



Indistinguishability Obfuscation, Range Avoidance, and Bounded Arithmetic

Rahul Ilango*

rilango@mit.edu

MIT

Cambridge, Massachusetts, USA

Jiafu Li

ljt714285@gmail.com

IIIS, Tsinghua University

Beijing, China

R. Ryan Williams†

rrw@mit.edu

MIT

Cambridge, Massachusetts, USA

ABSTRACT

The *range avoidance problem* (denoted by AVOID) asks to find a string outside of the range of a given circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, where $m > n$. Although at least half of the strings of length m are correct answers, it is not clear how to *deterministically* find one. Recent results of Korten (FOCS’21) and Ren, Wang, and Santhanam (FOCS’22) show that efficient deterministic algorithms for AVOID would have far-reaching consequences, including strong circuit lower bounds and explicit constructions of combinatorial objects (e.g., Ramsey graphs, extractors, rigid matrices). This strongly motivates the question: does an efficient deterministic algorithm for AVOID actually exist?

In this work, we prove under the existence of subexponentially secure *indistinguishability obfuscation* (*iO*) that polynomial-time deterministic algorithms for AVOID imply $\text{NP} = \text{coNP}$. Combining this with Jain, Lin, and Sahai’s recent breakthrough construction of *iO* from well-founded assumptions (STOC’21, EUROCRYPT’22), we provide the first plausible evidence that AVOID has no efficient deterministic algorithm. Moreover, we also prove the hardness of AVOID based on polynomially-secure *iO* and a weaker variant of the Nondeterministic Exponential Time Hypothesis (NETH).

Extending our techniques, we prove a surprising separation in *bounded arithmetic*, conditioned on similar assumptions. Assuming subexponentially secure *iO* and coNP is not infinitely often in AM , we show that AVOID has no deterministic polynomial-time algorithm even when we are allowed $O(1)$ queries to an oracle that can invert the given input circuit on an arbitrarily chosen m -bit string. It follows that the *dual Weak Pigeonhole Principle*, the combinatorial principle underlying AVOID, is not provable in Cook’s theory PV_1 . This gives (under plausible assumptions) the first separation of Cook’s theory PV_1 for polynomial-time reasoning and Jeřábek’s theory APC_1 for probabilistic polynomial-time reasoning.

CCS CONCEPTS

- Theory of computation → Complexity theory and logic; Proof complexity; Cryptographic primitives; Proof theory.

*Supported by an NSF Graduate Research Fellowship and NSF CCF-1909429.

†Supported by NSF CCF-1909429.



This work is licensed under a Creative Commons Attribution 4.0 International License.

STOC ’23, June 20–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9913-5/23/06.

<https://doi.org/10.1145/3564246.3585187>

KEYWORDS

circuit range avoidance, dual weak pigeonhole principle, indistinguishability obfuscation, bounded arithmetic

ACM Reference Format:

Rahul Ilango, Jiafu Li, and R. Ryan Williams. 2023. Indistinguishability Obfuscation, Range Avoidance, and Bounded Arithmetic. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC ’23), June 20–23, 2023, Orlando, FL, USA*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3564246.3585187>

1 INTRODUCTION

Given a circuit C mapping n -bit inputs to m -bit outputs, where $m > n$, at least half of the possible m -bit strings are never output by C . How efficiently can we find such a string? This meta-computational problem is known as RANGE AVOIDANCE:

Search Problem: RANGE AVOIDANCE (a.k.a. AVOID)

Input: A Boolean circuit C with n inputs, and $m > n$ outputs.

Output: A string $y \in \{0, 1\}^m$ such that for all $x \in \{0, 1\}^n$, $C(x) \neq y$.

There is a simple randomized algorithm for AVOID: a uniformly random $y \in \{0, 1\}^m$ will be outside the range of C with probability at least $1 - 2^n/2^m \geq 1/2$. Is there an efficient *deterministic* algorithm for AVOID?

This question is especially intriguing because it does not seem clear what the answer should be. Indeed, Ren-Wang-Santhanam [66] remark

“It is unknown whether $\text{AVOID} \in \text{FNP}$, $\text{AVOID} \in \text{FP}$, or their negations are implied by any plausible assumptions. As far as we know, we do not even have a good idea of what the ‘ground truth’ should be.”

(FNP and FP refer to the function versions of NP and P respectively.)

Under a plausible derandomization assumption, there is a deterministic polynomial-time algorithm for AVOID given access to an NP-oracle [49]. In randomized polynomial time with an NP oracle, one can repeatedly sample a uniformly random $y \in \{0, 1\}^m$ and verify it is not in the range with the NP oracle. This process can be derandomized, assuming E^{NP} requires $2^{\Omega(n)}$ -sized circuits [45]. In fact, Korten [49] shows that AVOID is in FP^{NP} if and only if E^{NP} does not have $2^o(n)$ -sized circuits infinitely often, so finding a deterministic algorithm for AVOID with a SAT oracle is equivalent to proving circuit lower bounds.

1.1 Background

Implications of Deterministic Algorithms for RANGE AVOIDANCE. Kleinberg, Korten, Mitropolsky, and Papadimitriou [44] initiated the study of AVOID (in their notation, AVOID is the problem α -EMPTY for $\alpha \geq 1$). They showed that various explicit construction problems can be reduced to AVOID, including the problem COMPLEXITY, of outputting a function with high A -oracle circuit complexity, given the truth table of the oracle A .¹

In follow-up work, Korten [49] convincingly demonstrated that deterministic algorithms for AVOID would have significant consequences for circuit complexity and combinatorics. For example, letting C be a $\text{poly}(2^\ell)$ -size circuit which takes as input descriptions of $2^{\ell/10}$ -size circuits and outputs their 2^ℓ -bit truth tables, solving RANGE AVOIDANCE on such C amounts to finding truth tables of *high* circuit complexity, a task that is widely believed to be solvable in deterministic polynomial time.² Korten extended this observation considerably, showing a deterministic AVOID algorithm would imply deterministic constructions of a variety of other objects (e.g., Ramsey graphs, extractors, rigid matrices) where a random choice suffices, but explicit constructions are longstanding open problems.

Ren-Santhanam-Wang [66] found further striking consequences of deterministic efficient AVOID algorithms. Among many other results, they show that a polynomial-time algorithm for NC_4^0 circuits (NC_k^0 denotes circuits in which each output only depends on k inputs) with stretch $m = n + n^{o(1)}$ would already yield functions in E that require circuits of depth $n^{1-o(1)}$, a major open problem in circuit complexity. Guruswami-Lyu-Wang [29] improve upon several reductions of Ren-Santhanam-Wang, and also show that RANGE AVOIDANCE is in fact solvable in deterministic polynomial time for NC_2^0 circuits. Gajulapalli-Golovnev-Nagargoje-Saraogi [24] show that a deterministic polynomial-time algorithm for AVOID on NC_3^0 circuits with $m = n + n^{2/3}$ implies breakthrough explicit constructions of rigid matrices. They also give deterministic polynomial-time algorithms for AVOID on NC_k^0 circuits with stretch $m \geq n^{k-1}/\log n$.

What to Believe? Arguments and Counterarguments. All the above results underscore the significance of finding nontrivial algorithms for RANGE AVOIDANCE and the importance of understanding how difficult RANGE AVOIDANCE really is. **Should we believe RANGE AVOIDANCE is in FP (the class of polynomial-time computable functions), or not?** To illustrate the depth of this question, we briefly consider some arguments and counterarguments.

From one point of view, it is natural to imagine a world in which $\text{AVOID} \in \text{FP}$. From Korten [49], we already know that if E^{NP} doesn't have subexponential-size circuits, then AVOID is in FP^{NP} . In light of this, it seems natural to believe that under the stronger assumption that E (without an NP oracle) doesn't have

¹Kleinberg, Korten, Mitropolsky, and Papadimitriou [44] also showed that an extremely low-stretch variant of AVOID, where one is given a circuit mapping from $[2^n - 1]$ to $[2^n]$, is NP-hard (their Theorem 1). However, because the stretch of this variant is exponentially small, the complexity of the problem is quite different: for example, this version can't be easily solved with randomness, as the total number of inputs and outputs only differ by one.

²In more detail, constructing a 2^n -bit truth table with circuit complexity $2^{\Omega(n)}$ in $\text{poly}(2^n)$ time is equivalent to showing that E requires $2^{\Omega(n)}$ circuit complexity, which is the main hypothesis powering the famous pseudorandom generators for $\text{BPP} = \text{P}$ [31, 70].

subexponential-size circuits (that is, the widely-believed conjecture $E \not\subseteq \text{io-SIZE}(2^{o(n)})$ [31]), one could show AVOID is in FP (without an NP oracle). Furthermore, standard methods from pseudorandomness [45] imply that there is a polynomial-time constructible *hitting set* for AVOID, assuming (for example) that E does not have subexponential-size SAT oracle circuits. That is, under plausible hypotheses, one *can* generate in FP a polynomial-size set $S_{s,m} \subseteq \{0,1\}^m$ such that for *every* circuit C of size s with m outputs, at least one $y \in S_{s,m}$ is not an output of C (see Appendix A of the full version). However, *checking* which y is a non-output apparently requires an NP oracle.

To add to these points, the existence of a randomized algorithm for AVOID seems to preclude a range of approaches to ruling out a deterministic (FP) algorithm for AVOID. For example, in Appendix B of the full version we present a barrier result *against* proving $\text{AVOID} \notin \text{FP}$ using standard *black-box* randomized Turing reductions, which exploits the fact that a random string is a correct answer with high probability. Still, the fact that AVOID has a fast randomized algorithm does not necessarily mean that we should believe it has a fast deterministic one. For an extreme example, one can easily sample strings of high Kolmogorov complexity *with randomness*, but one provably cannot do this deterministically at all: sufficiently long strings generated by fixed deterministic algorithms always have low Kolmogorov complexity [69, Chapter 6.4].

Indeed, one might believe AVOID $\notin \text{FP}$ because the opposite may seem “too good” to be true. The prior work mentioned above shows that, if $\text{AVOID} \in \text{FP}$, there are many interesting consequences for lower bounds and explicit constructions. However, as we expect all of those consequences to actually be *true*, these results alone don't give a strong argument that $\text{AVOID} \notin \text{FP}$. Rather, they indicate that the opposite may be hard to prove, as it would have significant consequences.

Another intuition for the difficulty of AVOID is that the generality of the problem allows for more power than merely generating varieties of hard functions and special combinatorial objects, each of which correspond to specific *structured* instances of AVOID. Solving AVOID for *all* circuits, even arbitrary ones whose descriptions may be very “scrambled” and complex, could be far more powerful than the generation of interesting mathematical objects.

In summary, it was entirely unclear whether RANGE AVOIDANCE should be solvable in deterministic polynomial time, or not. In this paper, we shall give evidence that is not, starting from the intuition of the previous paragraph.

Indistinguishability Obfuscation. Before stating our results, we take a quick detour to discuss one of our main tools: *Indistinguishability Obfuscation (iO)*, a notion first defined by Barak *et al.* [6, 7]. Roughly speaking, an *iO* is a polynomial-time probabilistic algorithm that given a circuit C , outputs an “obfuscated” circuit $iO(C)$ computing the same function. The security guarantee of *iO* is that, for any two circuits C and C' of the same size that compute the same function, $iO(C)$ and $iO(C')$ are computationally indistinguishable to a class of “adversaries” (for example, polynomial-sized circuits). See Section 2.1 for a formal definition.

For many years, *iO* had the dual deficiency of neither having any candidate constructions, nor having particularly interesting applications. However, in the two decades since its definition, both

of these statements have seen dramatic reversals. While iO 's security definition initially seems weak (as it only gives a guarantee about circuits computing the same function), it turns out to be extremely powerful. We now know that nearly every cryptographic primitive (e.g., one-way functions [48], public-key encryption [67], multi-party non-interactive key exchange [10], etc.) can be constructed assuming iO exists and NP is not in BPP infinitely often (see Section 1.3 of Jain-Lin-Sahai [33] for a more comprehensive list).

Similarly, while many initial candidate constructions required new assumptions that were later broken, groundbreaking work of Jain, Lin, and Sahai [33] showed that iO exists assuming four “well-founded” assumptions (this has been improved to three assumptions, in [35]). Moreover, the iO they construct has strong security properties: no polynomial-sized circuit adversary can distinguish the iO of equivalent circuits except with subexponentially small probability. We refer to iO with these security properties as *JLS-security* (a formal definition is in Section 2.1). Following the work of Jain, Lin, and Sahai, iO has now become a widely-believed assumption in cryptography (see for example the recent Quanta article of Klarreich [43]).

