P)^{\oplus P}$ if $f = \text{MCSP}[S_{\text{sparse}}]$.

Now we are ready to prove Theorem 1.3.

PROOF SKETCH OF THEOREM 1.3. The proof is a minor adaptation of the main results in [19]; we mainly have to observe that their proofs extend to superquadratic-size probabilistic formulas.

We first describe the proof of item (3). By Theorem 5.1, for a constant $\beta > 0$ and every 2^{n^β} -sparse NP language $L_\beta : \{0, 1\}^n \rightarrow \{0, 1\}$, there is an NP function $H : \{0, 1\}^\star \rightarrow \{0, 1\}$, and a (nonuniform) randomized algorithm O_{rand}^H with error probability 0.01 that can decide L_β using $O(1)$ non-adaptive oracle queries to H of length only $O(n^\beta)$.

Moreover, after we fix the randomness, each bit of each query to H has the form $\bigoplus_{i \in S} x_i$ for some subset $S \subseteq [n]$ of the input coordinates; i.e., each bit is a MOD-2 sum of some input bits. Since PARITY can be computed by De Morgan formulas of size $O(n^2)$, each bit of each query can be computed in $O(n^2)$ size.

Now, assume every problem in NP has n^k -size formulas, for some k . Then, the H -oracle on $O(n^\beta)$ -bit input can be computed by some $O(n^{\beta k})$ -size formula. Therefore the randomized algorithm O_{rand}^H deciding L_β can be implemented with a probabilistic formula of size $O(n^2) \cdot O(n^{\beta k}) \leq n^{2+\epsilon}$, by setting $\beta \leq \epsilon/2k$. This contradicts our assumption.

For Item (1) and (2), we can still apply Theorem 5.1, and use the “moreover” part. Note that, assuming $\oplus P \subset \text{NC}^1$ (respectively, $\text{EXP} \subset \text{NC}^1$), it follows that the H -oracle on $\tilde{O}(n^\alpha)$ -bit input can be computed by $n^{O(\alpha)}$ -size formulas, because the complexity class $(\Sigma_2P)^{\oplus P}$ collapses to NC^1 , under the assumption that $\oplus P \subset \text{NC}^1$. $\square$

We note that even slightly super-quadratic lower bounds would suffice.

PROOF SKETCH OF THEOREM 1.4. The proof is the same as the first two items of Theorem 1.3. Assuming the hypothesis $\oplus P \subset \text{NC}^1$ ($\text{EXP} \subset \text{NC}^1$, resp.), the H -oracle on $\tilde{O}((\log n)^d)$ -bit input can be computed by formulas of size $(\log n)^{O(d)} \ll (\log n)^{f(n)}$. $\square$

5.2 Nearly Matching Lower Bounds

Now we are ready to prove the nearly matching lower bounds. We first state the following connection between hitting sets for QD and sparse lower bounds. The promise problems Gap-MkTP and Gap-MCSP are defined in Section 2.2.

THEOREM 5.2 (HITTING SETS IMPLY SPARSE LOWER BOUNDS). *Let C be a circuit class that is closed under negation. Let $S \subseteq \{0, 1\}^n$ be a hitting set against C -circuit with $B(n)$ exceptional inputs.*

- *If every string $s \in S$ has Kt complexity at most $A(n)$, then $\text{Gap-MkTP}[A(n), \log B(n)] \notin C$.*
- *If every string $s \in S$ has circuit size at most $A(n)$, then $\text{Gap-MCSP}[A(n), \frac{\log B(n)}{10 \log \log B(n)}] \notin C$.*

PROOF. To prove the first item, suppose circuit $C \in C$ computes $\text{Gap-MkTP}[A(n), \log B(n)]$. Since there are at most $B(n)$ strings with Kt complexity $\leq \log B(n)$, C only accepts no more than $B(n)$ inputs. By the definition of S , there exists $s \in S$ such that C rejects s . However, every $s \in S$ has Kt complexity at most $A(n)$ and should be accepted by C , a contradiction.

The second item can be proved similarly. $\square$

Combined with Theorem 4.4, we immediately prove Theorem 1.5 as a corollary.

