A Generalization of Resource-Bounded Measure, With an Application (Extended Abstract)

Harry Buhrman[1], Dieter van Melkebeek[2], Kenneth W. Regan[3], D. Sivakumar[4], and Martin Strauss[5]

[1] CWI, Kruislaan 413, 1098SJ Amsterdam, The Netherlands. E-mail: buhrman@cwi.nl. Partly supported by the Dutch foundation for scientific research (NWO), by SION project 612-34-002, and by the European Union through NeuroCOLT ESPRIT Working Group Nr. 8556 and HC&M grant nr. ERB4050PL93-0516.

[2] Univ. of Chicago, Department of Computer Science, 1100 E. 58th St., Chicago, IL 60637 USA. E-mail: dieter@cs.uchicago.edu. Party supported by the European Union through Marie Curie Research Training Grant ERB-4001-GT-96-0783 at CWI and by NSF Grant CCR 92-53582.

[3] Computer Science, University at Buffalo, 226 Bell Hall, Buffalo, NY 14260-2000 USA. E-mail: regan@cs.buffalo.edu. Partly supported by NSF Grant CCR-9409104.

[4] Department of Computer Science, University of Houston, Houston, TX 77204-3475, USA. E-mail: siva@cs.uh.edu. Partly supported at Buffalo by NSF Grant CCR-9409104.

[5] AT&T Labs, Room C216, 180 Park Ave, Florham Park, NJ 07932-0971 USA. E-mail: mstrauss@research.att.com. Research performed at Rutgers University and Iowa State University, supported by NSF grants CCR-9204874 and CCR-9157382.

Abstract. We introduce *resource-bounded betting games*, and propose a generalization of Lutz's resource-bounded measure in which the choice of next string to bet on is fully adaptive. Lutz's martingales are equivalent to betting games constrained to bet on strings in lexicographic order. We show that if strong pseudo-random number generators exist, then betting games are equivalent to martingales, for measure on E and EXP. However, we construct betting games that succeed on certain classes whose Lutz measures are important open problems: the class of polynomial-time Turing-complete languages in EXP, and its superclass of polynomial-time Turing-autoreducible languages. If an EXP-martingale succeeds on either of these classes, or if betting games have the "finite union property" possessed by Lutz's measure, one obtains the non-relativizable consequence BPP $\neq$ EXP. We also show that if EXP $\neq$ MA, then the polynomial-time truth-table-autoreducible languages have Lutz measure zero, whereas if EXP $=$ BPP, they have measure one.

1 Introduction

Lutz's theory of measure on complexity classes is now usually defined in terms of resource-bounded martingales. A martingale can be regarded as a gambling game played on unseen languages A. Let $s_1, s_2, s_3, \ldots$ be the standard lexicographic ordering of strings. The gambler G starts with capital $C_0 = \$1$ and places a bet $B_1 \in [0, C_0]$ on either "$s_1 \in A$" or "$s_1 \notin A$." Given a fixed particular language A, whether $s_1 \in A$ is used to resolve the bet. If the bet won, then the new capital C_1 equals $C_0 + B_1$; if the bet lost, then $C_1 = C_0 - B_1$. The gambler then places a bet $B_2 \in [0, C_1]$ on (or against) membership of the string s_2, then on s_3, and so forth. The gambler *succeeds* if G's capital C_i grows toward $+\infty$. The class $\mathcal{C}$ of languages A

on which G succeeds (and any subclass) is said to have *measure zero*. One also says G *covers* $\mathcal{C}$. Lutz and others (see [13]) have developed a rich and extensive theory around this measure-zero notion, and have shown interesting connections to many other important problems in complexity theory.

We propose the generalization obtained by lifting the requirement that G must bet on strings in lexicographic order. That is, G may begin by choosing any string x_1 on which to place its first bet, and after the oracle tells the result, may choose any other string x_2 for its second bet, and so forth. Note that the sequences $x_1, x_2, x_3, \ldots$ (as well as $B_1, B_2, B_3, \ldots$) may be radically different for different oracle languages A—in complexity-theory parlance, G's queries are *adaptive*. The lone restriction is that G may not query (or bet on) the same string twice. We call G a *betting game*.