1.2 Our Results

In this paper, we give the first concrete evidence that RANGE AVOIDANCE is hard to solve deterministically when the number of outputs $m(n) = \text{poly}(n)$. In particular, our conditional lower bound for RANGE AVOIDANCE follows from indistinguishability obfuscation and various forms of $\text{NP} \neq \text{coNP}$. Our argument is quite general in that it holds for a variety of parameters with trade-offs on the assumptions, but we state a simple version first.

Theorem 1. *Assume that $\text{NP} \neq \text{coNP}$ and iO with JLS-security exists. Then for all $c \geq 1$, there is a $k \geq c$ such that there is no deterministic polynomial-time algorithm for AVOID on n^k -size circuits with n inputs and $m(n) = n^c$ outputs.*

That is, assuming $\text{NP} \neq \text{coNP}$ and JLS-secure iO exists, there are no efficient deterministic algorithms for AVOID, even if the number m of output bits is allowed to be an arbitrarily large polynomial in n . Note that when m is an arbitrarily large polynomial in n , all but an exponentially small fraction (2^{-m+n}) of length- m strings are outside of the range of C . Interestingly, the hard instances in our proof are circuits C with at most *two elements* in their range!³ In fact, one of those two elements is always the string 0^m .

Before we discuss extensions of Theorem 1, let us briefly motivate how the assumptions in Theorem 1 arise. Suppose a deterministic algorithm for AVOID exists, and one is aiming to show a contradiction. What can you do with a deterministic algorithm, that you could not do with a randomized algorithm (which we know exists)? With a deterministic efficient algorithm for AVOID, one can guarantee that for every circuit C , there is a short “proof” that a *specific* string y_C is outside the range of C . In particular, the computation history of the deterministic AVOID algorithm running on C and outputting y_C , constitutes such a “proof.” In contrast, it is unclear how to get such a guarantee from the simple randomized algorithm

³Having a range of only two elements is best-possible, in a sense. For circuits C with only one string in their range, there is a simple AVOID algorithm: output any m -bit string different from $C(0^m)$.

for AVOID that picks a string uniformly at random. A priori, it seems powerful that every circuit C has a short proof that a specific y_C is outside its range. If the description C was complex enough to function like a “black box,” then the shortest proof that y_C is not in the range may simply evaluate C on all inputs and observe that C never outputs y_C . Thus, at a very high level, a deterministic AVOID algorithm may provide short proofs for statements that might not have short proofs (motivating an assumption like $\text{NP} \neq \text{coNP}$) for circuits that behave like black boxes (motivating an assumption like iO). Of course, this is a very rough intuition; for more details we point the reader to the (relatively short) proof of Theorem 1 in Section 3. We remark that our work is certainly not the first time that iO is being applied in complexity theory (see Section 2.1 for references).

Extensions. We now discuss extensions of Theorem 1, which illustrate various tradeoffs between the time complexity of AVOID, simulations of coNP with nondeterminism, and the allowed stretch m . A more general version of our result rules out subexponential-time ($2^{n^{o(1)}}$ -time) deterministic algorithms for AVOID, as well as algorithms for AVOID where the number of outputs m can be subexponential in n . To rule out subexponential-time AVOID algorithms, we apparently require a stronger (but still standard) notion of security for iO than that of JLS-security. In particular, we need that no subexponential-size circuit adversary can distinguish the iO of equivalent circuits, except with subexponentially small probability. We refer to iO with this security as *subexponentially-secure iO* . We point the reader to Section 2.1 for formal definitions, but we stress that subexponentially-secure iO is very plausible and holds if, for example, corresponding security guarantees hold for the three well-founded assumptions used by Jain-Lin-Sahai [34].

In the statement below (and throughout this paper) we always assume the number of outputs $m(n)$ is a time-constructible function.

Theorem 2. *Assume subexponentially-secure iO exists. For every $m(n) > n$ there is an $s(n) = \text{poly}(m)$ such that, if there is a deterministic t -time algorithm for AVOID circuits with $m(n)$ outputs and size $s(n)$, then*

$$\text{coNP} \subseteq \bigcup_{k \in \mathbb{N}} \text{NTIME}[t(m^k(n))]$$

Here, $\text{NTIME}[t]$ refers to the set of languages computable by a nondeterministic Turing machine in time $O(t)$. Setting $t = 2^{n^{o(1)}}$, we conclude that RANGE AVOIDANCE cannot be solved in deterministic $2^{s^{o(1)}}$ time on circuits of size s , assuming subexponentially-secure iO and a rather weak exponential-time hypothesis. Let $\text{NSUBEXP} := \bigcap_{k \in \mathbb{N}} \text{NTIME}[2^{n^{1/k}}]$.

Corollary 3. *Assuming $\text{coNP} \not\subseteq \text{NSUBEXP}$ and subexponentially-secure iO , RANGE AVOIDANCE cannot be solved in deterministic $2^{s^{o(1)}}$ time on circuits of size s .*

The assumption $\text{coNP} \not\subseteq \text{NSUBEXP}$ is significantly weaker than what is often assumed in fine-grained complexity. For example, the Nondeterministic Exponential Time Hypothesis (NETH) [17, 18] states that there is an $\varepsilon > 0$ such that unsatisfiable 3SAT instances with n variables cannot be refuted in nondeterministic 2^{en} time. Note NETH is a much stronger hypothesis than $\text{coNP} \not\subseteq \text{NSUBEXP}$.

Setting $m(n)$ to be any constructible function that is $2^{n^{o(1)}}$ in Theorem 2, we conclude that even when the number of outputs m is close to exponential in n , AVOID still does not have a polynomial-time deterministic algorithm, assuming $\text{coNP} \not\subseteq \text{NSUBEXP}$ and subexponentially secure iO .

Corollary 4. *Assuming $\text{coNP} \not\subseteq \text{NSUBEXP}$ and subexponentially-secure iO , RANGE AVOIDANCE cannot be solved in deterministic polynomial-time for every constructible $m(n) = 2^{n^{o(1)}}$.*

We also prove that iO with security weaker than JLS-security or subexponential security can still be used to show the non-existence of deterministic algorithms for AVOID. We show $\text{AVOID} \notin \text{FP}$ assuming only polynomially-secure iO and a still weaker statement than the typical NETH.

Hypothesis 5 (NETH for Circuits). *There is an $\epsilon > 0$ such that Circuit Unsatisfiability on n -input circuits of size $2^{o(n)}$ cannot be solved nondeterministically in $2^{\epsilon n}$ time.*

Theorem 6. *Assuming NETH for Circuits and polynomially-secure iO , AVOID is not in FP. Moreover, under the assumptions, it follows that for all $b, c \geq 1$ there is an $\epsilon > 0$ such that AVOID cannot be solved in $O(2^{c\epsilon n})$ time on circuits of size $2^{\epsilon n}$ with n inputs and bn outputs.*

Note that, in the case of circuits with $2^{o(n)}$ size and bn outputs, exhaustive search solves RANGE AVOIDANCE in time about $2^{bn+o(n)}$. Under our assumptions, we rule out $2^{o(n)}$ time for $2^{o(n)}$ -size circuits, when the number of outputs m is linear in n . Therefore, even subexponential-time improvements over exhaustive search for RANGE AVOIDANCE should already be considered unlikely for these parameters, under our assumptions.

Witness Encryption Suffices. A careful inspection of our proofs reveals that our use of indistinguishability obfuscation can be replaced with a seemingly weaker cryptographic primitive called *witness encryption*. At a high level, witness encryption [26] allows one to, given a SAT formula φ , encrypt a message m such that only recipients who know a satisfying assignment to φ can decrypt m (the actual definition is more involved, but this is the basic idea). It is known that iO implies witness encryption as a special case [25]. We state our results in terms of iO instead of witness encryption for two reasons. First, currently the only known way to construct witness encryption with our desired parameters under well-founded assumptions is via iO . Second, this paper is aimed primarily at a complexity-theoretic audience, who is likely more familiar with the notion of iO than witness encryption.

Nevertheless, witness encryption is believed to be a weaker assumption compared to iO , and admits several plausible constructions. Chen, Vaikuntanathan, and Wee [19] proposed a simple construction from LWE-like problems, whose security was proved later by Vaikuntanathan, Wee, and Wichs [72] based on LWE-like assumptions. An alternative LWE-based construction was proposed by Tsabary [71] under similar assumptions. Barta, Ishai, Ostrovsky, and Wu [8] also gave a construction in “generic group model” based on an (unproven) hardness of approximation hypothesis of certain coding problems.

Application: Separations in Bounded Arithmetic. Bounded arithmetic refers to fragments of Peano arithmetic that aim to formalize the

(computational) complexity of reasoning. For instance, Cook’s theory PV_1 [22] corresponds to “reasoning in polynomial time”, Jeřábek’s theory APC_1 [36–38, 40]⁴ corresponds to randomized polynomial-time computation, and Buss’s theories S_2^1, S_2^2, \dots correspond to the polynomial-time hierarchy [12]. Indeed, theories corresponding to other complexity classes such as TC^0, NC^1 , and PSPACE have been studied (see, e.g., [21, 51]). From a proof complexity point of view, bounded theories can also be regarded as uniform versions of propositional proof systems (see, e.g., [53]).

One motivation to study bounded arithmetic is that it may explain why longstanding complexity-theoretic conjectures such as $\text{P} \neq \text{NP}$ and $\text{NEXP} \not\subseteq \text{P/poly}$ are hard to prove. In contrast to barriers such as relativization [5], natural proofs [65], and algebrization [2] that capture the limitation of specific techniques, it is more desirable to demonstrate the unprovability of these conjectures in *strong mathematical theories* with a solid logical foundation. Bounded arithmetic provides an ideal testing ground for this program. On the one hand, a rather large fragment of known algorithms and complexity theory results can be formalized in bounded theories such as PV_1 and APC_1 [58, 60, 61, 63]. On the other hand, connections between bounded arithmetic and complexity theory make it possible to employ complexity-theoretic techniques to obtain unprovability results, leading to exciting developments on the unprovability of complexity upper bounds [14–16, 56] and lower bounds [52, 62, 64] in PV .

To better understand the power of feasible reasoning, it is important to prove relations (separations or equivalences) among bounded theories. In particular, it has been open for about twenty years whether the *dual Weak Pigeonhole Principle* for PV functions⁵ (denoted by $\text{dWPHP}(\text{PV})$) is provable in PV_1 . In other words, the question is whether Jeřábek’s APC_1 , defined as $\text{PV}_1 + \text{dWPHP}(\text{PV})$, is the same as PV_1 . Here, $\text{dWPHP}(\text{PV})$ is the “logic version” of AVOID ⁶ that says for every PV function f and every n and z , there is a $y \in \{0, 1\}^{n+1}$ such that for every $x \in \{0, 1\}^n$, $f(z, x) \neq y$. Similar to the complexity of AVOID, there has been no strong evidence for either $\text{APC}_1 = \text{PV}_1$ and $\text{APC}_1 \neq \text{PV}_1$. Known results on this problem include the separation of the *relativized* versions of APC_1 and PV_1 [39] and a conditional separation based on the assumption that $\text{P} \subset \text{SIZE}[n^k]$ for some k [54]. (However, this assumption contradicts the widely-believed derandomization assumption that E requires circuits of size $2^{\Omega(n)}$.) Krajíček [55] recently proposed an open problem of showing a conditional separation of PV_1 and APC_1 under a “mainstream hypothesis” as the first step to understand the logical power of $\text{dWPHP}(\text{PV})$. Moreover, it might seem reasonable to believe that APC_1 is the same as PV_1 , since its complexity-theoretic counterpart $\text{BPP} = \text{P}$ follows from plausible circuit lower bounds [31, 70].

In this work, we provide the first plausible evidence that Jeřábek’s theory APC_1 is a *strict* extension of PV_1 .

Theorem 7 (Corollary 23, Informal). *Assuming iO with JLS-security and coNP is not infinitely often in AM , the dual Weak Pigeonhole*

⁴Note the terminology APC_1 was first used in [13].

⁵By Cobham’s characterization of the polynomial-time functions, PV functions (when interpreted in the standard model) are exactly polynomial-time computable functions (see Section 2.3 for details).

⁶Indeed, Korten’s investigation of the complexity of explicit constructions was inspired by the early developments of Jeřábek’s theory APC_1 , as noted in [49, 50].

Principle is not provable in PV (in particular, APC₁ is a strict extension of PV₁).

