Exact Learning Algorithms, Betting Games, and Circuit Lower Bounds

Ryan C. Harkins and John M. Hitchcock*

Abstract

This paper extends and improves work of Fortnow and Klivans [6], who showed that if a circuit class $\mathcal{C}$ has an efficient learning algorithm in Angluin's model of exact learning via equivalence and membership queries [2], then we have the lower bound $\mathrm{EXP}^{\mathrm{NP}} \not\subseteq \mathcal{C}$. We use entirely different techniques involving betting games [5] to remove the NP oracle and improve the lower bound to $\mathrm{EXP} \not\subseteq \mathcal{C}$. This shows that it is even more difficult to design a learning algorithm for $\mathcal{C}$ than the results of Fortnow and Klivans indicated. We also investigate the connection between betting games and natural proofs, and as a corollary the existence of strong pseudorandom generators.

1 Introduction

We continue the line of research basing hardness of learning results on computational complexity and cryptography (see for example [20, 9, 1, 10]). Fortnow and Klivans [6] consider Angluin's model of exact learning from equivalence and membership queries [2]. In an equivalence query, the learner presents a hypothesis and asks if it is equivalent to the unknown target concept. If the hypothesis is equivalent, the learner has succeeded; otherwise, the learner is given an example on which the hypothesis is incorrect. In a membership query, the learner chooses an example and asks the value of the target concept on that example. To succeed, the learner must exactly identify the target concept.

Fortnow and Klivans show that learning algorithms for a circuit class give lower bounds for that type of circuit in the class $\mathrm{EXP}^{\mathrm{NP}}$. Throughout this introduction, $\mathcal{C}$ denotes any nonuniform class of polynomial-size circuits. In the following theorem (as well as our results), it is most interesting to consider for $\mathcal{C}$ a class of circuits where no learning algorithm is yet known, nor is there is a known obstacle such as cryptographic hardness to learning. An example is polynomial-size depth-two threshold circuits, which corresponds to the complexity class $\mathcal{C} = \mathrm{TC}_0[2]$.

Theorem 1.1. (Fortnow and Klivans [6]) *If there is an efficient exact learning algorithm for $\mathcal{C}$ using equivalence and membership queries, then* $\mathrm{EXP}^{\mathrm{NP}} \not\subseteq \mathcal{C}$.

While "efficient" typically means a polynomial-time algorithm, the result of Fortnow and Klivans as well as our results allow for exponential time and subexponentially many queries (i.e. $(2^{O(n)})$ time and $2^{n^{o(1)}}$ queries).

*Department of Computer Science, University of Wyoming. This research was supported in part by NSF grants 0652601 and 0917417 and by an NWO travel grant. Part of this research was done while Hitchcock was on sabbatical at CWI.

For the case of exact learning algorithms that only make equivalence queries, a stronger lower bound follows immediately from results [11, 7] connecting resource-bounded measure and dimension with Littlestone's model of online mistake-bound learning [12]. Combining these results with the fact that learnability in the exact learning model with equivalence queries implies learnability in the online mistake-bound learning model, we have the following, which was also noted in [6].

Theorem 1.2. (Hitchcock [7, 6]) *If there is an efficient exact learning algorithm for $\mathcal{C}$ using equivalence queries, then* $\mathrm{EXP} \not\subseteq \mathcal{C}$.

Given these two theorems, it is natural to ask whether we can prove a lower bound in EXP assuming the learning algorithm makes both equivalence and membership queries. Where does the NP oracle come from in Theorem 1.1? The proof of Theorem 1.1 separates into two parts, giving an indirect diagonalization. Assume that $\mathrm{EXP}^{\mathrm{NP}} \subseteq \mathcal{C}$.

(i) Because $\mathcal{C} \subseteq \mathrm{P/poly}$, $\mathrm{EXP}^{\mathrm{NP}}$ collapses and reduces to the Permanent [4, 8, 19].

(ii) Use the exact learning algorithm to learn the $\mathcal{C}$ circuits for the Permanent. An NP oracle is used to answer the equivalence queries. This yields a P^{NP} algorithm for the Permanent.

Combining (i) and (ii) gives $\mathrm{EXP}^{\mathrm{NP}} \subseteq \mathrm{P}^{\mathrm{NP}}$, a contradiction. Therefore we see that the NP oracle in Theorem 1.1 is for the equivalence queries. In contrast, no NP oracle is required in Theorem 1.2 where equivalence queries are allowed. The proof of Theorem 1.2 relies on martingales and a more direct measure-theoretic diagonalization, which is very different than the double-collapse argument and indirect diagonalization used for Theorem 1.1. This suggests hope that a more direct diagonalization approach may yield the desired improvement.