Our betting games remedy a possible lack in the martingale theory, one best explained in the context of languages that are "random" for classes $\mathcal{D}$ such as E or EXP. A language L is *$\mathcal{D}$-random* if L cannot be covered by a $\mathcal{D}$-martingale. Based on one's intuition about random 0-1 sequences, the language $L' = \{ \mathit{flip}(x) : x \in L \}$ should likewise be $\mathcal{D}$-random, where $\mathit{flip}(x)$ changes every 0 in x to a 1 and vice-versa. However, this closure property is not known for E-random or EXP-random languages, because of the way martingales are tied to the fixed lex ordering of Σ^*. Betting games can adapt to easy permutations of Σ^* such as that induced by *flip*. Similarly, a class $\mathcal{C}$ that is *small* in the sense of being covered by a ($\mathcal{D}$-) betting game remains small if the languages $L \in \mathcal{C}$ are so permuted. In the r.e./recursive theory of random languages, our generalization is similar to "Kolmogorov-Loveland place-selection rules" (see [11]). We make this theory work for complexity classes via a novel definition of "running in time $t(n)$" for an infinite process.

We also provide a useful new angle on Lutz's theory, in which a major open question is whether the class of EXP-complete sets—under polynomial-time Turing reductions—has EXP-measure zero. If so (in fact if this set does not have measure one), then by results of Allender and Strauss [1], BPP $\neq$ EXP. Since there are oracles A relative to which $\mathrm{BPP}^A = \mathrm{EXP}^A$ [10], this kind of absolute separation would be a major breakthrough. We show that the EXP-complete sets *can* be covered by an EXP betting game—in fact, by an E-betting game. The one technical lack in our theory as a notion of measure is also interesting here: If the "finite unions" property holds for betting games (viz. $\mathcal{C}_1$ small $\wedge$ $\mathcal{C}_2$ small $\Longrightarrow \mathcal{C}_1 \cup \mathcal{C}_2$ small), then EXP $\neq$ BPP. Likewise, if Lutz's martingales do enjoy the permutation-invariance of betting games, then BPP $\neq$ EXP. Finally, we show that if a pseudorandom number generator (PRG) of security $2^{n^{\Omega(1)}}$ exists, then for every EXP-betting game G one can find an EXP-martingale that succeeds on all sets covered by G. PRGs of higher security $2^{\Omega(n)}$ likewise imply the equivalence of E-betting games and E-measure. Ambos-Spies and Lempp [4] proved that the EXP-complete sets have E-measure zero under a different hypothesis, namely P = PSPACE.

Measure theory and betting games help us to dig even deeper into questions about PRGs and complexity-class separations. Our pivot is the notion of an *autoreducible* set, whose importance in complexity theory was argued by Buhrman, Fortnow, and Torenvliet [7]. A language L is *$\leq_T^p$-autoreducible* if there is a polynomial-time oracle TM Q such that for all inputs x, Q^L correctly decides whether $x \in L$ without ever submitting x itself as a query to L. If Q is non-adaptive (i.e., computes a polynomial-time truth-table reduction), we say L is *$\leq_{tt}^p$-autoreducible*. We show that the class of

$\leq_T^p$-autoreducible sets is covered by an E-betting game. Since every EXP-complete set is $\leq_T^p$-autoreducible [7], this implies results given above. For the subclass of $\leq_{tt}^p$-autoreducible sets, we get a fairly tight picture:

- If MA $\neq$ EXP, then the $\leq_{tt}^p$-autoreducible sets have E-measure zero.
- If EXP = BPP, then the $\leq_{tt}^p$-autoreducible sets have E-measure one.

Here MA is the "Merlin-Arthur" class of Babai [5, 9], which contains BPP and NP.

In sum, the whole theory of resource-bounded measure has progressed far enough to wind the issues of (pseudo-)randomness and stochasticity within exponential time very tightly. We turn the wheels a few more notches, and seek greater understanding of complexity classes in the places where the boundary between "measure one" and "measure zero" seems tightest. A full version of this paper with the proofs omitted here is being submitted concurrently to the ECCC Technical Reports Series.

2 Martingales

A *martingale* is abstractly defined as a function d from $\{0,1\}^*$ into the nonnegative reals that satisfies $d(w) = (d(w0) + d(w1))/2$ for all $w \in \{0,1\}^*$. The interpretation in Lutz's theory is that a string $w \in \{0,1\}^*$ stands for an initial segment of a language over an arbitrary alphabet Σ as follows: Let $s_1, s_2, s_3, \ldots$ be the standard lexicographic ordering of Σ^*. Then for any language $A \subseteq \Sigma^*$, write $w \sqsubseteq A$ if for all i, $1 \leq i \leq |w|$, $s_i \in A$ iff the ith bit of w is a 1. We also regard w as a function with *domain* $\{s_1, \ldots, s_{|w|}\}$ and range $\{0,1\}$, writing $w(s_i)$ for the ith bit of w. A martingale d *succeeds on* a language A if the sequence of values $d(w)$ for $w \sqsubseteq A$ is unbounded. Let $S^\infty[d]$ stand for the (possibly empty, often uncountable) class of languages on which d succeeds.