Proof of Theorem 1.5. Every string in the hitting set from Theorem 3.2 is generated by a $\text{poly}(n)$ time algorithm on a $O(\log n)$ -bit seed, so they all have Kt complexity $O(\log n)$. And, these strings all have circuit size $\text{polylog } n$. The proof then follows from Theorem 5.2. $\square$

5.3 Trade-off Between Sparsity and Formula Lower Bounds

There is a gap between the known formula lower bounds for sparse and non-sparse NP languages (near n^2 versus near n^3). We can prove a trade-off between lower bounds and sparsity, in particular for the MCSP problem on different circuit size parameters.

THEOREM 5.3. *Let $0 < \alpha < 1$. $\text{MCSP}[n^\alpha]$ on input length $n = 2^m$ does not have $n^{2+\alpha-\varepsilon}$ -size probabilistic formulas, for any $\varepsilon > 0$.*

Our proof is based on the pseudorandom restriction from Theorem 3.2 and the Local PRG against formulas [23]. We provide a proof in the full version of this paper.

6 SHARP THRESHOLDS FOR EXPLICIT OBSTRUCTIONS

In this section we give our construction and sharp threshold for explicit obstructions against De Morgan formulas.

Reminder of Theorem 1.14. *For a universal constant $K > 0$, there is an explicit obstruction against De Morgan formulas of size $n^{2-K/\log \log n}$.*

PROOF. We apply the pseudorandom restriction generator from Theorem 3.2 with probability parameter $p(n) := 20/n$. More precisely, we will produce a so-called *anti-checker* for the PARITY function. We use Q to denote the set of restrictions generated by the pseudorandom restriction generator; note that $|Q| \leq \text{poly}(n)$.

For each restriction $\rho \in Q$ and each $i \in \{0\} \cup [n]$, we define an n -bit input $x^{\rho, i}$:

$$\text{for } j \in [n], \quad x_j^{\rho, i} := \begin{cases} \rho(j) & \rho(j) \neq \star, \\ 1 & \rho(j) = \star \text{ and } j = i, \\ 0 & \rho(j) = \star \text{ and } j \neq i. \end{cases}$$

Note that $x^{\rho, 0}$ corresponds to filling in all $\star$'s of ρ with 0's. We set our obstruction to be

$$S := \{(x^{\rho, i}, \text{PARITY}(x^{\rho, i})) : \rho \in Q, i \in \{0, 1, \dots, n\}\}.$$

Clearly, S is consistent (but it may contain repetitions), and it can be constructed in $\text{poly}(n)$ time.

Now let F be a formula of size $N^{2-K/\log \log N}$. Let $K = 4C$. Applying Markov's inequality, there is some restriction ρ from the generator such that $|\rho^{-1}(\star)| \geq pn/2 = 10$, and $L(F|_\rho) < 1$, i.e., $F|_\rho$ is constant (we argue such a ρ exists in the proof of Theorem 1.9). Now, let i be such that $\rho(i) = \star$, and consider two inputs $x^{\rho, 0}$ and $x^{\rho, i}$ from S . Since $F|_\rho$ is constant, F outputs the same value on these two inputs, yet they have different parities, which completes the proof. $\square$

Now we turn to showing how a slightly improved explicit obstruction implies exponential-size formula lower bounds. We will need a way to efficiently construct a perfect linear hash function for a given input set.

LEMMA 6.1 (EFFICIENT LINEAR HASHING). *There is a deterministic algorithm which takes m distinct strings $s_1, \dots, s_m \in \{0, 1\}^n$ as input, and in $\text{poly}(mn)$ time outputs $t = O(\log m)$ strings $w_1, \dots, w_t \in \{0, 1\}^n$, such that $H(s_i) \neq H(s_j)$ for all $i \neq j$, where $H : \{0, 1\}^n \rightarrow \{0, 1\}^t$ is defined as*

$$H(x) = (\langle x, w_1 \rangle, \langle x, w_2 \rangle, \dots, \langle x, w_t \rangle).$$

The construction of this hash function is standard. We define the hash function as

$$H_v(x) := (E(x)_{v_1}, E(x)_{v_2}, \dots, E(x)_{v_t}) \in \{0, 1\}^t,$$

where $t = c \cdot \log m$ for a sufficiently large constant c , E is an asymptotically good linear error-correcting code (see e.g. [57]), and $v = (v_1, \dots, v_t)$ is an array of indices obtained from expander-walk sampling [25] using only $O(\log n + t) = O(\log(mn))$ bits of seed. The construction can be derandomized by enumerating all possible seeds. We provide a detailed proof of this lemma in the full version of this paper.