To improve Theorems 1.1 and 1.2, we must simulate the learning algorithm's queries while performing our diagonalization. Following [7], consider implementing this in the form of a martingale. The equivalence queries are no problem. We simply use the transformation in [12] that converts an equivalence query algorithm to a mistake-bound learning algorithm and apply the technique from [7]. We can therefore focus our effort on the membership queries. Unfortunately, we have been unable to discover a method for simulating membership queries in a martingale, due to the stringent requirement that a martingale must bet on all strings in lexicographic order. However, there is an extension of martingales called betting games that do help.

1.1 Betting Games

Buhrman et al. [5] introduced betting games to define a generalization of resource-bounded measure with applications to autoreducibility and the BPP vs. EXP problem. Betting games and martingales are similar; the difference is how the betting strategy is allowed to operate. A martingale is required to consider the strings at each length in lexicographic order. Betting games lift this requirement by allowing the betting strategy to pick the next string that it will bet on. We can easily simulate a membership query in a betting game – the strategy simply asks to make a prediction on the queried string and then it gets to see the answer for that string. A betting game succeeds on a class $\mathcal{C}$ if it able to attain unbounded amounts of capital betting on languages in $\mathcal{C}$.

Theorem 1.3. *If there is an exact learning algorithm for $\mathcal{C}$ using equivalence and membership queries, then there is a betting game which succeeds on $\mathcal{C}$.*

Buhrman et al. showed that it is also possible to diagonalize within EXP against betting games [5], just as is the case for martingales [13]. Formally, no betting game can succeed on all of EXP. Hence the desired improvement, our main result, follows from Theorem 1.3:

Theorem 1.4. *If there is an exact learning algorithm for $\mathcal{C}$ using equivalence and membership queries, then* $\mathrm{EXP} \not\subseteq \mathcal{C}$.

These results show that designing an exact learning algorithm for many circuit classes $\mathcal{C}$ will be difficult, as for most classes of interest it is an open problem whether $\mathrm{EXP} \subseteq \mathcal{C}$.

1.2 Natural Proofs

Razborov and Rudich [17] introduced natural proofs in order to explain the difficulty of proving circuit lower bounds. Regan, Sivakumar, and Cai [18] connected natural proofs to martingales, showing that if there is a martingale that succeeds on P/poly, then there is a natural proof against P/poly. It is a basic question to ask whether this extends to betting games as well. We show that this extends to "honest" betting games:

Theorem 1.5. *If there is an honest betting game which succeeds on* P/poly*, then there is a natural proof against* P/poly.

The betting game in the proof of Theorem 1.3 is honest. Combining Theorem 1.5 with Theorem 1.3, we have the following corollary:

Corollary 1.6. *If there is an efficient exact learning algorithm for* P/poly *using equivalence and membership queries, then there is a natural proof against* P/poly.

Razborov and Rudich showed that a natural proof against P/poly implies that cryptographically secure pseudorandom generators do not exist. Therefore, we have the following:

Corollary 1.7. *If strong pseudorandom generators exist, then there is no honest betting game that succeeds on* P/poly.

It is also interesting to compare this sequence of implications to known hardness of learning results based on cryptographic assumptions (e.g. [9]). We find the connection between exact learning of Boolean circuits and the existence of strong pseudorandom generators via betting games and natural proofs to be elegant.

1.3 Organization

This paper is organized as follows. Precise technical definitions for betting games, exact learning, and natural proofs are given in section 2. We construct betting games from exact learning algorithms in section 3. The connections with natural proofs and pseudorandom generators are presented in section 4.

2 Preliminaries

We use standard notation. The binary alphabet is $\Sigma = \{0, 1\}$, the set of all binary strings is Σ^*, the set of all binary strings of length n is Σ^n, and the set of all infinite binary sequences is Σ^∞.

The empty string is denoted by λ. We use the standard enumeration of strings, $\lambda, 0, 1, 00, 01, \ldots$, and a total ordering of strings corresponding to this enumeration. A language A can alternatively be seen as a subset of Σ^*, or as an element of Σ^∞ via identification with its characteristic sequence.