Definition 1 (cf. [12, 14]). Let Δ be a complexity class of functions. A class C of languages *has Δ-measure zero*, written $\mu_\Delta(C) = 0$, if there is a martingale d computable in Δ such that $C \subseteq S^\infty[d]$. One also says that d *covers* C.

For example, P has E-measure zero. Indeed, for any fixed $c > 0$, DTIME$[2^{cn}]$ has E-measure zero, and DTIME$[2^{n^c}]$ has EXP-measure zero [12].

Lutz defined complexity bounds in terms of the length of the argument w to d, which we denote by N. However, we prefer to work in terms of the largest length n of a string in the domain of w. For $N > 0$, n equals $\lfloor \log N \rfloor$; all we care about is that $n = \Theta(\log N)$ and $N = 2^{\Theta(n)}$. Because complexity bounds on languages we want to analyze will naturally be stated in terms of n, we prefer to use n for martingale complexity bounds. The following correspondence is helpful:

$$\begin{array}{llll}\text{Lutz's ``}p\text{''} & \sim \quad N^{O(1)} = 2^{O(n)} & \sim & \text{measure on E} \\ \text{Lutz's ``}p_2\text{''} & \sim 2^{(\log N)^{O(1)}} = 2^{n^{O(1)}} & \sim & \text{measure on EXP}\end{array}$$

Our convention lets us simply write "μ_{E}" for E-measure (regarding Δ as E for functions), similarly "μ_{EXP}" for EXP-measure, and generally μ_Δ for any Δ that names both a language and function class. Abusing notation similarly, we define:

Definition 2 (after [12]). A class $\mathcal{C}$ has Δ-measure one, written $\mu_\Delta(\mathcal{C}) = 1$, if $\mu_\Delta(\Delta \setminus \mathcal{C}) = 0$.

The following lemma has appeared in various forms [14, 8]. It essentially says that we can assume a martingale grows almost monotonically (sure winnings) and not too fast (slow winnings).

Lemma 3. *Let d be a time-$t(n)$ computable martingale, with $d(\lambda) = 1$. Then we can compute in time $O(2^n t(n))$ a martingale d' with $S^\infty[d] \subseteq S^\infty[d']$ such that*

$$(\forall w)(\forall u) : d'(wu) > d'(w) - 2, \quad \text{and} \tag{1}$$

$$(\forall w) : d'(w) < 2(|w| + 1) \ . \tag{2}$$

3 Betting Games

To capture intuitions that have been expressed not only for Lutz measure but also in many earlier papers on random sequences, we formalize a betting game as an *infinite* process, rather than as a Turing machine that has *finite* computations on string inputs.

Definition 4. A *betting game* G is an oracle Turing machine that maintains a "capital tape" and a "bet tape," in addition to its standard query tape and worktapes, and works in *stages* $i = 1, 2, 3 \ldots$ as follows: Beginning each stage i, the capital tape holds a nonnegative rational number C_{i-1} — initially $C_0 = 1$. G computes a query string x_i to bet on, a *bet amount* B_i, $0 \leq B_i \leq C_{i-1}$, and a *bet sign* $b_i \in \{-1, +1\}$. The computation is *legal* so long as x_i does not belong to the set $\{x_1, \ldots, x_{i-1}\}$ of strings queried in earlier stages. G ends stage i by entering a special query state. For a given oracle language A, if $x_i \in A$ and $b_i = +1$, or if $x_i \notin A$ and $b_i = -1$, then the new capital is given by $C_i := C_{i-1} + B_i$, else by $C_i := C_{i-1} - B_i$. The query and bet tapes are blanked, and G proceeds to stage $i + 1$.

Since we require that G spend the time to write each bet out in full, it does not matter whether we suppose that the new capital is computed by G itself or updated instantly by the oracle. Note that every oracle set A determines a unique infinite computation of G, which we denote by G^A. This includes a unique infinite sequence $x_1, x_2, \ldots$ of query strings, and a unique sequence $C_0, C_1, C_2, \ldots$ telling how the gambler fares against A .

Definition 5. A betting machine G *runs in time* $t(n)$ if for all oracles A, every query of length n made by G^A is made in the first $t(n)$ steps of the computation.

A similar definition can be made for space usage, taking into account standard issues such as whether the query tape counts against the space bound, or whether the query itself is preserved in read-only mode for further computation by the machine.