Reminder of Theorem 1.15. *The following are equivalent:*

- (1) *There is an $\alpha > 0$ and an explicit obstruction against formulas of size $n^{2+\alpha}$.*
- (2) *For all k , there is an explicit obstruction against formulas of size n^k .*
- (3) *There is an $\varepsilon > 0$ and a function in E that does not have $2^{\varepsilon n}$ -size formulas, even infinitely often.*

PROOF. (1) $\implies$ (3). By assumption, there is an explicit obstruction against $n^{2+\alpha}$ -size formulas, consisting of n^c input/output pairs, for some constants $\alpha > 0, c \geq 1$. Suppose for contradiction that E has $2^{\varepsilon n}$ size formulas, for all $\varepsilon > 0$.

We run the poly-time algorithm computing the obstruction set S on 1^n , then feed S to the algorithm of Lemma 6.1 to choose a perfect linear hash $H : \{0, 1\}^n \rightarrow \{0, 1\}^{dc \log n}$ that maps all n^c inputs from S to distinct strings, for some universal constant d .

For all n , define the language $L_n \subseteq \{0, 1\}^{dc \log n}$, where a string $y \in L_n$ if and only if there is an $x \in \{0, 1\}^n$ with $H(x) = y$ such that $(x, 1) \in S$. Define $L := \bigcup L_n$. Because the obstruction can be produced in polynomial time, and the hash H is computable in polynomial time, L can be decided in $\text{poly}(n^c)$ time. Hence $L \in E$. By assumption, L has formulas of size $2^{\varepsilon(dc \log n)} \leq n^{\alpha/2}$, for $\varepsilon > 0$ sufficiently small.

Now define an n -variable function F by $F(x) := L(H(x))$. Each output bit of H is a PARITY of a subset of n input bits, which can be implemented by a formula of size n^2 . Thus F has formula complexity at most $\leq n^{\alpha/2} \cdot n^2$. Because H maps all of S to distinct strings, the function F defined above agrees with all input/output pairs in S . This implies that F does not have formulas of size $n^{2+\alpha}$, a contradiction.

(3) $\implies$ (2). Let L be a function in E that does not have $2^{\varepsilon n}$ -size formulas, even infinitely often. In other words, for all but finitely many n , $L \cap \{0, 1\}^n$ has formula complexity greater than $2^{\varepsilon n}$.

For every $k > 1$, we can construct an explicit obstruction against n^k -size formulas as follows. Set $m := (2k/\varepsilon) \log n$. For all $x \in \{0, 1\}^m$, include the pair $(x0^{n-m}, L(x))$ in the obstruction, which can be computed in $\text{poly}(n^c)$ time. Now suppose a formula F agrees with all such input-output pairs. Then for all x , the formula $F(x0^{n-m})$ computes $L(x)$ on input length m , implying that F itself must have size at least $2^{\varepsilon m} = n^{2k}$.

(2) $\implies$ (1). Trivial. $\square$

Reminder of Theorem 1.16. *If there is an unbounded function $f(n)$ and an explicit obstruction against formulas of size $n^2 \cdot (\log n)^{f(n)}$, then $E \not\subseteq \text{NC}^1$.*

PROOF SKETCH. The proof is essentially the same as the ((1) $\implies$ (3)) case of Theorem 1.15. If we assume $E \subseteq \text{NC}^1$ (instead merely assuming that E has $2^{\varepsilon n}$ -size formulas), then the function F constructed has formula complexity at most $n^2(c \log n)^t$ for some constant t , which is smaller than $n^2 \cdot (\log n)^{f(n)}$ for sufficiently large n . But F agrees with all input/output pairs in the obstruction, a contradiction. $\square$

Reminder of Proposition 1.17. *There are explicit obstructions against B_2 -formulas of size $o(n)$, and branching programs of size $o(n)$.*

PROOF. Note that an $o(n)$ -size B_2 -formula (or $o(n)$ -size branching program) does not depend on all of its inputs. In such a case, our obstruction can simply be the set $S = \{(0^n, 0)\} \cup \{(0^{i-1}10^{n-i}, 1) \mid i \in [n]\}$. $\square$

It is an interesting question to ask: to what degree can this simple proposition be improved? One interesting open problem is whether we can get explicit obstructions against cn -size B_2 -formulas or branching programs, for every constant $c \geq 1$. We can show that improving the proposition in the exponent would imply a breakthrough:

Reminder of Theorem 1.18. *The following are equivalent:*

- There is an $\alpha > 0$ and an explicit obstruction against B_2 -formulas of size $n^{1+\alpha}$.
- For all k , there is an explicit obstruction against B_2 -formulas of size n^k .
- There is an $\varepsilon > 0$ such that E cannot be computed by $2^{\varepsilon n}$ -size formulas, even infinitely often.