Our proof utilizes the standard KPT Witnessing Theorem (see Theorem 18) which extracts an algorithm for Avoid that is allowed to *invert* the input circuit on a *constant* number of m -bit strings, assuming APC₁ is the same as PV₁. Corollary 23 then follows from an extension of the conditional lower bound for Avoid which holds against polynomial-time algorithms with such circuit-inversion oracles.⁷

Theorem 8 (Theorem 21, Informal). *Assuming the existence of iO with JLS-security and NP is not infinitely often in AM, there is no polynomial-time algorithm for AVOID with $O(1)$ queries to a circuit-inversion oracle.*

This conditional lower bound stands in interesting contrast to the fact that, under standard derandomization hypotheses (E does not have $2^{\Omega(n)}$ -size circuits), there is a deterministic polynomial-time algorithm for Avoid that makes *polynomially-many* circuit inversion queries (follows from [49]; see Appendix C of the full version). Alternatively, under a stronger assumption, namely E does not have $2^{\Omega(n)}$ -size SAT-oracle circuits, we can construct a hitting set of size $\text{poly}(n)$ (see Appendix A of the full version) and then find a non-output of the given circuit with $\text{poly}(n)$ circuit-inversion queries.

Under similar assumptions, we also demonstrate a separation of APC₁ and its fragment UAPC₁ that is strong enough to prove interesting results in complexity theory and formalize approximate counting in Jeřábek's framework [38], see Section 4.3 of the full version for more details.

Application: the Oracle Derandomization Hypothesis for Time-Bounded Kolmogorov Complexity. Our results also have bearing on other hypotheses regarding derandomization. Motivated by applications in parameterized complexity and questions related to “instance compressibility,” Fortnow and Santhanam [23] introduced the Oracle Derandomization Hypothesis (ODH). Roughly speaking, ODH says that, given a length- n truth table z of an arbitrary Boolean function, one can efficiently deterministically generate a truth table y of length at least $n^{.01}$ such that the function represented by y has circuit complexity $n^{\Omega(1)}$ even when the circuits are given oracle access to the function represented by z . (See Hypothesis 24 for a formal definition, and Section 5 for a comparison to the related problem COMPLEXITY.)

This hypothesis is especially intriguing because it is unclear whether it should be true or false. Indeed, Fortnow and Santhanam remark:

“In our opinion, quite apart from its relevance to compressibility, the Oracle Derandomization Hypothesis is interesting in its own right because it tests our intuitions of which kinds of derandomization are plausible and which are not... We do not have a strong belief

⁷The drawback of the lower bound result comparing with Theorem 2 is that the assumption $\text{NP} \neq \text{coNP}$ is strengthened to $\text{coNP} \not\subseteq \text{i.o. AM}$ due to technical reasons. Intuitively, we need the Goldwasser-Sipser protocol in AM (see Lemma 13) for approximate counting, and we only know how to “eliminate” oracle queries in the range avoidance algorithm assuming infinitely often lower bounds for coNP.

about the truth of our derandomization assumption, but we do believe it is hard to refute.”

Using essentially the same proof as in Theorem 1, we rule out a related *time-bounded Kolmogorov complexity* version of ODH (formally, Hypothesis 25) under plausible assumptions. Roughly speaking, Hypothesis 25 says that given a string z of length n , one can efficiently deterministically generate a string y of length $n^{.01}$ such that the conditional polynomial-time bounded Kolmogorov complexity of y given z is at least $n^{\Omega(1)}$. Here, the conditional t -time bounded Kolmogorov complexity of y given z refers to the length of the shortest program that outputs x on input y in time $t(|x|)$ (see Section 2.4 for a formal definition) [46, 47, 68].

Theorem 9 (Informal version of Theorem 26). *The time-bounded Kolmogorov complexity ODH is false assuming $\text{NP} \neq \text{coNP}$ and subexponentially secure iO exists.*

It is a tantalizing open question as to whether one can extend this result to rule out ODH itself. To prove Theorem 9 we crucially make use of the fact that, in this version of ODH, the computational model is able to read the entire string z . In contrast, in the (original) ODH setting, the computational model is only allowed to make a limited number of queries to the string z .

2 PRELIMINARIES

We assume basic familiarity with notions in computational complexity [3]. We first review two extensively used tools: Indistinguishability Obfuscation and interactive (Arthur-Merlin) protocols. We also provide a brief introduction on bounded theories PV₁ and APC₁ (see [21, 51, 53] for more detailed expositions), as well as the time-bounded Kolmogorov complexity.

2.1 Indistinguishability Obfuscation

Definition 10 (Indistinguishability Obfuscation). A polynomial-time randomized algorithm iO that takes as input a security parameter λ and a circuit C , and randomness r is an indistinguishability obfuscator with security (S, ϵ) if both of the following hold:

- **Perfect Functionality:** For all C and λ , $iO(1^\lambda, C)$ outputs a circuit computing the same function as C with probability one over its randomness.
- **Indistinguishability:** For all λ , any two circuits C and C' of size at most λ computing the same function, and any $S(\lambda)$ -sized adversary circuit A , we have that

$$|\Pr[A(iO(1^\lambda, C) = 1] - \Pr[A(iO(1^\lambda, C') = 1)]| \leq \epsilon(\lambda)$$

When λ is clear from the context, we write $iO(C)$ instead of $iO(1^\lambda, C)$. When we want to specify the randomness r used by the iO algorithm, we write $iO(1^\lambda, C; r)$.

We say an iO is *polynomially-secure* if it is secure for some $S(n) = n^{\omega(1)}$ and $\epsilon(n) < 1/n^{\omega(1)}$. We say it is *subexponentially-secure* if it is secure with $S(n) = 2^{n^\delta}$ and $\epsilon(n) = 2^{-n^\delta}$ for some $\delta > 0$. We say it is *JLS-secure* if it is secure with $S(n) = n^{\omega(1)}$ and $\epsilon(n) = 2^{-n^\delta}$ for some $\delta > 0$. For our results, it is important that our adversaries are non-uniform circuits.

The breakthrough results of Jain-Lin-Sahai [33, 35] give constructions of both JLS-secure and subexponentially secure iO .

Theorem 11 (Informal version of Jain-Lin-Sahai [35]). *If three “well-founded” cryptographic assumptions hold, then \mathcal{JL} -secure iO exists.*

Theorem 12 (Informal version of Jain-Lin-Sahai [34]). *Assuming three “well-founded” cryptographic assumptions are secure against subexponential-sized adversaries with subexponential advantage, then subexponentially secure iO exists.*

Other Complexity Theory Work Using iO . Building on the work of Bitansky, Paneth, and Rosen [9], Garg, Pandey, and Srinivasan [27] show that computing Nash Equilibria (and thus the TFNP class PPAD) is intractable assuming one-way permutations and iO exist. Impagliazzo, Kabanets, and Volkovich [30] show that if iO exists then the Minimum Circuit Size Problem [41] is in ZPP if and only if $\text{NP} \subseteq \text{ZPP}$.

2.2 Arthur-Merlin Protocols

An Arthur-Merlin protocol [4] for a language $L \subseteq \{0, 1\}^*$ is defined as a constant-round interactive protocol between a computationally unbounded prover (Prover) and a polynomial-time probabilistic verifier (Verifier). For a given input $x \in \{0, 1\}^n$ that is accessible by both Prover and Verifier, Prover wants to convince Verifier that $x \in L$ (even if $x \notin L$) with $\text{poly}(n)$ bits of communication, whereas the Verifier needs to decide whether $x \in L$ based on the information provided by Prover. Formally, the protocol needs to satisfy the following properties.

- **Completeness:** If $x \in L$, it is possible for Prover to send $\text{poly}(n)$ -bits of messages in constant rounds such that Verifier accepts with probability at least 2/3.
- **Soundness:** If $x \notin L$, given any messages from Prover, Verifier accepts with probability at most 1/3.

The complexity class AM is defined as the languages that have a sound and complete Arthur-Merlin protocol. As $\text{coNP} \subseteq \text{AM}$ implies the collapse of PH [11], it is widely believed that $\text{coNP} \not\subseteq \text{AM}$ and therefore $\text{UNSAT} \notin \text{AM}$. Indeed, we know that $\text{NP} = \text{AM}$ assuming a standard derandomization hypothesis: namely, there exists a language $L \in \text{NE} \cap \text{coNE}$ requiring nondeterministic circuits of size $2^{\Omega(n)}$ [45, 59]. This implies that coNP is unlikely to even be infinitely often in AM.

Goldwasser-Sipser Set Lowerbound Protocol. We will need a well-known Arthur-Merlin protocol for approximately counting the size of a set that has efficiently computable membership queries.

Lemma 13 ([28], also see [3, Section 8.4]). *There is an Arthur-Merlin protocol such that the following holds. Suppose that both Prover and Verifier receive a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ and a number $s \leq 2^n$. Let $S = \{x \in \{0, 1\}^n \mid C(x) = 1\}$. Then*

- **Completeness:** If $|S| \geq s$, then there exist messages Prover can send such that Verifier accepts with probability at least 2/3.
- **Soundness:** If $|S| \leq s/2$, then regardless of what Prover sends, Verifier accepts with probability at most 1/3.

Moreover, the protocol is a two-round public-key protocol: Verifier first sends a random seed r and receives a message m ; then it deterministically decides whether to accept based on r and m .

2.3 Bounded Theories PV_1 and APC_1

Cook [22] defined the theory PV related to polynomial-time complexity as an *equational theory* (i.e. sentences are of the form $t = u$ for terms t and u). Based on a machine-independent characterization of FP due to Cobham [20], it can be shown that the set of function symbols introduced in PV (when interpreted in the standard model) is exactly FP. PV_1 is defined as a first-order theory that is a conservative extension of PV axiomatized by universal sentences. The formal definitions of PV and PV_1 are tedious and we refer interested readers to [53, Chapter 12]. We define TPV to be the universal true theory over the language $\mathcal{L}(\text{PV})$ of PV over the standard model \mathbb{N} .⁸

To formalize the probabilistic methods that are widely used in complexity theory and combinatorics, Jeřábek [36–38, 40] introduced an extension of Cook’s PV by including the *dual Weak Pigeonhole Principle* for PV function symbols as axioms, which is now known as APC_1 (stands for approximate counting). Let $f(\vec{w}, x)$ be function symbols⁹, and let $m(n) > n$ be a function. We define the *dual weak Pigeonhole Principle* for $f(\vec{w}, \cdot)$ with stretch function m ,¹⁰ denoted by $\text{dWPHP}_m(f)$, as

$$\begin{aligned} \text{dWPHP}_m(f) := & \forall n \in \text{Log} \ \forall \vec{w} \ \exists y \in \{0, 1\}^{m(n)} \\ & \forall x \in \{0, 1\}^n \ f(\vec{w}, x) \neq y, \end{aligned} \quad (1)$$

which claims that $f(\vec{w}, \cdot) : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ cannot be surjective. Here, $\forall n \in \text{Log}$ is short for $\forall N \ \forall n = |N|$, which means that the feasible reasoning is with respect to strings of length n instead of $\log n$. We use $\text{dWPHP}(f)$ (where m is omitted in the subscript) to mean $\text{dWPHP}_m(f)$ for $m(n) = n + 1$. We define

$$\text{dWPHP}(\text{PV}) := \{ \text{dWPHP}(f) \mid f \in \mathcal{L}(\text{PV}) \},$$

and the theory APC_1 is defined as $\text{PV}_1 + \text{dWPHP}(\text{PV})$. Since any polynomial-time function can be computed with a multi-output polynomial-size circuit, APC_1 can also be defined equivalently as $\text{PV}_1 + \text{dWPHP}(\text{Eval})$, where $\text{Eval}(C, x)$ evaluates the circuit C on the input x .

Remark 14. The stretch function m for $\text{dWPHP}(\text{PV})$ used to define APC_1 can be a subtle issue, as we cannot prove an equivalence between $\text{dWPHP}(\text{PV})$ with different stretch functions within PV_1 (such equivalence can be proved within Buss’s theory S_2^1 for polynomial computation, see, e.g., [39, Theorem 3.1]). Jeřábek [39] also proved that $\text{PV}_1(\alpha)$ (a relativised version of PV_1) cannot prove the equivalence of $\text{dWPHP}(\alpha)$ between different parameters. This will not be a problem for us, since our unprovability result works for the weakest version of dWPHP (i.e. with stretch function $m(n)$ allowed to be an arbitrarily large polynomial). \triangle

We remark that, besides the application of bounded arithmetic to understanding complexity barriers, there are also recent interesting applications in propositional proof complexity [42] and cryptography [32].

⁸That is, TPV contains all sentences of the form $\forall \vec{x} \ \beta$ for some quantifier-free formula β that are true in the standard model \mathbb{N} .

⁹Note that the inclusion of \vec{w} is crucial; if we remove \vec{w} in the definition of APC_1 , we will obtain a (possibly) weaker fragment of APC_1 (see, e.g., [64, Section 2]).

¹⁰In the rest of the paper, we assume that the stretch function $m(n)$ is a PV-function. This is without loss of generality, as the set of PV functions (when interpreted in the standard model) is exactly FP, the class of polynomial-time computable functions.