2.1 Betting Games

Betting games, which are also called nonmonotonic martingales, originated in the field of algorithmic information theory. In that setting they yield the notion of Kolmogorov-Loveland randomness (generalizing Kolmogorov-Loveland stochasticity) [16, 15]. The concept was introduced to computational complexity by Buhrman et al. [5]. Our notation for betting games is taken predominantly from Merkle et al. [15]. First, for comparison, we recall the definition of a martingale:

Definition. A *martingale* is a function $d : \Sigma^* \to [0, \infty)$ such that for all $w \in \Sigma^*$, we have the following averaging condition:

$$d(w) = \frac{d(w0) + d(w1)}{2}.$$

Intuitively, a martingale is betting in order on the characteristic sequence of an unknown language. The martingale starts with finite initial capital $d(\lambda)$. The quantity $d(w)$ represents the current capital the martingale has after betting on the first $|w|$ bits of a sequence that begins with w. The quantities $\pi(w, 0) = d(w0)/2d(w)$ and $\pi(w, 1) = d(w1)/2d(w)$ represent the fraction of its current capital that the martingale is waging on 0 and 1, respectively, being the next bit of the sequence. This next bit is revealed and the martingale has $d(w0) = 2\pi(w, 0)d(w)$ in the case of a 0 and $d(w1) = 2\pi(w, 1)d(w)$ in the case of a 1.

Betting games are similar to martingales but have an additional capability of selecting which position in a sequence, or equivalently, which string in a language, to bet upon next. A betting game is permitted to select strings in a nonmonotone order, that is, it may bet on longer strings, then shorter strings, then longer strings again (with the important restriction that it may not bet on the same string twice). Because of this added complexity, it is simpler to break the description of a betting game into pieces.

Definition. A *betting game* is a system that bets nonmonotonically on an infinite sequence (or equivalently a language) and formally is a triple $G = (s, \pi, V)$ consisting of a scan rule s, a stake function π, and a capital function V.

1. A *finite assignment* is a sequence $w \in (\Sigma^* \times \Sigma)^*$. In essence, it is a list of strings examined thus far, each string coupled with an assignment, saying whether or not it is included in the language being bet upon. The set of all finite assignments is denoted FA. We use the $\cdot$ symbol for concatenation of finite assignments.

2. The *scan rule* is a (partial) computable function $s : FA \to \Sigma^*$ from finite assignments to strings that looks at a finite assignment and determines the next string (or bit of a sequence) to bet on. The scan rule is limited in that it cannot select a string that already appears in the current finite assignment.

3. The *stake function* is a partial function $\pi : FA \times \Sigma \to [0, 1]$. Its function is to examine the current finite assignment and determine what fraction of the capital to bet on either side. It carries a condition that $\pi(w, 0) + \pi(w, 1) = 1$ for all $w \in FA$.

4. The *capital function* is a partial function $V : FA \to [0, \infty)$ from finite assignments to nonnegative reals, and utilizes the stake function π. Its initial capital $V(\lambda)$ is finite. For $w \in FA$, when $V(w)$ is defined, the scan rule $s(w)$ determines the next string to bet on, and $\pi(w, b)$ is the stake amount, the capital is updated according to the rule

$$V(w \cdot (s(w), b)) = 2\pi(w, b)V(w). \tag{2.1}$$

A betting game's capital function also satisfies an averaging condition, in analogy with the definition of a martingale:

$$V(w) = \frac{V(w \cdot (s(w), 0)) + V(w \cdot (s(w), 1))}{2}.$$

Note that a betting game is a martingale when the scan rule always selects the next string in the lexicographic order.

The play of the betting game on a language A is defined by recurrence $w_0^A = \lambda$ and $w_{n+1}^A = w_n^A \cdot (s(w_n), A(s(w_n)))$ for $n \geq 0$, which may be a finite or infinite sequence in general (depending on the definition of the scan rule). For the definition of success we are concerned with the sequence of capital values $V_n^A = V(w_n^A)$.

Definition. If a betting game G earns unbounded capital on a language A (in the sense that for every constant c there is a point at which the capital V_n^A exceeds c when betting on A), we say that G *succeeds on* A. The *success set* of a betting game G, denoted $S^\infty[G]$, is the set of all languages on which G succeeds. A betting game G *succeeds on* a class X of languages if $X \subseteq S^\infty[G]$.

The ability of the betting game to examine a sequence nonmonotonically makes determining its running time complicated, since each language can induce a unique computation of the betting game. In other words, the betting game may choose to examine strings in different orders depending upon the language it is wagering against. Buhrman et al. looked at a betting game as an infinite process on a language, rather than a finite process on a string. They used the following definition:

Definition. A betting game G runs in time $t(n)$ if for all languages A, every query of length n made by G occurs in the first $t(n)$ steps of the computation.

