

Student-Teacher Constructive Separations and (Un)Provability in Bounded Arithmetic: Witnessing the Gap

Marco Carmosino, Stefan Grosser

November 16, 2024

Abstract

Let \mathcal{C} be a complexity class and A be a language. The statement “ $A \notin \mathcal{C}$ ” is a *separation* of A from \mathcal{C} . A separation is *constructive* if there is an efficient algorithm called a *refuter* that prints counterexamples to the statement “ M decides A ” for every \mathcal{C} -algorithm M . Concretely, refuters witness errors of M on A by printing, on input 1^n , an n -bit string x such that $M(x) \neq A(x)$. Many recent breakthroughs in lower bounds and derandomization, like the algorithmic method [13], rely on constructive separations as a core component. Chen, Jin, Santhanam, and Williams [11] studied the consequences of constructivizing classical non-constructive lower bounds in complexity theory. They showed that (1) constructivizing many known separations would imply breakthrough lower bounds, and (2) some separations are impossible to constructivize.

We study a more general notion of “efficient refutation” in terms of *\mathcal{C} -Student-Teacher Games*, where the \mathcal{C} -refuter (Student) is allowed to adaptively propose candidate counterexamples x_i to an omniscient Teacher. If x_i fails to witness an error, Teacher reveals a counterexample y_i to the statement “ x_i is a counterexample to the statement ‘ M decides A ’ ” — the nature of y_i depending on how the separated language A and complexity class \mathcal{C} are defined. We show:

- If there is a P-Student-Teacher constructive separation of Palindromes from one-tape nondeterministic $O(n^{1+\varepsilon})$ time [38], then $\text{NP} \not\subseteq \text{SIZE}[n^k]$ for every k .
- If there is a uniform $\text{AC}^0[\text{qpoly}]$ -Student-Teacher protocol generating truth tables of super fixed polynomial circuit complexity, then $\text{P} \neq \text{NP}$.
- There is no P-Student-Teacher protocol which for infinitely many $c > 0$, generates high- K^{n^c} strings.

Our results imply a conditional separation of Jeřábek’s theory VAPC from V^1 , a theory equivalent to Buss’s theory S_2^1 . This improves and significantly simplifies the work of Ilango, Li, and Williams [24], who separate VAPC from the weaker theory VPV under the existence of indistinguishability obfuscation. We do not use cryptographic assumptions in our separation. Instead we introduce a natural and plausible conjecture on the uniformity of proofs in bounded arithmetic, inspired by Kreisel’s Conjecture in logic. We believe this conjecture to be of independent interest.

1 Introduction

Constructive lower bounds are a key concern of complexity theory. We know that hard functions *exist*, but not how to *exhibit* them efficiently. There are two ways to formalize the notion of constructivity. The *algorithmic* perspective asks for the computational complexity of searching for witnesses to a complexity lower bound, like hard truth tables. The *proof-theoretic* perspective asks for the weakest logical theory that proves complexity lower bounds. This paper obtains new results about both formulations of constructivity and the relationship between them.

In particular, we study the computational model of *Student-Teacher Games*, which links the proof-theoretic and algorithmic perspectives on constructivity. Roughly, if a complexity lower bound LB is provable in a “bounded” logical theory, then “efficient” Student-Teacher games witnessing LB follow. We obtain new results about Student-Teacher games witnessing: lower bounds for deciding Palindromes, existence of time-bounded Kolmogorov-random strings, and existence of Boolean functions of high circuit complexity. From these results, we conditionally derive (1) consequences about the provability of these statements and (2) separations of expressive and well-studied logical theories. We structure our introduction as follows:

- (i) Building on prior work studying the consequences of algorithmic constructivity in complexity theory [11, 27, 49], we show that efficient Student-Teacher games witnessing known complexity lower bounds imply breakthrough lower bounds.
- (ii) We scrutinize the relationship between algorithmic constructivity and proof-theoretic constructivity. The natural translation into bounded arithmetic of complexity-theoretic statements that mention “polynomially bounded resources” results in a *schema* of logical sentences, one for each *fixed* polynomial. This disrupts the well-known connection between provability of lower bounds and witnessing Student-Teacher games.
- (iii) We identify a new family of conjectures called *Witnessing Hypotheses for Uniform Proofs* which give Student-Teacher witnessing from a *schema* of lower bounds. We demonstrate that these conjectures are well-founded, and connect them to the famous Kreisel Conjecture in mathematical logic.
- (iv) As a consequence of these conjectures, we give the first known conditional separation between the well-studied bounded arithmetic theories VAPC and V^1 (equivalently APC_1 and S_2^1). Moreover, we do so *without* cryptographic assumptions. This constitutes substantial progress towards understanding the necessary tools needed to show unprovability in bounded arithmetic.

In the remainder of this introduction we give context, motivation, a more detailed description of our results, and a list of open problems about constructive complexity theory.

1.1 Algorithmic Constructivity in Complexity Theory

Underpinning many recent developments in complexity theory is the notion of *constructive* lower bounds. Namely, algorithms for solving refutation problems and avoidance search problems.

1. *Refutation* — If a lower bound holds for a problem Π against a model of computation M , then the Π -*Refutation for M* problem is: given an algorithm A from M and a number n , print a string of length n for which A fails to solve Π correctly; i.e. a counterexample to the claim that “ A solves Π .”
2. *Avoidance* — Fix $\Lambda = \bigcup_{n \in \mathbb{N}} \Lambda_n$ an infinite set of compressible strings, where each Λ_n denotes the n -bit strings described by a particular set of bounded-complexity devices, such as Boolean circuits of size at most $\log(n)^2$. The Λ -*Avoid problem* is then: given a number n , print an n -bit string outside Λ_n . That is, print a counterexample to the claim “every n -bit string is compressed by a Λ -device.”

Refutation has been an explicit object of study since Kabanets [26] introduced refuters to give an unconditional weak derandomization of RP. Since then, upper bounds on refuters have been a driving force behind derandomization and lower bounds. A seminal example is the algorithmic method of Williams to give lower bounds against ACC^0 . These lower bounds [51, 13] crucially use a refuter for the NTIME hierarchy theorem, with [13] in particular using an almost-everywhere refuter of Fortnow and Santhanam [19] against

47 NTIMEGUESS $[T(n), O(n)]$, for time-constructive $T(n)$. Refutation has also been recently shown by Chen,
 48 Tell, and Williams [14] to unify many previous techniques that give conditional derandomization results.

49 Avoidance search problems are also intimately tied with circuit lower bounds and derandomization.
 50 Korten [27] showed that many important explicit construction problems (e.g. computing hard truth tables,
 51 rigid matrices, and high time-bounded Kolmogorov complexity strings) are all reducible to the avoidance
 52 problem EMPTY.

53 EMPTY: Given circuit $C: \{0, 1\}^m \rightarrow \{0, 1\}^n$, with $n > m$ output a string $x \in \{0, 1\}^n$ outside the range of C .

54 This shows that circuit lower bounds and derandomization are implied by fast algorithms for EMPTY.
 55 Ren, Santhanam, and Wang [49] deepened this connection by studying algorithms for \mathcal{C} -Avoid, parametrized
 56 by a circuit class \mathcal{C} . Finally, Chen, Hirahara, and Ren [10], and a follow-up by Li [35], used Korten’s reduction
 57 from finding hard truth tables to EMPTY in order to give truly exponential circuit lower bounds for $\mathsf{S}_2\mathsf{E}$.

58 **Constructive Separations.** To formally study the power and limitations of constructive lower bounds,
 59 Chen, Jin, Santhanam, and Williams [11] asked what happens if you can convert several classical *non-*
 60 *constructive* lower bounds into constructive ones? Their definition of constructivity goes through efficient
 61 Refutation algorithms, and implicitly Avoidance algorithms.¹

62 **Definition 1.1** (\mathcal{C} -Refuter). Let $f: \{0, 1\}^* \rightarrow \{0, 1\}$ be a function and let A be an algorithm. The refutation
 63 search problem $\text{Ref}_{f,A} := \{(n, x) \mid x \in \{0, 1\}^n \text{ and } f(x) \neq A(x)\}$ asks to find an input x where A disagrees
 64 with function f . An algorithm $R(1^n)$ is a \mathcal{C} -refuter against A if $R \in \mathcal{C}$ and for infinitely many n , $(n, R(1^n)) \in$
 65 $\text{Ref}_{R,A}$.

66 **Definition 1.2** (\mathcal{C} -Constructive Separation). For complexity classes $\mathcal{A}, \mathcal{B}, \mathcal{C}$, a separation $\mathcal{B} \not\subseteq \mathcal{A}$ is called
 67 \mathcal{C} -constructive if for some language \mathcal{L}_B decidable in \mathcal{B} and any proposed algorithm $A \in \mathcal{A}$ that decides \mathcal{L}_B ,
 68 there is a \mathcal{C} -refuter R_A .

69 Chen et. al [11] gave several insights on constructive separations. First, they showed that a P -constructive
 70 separation for the classic Palindromes lower bound of Maass [38] implies a major complexity separation.

71 **Theorem 1.3** (Theorem 3.4, [11]). If Maass’s lower bound against deciding Palindromes with one-tape
 72 nondeterministic Turing machines of subquadratic time can be made P -constructive, then $\mathsf{E} \not\subseteq \text{SIZE}[2^{\delta n}]$, for
 73 some $\delta > 0$.

74 They also showed that efficient Avoid algorithms imply complexity separations.

75 **Theorem 1.4** (Implicit to Theorem 1.7(i), [11]). If there is a uniform $\text{AC}^0[\text{qpoly}]$ algorithm solving Avoid
 76 for circuits of size $s(n) = n^{(\log n)^{\omega(1)}}$, then $\mathsf{P} \neq \mathsf{NP}$.

77 However, not all known lower bounds can be made constructive. Chen et. al. observed that there can
 78 be no polynomial time algorithm which on input 1^n , outputs an n -bit string of high- K^{poly} complexity (see
 79 Proposition 4.4). This contrasts in a peculiar way with the lower bounds of Theorem 1.3 and Theorem
 80 1.4. Chen et. al. argue that understanding better which lower bounds are likely to be constructive or
 81 non-constructive will be key to progress in complexity theory. See their paper for more details.

82 1.2 Our Results: Student-Teacher Constructive Separations

83 In this paper, we generalize the results of Chen et. al. to the setting of *Student-Teacher* refuters.

84 **Definition 1.5** (Student-Teacher Game (*Informal*)). Let $\varphi(\overline{X}) = \forall Y \theta(\overline{X}, Y)$, for θ a quantifier-free formula,
 85 and let $p(n), q(n)$ be polynomials. We say $S(1^n)$ is a \mathcal{C} -*Student-Teacher game* if S is an algorithm in \mathcal{C} with
 86 access to a *counterexample oracle* $CX[\varphi]$ which given an $\overline{X} \in \{0, 1\}^{p(n)}$, either responds “YES” or returns
 87 a $Y \in \{0, 1\}^{q(n)}$ such that $\theta(\overline{X}, Y)$ is false. We further write $CX[\varphi, r(n)]$ for a function $r(n)$ to indicate S
 88 gets access to $r(n)$ calls to the oracle CX .

¹An avoidance algorithm can be thought of as a refuter for the “always-YES” algorithm against some hard language. For example, a polynomial time algorithm for solving Avoid for circuits of size $s(n)$ gives a polynomial time algorithm refuting the “always-YES” algorithm for $\text{MCSP}[s(n)]$.

89 Student-Teacher games are a natural model of computation which appear in learning algorithms and
90 bounded arithmetic. For many complexity lower bounds, provability in a weak logical theory implies an
91 efficient Student-Teacher refuter, rather than a standard refuter as seen in [11]. This means that to study
92 the (un)provability of complexity lower bounds, it is necessary to study Student-Teacher games. We make
93 this connection more explicit in the following section. See a more detailed definition of Student-Teacher
94 games in Section 2.6.

95 As an example of a Student-Teacher game, consider a P-Student-Teacher refuter S_M solving the refutation
96 search problem $\text{Ref}_{\text{Pal}, M}$, with M a one-tape subquadratic time nondeterministic Turing machine. Here,
97 $\varphi_{\text{Pal}}(X, W^*)$ expresses the following formula:

$$98 \quad \varphi_{\text{Pal}}(X, W^*) \triangleq \text{“For every witness } W \text{ of length } |X|^{1.1}, [M(X, W) = 0 \text{ and } X \text{ is a palindrome}] \text{ or} \\ 99 \quad [M(X, W^*) = 1 \text{ and } X \text{ is not a palindrome.]}”$$

100 Each round, S would propose X, W^* to the counterexample oracle CX , where either CX says “YES” if
101 X is an input that M fails to decide whether X is a palindrome, or CX responds to S with a witness W
102 such that M does correctly decide X .

103 **Palindromes.** We generalize Theorem 1.3 to P-Student-Teacher refuters.

104 **Theorem 1.6.** If for any nondeterministic one-tape subquadratic time Turing machine M there is a P-
105 Student-Teacher game $S_M(1^n)$ with counterexample oracle $CX[\varphi_{\text{Pal}}, O(1)]$ solving $\text{Ref}_{\text{Pal}, M}$ for n -bit inputs,
106 then $\text{NP} \not\subseteq \text{SIZE}[n^k]$ for any $k \geq 0$.

107 A P-Student-Teacher refuter is considerably stronger than a P-refuter in the context of Palindromes
108 lower bounds. The oracle $CX[\varphi_{\text{Pal}}]$ acts as a restricted NP-oracle, as finding a witness W where $M(X, W)$ is
109 correct is an NP-language. Nevertheless, we show that P-Student-Teacher refuters for one-tape subquadratic
110 NTMs deciding Pal still imply breakthrough circuit lower bounds.

111 **Weak Shannon Counting.** We give a slightly orthogonal result to Theorem 1.4. Here, we consider
112 avoidance algorithms for weak Shannon counting. Namely, for a fixed $b \in \mathbb{N}$ and given an input 1^N , output
113 a truth-table of length N which is not computed by a size $(\log N)^b$ Boolean circuit. Let $\varphi_{WSC}(X, b)$ express
114 the following formula:

$$115 \quad \varphi_{WSC}(X, b) \triangleq \text{“For every circuit } C \text{ of size } (\log |X|)^b, \text{ the truth table generated by } C \text{ disagrees with } X \text{ on} \\ 116 \quad \text{some bit.}”$$

117 We introduce the notion of an *absolute* Student-Teacher game, which solves $\text{SIZE}[n^b]$ -Avoid for *infinitely*
118 *many* b . Meaning, the student S takes two inputs: 1^N and b (represented in binary) and solves $\text{SIZE}[n^b]$ -
119 Avoid.

120 **Theorem 1.7.** If there is an absolute poly log-uniform $\text{AC}^0[\text{qpoly}]$ -Student-Teacher protocol $S(1^N, b)$ with
121 oracle $CX[\varphi_{WSC}, O(1)]$ solving $\text{SIZE}[n^b]$ -Avoid for infinitely many $b \in \mathbb{N}$, then $\text{P} \neq \text{NP}$.

122 If you simply fix a $b \in \mathbb{N}$ instead of having an absolute Student-Teacher protocol, then such a Student-
123 Teacher protocol *does* exist (see Section 5). However, we only get our consequence $\text{P} \neq \text{NP}$ for an absolute
124 Student-Teacher protocol.

125 Our notion of an absolute Student-Teacher protocol appears naturally in the context of bounded arith-
126 metic and witnessing theorems. See Section 1.4 and Section 4.3 for a discussion.

127 **High- K^{poly} Strings.** We now give an avoidance problem which provably has no Student-Teacher game.
128 Let $\varphi_{K^\epsilon}(X, b)$ express the following formula:

$$129 \quad \varphi_{K^\epsilon}(X, b) \triangleq \text{“For every Turing machine and advice pair } (M, \alpha) \text{ of description length } |X|/4, \text{ running } M \text{ for} \\ 130 \quad n^b \text{ steps with advice } \alpha \text{ has output disagreeing with } X.”$$

131 Let $\text{BHaltDesc}[n^c, p(n)]$ be the class of Turing machine and advice pairs (M, α) of total description length
132 $p(n)$ which, starting with α on the tape of M , runs for n^c time. We then have the following.

133 **Theorem 1.8.** There is no absolute P-Student-Teacher protocol $S(1^n, b)$ (b given in unary) with oracle
134 $CX[\varphi_{K^\epsilon}, \text{poly}(n)]$ solving $\text{BHaltDesc}[n^b, n/4]$ -Avoid.

1.3 Proof Theoretic Constructivity in Complexity Theory

Our results on Student-Teacher constructive separations have several consequences on the provability of lower bounds in bounded arithmetic. Bounded arithmetic is a related and more fine-grained notion of constructivity that comes from not just the algorithms for refutation and avoidance, but also from the logical expressivity needed to prove the correctness of these algorithms.

Bounded Arithmetic. Bounded arithmetic studies fragments of Peano Arithmetic (PA) which use reasoning inherent to computational complexity classes. The earliest example is the theory $\text{I}\Delta_0$, introduced by Parikh [44]. He showed that reasoning in $\text{I}\Delta_0$ corresponds to the Linear Time Hierarchy (LTH), and that certain operations like exponentiation are infeasible in this theory. One of the most important and well-studied bounded arithmetic theories is Cook’s theory VPV. It is an equivalent version of Cook’s original theory, PV_1 , defined in his seminal 1975 paper [18]. This theory was the first proposed to exactly characterize polynomial-time computation and reasoning, and more generally was the first theory introduced to explicitly connect standard complexity classes and bounded arithmetic. PV stands for *polynomially verifiable*, and one of Cook’s original motivations for defining this theory is that when VPV proves a statement like $\forall X \varphi(X)$, there is a polynomial time algorithm VERIFY_φ which on input Y verifies $\varphi(Y)$ holds.

This constructive property is known as witnessing. If a theory T proves the existence of some object, then this implies there is an efficient algorithm that generates this object. As an example, suppose VPV were to prove the following Π_2 statement describing a circuit lower bound for language \mathcal{L} , which is decided by machine M :

“For every input length n and circuit $C \in \mathcal{C}$, there exists an input x of length n such that $C(x) \neq M(x)$ ”

Then, the witnessing property for VPV says there is a polynomial time algorithm which finds an incorrect x when given n, C as input. Notice that this is a P-refuter!

While a provable Π_2 statement directly translates into a refuter, the situation is more complicated for Π_i statements with $i \geq 3$. Focusing on $i = 3$, witnessing properties give Student-Teacher protocols. For a Π_3 statement $\forall n \exists X \forall Y \theta(n, X, Y)$, a witnessing Student-Teacher protocol S would take as input 1^n , and query the counterexample oracle $CX[\varphi]$ on guesses for a satisfying X , for $\varphi = \forall Y \theta(n, X, Y)$. See Sections 2.4 and 2.5 for more details on witnessing theorems.