2.4 Time-Bounded Kolmogorov Complexity

There are multiple notions of time-bounded Kolmogorov complexity in the literature; in this paper, we consider the following. Let \mathcal{U} be a fixed universal Turing machine such that $\mathcal{U}(M, x, 1^t)$ runs the Turing machine encoded by M for t steps on the input x and outputs the string on the tape.

Let $t : \mathbb{N} \rightarrow \mathbb{N}$ be a function.

Definition 15 (Time-Bounded Kolmogorov Complexity [46, 47, 68]). The K^t -complexity of a string $x \in \{0, 1\}^n$, denoted by $K^t(x)$, is the minimum ℓ such that for some $M \in \{0, 1\}^\ell$, $\mathcal{U}(M, x, 1^{t(n)}) = x$.

Definition 16 (Conditional K^t -Complexity). The K^t -complexity of a string $x \in \{0, 1\}^n$ conditioned on a string y , denoted by $K^t(x|y)$, is the minimum ℓ such that for some $M \in \{0, 1\}^\ell$, $\mathcal{U}(M, y, 1^{t(n)}) = x$.

Although the definitions of K^t -complexity and conditional K^t -complexity depend on the universal Turing machine \mathcal{U} , the results in this paper (as well as most results on these notions) are not sensitive to the choice of \mathcal{U} . Therefore we will omit \mathcal{U} and simply use “the encoding of M ” to denote the encoding of M with respect to \mathcal{U} .

3 NO EFFICIENT DETERMINISTIC ALGORITHMS FOR RANGE AVOIDANCE

In this section, we prove Theorem 1 and Theorem 2. We restate both of these theorems below.

REMINDER OF THEOREM 1. Assume that $\text{NP} \neq \text{coNP}$ and iO with JLS-security exists. Then for all $c \geq 1$, there is a $k \geq c$ such that there is no deterministic polynomial-time algorithm for *Avoid* on n^k -size circuits with n inputs and n^c outputs.

REMINDER OF THEOREM 2. Assume subexponentially-secure iO exists. For every $m(n) > n$ there exists an $s(n) = \text{poly}(m)$ such that if there is a deterministic t -time algorithm for *Avoid* circuits with $m(n)$ outputs and size $s(n)$, then

$$\text{coNP} \subseteq \bigcup_{k \in \mathbb{N}} \text{NTIME}[t(m^k(n))]$$

We prove Theorem 2 and remark in the proof how to modify it prove Theorem 1.

PROOF OF THEOREM 2. We will show that, under the assumptions, there is a $t(\text{poly}(s))$ -time nondeterministic algorithm \mathcal{A} for the coNP-complete problem of checking whether a propositional formula with n variables and $O(n)$ size¹¹ is unsatisfiable. Our algorithm \mathcal{A} takes as input an $O(n)$ -size formula φ with n variables, and \mathcal{A} accepts φ if and only if φ is unsatisfiable. \mathcal{A} works as follows:

- (1) Nondeterministically guess a $y \in \{0, 1\}^{m(n)}$ and an $r \in \{0, 1\}^{\text{poly}(n+\lambda)}$ where $\lambda = \text{poly}(m(n))$.
(The degree of the polynomial for λ will be chosen later.)

¹¹To see this problem is coNP-complete, note that an arbitrary formula can be made linear-sized with polynomial blowup while preserving unsatisfiability: simply add extra variables that do nothing.

- (2) Let $C[\varphi, y]$ denote a circuit¹² that takes n input bits, outputs $m(n) > n$ bits, and satisfies

$$C[\varphi, y](x) = \begin{cases} 0^{m(n)}, & \text{if } \varphi(x) = 0 \\ y, & \text{if } \varphi(x) = 1 \end{cases}.$$

- (3) Accept if and only if $y = \text{AVOID}(iO(C[\varphi, y]; r))$.
(Here, we abuse notation and let *AVOID* and iO denote their corresponding algorithms.)

This completes the description of \mathcal{A} .

We now argue the correctness of the reduction \mathcal{A} . Observe that, by construction, $iO(C[\varphi, y]; r)$ is a circuit with n -inputs, $m(n)$ -outputs, and size $\text{poly}(m(n))$. (To see this size bound, note that φ has size $O(n)$, so $C[\varphi, y]$ has size $O(m(n)^c)$ for some constant $c \geq 1$, and iO blows up this size by at most a fixed polynomial in $\lambda = \text{poly}(m(n))$.) Thus, by setting s to be a sufficiently large polynomial in m , the input circuit $iO(C[\varphi, y]; r)$ to the *AVOID* algorithm in item 3 is indeed an instance where the algorithm is assumed to work. Furthermore, it is easy to see that \mathcal{A} runs in nondeterministic time $t(\text{poly}(m(n))) + \text{poly}(m(n)) = t(\text{poly}(m(n)))$. Hence, to prove the theorem we just need to show soundness (if \mathcal{A} accepts, then φ is unsatisfiable) and completeness (if φ is unsatisfiable, then \mathcal{A} accepts).

First we show soundness. Suppose \mathcal{A} accepts φ . Then there exists y and r such that $y = \text{AVOID}(iO(C[\varphi, y]; r))$. By the correctness of *AVOID* and the perfect functionality of iO , we know that y is not in the range of $C[\varphi, y]$. By construction of $C[\varphi, y]$, this means that φ is not satisfiable.

Now we show completeness. Suppose the formula φ is unsatisfiable. For simplicity of notation, we let $m = m(n)$ in the following. We begin by considering the output distribution of the *AVOID* algorithm on $iO(C[\varphi, 0^m]; r)$ for uniformly random r (notice we have set y here to be 0^m). Since *AVOID* always outputs a string in $\{0, 1\}^m$, there exists a “frequent” string y^* such that

$$\Pr_r[y^* = \text{AVOID}(iO(C[\varphi, 0^m]; r))] \geq 2^{-m}.$$

We now consider what happens when we set $y = y^*$. The crucial point is this: because φ is unsatisfiable, observe that $C[\varphi, 0^m]$ and $C[\varphi, y^*]$ compute the same function! Therefore, by the subexponential security of iO , there is an $\epsilon > 0$ such that for every adversary B of size 2^{λ^ϵ} taking input of length $\text{poly}(m(n))$,

$$\left| \Pr_r[B(iO(C[\varphi, 0^m]) = 1) - \Pr_r[B(iO(C[\varphi, y^*]) = 1)] \right| < 2^{-\lambda^\epsilon}.$$

In particular, we can consider the *non-uniform*¹³ circuit adversary $B(X)$ that outputs 1 if and only if $\text{AVOID}(X) = y^*$. Without loss of generality, we can assume the size of B is at most $2^{\text{poly}(m(n))}$, as the exhaustive search algorithm for *AVOID* would provide such a size bound. Thus the size of B is at most

$$2^{\text{poly}(m(n))} \leq 2^{\lambda^\epsilon}$$

when λ is a sufficiently large polynomial in m . (To modify this proof to prove Theorem 1 instead, the only difference is to observe that, if $t(q) = \text{poly}(q)$, then it suffices to have iO with JLS-security in this step, since the circuit checking if the output of *AVOID* equals

¹²It does not matter precisely how we implement $C[\varphi, y]$; the argument will work as long as our $C[\varphi, y]$ satisfies the specification and has size $\text{poly}(m(n))$.

¹³The non-uniformity comes from y^* .

y^* has polynomial size.) Consequently, applying iO security to B , we derive

$$\begin{aligned} \Pr_r[y^* = \text{AVOID}(iO(C[\varphi, y^*]; r))] \\ \geq \Pr_r[y^* = \text{AVOID}(iO(C[\varphi, 0^m]; r))] - 2^{-\lambda^\epsilon} \\ \geq 2^{-m} - 2^{-\lambda^\epsilon} > 0, \end{aligned}$$

by setting λ to be a sufficiently large polynomial in m . Hence, we can conclude there is an r such that

$$y^* = \text{AVOID}(iO(C[\varphi, y^*]; r)),$$

so \mathcal{A} will accept φ . \square

Remark 17. For readers familiar with the cryptographic notion of *witness encryption* [26], we note that in Theorem 2 (and in fact all theorems in this paper that assume iO), we can relax the assumption that iO exists to the potentially weaker assumption that witness encryption with similar security and a deterministic decryption algorithm exists. Informally, witness encryption allows one to encrypt a string y with a SAT formula φ such that

- **(Correctness).** If φ is satisfiable, then one can efficiently decrypt y given a satisfying assignment.
- **(Security).** If φ is unsatisfiable, then the encryption of y and $0^{|y|}$ are computationally indistinguishable.

Garg et al. [25] observed that iO implies witness encryption (with a deterministic decryption algorithm) as a special case. Indeed, in Garg et al.'s construction, one can witness-encrypt a message y with a formula φ by outputting $iO(C[\varphi, y])$. We are implicitly using this construction in our proofs.

To modify the proof of Theorem 2 to use witness encryption instead, one modifies the algorithm \mathcal{A} to the nondeterministic algorithm below:

- (1) Nondeterministically guess a $y \in \{0, 1\}^{m(n)}$ and an $r \in \{0, 1\}^{\text{poly}(n+\lambda)}$
- (2) Let e be the witness encryption of the string y according to φ using randomness r and security parameter λ .
- (3) Let C be the circuit that takes as input a string $x \in \{0, 1\}^n$ and attempts to output the decryption of e using the purported witness x (if decryption fails, output 0^m).
- (4) Accept if and only if $y = \text{AVOID}(C)$.

The analysis of the new algorithm is essentially the same as the analysis of the original \mathcal{A} , where the *perfect functionality* and *indistinguishability* properties of iO are now replaced by the *correctness* and *security* properties of the witness encryption scheme respectively. Our other proofs using iO can be similarly modified. \triangle

The proof of Theorem 2 can be generalized in several other ways. For one, the same proof also works to rule out zero-error randomized algorithms (although zero-error randomized algorithms also imply deterministic algorithms under a derandomization assumption).

The proof can also be generalized to work with just polynomially-secure iO . To do this, instead of $\text{NP} \neq \text{coNP}$, we consider a nondeterministic version of the exponential time hypothesis:

REMINDER OF HYPOTHESIS 5. (*NETH for Circuits*) *There is an $\epsilon > 0$ such that Circuit Unsatisfiability problem on n -input circuits of size $2^{o(n)}$ cannot be solved nondeterministically in $2^{\epsilon n}$ time.*

Hypothesis 5 is in fact a much weaker statement than the usual NETH, which posits that nondeterministically refuting unsatisfiable 3-CNFs requires $2^{\epsilon n}$ -size proofs verifiable in $2^{\epsilon n}$ time. We only require an exponential lower bound in the case of refuting subexponential-size circuits.

REMINDER OF THEOREM 6. *Assuming NETH for Circuits and the existence of polynomially-secure iO , AVOID is not in FP. Moreover, under the assumptions, it follows that for all $b, c \geq 1$ there is an $\epsilon > 0$ such that AVOID cannot be solved in $O(2^{\epsilon n})$ time on circuits of size $2^{\epsilon n}$ with n inputs and bn outputs.*

PROOF. The algorithm for UNSAT in our proof is essentially the same as that of Theorem 2, except we analyze the case of a very large security parameter λ .

Assume there is a universal constant $d \geq 1$ and an algorithm B for polynomially-secure iO which runs in time $O((s \cdot \lambda)^d)$ on circuits of size s with security parameter λ . Furthermore, assume that there are universal constants $b, c \geq 1$ such that for all sufficiently small $\alpha \in (0, 1)$, RANGE AVOIDANCE can be solved in $O(2^{\alpha n})$ time on circuits of size $2^{\alpha n}$ with n inputs and bn outputs. Given the assumptions, we show how to construct a nondeterministic algorithm for proving the unsatisfiability of arbitrary subexponential-size circuits in subexponential time.

Fix $b, c \geq 1$. Let $m = bn$, $\alpha \in (0, 1)$ be a constant to be chosen later, and A be an algorithm for RANGE AVOIDANCE as described above. Let $\epsilon > 0$ be an arbitrarily small constant. Given a circuit φ with n inputs and size $2^{o(n)}$, we run precisely the same reduction as Theorem 2, except with an exponentially large value of λ .

- (1) Nondeterministically guess $y \in \{0, 1\}^m$ and $r \in \{0, 1\}^\ell$, where $\ell = (\text{poly}(|\varphi|)\lambda)^d$ and $\lambda = 2^{\epsilon \cdot n}$. (The parameter ℓ is upper-bounded by the running time of B .)
- (2) Let $C[\varphi, y]$ be a circuit taking n input bits and outputting m bits with the specification:

$$C[\varphi, y](x) = \begin{cases} 0^m, & \text{if } \varphi(x) = 0 \\ y, & \text{if } \varphi(x) = 1. \end{cases}$$

- (3) Accept if and only if $y = A(B(C[\varphi, y]; r))$.