Specifically, once a $t(n)$-time-bounded betting game uses $t(n)$ computational steps, its scan rule cannot go back and select any string of length n. We remark that in the results of this paper all betting games have the special form that they bet on all strings of each length before moving on to the next length, so the technical issue of measuring the run time is not important for this paper. In any case, the crucial result is that exponential-time betting games cannot succeed on the classes $\mathrm{E} = \mathrm{DTIME}(2^{O(n)})$ and $\mathrm{EXP} = \mathrm{DTIME}(2^{n^{O(1)}})$.

Theorem 2.1. (Buhrman et al. [5])

1. *No* $2^{O(n)}$*-time betting game succeeds on* E.
2. *No* $2^{n^{O(1)}}$*-time betting game succeeds on* EXP.

2.2 Exact Learning

In general, a learning algorithm seeks to identify an unknown concept from some known class of concepts. We now review the basic notation and definitions for the exact learning model.

A *concept* is a Boolean function $c_n : \Sigma^n \to \Sigma$. For any string $x \in \Sigma^n$, if $c_n(x) = 1$, then x is positively classified as belonging to the concept, while if $c_n(x) = 0$, then x is classified as not belonging to the concept. A string x paired with the classification $c_n(x)$ is called an *example*. A concept c_n is often identified with the set of positive examples $\{x \mid c_n(x) = 1\} \subseteq \Sigma^n$. A *concept class* $\mathcal{C}_n$ is a set of concepts over Σ^n. A *concept class family* is a sequence $\mathcal{C} = \{\mathcal{C}_n\}_{n \geq 0}$ of concept classes.

A learning algorithm tries to identify a *target concept* drawn from $\mathcal{C}_n$, and often does this by forming a *hypothesis*, which is typically some concept in $\mathcal{C}_n$ that is consistent with (i.e. classifies correctly) all the examples seen thus far. In the exact learning paradigm, a learner $\mathcal{A}$ may make various sorts of queries to a teacher, and then, depending on the answers, formulate a hypothesis. This process repeats until $\mathcal{A}$ has successfully discovered the target concept. We will focus on two types of queries: equivalence queries and membership queries.

Definition. An *equivalence query* is a request to the teacher to know if the current hypothesis matches the target concept. If the answer is yes, the teacher responds accordingly. If the answer is no, then the teacher provides the learner with a counterexample (an example that is incorrectly classified by the current hypothesis).

Definition. A *membership query* is a request to the teacher to know the classification of a specific string x. The teacher responds with $c_n(x)$, where c_n is the target concept.

2.3 Natural Proofs and Pseudorandom Generators

The natural proofs of Razborov and Rudich are combinatorial properties which can be shown to be "useful," i.e. diagonalize against, certain specific classes. Formally, let F_n be the set of all n-variable Boolean functions. A combinatorial property is a sequence $\Pi = \{\Pi_n\}_{n \geq 0}$ where each Π_n is a subset of F_n. For classes $\mathcal{C}$ and $\mathcal{D}$, Π is $\mathcal{C}$-natural against $\mathcal{D}$ if it satisfies the following conditions:

- Constructivity: the decision problem $f_n \in \Pi_n$ belongs $\mathcal{C}$.
- Largeness: $|\Pi_n| \geq 2^{-O(n)} \cdot |F_n|$ for all n.
- Usefulness: for any $A \in \mathcal{D}$, for infinitely many n, $A_{=n} \notin \Pi_n$.

For more information on natural proofs, we refer the reader to [17, 18, 14]. We now recall the connection between natural proofs and pseudorandom generators. A language A belongs to the class P/poly if there is a polynomial $q(n)$ such that for all n, $A_{=n}$ has a Boolean circuit of size at most $q(n)$. Similarly, A belongs to the class QP/qpoly if the circuits have size bounded by $2^{(\log n)^{O(1)}}$.

Theorem 2.2. (Razborov and Rudich [17]) *If there exists a combinatorial property that is* QP/qpoly*-natural against* P/poly*, then pseudorandom generators of exponential hardness against nonuniform adversaries do not exist.*

We refer to [17] for precise definitions about pseudorandom generators.