Π_3 formulas naturally encode many refutation and avoidance type statements. Maass’s Palindromes lower bound, Shannon counting, and the existence of High- K^{poly} strings are all examples. Hence, the (un)provability of these statements in bounded arithmetic is closely tied to Student-Teacher constructive separations.

A Gap Between Constructive Separations and Provability. In Chen et. al. [11], they noted a gap between constructive separations and provability. A P-constructive separation of $\mathcal{B} \not\leq \mathcal{A}$ does not at all guarantee that $\text{VPV} \vdash \text{“}\mathcal{B} \not\leq \mathcal{A}\text{”}$. This can be for several reasons: the refuter might not be *provably correct* inside VPV, or $\mathcal{B} \not\leq \mathcal{A}$ might only be formalizable as a Π_3 formula, which by witnessing gives a Student-Teacher game with an polynomial time student, rather than just a P-refuter. For these reasons, the consequences of a (non)constructive separation of $\mathcal{B} \not\leq \mathcal{A}$ may not have any bearing on the (un)provability of the same lower bound $\mathcal{B} \not\leq \mathcal{A}$.

An explicit example of this gap, given by [17, 12], is proving the correctness of the AKS primality testing algorithm [1]. We can formalize the correctness of it as the following formula:

$$\forall n [\text{AKS}(n) = 1 \leftrightarrow \forall 1 < d < n, d \nmid n],$$

where AKS is a function symbol for the AKS primality testing algorithm. If this statement were provable in VPV, then there would be a polynomial time algorithm which on input 1^n , n a composite number, would be able to determine a factor of n . Hence, VPV proving the correctness of AKS would imply that factoring has a polynomial time algorithm. This shows that the proof of correctness in [1] of their polynomial time algorithm AKS uses functions which are themselves not polynomial time computable.

179 **Formalizing Lower Bounds as Schemas.** Straightforward translations of many complexity lower bounds
 180 into the language of bounded arithmetic require *schemas* of formulas: a sequence of formulas indexed by
 181 substitution of fixed polynomials $\{n^c\}_{c \in \mathbb{N}}$. Theories of bounded arithmetic cannot quantify over arbitrary
 182 polynomials. A simple example of this would be the Deterministic Time Hierarchy theorem.

183 $\text{DTIMEH}(c) \triangleq$ “For every sufficiently large n and Turing machine M , there exists an input X of length n
 184 where simulating M on X for n^c time incorrectly decides hard language \mathcal{L}_{c+1} ”

185 For each $c \in \mathbb{N}$, you get a new formula $\text{DTIMEH}(c)$ describing that $\text{DTIME}[n^{c+1}] \not\subseteq \text{DTIME}[n^c]$. This is
 186 necessary in bounded arithmetic as exponentiation is not a feasible operation, so a single formula DTIMEH
 187 quantifying over all c is not possible in a theory like VPV . The same issue also occurs when writing down,
 188 say, $\text{P} \neq \text{NP}$ in the language of VPV .

189 This poses a problem when trying to study the provability of lower bounds. Applying witnessing to a
 190 schema of Π_3 formulas $\Psi[c]$ would result in a schema of Student-Teacher games, each solving a different
 191 search problem parametrized by c . Further, a Student-Teacher game for $\Psi[c_0]$ is not required to share any
 192 structure or runtime with a Student-Teacher game for $\Psi[c_1]$, with $c_0 \neq c_1$. This means that our results on
 193 *absolute* Student-Teacher games do not automatically imply provability consequences.

194 1.4 Our Results: Consequences in Bounded Arithmetic

195 We initiate the study of the following natural question about lower bound schemas.

196 **Question 1.9.** Let $\text{LB}(c)$ be the logical translation of some complexity theoretic lower bound, parametrized
 197 by $c \in \mathbb{N}$. If $\text{VPV} \vdash \text{LB}(c)$, for every c , then does VPV use the same “proof” for every c ?

198 While it is unclear at all if $\text{VPV} \vdash \text{“P} \neq \text{NP”}$, it is known that $\text{VPV} \vdash \text{DTIMEH}(c)$, for every $c \in \mathbb{N}$.
 199 Amazingly, VPV could use “the same” proof for every c ! This follows as the *refuter* for DTIMEH is completely
 200 agnostic to c , and hence is the same regardless of the value of c . Specifically, there is a hard language
 201 $\mathcal{L} \in \text{DTIME}[n^{c+1}] \setminus \text{DTIME}[n^c]$ where a refuter runs linearly in n to construct a counterexample of a proposed
 202 machine $M \in \text{DTIME}[n^c]$ deciding \mathcal{L} . Does this property of the Deterministic Time Hierarchy theorem hold
 203 more generally for other complexity lower bounds?

204 We denote by ‘Witnessing Hypothesis for Uniform Proofs’ (WHUP) that such a phenomenon in fact
 205 holds, and that for certain classes of lower bound schemas, if a theory \mathcal{T} proves the schema, then it does so
 206 with the same proof.

207 **Hypothesis 1.10** (WHUP for theory VPV (Informal)). Let $\text{VPV} \vdash \forall n \exists X \forall Y \theta(n, X, Y, n^c)$, for a quantifier-
 208 free θ and for infinitely many $c \in \mathbb{N}$, and let $\varphi(n, X, c) = \forall Y \theta(n, X, Y, n^c)$. Then there is a witnessing *absolute*
 209 Student-Teacher game $S(1^n, c)$ which, for infinitely many c , finds a satisfying X of length n using $O(1)$ oracle
 210 calls to $CX[\varphi]$.

211 WHUPs can be used to connect the (un)provability of *schemas* of formulas in bounded arithmetic with
 212 *absolute* Student-Teacher constructive separations. Our Witnessing Hypotheses are similar to and inspired by
 213 a conjecture of Kreisel for Peano Arithmetic. See Section 4 for a detailed discussion where we carefully define
 214 Witnessing Hypotheses and provide many supporting examples for the validity of WHUPs. We conclude
 215 this section with a sketch of each of our results on provability.

216 **Palindromes.** First, we extend Theorem 1.6 to provability in VPV . Let Pal be a Π_3 formula expressing
 217 Maass’s lower bound.

218 **Theorem 1.11.** If $\text{VPV} \vdash \text{Pal}$, then $\text{NP} \not\subseteq \text{SIZE}[n^k]$.

219 This complements recent work of Chen et. al [12], showing that if Maass’s lower bound is provable in
 220 VPV , then collision resistant hash functions do not exist.

221 **Weak Shannon Counting.** Building on the work of Thapen [50], Jeřábek [25], studied the theory $VAPC :=$
 222 $VPV + dWPHP(VPV)$ which adds the *dual weak pigeonhole principle*, the combinatorial principle behind
 223 $EMPTY$ and \mathcal{C} -Avoid. He showed $VAPC$ has an intimate relationship with randomized complexity (ZPP) and
 224 approximate counting. Namely, all provably total functions in $VAPC$ are contained in ZPP , and conversely,
 225 a large natural subclass of ZPP is definable in $VAPC$. It is possible that $VAPC$ may completely characterize
 226 ZPP , but this would require showing that ZPP has a complete problem [50].

227 $VAPC$ is interesting because of its ability to formalize most known complexity lower bounds. Jeřábek
 228 showed that it can formalize Shannon counting arguments, and Müller and Pich [41] further illustrated its
 229 power by formalizing Parity lower bounds and the method of approximations. Recent results in *unprovability*
 230 have also been shown. Chen et. al [12] showed that if collision resistant hash functions exist, then Maass's
 231 palindromes lower bound is not formalizable in $VAPC$.

232 We give a result orthogonal to Jeřábek's provability of Shannon counting in $VAPC$. We introduce a theory
 233 $V_{\#}^0$ to characterize quasipolynomial AC^0 reasoning. This theory is incomparable to $VAPC$, but we show it is
 234 also capable of proving weak Shannon counting.

235 **Lemma 1.12.** Let $b \in \mathbb{N}$. $V_{\#}^0$ proves the existence of truth tables not computable by Boolean circuits of
 236 size n^b .

237 Further, under a WHUP for $V_{\#}^0$, we have consequences for the provability of hard truth tables.

238 **Theorem 1.13.** Assuming a WHUP for the theory $V_{\#}^0$, if $V_{\#}^0$ proves for every $b \in \mathbb{N}$ that there are truth
 239 tables of hard for size n^b circuits, then $P \neq NP$.

240 We then get as a clear corollary,

241 **Corollary 1.14.** A WHUP for $V_{\#}^0$ implies that $P \neq NP$.

242 See Theorem 5.10 for full details.

243 **Conditional Separation of V^1 and $VAPC$** We show the following surprising unprovability result.

244 **Theorem 1.15.** Under a Witnessing Hypothesis, VPV (or even V^1) cannot show the existence of High- K^{n^b}
 245 strings, for almost every $b \in \mathbb{N}$.

246 As a corollary, we conditionally separate theories V^1 and $VAPC$ (equivalently S_2^1 and APC_1).

247 **Theorem 1.16.** Under a Witnessing Hypothesis, $VAPC$ is not equivalent to V^1 .

248 *Proof.* In the work of Korten [27], it was shown that APC_1 proves the existence of high- K^{poly} strings. However,
 249 under a Witnessing Hypothesis (Hypothesis 4.11), V^1 does not show these strings exist. \square

250 This improves and greatly simplifies the result of Ilango et. al [24] separating $VAPC$ and VPV under the
 251 existence of indistinguishability obfuscation and $NP \neq coNP$. Our result is also the first such conditional
 252 separation between any bounded arithmetic theory \mathcal{T} and $VAPC$ which uses a plausible² non-cryptographic
 253 assumption.

254 Separating theories of bounded arithmetic should be far easier to prove than demonstrating the existence
 255 of cryptographic objects like collision resistant hash functions or indistinguishability obfuscation. It is then
 256 desirable to give conditional separations of theories using assumptions much weaker than cryptography. We
 257 believe Witnessing Hypotheses are such an assumption. However, Corollary 1.14 indicates that WHUPs are
 258 still quite strong, as some will imply major complexity separations. Is this the case with a WHUP for VPV ?
 259 Are WHUPs in fact $EQUAL$ to the existence of some cryptographic object?

²In [29], Krajíček showed that assuming Kolmogorov's Conjecture, $P \subset SIZE[n^k]$ for some fixed k , then $VAPC$ is strictly stronger than VPV . However, Kolmogorov's Conjecture is widely believed to be false.

1.5 Our Techniques.

Constructive Separations. We employ the general strategy of Chen et. al. [11] to show that efficient refuters imply circuit lower bounds.

- (i) Assume, for sake of contradiction, a complexity collapse (eg. $P = NP$ or $P \subset \text{SIZE}[n^k]$). Show that a \mathcal{C} -refuter from a \mathcal{C} -constructive separation of $\mathcal{B} \not\subseteq \mathcal{A}$ produces outputs of very small circuit complexity.
- (ii) Show that there exists a too efficient algorithm $M \in \mathcal{A}$ for a hard language $\mathcal{L}_{\mathcal{B}}$ which is correct on all inputs of low circuit complexity. This forms a contradiction.

As mentioned in Section 1.4, this argument is not sufficient on its own to discuss the consequences of provability of lower bounds, as provability and witnessing implies Student-Teacher refuters instead of standard refuters. To handle this, we introduce several novel *round collapse* techniques to remove the Teacher from Student-Teacher protocols. This gives a reduction to the proof strategy of [11].

Round Collapses. Round collapse techniques have seen widespread recent study to show the unprovability of Π_3 sentences in theories of bounded arithmetic [8, 24, 29, 9, 16, 46]. We continue this line of work by introducing three novel round collapse arguments.

A common issue with round collapse techniques is that they are very ad hoc and strongly depend on the discussed lower bound. In the case of Maass’s Palindromes lower bound, we introduce in Section 3 a very general technique to deal with collapsing a P-Student-Teacher protocol whose counterexample oracle $CX[\varphi, O(1)]$ solves an NP-language. Recall for a one-tape subquadratic time NTM M , $\varphi_{\text{Pal}}(X)$ certifies that for every witness W , $M(X, W) = 0$ when X is a palindrome. Hence the counterexample oracle $CX[\varphi_{\text{Pal}}]$ solves the NP language of determining a witness W' where $M(X, W') = 1$. Assuming $NP \subset \text{SIZE}[n^k]$, we may use the Easy Witness Lemma of Murray and Williams [42] to give a compressed description of W' . By providing this compression of W' as advice to the Student, we can replace a single query to $CX[\varphi_{\text{Pal}}]$. Repeating this argument allows the conversion of a P-Student-Teacher refuter into a $P/o(n)$ -refuter.

Our other round collapses are much more ad hoc. For weak Shannon counting and Theorem 1.7, we generalize the technique of Chen et. al. [11] to efficiently simulate a polylog-uniform $\text{AC}_d^0[\text{qpoly}]$ refuter $C(1^n)$ with a general sublinear size Boolean circuit, assuming $P = NP$. Their idea is to show that computing the i -th bit of $C(1^n)$ is a $\Sigma_d^P[\text{polylog}(n)]$ problem, which collapses to $\text{DTIME}[\text{polylog}(n)]$ under $P = NP$. We show this argument completely in Lemma 5.5. Where we must generalize this argument is to further allow a polylog-uniform $\text{AC}^0[\text{qpoly}]$ Student-Teacher refuter, and provide a method to remove the CX oracle gates. We do so in Section 5.

The round collapse for high- K^{poly} strings and Theorem 1.8 is conceptually the most natural. We take direct inspiration from the $\text{DTIME}[n]$ -constructive separation of $\text{DTIME}[n^{c+1}] \not\subseteq \text{DTIME}[n^c]$. The linear time refuter $R_M(1^n)$ simply outputs the padded source code of M , $\langle M \rangle \circ 0^{n-|\langle M \rangle|}$. Our observation in Section 4 is that for an absolute P-Student-Teacher protocol solving $\text{BHaltDesc}[n^c, n/4]$ -Avoid for all $c \in \mathbb{N}$, the source code of the Student is a valid response for the counterexample oracle $CX[\varphi_{\mathcal{K}^c}]$. We give a fine-grained reflection argument to generally transform a P-Student-Teacher protocol for $\text{BHaltDesc}[n^c, n/4]$ -Avoid into a polynomial time algorithm, even when polynomially many Teacher queries are made by the Student-Teacher protocol.

1.6 Comparison to Other Work.

AC^0 reasoning and Provable Circuit Lower Bounds. Several previous works have studied the provability of circuit lower bounds in bounded arithmetic via round collapses. Pich [46] showed unconditionally that the theory V^0 , corresponding to log-uniform $\text{AC}^0[\text{poly}]$ reasoning, cannot prove superpolynomial size circuit lower bounds. This contrasts with our Lemma 1.12, where we show that the theory $V_{\#}^0$ proves fixed polynomial size circuit lower bounds. This suggests an intriguing question of finding the exact logical strength necessary for proving fixed polynomial size circuit lower bounds.

305 **Separating VAPC and VPV.** Krajíček [29] gave the first conditional separation of VPV and VAPC via round
 306 collapse techniques. Unfortunately, his round collapse required the unlikely assumption that $P \subset \text{SIZE}[n^k]$.
 307 In [30], Krajíček called for reasonable assumptions under which VAPC is strictly stronger than VPV. This
 308 was achieved by Ilango et. al. [24], who showed that under indistinguishability obfuscation and $\text{NP} \neq \text{coNP}$,
 309 these theories are indeed separated. These conditional separations of [29, 24] were accomplished by studying
 310 the dual weak pigeonhole principle $\text{dWPHP}(\text{VPV})$ and Student-Teacher protocols for solving EMPTY .

311 Using a Witnessing Hypothesis, we conditionally separate VAPC from an even stronger theory V^1 , which
 312 contains VPV. Further, we make use of a weaker, *uniform* version of the dual weak pigeonhole principle.

313 1.7 Open Problems.

314 This work suggests several continuations and open problems. We provide two directions pertaining to
 315 (un)provability, and several towards understanding and proving WHUPs.

316 **Improving the Palindromes round collapse.** While we show that a constant round Student-Teacher
 317 refuter for Palindromes would imply $\text{NP} \not\subset \text{SIZE}[n^k]$ for any $k > 0$, we fall short of proving this for $\omega(1)$
 318 rounds. Is there a polynomial round Student-Teacher refuter for Palindromes? This could be used to extend
 319 our provability consequence to theories V^1/S_2^1 .

320 **Unprovability for APC^1 .** In Section 4, we show that under a Witnessing Hypothesis for VPV, generating
 321 high K^{poly} strings is not feasible in VPV. Can this be generalized to unprovability in APC^1 . This would likely
 322 have to be for a notion of *zero-error* time-bounded Kolmogorov complexity, which the authors are unaware
 323 of appearing in the present meta complexity literature.

324 **What do proofs look like?** Amongst our examples of the “absolute” witness phenomenon, like the refuter
 325 for DTIMEH , what do the VPV proofs *actually* look like? This would be a basic building block to understand
 326 before attempting to prove a WHUP. We emphasize that we know the proofs, but not a structural measure
 327 or property that makes it clear they are “the same” across different values of the parameter. Surprisingly
 328 simple polynomial schemas have proofs in VPV where we do not have a solid understanding of their structure.
 329 One example is,

$$\varphi(b, c) \triangleq \forall n \ c > b \rightarrow n^c > n^b.$$

330 As VPV is defined for the purpose of encapsulating polynomial time computation (rather than performing
 331 arithmetic), even simple arithmetic identities can have “complicated” proofs. Showing that the sequent
 332 calculus proofs of $\varphi(b, c)$ over VPV are the same for all $b, c \in \mathbb{N}$ would be of interest.

333 **Correct Formulation of “Same” Proofs.** We phrase our WHUPs based on the notion of Herbrand
 334 proofs from the famous Herbrand’s Theorem in first order logic. This allows us to interplay with witnessing
 335 theorems nicely. However, it is possible that our notion of “same proof” is still too coarse, and that WHUPs
 336 would be more appropriately phrased in another way. One potential example would be the notion of *uniform*
 337 proofs, proposed by Buss [7], where proofs are given an efficiently decidable direct connection language as
 338 you would a circuit.

339 **Proving a WHUP.** Perhaps the most obvious would be actually showing a WHUP to be true for VPV,
 340 $V_{\#}^0$, or any other bounded arithmetic theory. We believe that past work on Kreisel’s Conjecture [22, 23, 32]
 341 serves as an excellent starting point. For example, Krajíček and Pudlák [32] show that Kreisel’s Conjecture
 342 is true over any theory which is finitely axiomatizable, of which many theories of bounded arithmetic are
 343 (including VPV and $V_{\#}^0$).