Definition 6. A betting game G *succeeds* on a language A, written $A \in S^\infty[G]$, if the sequence of values C_i in the computation G^A is unbounded. If $A \in S^\infty[G]$, then we also say G *covers* A.

Our main motivating example where one may wish not to bet in lexicographic order, or according to any fixed ordering of strings, is deferred to Sect. 6.

We now want to argue that the more liberal requirement of being covered by a time $t(n)$ betting game, still defines a smallness concept for subclasses of DTIME$[t(n)]$ in the intuitive sense Lutz established for his measure-zero notion. The following result is a good beginning.

Theorem 7. *For every time-$t(n)$ betting game G, we can construct a language in* DTIME$[t(n)]$ *that is not covered by G.*

In particular, the class E cannot be covered by an E-betting game, nor EXP by an EXP-betting game. Put another way, the "measure conservation axiom" [12] of Lutz's measure carries over to betting games.

To really satisfy the intuition of "small," however, it should hold that the union of two small classes is small (moreover, "easy" countable unions of small classes should be small, as in [12]). Our lack of meeting this "finite union axiom" will later be excused insofar as it has the non-relativizing consequence BPP $\neq$ EXP. Theorem 7 is still good enough for the "measure-like" results in this paper.

To begin comparing betting games and martingales, we note first that the latter can be considered a direct special case of betting games. Say a betting game G is *lex-limited* if for all oracles A, the sequence $x_1, x_2, x_3 \ldots$ of queries made by G^A is in lex order. (It need not equal the lex enumeration $s_1, s_2, s_3, \ldots$ of Σ^*.)

Theorem 8. *Let $\mathcal{T}(n)$ be a collection of time bounds that is closed under multiplication by 2^n, such as $2^{O(n)}$ or $2^{n^{O(1)}}$. Then a class C has time-$\mathcal{T}(n)$ measure zero iff C is covered by a time-$\mathcal{T}(n)$ lex-limited betting game.*

Hence in particular for measure on E and EXP, martingales are equivalent to betting games constrained to bet in lex order.

A general betting game embodies a martingale in a different sense given by the following definition:

Definition 9. Let G be a betting games, and $i \geq 0$ an integer.

(a) A *play* α of length i is a sequence of i oracle answers. Note that α determines the first i-many stages of G, together with the query and bet for the next stage.
(b) $c_G(\alpha)$ is the capital C_i that G has at the end of the play α.

Note that the function c_G is a martingale over plays α. The following carryover of Lemma 3 is important.

Lemma 10 ("Slow-But-Sure" Lemma for betting games). *Let G be a betting game that runs in time $t(n)$. Then we can construct a betting game G' running in time $O(t(n))$ that always makes the same queries in the same order as G, such that $S^\infty[G] \subseteq S^\infty[G']$ and:*

$$(\forall\alpha)(\forall\beta) : c_{G'}(\alpha\beta) > c_{G'}(\alpha) - 2, \quad \text{and} \tag{3}$$

$$(\forall\alpha) : c_{G'}(\alpha) < 2|\alpha| + 2\ . \tag{4}$$

4 From Betting Games to Martingales

This section associates to every betting game G a martingale d_G such that $S^\infty[G] \subseteq S^\infty[d_G]$, and begins examining the complexity of d_G. Before defining d_G, however, we discuss some tricky subtleties of betting games and their computations.

Given a finite initial segment w of an oracle language A, one can define the partial computation G^w of the betting game up to the stage i at which it first makes a query x_i that is not in the domain of w. Define $d(w)$ to be the capital C_{i-1} that G had entering this stage. It is tempting to think that d is a martingale and that d succeeds on all A for which G succeeds—but neither statement is true in general.

To see this, suppose x_i itself is the lexicographically least string not in the domain of w. That is, x_i is indexed by the bit b of wb, and $w1 \sqsubseteq A$ iff $x_i \in A$. It is possible that G^A makes a small (or even zero) bet on x_i, *and then goes back to make more bets in the domain of w, winning lots of money on them.* The definitions of both $d(w0)$ and $d(w1)$ will then reflect these added winnings, and both values will be greater than $d(w)$. For example, suppose G^A first puts a zero bet on $x_i = s_j$, then bets all of its money on $x_{i+1} = s_{j-1}$ not being in A, and then proceeds with $x_{i+2} = s_{j+1}$. If $w(s_{j-1}) = 0$, then $d(w0) = d(w1) = 2d(w)$.