3 Exact Learning and Betting Games

Theorem 3.1. *Let $\mathcal{C} = \{\mathcal{C}_n \mid n \in \mathbb{N}\}$ be a concept class family, and let*

$$X = \{A \mid (\exists^\infty n) A_{=n} \in \mathcal{C}_n\}.$$

1. *If there is an exact learning algorithm for $\mathcal{C}$ that learns each $\mathcal{C}_n$ in time 2^{cn} and makes no more than 2^{n-2} equivalence and membership queries, then there exists a betting game G that runs in time $O(2^{(c+2)n})$, such that $X \subseteq S^\infty[G]$.*

2. *If there is an exact learning algorithm for $\mathcal{C}$ that learns each $\mathcal{C}_n$ in time 2^{n^c} and makes no more than 2^{n-2} equivalence and membership queries, then there exists a betting game G that runs in time $O(2^{n^{c+2}})$, such that $X \subseteq S^\infty[G]$.*

Proof. We prove the first item – the proof of the second item is analogous.

Let $\mathcal{A}$ be the learning algorithm that learns concepts in $\mathcal{C}$. In other words, for each $n \in \mathbb{N}$, and for any target concept $c_n \in \mathcal{C}_n$, $\mathcal{A}$ can learn c_n using no more than 2^{cn} time, and making at most 2^{n-2} equivalence and membership queries.

Let $G(s, \pi, V)$ be as follows. G effectively runs in stages, examining strings by length, betting on all strings of length n before betting on any string of size $n+1$. This is proper, since $\mathcal{A}$ learns concepts taken from $\mathcal{C}_n$, whose concepts only classify strings of length n. Therefore, we will apply two indices to the stages of calculation, the first to indicate string length, and the second to indicate how many stages have been executed at the string length.

G starts with capital $V_{0,0} = 2$, but it treats its capital as divided up into an infinite number of amounts, 2^{-n} for each n. Thus at each stage $(n, 0)$, the capital effectively is $V_{n,0} = 2^{-n}$ (with all previous winnings "banked" and untouchable). To reflect this, we will divide π in a class of functions $\{\pi_n\}_{n \geq 0}$, so that π_n only touches the capital $V_{n,i}$ for $0 \leq i \leq 2^n$.

At each stage (n, i), the scan rule s runs the learning algorithm $\mathcal{A}$, which updates its current hypothesis $h_{n,i}$. In the process of formulating $h_{n,i+1}$, one of two things will happen:

- $\mathcal{A}$ will make a membership query to some string $x \in \Sigma^n$
- $\mathcal{A}$ will make an equivalence query

If $\mathcal{A}$ makes a membership query to x, s then checks to see if x occurs in w, the current finite assignment. If so, then s answers the query according to the label. If not, then we set $s(w) = x$ and $\pi_n(w, 0) = \pi_n(w, 1) = 1/2$. In other words, the betting game selects x and makes no bet on the outcome. Once the label for x is revealed, the finite assignment is updated ($w = w \cdot (x, b)$, where b is the label for x). The scan rule then provides the correct classification of x to $\mathcal{A}$ as the answer to the membership query. The betting game's capital is unchanged and the computation then proceeds onto stage $(n, i+1)$.

If $\mathcal{A}$ makes an equivalence query, then s proceeds in an online fashion. First, the scan rule selects $s(w) = x$, where x is the lexicographically least string in Σ^n that does not appear in w. The stake function computes the prediction $b = h_{n,i}(x)$ using the current hypothesis $h_{n,i}$ and sets $\pi_n(w, b) = 3/4$ and $\pi_n(w, 1-b) = 1/4$. The true classification of x is revealed and $V_{n,i}$ is updated according to (2.1). If $c_n(x) \neq h_{n,i}(x)$, then $(x, 1-b)$ is presented to $\mathcal{A}$ as a counterexample, and computation proceeds to stage $(n, i+1)$. If $c_n(x) = h_{n,i}(x)$, then computation proceeds to stage $(n, i+1)$, and $\mathcal{A}$ is treated as still making an equivalence query. This process repeats until either a counterexample is discovered or the strings of size n are exhausted.

Without loss of generality, we will assume that $\mathcal{A}$ always finishes with an equivalence query to ensure correctness of its hypothesis. Thus if $\mathcal{A}$ has formed the correct hypothesis, G will continue searching for a counterexample until all strings of length n are exhausted, and G moves onto $(n+1, 0)$, where it utilizes $\mathcal{A}$ to learn a new concept.

To examine the running time of G, we note first that $\mathcal{A}$ updates its hypothesis in 2^{cn} time. The remaining time is allotted to the scan rule, which takes only time linear in the size of the current finite assignment (which has size at most 2^{n+1}) to determine a string to select, and to updating the capital function, which can be done in time $O(2^n)$. Hence the aggregate time for G to make all queries of size n is $O(2^n \cdot 2^{cn} \cdot 2^n)$. Therefore G is an $O(2^{(c+2)n})$-time betting game.