344 **Consequences of WHUPs.** Corollary 1.14 shows that WHUPs can have immediate consequences if true.
 345 Are there more examples of WHUP consequences, but for standard theories like VPV and V^1 ? Further, are
 346 there consequences if WHUPs are *false*?

1.8 Paper Organization.

In Section 2, we give the requisite preliminaries in bounded arithmetic and complexity theory. In Section 3, we give our results on Student-Teacher constructive separations for Palindromes, and its applications to provability in VPV. In Section 4, we introduce Witnessing Hypotheses for Uniform Proofs and apply them to get the unprovability of finding high- K^{poly} strings in VPV. We further give a detailed discussion of the viability of WHUPs and their inspiration from the famous Kreisel Conjecture in logic. Finally, in Section 5, we introduce the theory $V_{\#}^0$ and show that a WHUP for this theory would imply $P \neq NP$.

2 Preliminaries

Basic knowledge of complexity classes is assumed. See [3] or any text on complexity theory for a reference of the standard definitions. We attempt to keep this paper as self-contained as possible for mathematical logic and bounded arithmetic; however, we recommend seeing the SIGACT column of Oliveira [43] which surveys much of the recent work on the provability of complexity theory. This survey provides invaluable context to the motivations of this paper.

2.1 Circuit Uniformity

A family of circuits $C = \{C_n\}_{n \geq 1}$ is called uniform if some uniform algorithm is able to, on input n , compute a fixed binary encoding of $\langle C_n \rangle$. We will use the *direct connection* encoding of circuits, where $\langle C_n \rangle_i = 1$ if and only if i encodes a triple (g, h, r) with g and h being gate indices, and r indicating the type of g (namely, one of NOT/AND/OR/INPUT/OUTPUT). In the case that r is an INPUT type, it must also indicate which input bit out of n . Topologically, h feeds into g as an input, unless r indicates that g is an INPUT type. We will also need an *oracle direct connection* encoding. This is a slight modification where we add two types of gates: ORACLE, and Oracle OUTPUT, where ORACLE represents a black-box oracle that takes in $p(n)$ bits and outputs $q(n)$ bits. For both of these gate types, r must also indicate which of the $p(n)$ input bits a gate h is feeding into ORACLE, or which of the $q(n)$ output bits an OUTPUT gate g is. Note that given a circuit with $s(n)$ gates, its (oracle) direct connection encoding will be of length at most $s(n)^3$.

Definition 2.1 (LOGTIME-uniformity). We say that a circuit family $C = \{C_n\}_n$, where C_n is of size $s(n)$, is *logtime uniform* if there is a linear time algorithm U which on input n and an index $i < |\langle C_n \rangle|$, both represented in binary, outputs the i -th bit of $\langle C_n \rangle$. Similarly, such a circuit family is *polylogtime uniform* if the uniform algorithm U runs in time polynomial in the input size.

2.2 Basic Logic and Terminology

We will assume basic knowledge of propositional and first-order logic, as well as Gentzen's sequent calculus. We remind the reader of some of the standard syntax below. For a concise and complete introduction to the necessary logic and proof theory, see Chapters I-III of [17] or Chapters I and II of [6].

Definition 2.2 (Syntax).

Symbols and Terms: The symbols appearing in first-order logic are the usual logical connectives ($\neg, \wedge, \vee, \rightarrow$), quantifiers (\forall, \exists), specified function and predicate symbols, and constants (0-arity functions). As well, arbitrary names for variables are allowed. A *term* is inductively defined: any variable x is a term, and for any function symbol f of arity k and terms t_1, \dots, t_k , $f(t_1, \dots, t_k)$ is a term.

Formulas: A *formula* is also inductively defined. Atomic formulas are of the form $P(t_1, \dots, t_k)$ for a predicate P of arity k , and general formulas are built up from atomic ones by applying logical connectives and quantifiers. We say a variable x in a formula is bound if it is in the scope of a quantifier Qx . Otherwise, it is free. A formula with no free variables is called a *sentence*.

Substitution: Let $A(x)$ be a formula with x a free variable. For a term t , we denote $A(t/x)$ to be the *substitution* of t for x in A , where we replace every occurrence of the free variable x in A with t .

Definition 2.3. A *first-order theory* T is a set of sentences which is closed under logical implication. Specifically, if T derives via a sequent calculus proof the sentence φ , then $\varphi \in T$. A set of sentences Γ are

392 an *axiomatization* of T if $\Gamma \subset T$ and all of T is derivable from Γ via sequent calculus proofs. The *language*
 393 of a theory T , $\mathcal{L}(T)$, is the set of symbols for functions, predicates, and constants (0-ary functions) used in
 394 the logical sentences contained in T . A theory is said to be *universal* if it has an axiomatization with only
 395 universally quantified sentences in prenex normal form.

396 We can compare theories by considering the set of theorems that they prove. The appropriate notion is

397 **Definition 2.4** (Conservative Extension). Suppose that \mathcal{T}_1 and \mathcal{T}_2 are two theories where $\mathcal{T}_1 \subseteq \mathcal{T}_2$ and the
 398 vocabulary of \mathcal{T}_2 may contain function or predicate symbols not in \mathcal{T}_1 . We say \mathcal{T}_2 is a *conservative extension*
 399 of \mathcal{T}_1 if for every formula φ in the vocabulary of \mathcal{T}_1 , if $\mathcal{T}_2 \vdash \varphi$ then $\mathcal{T}_1 \vdash \varphi$.

400 In other words, the second theory proves nothing new over the original vocabulary.

401 In this paper, we study the first-order theory of arithmetic, Peano Arithmetic (PA), as well as its sub-
 402 theories. We denote by \mathbb{N} the standard model of PA, which should be interpreted as the ‘real world’. The
 403 defining feature of Peano Arithmetic (and its intended model \mathbb{N}) is induction: for a formula $\varphi(x, \bar{y})$, the
 404 *axiom of induction*, $I_x\varphi$, is the sentence:

$$\forall \bar{y} (\varphi(0, \bar{y}) \wedge \forall x (\varphi(x, \bar{y}) \rightarrow \varphi(x + 1, \bar{y})) \rightarrow \forall x \varphi(x, \bar{y})).$$

405 Peano Arithmetic is defined by basic arithmetic axioms and the axiom of induction for every formula
 406 φ . For a restricted class of formulas Φ , we define $\text{I}\Phi$ as the subtheory of PA with induction restricted to
 407 formulas $\varphi \in \Phi$.

408 2.3 Peano Arithmetic

409 We recall the characterization of provably recursive functions of $\text{I}\Sigma_n$ [5].

410 **Definition 2.5.** Let T be a subtheory of PA and $f : \mathbb{N}^k \rightarrow \mathbb{N}$. The function f is Σ_i -definable in T iff there
 411 is a formula $\varphi(x_1, \dots, x_k, y) \in \Sigma_i$ such that:

- 412 1. $T \vdash (\forall \vec{x})(\exists! y)\varphi(\vec{x}, y)$
- 413 2. $\{(\vec{a}, b) : \mathbb{N} \models \varphi(\vec{a}, b)\}$ is the graph of f , i.e. $\varphi(\vec{a}, b)$ holds iff $f(\vec{a}) = b$ for all naturals \vec{a}, b .

414 Σ_1 -definable functions in a theory \mathcal{T} are also commonly called the *provably recursive* functions of \mathcal{T} .

415 **Lemma 2.6** (Informal). Let f be a function that is provably recursive in PA. Then we can freely add the
 416 function symbol f to $\mathcal{L}(\text{PA})$ and the defining axioms of f to PA without modifying the strength of PA.

417 **Definition 2.7.** Let $n \geq 1$. The set of functions which are *primitive recursive in Σ_n* is defined inductively
 418 by:

- 419 1. Constant function 0, successor function, and all projection functions are primitive recursive in Σ_n .
- 420 2. Closure under composition.
- 421 3. If $g : \mathbb{N}^k \rightarrow \mathbb{N}$ and $h : \mathbb{N}^{k+2} \rightarrow \mathbb{N}$ are primitive recursive in Σ_n , then so is the function f defined by

$$\begin{aligned} f(0, \vec{z}) &= g(\vec{z}) \\ f(m + 1, \vec{z}) &= h(m, \vec{z}, f(m, \vec{z})) \end{aligned}$$

- 422 4. If $\varphi(\vec{z})$ is a formula $(\exists x)\psi(x, \vec{z})$ where $B \in \Pi_{n-1}$ then U_A is primitive recursive in Σ_n .

423 **Theorem 2.8** (Theorem 12, [5]). The Σ_n -definable functions of $\text{I}\Sigma_n$ are the functions which are primitive
 424 recursive in Σ_n .

2.4 Theories of Bounded Arithmetic

We will be working with two-sorted theories, which deal with both a number-type (think in \mathbb{N}) and a finite binary string type. The binary string type has an equivalent interpretation as a set type, where the i -th index of a string X being 1 indicates that i is in the set X . We follow the convention of denoting numbers in lower case (x, y, z, \dots) and strings in upper case (X, Y, Z, \dots). All theories in this paper are theories of arithmetic, and all share the language of arithmetic (\mathcal{L}_A^2), which contains the set of first-sort functions and predicates, $\{0, 1, +, \cdot, S, |\cdot|; =, \leq\}$ and the set of second-sort functions and predicates, $\{X(\cdot), |\cdot|; =_2\}$. Here, S refers to the successor function of a number, $X(i)$ outputs in number type the i -th bit of string X , and $|\cdot|$ on a string-type variable outputs a number-type which is the length of the string.

In two-sorted bounded arithmetic theories, function symbols can be thought of as the run of some resource-bounded computational model (eg. Turing machines or uniform circuits). As such, the representation of its inputs becomes important. We will take the standard convention that the string-type is presented “as itself” in binary and a number-type is represented in unary when supplied as input to a function symbol. A feature of Peano Arithmetic and its subtheories is that any function f which is “easily” definable and provably total may be freely added to the language without changing the provability of any sentences. Below, we will specify exactly what these functions are for each theory we use.

Definition 2.9. We denote a number quantifier as *bounded* by writing $\forall x < t\theta(x)$ or $\exists x < t\theta(x)$, for a term t not using x . This is syntactic shorthand for $\forall x [x < t \implies \theta(x)]$ and $\exists x [x < t \wedge \theta(x)]$ respectively. Similarly for quantifiers over strings, we say write $\forall X < t\theta(X)$, and $\exists X \leq t\theta(X)$ to indicate $\forall X (|X| < t \implies \theta(X))$ and $\exists X (|X| < t \wedge \theta(X))$. We say that a formula φ is $\Sigma_0^B = \Pi_0^B$ if the only quantifiers are bounded quantifiers over the number type (though there may be free string variables). A formula φ is $\Sigma_{i+1}^B/\Pi_{i+1}^B$, for $i \geq 0$, if φ is of the form, $\exists X < t\theta(X)$, for $\theta(X)$ a Π_i^B formula, or respectively, $\forall X < t\theta(X)$, for $\theta(X)$ a Σ_i^B formula.

Σ_i^B formulas can be thought of as an effective version of the arithmetic hierarchy, and bears many similarities and connections to the polynomial hierarchy.

Definition 2.10 (Provably Total Functions). Let T be a two-sorted subtheory of PA and $f : \mathbb{N}^k \rightarrow \mathbb{N}$. The function f is Σ_i^B -definable in T iff there is a Σ_i^B -formula $\varphi(x_1, \dots, x_k, y)$ such that:

1. $T \vdash (\forall \vec{x})(\exists! y)\varphi(\vec{x}, y)$
2. $\{(\vec{a}, b) : \mathbb{N} \models \varphi(\vec{a}, b)\}$ is the graph of f , i.e. $\varphi(\vec{a}, b)$ holds iff $f(\vec{a}) = b$ for all naturals \vec{a}, b .

Σ_1^B -definable functions in a theory \mathcal{T} are also commonly called the *provably total* functions of \mathcal{T} .

We may give a lemma similar to Lemma 2.6 for provably total functions.

Lemma 2.11 ((Informal)). Let f be a function that is provably total in a two-sorted theory T . Then we can freely add the function symbol f to $\mathcal{L}(T)$ and the defining axioms of f to T without modifying the strength of T .

Theory V^0 . One of the weakest and most basic of theories in bounded arithmetic that is studied is Cook and Nguyen’s theory V^0 , which captures uniform- AC^0 reasoning. It is a uniform version of the propositional proof system AC^0 -Frege, and superpolynomial lower bounds for AC^0 -Frege imply unprovability in V^0 .

At the base of V^0 are the so-called 2-BASIC axioms, which define the basics of how each function and predicate in \mathcal{L}_A^2 behaves. This includes statements like $x \cdot 0 = 0$, distributivity of addition over multiplication, and many others. See [17] for the full list of axioms. In addition to 2-BASIC are the comprehension axioms Σ_0^B -COMP, where for any Σ_0^B -formula φ , we get the axiom,

$$\exists X \leq y \forall z < y X(z) \leftrightarrow \varphi(z).$$

Σ_0^B -COMP axioms should be thought of as giving V^0 the power to generate truth tables of AC^0 -computable functions. V^0 will, in addition to \mathcal{L}_A^2 , have a function symbol f in its language for every LOGTIME-uniform AC^0 function f , and the Σ_1^B -defining axiom of f added to V^0 . Note it is well-known that LOGTIME-uniform AC^0 is equivalent to the LOGTIME Hierarchy, so we may include functions symbols for either.

V^0 is surprisingly powerful and expressive. It is capable of proving many elementary theorems about number theory and combinatorics and can perform Gödel numbering and coding of sequences. It is known that V^0 cannot reason about the Parity function (\oplus) or other functions which have AC^0 lower bounds.

472 **Theory VPV.** The full definition of VPV is involved, and the details do not matter here outside of its
473 correspondence with polynomial time functions. To see a detailed definition of VPV, see [17]. The language
474 of VPV is \mathcal{L}_A^2 along with a symbol f for any polynomial-time computable function f . The theory is defined
475 by initially adding the defining axioms of five uniform-AC⁰ functions, and then using Cobham’s recursive
476 definition of polynomial time functions [15] within the theory to build out the rest of FP.

477 **Theory V¹.** Adding the comprehension axioms Σ_1^B -COMP to 2-BASIC, we go from V⁰ to V¹. As every
478 polynomial time function is Σ_1^B -definable in V¹, we may freely add their defining axioms to the theory
479 and add a function symbol for every $f \in \text{FP}$. This theory characterizes polynomial time computation and
480 reasoning, similarly to VPV. It has the benefit of being much easier to define, and is more easily generalizable
481 to reflect reasoning in the i -th level of the polynomial hierarchy (theory V ^{i}). It is known that VPV \subseteq V¹, but
482 it is open if VPV and V¹ are in fact equal; under cryptographic assumptions like the hardness of factoring,
483 Thapen showed that V¹ is strictly stronger [50]. As we shall also see, there is an important difference in the
484 witnessing theorems for VPV compared to the witnessing theorems for V¹.

485 V¹ (and more generally V ^{i} , for $i > 0$), are equivalent to the single-sorted theories S₂ ^{i} introduced by Buss
486 in his seminal PhD Thesis [4].

487 **VPV Function Symbols** We will be translating several lower bounds against Turing machines of different
488 resource bounds. In order to give VPV-translations of these statements, we must introduce some preliminary
489 function symbols.

490 Let $\text{Run}_M(X, n)$ be the VPV function symbol that on input X and clock bound n , runs M for n steps
491 on input X and outputs the contents of its tape. Similarly for a nondeterministic machine M , an input
492 X , clock n , and witness W supplied on a separate read-only witness tape, we have a VPV function symbol
493 $\text{Run}_M(X, n, W)$ which run M for n steps on input X and nondeterminism W and outputs the contents of
494 its tape input/work tape.

495 **Lemma 2.12** (Implicit in [18, 4]). There is a paddable encoding of one-tape deterministic Turing machines
496 $\mathcal{L}_{TM} \subset \{0, 1\}^*$ which is *decodable* in VPV. Specifically, there is a VPV function symbol $\text{Run}(M, X, n)$
497 where for every Turing machine M and its binary encoding $E_M \in \mathcal{L}_{TM}$, $\text{VPV} \vdash \forall X \forall n \text{Run}_M(X, n) =$
498 $\text{Run}(E_M, X, n)$.

499 Similarly for one-tape *nondeterministic* Turing machines, we can give an encoding language $\mathcal{L}_{NTM} \subset$
500 $\{0, 1\}^*$ which is decodable. Specifically, there is a VPV function symbol $\text{Run}(M, X, n, W)$ where for every non-
501 deterministic Turing machine M and its binary encoding $E_M \in \mathcal{L}_{NTM}$, $\text{VPV} \vdash \forall X \forall W \forall n \text{Run}_M(X, n, W) =$
502 $\text{Run}(E_M, X, n, W)$.

503 The above lemma can be reformulated for k -tape Turing machines for any number k , but we will only
504 be concerned with one-tape machines in this paper. We will always assume Turing machines are encoded as
505 \mathcal{L}_{TM} from Lemma 2.12.

506 2.5 Witnessing Theorems in Bounded Arithmetic

507 Witnessing theorems broadly show that if a theory \mathcal{T} proves a $\forall \Sigma_i^B$ formula φ , then there is a function f_φ
508 computable in a complexity class $\mathcal{C}_\mathcal{T}$ which finds a witness to the existential quantifiers in φ . We will largely
509 work only with $\forall \Sigma_1^B$ and $\forall \Sigma_2^B$ formulas, which make witnessing conceptually simpler due to there being a
510 single existential quantifier.

511 The most classical example of witnessing in Bounded Arithmetic is Buss Witnessing [4], which is written
512 in the language of two-sorted theories in [17].

513 **Theorem 2.13** (Buss Witnessing, [4, 17]). Let \mathcal{T} be either V¹ or VPV, and let φ be a Σ_1^B formula. Suppose
514 that

$$\mathcal{T} \vdash \forall X \exists Y \varphi(X, Y).$$

515 Then there exists a function $F \in \text{FP}$ such that $\mathbb{N} \models \forall X \varphi(X, F(X))$.

516 Krajíček, Pudlák, and Takeuti generalized Buss Witnessing to $\forall \Sigma_2^B$ formulas as follows.

517 **Theorem 2.14** (KPT Witnessing Theorem, [33]). Let T be a universal theory with language \mathcal{L} . Suppose
518 that for a Σ_0^B formula φ ,

$$T \vdash \forall X \exists Y \forall Z \varphi(X, Y, Z).$$

519 Then for a constant $k \geq 1$ and a sequence C_1, \dots, C_k of \mathcal{L} -string terms,

$$T \vdash \forall X \forall \bar{Z} [\varphi(X, C_1(X), Z_1) \vee \varphi(X, C_2(X, Z_1), Z_2) \vee \dots \vee \varphi(X, C_k(X, Z_1, \dots, Z_{k-1}), Z_k)].$$

520 This theorem applies to VPV, as VPV is a universal theory. For V^0 and V^1 , KPT Witnessing as above
521 cannot be immediately applied as neither theory is universal. There are several ways around this. One way
522 is to universalize the axioms of V^0 and V^1 to give conservative extensions \bar{V}^0 and \bar{V}^1 , where KPT Witnessing
523 can be applied. The other way is to prove KPT Witnessing directly using proof theoretic arguments and
524 Buss Witnessing. For V^0 , we will use the former method and apply the above KPT Witnessing Theorem to
525 the universal \bar{V}^0 . For V^1 , we present its own KPT Witnessing Theorem below.