The same equivalence also holds with “branching programs” in place of “ B_2 -formulas.”

PROOF SKETCH. The proof is analogous to the equivalence proved in Theorem 1.15. The key difference is that PARITY has $O(n)$ -size branching programs and B_2 -formulas. Therefore the function F implemented in that proof has B_2 -formula complexity at most $n^{1+\alpha/2}$, and branching program complexity at most $n^{1+\alpha/2}$. $\square$

By standard hardness-to-randomness connections, we also have:

PROPOSITION 6.2 ([37, 52, 59]). *E does not have $2^{\varepsilon n}$ -size formulas (respectively, branching programs) for some $\varepsilon > 0$ (even infinitely often) if and only if for all k , there is an $O(\log n)$ -seed PRG fooling n^k -size formulas (respectively, branching programs) that is computable in $\text{poly}(n)$ -time.*

Thus the problem of computing non-trivial explicit obstructions for these computational models is already equivalent to constructing pseudorandom generators.

REFERENCES

- [1] Scott Aaronson and Avi Wigderson. 2009. Algebrization: A New Barrier in Complexity Theory. *TOCT* 1, 1 (2009), 2:1–2:54. <https://doi.org/10.1145/1490270.1490272>
- [2] Amir Abboud and Arturs Backurs. 2017. Towards Hardness of Approximation for Polynomial Time Problems. In *ITCS*. 11:1–11:26. <https://doi.org/10.4230/LIPIcs.ITCS.2017.11>
- [3] Amir Abboud and Karl Bringmann. 2018. Tighter Connections Between Formula-SAT and Shaving Logs. In *ICALP*. 8:1–8:18. <https://doi.org/10.4230/LIPIcs.ICALP.2018.8>
- [4] Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. 2016. Simulating branching programs with edit distance and friends: or: a polylog shaved is a lower bound made. In *Proc. of the 48th STOC*. 375–388.
- [5] Amir Abboud and Aviad Rubinfeld. 2018. Fast and Deterministic Constant Factor Approximation Algorithms for LCS Imply New Circuit Lower Bounds. In *ITCS*. 35:1–35:14. <https://doi.org/10.4230/LIPIcs.ITCS.2018.35>
- [6] Eric Allender, Harry Buhrman, Michal Koucký, Dieter van Melkebeek, and Detlef Ronneburger. 2006. Preliminary version in FOCS’02. Power from Random Strings. 35, 6 (2006). Preliminary version in FOCS’02, 1467–1493. <https://doi.org/10.1137/050628994>
- [7] Eric Allender and Shuichi Hirahara. 2019. New Insights on the (Non-)Hardness of Circuit Minimization and Related Problems. *TOCT* 11, 4 (2019), 27:1–27:27. <https://doi.org/10.1145/3349616>
- [8] Eric Allender, Dhiraj Holden, and Valentine Kabanets. 2017. The Minimum Oracle Circuit Size Problem. *Computational Complexity* 26, 2 (2017), 469–496. <https://doi.org/10.1007/s00037-016-0124-0>
- [9] Eric Allender and Michal Koucký. 2010. Amplifying lower bounds by means of self-reducibility. *J. ACM* 57, 3 (2010), 14:1–14:36. <https://doi.org/10.1145/1706591.1706594>
- [10] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. 1990. Simple Constructions of Almost k -Wise Independent Random Variables. In *FOCS*. 544–553. <https://doi.org/10.1109/FSCS.1990.89575>
- [11] Noga Alon, Raphael Yuster, and Uri Zwick. 1995. Color-coding. *JACM* 42, 4 (1995), 844–856.
- [12] Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity - A Modern Approach*. Cambridge University Press. <http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264>
- [13] Arturs Backurs and Piotr Indyk. 2015. Edit Distance Cannot Be Computed in Strongly Subquadratic Time (unless SETH is false). In *Proc. of the 47th STOC*. 51–58.