To see that G succeeds on X, it suffices to show that for infinitely many n, $V_{n,2^n} \geq 1$. We know that for any $A \in X$, there exist infinitely many n such that $A_{=n} \in \mathcal{C}_n$, and hence for infinitely many n, $\mathcal{A}$ will either correctly learn $A_{=n}$, or at least will classify correctly a sufficiently large number of strings in $A_{=n}$. Thus it suffices to show that for any sufficiently large n, $\mathcal{A}$ learns sufficiently quickly enough to bring in large earnings.

The worst case occurs if all 2^{n-2} queries are equivalence queries, to which a counterexample is ultimately found for each query. (If we allow for membership queries, the bet for each query is 0, and so the capital does not change regardless of the true label. The counterexample to the equivalence query, though, will decrease capital.) Since, by definition, each string of length n will be queried, we have:

$$V_{n,2^n} = V_{n,0} \cdot \left(\frac{1}{2}\right)^{2^{n-2}} \cdot \left(\frac{3}{2}\right)^{3 \cdot 2^{n-2}} = 2^{-n} \left(\frac{27}{16}\right)^{2^{n-2}} \geq 1$$

Hence for infinitely many n, G will "bank" one dollar, and therefore its earnings will increase unbounded. Therefore, $X \subseteq S^{\infty}[G]$. □

Our improvement to the result of Fortnow and Klivans now follows as an immediate corollary to Theorems 2.1 and 3.1:

Corollary 3.2. *Let $\mathcal{C} = \{\mathcal{C}_n \mid n \in \mathbb{N}\}$ be a concept class family, and let*

$$X = \{A \mid (\exists^{\infty} n) A_{=n} \in \mathcal{C}_n\}.$$

1. *If there is an exact learning algorithm for $\mathcal{C}$ that learns each $\mathcal{C}_n$ in time $2^{O(n)}$ and makes no more than 2^{n-2} equivalence and membership queries, then $\mathrm{E} \not\subseteq X$.*
2. *If there is an exact learning algorithm for $\mathcal{C}$ that learns each $\mathcal{C}_n$ in time $2^{n^{O(1)}}$ and makes no more than 2^{n-2} equivalence and membership queries, then $\mathrm{EXP} \not\subseteq X$.*

Corollary 3.3. *Let $\mathcal{C}$ be any subclass of* P/poly *that is exactly learnable in time $2^{O(n)}$, making at most 2^{n-2} equivalence and membership queries. Then $\mathrm{E} \not\subseteq \mathcal{C}$. In particular, if* P/poly *is learnable under these assumptions, then* $\mathrm{EXP} \not\subseteq \mathrm{P/poly}$.

4 Betting Games and Natural Proofs

A basic question to ask regarding betting games is if they, like martingales as shown in Regan, Sivakumar, and Cai [18], give rise to natural proofs, as defined by Razborov and Rudich [17]. While a general answer proves elusive, we can restrict our focus to "honest" betting games, defined analagously to the "honest" martingales in [18]:

Definition. An *honest* betting game $G = (s, \pi, V)$ is one in which, for each n, the scan rule s selects every string of length n before selecting any strings of greater length, and the computation to select a string of length n depends only on the portion of the finite assignment dealing with strings of length n.

In other words, an honest betting game is one that bets entirely within a string length before betting on the next string length, and the only history that matters is the history within the current string length. Our betting game in Theorem 3.1 is an honest betting game.

We will now provide definitions and lemmas much akin to those of [18] and show that we can construct a P-natural property that is useful against P/poly.

Notation. For an honest betting game $G = (s, \pi, V)$, let $V_n(w)$ denote the fractional capital accrued while betting on strings of length n, for a language whose characteristic sequence at length n is determined by the string $w \in \Sigma^{2^n}$, and let $V_{\mathrm{cur}}(n)$ be the current capital before betting on strings of length n.

In other words, for any language A, we have $V_{\mathrm{cur}}(0) = V(\lambda)$ and

$$V_{\mathrm{cur}}(n+1) = V_n(\chi_{A=n}) \cdot V_{\mathrm{cur}}(n).$$

Also, note that if $\limsup_{n\to\infty} \prod_{i=0}^{n} V_i(\chi_{A=n}) = \infty$, then $A \in S^\infty[G]$.