526 **Theorem 2.15** (KPT Witnessing Theorem for V^1 , [28]). Suppose that for a Σ_0^B formula φ ,

$$V^1 \vdash \forall X \exists Y \forall Z. (|Z| < |X|) \varphi(X, Y, Z).$$

527 Then there is an FP function F such that,

$$\mathbb{N} \models \forall X \forall Z. (|Z| < |X|) \varphi(X, F(X), Z),$$

528 where F has access to the counterexample oracle $CX[\varphi]$ which on query (X, Y) outputs a string Z of length
529 at most $|X|$ such that $\mathbb{N} \models \neg\varphi(X, Y, Z)$ or “yes” otherwise.

530 The Student-Teacher game interpretation of KPT Witnessing is very useful. A Student F , which is
531 a search algorithm of some complexity class \mathcal{C} , will take in X as input and want to find a Y such that
532 $\forall Z \varphi(X, Y, Z)$. They start by proposing $F_1(X)$ to the Teacher, the counterexample oracle, who either says
533 $F_1(X)$ is correct or gives a Z_1 back to the Student as a counterexample. This repeats for r rounds until the
534 Student proposes a correct Y .

535 A difference between VPV and V^1 is revealed here: the Student-Teacher game from the KPT Witnessing
536 for VPV ends in constantly many rounds, while the Student-Teacher game for V^1 ends in polynomially many
537 rounds. This makes unprovability of $\forall \Sigma_2^B$ formulas in V^1 potentially much harder than in VPV. Unprovability
538 of $\forall \Sigma_2^B$ formulas usually goes by applying KPT Witnessing and showing the resulting Student-Teacher game
539 can collapse into an impossibly fast/small algorithm *without* the counterexample oracle. The more rounds
540 of a Student-Teacher game, the harder it is to prove that the oracle may be removed.

541 2.6 Student Teacher Games and Refuters

542 We formally introduce the Student-Teacher game framework which witnesses the KPT Witnessing Theorem.

543 **Definition 2.16** (\mathcal{C} -ST $^{CX[\varphi, r]}$ uniformity). Let \mathcal{C} be a complexity class, and for a term t , let

$$\psi := \forall n \exists Y (|Y| < t(n)) \forall Z (|Z| = n) \varphi(n, Y, Z)$$

544 be a formula with $\varphi \in \Sigma_0^B$ and $\mathbb{N} \models \psi$. As well, let $r(n)$ be a time-constructible function. Define Search_φ to
545 be the total search problem $\text{Search}_\varphi := \{(n, Y) \mid Y \text{ a binary string such that } \mathbb{N} \models \forall Z (|Z| = n) \varphi(n, Y, Z)\}$.

546 We say that \mathcal{A} is a \mathcal{C} -ST $^{CX[\varphi, r]}$ search algorithm for Search_φ if $\mathcal{A} \in \mathcal{C}$ and on input 1^n , \mathcal{A} outputs a
547 Y satisfying $\forall Z (|Z| = n) \varphi(n, Y, Z)$ using at most $r(n)$ many oracle queries to the counterexample oracle
548 $CX[\varphi]$.

549 Many complexity lower bounds are easily formalizable as either $\forall \Sigma_1^B$ or $\forall \Sigma_2^B$ formulas in $\mathcal{L}(\text{VPV})$ and
550 $\mathcal{L}(V^0)$, where the existential quantifier witnesses a mistake that some Turing machine or algorithm has made
551 when attempting to decide a hard language. Applying witnessing theorems to these lower bounds when they
552 are provable in bounded arithmetic gives us *refuters*.

553 Suppose, say, VPV were to prove a complexity lower bound formalizable as $\forall \Sigma_2^B$ formula ψ . Applying
554 KPT Witnessing, we would then get an P-ST $^{CX[\varphi, r]}$ constructive separation. For a $\forall \Sigma_1^B$ formalizable lower
555 bound, Buss Witnessing then directly gives an P-refuter and a P-constructive separation.

2.7 Time-Bounded Kolmogorov Complexity

There are many ways to define time-bounded Kolmogorov complexity [2, 34]. Some choices made in these definitions are essentially arbitrary, like which efficient universal Turing Machine to use. We will specify these choices carefully enough to give a particular translation of time-bounded Kolmogorov complexity into theories of (bounded) arithmetic, but our results will not depend on the precise formalization. We follow Section 2.2 of [36], elaborating on some details.

Fix a *string pair encoding function* $\langle \cdot, \cdot \rangle : \{0, 1\}^+ \times \{0, 1\}^+ \rightarrow \{0, 1\}^+$ defined by the map $\langle u, v \rangle \mapsto \text{dbl}(u) \circ 01 \circ v$, where $\text{dbl}(u) = u_1 u_1 \circ u_2 u_2 \circ \dots \circ u_{|u|} u_{|u|}$ simply double-prints each bit of u . Denote by π_1 and π_2 the left and right extraction functions, so $\pi_1(\langle u, v \rangle) = u$ and $\pi_2(\langle u, v \rangle) = v$. These pair encoding and element extraction function are linear-time computable and well-defined for all non-empty binary strings. Furthermore, delimiter overhead is only incurred for the length of the first string, plus an additive constant: $\forall u, v \ |\langle u, v \rangle| = 2|u| + 2 + |v|$.

Fix U a Universal Turing machine that can emulate any single-tape Turing Machine M with at most polynomial-time overhead. Let $\text{run}_U(M, x, 1^t)$ denote the function that outputs the entire non-blank contents of the tape of M simulated on input x for t steps of U . By the assumption that U is efficient, run_U can be computed in time $\text{poly}(|M|, |x|, t)$.

Finally, the *t-time bounded Kolmogorov Complexity* $K^t(x)$ of a string x is the length of the shortest two-part description d of x such that U decodes d into x :

$$K^t(x) = \min_{d \in \{0, 1\}^*} \{|d| : U(\pi_1(d), \pi_2(d), 1^{t(|x|)}) = x\}$$

The K^t complexity of any string x is at most $|x|$, because the two-part description can simply “memorize” x . Consider the description $d = \langle H, x \rangle$ where H is the constant-length description of a Turing Machine that immediately halts. Because run outputs the contents of the tape of H , this is simply x . Thus we have the following

Fact 2.17. There is an absolute constant c such that for every function $t(n) > 0$ and every $x \in \{0, 1\}^+$ it holds that $K^t(x) \leq |x| + c$.

Observe that it is important to pay delimiter overhead for the constant-length machine H instead of the variable-length string x to obtain the basic fact above. This is implicit in every reasonable definition of time-bounded Kolmogorov complexity.

3 Provability of Palindromes Lower Bounds

In this section, we generalize the work of Chen et. al. [11] and show that provability of the palindromes lower bound in VPV implies circuit lower bounds.

To do this, we formalize Maass’s lower bound as a $\forall \Sigma_2^B$ $\mathcal{L}(\text{VPV})$ -sentence and, assuming $\text{VPV} \vdash$ “Maass”, apply the KPT Witnessing theorem. We then assume a complexity upper bound that *both* collapses the Student-Teacher refuter into a P-refuter and causes a contradiction via the argument of [11].

In Section 3.1, we give a formalization of palindrome lower bounds and discuss its witnessed Student-Teacher refuter under VPV-provability. We then give a slightly generalized version of the constructive separations argument of [11]. Finally, in Section 3.3, we identify a complexity assumption that both collapses the Student-Teacher refuter and allows a standard constructive separations argument to go through.

3.1 Formalization of One-Tape Nondeterministic Turing Machine Lower Bounds

First, we state Maass’s theorem in plain English.

Theorem 3.1 ([38]). The language $\text{PAL} := \{p \in \{0, 1\}^* \mid p \text{ a palindrome}\}$ is not computable by any one-tape nondeterministic Turing machine in $n^{1.1}$ steps.

To formalize Theorem 3.1, we will need to introduce several functions, all of which are clearly VPV function symbols. The symbol $\text{ValNTM}(\cdot)$ takes in a string M and outputs 1 if and only if M is a valid

599 encoding $(M \in \mathcal{L}_{NTM})$ of a one-tape nondeterministic Turing machine. We define $\text{IsPal}(X)$ to output 1 if
600 the string X is a palindrome, and 0 otherwise. Recall that $\text{Run}(M, X, t, W)$ outputs 1 if nondeterministic
601 machine M on input X with guess bits W ACCEPTS within t steps. Finally,

$$\text{Err}_{PAL}^i(M, X, t, W) \triangleq (\text{IsPal}(X) = i) \wedge (\text{Run}(M, X, t, W) = 1 - i).$$

602 Let $\text{Pal}(n_0)$ denote the following sentence.

$$\begin{aligned} \text{Pal}(n_0) \triangleq \forall n (n > n_0) \forall M (|M| \leq n/2) \exists X (|X| = n) \exists W_X (|W_X| \leq n^{1.1}) \forall W (|W| \leq n^{1.1}) \\ \text{ValNTM}(M) \wedge (\text{Err}_{PAL}^1(M, X, n^{1.1}, W) \vee \text{Err}_{PAL}^0(M, X, n^{1.1}, W_X)) \end{aligned}$$

603 The formalization covers two cases: either the machine M claims an input X is a palindrome when it is
604 not (captured by Err^0), or it claims X is not a palindrome when it in fact is (captured by Err^1).

605 **The Student-Teacher Refuter.** Assuming $\text{VPV} \vdash \text{Pal}(n_0)$, for some n_0 , we have the following Student-
606 Teacher game interpretation via the KPT Witnessing theorem.

607 Let φ be the innermost Σ_0^B formula of $\text{Pal}(n_0)$, and r be the fixed constant many rounds of the Student-
608 Teacher game. A P-Student will take as input 1^n and a machine M . In round one, they will query the
609 Teacher on a string X and witness W_X where it thinks M incorrectly decides X is or isn't a palindrome. The
610 Teacher will respond with a witness W that either shows the machine M correctly accepts the palindrome X
611 on $M(X, W)$, or that the proposed witness W_X actually rejects a non-palindrome X .³ This is an $\text{P-ST}^{CX[\varphi, r]}$
612 constructive separation for Maass's lower bound.

613 3.2 Constructive Separations for Palindromes

614 In order to collapse Student-Teacher games, we will need small amounts of nonuniformity to replace the
615 Teacher's responses. This generalizes the argument of [11] that P-constructive proofs of Maass's lower
616 bound imply breakthrough circuit lower bounds. Here, we will need $\text{P}/o(n^\varepsilon)$ -constructivity.

617 **Lemma 3.2** (Lemma 3.3, [11]). There exists a one-tape nondeterministic Turing Machine M running in
618 subquadratic time that acts correctly on all inputs x with circuit complexity $|x|^\delta$, for a fixed $0 < \delta < 1$.

619 *Proof sketch.* First, on input x , M will guess a $\log n$ -input circuit C_x of size n^δ and evaluate it on all n
620 possible inputs to verify that C_x succinctly represents x . Next, for each $0 \leq i \leq n/2$, M will evaluate C_x on
621 i and $n - i$ and ensure that $C(i) = C(n - i)$. In total, M will run in time $n \cdot n^{O(\delta)} = o(n^2)$ for a sufficiently
622 small constant δ . \square

623 **Lemma 3.3** (Generalization of Lemma 2.3, [11]). Assume that $\text{P} \subset \text{SIZE}[n^k]$ for some $k \geq 1$. Let $\varepsilon > 0$. Then
624 for any $\text{P}/o(n^\varepsilon)$ -algorithm $R(1^n)$ with n output bits, we have that the string $R(1^n)$ has circuit complexity
625 $o(n^\varepsilon)$.

626 *Proof.* Assume that $\text{P} \subset \text{SIZE}[n^k]$ for some $k \geq 1$, and let R be a P-algorithm with advice α of length
627 $|\alpha| = o(n^\varepsilon)$ which takes in a unary input 1^n and outputs an n -bit string. For any $\varepsilon' > 0$, we can construct a
628 new $\text{P}/(|\alpha| + O(\log n))$ -algorithm R' where R' takes as input $1^{n^{\varepsilon'}}$ and $i \in [n]$ in binary, is given n in binary as
629 advice, and outputs the i -th bit of $R(1^n)$. This is clearly still a polynomial time algorithm, and by the above
630 assumption, has a circuit of size $O(n^{\varepsilon'/k} + o(n^\varepsilon))$. Set $\varepsilon' = \varepsilon/2k$ to achieve the desired circuit complexity. \square

631 The following is a straightforward generalization of the second item of Theorem 3.4 in [11].

632 **Theorem 3.4.** Let $0 < \varepsilon < 1$. A P/n^ε -constructive proof of Maass' bound implies that $\text{P} \not\subset \text{SIZE}[n^k]$.

633 *Proof.* Suppose that $\text{P} \subset \text{SIZE}[n^k]$. Then by combining the above two lemmas, there is a one-tape NTM
634 M running in subquadratic time that is correct on all strings which could be output by refuters. This
635 contradicts Maass' lower bound being P/n^ε -constructive. \square

³Note that in the second case, no response from the Teacher is actually needed as a polynomial time Student can check this condition for themselves.

636 3.3 Round Elimination for the Student-Teacher Refuter

637 Similar to the round elimination of [8], we show that every query to the counterexample oracle CX can be
 638 replaced by a sublinear advice string. There are two new ideas compared to previous work.

639 (i) Recognize that Teacher in the Student-Teacher refuter is just an NP predicate.

640 (ii) By assuming (towards a contradiction) that $\text{NP} \subset \text{SIZE}[n^k]$, we can use the Easy Witness Lemma for
 641 NP to “compress away” Teacher into sub-linear advice, round-by-round.

642 **Theorem 3.5** (Easy Witness Lemma, [42]). Let $k > 0$. Suppose that $\text{NP} \subset \text{SIZE}[n^k]$. Then there is a
 643 constant $d > 0$ where for any $\mathcal{L} \in \text{NP}$ and Yes-input X , there is a witness W succinctly represented by a
 644 circuit of size n^{dk^3} .

645 **Theorem 3.6.** Let r, k be positive integers. Assume that $\text{NP} \subseteq \text{SIZE}[n^k]$. Then an $\text{P-ST}^{CX[\varphi, r]}/a(n)$ refuter
 646 for Maass implies an $\text{P-ST}^{CX[\varphi, r-1]}/a'(n)$ refuter for $a(n) = O(n^\delta)$ with $\delta < 1$ and $a'(n) = C \cdot a(n)^{O(k^3)}$,
 647 with $C > 0$ a constant.

648 *Proof.* Let M be a nondeterministic Turing machine clocked to run in time $n^{1.1}$, and let d be the constant
 649 appearing in Theorem 3.5. First, we note that without loss of generality, the Student will only propose a
 650 palindrome to the counterexample oracle. This is because if the Student proposes a non-palindrome, then
 651 the oracle response can be compressed to 0 bits and completely removed; the Student can check for itself in
 652 linear time⁴ that its proposed string X is not a palindrome, and in $n^{1.1}$ time to simulate M on X and the
 653 proposed witness W_X .

654 Let $p \in \{0, 1\}^n$ be the first palindrome that the student queries the teacher. As no teacher queries are
 655 made yet, p is computable in $\text{P}/a(n)$. Consider the following NP-language \mathcal{L}_{wit} :

$$\mathcal{L}_{wit}^n := \{x : x \in \{0, 1\}^{n^{1.1}} \text{ and } M(p, x) = 1\}.$$

656 Note that \mathcal{L}_{wit}^n is the set of witnesses to the nondeterministic machine W that takes in 1^n as input and
 657 a string of length $a(n)$ as advice and decides if p is a 1-input to M . Further, we can pad down the input
 658 to 1^{n^ε} , for any constant $\varepsilon > 0$, and add n in binary as advice. Pick $\varepsilon < \delta$. Hence by Theorem 3.5, there
 659 exists an $x \in \mathcal{L}_{wit}^n$ that has circuit complexity $(n^\varepsilon + \log n + a(n))^{dk^3} \leq (2a(n))^{dk^3}$. We replace the teacher
 660 by instead giving the student this witness circuit at the beginning of the Student-Teacher game. As a result,
 661 we change the protocol to have $r - 1$ rounds of interaction and $a(n) + (a(n) + n^\delta + \log n)^{dk^3} < (4a(n))^{dk^3}$
 662 bits of advice. \square

663 We then have the following corollaries.

664 **Theorem 3.7** (Theorem 1.6). If for any nondeterministic one-tape subquadratic time Turing machine M
 665 there is a P-Student-Teacher game $S_M(1^n)$ with counterexample oracle $CX[\varphi_{\text{Pal}}, O(1)]$ solving $\text{Ref}_{\text{Pal}, M}$ for
 666 n -bit inputs, then $\text{NP} \not\subset \text{SIZE}[n^k]$ for any $k \geq 0$.

667 *Proof.* Suppose there is an $\text{P-ST}^{CX[\varphi, r]}$ refuter of constantly many rounds r for Palindromes. Apply Theorem
 668 3.6 to remove the first teacher query, adding n^ε bits of advice, for any $\varepsilon > 0$ we desire. Pick $\varepsilon < 1/(100dk^3)^{2r}$.
 669 Repeatedly apply Theorem 3.6 another $r - 1$ times to have a $\text{P}/o(n^{1/100})$ refuter, we contradict Theorem 3.4.
 670 \square

671 **Theorem 3.8** (Theorem 1.11). If $\text{VPV} \vdash \text{Pal}(n_0)$, for any $n_0 > 0$, then $\text{NP} \not\subset \text{SIZE}[n^k]$.

672 *Proof.* Suppose $\text{VPV} \vdash \text{Pal}(n_0)$ and that $\text{NP} \subset \text{SIZE}[n^k]$ for some $k > 0$. Then by the KPT-witnessing
 673 theorem, we get an $\text{P-ST}^{CX[\varphi, r]}$ refuter of constantly many rounds r . Applying Theorem 3.7, we have a
 674 contradiction. \square

⁴Student need not be a one-tape TM, so checking PALINDROME can indeed be linear time.