It takes $\text{poly}(|\varphi|)$ time to construct $C[\varphi, y]$. Given our assumption on algorithm B , the output of $B(C[\varphi, y])$ is a circuit of size $(\text{poly}(|\varphi|)\lambda)^d = 2^{\epsilon dn+o(n)}$. Thus our algorithm A for RANGE AVOIDANCE applies to the circuit $B(C[\varphi, y])$, by setting $\alpha = 2\epsilon d$ so that $(\text{poly}(|\varphi|)\lambda)^d = 2^{\epsilon dn+o(n)} \leq 2^{\alpha n}$. Therefore the above algorithm runs in nondeterministic time $2^{\epsilon cn} = 2^{2\epsilon cdn}$. As the input circuit φ has size $2^{o(n)}$, c and d are fixed, and $\epsilon > 0$ can be made arbitrarily small, this would refute NETH for Circuits.

As in the proof of Theorem 2, it remains to show that φ is unsatisfiable if and only if there is a y and r such that y is the output of A on $B(C[\varphi, y]; r)$. First, if such y and r exist, then analogously to the proof of Theorem 2, we conclude that φ is unsatisfiable by the construction of $C[\varphi, y]$.

For the other direction, suppose that φ is unsatisfiable. As in the proof of Theorem 2, we know that there is some $y^* \in \{0, 1\}^m$ such that

$$\Pr_r[y^* = A(B(C[\varphi, 0^m]; r))] \geq 2^{-m}.$$

Given y^* , we may define an “adversary” circuit D which has $y^* \in \{0, 1\}^m$ hard-coded, takes in an input circuit C' (ostensibly simulating $B(C[\varphi, y]; r)$ on some y and r), and outputs 1 if and only if $y^* = A(C')$. Using a standard translation of t -time algorithms into $O(t^3)$ -size circuits, the input circuit C' only has to have size at most $2^{3edn+o(n)}$. Translating the composition of A , B , and $C[\varphi, y]$ into circuits, the size of the adversary circuit D is at most $2^{6ecd\lambda}$, which is only polynomially larger than its input circuit C' .

Assuming polynomially-secure iO , since $C[\varphi, 0^m]$ and $C[\varphi, y^*]$ compute the same function, we have for every constant $K \geq 1$ that

$$\left| \Pr_r[D(B(C[\varphi, 0^m]; r)) = 1] - \Pr_r[D(B(C[\varphi, 0^m]; r)) = 1] \right| < \frac{1}{\lambda^K},$$

which implies

$$\begin{aligned} & \Pr_r[y^* = A(B(C[\varphi, y^*]; r))] \\ & \geq \Pr_r[y^* = A(B(C[\varphi, y^*]; r))] - \frac{1}{\lambda^K} \\ & \geq \frac{1}{2^m} - \frac{1}{\lambda^K}. \end{aligned} \tag{2}$$

The above probability is greater than 0, provided that $2^m < \lambda^K$, i.e.,

$$\lambda > 2^{m/K}. \tag{3}$$

Recall that we set $\lambda = 2^{\varepsilon n}$ for an arbitrarily small fixed $\varepsilon > 0$, and $m = bn$ for a fixed $b \geq 1$, while the constant $K \geq 1$ can be as arbitrarily large as needed (independently of all other constants). Therefore (3) holds.

By (2), there is an r such that $y^* = A(B(C[\varphi, y^*]; r))$, which completes the proof. \square

4 APPLICATION IN BOUNDED ARITHMETIC: SEPARATING APC₁ AND PV₁

In this section, we prove a conditional separation of the bounded theories APC₁ and PV₁, assuming that iO with JLS-security exists and coNP is not infinitely often in AM.

The only result from logic that we need is the standard KPT Witnessing Theorem for $\forall\exists\forall$ formulas, which connects the provability of any $(\forall\exists\forall)$ -sentence with a Student-Teacher game for interactively computing a witness to the existential quantifier.

Theorem 18 (KPT Witnessing Theorem for TPV [57]). *For every quantifier-free formula $\varphi(\vec{x}, y, z)$ in the language $\mathcal{L}(\text{PV})$, if $\text{TPV} \vdash \forall\vec{x} \exists y \forall z \varphi(\vec{x}, y, z)$, then there is a $k \in \mathbb{N}$ and $\mathcal{L}(\text{PV})$ -terms t_1, t_2, \dots, t_k such that*

$$\text{TPV} \vdash \forall\vec{x} \forall z_1 \forall z_2 \dots \forall z_k \bigvee_{i=1}^k \varphi(\vec{x}, t_i(\vec{x}, z_1, \dots, z_{i-1}), z_i). \tag{4}$$

It is well-known that the terms t_1, t_2, \dots, t_k extracted from the proof in the KPT Witnessing Theorem can be interpreted as an k -round interactive computation of a witness y such that $\forall z \varphi(\vec{x}, y, z)$ given the input \vec{x} . Consider the following game between a *Student* who wants to find a correct witness y and a *Teacher* who will provide help. In the first round, the Student proposes $y_1 := t_1(\vec{x})$

as a candidate, and, in the case, y_1 is not a correct witness of $\exists y \forall z \varphi(\vec{x}, y, z)$, the Teacher provides a *counterexample* z_1 such that $\varphi(\vec{x}, y_1, z_1)$ is false. The Student then proposes a new candidate $y_2 := t_2(\vec{x}, z_1)$ based on the counterexample given in the first round, and the Teacher, again, provides a counterexample z_2 if y_2 is not a correct witness, etc. Equation (4) means that after k rounds of interaction between the Student and the Teacher, at least one of the y_1, y_2, \dots, y_k proposed by the Student has to be a correct witness of $\exists y \forall z \varphi(\vec{x}, y, z)$. We refer the readers to [53] for more discussion about the KPT Witnessing Theorem and the Student-Teacher game.

4.1 Provability of dWPHP and the Tractability of Avoid

The KPT Witnessing Theorem (Theorem 18) provides a connection between the provability of dWPHP and the tractability of Avoid, in the sense of the Student-Teacher game. Let $\text{Eval}(C, x) := C(x)$ be the circuit evaluation function. Recall that $\text{dWPHP}_\ell(\text{Eval})$ refers to the following sentences:

$$\begin{aligned} \text{dWPHP}_\ell(\text{Eval}) &:= \forall n \in \text{Log} \forall \text{circuits } C : \{0, 1\}^n \rightarrow \{0, 1\}^n \\ &\quad \exists y \in \{0, 1\}^\ell \forall x \in \{0, 1\}^n [\text{Eval}(C, x) \neq y]. \end{aligned}$$

Suppose that $\text{dWPHP}_\ell(\text{Eval})$ is provable in TPV . Then there is a constant-round Student-Teacher game that finds a $y \in \{0, 1\}^\ell$ witnessing the existential quantifier, where the Teacher provides counterexamples for the universal quantifier over $x \in \{0, 1\}^n$. Taking a closer look at this game, we can see that this corresponds to an algorithm for Avoid with circuit-inversion oracle queries, which is formally defined as follows.

Definition 19 (Solving Avoid with a Circuit-Inversion Oracle). Let $m = m(n)$ and $k = k(n)$. A *polynomial-time algorithm with k circuit-inversion oracle queries for Avoid with m outputs* is a polynomial-time oracle algorithm A such that given a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and access to an oracle $O(\cdot) : \{0, 1\}^m \rightarrow \{0, 1\}^n$ with at most k queries, $A^{O(\cdot)}(C)$ outputs a $y \in \{0, 1\}^m$ such that $C(O(y)) \neq y$. Furthermore, $O(y)$ always returns an x such that $C(x) = y$, when such an x exists.

Theorem 20. *For every constructive function $m(n) \leq \text{poly}(n)$ such that $m(n) > n$, if $\text{TPV} \vdash \text{dWPHP}_\ell(\text{Eval})$, then Avoid with m outputs has a polynomial-time algorithm with $O(1)$ circuit-inversion oracle queries.*

PROOF. Let $m(n) = \text{poly}(n)$ be some constructive function. By the assumptions, we know that

$$\text{TPV} \vdash \forall n \in \text{Log} \forall C \exists y \in \{0, 1\}^{m(n)} \forall x \in \{0, 1\}^n [\text{Eval}(C, x) \neq y].$$

By the KPT Witnessing Theorem (Theorem 18), there is a $k = O(1)$ and $\mathcal{L}(\text{PV})$ -terms t_1, t_2, \dots, t_k such that

$$\begin{aligned} \text{TPV} \vdash \forall n \forall C \forall z_1 \dots \forall z_n & \left([\text{Eval}(C, z_1) \neq t_1(n, C)] \vee \right. \\ & \quad [\text{Eval}(C, z_2) \neq t_2(n, C, z_1)] \vee \dots \vee \\ & \quad \left. [\text{Eval}(C, z_k) \neq t_k(n, C, z_1, \dots, z_{k-1})] \right). \end{aligned}$$

Now we show that given any circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ and access to an oracle $O : \{0, 1\}^{m(n)} \rightarrow \{0, 1\}^n$, there is a polynomial time algorithm that makes k queries to $O(\cdot)$ and finds a

$y \in \{0,1\}^{m(n)}$ such that $C(O(y)) \neq y$. Let $f_1, f_2, \dots, f_k \in \text{FP}$ be the functions that are the interpretations of t_1, t_2, \dots, t_k , respectively, in the standard model. We play the aforementioned Student-Teacher game, where in the i -th round, the Student proposes $y_i := f_i(n, C, z_1, \dots, z_{i-1})$, and the Teacher responds with $z_i := O(y_i)$. By Equation (4) and the soundness of TPV in the standard model, we know that for some $i \in [k]$, the Teacher fails to provide a correct counterexample z_i in the i -th round, i.e., $\text{Eval}(C, z_i) \neq y_i$. The algorithm can then output y_i . \square

4.2 Impossibility of Solving AVOID with a Circuit-Inversion Oracle

Now we show that AVOID has no polynomial-time algorithm with $O(1)$ circuit-inversion oracle queries under plausible assumptions by generalizing the proof of Theorem 2.

Theorem 21. *Let $m = m(n) = \text{poly}(n)$ and $k = O(1)$ such that $m \geq n + 1$. Assume that coNP is not infinitely often in AM and iO with JLS -security exists. Then there is no polynomial-time deterministic algorithm for AVOID on circuits with m outputs using k circuit-inversion oracle queries.*

PROOF. Let m and k be defined as above and let iO be a JLS -secure indistinguishability obfuscator. To prove that AVOID with the given parameters cannot be solved in deterministic polynomial time, it suffices to show that for every k -query oracle algorithm A , there exists a circuit C mapping n -bits to m -bits and a (consistent) inversion oracle $O(\cdot) : \{0,1\}^m \rightarrow \{0,1\}^n$ such that $A^{O(\cdot)}(C)$ outputs a y where $C(O(y)) = y$ (in which case, A fails to solve AVOID on C). For any polynomial-time algorithm A that makes k queries to $O(\cdot) : \{0,1\}^m \rightarrow \{0,1\}^n$, we can decompose A into $k+1$ polynomial-time algorithms A_1, A_2, \dots, A_{k+1} (without oracle queries) that work as follows:

- A_1 : Given the input circuit C , it computes $y_1 = A_1(C)$ and queries $O(y_1)$.
- A_2 : Letting x_1 be the answer to the last query, it computes $y_2 = A_2(C, x_1)$ and queries $O(y_2)$.
- A_3 : Letting x_2 be the answer to the last query, it computes $y_3 = A_3(C, x_1, x_2)$ and queries $O(y_3)$.
- \dots
- A_{k+1} : Letting x_k be the answer to the last query, it computes $y_{k+1} = A_{k+1}(C, x_1, \dots, x_k)$ and outputs y_{k+1} .

Therefore to rule out the existence of the oracle algorithm A as described above, it suffices to show that for all deterministic polynomial-time algorithms A_1, \dots, A_{k+1} , there is a circuit $C : \{0,1\}^n \rightarrow \{0,1\}^m$, strings $y_1, \dots, y_{k+1} \in \{0,1\}^m$, and strings $x_1, \dots, x_{k+1} \in \{0,1\}^n$ such that:

- **(Oracle Consistency).** For all $i, j \in [k+1]$ such that $y_i = y_j$, we have $x_i = x_j$.
(That is, the oracle gives consistent answers x_i to input strings y_i .)
- **(Oracle Inverting).** For every $i \in [k+1]$, we have $C(x_i) = y_i = A_i(C, x_1, \dots, x_{i-1})$.
(That is, given y_i , the oracle indeed provides an x_i such that $C(x_i) = y_i$.)