- [14] Theodore P. Baker, John Gill, and Robert Solovay. 1975. Relativizations of the $P = ?NP$ Question. *SIAM J. Comput.* 4, 4 (1975), 431–442. <https://doi.org/10.1137/0204037>
- [15] Marco L. Carmosino, Russell Impagliazzo, Shachar Lovett, and Ivan Mihajlin. 2018. Hardness Amplification for Non-Commutative Arithmetic Circuits. In *33rd Computational Complexity Conference, CCC*. 12:1–12:16. <https://doi.org/10.4230/LIPIcs.CCC.2018.12>
- [16] J. Lawrence Carter and Mark N. Wegman. 1979. Universal classes of hash functions. *J. Comput. System Sci.* 18, 2 (1979), 143–154.
- [17] Lijie Chen, Shafi Goldwasser, Kaifeng Lyu, Guy N. Rothblum, and Aviad Rubinfeld. 2019. Fine-grained Complexity Meets $IP = PSPACE$. In *SODA*. 1–20. <https://doi.org/10.1137/1.9781611975482.1>
- [18] Lijie Chen, Shuichi Hirahara, Igor Carboni Oliveira, Ján Pich, Ninad Rajgopal, and Rahul Santhanam. 2020. Beyond Natural Proofs: Hardness Magnification and Locality. In *11th Innovations in Theoretical Computer Science Conference, ITCS*. 70:1–70:48. <https://doi.org/10.4230/LIPIcs.ITCS.2020.70>
- [19] Lijie Chen, Ce Jin, and Ryan Williams. 2019. Hardness Magnification for all Sparse NP Languages. In *FOCS*. 1240–1255.
- [20] Lijie Chen, Dylan M. McKay, Cody D. Murray, and R. Ryan Williams. 2019. Relations and Equivalences Between Circuit Lower Bounds and Karp-Lipton Theorems. In *CCC*. 30:1–30:21. <https://doi.org/10.4230/LIPIcs.CCC.2019.30>
- [21] Lijie Chen and Roei Tell. 2019. Bootstrapping results for threshold circuits “just beyond” known lower bounds. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 34–41. <https://doi.org/10.1145/3313276.3316333>
- [22] Ruiwen Chen, Rahul Santhanam, and Srikanth Srinivasan. 2016. Average-Case Lower Bounds and Satisfiability Algorithms for Small Threshold Circuits. In *31st Conference on Computational Complexity, CCC*. 1:1–1:35. <https://doi.org/10.4230/LIPIcs.CCC.2016.1>
- [23] Mahdi Cheraghchi, Valentine Kabanets, Zhenjian Lu, and Dimitrios Myrasiotis. 2019. Circuit Lower Bounds for MCSP from Local Pseudorandom Generators. In *ICALP*. 39:1–39:14. <https://doi.org/10.4230/LIPIcs.ICALP.2019.39>
- [24] Magnus Gausdal Find, Alexander Golovnev, Edward A. Hirsch, and Alexander S. Kulikov. 2016. A Better-Than- $3n$ Lower Bound for the Circuit Complexity of an Explicit Function. In *FOCS*. 89–98. <https://doi.org/10.1109/FOCS.2016.19>
- [25] David Gillman. 1998. A Chernoff bound for random walks on expander graphs. *SIAM J. Comput.* 27, 4 (1998), 1203–1220.
- [26] Oded Goldreich. 2008. *Computational complexity - a conceptual perspective*. Cambridge University Press.
- [27] Oded Goldreich, Salil Vadhan, and Avi Wigderson. 2011. Simplified derandomization of BPP using a hitting set generator. In *Studies in Complexity and Cryptography. Miscellaneous on the Interplay between Randomness and Computation*. Springer, 59–67. Preliminary version in *ECCC*, TR00-004, 2000.
- [28] Oded Goldreich and Avi Wigderson. 2014. On Derandomizing Algorithms That Err Extremely Rarely. In *Proc. of the 47th STOC*. 109–118. <https://doi.org/10.1145/2591796.2591808>
- [29] Johan Hästad. 1998. The Shrinkage Exponent of de Morgan Formulas is 2. *SIAM J. Comput.* 27, 1 (1998), 48–64. <https://doi.org/10.1137/S0097539794261556>
- [30] Shuichi Hirahara. 2018. Non-Black-Box Worst-Case to Average-Case Reductions within NP. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*. 247–258. <https://doi.org/10.1109/FOCS.2018.00032>