675 4 Existence of K^t -Random Strings

676 Hirahara’s lower bound $R_{K^t} \notin P$ for $t = \text{qpoly}$ is **unconditionally** non-constructive [21, 11]. Could we
 677 extract a related unprovability result for VPV? Non-constructivity was established by using assumed P-
 678 refuters to print high- K^t strings for $t = \text{qpoly}$ in only **poly**-time — a contradiction [11]. This suggests to
 679 begin studying VPV-provability of the lower bound “ $R_{K^t} \notin P$ ” by considering first the simpler statement
 680 “there exist K^t -random strings,” abbreviated informally as $\exists R_{K^t}$ below.

681 Even $\exists R_{K^t}$ requires some care to express in VPV. Straightforward (i.e., without padding) translation
 682 of $\exists R_{K^t}$ into VPV with $t = \text{qpoly}$ is impossible, because VPV-number-terms must have fixed polynomial
 683 growth. So we study instead provability of a sequence of statements asserting that high- K^{n^c} strings exist:
 684 “for sufficiently large n , there is an n -bit string X with $K^t(X) > n/2$ ” where $t = n^c$ for each c .

685 *Formalization 4.1* (HiK^t for VPV). Fixing n_0 , define the following sequence of VPV sentences for each $c \in \mathbb{N}$.

$$\text{HiK}^t[c] := \forall n.(n > n_0) \exists X.(|X| = n) \forall D.(|D| < n/2) \text{run}(\pi_1(D), \pi_2(D), n^c) \neq X$$

686 *Remark 4.2.* The symbol c is *not* a free variable in a VPV-formula called HiK^t . It is rather the parameter of
 687 a sequence of formulas where “ n^c ” abbreviates the constant-length term $\underbrace{n \cdot n \cdot n \cdots n}_{c \text{ occurrences of } n}$.

688 Fixing sufficiently large n_0 , each statement $\text{HiK}^t[c]$ is true in the standard model by simple counting.⁵
 689 Furthermore, the argument is essentially identical for each c , differing only by a substitution of numeric
 690 terms. Can VPV carry it out? Can VPV prove HiK^t via a “uniform” argument, such that the proofs for
 691 $\text{HiK}^t[c]$ and $\text{HiK}^t[c']$ with $c \neq c'$ have a clean quantitative relationship as syntactic objects?

692 We make some progress towards answering these questions about provability of the schema $\text{HiK}^t[c]$ by
 693 giving lower bounds on Student-Teacher search for K^t -random strings for each fixed $t \in \text{poly}$ (Section 4.1)
 694 and proof-theoretic hypotheses under which these lower bounds imply unprovability (Section 4.3).

695 4.1 Student-Teacher-Search Lower Bounds for K^t -Random Strings

696 First we derive a sequence of search problems from the schema HiK^t as described in Section 2.6. Extract the
 697 quantifier-free part of $\text{HiK}^t[c]$ for each c as:

$$\psi_c(n, X, D) := (|D| \leq n/2 \wedge n > n_0) \rightarrow (\text{run}(\pi_1(D), \pi_2(D), n^c) \neq X \wedge |X| = n)$$

698 Because $\text{HiK}^t[c]$ is true in the standard model for every c , the problem Search_{ψ_c} is total and well-defined for
 699 every c . To ease notation, we spell out and abbreviate these search problems below.

700 **Definition 4.3** (Search for K^t -Random Strings). For each $c \in \mathbb{N}$, abbreviate the problem Search_{ψ_c} by

$$\exists \text{HiK}^t[c] := \{(1^n, X) \mid K^{n^c}(X) > n/2 \wedge |X| = n\}$$

701 An answer to the counterexample query X for $\exists \text{HiK}^t[c]$ is binary string D that is

- 702 1. *short*, so $|D| < n/2$ and
- 703 2. *describes* X , so $D = \langle M, A \rangle$ with M run on input A for at most n^c steps halts with X on the tape.

704 Any such D is a valid counterexample to the claim “ $K^{n^c}(X) \geq n/2$.” Having fixed terminology, we are ready
 705 to state and prove our lower bounds against Student-Teacher search for K^t -random strings.

706 The base case — Students that make no queries — is implicit in Proposition 1.8 of [11]. Generalizing the
 707 “indexing template” embedded in that proof yields our construction. Their argument is paraphrased below.

708 **Proposition 4.4.** For $c \geq 1$, no student running in time $\tilde{O}(n^c)$ and making zero queries solves $\exists \text{HiK}^t[c+1]$.

⁵The constant n_0 need only be large enough to ensure that $\text{run}(\pi_1(D), \pi_2(D), n_0^c)$ is well-defined for $|D| \geq n_0/2$. Thus n_0 can be fixed to an absolute constant depending only on the machine and pair encoding implicit in the run and π functions.

709 *Proof.* Suppose S is a student that runs in time $\tilde{O}(n^c)$ and solves $\exists\text{HiK}^t[c+1]$ without making any queries.
710 Denote by ℓ the description length of S , fix arbitrary $n \in \mathbb{N}$, and let $x_n = S(1^n)$ be the n -bit $K^{n^{c+1}}$ -random
711 string found by S . Define the *indexing* of S to be the standard, one-tape Turing Machine $\text{ix}(S)$ that results
712 from substituting S into the Indexing Template (Algorithm 1). Because S makes no queries, it can indeed
713 be simulated by a standard one-tape Turing Machine.

714 By construction, $\text{ix}(S)$, given input n encoded in binary, prints x_n . This takes $\tilde{O}(n^c)$ steps for a larger
715 **polylog** factor than in the original runtime of S , accounting for time to print 1^n onto the worktape and
716 to run $S(1^n)$. The description length of $\text{ix}(S)$ is just $\ell + a$ for an absolute constant a depending on the
717 universal machine and book-keeping code to expand the binary representation of n into 1^n . Therefore, the
718 pair $\langle \text{ix}(S), \text{bin}(n) \rangle$ witnesses $\mathsf{K}^{n^{c+1}}(x_n) \leq 2(\ell + a) + \log n + 2$ — a contradiction for sufficiently large n . \square

Algorithm 1 Indexing Template $\text{ix}(S)$

Parameters S the description of a Turing machine

- 1: On input $\text{bin}(n)$
 - 2: **output** $S(1^n)$
-

724 Observe that $a_n = \langle \text{ix}(S), \text{bin}(n) \rangle$ is a uniform counterexample to the claim “ x_n is a K^t random string”
725 for any zero-query student and sufficiently large n . This suggests that even if a student S for $\exists\text{HiK}^t$ does
726 make queries, the description of S could be used to answer and eliminate them. Two-parameter indexing —
727 tracking both n and number of queries made by $S(1^n)$ — suffices to realize this intuition (Algorithm 2).
728

729 **Theorem 4.5.** For $c \geq 1$, no student running in time n^c solves $\exists\text{HiK}^t[c+1]$.

730 *Proof.* Suppose S is a student of description length ℓ running in time n^c that solves $\exists\text{HiK}^t[c+1]$ using at
731 most $r(n) < n^c$ queries. By Proposition 4.4, it is immediate that $r(n) \geq 1$. We will eliminate these queries
732 by constructing a uniform sequence of valid answers — derived from S itself — that are easy to produce
733 without a teacher. Before arguing for validity, we show that such a “reflection exchange” of answers and
734 queries is well-defined and establish some basic properties (Claim 4.6).

735 More precisely, to generate counter-examples for S from the description of S , we must convert S into a
736 standard, one-tape Turing machine (TM) — because $\exists\text{HiK}^t$ is defined with respect to this *particular* model
737 of computation. The *Reflection Template* transforms any student S into a standard Turing machine $\text{rf}(S)$
738 by substituting the description of S into Algorithm 2 below. For each standard one-tape Turing machine M ,
739 write $\ulcorner M \urcorner$ for the binary encoding of M induced by the *particular* universal machine used to define $\exists\text{HiK}^t$.
740 We can now state

741 **Claim 4.6.** There is a standard one-tape Turing machine $\text{rf}(S)$ such that, fixing the sequence of answers
742 $a_{n,j} = \langle \ulcorner \text{rf}(S) \urcorner, \langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), \text{bin}(j) \rangle \rangle \rangle$ and denoting by $q_{n,i}$ the induced sequence of queries
743 $q_{n,i} =$ “the i -th query made by $S(1^n)$ after getting $a_{n,j}$ in response to the j -th query for $j \in \{1, \dots, (i-1)\}$,”
744 the following properties hold:

- 745 1. $\text{rf}(S)$ on input $\langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), 1 \rangle \rangle$ prints the first query made by $S(1^n)$.
- 746 2. $\text{rf}(S)$ on input $\langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), \text{bin}(i) \rangle \rangle$ prints $q_{n,i}$.
- 747 3. $\text{rf}(S)$ runs in time $O(\ell + n^c \log n)$ on all inputs of the form $\langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), \text{bin}(j) \rangle \rangle$.
- 748 4. The description length of $\text{rf}(S)$ is $\ell + a_{\text{rf}}$ for some absolute constant a_{rf} .

749 *Proof.* Observe that “running $\text{rf}(S)$ on appropriate inputs” is exactly a constructive definition of the queries
750 $q_{n,i}$ for each n and $i < r(n)$. All claimed properties follow by inspection and simulation of $\text{rf}(S)$ because S is
751 deterministic and time-bounded. None of these assertions are about the *validity* of answers $a_{n,j}$ as responses
752 to queries $q_{n,i}$ — they assert only that both sequences are well-defined and can be obtained in bounded time
753 by running and manipulating the description of $\text{rf}(S)$. \square

754 To eliminate $q \geq 1$ queries from S , we answer them with a description of the reflection template applied
755 to S — the standard machine $\text{rf}(S)$. The *Autodidact Template* transforms any student S making at most
756 $r(n)$ queries into a student $\text{ad}(S, q)$ making at most $r(n) - q$ queries by substituting the description of S and
757 $\text{bin}(q)$ into Algorithm 3 below. Preservation of correctness and runtime guarantees is

Algorithm 2 Reflection Template $\text{rf}(S)$

Parameters S a student

```
1: On input  $\langle D, \langle \text{bin}(n), \text{bin}(q) \rangle \rangle$ 
2:  $i \leftarrow 1$   $\triangleright$  assumption:  $S$  makes at least one query
3: loop
4:    $q_{n,i} \leftarrow$  Simulate  $S(1^n)$  until it queries teacher
5:   if  $i < q$  then
6:     Answer the simulated query  $q_{n,i}$  with  $\langle D, \langle D, \langle \text{bin}(n), \text{bin}(i) \rangle \rangle \rangle$   $\triangleright$  exactly  $a_{n,i}$  when  $D = \ulcorner \text{rf}(S) \urcorner$ 
7:      $i \leftarrow i + 1$ 
8:   else
9:     break the loop
10: output  $q_{n,i}$   $\triangleright$  the last query from simulated  $S(1^n)$ 
```

752 *Claim 4.7.* Student $\text{ad}(S, q)$ runs in time $O(n^c \log n)$ and solves $\exists\text{HiK}^t[c + 1]$ using at most $r(n) - q$ queries.

753 *Proof.* We argue by induction, showing first that student $\text{ad}(S, 1)$ solves $\exists\text{HiK}^t[c + 1]$ within the claimed
754 runtime and makes at most $r(n) - 1$ queries. Consider the set of first queries $q_{n,1}$ asked by $S(1^n)$ for each
755 n . These strings depend only on S and n — so intuitively, their $\mathbb{K}^{n^{c+1}}$ -complexity is bounded. Formally, the
756 machine $\text{rf}(S)$ on input $\langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), 1 \rangle \rangle$ prints $q_{n,1}$ for each n in at most $O(n^c \log n)$ steps (items 1 and
757 3 of Claim 4.6). Therefore, the machine-input pair

$$\langle \ulcorner \text{rf}(S) \urcorner, \langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), 1 \rangle \rangle \rangle = a_{n,1}$$

758 of length $O(\ell) + O(\log n)$ witnesses $\mathbb{K}^{n^{c+1}}(q_{n,1}) < n/2$ for all sufficiently large n . Thus, for sufficiently large
759 n , the string $a_{n,1}$ supplied to $S(1^n)$ by line 6 of $\text{ad}(S, 1)$ is a valid answer to query $q_{n,1}$. By the assumption
760 that S solves $\exists\text{HiK}^t[c + 1]$, it must produce an element of $R_{\mathbb{K}^{n^{c+1}}}$ given *any* sequence of valid answers from
761 teacher of length at most $r(n)$. Therefore, the simulation of $S(1^n)$ executed by $\text{ad}(S, 1)$ will solve $\exists\text{HiK}^t[c + 1]$
762 using at most $r(n) - 1$ queries to a real teacher, because $a_{n,1}$ is a valid answer to $q_{n,1}$. Accounting for the
763 time complexity of simulation and string manipulation, $\text{rf}(S, 1)$ takes at most $O(n^c \log n)$ steps on inputs 1^n .
764 This concludes the base case.

765 For the inductive step, suppose that student $\text{ad}(S, i)$ solves $\exists\text{HiK}^t[c + 1]$ using at most $r(n) - i$ queries.
766 Inspecting the autodidact template we have that, when running $\text{ad}(S, i)$: (1) all queries made by S until the
767 $(i + 1)$ -th query are answered by $a_{n,j}$ for $j \in \{1, \dots, i\}$ and (2) query $q_{n,(i+1)}$ is the *first* query answered by
768 teacher. Because $\text{ad}(S, i)$ is a student solving $\exists\text{HiK}^t[c + 1]$, it must produce an element of $R_{\mathbb{K}^{n^{c+1}}}$ given *any*
769 sequence of valid answers from teacher of length at most $r(n) - i$. We argue that $a_{n,(i+1)}$ is a valid answer
770 to query $q_{n,(i+1)}$.

771 The standard, one-tape machine $\text{rf}(S)$ on input $\langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), \text{bin}(i + 1) \rangle \rangle$ prints $q_{n,(i+1)}$ in at most
772 $O(\ell + n^c \log n)$ steps (items 2 and 3 of Claim 4.6). Therefore, the machine-input pair

$$\langle \ulcorner \text{rf}(S) \urcorner, \langle \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), \text{bin}(i + 1) \rangle \rangle \rangle = a_{n,(i+1)}$$

773 of length at most $O(\ell) + O(\log n) + O(\log r(n))$ (by item 4 of Claim 4.6) witnesses $\mathbb{K}^{n^{c+1}}(q_{n,(i+1)}) < n/2$
774 for all sufficiently large n , because we know $r(n) < n^c$ from the runtime bound of S . Therefore, student
775 $\text{ad}(S, i + 1)$ correctly simulates one additional teacher response for S compared to $\text{ad}(S, i)$ and so solves
776 $\exists\text{HiK}^t[c + 1]$ using at most $r(n) - (i + 1)$ queries. Induction on i now proves Claim 4.7. \square

777 Now conclude the proof of Theorem 4.5 by substituting $q = r(n)$ into Claim 4.7 to get that $\text{ad}(S, q)$ solves
778 $\exists\text{HiK}^t[c + 1]$ using zero queries in $\tilde{O}(n^c)$ time, contradicting Proposition 4.4. \square

779 4.2 Gap Between Student-Teacher Search Lower Bounds & VPV-Unprovability

780 The Student-Teacher search lower bounds above do not suffice to obtain VPV-unprovability. Suppose VPV
781 proves $\text{HiK}^t[c]$ for every c . Applying KPT-witnessing, we would obtain for every c a $\text{DTIME}[q_c]$ Student-
782 Teacher search solving $\exists\text{HiK}^t[c]$, for some arbitrary polynomial q_c . There is no contradiction to Theorem 4.5,

Algorithm 3 Autodidact Template $\text{ad}(S, q)$

Parameters $q \in \mathbb{N}$ and S a student

```
1: On input  $1^n$ 
2:  $i \leftarrow 1$   $\triangleright$  assumption: reflect at least one query
3: loop
4:    $q_{n,i} \leftarrow$  Simulate  $S(1^n)$  until it queries teacher
5:   if  $i \leq q$  then
6:     Answer the simulated query  $q_{n,i}$  with  $a_{n,i} = \langle \ulcorner \text{rf}(S) \urcorner, \ulcorner \text{rf}(S) \urcorner, \langle \text{bin}(n), \text{bin}(i) \rangle \rangle \rangle$ 
7:      $i \leftarrow i + 1$   $\triangleright$  increment #queries reflected
8:   else
9:     break the loop
10: Continue simulating  $S(1^n)$  but answer all subsequent queries by asking teacher
11: output the output of simulated  $S(1^n)$ 
```

783 because it does not control the relationship between n^c and q_c . However, if $q_c = \tilde{O}(n^c)$ could be guaranteed
784 for even a single c , then unprovability of the HiK^t schema in VPV would follow.

785 One way forward is to make a stronger assumption about the supposed VPV-proofs of $\text{HiK}^t[c]$. In larger
786 theories than VPV that are known to prove $\text{HiK}^t[c]$ for each c , the proofs are *uniform* — essentially the same
787 for each c . An assumption like “VPV proves $\text{HiK}^t[c]$ for each c and furthermore the proofs are structurally
788 uniform” could enable control over Student runtime, such that a *single* polynomial-time algorithm witnesses
789 K^t -random strings for *every* $t \in \text{poly}$.

790 Such dramatic consequences of uniform proofs might seem unrealistic; the term n^c appears in the
791 quantifier-free part of $\text{HiK}^t[c]$, so shouldn’t any student witnessing $\text{HiK}^t[c]$ take time at least n^c ? This
792 appealing but flawed intuition presumes that witnessing requires simulation of an n^c -time machine. In re-
793 ality, Teacher may be the only party responsible for an n^c -time computation — it depends on the scheme.
794 In Section 4.5 we give several examples of VPV schemata $\Phi[c]$ parameterized by arbitrary polynomial time
795 bounds n^c — with quantifier prefix identical or similar to $\text{HiK}^t[c]$ — where both (1) each statement is prov-
796 able in VPV for every c by the “same” proof and (2) witnessing the statement takes **absolute** polynomial
797 time — **not** n^c for each c .

798 Summarizing the above, it is both plausible and well-motivated to ask for better control over the com-
799 plexity of witnessing terms when VPV proves a parameterized sequence of theorems by “essentially the same”
800 proof. This requires a definition of uniform proofs. Towards this end, we discuss next a similar question
801 about Peano Arithmetic (PA) and extract a witnessing hypothesis for uniform VPV-proofs by analogy.

802 4.3 Kriesel’s Conjecture & Witnessing Hypotheses for Uniform Proofs

803 A fundamental question about “merging” a sequence of theorems into a single theorem appeared in 1975 as
804 Problem 34 on Friedman’s list of One Hundred and Two Problems in Logic, attributed to Kriesel [20]. For
805 some theories, a positive answer to this question would imply uniform witnessing.