To add more detail, the aforementioned circuit C and any oracle satisfying $O(y_i) = x_i$ for every $i \in [k+1]$ can force each A_i (equivalently, the i -th oracle query of A) to output y_i . In such a case, A_{k+1} (equivalently, the oracle algorithm A) outputs a string $y_{k+1} = C(x_{i+1})$ that is in the output range of the input circuit C , and thus **does not** solve AVOID.

We first introduce some notation. For n -variable 3-CNF formulas $\varphi_1, \dots, \varphi_k$ of size n and strings $y_1, \dots, y_k \in \{0,1\}^m$, we let $C[\varphi_1, \dots, \varphi_k; y_1, \dots, y_k]$ denote a polynomial-size circuit that takes an input $(x, i) \in \{0,1\}^n \times [k]$ and outputs

$$C[\varphi_1, \dots, \varphi_k; y_1, \dots, y_k](x, i) := \begin{cases} 0^m, & \text{if } \varphi_i(x) = 0 \\ y_i, & \text{if } \varphi_i(x) = 1 \end{cases}.$$

In the case that we do not specify a φ_i and its corresponding y_i , such as in $C[\varphi_1, \varphi_2; y_1, y_2]$ for $k = 3$, we adopt the convention that any missing φ_i is the trivial unsatisfiable formula \perp and $y_i = 0^m$. For instance when $k = 3$, $C[\varphi_1, \varphi_2; y_1, y_2] := C[\varphi_1, \varphi_2, \perp; y_1, y_2, 0^m]$.

Fix any polynomial-time algorithm A with k oracle queries and its decomposition as polynomial-time algorithms A_1, A_2, \dots, A_{k+1} . We begin by making the following claim.

Claim 22. *For all $j \in \{0, \dots, k+1\}$ and for all sufficiently large n , there exist n -variable satisfiable 3-CNF formulas $\varphi_1, \dots, \varphi_j$ of size n , strings $x_1, \dots, x_j \in \{0,1\}^n$, and strings $y_1, \dots, y_j \in \{0,1\}^m$ such that the following holds.*

- For all distinct $i_1, i_2 \in [j]$, if $\varphi_{i_1} = \varphi_{i_2}$, then $x_{i_1} = x_{i_2}$.
- For every $i \in [j]$, $\varphi_i(x_i) = 1$.
- Let $\hat{C}_j = iO(C[\varphi_1, \dots, \varphi_j; y_1, \dots, y_j]; r)$. Over the random seed r of iO , it holds with probability at least $2^{-\Omega((2k-j)m)}$ that for every $i \in [j]$, $y_i = A_i(\hat{C}_j, x_1, \dots, x_{i-1})$.

Observe that when $j = k+1$, Claim 22 implies the existence of $\varphi_1, \dots, \varphi_{k+1}, x_1, \dots, x_{k+1}$, and y_1, \dots, y_{k+1} that satisfy the aforementioned *Oracle Consistency* property (which follows from the first bullet of Claim 22) and *Oracle Inverting* property (which follows from the second and the third bullet of Claim 22 and the perfect functionality of iO). Indeed, this claim shows that the polynomial-time algorithm A with k circuit-inversion oracle queries will **fail** on \hat{C}_{k+1} (A will output a string in the range of \hat{C}_{k+1}) with probability at least $2^{-\Omega((k-1)m)}$. Therefore, to prove the theorem, it remains to prove Claim 22.

Before we start, one crucial definition is in order. Define a circuit D to be *j-good* if for every $i \in [j]$, $y_i = A_i(D, x_1, \dots, x_{i-1})$. In other words, D is *j-good* if it satisfies the property in the third bullet of Claim 22.

We prove Claim 22 by induction on j . The base case $j = 0$ is trivially true. Now we assume the claim is true for $j - 1$, which gives 3-CNF formulas $\varphi_1, \dots, \varphi_{j-1}$ of size n , strings x_1, \dots, x_{j-1} , and strings y_1, \dots, y_{j-1} that make the claim true for $j - 1$. Let $\lambda = \text{poly}(m)$ be the security parameter to be determined later, and let $\ell = \text{poly}(n, \lambda)$ be the randomness complexity of iO with security parameter λ . We define an AM protocol \mathcal{P} , detailed in Algorithm 1, that attempts to solve UNSAT (i.e., the Prover aims to convince the Verifier that a given formula is unsatisfiable). For simplicity, we assume without loss of generality that our formulas are 3-CNFs with n clauses and n variables.

```

Input: A 3-CNF formula  $\varphi(x)$  on  $n$  variables and clauses.
1 Prover sends 3-CNF formulas  $\varphi_1, \dots, \varphi_{j-1}$ , strings
 $x_1, \dots, x_{j-1} \in \{0, 1\}^n$ , and  $y_1, \dots, y_j \in \{0, 1\}^m$ ;
2 Verifier rejects if  $\varphi_i(x_i) \neq 1$  for some  $i \in [j]$ , or  $\varphi_{i_1} = \varphi_{i_2}$ 
and  $x_{i_1} \neq x_{i_2}$  for distinct  $i_1, i_2 \in [j]$ ;
// For  $r \in \{0, 1\}^\ell$ , let
 $C_j^r := iO(C[\varphi_1, \dots, \varphi_{j-1}, \varphi; y_1, \dots, y_j]; r)$ .
// Let  $E : \{0, 1\}^\ell \rightarrow \{0, 1\}$  be a circuit such that
 $E(r) = 1$  if and only if  $C_j^r$  is  $j$ -good.
3 Prover and Verifier run the Goldwasser-Sipser protocol
(Lemma 13) on the instance  $(E, \delta)$ ;
// The parameter  $\delta = 2^{\ell - \Omega((2k-j)m)}$ .
// Prover aims to convince Verifier that
 $|\{r \in \{0, 1\}^\ell \mid E(r) = 1\}| \geq \delta$ .

```

Algorithm 1: The AM protocol \mathcal{P} that aims to solve UNSAT.

Completeness of \mathcal{P} . We now show that the AM protocol \mathcal{P} is complete; the Verifier accepts every unsatisfiable 3-CNF formula φ on n variables with probability at least $2/3$. Let φ be an arbitrary unsatisfiable formula, and let \hat{C}_{j-1} be the random variable $\hat{C}_{j-1} := iO(C[\varphi_1, \dots, \varphi_{j-1}; y_1, \dots, y_{j-1}]; r)$ defined over $r \in \{0, 1\}^\ell$. By the induction hypothesis, we know that

$$\Pr_{r \in \{0, 1\}^\ell} [\hat{C}_{j-1} \text{ is } (j-1)\text{-good}] \geq 2^{-\Omega((2k-j+1)m)}.$$

By an averaging argument, there is a $y \in \{0, 1\}^m$ such that

$$\begin{aligned} & \Pr_{r \in \{0, 1\}^\ell} [(\hat{C}_{j-1} \text{ is } (j-1)\text{-good}) \wedge [y = A_j(\hat{C}_{j-1}, x_1, \dots, x_{j-1})]] \\ & \geq 2^{-\Omega((2k-j+1)m)} \cdot 2^{-m} = 2^{-\Omega((2k-j)m)}. \end{aligned} \quad (5)$$

Let $y_j \in \{0, 1\}^m$ be one such y , and let \hat{C}_j be the random variable $\hat{C}_j := iO(C[\varphi_1, \dots, \varphi_{j-1}, \varphi; y_1, \dots, y_j]; r)$ defined over $r \in \{0, 1\}^\ell$. Since φ is unsatisfiable, it follows that

$$C[\varphi_1, \dots, \varphi_{j-1}, \varphi; y_1, \dots, y_j] \text{ and } C[\varphi_1, \dots, \varphi_{j-1}; y_1, \dots, y_{j-1}]$$

compute the same function. By the JLS-security of iO , we know that \hat{C}_j and \hat{C}_{j-1} are $2^{-\lambda^\ell}$ -indistinguishable against any polynomial-size adversary. To verify that a circuit \hat{C} of $\text{poly}(n)$ size is j -good, we need to check $y_i = A_i(\hat{C}, x_1, \dots, x_{i-1})$ for every $i \in [j]$, which can be done by a circuit of $\text{poly}(n)$ size. This means by Equation (5) that

$$\begin{aligned} & \Pr_{r \in \{0, 1\}^\ell} [\hat{C}_j \text{ is } j\text{-good}] \\ & = \Pr_{r \in \{0, 1\}^\ell} [(\hat{C}_j \text{ is } (j-1)\text{-good}) \wedge [y_j = A_j(\hat{C}_j, x_1, \dots, x_{j-1})]] \\ & \geq 2^{-\Omega((2k-j)m)} \cdot 2^{-\lambda^\ell}. \end{aligned} \quad (6)$$

Let $\lambda := m^{2/\ell} = \text{poly}(n)$ and $\delta := 2^\ell \cdot (6) = 2^{\ell - \Omega(2k-j)m}$. The Prover will work as follows. In the first step, the Prover sends 3-CNF formulas $\varphi_1, \dots, \varphi_{j-1}$, strings $x_1, \dots, x_{j-1} \in \{0, 1\}^n$, and $y_1, \dots, y_j \in \{0, 1\}^m$. By Equation (6), we know that \hat{C}_j is j -good with probability at least $\delta/2^\ell$, which means by the definition of $E : \{0, 1\}^\ell \rightarrow \{0, 1\}$ (which can be implemented in $\text{poly}(n)$ size) that there is a Prover for the Goldwasser-Sipser protocol in Line

3 of Algorithm 1 such that the Verifier accepts with probability at least $2/3$. This concludes the completeness of the protocol \mathcal{P} .

Employ the Lack of Soundness. At this point, we have shown that \mathcal{P} is a polynomial-time AM protocol attempting to check unsatisfiability, and that \mathcal{P} has the completeness property. By the assumption that coNP is not infinitely often in AM, \mathcal{P} cannot solve unsatisfiability even infinitely often, which means that this \mathcal{P} does not have soundness for *all* sufficiently large n . In other words, there is a Prover such that for sufficiently large n , the Verifier accepts some satisfiable 3-CNF formula φ on n variables with probability $> 1/3$. Let φ_j be this formula φ and let

- 3-CNF formulas $\varphi_1, \varphi_2, \dots, \varphi_{j-1}$,
- strings $x_1, x_2, \dots, x_{j-1} \in \{0, 1\}^n$, and
- strings $y_1, y_2, \dots, y_j \in \{0, 1\}^m$

be the message sent in Line 1 (of Algorithm 1) by this Prover on the input φ_j . Since the Verifier does not reject in Line 2, we have

- $\varphi_i(x_i) = 1$ for every $i \in [j-1]$, and
- for all distinct $i_1, i_2 \in [j]$, if $\varphi_{i_1} = \varphi_{i_2}$, then $x_{i_1} = x_{i_2}$.

We define the string $x_j \in \{0, 1\}^n$ to be x_i if there is some $i \in [j-1]$ such that $\varphi_i = \varphi_j$; otherwise, we set x_j to be an arbitrary n -bit string such that $\varphi_j(x_j) = 1$. Now we show that the formulas $\varphi_1, \dots, \varphi_j$, strings $x_1, \dots, x_j \in \{0, 1\}^n$, and strings $y_1, \dots, y_j \in \{0, 1\}^m$ satisfy the conditions of Claim 22, which will conclude the proof.

Let \hat{C}_j be the random variable defined as

$$\hat{C}_{j-1} := iO(C[\varphi_1, \dots, \varphi_j; y_1, \dots, y_j]; r).$$

Since the Verifier accepts with probability $> 1/3$, by the soundness of the Goldwasser-Sipser protocol (Lemma 13), we know that

$$\Pr_{r \in \{0, 1\}^\ell} [\hat{C}_j \text{ is } j\text{-good}] \geq \frac{1}{3} \cdot 2^{-\ell} \cdot \delta = 2^{-\Omega((2k-j)m)}.$$

This implies the third bullet of Claim 22. The first two bullets hold by the definition of φ_i , x_i , and y_i . \square

By combining Theorem 20 and Theorem 21, we conclude that $\text{dWPHP}_\ell(\text{Eval})$ is not provable in TPV based on the assumptions of Theorem 21. Since TPV is an extension of PV_1 , this further means that $\text{dWPHP}(\text{PV})$ is not provable in PV , which separates PV_1 and APC_1 . We summarize the results as follows.

Corollary 23. *Assume the existence of JLS-secure iO and coNP is not infinitely often in AM. For every constructive function $\ell(n) \leq \text{poly}(n)$ such that $\ell(n) > n$, $\text{TPV} \not\models \text{dWPHP}_\ell(\text{Eval})$. In particular, APC_1 is a strict extension of PV_1 .*

5 THE ORACLE DERANDOMIZATION HYPOTHESIS FOR TIME-BOUNDED KOLMOGOROV COMPLEXITY

In this section, we investigate the Oracle Derandomization Hypothesis (ODH) and its variants.