- [31] Shuichi Hirahara, Igor Carboni Oliveira, and Rahul Santhanam. 2018. NP-hardness of Minimum Circuit Size Problem for OR-AND-MOD Circuits. In *33rd Computational Complexity Conference, CCC*. 5:1–5:31. <https://doi.org/10.4230/LIPIcs.CCC.2018.5>
- [32] Shuichi Hirahara and Rahul Santhanam. 2017. On the Average-Case Complexity of MCSP and Its Variants. In *32nd Computational Complexity Conference, CCC*. 7:1–7:20. <https://doi.org/10.4230/LIPIcs.CCC.2017.7>
- [33] Shuichi Hirahara and Osamu Watanabe. 2016. Limits of Minimum Circuit Size Problem as Oracle. In *31st Conference on Computational Complexity, CCC*. 18:1–18:20. <https://doi.org/10.4230/LIPIcs.CCC.2016.18>
- [34] John M. Hitchcock and Aduri Pavan. 2015. On the NP-Completeness of the Minimum Circuit Size Problem. In *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS*. 236–245. <https://doi.org/10.4230/LIPIcs.FSTTCS.2015.236>
- [35] Russell Impagliazzo, Raghu Meka, and David Zuckerman. 2019. Preliminary version in *FOCS'12*. Pseudorandomness from Shrinkage. *J. ACM* 66, 2, Article 11 (April 2019). Preliminary version in *FOCS'12*, 16 pages. <https://doi.org/10.1145/3230630>
- [36] Russell Impagliazzo, Ramamohan Paturi, and Michael E. Saks. 1997. Size-Depth Tradeoffs for Threshold Circuits. *SIAM J. Comput.* 26, 3 (1997), 693–707. <https://doi.org/10.1137/S0097539792282965>
- [37] Russell Impagliazzo and Avi Wigderson. 1997. $P = BPP$ if E Requires Exponential Circuits: Derandomizing the XOR Lemma. In *STOC*. 220–229.
- [38] Kazuo Iwama and Hiroki Morizumi. 2002. An Explicit Lower Bound of $5n - o(n)$ for Boolean Circuits. In *MFCS*. 353–364. https://doi.org/10.1007/3-540-45687-2_29
- [39] Valentine Kabanets and Jin-Yi Cai. 2000. Circuit minimization problem. In *STOC*. 73–79. <https://doi.org/10.1145/335305.335314>
- [40] Ilan Komargodski, Ran Raz, and Avishay Tal. 2017. Improved average-case lower bounds for de Morgan formula size: Matching worst-case lower bound. *SIAM J. Comput.* 46, 1 (2017), 37–57.
- [41] Oded Lachish and Ran Raz. 2001. Explicit lower bound of $4.5n - o(n)$ for boolean circuits. In *Proc. of 33rd STOC*. 399–408. <https://doi.org/10.1145/380752.380832>
- [42] Leonid A. Levin. 1984. Randomness conservation inequalities; information and independence in mathematical theories. *Information and Control* 61, 1 (1984), 15–37. [https://doi.org/10.1016/S0019-9958\(84\)80060-1](https://doi.org/10.1016/S0019-9958(84)80060-1)
- [43] Xin Li. 2016. Improved Two-Source Extractors, and Affine Extractors for Poly-logarithmic Entropy. In *FOCS*. 168–177.
- [44] Richard J. Lipton and Ryan Williams. 2013. Amplifying circuit lower bounds against polynomial time, with applications. *Computational Complexity* 22, 2 (2013), 311–343. <https://doi.org/10.1007/s00037-013-0069-5>
- [45] Richard J. Lipton and Neal E. Young. 1994. Simple strategies for large zero-sum games with applications to complexity theory. In *STOC*. 734–740.
- [46] Dylan M. McKay, Cody D. Murray, and R. Ryan Williams. 2019. Weak lower bounds on resource-bounded compression imply strong separations of complexity classes. In *STOC*. 1215–1225. <https://doi.org/10.1145/3313276.3316396>
- [47] Moritz Müller and Ján Pich. 2020. Feasibly constructive proofs of succinct weak circuit lower bounds. *Ann. Pure Appl. Log.* 171, 2 (2020). <https://doi.org/10.1016/j.apal.2019.102735>