806 **Conjecture 4.8** (Kriesel’s Conjecture, §4.4 of [48]). Suppose for a formula $\varphi(x)$ and a number k , one can
807 prove $\varphi(S^c(0))$ in Peano Arithmetic using $\leq k$ steps for every c . Then $\forall c \varphi(c)$ is provable in Peano Arithmetic.

808 Efforts to resolve Kriesel’s Conjecture (KC) uncovered a peculiar situation: KC is very sensitive to how
809 PA is axiomatized! For example, KC is true when PA is axiomatized with a ternary relation for multiplication
810 [45, 40] or with minimality instead of induction [23]. But KC is false when PA has a function symbol for
811 subtraction [22], and remains open for the “textbook” presentation of PA using function symbols $\{S, +, \times\}$.
812 Hrubeš discusses these issues in detail [23].

813 Let PA_L denote the theory of Peano Arithmetic with symbols for every primitive recursive function,
814 axiomatized by a list L of formulas. If KC is true for PA_L , then it is straightforward to extract parameter-
815 independent witnessing terms from a sequence of proofs: just apply KC followed by KPT witnessing. This
816 interchanges the order of quantifiers as desired: a **single** sequence of witnessing terms that works **for all**
817 sentences in the schema. One intermediate step is required — PA is not a universal theory, and so KPT
818 does not apply directly. We work out the details below for $\exists \text{K}^t \text{R}$, towards developing a uniform witnessing

819 hypothesis for the weaker theory VPV by analogy. First recall the KPT theorem, stated below for single-
820 sorted theories.

821 **Theorem 4.9** (Single-Sorted KPT). Let T be a universal theory with vocabulary L . Let φ be an open
822 L -formula, and suppose that $T \vdash \forall \vec{x} \exists y \forall z \varphi(\vec{x}, y, z)$. Then there is a finite sequence t_1, \dots, t_r of L -terms
823 such that

$$T \vdash \forall \vec{x} \forall z_1, \dots, z_r [\varphi(\vec{x}, t_1(\vec{x}), z_1) \vee \varphi(\vec{x}, t_2(\vec{x}, z_1), z_2) \vee \dots \vee \varphi(\vec{x}, t_r(\vec{x}, z_1, \dots, z_{r-1}), z_r)]$$

824 Now translate “for every c and almost every n , there exists a K^{n^c} -random string of length n ” into a
825 PA-formula. Because PA-terms are not bounded by polynomials, here we *can* admit the runtime exponent c
826 as a *free variable*. Fixing sufficiently large n_0 , define

$$\text{HiK}^t(c) := \forall n. (n > n_0) \exists x. (x < 2^n) \forall d. (d < 2^{n/2}) \text{run}(\pi_1(d), \pi_2(d), \text{unary}(\exp(n, c))) \neq x$$

827 For every reasonable list of axioms L , if PA_L includes all primitive recursive functions, it includes the
828 necessary function symbols and proves their relevant properties.

- 829 • $\text{unary}(w)$ is the PA_L symbol for the function that outputs z such that $\text{bin}(z) = 1^w$, and
- 830 • $\text{exp}(n, c)$ is the PA_L symbol for the exponentiation function n^c .
- 831 • run is the PA_L function symbol for run_U from the definition of time-bounded Kolmogorov complexity,
- 832 • $\pi_1(z)$ and $\pi_2(z)$ are the PA_L symbols for the pair decoding functions (see Section 2.7).

833 Again, this translation of $\exists K^t\text{R}$ exploits the power of PA to admit c as a variable of the object language.

834 **Proposition 4.10.** Suppose KC is true for PA_L and there exist absolute constants n_0 and k such that one
835 can prove $\text{HiK}^t(S^c(0))$ in PA_L using $\leq k$ steps for every c . Then, letting PA'_L be any universal conservative
836 extension of PA_L , there is a finite sequence of PA'_L -terms q_1, \dots, q_r such that

$$\text{PA}'_L \vdash \forall c \forall n. (n > n_0) \forall d_1, \dots, d_r \left[\begin{aligned} &(\text{run}(\pi_1(d_1), \pi_2(d_1), \text{unary}(\exp(n, c))) \neq q_1(n, c)) \vee \\ &(\text{run}(\pi_1(d_2), \pi_2(d_2), \text{unary}(\exp(n, c))) \neq q_2(n, c, d_1)) \vee \\ &\dots \vee \\ &(\text{run}(\pi_1(d_r), \pi_2(d_r), \text{unary}(\exp(n, c))) \neq q_r(n, c, d_1, \dots, d_{r-1})) \end{aligned} \right]$$

837 *Proof.* Assume that PA_L proves HiK^t as in the statement of the lemma, and KC is true of PA_L . Applying
838 KC, we have $\text{PA}_L \vdash \forall c \text{HiK}^t(c)$ for some absolute constant n_0 . Now let PA'_L be any universal conservative
839 extension of PA_L . Because PA'_L extends PA_L , we also have $\text{PA}'_L \vdash \forall c \text{HiK}^t(c)$. Because PA'_L is universal,
840 appeal to KPT witnessing (Theorem 4.9) concludes this proof. \square

841 4.4 Conditional Unprovability of $\text{HiK}^t[c]$ in VPV and V^1

842 By analogy to the outcome of assuming KC and applying KPT to a conservative universal extension of PA,
843 introduce the following

844 **Hypothesis 4.11** (Witnessing for Linecount-Uniform VPV-Proofs). Let $\varphi(n, p, X, Y)$ be a Σ_0^B (VPV) formula
845 with all free variables displayed. Suppose there is an absolute constant n_0 , number k , and VPV-term t such
846 that one can prove $\forall n. (n > n_0) \exists X. (|X| < t(n)) \forall Y \varphi(n, n^c, X, Y)$ in VPV using $\leq k$ steps for every c . Then
847 there is a finite sequence F_1, \dots, F_r of VPV-function symbols that are *absolutely witnessing*:

$$\text{for every } c, \text{VPV} \vdash \forall n. (n > n_0) \forall Y_1, \dots, Y_r \left[\begin{aligned} &\varphi(n, n^c, F_1(n, c), Y_1) \vee \\ &\varphi(n, n^c, F_2(n, c, Y_1), Y_2) \vee \\ &\dots \vee \\ &\varphi(n, n^c, F_r(n, c, Y_1, \dots, Y_{r-1}), Y_r) \end{aligned} \right]$$

848 The asymmetry in how c is given to φ compared to how c is given to each F_i — n^c vs. c — is crucial for
849 our applications. If Hypothesis 4.11 holds, then any student derived from the hypothesis takes arguments 1^n
850 and 1^c because numeric terms are supplied in unary for two-sorted complexity classes (see Section 2.4). If c
851 were instead given to F_i as n^c , the implicit student would take $\text{poly}(n^c)$ time to print K^{n^c} -random strings of
852 length n — and no contradiction would arise. However, combining Hypothesis 4.11 with the Student-Teacher
853 lower bounds for $\exists\text{HiK}^t$ from the last section (Theorem 4.5), we have

854 **Corollary 4.12.** Under the Witnessing Hypothesis for Linecount-Uniform VPV-Proofs, there is no fixed k
855 such that one can prove $\text{HiK}^t[c]$ in VPV using $\leq k$ steps for each c .

856 This would rule out *linecount uniform* proofs of $\text{HiK}^t[c]$. However, linecount uniformity — though well-
857 motivated by Kriesel’s Conjecture — is certainly not the only reasonable notion of uniformity in proofs. We
858 hope that a deeper understanding of uniform VPV-proofs will emerge by studying witnessing hypotheses
859 that emphasize different aspects of common structure in theorems and proofs. To begin the investigation,
860 we introduce a strong witnessing hypothesis that emphasizes the common element in statements like $\text{HiK}^t[c]$
861 — substitution of polynomial time-bounds into the execution of Turing machines, formalized as

862 **Definition 4.13** (poly-Runtime Schema). Fix a universal function symbol $\text{run}(M, A, s)$ to output the tape
863 of machine M run on input A for s steps. An infinite sequence of formulas Φ is a *poly-runtime schema* if Φ is
864 obtained by taking an infinite union over substitution of polynomial runtimes. Formally, let φ be a formula
865 with free a variable p occurring only in terms of the form $\text{run}(M, A, p)$ — as the time bound. Then,

$$\Phi = \bigcup_{c \in \mathbb{N}} \varphi(p/n^c)$$

866 We refer to the c -th sentence in such a schema by Φ_c .

867 **Hypothesis 4.14** (Witnessing for poly-Runtime Schema in VPV). Suppose Φ is a poly-runtime schema with
868 $\varphi = \forall n.(n > n_0) \exists X.(|X| < t(n)) \forall Y \psi(n, p, X, Y)$ for ψ a $\Sigma_0^B(\text{VPV})$ formula and t a VPV-term, and there
869 is an absolute constant n_0 such that $\text{VPV} \vdash \Phi$. Then there is a finite sequence F_1, \dots, F_r of VPV-function
870 symbols that are *absolutely witnessing*:

$$\text{for infinitely many } c, \text{VPV} \vdash \forall n.(n > n_0) \forall Y_1, \dots, Y_r \left[\begin{aligned} &\psi(n, n^c, F_1(n, c), Y_1) \vee \\ &\psi(n, n^c, F_2(n, c, Y_1), Y_2) \vee \\ &\dots \vee \\ &\psi(n, n^c, F_r(n, c, Y_1, \dots, Y_{r-1}), Y_r) \end{aligned} \right]$$

871 The conclusion is essentially identical to that of the linecount WHUP. However Hypothesis 4.16 is much
872 stronger: it asserts that VPV cannot help but give absolute witnessing if it proves a poly-runtime schema.
873 Therefore, combining Hypothesis 4.16 with the Student-Teacher lower bounds for $\exists\text{HiK}^t$ from the last section
874 (Theorem 4.5), we have

875 **Corollary 4.15.** Under the Witnessing Hypothesis for poly-Runtime Schemas in VPV, there are infinitely
876 many c such that VPV does not prove $\text{HiK}^t[c]$.

877 Under Hypothesis 4.16, we get VPV-unprovability of $\text{HiK}^t[c]$, but not V^1 unprovability. This under-
878 exploits our Student-Teacher lower bounds for $\exists\text{HiK}^t$, which can eliminate poly-many rounds from Student.
879 So, we introduce an appropriate WHUP for V^1 — derived from the KPT Theorem for V^1 (Theorem 2.15).

880 **Hypothesis 4.16** (Witnessing for poly-Runtime Schema in V^1). Suppose Φ is a poly-runtime schema with
881 $\varphi = \forall n.(n > n_0) \exists X.(|X| \leq t(n)) \forall Y \psi(n, p, X, Y)$ for ψ a $\Sigma_0^B(V^1)$ formula and t a V^1 -term, and there is an
882 absolute constant n_0 such that $V^1 \vdash \Phi$. Then there is an *absolutely witnessing* FP function F such that for
883 infinitely many c ,

$$\mathbb{N}_2 \models \forall n.(n > n_0) \forall Y \psi(n, n^c, F^{CX[\Phi_c]}, Y)$$

884 **Corollary 4.17.** Under the Witnessing Hypothesis for poly-*Runtime Schemas* in V^1 , there are infinitely
885 many c such that V^1 does not prove $\text{HiK}^t[c]$.

886 These two corollaries imply separations with Jeřábek’s theory VAPC .

887 **Theorem 4.18.** $\text{VAPC} \vdash \text{HiK}^t[c]$, for all $c \in \mathbb{N}$. Further, under Witnessing Hypotheses, $\text{VPV} \not\vdash \text{HiK}^t[c]$ and
888 $V^1 \not\vdash \text{HiK}^t[c]$.

889 *Proof.* It was shown by Korten [27] that $\text{VAPC} \vdash \text{HiK}^t[c]$. □

890 We spend the remainder of this section addressing the plausibility of these hypotheses, by giving examples
891 of VPV -theorems that do enjoy absolute witnessing despite varying polynomial bounds.

892 4.5 Examples of Schemata With “Uniform” Proofs & Absolute Witnessing

893 The WHUPs discussed in this section apply to VPV -schemata of the form

$$\Phi[c] := \forall n.(n > n_0) \exists X \forall Y \varphi(n^c, X, Y)$$

894 where $\varphi(p, X, Y)$ is Σ_0^B for each $c \in \mathbb{N}$. Here we give examples of simple VPV -theorems to illustrate that
895 this class of schemata is non-trivial. All these examples have both proofs that are identical up to numeric
896 substitutions and witnessing algorithms that run in some **absolute** polynomial time — **not** n^c for each c .
897 Therefore, no contradiction can arise from assuming a WHUP (and constant-line proofs) for any of these
898 theorems. The WHUP would just “automatically” transform proofs into witnessing algorithms that meet
899 known complexity upper bounds. The common element in all these examples is efficient transformation
900 of encoded Turing Machines. For each example we describe the VPV -translation and carefully discuss the
901 complexity of witnessing. We do not argue for VPV -provability, because all these theorems follow from
902 properties of universal machines and lemmas about efficient string manipulation that are readily available
903 in VPV — see the discussion in Sections 2.1 and 4 of [47].

904 4.5.1 Machine Templates

905 The first three examples give basic properties of machine-only Kolmogorov complexity. Fix a universal
906 Turing machine U and define the *machine-only t -time bounded Kolmogorov Complexity* $\text{moK}_U^t(x)$ of a string
907 x as the length of the shortest encoded machine that prints x when simulated by U :

$$\text{moK}_U^t(x) = \min_{d \in \{0,1\}^*} \{|d| : U(d, \varepsilon, 1^{t(|x|)}) = x\}$$

908 This definition is brittle compared to standard time-bounded Kolmogorov complexity. The UTM never
909 provides any input to the encoded machine d , forcing d to “hardcode” useful strings instead of reading them
910 from an input tape. Therefore the basic fact about K^t — $\forall x \text{K}^t(x) < |x| + a$ for an absolute constant a
911 — fails. However, we can recover something similar for moK^t , even in VPV : an uniform upper bound on
912 $\text{moK}^t(x)$ for every x .

913 **Memorization Templates.** For every polynomial time bound t , for every string length n , there is a
914 hardcoded-string “template” machine M of length n , such that any string X of “sufficiently smaller” length
915 can be pasted into the template to produce a new machine M' . The machine M' prints X in less than t
916 time. Pasting is a polynomial-time string function that copies the bits of Y into a sequence of states of M .
917 We formalize this as a VPV -schema below, varying the polynomial time bound.

$$\text{MEMT}[n_0, c] := \forall n.(n > n_0) \exists M.(|M| = n) \forall X.(|X| \leq n/16) \text{run}_U(\text{paste}_U(M, X), n^c) = X$$

918 VPV cannot quantify over arbitrary polynomial time bounds, but it can prove the MEMT schema via an
919 essentially-identical proof for each c . However, no contradiction can arise from a WHUP because it is easy
920 to witness M : print the U -encoding of a Turing Machine that prints an *explicit* all-zero string instead of an
921 *implicit* all-zero string. That is, the i th state of M is “write 0 to the tape, move the head right, transition
922 to state $i+1$.” The *paste* function replaces the “write 0” element of state i of M with bit $X(i)$. The content

923 of this simple theorem is the gap between n and $|X|$ — it asserts an upper bound on the cost of memorizing
 924 a string relative to some fixed model of computation and encoding of machines shared by run_U and paste_U .

925 Notice that witnessing M in this example takes linear time completely independant of c . This is more
 926 restrictive than the consequences of a WHUP, which allows witnessing algorithms to take 1^c as an argument.
 927 Our next example actually exploits this dependence.

928 **Clocking Templates.** For every polynomial time bound p , for each sufficiently large n , a “template”
 929 machine M of length n enforces a p -step timeout on shorter machines, making sure they halt in time p
 930 and signalling a fault if they run too long. We’ll formalize this in VPV using a pair encoding function: the
 931 clocking template applied to machine description D outputs $\langle h, \text{run}(D, \varepsilon, n^c) \rangle$ where h is 1 if D halted within
 932 n^c steps and zero otherwise. Consider the following collection of VPV-theorems $\text{CLOCKT}[n_0, c] :=$

$$\forall n. (n > n_0) \exists M. (|M| = n) \forall D. (|X| \leq n/16) (\text{halt}(D, \varepsilon, n^c) \rightarrow \text{run}(\text{paste}(M, D), \varepsilon, n^{2c}) = \langle 1, \text{run}(D, \varepsilon, n^c) \rangle) \\ \wedge (\neg \text{halt}(D, \varepsilon, n^c) \rightarrow \text{run}(\text{paste}(M, D), \varepsilon, n^{2c}) = \langle 0, \text{run}(D, \varepsilon, n^c) \rangle)$$

933 Once again, there is a straightforward witnessing for M : print the U -encoding of a machine that *explicitly*
 934 prints the all-zero string Z of length $n/16$ to the worktape (as in the memorization template), and then runs
 935 a n^c -clocked U to simulate Z . Pair the worktape contents of the results with 0 or 1 depending on if Z halted.
 936 The paste function then replaces the explicitly-coded Z with the encoding of D , resulting in a template with
 937 the desired behaviour.

938 Witnessing this template actually depends on c : the clock requires $c \log(n)$ hardcoded bits in the descrip-
 939 tion of M . However, this dependence is *not* polynomial: for sufficiently large n , $c \log(n) < n$. Inspecting the
 940 WHUPs for VPV (Hypotheses 4.11, 4.16) we see that the witnessing function symbols occur as $F(n, c, \dots)$,
 941 meaning that n and c are given in unary to the witnessing algorithm. Therefore, in fixed $\text{poly}(n, c)$ time we
 942 can hardcode the binary representation of n^c into a clock. This is an example where the straightforward
 943 witnessing has *exactly* the complexity implied by a WHUP.

944 **Clocked Unrolling Templates.** VPV can also discuss a local formulation⁶ of machine-only $K^t(x)$, which
 945 bounds the time complexity of producing each individual bit x_i of x given i in binary. Consider the following
 946 polynomial-time function, which “unrolls” a given machine into an n -bit string — essentially a machine
 947 analog of the truth-table generator for circuits [30].