Hypothesis 24 (Oracle Derandomization Hypothesis [23]). *For any m with $n \leq m \leq \text{poly}(n)$, there is a deterministic algorithm A mapping m bits to n bits such that for all $z \in \{0, 1\}^m$ we have*

that $y = A(z)$ satisfies that the circuit complexity¹⁴ of y given oracle access to z is at least $n^{\Omega(1)}$ (when y and z are viewed as truth tables in the natural way).

We note that ODH is closely related to the COMPLEXITY problem studied by Kleinberg, Korten, Mitropolsky, and Papadimitriou [44], which is defined as follows:

Search Problem: COMPLEXITY

Input: A length- n truth table z of a Boolean function.

Output: A length- n truth table y such that the circuit complexity of the function represented by y given oracle access to the function represented by z is at least $\Omega(\frac{n}{\log^2 n})$.

There are two main differences between ODH, and having a deterministic polynomial-time algorithm for COMPLEXITY. First, to solve COMPLEXITY, one is interested in truth tables with near-maximal conditional circuit complexity, while in ODH it suffices to output truth tables that have conditional circuit complexity $n^{\Omega(1)}$. Second, in COMPLEXITY, $|z| = |y|$, while in ODH one needs to handle cases where $|y|$ is polynomially smaller than $|z|$.

We consider a time-bounded Kolmogorov complexity version of ODH. Roughly speaking, this version says that given a string z of length n , one can efficiently deterministically generate a string y of length $n^{0.1}$ such that the time-bounded Kolmogorov complexity of y given z is large. We give a formal definition below.

Hypothesis 25 (Time-Bounded Kolmogorov Complexity Oracle Derandomization Hypothesis). *For every m with $n \leq m \leq \text{poly}(n)$ and for any $t = \text{poly}(n)$, there is a deterministic algorithm A mapping m bits to n bits such that for all $z \in \{0, 1\}^m$ we have that $y = A(z)$ is a string such that $K^t(y|z) = n^{\Omega(1)}$.*

Theorem 26. *Hypothesis 25 is false assuming $\text{NP} \neq \text{coNP}$ and JLS-secure iO exists.*

PROOF. The proof is very similar to the proof of Theorem 2.

Let $m = \text{poly}(n)$ and $t = \text{poly}(n)$ be parameters we set later. For contradiction, let A be an algorithm mapping m bits to n bits such that for all $z \in \{0, 1\}^m$ we have $K^t(y|z) = \Omega(n^\epsilon)$ for some $\epsilon > 0$, where $y = A(z)$.

We give a polynomial-time nondeterministic algorithm for checking if an $n^{\epsilon/2}$ -variable formula φ is unsatisfiable.

- (1) Nondeterministically guess strings $y \in \{0, 1\}^n$ and $r \in \{0, 1\}^{\text{poly}(n+\lambda)}$ where $\lambda = \text{poly}(n)$ is the security parameter of the JLS-secure iO to be determined later.
- (2) Let $C[\varphi, y]$ denote a $\text{poly}(n)$ -size circuit that takes n input bits and outputs m bits and satisfies

$$C[\varphi, y](x) = \begin{cases} 0^m, & \text{if } \varphi(x) = 0 \\ y, & \text{if } \varphi(x) = 1. \end{cases}$$

- (3) Set z to be the m -bit string given by the description of the circuit $iO(C[\varphi, y]; r)$. (We set $m > n$ so that this is possible. We can also pad the description with zeroes if necessary.)
- (4) Accept if and only if $y = A(z)$.

¹⁴Actually, [23] considers nondeterministic circuit complexity instead of the usual circuit complexity.

It is easy to see that this is a polynomial-time nondeterministic algorithm. We now will show soundness and completeness for sufficiently large n .

First, we show soundness. If φ is satisfiable, then we claim that for all choices of y and r , we have $K^t(y|z) \leq O(n^{\epsilon/2})$. This is because if φ is satisfiable, then there is an input $x^* \in \{0, 1\}^{n^{\epsilon/2}}$ such that $\varphi(x^*) = 1$. Consequently, $C[\varphi, y](x^*) = y$. Thus, when z is interpreted as the *description* of the circuit $iO(C[\varphi, y]; r)$, we have $z(x^*) = y$. This shows that $K^t(y|z) \leq O(n^{\epsilon/2})$ (setting t to be a sufficiently large polynomial). This completes our proof of soundness.

Now we show completeness. Suppose φ is unsatisfiable. As in the proof of Theorem 2, we consider A being run on $iO(C[\varphi, 0^n]; r)$ for uniformly random r . As A must output an element of $\{0, 1\}^n$, we know there's a y^* satisfying

$$\Pr_r[y^* = A(iO(C[\varphi, 0^n]; r))] \geq 2^{-n}.$$

To this end, we (just as in Theorem 2) consider an adversary circuit B which takes in a circuit C' (ostensibly of the form $C' = iO(C[\varphi, 0^m])$) and outputs 1 if and only if $y^* = A(C')$. As in Theorem 2, we can argue that the size of the adversary B is small enough that our iO assumption applies.

As in Theorem 2, $C[\varphi, 0^n]$ and $C[\varphi, y^*]$ compute the same function. Thus, by the JLS-security of iO , applied to the adversary circuit B on the input circuits $iO(C[\varphi, 0^n])$ and $iO(C[\varphi, y^*])$, we have

$$\begin{aligned} & \Pr_r[y^* = A(iO(C[\varphi, y^*]; r))] \\ & \geq \Pr_r[y^* = A(iO(C[\varphi, 0^n]; r))] - 2^{-\lambda^\epsilon} \\ & \geq 2^{-n} - 2^{-\lambda^\epsilon} > 0 \end{aligned}$$

by setting λ to be a sufficiently large polynomial in n . Thus, there is an r such that

$$y^* = A(iO(C[\varphi, y^*]; r)),$$

so \mathcal{A} accepts φ . □

6 CONCLUSION

We conclude with several open questions and directions of particular interest.

Intractability of C -Avoid? For a given circuit class C (e.g., AC^0 , TC^0 , NC^1), Ren, Santhanam, and Wang [66] introduce the C -Avoid problem, which considers Avoid over circuits drawn from C . They showed many interesting lower bound consequences from showing C -AVOID is in FP (or even FP^{NP}).

Our work suggests the following natural question:

What is the “weakest” circuit class C such that, under plausible assumptions, C -AVOID is not in FP?

It seems reasonable that NC -AVOID (i.e., RANGE AVOIDANCE over $\text{poly}(\log n)$ -depth circuits of $\text{poly}(n)$ size) is not in FP, under similar assumptions to ours.

Intractability of Range Avoidance on Uniform Circuits? The specific instances used to show lower bound and combinatorial consequences of $\text{AVOID} \in \text{FP}$ [29, 50, 66] arise from *uniform* circuits. Formally, for these instances, there is a deterministic machine M with an $O(1)$ -bit description such that, given 1^n , M runs in $\text{poly}(n)$

time and prints the description of a C_n on n inputs and $m > n$ outputs. Could **Avoid** on such uniformly-generated circuits be in FP? On the one hand, our arguments concluding **Avoid** \notin FP evidently rely on the non-uniformity of input circuits in a crucial way. On the other hand, since (for example) NP-complete problems on uniformly-generated instances typically remain hard in a different way (e.g., NEXP-complete), some of the authors are skeptical that this special case of **Avoid** is easy.

The Landscape Around Range Avoidance. Prior to our work, there were no examples of a relational problem that has both an efficient randomized algorithm (i.e. in the class FBPP) and an inefficient deterministic algorithm (for example, in FPH), but was unlikely to have an efficient deterministic algorithm. (For comparison, Aaronson, Buhrman, and Kretschmer [1] give an example of problem that is in $\text{FEXP} \cap \text{FBPP}$ but unconditionally not in FP.) **Avoid** has these properties, and (for large stretch) apparently lies in $(\text{FBPP} \cap \text{FPH}) \setminus \text{FP}$, under the assumptions of this paper. What are other examples of such problems? The space of such problems seems interesting to study, in itself.

The Structure of Probabilistic Feasible Reasoning. We have demonstrated that under plausible assumptions, PV_1 is a strict sub-theory of APC_1 . Furthermore, the fragment UAPC_1 of APC_1 that sustains the basic mechanism of Jeřábek's approximate counting framework [38] is strictly weaker than APC_1 under similar assumptions. This motivates revisiting the question of what is the “right” theory to capture probabilistic feasible reasoning. For instance, we may ask the following question: Is UAPC_1 conservative over PV_1 ? Is there a strict fragment of UAPC_1 that (in some sense) captures probabilistic feasible reasoning? Are there interesting mathematical theorems that are provable in APC_1 but not provable in UAPC_1 (or its weaker fragments)?

ACKNOWLEDGMENTS

We thank Lijie Chen, Yilei Chen, Aayush Jain, Antonina Kolokolova, Oliver Korten, Rachel Huijia Lin, Zhenjian Lu, Pavel Pudlák, Igor Carboni Oliveira, Hanlin Ren, Amit Sahai, Rahul Santhanam, Neekon Vafa, Vinod Vaikuntanathan, and Tianqi Yang for discussions on this work. Hanlin Ren pointed us to the Oracle Derandomization Hypothesis. Part of the work was done when Jiatu Li was visiting the University of Warwick as an intern of Igor Carboni Oliveira.