Put another way, a finite initial segment w may carry much more "winnings potential" than the above definition of $d(w)$ reflects. To capture all of it, one needs to consider potential plays of the betting game outside the domain of w. Happily, one can bound the length of the considered plays via the running time function t of G. Let n be the maximum length of a string indexed by w; i.e., $n = \lfloor \log_2(|w|) \rfloor$. Then after $t(n)$ steps, G cannot query any more strings in the domain of w, so w's potential is exhausted. We will hence define $d_G(w)$ as an *average* value of those plays that can happen, given the query answers fixed by w. We use the following definitions and notation:

Definition 11. For any $t(n)$ time-bounded betting game G and string $w \in \Sigma^*$:

(a) A play α is *t-maximal* if G completes the first $|\alpha|$ stages, but *not* the query and bet of the next stage, within t steps.

(b) A play α is *G-consistent with* w, written $\alpha \sim_G w$, if for all stages j such that the queried string x_j is in the domain of w, $\alpha_j = w(x_j)$. That is, α is a play that could possibly happen given the information in w. Also let $m(\alpha, w)$ stand for the number of such stages j whose query is answered by w.

(c) Finally, put $d_G(\lambda) = 1$, and for nonempty w, with $n = \lfloor \log_2(|w|) \rfloor$ as above, let

$$d_G(w) = \sum_{\alpha\ t(n)-maximal, \alpha \sim_G w} c_G(\alpha)\, 2^{m(\alpha,w)-|\alpha|} \ . \tag{5}$$

The weight $2^{m(\alpha,w)-|\alpha|}$ in (5) has the following meaning: Suppose we extend the simulation of G^w by flipping a coin for every query outside the domain of w, for exactly i stages. Then the number of coin-flips in the resulting play α of length i is $i - m(\alpha, w)$, so $2^{m(\alpha,w)-i}$ is its probability. Thus $d_G(w)$ returns the suitably-weighted average of $t(n)$-step computations of G with w fixed. The interested reader may verify that this is the same as averaging $d(wv)$ over all v of length $2^{t(n)}$ (or any fixed longer length), where d is the non-martingale defined above.

Lemma 12. *The function $d_G(w)$ is a martingale.*

To ensure that d_G succeeds on all languages covered by G, however, we must first arrange that G satisfies the sure-winnings condition (3) of Lemma 10.

Lemma 13. *If G satisfies (3), then $S^\infty[G] \subseteq S^\infty[d_G]$.*

Now we turn our attention to the complexity of d_G. If G is a time-$t(n)$ betting game, it is clear that d_G can be computed deterministically in $O(t(n))$ *space*, because we need only cycle through all α of length $t(n)$, and all the items in (5) are computable in space $O(t(n))$. In particular, every E-betting game can be simulated by an ESPACE-martingale, and every EXP-betting game by an EXPSPACE-martingale. However, we show in the next section that one can *estimate* $d_G(w)$ well without having to cycle through all the α, using a pseudo-random generator to "sample" only a very small fraction of them.

5 Sampling Results

First we determine the accuracy to which we need to estimate the values $d(w)$ of a hard-to-compute martingale. Recall $N = 2^n$.

Lemma 14. *Let d be a martingale and $[\epsilon(i)]$ a sequence whose sum K converges. Suppose we can compute in time $t(n)$ the partial sum $\sum_{i=0}^{N} \epsilon(i)$ and a function $g(w)$ such that $|g(w) - d(w)| \leq \epsilon(N)$ for all w of length N. Then there is a martingale d' computable in time $O(2^n t(n))$ such that for all w, $|d'(w) - d(w)| \leq 2K$.*

Next, we will specify the function f_G that we will sample in order to estimate d_G.

Let G be a $t(n)$ time-bounded betting game. Consider a prefix w and let n denote the largest length of a string in the domain of w. With any string ρ of length $t(n)$, we can associate a unique "play of the game" G defined by using w to answer queries in the domain of w, and the successive bits of ρ to answer queries outside it. We can stop this play after $t(n)$ steps — so the stopped play is a $t(n)$-maximal α — and we define $f_G(w, \rho)$ to be the capital $c_G(\alpha)$. Note that we can compute $f_G(w, \rho)$ in linear time. The proportion of strings ρ of length $t(n)$ that map to the same play α is exactly the weight $2^{m(\alpha,w)-|\alpha|}$ in (5) for $d_G(w)$. Letting E stand for mathematical expectation, we have:

$$d_G(w) = E_{|\rho|=t(n)}[f_G(w, \rho)].$$

To estimate this mean, we apply concepts from pseudo-random generators. Alternatively, we can assume P = NP and apply Stockmeyer's method of approximate counting using alternation [16]—this yields results similar to Theorem 18 given in our full paper that improve those of [4], where P = PSPACE is assumed.