Algorithm 4 Unrolling a Machine, $\text{Unroll}(D, n, n^c)$

Parameters n in unary, n^c in unary

```

1: for all  $i \in \{0, \dots, n\}$  do
2:   if  $D$  run on input  $\text{bin}(i)$  accepts within  $n^c$  steps then
3:     | Print 1
4:   else
5:     | Print 0
```

948 For every polynomial time bound p , for each sufficiently large n , a “template” machine M of length n
 949 can extract an n -bit vector of p -step decisions from sufficiently shorter machines. That is, pasting a shorter
 950 machine D into M and running the result agrees with $\text{Unroll}(D, n, p)$. Translating into VPV define the
 951 schemea $\text{UNROLLT}[n_0, c] :=$

$$\forall n. (n > n_0) \exists M. (|M| = n) \forall D. (|D| \leq n/16) \text{run}_U(\text{paste}_U(M, D), \varepsilon, n^{2c+1}) = \text{Unroll}(D, n, n^c)$$

952 Witness M in $\text{poly}(n, c)$ time by printing an appropriate U -encoding of Algorithm 5 below.
 953 To accomodate paste , implement line 1 of M by *explicitly* printing 0 symbols — one state per symbol, exactly
 954 as in the previous two templates. Implement lines 3 and 4 by maintaining binary counters on the worktape.
 955 This requires $O(\log n)$ and $O(c \log n)$ bits to be hardcoded in M , respectively. Finally, the code of Unroll
 956 takes some absolute constant number of bits in the encoding of M . Just as above, printing M takes fixed
 957 polynomial time given $(1^n, 1^c)$ as input.

⁶A local formulation of standard K^t complexity appears, for example, as Definition 3 of [37] where a hardness assumption about deciding local K^t is used in a direct and elegant construction of pseudo-random functions.

Algorithm 5 Unrolling Template

Parameters D the description of a machine, n in unary, n^c in unary

- 1: Write $0^{n/16}$ to the worktape
 - 2: Move the head two cells right — leaving a blank
 - 3: Write 1^n to the worktape
 - 4: Move the head two cells right — leaving a blank
 - 5: Write 1^{n^c} to the worktape
 - 6: Run Unroll on the contents of the worktape, with arguments separated by blanks
-

958 4.5.2 Deterministic Time Hierarchy Theorem

959 Consider the compressible-counterexample deterministic time hierarchy theorem, used to obtain Student-
960 Teacher lower bound for constructing circuits [8].

961 **Lemma 4.19.** For every $c \in \mathbb{N}$, there is a language $H_c \in \text{DTIME}[n^{c+1}]$ satisfying the following:

- 962 • **COUNTEREXAMPLES:** Every candidate n^c -time TM M that tries to compute H_b will make a mistake
963 on an n -bit input $x_{\text{error}} = \ulcorner M \urcorner \circ \pi$ where \circ denotes concatenation and $\pi \in 0^*$ is a padding string
964 chosen to make $|x_{\text{error}}| = n$ for all sufficiently large n .
- 965 • **COMPRESSIBILITY OF COUNTEREXAMPLES:** The counterexamples x_{error} are efficiently compressible to
966 $O(\log(n))$ bits by recording both the constant-length description M and n in binary, by just padding
967 M to the appropriate length.

968 Though the diagonalization machine H_c uses time $O(n^{c+1})$, the implicit refuter uses only time $O(n)$ —
969 and is the same regardless of which polynomial “slice” of the hierarchy is being refuted! The deterministic
970 time hierarchy theorem has a straightforward translation into a sequence of VPV-sentences.

Formalization 4.20.

$$\text{DTIMEH}[c] := \forall n \forall M. (|M| < n/16) \exists X. (|X| = n) \text{run}(M, X, n^c) \neq \text{run}(H_c, X, n^{c+1})$$

971 This is a simpler formula than the VPV-schemata $\Phi[c]$ used in WHUPs, because the quantifier prefix is
972 $\forall \exists$ instead of $\forall \exists \forall$. We know that $\text{VPV} \vdash \text{DTIMEH}[c]$ for each c , and each proof is “essentially the same”
973 up to substitution of n^c (Lemma 3.1 of [31]). Therefore, Buss Witnessing (Theorem 2.13) applies and we
974 immediately get refuters for H_c . But the uniformity is not exploited by Buss Witnessing – we get a sequence
975 of refuters with arbitrary and unrelated polynomial runtime for each c , and indeed each runtime may be
976 much larger than n^c . This is much worse than the absolute refuter obtained outside VPV.

977 In this simpler setting, is there a generic way to convert such uniform collections of proofs into absolute
978 witnessing that we *already know exists*? To further assess the plausibility of such convenient witnessing, we
979 let w_H be the VPV-term given by the compressible counterexamples of Lemma 4.19 and ask the following

980 **Question 4.21.** Does VPV prove that w_H witnesses the $\text{DTIMEH}[c]$ errors for each c ?

981 4.5.3 Efficient Conversion From Multi-Tape to One-Tape Turing Machines

982 It is a classical theorem that for $k \geq 2$ any k -tape Turing Machine can be simulated by a one-tape Turing
983 Machine with at most quadratic overhead.

984 **Theorem 4.22** (Claim 1.6 of [3]). If the language L can be decided in time n^c on a k -tape Turing Machine,
985 then L can be decided in time $16kn^{2c}$ on a single-tape Turing Machine.

986 Formalize this in VPV by defining the function symbols run_k and run_1 to simulate k -tape and single-tape
987 Turing Machines, respectively. Then consider the following collection of VPV-theorems $\text{ONE.TAPE}[n_0, k, c] :=$

$$\forall M \exists M' \forall n. (n \geq n_0) \forall X. (|X| = n) \text{run}_k(M, X, n^c) = \text{run}_1(M', X, 16kn^{2c})$$

988 This quantifier prefix is identical to that of $\Phi[c]$, but the types are different: strings instead of numbers.
989 Therefore, a witnessing algorithm for ONE.TAPE is given M encoded in binary and must print the encoding

990 of M' . The encoding length $|M|$ under any reasonable encoding — which we fix using run_k — is determined
 991 by the number of states and alphabet size. The number of states in M' given by straightforward proofs
 992 of Theorem 4.22 is exponential in k but linear in the states and alphabet-size of M . Therefore, in a fixed
 993 polynomial time in $|M|$, a witnessing algorithm prints M' . Only the transformation of the “code” of M
 994 matters to the witnessing algorithm — **not** the runtime bound on M .

995 5 Consequences of Provably Hard Truth Tables

996 Under a natural witnessing hypothesis for a theory corresponding to uniform $\text{AC}^0[\text{qpoly}]$, we have $\text{P} \neq \text{NP}$.

997 5.1 A Theory for $\text{AC}^0[\text{qpoly}]$ -Reasoning

998 In Section 2.4, we recalled the two-sorted theory \mathbf{V}^0 which corresponds to log-uniform AC^0 circuits. Here,
 999 we extend the definition to a new theory denoted $\mathbf{V}^0_{\#}$, corresponding to $\text{polylog-uniform AC}^0[\text{qpoly}]$. The
 1000 definition of $\mathbf{V}^0_{\#}$ is very simple: starting from the axioms of \mathbf{V}^0 and language $\mathcal{L}(\mathbf{V}^0)$, add the function symbol
 1001 $\#$, commonly known as the *smash* operator, to the language and defining axioms of \mathbf{V}^0 . Smash is defined
 1002 by axioms stating $x\#y = 2^{|x|\cdot|y|}$, for numbers x, y . It is used to give quasipolynomial growth rates of the
 1003 number type.

1004 The characterization of \mathbf{V}^0 with uniform AC^0 is carried out in detail in Chapters IV and V of [17]. They
 1005 treated AC^0 as the logtime-hierarchy LH, known to be equivalent to logtime-uniform AC^0 .

1006 **Theorem 5.1** (Folklore). $\text{log-uniform AC}^0[\text{poly}] = \text{LH}$.

1007 Importantly, this is generalizable to the $\text{polylog-hierarchy polyLH}$.

1008 **Theorem 5.2** (Folklore). $\text{polyLH} = \text{polylog-uniform AC}^0[\text{qpoly}]$

1009 Cook and Nguyen showed that all logtime-uniform AC^0 -functions are Σ_1^B -definable in \mathbf{V}^0 , as well as the
 1010 converse witnessing theorem that any $\forall\Sigma_1^B$ sentence provable in \mathbf{V}^0 has its existential quantifier witnessed
 1011 by a logtime-uniform AC^0 function.

1012 This correspondence holds for uniform $\text{AC}^0[\text{poly}]$, but generalizes to any class of circuit sizes that is closed
 1013 under composition, with the appropriate modification of the language and axioms. By adding the smash
 1014 operator $\#$, it is standard to get an identical correspondence between $\text{polylog-uniform AC}^0[\text{qpoly}]$ and $\mathbf{V}^0_{\#}$.

1015 We believe this theory is of independent interest, and will discuss its strength at the end of the section.

1016 5.2 Stating Existential Circuit Lower Bounds in $\mathcal{L}(\mathbf{V}^0_{\#})$

1017 We first give a logical translation of the classical lower bound due to Shannon.

1018 **Theorem 5.3** (Shannon Counting). Let $b > 0$. For every sufficiently large N , there exists a truth table X
 1019 of length N which is not succinctly represented by any $|N|$ -input circuit of size $|N|^b$.

1020 Normally, it would be impossible to describe such a lower bound in $\mathcal{L}(\mathbf{V}^0)$ or $\mathcal{L}(\mathbf{V}^0_{\#})$, as it involves
 1021 evaluating *general* circuits, instead of AC^0 circuits. However, because our feasible objects will be truth
 1022 tables of length 2^n , we can evaluate general circuits of each fixed polynomial size n^c in size $\text{qpoly}(2^n)\text{-AC}^0$.
 1023 While we normally reserve capital letters for string-types, we will use N to refer to 2^n in this section. More
 1024 formally, we use the following folklore lemma about AC^0 evaluation of general circuits.

1025 **Lemma 5.4** (Folklore). Let $k > 0$. There is a $\text{polylog-uniform AC}^0$ circuit of size $N^{\log(N)^{3k}}$ which on input
 1026 the DCL encoding of a general circuit C of size n^k and an input x of n bits, outputs $C(x)$.

1027 *Proof.* By a standard counting argument, it is known that there are at most $2^{O(s(n)\log n)}$ circuits of size $s(n)$.
 1028 As a DCL representation of a size $s(n)$ circuit is a string of length at most $s(n)^2$, we have that there are
 1029 at most $2^{O(s(n)^2 \log n)}$ DCL strings of size $s(n)$ circuits. Plugging in $s(n) = n^k$, we get that there are up to
 1030 $2^{O(n^{2k} \log n)} = O(N^{(\log N)^3})$ DCL strings of size n^k circuits.

1031 We construct an AC^0 circuit \mathcal{E} to evaluate any size n^k general circuit as follows:

- 1032 1. CIRCUIT LOOKUP LAYER: Have a multiplexer identify the which circuit C has been input
- 1033 2. INPUT LOOKUP LAYER: Have a multiplexer identify the circuit input x which has been specified.
- 1034 3. EVALUTATION LAYER: Output the memorized evaluation of indentified circuit C and input x .

1035 **Uniformity.** Normally, the above circuit would be highly non-uniform. However, in the size regime of
 1036 $N = 2^n$, this becomes feasible. A polylog-uniformity algorithm of \mathcal{E} would have runtime $\text{polylog}(N^{\log(N)^{2k}})$,
 1037 which is $\text{DTIME}[\text{poly}(n)]$. This means that a uniformity algorithm $A_{\mathcal{E}}$ running in time $\log(N)^c = n^c$ for
 1038 circuit \mathcal{E} has time to evaluate circuits of size $n^{c/3}$. Setting $c > 3k$ would allow for the uniformity algorithm
 1039 to, after stages (1) and (2), evaluate the input (C, x) and give the output bits. \square

1040 By Lemma 5.4, we have in $V_{\#}^0$ a sequence of function symbols for generating the truth tables of fixed-
 1041 polynomial size *general* circuits given as input. This is feasible due to the input length N , where a fixed
 1042 polynomial is only polylog. Define the sequence of symbols $\text{TT}_b(C, N)$ as functions that take a number N
 1043 and circuit C of length $|C| = |N|^b$ and output the truth table of C of length N . These operations have
 1044 function symbols and defining axioms in $V_{\#}^0$ because the polylog-uniform AC^0 complexity will be $N \cdot \text{qpoly}(N)$
 1045 to evaluate the circuit C on each of the N possible inputs (by Lemma 5.4). We will also assert that $\text{TT}_b(\cdot, \cdot)$
 1046 checks if the input circuit is valid and of length $|N|^b$, and treats it as the constant 0 function if it is not.
 1047 This is because verifying a DCL encoding can be done efficiently in AC^0 . We can now give the following
 1048 translation,

$$\text{Hard}(b) \triangleq \forall N \exists X (|X| = N) \forall D (|D| < N) \text{TT}_b(D, N) \neq X$$

1049 The above formula makes a choice to rely on the function symbol TT_b verifying that N is a power of two and
 1050 the circuit D is a valid circuit of size $|N|^b$ instead of explicitly verifying this outside of the function symbol.
 1051 We make this choice because we will need a WHUP for $V_{\#}^0$, and the cleanest presentation of such a WHUP
 1052 is given when TT_b absorbs the circuit verification procedure. See the discussion on WHUPs below for more
 1053 detail.

1054 **Comparison to VAPC and Shannon Counting.** The typical description of Shannon counting is that
 1055 there exists a truth table x of length N , which has circuit complexity $N/\log N$. It is this formulation
 1056 which Jeřábek showed is provable in VAPC. Our logical translation, however, is weaker: we only require
 1057 proving a *schema* that asserts a truth table with *super-fixed-polynomial* circuit complexity exists, rather
 1058 than exponential.

1059 5.3 Round Elimination of the Student-Teacher Refuter

1060 **Student-Teacher Interpretation** The structure of the Student-Teacher game is very similar to previous
 1061 sections. In each round, a polylog-uniform $\text{AC}^0[\text{qpoly}(N)]$ Student constructs a truth table X and queries
 1062 the Teacher. Every round that the Student is not correct, the Teacher will respond with a small circuit D
 1063 that succinctly represents X . Crucially, Teacher's response (to be replaced with a SearchMCSP oracle) is of
 1064 length $\text{polylog}(N)$ and computable in PH.

1065 The round elimination strategy will be different from previous sections. We will show that the problem
 1066 of outputting the i -th bit of the Student-Teacher game is in fact in the polylog hierarchy polyLH, and use the
 1067 assumption $\text{P} = \text{NP}$ to show that the output of the Student-Teacher game will have small circuit complexity.
 1068 We begin with a warm-up lemma.

1069 **Lemma 5.5** (Lemma 2.5, [11]). Assume $\text{P} = \text{NP}$. Then for every polylogtime-uniform AC^0 algorithm A
 1070 which outputs n bits on input 1^n , the output $A(1^n)$ has circuit complexity at most $\text{polylog}(n)$.

1071 *Proof.* Let D be the uniformity machine for AC^0 algorithm A which on input n in binary and index i in
 1072 binary, reports the i -th bit of wire and gate information of A_n , the n -th AC^0 circuit of family A . Let $f(n, i)$
 1073 be the function that outputs the i -th output bit of $A_n(1^n)$. Notice that f is in PH: due to A_n being constant
 1074 depth, one can existentially and universally guess gate/wire information and verify it due to D . This means
 1075 the evaluation of f is in $\Sigma_d\text{-TIME}[O(\log^d n)]$ for some constant d depending on the depth of circuit A and

1076 the polynomial time SAT algorithm. By assumption, $P = PH$, hence the evaluation of $A_n(1^n)$ may be done
 1077 in deterministic $\text{polylog}n$ time. It is standard to convert such a program to a circuit of polylog size. \square

1078 The above lemma is a blueprint for the generalization to Student-Teacher games. Let φ_b denote the
 1079 quantifier-free part of $\text{Hard}[b]$, to state

1080 **Lemma 5.6.** Assume $P = NP$. Let $r \in \mathbb{N}$ be a constant. Then any polylog-uniform $\text{AC}^0[N^{(\log N)^k}]$ - $\text{ST}^{CX[\varphi_b, r]}$
 1081 game for $\text{Hard}[b]$ has output of circuit complexity $\log(N)^m = n^m$ on inputs 1^N , for some $m = m(k) > 0$.
 1082 Furthermore, if $k = O(1)$, then $m = O(1)$.

1083 *Proof.* We exactly follow the proof of Lemma 5.5 with one major modification: we must replace the oracle
 1084 gates/Teacher's responses.

1085 Let the Student S be an $\text{AC}_d^0[N^{(\log N)^k}]$ counterexample oracle circuit with $d > r$ and a uniformity
 1086 algorithm $A(i, n)$ which runs in time $(\log N)^{q_1}$. As well, let $P = NP$ be realized by a polynomial time SAT
 1087 algorithm of time n^α . As $\text{SearchMCSP} \in \text{FNP}$ there is by assumption a fixed polynomial $p = p(N, |s(n)|)$,
 1088 for $s(n)$ a size function $s(n) < 2^n/n$, where $\text{SearchMCSP} \in \text{DTIME}[p(N, s(n))]$. With $s(n) = n^b$, we have
 1089 that there is a LOGTIME-uniform circuit family $\{C_N\}_N$ of size $p(N)^k$ solving SearchMCSP . By Lemma 5.4,
 1090 we can evaluate C_N by a polylog-uniform $\text{AC}_3^0[N^{\log(N)^{3k}}]$ circuit \mathcal{E}_N . Let the uniformity algorithm for \mathcal{E} , $A_{\mathcal{E}}$,
 1091 run in time $(\log N)^{q_2}$ for some constant q_2 . From student S , we modify the oracle circuit, by replacing any
 1092 oracle gate by the circuit solving $\text{SearchMCSP}(\cdot, n^b)$, \mathcal{E}_N , and all oracle output bits by the output bits of \mathcal{E}_N .
 1093 Denote this new circuit S^* .

1094 By the same proof of Lemma 5.5, the output $S^*(1^N)$ has circuit complexity $(\log N)^m = n^m$, where
 1095 $m = 100d\alpha k \max(q_1, q_2)$. This follows by the repeated application of the polynomial time SAT algorithm,
 1096 and the substitution of \mathcal{E}_N into S for each Teacher oracle. \square

1097 Finally, we will introduce a WHUP for $\mathbb{V}_{\#}^0$ in order to obtain an absolute Student-Teacher game for any
 1098 $b \in \mathbb{N}$ and $\text{Hard}[b]$. Proof-theoretic consequences will follow.

1099 **Witnessing Hypothesis.** The structure of the schema $\text{Hard}[b]$ is somewhat different from the schema
 1100 $\text{HiK}^t[c]$ for the existence of high K^{poly} strings seen in Section 4. For $\text{Hard}[b]$, we are substituting *function*
 1101 *symbols* instead of substituting runtimes, contrasting with our WHUP for VPV which substitutes runtimes
 1102 into a universal machine. Such a difference is natural when we go from reasoning with Turing machines to
 1103 reasoning with circuits. Another variation of WHUPs is required to handle varying function symbols.