Lemma 4.1. *Let $G = (s, \pi, V)$ be an honest betting game and $n \in \mathbb{N}$. Then*

$$\left\| \left\{ v \in \Sigma^{2^n} \;\middle|\; V_n(v) \leq \left(1 + \frac{1}{n^2}\right) \right\} \right\| \geq 2^{2^n} \left(\frac{1}{n^2+1}\right).$$

Proof. This follows from the averaging condition of an honest betting game. Like a martingale, an honest betting game must average $V_{\mathrm{cur}}(n)$ across all possible $v \in \Sigma^{2^n}$. That is,

$$\sum_{v \in \Sigma^{2^n}} V_n(v) V_{\mathrm{cur}}(n) = V_{\mathrm{cur}}(n),$$

which implies

$$\sum_{v \in \Sigma^{2^n}} V_n(v) = 1.$$

An application of Markov's inequality yields the bound in the lemma. □

Lemma 4.2. *If an honest 2^{cn}-time-bounded betting game $G = (s, \pi, V)$ succeeds on* P/poly*, then for every polynomial q there exist infinitely many n such that for all circuits C_n of size at most $q(n)$,*

$$V_n(v) \geq \left(1 + \frac{1}{n^2}\right)$$

where $v \in \Sigma^{2^n}$ is the characteristic sequence that agrees with the circuit C_n.

Proof. Suppose not. There exists some N such that for all $n > N$, there exists some circuit C_n of size at most $q(n)$ such that $V_n(v) < (1 + 1/n^2)$.

We will construct a language L as follows: arbitrarily fix $L^{\leq N}$. Then for each $n > N$, let C_n be a circuit fitting the criteria above. Then fix $L_{=n}$ to be v, where v is the characteristic sequence that

agrees with C_n. Then it follows that $L \in \mathrm{P/poly}$ (we have a family of circuits $\{C_i\}_{i\geq 0}$ where each C_i has size at most $q(i)$), and that $L \notin S^\infty[G]$. This last fact is because $\prod_{n=1}^{\infty}(1+1/n^2)$ converges, and

$$\lim_{n\to\infty} V_{\mathrm{cur}}(n) = \lim_{n\to\infty} V(\lambda)\prod_{i=1}^{n} V_n(v) \leq V(\lambda)\lim_{n\to\infty}\prod_{i=1}^{n}\left(1+\frac{1}{n^2}\right) < \infty.$$

This provides the necessary contradiction. □

Theorem 4.3. *If there is an honest, 2^{cn}-time-bounded betting game $G = (s, \pi, V)$ that succeeds on* P/poly*, then there is a* P*-natural property against* P/poly*.*

Proof. For all n, we define

$$\Pi_n = \left\{ v \in \Sigma^{2^n} \;\middle|\; V_n(v) \leq \left(1+\frac{1}{n^2}\right) \right\}.$$

We note that we do not worry here about the potential of G having depleted all of its capital prior to betting on strings of length n because we only consider the fractional amount accrued when calculating V_n. It follows that computing $v \in \Pi_n$ is in P: it requires at most 2^{cn} time to determine $V_n(v)$ by simulating the computation of G, but $|v| = 2^n$, which means the time required is $|v|^c$.

It follows from Lemma 4.1 that Π_n is large: $|\Pi_n| \geq 1/(n^2+1)\cdot|F_n|$, which is significantly bigger than $2^{-O(n)}\cdot|F_n|$. Finally, usefulness follows from Lemma 4.2: for any $L \in \mathrm{P/poly}$, there exist infinitely many n such that $L_{=n} \notin \Pi_n$. If this were not so, then for all but finitely many n, $L_{=n} \in \Pi_n$, but the capital accrued on all these segments would be bounded, and thus $L \notin S^\infty[G]$. □

Combining Theorems 2.2 and 4.3, we have the following:

Corollary 4.4. *If there is an honest, 2^{cn}-time-bounded betting game that succeeds on* P/poly*, then pseudorandom generators of exponential hardness against nonuniform adversaries do not exist.*

We remark that in this section, we could also use $2^{n^{O(1)}}$-time betting games and learning algorithms, obtaining the same result through a QP-natural property.