- [48] Ketan Mulmuley. 2011. On P vs. NP and geometric complexity theory: Dedicated to Sri Ramakrishna. *J. ACM* 58, 2 (2011), 5:1–5:26. <https://doi.org/10.1145/1944345.1944346>
- [49] Cody Murray and R. Ryan Williams. 2018. Circuit lower bounds for nondeterministic quasi-polynomial: an easy witness lemma for NP and NQP . In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 890–901. <https://doi.org/10.1145/3188745.3188910>
- [50] Cody D. Murray and R. Ryan Williams. 2017. On the (Non) NP -Hardness of Computing Circuit Complexity. *Theory of Computing* 13, 1 (2017), 1–22. <https://doi.org/10.4086/toc.2017.v013a004>
- [51] E. I. Nechiporuk. 1966. On a Boolean function. *Doklady of the Academy of Sciences of the USSR* 169, 4 (1966), 765–766. English translation in *Soviet Mathematics Doklady* 7:4, pages 999–1000.
- [52] Noam Nisan and Avi Wigderson. 1994. Hardness vs Randomness. *J. Comput. Syst. Sci.* 49, 2 (1994), 149–167. [https://doi.org/10.1016/S0022-0000\(05\)80043-1](https://doi.org/10.1016/S0022-0000(05)80043-1)
- [53] Igor Carboni Oliveira. 2019. Randomness and Intractability in Kolmogorov Complexity. In *ICALP*. 32:1–32:14. <https://doi.org/10.4230/LIPIcs.ICALP.2019.32>
- [54] Igor Carboni Oliveira, Ján Pich, and Rahul Santhanam. 2019. Hardness Magnification near State-Of-The-Art Lower Bounds. In *Proc. of the 34th CCC*. 27:1–27:29. <https://doi.org/10.4230/LIPIcs.CCC.2019.27>
- [55] Igor Carboni Oliveira and Rahul Santhanam. 2018. Hardness Magnification for Natural Problems. In *FOCS*. 65–76. <https://doi.org/10.1109/FOCS.2018.00016>
- [56] Alexander A. Razborov and Steven Rudich. 1997. Natural Proofs. *J. Comput. Syst. Sci.* 55, 1 (1997), 24–35. <https://doi.org/10.1006/jcss.1997.1494>
- [57] Daniel A. Spielman. 1996. Linear-time encodable and decodable error-correcting codes. *IEEE Trans. Information Theory* 42, 6 (1996), 1723–1731. <https://doi.org/10.1109/18.556668>
- [58] Aravind Srinivasan. 2003. On the approximability of clique and related maximization problems. *J. Comput. Syst. Sci.* 67, 3 (2003), 633–651. [https://doi.org/10.1016/S0022-0000\(03\)00110-7](https://doi.org/10.1016/S0022-0000(03)00110-7)
- [59] Madhu Sudan, Luca Trevisan, and Salil P. Vadhan. 2001. Pseudorandom Generators without the XOR Lemma. *J. Comput. Syst. Sci.* 62, 2 (2001), 236–266. <https://doi.org/10.1006/jcss.2000.1730>
- [60] Avishay Tal. 2014. Shrinkage of De Morgan Formulae by Spectral Techniques. In *Proc. of 55th FOCS*. 551–560. <https://doi.org/10.1109/FOCS.2014.65>
- [61] Roei Tell. 2018. Quantified derandomization of linear threshold circuits. In *Proc. of the 50th STOC*. 855–865.
- [62] Roei Tell. 2019. Improved Bounds for Quantified Derandomization of Constant-Depth Circuits and Polynomials. *Computational Complexity* 28, 2 (2019), 259–343. <https://doi.org/10.1007/s00037-019-00179-2>
- [63] B. A. Trakhtenbrot. 1984. A Survey of Russian Approaches to Perebor (Brute-Force Searches) Algorithms. *IEEE Ann. Hist. Comput.* 6, 4 (Oct. 1984).
- [64] Salil P. Vadhan. 2012. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science* 7, 1–3 (2012), 1–336.
- [65] Ingo Wegener. 1987. *The Complexity of Boolean Functions*. John Wiley & Sons Ltd.
- [66] Ryan Williams. 2013. Improving exhaustive search implies superpolynomial lower bounds. *SIAM J. Comput.* 42, 3 (2013), 1218–1244.
- [67] Ryan Williams. 2014. Nonuniform ACC circuit lower bounds. *Journal of the ACM (JACM)* 61, 1 (2014), 2.