1104 **Definition 5.7.** A *parametrized uniform circuit family* $\{C_n(b)\}_n$ is a circuit family where the uniformity
 1105 algorithm $A(i, n, b)$ takes in an index to the DCL i , the input length n , and an additional parameter b , all
 1106 represented in binary. Unless stated otherwise, we will only consider polylog-uniformity.

1107 **Definition 5.8.** Fix a parametrized uniform $\text{AC}^0[\text{qpoly}]$ family $C(b) = \{C_n(b)\}_n$ and let $f_b(\bar{X})$ be the $\mathbb{V}_{\#}^0$
 1108 function symbol which evaluates C on input \bar{X} with parameter b for the uniformity machine. We say that
 1109 an infinite sequence of formulas Φ is a *parametrized uniform schema* if Φ is obtained by taking an infinite
 1110 union over parameter values of a parametrized uniform circuit family. Formally, let φ be a formula which
 1111 has a parametrized uniform function symbol f_p . We denote $\varphi(f_{p/b})$, $b \in \mathbb{N}$ to be “substituting” the value b
 1112 for the parameter p , where we use the function f_b wherever f is named in φ . We set

$$\Phi = \bigcup_{c \in \mathbb{N}} \varphi(f_{p/c}).$$

1113 Outside the theory, this can be thought of as a substitution; all we are doing is substituting numerals into
 1114 the parameter of a uniformity algorithm. However, it is **not** a true term substitution in the object language,
 1115 as a first order theory cannot treat functions as free variables.

1116 The following WHUP for $\mathbb{V}_{\#}^0$ strengthens the WHUP for VPV using parametrized uniform circuits.

1117 **Hypothesis 5.9** ($\mathbb{V}_{\#}^0$ Witnessing Hypothesis for Uniform Proofs). Suppose $\Phi = \bigcup_{c \in \mathbb{N}} \varphi(f_{p/c})$ is a parametrized
 1118 uniform schema with $\varphi \in \Sigma_2^B$ and $\mathbb{V}_{\#}^0 \vdash \Phi$. Then there is a finite sequence F_1, \dots, F_r of $\mathbb{V}_{\#}^0$ -function symbols
 1119 such that, for infinitely many c ,

$$\begin{aligned} \mathbb{V}_{\#}^0 \vdash \forall n \forall \vec{Y}_1, \dots, \forall \vec{Y}_r \left(\varphi(n, n^c), F_1(n, c), \vec{Y}_1 \right) \vee \varphi(n, n^c, F_2(n, c, \vec{Y}_1), \vec{Y}_2) \vee \\ \dots \vee \varphi(n, n^c, F_r(n, c, \vec{Y}_1, \dots, \vec{Y}_{r-1}), \vec{Y}_r) \end{aligned}$$

1120 We are now ready to show our provability consequences for $\mathbb{V}_{\#}^0$.

1121 **Theorem 5.10.** Assume Hypothesis 5.9. If $\mathbb{V}_{\#}^0 \vdash \text{Hard}(b)$, for every $b \in \mathbb{N}$, then $\text{P} \neq \text{NP}$.

1122 *Proof.* Because $\text{Hard}[b]$ is a parametrized-uniform schema, under Hypothesis 5.9, there is a *fixed* polylog-
1123 uniform $\text{AC}^0[\mathbf{qpoly}(N)]$ student S for the Student-Teacher game of $\text{Hard}(b)$, for infinitely many $b \in \mathbb{N}$. For
1124 sake of contradiction, assume $\text{P} = \text{NP}$. Let $m = m_{(5.6)}$ be the exponent in the circuit size of $S(1^N)$, per
1125 Lemma 5.6. Note that m is a constant. If $b > m$, then we have a contradiction, as the output $S(1^N)$ will
1126 have circuit complexity smaller than n^b . \square

1127 5.4 $\text{VPV}_{\#}$ Proves $\text{Hard}(b)$ for Every b

1128 Before demonstrating that $\mathbb{V}_{\#}^0 \vdash \text{Hard}(b)$, we show as a conceptually simpler task that $\text{VPV}_{\#} \vdash \text{Hard}(b)$, for
1129 an appropriate logical translation of weak Shannon counting in $\mathcal{L}(\text{VPV}) \cup \{\#\}$. The theory $\text{VPV}_{\#}$ takes the
1130 theory VPV and adds the functions symbols and defining axioms for $|\cdot|_1$ and $\#$, where $|\cdot|_1$ gives the length
1131 of a *number*, and $\#$ is the smash operator. With these additions, we have the following $\mathcal{L}(\text{VPV}_{\#})$ logical
1132 translation $\text{Hard}(b)$,

$$\forall n, N. (n = |N|_1) \exists F. (|F| = N) \forall D. (|D| = n^b) \text{TT}(D, N) \neq F$$

1133 We can prove weak Shannon counting in this theory by *iterated halving*, the classic procedure common
1134 in learning algorithms. Let \mathcal{C}_b be the set of Boolean circuits of size n^b with n input bits. There are at most
1135 $2^{O(n^b \log n)}$ such circuits. If $N = 2^n$, then this number is bounded by $N\#N\#N \cdots \#N$, $b+1$ many times.
1136 Set $\mathcal{C}_b^0 = \mathcal{C}_b$. We will proceed in rounds; in round i , we set

$$\begin{aligned} b_i &= \arg \min_{b \in \{0,1\}} |\{C \in \mathcal{C}_k^{i-1} \mid \text{TT}(C)_i = b\}| \\ \mathcal{C}_k^i &= \{C \in \mathcal{C}_k^{i-1} \mid \text{TT}(C)_i = b_i\}. \end{aligned}$$

1137 Clearly, $|\mathcal{C}_b^i| \leq |\mathcal{C}_b^{i-1}|/2$, hence halving our search space. After $r = |\mathcal{C}_b|$ rounds, we will have 0 re-
1138 maining circuits matching the truth table prefix $b_1 b_2 \dots b_r$. If 0 circuits remain, then we have the answer
1139 $b_1 b_2 \dots b_r 0^{N-r}$. As N is feasible in the translation $\text{Hard}(b)$, computing b_i each round is a feasible minimiza-
1140 tion in $\text{VPV}_{\#}$. This step is feasible in $\text{VPV}_{\#}$, but *not* in $\mathbb{V}_{\#}^0$. Finally, the rounds can be expressed as a Σ_0^B
1141 induction over the bits of the resulting truth table, where one existentially guesses a *number* $b \in \{0,1\}^i$ in
1142 round $i < |N\#N\#N \cdots \#N|_1$ which is the least popular truth table prefix generated by size n^b circuits.

1143 5.5 $\mathbb{V}_{\#}^0 \vdash \text{Hard}(b)$

1144 We give a proof sketch that, in fact, $\mathbb{V}_{\#}^0 \vdash \text{Hard}(b)$, for every $b \in \mathbb{N}$. We need the following folklore theorem,

1145 **Theorem 5.11.** [17] \mathbb{V}^0 proves the soundness of bounded depth Frege

1146 **Corollary 5.12.** $\mathbb{V}_{\#}^0$ proves the soundness of quasipolynomial size bounded depth Frege.

1147 **Lemma 5.13.** $\mathbb{V}_{\#}^0 \vdash \text{Hard}(b)$, for every $b \in \mathbb{N}$.

1148 *Proof Sketch.* It is known that depth-(0.5) Frege has quasipolynomial sized propositional proofs of the
1149 dWPHP [39]. Further, these proofs are highly uniform, in the sense that you can give a direct connection
1150 language for the proofs, as you would for circuits, and this language would be polylog-uniform $\text{AC}^0[\mathbf{poly}]$.
1151 A $\mathbb{V}_{\#}^0$ proof of $\text{Hard}(b)$ would simply verify the soundness of the bounded depth Frege proof of dWPHP,
1152 substituting each propositional variable $x_{C,T}$ with the assertion that the truth table of the circuit C of size
1153 n^b is T . \square

1154 This proof method is likely not the easiest method; a more straightforward way would be directly taking
1155 the bounded arithmetic proof of $\text{dWPHP}(\Sigma_1^b(\alpha))$ in Buss’s theory $T_2^2(\alpha)$, and reformulating it as a proof in
1156 V^0 , replacing the uninterpreted oracle symbol α with V^0 function symbols.

1157 We then get as a corollary the surprising consequence,

1158 **Corollary 5.14.** If Hypothesis 5.9 is true, then $P \neq NP$.

1159 6 Acknowledgements

1160 Marco Carmosino was. Stefan Grosser was supported by the NSERC PGS D Scholarship. The authors
1161 thank Robert Robere, Sam Buss, Erfan Khaniki, and Jan Pich for helpful discussions. The authors also
1162 thank Robert Robere, Sam Buss, and Noel Arteché for suggesting improvements to an earlier draft of the
1163 paper. As well, the authors thank Erfan Khaniki for pointing out the necessity for bounding Skolem terms
1164 in our WHUPs.

1165 Part of this work was completed while the authors attended the Simons Institute “Meta Complexity”
1166 program reunion in April 2024.

1167 References

- 1168 [1] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. Primes is in p . *Annals of mathematics*, pages
1169 781–793, 2004.
- 1170 [2] Eric Allender, Michal Koucký, Detlef Ronneburger, and Sambuddha Roy. The pervasive reach of
1171 resource-bounded kolmogorov complexity in computational complexity theory. *J. Comput. Syst. Sci.*,
1172 77(1):14–40, 2011.
- 1173 [3] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University
1174 Press, 2009.
- 1175 [4] Samuel R. Buss. *Bounded arithmetic*. Bibliopolis, 1986.
- 1176 [5] Samuel R. Buss. The witness function method and provably recursive functions of peano arithmetic.
1177 In Dag Prawitz, Brian Skyrms, and Dag Westerståhl, editors, *Logic, Methodology and Philosophy of*
1178 *Science IX*, volume 134 of *Studies in Logic and the Foundations of Mathematics*, pages 29–68. Elsevier,
1179 1995.
- 1180 [6] Samuel R Buss. *Handbook of proof theory*. Elsevier, 1998.
- 1181 [7] Samuel R Buss. Bounded arithmetic and constant depth frege proofs. *Complexity of computations and*
1182 *proofs*, 13:153–174, 2004.
- 1183 [8] Marco Carmosino, Valentine Kabanets, Antonina Kolokolova, and Igor C Oliveira. Learn-uniform
1184 circuit lower bounds and provability in bounded arithmetic. In *2021 IEEE 62nd Annual Symposium on*
1185 *Foundations of Computer Science (FOCS)*, pages 770–780. IEEE, 2022.
- 1186 [9] Marco Carmosino, Valentine Kabanets, Antonina Kolokolova, Igor C Oliveira, and Dimitrios Tsintsili-
1187 das. Provability of the circuit size hierarchy and its consequences. 2024.
- 1188 [10] Lijie Chen, Shuichi Hirahara, and Hanlin Ren. Symmetric exponential time requires near-maximum
1189 circuit size. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1990–
1190 1999, 2024.
- 1191 [11] Lijie Chen, Ce Jin, Rahul Santhanam, and R Ryan Williams. Constructive separations and their
1192 consequences. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*,
1193 pages 646–657. IEEE, 2022.
- 1194 [12] Lijie Chen, Jiayu Li, and Igor Carboni Oliveira. Reverse mathematics of complexity lower bounds, Apr
1195 2024.

- 1196 [13] Lijie Chen, Xin Lyu, and R Ryan Williams. Almost-everywhere circuit lower bounds from non-trivial
1197 derandomization. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*,
1198 pages 1–12. IEEE, 2020.
- 1199 [14] Lijie Chen, Roei Tell, and Ryan Williams. Derandomization vs refutation: A unified framework for
1200 characterizing derandomization. In *2023 IEEE 64th Annual Symposium on Foundations of Computer
1201 Science (FOCS)*, pages 1008–1047. IEEE, 2023.
- 1202 [15] Alan Cobham. The intrinsic computational difficulty of functions. 1965.
- 1203 [16] Stephen Cook and Jan Krajíček. Consequences of the provability of np in p/poly. *The Journal of
1204 Symbolic Logic*, 72(4):1353–1371, 2007.
- 1205 [17] Stephen Cook and Phuong Nguyen. *Logical foundations of proof complexity*, volume 11. Cambridge
1206 University Press Cambridge, 2010.
- 1207 [18] Stephen A Cook. Feasibly constructive proofs and the propositional calculus (preliminary version). In
1208 *Logic, Automata, and Computational Complexity: The Works of Stephen A. Cook*, pages 193–218. 2023.
- 1209 [19] Lance Fortnow and Rahul Santhanam. New non-uniform lower bounds for uniform classes. In *31st Con-
1210 ference on Computational Complexity (CCC 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik,
1211 2016.
- 1212 [20] Harvey Friedman. One hundred and two problems in mathematical logic. *J. Symb. Log.*, 40(2):113–129,
1213 1975.
- 1214 [21] Shuichi Hirahara. Unexpected hardness results for kolmogorov complexity under uniform reductions.
1215 In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy,
1216 editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC
1217 2020, Chicago, IL, USA, June 22–26, 2020*, pages 1038–1051. ACM, 2020.
- 1218 [22] Pavel Hrubes. Theories very close to PA where kreisel’s conjecture is false. *J. Symb. Log.*, 72(1):123–137,
1219 2007.
- 1220 [23] Pavel Hrubes. Kreisel’s conjecture with minimality principle. *J. Symb. Log.*, 74(3):976–988, 2009.
- 1221 [24] Rahul Ilango, Jiayu Li, and R Ryan Williams. Indistinguishability obfuscation, range avoidance, and
1222 bounded arithmetic. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*,
1223 pages 1076–1089, 2023.
- 1224 [25] Emil Jeřábek. Dual weak pigeonhole principle, boolean complexity, and derandomization. *Annals of
1225 Pure and Applied Logic*, 129(1-3):1–37, 2004.
- 1226 [26] Valentine Kabanets. Easiness assumptions and hardness tests: Trading time for zero error. *Journal of
1227 Computer and System Sciences*, 63(2):236–252, 2001.
- 1228 [27] Oliver Korten. The hardest explicit construction. In *2021 IEEE 62nd Annual Symposium on Founda-
1229 tions of Computer Science (FOCS)*, pages 433–444. IEEE, 2022.
- 1230 [28] Jan Krajíček. No counter-example interpretation and interactive computation. In *Logic from Computer
1231 Science: Proceedings of a Workshop held November 13–17, 1989*, pages 287–293. Springer, 1992.
- 1232 [29] Jan Krajíček. Small circuits and dual weak php in the universal theory of p-time algorithms. *ACM
1233 Transactions on Computational Logic (TOCL)*, 22(2):1–4, 2021.
- 1234 [30] Jan Krajíček. On the existence of strong proof complexity generators. *Bulletin of Symbolic Logic*,
1235 30(1):20–40, 2024.
- 1236 [31] Jan Krajíček and Igor C. Oliveira. Unprovability of circuit upper bounds in cook’s theory PV. *Log.
1237 Methods Comput. Sci.*, 13(1), 2017.

- 1238 [32] Jan Krajíček and Pavel Pudlák. The number of proof lines and the size of proofs in first order logic.
1239 *Arch. Math. Log.*, 27(1):69–84, 1988.
- 1240 [33] Jan Krajíček, Pavel Pudlák, and Gaisi Takeuti. Bounded arithmetic and the polynomial hierarchy.
1241 *Annals of pure and applied logic*, 52(1-2), 1991.
- 1242 [34] Ming Li and Paul M. B. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications, 4th*
1243 *Edition*. Texts in Computer Science. Springer, 2019.
- 1244 [35] Zeyong Li. Symmetric exponential time requires near-maximum circuit size: Simplified, truly uniform.
1245 In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 2000–2007, 2024.
- 1246 [36] Yanyi Liu and Rafael Pass. On one-way functions and kolmogorov complexity. In Sandy Irani, editor,
1247 *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA,*
1248 *November 16-19, 2020*, pages 1243–1254. IEEE, 2020.
- 1249 [37] Yanyi Liu and Rafael Pass. A direct PRF construction from kolmogorov complexity. In Marc Joye and
1250 Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International*
1251 *Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-*
1252 *30, 2024, Proceedings, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 375–406.
1253 Springer, 2024.
- 1254 [38] Wolfgang Maass. Quadratic lower bounds for deterministic and nondeterministic one-tape turing ma-
1255 chines. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 401–408,
1256 1984.
- 1257 [39] Alexis Maciel, Toniann Pitassi, and Alan R Woods. A new proof of the weak pigeonhole principle. In
1258 *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 368–377,
1259 2000.
- 1260 [40] Tohru Miyatake. On the length of proofs in formal systems. *Tsukuba Journal of Mathematics*, 4(1):115–
1261 125, 1980.
- 1262 [41] Moritz Müller and Ján Pich. Feasibly constructive proofs of succinct weak circuit lower bounds. *Annals*
1263 *of Pure and Applied Logic*, 171(2):102735, 2020.
- 1264 [42] Cody Murray and Ryan Williams. Circuit lower bounds for nondeterministic quasi-polytime: an easy
1265 witness lemma for np and nqp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory*
1266 *of Computing*, pages 890–901, 2018.
- 1267 [43] Igor C Oliveira. Meta-mathematics of computational complexity theory. *SIGACT News Complexity*
1268 *Theory Column (DRAFT)*.
- 1269 [44] Rohit Parikh. Existence and feasibility in arithmetic. *The journal of symbolic logic*, 36(3):494–508,
1270 1971.
- 1271 [45] Rohit Parikh. Some results on the length of proofs. *Transactions of The American Mathematical Society*
1272 *- TRANS AMER MATH SOC*, 177:29–29, 03 1973.
- 1273 [46] Ján Pich. Circuit lower bounds in bounded arithmetics. *Annals of Pure and Applied Logic*, 166(1):29–45,
1274 2015.
- 1275 [47] Ján Pich. Logical strength of complexity theory and a formalization of the PCP theorem in bounded
1276 arithmetic. *Log. Methods Comput. Sci.*, 11(2), 2015.
- 1277 [48] Pavel Pudlák. Chapter viii - the lengths of proofs. In Samuel R. Buss, editor, *Handbook of Proof Theory*,
1278 volume 137 of *Studies in Logic and the Foundations of Mathematics*, pages 547–637. Elsevier, 1998.
- 1279 [49] Hanlin Ren, Rahul Santhanam, and Zhikun Wang. On the range avoidance problem for circuits. In
1280 *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 640–650.
1281 IEEE, 2022.

- 1282 [50] Neil Thapen. *The weak pigeonhole principle in models of bounded arithmetic*. PhD thesis, University of
1283 Oxford, 2002.
- 1284 [51] Ryan Williams. Nonuniform acc circuit lower bounds. *Journal of the ACM (JACM)*, 61(1):1–32, 2014.