Definition 15 ([15]). (a) A *pseudo-random generator* (PRG) is a function D that, for each n, maps Σ^n into $\Sigma^{r(n)}$ where $r(n) > n$. The function r is called the *stretching* of D. We say that D is computable in a class C if every bit of $D(y)$ is computable in C, given y and the index of the bit in binary.

(b) The *security* $S_D(n)$ of D at length n is the largest integer s such that for any circuit C of size at most s with $r(n)$ inputs

$$|\mathrm{Pr}_x[C(x) = 1] - \mathrm{Pr}_y[C(D(y)) = 1]| \leq \frac{1}{s} ,$$

where x is uniformly distributed over $\Sigma^{r(n)}$ and y over Σ^n.

For our purposes, we will need a PRG computable in E that stretches seeds super-polynomially and has super-polynomial security at infinitely many lengths. Combining the results of Babai et al. [6] and of Nisan and Wigderson [15] with some padding yields:

Theorem 16. *If* MA $\neq$ EXP, *there is a PRG D computable in* E *with stretching $n^{\theta(\log n)}$ such that for any integer k, $S_D(n) \geq n^k$ for infinitely many n.*

We will also use PRGs with exponential security that are computable in exponential time.

Now we can apply PRGs to provide the accuracy and time bounds needed to get the desired martingale from Lemma 14.

Theorem 17. *Let D be a PRG computable in time $\delta(n)$ and with stretching $r(n)$. Let $f : \Sigma^* \times \Sigma^* \to (-\infty, \infty)$ be a linear-time computable function, and $s, R, m : \mathbb{N} \to \mathbb{N}$ be constructible functions such that $s(N) \geq N$ and the following relations hold for any integer $N \geq 0$, $w \in \Sigma^N$, and $\rho \in \Sigma^{s(N)}$:*

$$\begin{aligned} |f(w,\rho)| &\leq R(N) \\ r(m(N)) &\geq s(N) \\ S_D(m(N)) &\geq (s(N) + R(N))^6 \ . \end{aligned} \tag{6}$$

Then we can approximate

$$h(w) = E_{|\rho|=s(N)}[f(w,\rho)] \tag{7}$$

to within N^{-2} in time $O(2^{m(N)} \cdot (s(N) + R(N))^4 \cdot \delta(m(N)))$.

Now, we would like to apply Theorem 17 to approximate efficiently $h = d_G$ given by (5) to within N^{-2}, by setting $f = f_G$ and $s(N) = t(\log N)$. The problem is that a given betting game G running in time $t(n)$ may only guarantee an upper bound of $R(N) = 2^{t(\log N)}$ on $|f(w,\rho)|$. Since S_D can be at most exponential, condition (6) would force $m(N)$ to be $\Omega(t(\log N))$, and Theorem 17 would only yield an approximation computable in time $2^{O(t(\log N))}$. *However*, we can assume wlog. that G satisfies the slow-winnings condition (4) of Lemma 10, in which case an upper bound of $R(N) \in O(N)$ holds. Then the term $s(N)$ in the right-hand side of (6) dominates, provided $t(n) \in 2^{\Omega(n)}$.

Taking everything together, we obtain the following result about transforming E- and EXP-betting games into equivalent E- respectively EXP-martingales:

Theorem 18. *If there is a PRG computable in* E *with security $2^{\Omega(n)}$, then for every* E*-betting game G, there exists an* E*-martingale d such that $S^\infty[G] \subseteq S^\infty[d]$. If there is a PRG computable in* EXP *with security $2^{n^{\Omega(1)}}$, then for every* EXP*-betting game G, there exists an* EXP*-martingale d such that $S^\infty[G] \subseteq S^\infty[d]$.*

6 Autoreducible Sets

An oracle Turing machine M is said to *autoreduce* a language A if $L(M^A) = A$, and for all strings x, M^A on input x does not query x. That is, one can learn the membership of x by querying strings other than x itself. If M runs in polynomial time, then A is P-*autoreducible*—we also write $\leq_T^p$-autoreducible. If M is also non-adaptive, then A is *$\leq_{tt}^p$-autoreducible.*

Autoreducible sets were brought to the polynomial-time context by Ambos-Spies [3]. Their importance was further argued by Buhrman, Fortnow, and Torenvliet [7], who showed that all $\leq_T^p$-complete sets for EXP are $\leq_T^p$-autoreducible (while some complete sets for other classes are not). Here we demonstrate that autoreducible sets are important for testing the boundaries of Lutz's measure theory. As stated in the Introduction, if the $\leq_T^p$-autoreducible sets in EXP (or sufficiently the $\leq_T^p$-complete sets for EXP) are covered by an EXP-martingale, then EXP $\neq$ BPP, a non-relativizing consequence. However, it is easy to cover them by an E-betting game. Indeed, the betting game uses its adaptive freedom only to "look ahead" at the membership of lexicographically greater strings, betting nothing on them.