REFERENCES

- Scott Aaronson, Harry Buhrman, and William Kretschmer. 2023. A Qubit, a Coin, and an Advice String Walk Into a Relational Problem. *Electron. Colloquium Comput. TR23-015* (2023). <https://eccc.weizmann.ac.il/report/2023/015>
- Scott Aaronson and Avi Wigderson. 2009. Algebraization: A New Barrier in Complexity Theory. *ACM Trans. Comput. Theory* 1, 1 (2009), 2:1–2:54. <https://doi.org/10.1145/1490270.1490272>
- Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity - A Modern Approach*. Cambridge University Press. <http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264>
- László Babai. 1985. Trading Group Theory for Randomness. In *STOC*. ACM, 421–429. <https://doi.org/10.1145/22145.22192>
- 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>
- Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. 2001. On the (Im)possibility of Obfuscating Programs. In *CRYPTO (Lecture Notes in Computer Science, Vol. 2139)*. Springer, 1–18. https://doi.org/10.1007/3-540-44647-8_1
- Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. 2012. On the (Im)possibility of Obfuscating Programs. *J. ACM* 59, 2 (2012), 6:1–6:48. <https://doi.org/10.1145/2160159>
- Ohad Barta, Yuval Ishai, Rafail Ostrovsky, and David J. Wu. 2020. On Succinct Arguments and Witness Encryption from Groups. In *CRYPTO (1) (Lecture Notes in Computer Science, Vol. 12170)*. Springer, 776–806. https://doi.org/10.1007/978-3-030-56784-2_26
- Nir Bitansky, Omer Paneth, and Alon Rosen. 2015. On the Cryptographic Hardness of Finding a Nash Equilibrium. In *FOCS*. IEEE Computer Society, 1480–1498. <https://doi.org/10.1109/FOCS.2015.94>
- Dan Boneh and Marik Zhandry. 2017. Multiparty Key Exchange, Efficient Traitor Tracing, and More from Indistinguishability Obfuscation. *Algorithmica* 79, 4 (2017), 1233–1285. <https://doi.org/10.1007/s00453-016-0242-8>
- Ravi B. Boppana, Johan Håstad, and Stathis Zachos. 1987. Does co-NP Have Short Interactive Proofs? *Inf. Process. Lett.* 25, 2 (1987), 127–132. [https://doi.org/10.1016/0020-0190\(87\)90232-8](https://doi.org/10.1016/0020-0190(87)90232-8)
- Samuel R Buss. 1985. *Bounded arithmetic*. Princeton University.
- Samuel R. Buss, Leszek Aleksander Kolodziejczyk, and Neil Thapen. 2014. Fragments of Approximate Counting. *J. Symb. Log.* 79, 2 (2014), 496–525. <https://doi.org/10.1017/jsl.2013.37>
- Jan Bydzovsky, Jan Krajicek, and Igor Carboni Oliveira. 2020. Consistency of circuit lower bounds with bounded theories. *Log. Methods Comput. Sci.* 16, 2 (2020). [https://doi.org/10.23638/LMCS-16\(2:12\)2020](https://doi.org/10.23638/LMCS-16(2:12)2020)
- Jan Bydzovsky and Moritz Müller. 2020. Polynomial time ultrapowers and the consistency of circuit lower bounds. *Arch. Math. Log.* 59, 1–2 (2020), 127–147. <https://doi.org/10.1007/s00153-019-00681-y>
- Marco Carmosino, Valentine Kabanets, Antonina Kolokolova, and Igor Carboni Oliveira. 2021. LEARN-Uniform Circuit Lower Bounds and Provability in Bounded Arithmetic. In *FOCS*. IEEE, 770–780. <https://doi.org/10.1109/FOCS52979.2021.00080>
- Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. 2016. Nondeterministic Extensions of the Strong Exponential Time Hypothesis and Consequences for Non-reducibility. In *ITCS*. ACM, 261–270. <https://doi.org/10.1145/2840728.2840746>
- Lijie Chen, Ron D. Rothblum, Roei Tell, and Eylon Yogev. 2020. On Exponential-Time Hypotheses, Derandomization, and Circuit Lower Bounds: Extended Abstract. In *FOCS*. IEEE, 13–23. <https://doi.org/10.1109/FOCS46700.2020.00010>
- Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. 2018. GGH15 Beyond Permutation Branching Programs: Proofs, Attacks, and Candidates. In *CRYPTO (2) (Lecture Notes in Computer Science, Vol. 10992)*. Springer, 577–607. https://doi.org/10.1007/978-3-319-96881-0_20
- Alan Cobham. 1965. The intrinsic computational difficulty of functions. *Proc. Logic, Methodology and Philosophy of Science* (1965), 24–30.
- Stephen Cook and Phuong Nguyen. 2010. *Logical foundations of proof complexity*. Vol. 11. Cambridge University Press Cambridge.
- Stephen A. Cook. 1975. Feasibly Constructive Proofs and the Propositional Calculus (Preliminary Version). In *STOC*. ACM, 83–97. <https://doi.org/10.1145/800116.803756>
- Lance Fortnow and Rahul Santhanam. 2011. Infeasibility of instance compression and succinct PCPs for NP. *J. Comput. Syst. Sci.* 77, 1 (2011), 91–106. <https://doi.org/10.1016/j.jcss.2010.06.007>
- Karthik Gajulapalli, Alexander Golovnev, Satyajeet Nagargoje, and Sidhanth Saraoji. 2023. Range Avoidance for Constant-Depth Circuits: Hardness and Algorithms. *CoRR* abs/2303.05044 (2023). <https://doi.org/10.48550/arXiv.2303.05044> arXiv:2303.05044
- Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. 2016. Candidate Indistinguishability Obfuscation and Functional Encryption for All Circuits. *SIAM J. Comput.* 45, 3 (2016), 882–929. <https://doi.org/10.1137/14095772X>
- Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. 2013. Witness encryption and its applications. In *STOC*. ACM, 467–476. <https://doi.org/10.1145/2488608.2488667>
- Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. 2016. Revisiting the Cryptographic Hardness of Finding a Nash Equilibrium. In *CRYPTO (2) (Lecture Notes in Computer Science, Vol. 9815)*. Springer, 579–604. https://doi.org/10.1007/978-3-662-53008-5_20
- Shafi Goldwasser and Michael Sipser. 1989. Private Coins versus Public Coins in Interactive Proof Systems. *Adv. Comput. Res.* 5 (1989), 73–90.
- Venkatesan Guruswami, Xin Lyu, and Xiuhan Wang. 2022. Range Avoidance for Low-Depth Circuits and Connections to Pseudorandomness. In *APPROX/RANDOM (LIPIcs, Vol. 245)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 20:1–20:21. <https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2022.20>
- Russell Impagliazzo, Valentine Kabanets, and Ilya Volkovich. 2018. The Power of Natural Properties as Oracles. In *Computational Complexity Conference (LIPIcs, Vol. 102)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:20. <https://doi.org/10.4230/LIPIcs.CCC.2018.7>
- Russell Impagliazzo and Avi Wigderson. 1997. $P = BPP$ if E Requires Exponential Circuits: Derandomizing the XOR Lemma. In *STOC*. ACM, 220–229. <https://doi.org/10.1145/258533.258590>

[32] Abhishek Jain and Zhengzhong Jin. 2022. Indistinguishability Obfuscation via Mathematical Proofs of Equivalence. In *FOCS*. IEEE, 1023–1034. <https://doi.org/10.1109/FOCS54457.2022.00100>

[33] Aayush Jain, Huijia Lin, and Amit Sahai. 2021. Indistinguishability obfuscation from well-founded assumptions. In *STOC*. ACM, 60–73. <https://doi.org/10.1145/3406325.3451093>

[34] Aayush Jain, Huijia Lin, and Amit Sahai. 2022. Personal Communication.

[35] Aayush Jain, Huijia Lin, and Amit Sahai. 2022. Indistinguishability Obfuscation from LPN over \mathbb{F}_p , DLIN, and PRGs in NC^0 . In *EUROCRYPT (1) (Lecture Notes in Computer Science, Vol. 13275)*. Springer, 670–699. https://doi.org/10.1007/978-3-031-06944-4_23

[36] Emil Jeřábek. 2004. Dual weak pigeonhole principle, Boolean complexity, and derandomization. *Ann. Pure Appl. Log.* 129, 1–3 (2004), 1–37. <https://doi.org/10.1016/j.apal.2003.12.003>

[37] Emil Jeřábek. 2005. *Weak pigeonhole principle, and randomized computation*. Ph. D. Dissertation. Ph. D. thesis, Faculty of Mathematics and Physics, Charles University, Prague.

[38] Emil Jeřábek. 2007. Approximate counting in bounded arithmetic. *J. Symb. Log.* 72, 3 (2007), 959–993. <https://doi.org/10.2178/jsl/1191333850>

[39] Emil Jeřábek. 2007. On Independence of Variants of the Weak Pigeonhole Principle. *J. Log. Comput.* 17, 3 (2007), 587–604. <https://doi.org/10.1093/logcom/exm017>

[40] Emil Jeřábek. 2009. Approximate counting by hashing in bounded arithmetic. *J. Symb. Log.* 74, 3 (2009), 829–860. <https://doi.org/10.2178/jsl/1245158087>

[41] Valentine Kabanets and Jin-yi Cai. 2000. Circuit minimization problem. In *STOC*. ACM, 73–79. <https://doi.org/10.1145/335305.335314>

[42] Erfan Khaniki. 2022. Nisan-Wigderson Generators in Proof Complexity: New Lower Bounds. In *CCC (LIPIcs, Vol. 234)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 17:1–17:15. <https://doi.org/10.4230/LIPIcs.CCC.2022.17>

[43] Erica Klarreich. 2020. Computer Scientists Achieve ‘Crown Jewel’ of Cryptography. *Quanta Magazine* (Nov 2020).

[44] Robert Kleinberg, Oliver Korten, Daniel Mitropolsky, and Christos H. Papadimitriou. 2021. Total Functions in the Polynomial Hierarchy. In *ITCS (LIPIcs, Vol. 185)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 44:1–44:18. <https://doi.org/10.4230/LIPIcs.ITCS.2021.44>

[45] Adam R. Klivans and Dieter van Melkebeek. 2002. Graph Nonisomorphism Has Subexponential Size Proofs Unless the Polynomial-Time Hierarchy Collapses. *SIAM J. Comput.* 31, 5 (2002), 1501–1526. <https://doi.org/10.1137/S0097539700389652>

[46] Ker-I Ko. 1986. On the Notion of Infinite Pseudorandom Sequences. *Theor. Comput. Sci.* 48, 3 (1986), 9–33. [https://doi.org/10.1016/0304-3975\(86\)90081-2](https://doi.org/10.1016/0304-3975(86)90081-2)

[47] Andrei N Kolmogorov. 1965. Three approaches to the quantitative definition of information. *Problems of information transmission* (1965). <https://doi.org/10.1080/00207166808803030>

[48] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yoge. 2022. One-Way Functions and (Im)perfect Obfuscation. *SIAM J. Comput.* 51, 6 (2022), 1769–1795. <https://doi.org/10.1137/15m1048549>

[49] Oliver Korten. 2021. The Hardest Explicit Construction. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS*. IEEE, 433–444. <https://doi.org/10.1109/FOCS52979.2021.00051>

[50] Oliver Korten. 2022. Derandomization from Time-Space Tradeoffs. In *CCC (LIPIcs, Vol. 234)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 37:1–37:26. <https://doi.org/10.4230/LIPIcs.CCC.2022.37>

[51] Jan Krajíček. 1995. *Bounded arithmetic, propositional logic, and complexity theory*. Encyclopedia of mathematics and its applications, Vol. 60. Cambridge University Press.

[52] Jan Krajíček. 2011. On the Proof Complexity of the Nisan-Wigderson Generator based on a Hard $\text{NP} \cap \text{coNP}$ function. *J. Math. Log.* 11, 1 (2011). <https://doi.org/10.1142/S0219061311000979>

[53] Jan Krajíček. 2019. *Proof complexity*. Vol. 170. Cambridge University Press.

[54] Jan Krajíček. 2021. Small Circuits and Dual Weak PHP in the Universal Theory of p -time Algorithms. *ACM Trans. Comput. Log.* 22, 2 (2021), 11:1–11:4. <https://doi.org/10.1145/3446207>

[55] Jan Krajíček. 2022. On the existence of strong proof complexity generators. *Electron. Colloquium Comput. Complex.* TR22-120 (2022). <https://eccc.weizmann.ac.il/report/2022/120>

[56] Jan Krajíček and Igor Carboni Oliveira. 2017. Unprovability of circuit upper bounds in Cook’s theory PV. *Log. Methods Comput. Sci.* 13, 1 (2017). [https://doi.org/10.23638/LMCS-13\(1:4\)2017](https://doi.org/10.23638/LMCS-13(1:4)2017)

[57] Jan Krajíček, Pavel Pudlák, and Gaisi Takeuti. 1991. Bounded Arithmetic and the Polynomial Hierarchy. *Ann. Pure Appl. Log.* 52, 1–2 (1991), 143–153. [https://doi.org/10.1016/0168-0072\(91\)90043-L](https://doi.org/10.1016/0168-0072(91)90043-L)

[58] Dai Tri Man Le and Stephen A. Cook. 2011. Formalizing Randomized Matching Algorithms. *Log. Methods Comput. Sci.* 8, 3 (2011). [https://doi.org/10.2168/LMCS-8\(3:5\)2011](https://doi.org/10.2168/LMCS-8(3:5)2011)

[59] Peter Bro Miltersen and N. V. Vinodchandran. 2005. Derandomizing Arthur-Merlin Games using Hitting Sets. *Comput. Complex.* 14, 3 (2005), 256–279. <https://doi.org/10.1007/s00037-005-0197-7>

[60] 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>

[61] Kerry Ojakian. 2004. *Combinatorics in bounded arithmetic*. Ph. D. Dissertation. Carnegie Mellon University.

[62] Ján Pich. 2015. Circuit lower bounds in bounded arithmetics. *Ann. Pure Appl. Log.* 166, 1 (2015), 29–45. <https://doi.org/10.1016/j.apal.2014.08.004>

[63] Ján Pich. 2015. Logical strength of complexity theory and a formalization of the PCP theorem in bounded arithmetic. *Log. Methods Comput. Sci.* 11, 2 (2015). [https://doi.org/10.2168/LMCS-11\(2:8\)2015](https://doi.org/10.2168/LMCS-11(2:8)2015)

[64] Ján Pich and Rahul Santhanam. 2021. Strong co-nondeterministic lower bounds for NP cannot be proved feasibly. In *STOC*. ACM, 223–233. <https://doi.org/10.1145/3406325.3451117>

[65] 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>

[66] Hanlin Ren, Rahul Santhanam, and Zhiqun Wang. 2022. On the Range Avoidance Problem for Circuits. In *FOCS*. IEEE, 640–650. <https://doi.org/10.1109/FOCS54457.2022.000067>

[67] Amit Sahai and Brent Waters. 2021. How to Use Indistinguishability Obfuscation: Deniable Encryption, and More. *SIAM J. Comput.* 50, 3 (2021), 857–908. <https://doi.org/10.1137/15M1030108>

[68] Michael Sipser. 1983. A Complexity Theoretic Approach to Randomness. In *STOC*. ACM, 330–335. <https://doi.org/10.1145/800061.808762>

[69] Michael Sipser. 1997. *Introduction to the theory of computation*. PWS Publishing Company.

[70] 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>

[71] Rotem Tsabary. 2022. Candidate Witness Encryption from Lattice Techniques. In *CRYPTO (1) (Lecture Notes in Computer Science, Vol. 13507)*. Springer, 535–559. https://doi.org/10.1007/978-3-031-15802-5_19

[72] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. 2022. Witness Encryption and Null-IO from Evasive LWE. In *ASIACRYPT (1) (Lecture Notes in Computer Science, Vol. 13791)*. Springer, 195–221. https://doi.org/10.1007/978-3-031-22963-3_7

Received 2022-11-07; accepted 2023-02-06