We conclude with the following diagram summarizing the implications for P/poly:

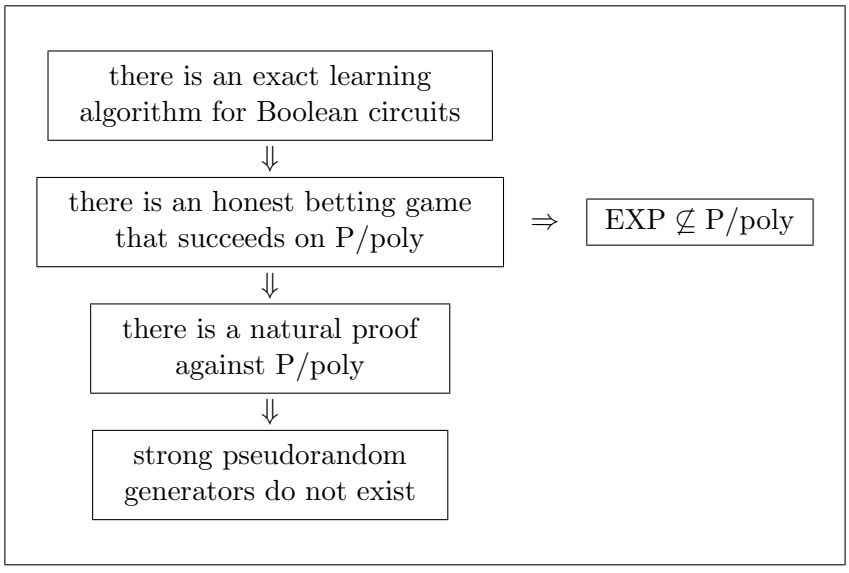

References

[1] H. Aizenstein, T. Hegeds, L. Hellerstein, and L. Pitt. Complexity theoretic hardness results for query learning. *Computational Complexity*, 7:19–53, 1998.

[2] D. Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988.

[3] N. H. Bshouty, R. Cleve, S. Kannan, R. Gavaldà, and C. Tamon. Oracles and queries that are sufficient for exact learning. *Journal of Computer and System Sciences*, 52(3):421–433, 1996.

[4] H. Buhrman and S. Homer. Superpolynomial circuits, almost sparse oracles and the exponential hierarchy. In *Proceedings of the 12th Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 116–127. Springer, 1992.

[5] H. Buhrman, D. van Melkebeek, K. W. Regan, D. Sivakumar, and M. Strauss. A generalization of resource-bounded measure, with application to the BPP vs. EXP problem. *SIAM Journal on Computing*, 30(2):576–601, 2001.

[6] L. Fortnow and A. R. Klivans. Efficient learning algorithms yield circuit lower bounds. *Journal of Computer and System Sciences*, 75(1):27–36, 2009.

[7] J. M. Hitchcock. Online learning and resource-bounded dimension: Winnow yields new lower bounds for hard sets. *SIAM Journal on Computing*, 36(6):1696–1708, 2007.

[8] R. M. Karp and R. J. Lipton. Some connections between nonuniform and uniform complexity classes. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing*, pages 302–309, 1980.

[9] M. Kearns and L. Valiant. Cryptographic limitations on learning Boolean formulae and finite automata. *Journal of the ACM*, 41(1):67–95, 1994.

[10] M. Kharitonov. Cryptographic lower bounds for learnability of Boolean functions on the uniform distribution. *Journal of Computer and System Sciences*, 50(3):600–610, 1995.

[11] W. Lindner, R. Schuler, and O. Watanabe. Resource-bounded measure and learnability. *Theory of Computing Systems*, 33(2):151–170, 2000.

[12] N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1988.

[13] J. H. Lutz. Almost everywhere high nonuniform complexity. *Journal of Computer and System Sciences*, 44(2):220–258, 1992.

[14] J. H. Lutz. The quantitative structure of exponential time. In L. A. Hemaspaandra and A. L. Selman, editors, *Complexity Theory Retrospective II*, pages 225–254. Springer-Verlag, 1997.

[15] W. Merkle, J. S. Miller, A. Nies, J. Reimann, and F. Stephan. Kolmogorov-Loveland randomness and stochasticity. *Annals of Pure and Applied Logic*, 138(1–3):183–210, 2006.

[16] A. A. Muchnik, A. L. Semenov, and V. A. Uspensky. Mathematical metaphysics of randomness,. *Theoretical Computer Science*, 207(2):263 – 317, 1998.

[17] A. A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1):24–35, 1997.

[18] K. W. Regan, D. Sivakumar, and J. Cai. Pseudorandom generators, measure theory, and natural proofs. In *Proceedings of the 36th Symposium on Foundations of Computer Science*, pages 26–35. IEEE Computer Society, 1995.

[19] S. Toda. On the computational power of PP and $\oplus$P. *SIAM Journal on Computing*, 20(5):865–877, 1991.

[20] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.