Theorem 19. *There is an* E*-betting game that covers all $\leq_T^p$-autoreducible sets.*

Proof. One can effectively enumerate oracle TMs $M_1, M_2, \ldots$ that never query their input, with each M_i running in time $n^i + i$. Our betting game G regards its capital as composed of infinitely many "shares" c_i, one for each M_i. Initially, $c_i = 1/2^i$. Letting $\langle\cdot,\cdot\rangle$ be a standard pairing function, inductively define $n_0 = 0$ and $n_{\langle i,j\rangle+1} = (n_{\langle i,j\rangle})^i + i$.

During a stage $s = \langle i,j\rangle$, G simulates M_i on input 0^{n_s-1}. Whenever M_i makes a query of length less than n_{s-1}, G looks up the answer from its table of past queries. Whenever M_i makes a query of length n_{s-1} or more, G places a bet of zero on that string and makes the same query. Then G bets all of the share c_i on 0^{n_s-1} according to the answer of the simulation of M_i. Finally, G "cleans up" by putting zero bets on all strings with length in $[n_{s-1}, n_s)$ that were not queries in the previous steps.

If M_i autoreduces A, then share c_i doubles in value at each stage $\langle i,j\rangle$, and makes the total capital grow to infinity. And G runs in time $2^{O(n)}$—indeed, only the "cleanup" phase needs this much time. □

Corollary 20. *Each of the following statements implies* BPP $\neq$ EXP*, as do the statements obtained on replacing "*E*" by "*EXP*."*

1. *The class of $\leq_T^p$-autoreducible sets has* E*-measure zero.*
2. *The class of $\leq_T^p$-complete sets for* EXP *has* E*-measure zero.*
3. E*-betting games and* E*-martingales are equivalent.*
4. E*-betting games have the finite union property.*

Since there is an oracle A giving $\text{EXP}^A = \text{BPP}^A$ [10], this shows that relativizable techniques cannot establish the equivalence of E-martingales and E-betting games, nor of EXP-martingales and EXP-betting games. They cannot refute it either, since there are oracles relative to which strong PRGs exist—all "random" oracles, in fact.

It is tempting to think that the *non*-adaptively P-autoreducible sets should have E-measure zero, or at least EXP-measure zero, insofar as betting games are the adaptive cousins of martingales. However, it is not just adaptiveness but also the freedom to bet *out of the fixed lexicographic order* that adds power to betting games. If one carries out the proof of Theorem 19 to cover the class of $\leq_{tt}^{p}$-autoreducible sets, using an enumeration $[M_i]$ of $\leq_{tt}^{p}$-autoreductions, one obtains a *non-adaptive* E-betting game that (independent of its oracle) bets on all strings in order given by a single permutation of Σ^*. The permutation itself is E-computable. It might seem that an E-martingale should be able to "un-twist" the permutation and succeed on all these sets. However, our next results, which strengthen the above corollary, close the same "non-relativizing" door on proving this with current techniques.

Theorem 21. *For any k, the $\leq_{tt}^{p}$-complete sets for Δ_k^{p} are $\leq_{tt}^{p}$-autoreducible.*

Corollary 22. *Each of the following statements implies* BPP $\neq$ EXP, *as do the statements obtained on replacing "*E*" by "*EXP*."*

1. *The class of $\leq_{tt}^{p}$-autoreducible sets has* E*-measure zero.*
2. *The class of $\leq_{tt}^{p}$-complete sets for* EXP *has* E*-measure zero.*
3. *Non-adaptive* E*-betting games and* E*-martingales are equivalent.*
4. *If two classes can be covered by non-adaptive* E*-betting games, then their union can be covered by an* E*-betting game.*

This puts the spotlight on the question: Under what hypotheses can we show that the $\leq_{tt}^{p}$-autoreducible sets have E-measure zero? Our final results show that the hypothesis MA $\neq$ EXP suffices. This assumption is only known to yield PRGs of super-polynomial security (at infinitely many lengths) rather than exponential security (at almost all lengths). On the other hand, there exist oracles relative to which exponentially strong PRGs exist, but EXP = MA.

Theorem 23. *If* MA $\neq$ EXP, *then the class of languages A autoreducible by polynomial-time OTMs that make their queries in lex order has* E*-measure zero.*

Corollary 24. *If* MA $\neq$ EXP, *then the $\leq_{tt}^{p}$-autoreducible sets have* E*-measure zero.*

7 Conclusions

The initial impetus for this work was a simple question about measure: is the pseudo-randomness of a characteristic sequence invariant under simple permutations such as that induced by *flip* in the Introduction? Our "betting games" preserve Lutz's original idea of "betting" as a means of "predicting" membership in a language, without being tied to a fixed order of which instances one tries to predict, or to a fixed order of how one goes about gathering information on the language. We have shown some senses in which betting games are robust and well-behaved. We also contend that some current defects in the theory of betting games, notably the lack of a finite-unions theorem pending the status of pseudo-random generators, trade off with lacks in the resource-bounded measure theory, such as being tied to the lexicographic ordering of strings.

The research problems left in this paper that are most open to attack are to tighten even further the connections among PRGs, separation of classes within

EXP, and resource-bounded measure. Does EXP $\neq$ MA suffice to make the $\leq_T^p$-autoreducible sets have E-measure zero? Does that suffice to simulate every betting game by a martingale of equivalent complexity?

Another challenge is to determine how well these ideas work for measures on classes below E. Here even straightforward attempts to carry over Lutz's definitions run into difficulties, as described in [14] and [1, 2]. Perhaps the results in these papers can be re-cast in terms of betting games in ways that release new insights.

Acknowledgments The authors specially thank Klaus Ambos-Spies, Ron Book (pace), and Jack Lutz for organizing a special Schloss Dagstuhl workshop in July 1996, where preliminary versions of results and ideas in this paper were presented and extensively discussed. We also thank the STACS'98 referees for helpful comments.

References

1. Allender, E., Strauss, M.: Measure on small complexity classes, with applications for BPP. DIMACS TR 94-18, Rutgers University and DIMACS, April 1994.
2. Allender, E., Strauss, M.: Measure on P: Robustness of the notion. In *Proc. 20th International Symposium on Mathematical Foundations of Computer Science*, volume 969 of *Lect. Notes in Comp. Sci.*, pages 129–138. Springer Verlag, 1995.
3. Ambos-Spies, K.: P-mitotic sets. In E. Börger, G. Hasenjäger, and D. Roding, editors, *Logic and Machines, Lecture Notes in Computer Science 177*, pages 1-23. Springer-Verlag, 1984.
4. Ambos-Spies, K., Lempp, S.: Presentation at a Schloss Dagstuhl workshop on "Algorithmic Information Theory and Randomness," July 1996.
5. Babai, L.: Trading group theory for randomness. In *Proc. 17th Annual ACM Symposium on the Theory of Computing*, pages 421–429, 1985.
6. Babai, L., Fortnow, L., Nisan, N., Wigderson, A.: BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3, 1993.
7. Buhrman, H., Fortnow, L., Torenvliet, L.: Using autoreducibility to separate complexity classes. In *36th Annual Symposium on Foundations of Computer Science*, pages 520–527, Milwaukee, Wisconsin, 23–25 October 1995. IEEE.
8. Buhrman, H., Longpré, L.: Compressibility and resource bounded measure. In *13th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1046 of *lncs*, pages 13–24, Grenoble, France, 22–24 February 1996. Springer.
9. Babai, L., Moran, S.: Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. *J. Comp. Sys. Sci.*, 36:254–276, 1988.
10. Heller, F.: On relativized exponential and probabilistic complexity classes. *Inform. and Control*, 71:231-243, 1986.
11. Loveland, D. W.: A variant of the Kolmogorov concept of complexity. *Inform. and Control*, 15:510–526, 1969.
12. Lutz, J.: Almost everywhere high nonuniform complexity. *J. Comp. Sys. Sci.*, 44:220–258, 1992.
13. Lutz, J.: The quantitative structure of exponential time. In L. Hemaspaandra and A. Selman, eds., *Complexity Theory Retrospective II*. Springer Verlag, 1997.
14. Mayordomo, E.: *Contributions to the Study of Resource-Bounded Measure*. PhD thesis, Universidad Polytécnica de Catalunya, Barcelona, April 1994.
15. Nisan, N., Wigderson, A.: Hardness versus randomness. *J. Comp. Sys. Sci.*, 49:149–167, 1994.
16. Stockmeyer, L.: The complexity of approximate counting. In *Proc. 15th Annual ACM Symposium on the Theory of Computing*, pages 118-126, Baltimore, USA, April 1983. ACM Press.