The Round Complexity of Two-Party Random Selection*

Saurabh Sanghvi† Salil Vadhan‡
Division of Engineering & Applied Sciences
Harvard University, Cambridge, MA
ssanghvi@post.harvard.edu, salil@eecs.harvard.edu

February 6, 2007

Abstract

We study the round complexity of two-party protocols for generating a random n-bit string such that the output is guaranteed to have bounded "bias" even if one of the two parties deviates from the protocol (possibly using unlimited computational resources). Specifically, we require that the output's statistical difference from the uniform distribution on $\{0,1\}^n$ is bounded by a constant less than 1.

We present a protocol for the above problem that has $2\log^* n + O(1)$ rounds, improving a previous $2n$-round protocol of Goldreich, Goldwasser, and Linial (FOCS '91). Like the GGL Protocol, our protocol actually provides a stronger guarantee, ensuring that the output lands in any set $T \subseteq \{0,1\}^n$ of density μ with probability at most $O(\sqrt{\mu + \delta})$, where δ may be an arbitrarily small constant.

We then prove a nearly matching lower bound, showing that any protocol guaranteeing bounded statistical difference requires at least $\log^* n - \log^* \log^* n - O(1)$ rounds.

We also prove several results for the case when the output's bias is measured by the maximum *multiplicative* factor by which a party can increase the probability of a set $T \subseteq \{0,1\}^n$.

Keywords: cryptography, distributed computing, coin-flipping

*Preliminary versions of this work have appeared in the first author's undergraduate thesis [San04], and in the conference paper [SV05].

†Supported in part by a grant from the Harvard College Research Program and NSF grant CCR-0133096.

‡Work done in part while a Fellow at the Radcliffe Institute for Advanced Study at Harvard University. Also supported by NSF grants CNS-0430336, ONR grant N00014-04-1-0478, and a Sloan Research Fellowship.

1 Introduction

One of the most basic protocol problems in cryptography and distributed computing is that of *random selection*, in which several mutually distrusting parties aim to generate an n-bit random string jointly. The goal is to design a protocol such that even if a party cheats, the outcome will still not be too "biased". (There are many different choices for how to measure the "bias" of the output; the one we use will be specified later.) Random selection protocols can dramatically simplify the design of protocols for other tasks via the following common methodology: first design a protocol in a model where truly random strings are provided by a trusted third party (generally a much easier task), and then use the random selection protocol to eliminate the trusted third party. Specific applications of this paradigm often require random selection protocols with specific additional properties (such as "simulatability"), but the basic requirement of bounded "bias" in the face of adversarial behavior is always present in some form and thus merits study on its own.

Because of their wide applicability, there is a large literature on random selection protocols, both in the computational setting, where cheating parties are restricted to polynomial time (starting with Blum's "coin flipping by telephone" [Blu82]), and in the information-theoretic setting, where security is provided even against computationally unbounded adversaries. There has also been a significant amount of recent work on random selection in the quantum setting, where the communication consists of quantum bits (qubits) and security is provided against a computationally unbounded quantum adversary; see [Amb04] and the references therein.

In this paper, we focus on *two-party* protocols in the (classical) *information-theoretic setting* (also known as the "full information model"). In addition to its stronger security guarantees, the information-theoretic setting has the advantage that protocols typically do not require complexity-theoretic assumptions (such as the existence of one-way functions). Various such random selection protocols have been used to construct perfectly hiding bit-commitment schemes [NOVY98], to convert honest-verifier zero-knowledge proofs into general zero-knowledge proofs [Dam94, DGW94, GSV98], to construct oblivious transfer protocols in the bounded storage model [CCM98, DHRS04], and to perform general fault-tolerant computation [GGL98]. There has also been substantial work in the k-party case for $k \geq 3$, where the goal is to tolerate coalitions of a minority of cheating players. This body of work includes the well-studied "collective coin-flipping" problem [BL89, Sak89, AN93, ORV94, RZ98, Fei99] (closely related to the "leader election" problem) and again the use of random selection as a tool for general fault-tolerant computation [GGL98].

In most of the lines of work mentioned above (computational and information-theoretic, two-party and k-party), the *round complexity* has been a major parameter of interest. For some forms of random selection and their applications, constant-round protocols have been found (e.g. [DGW94, GSV98] improving [Dam94], [DHRS04] improving [CCM98], and [Lin01, KO04] improving [Blu82, Yao86]), but for others the best known protocols have a nonconstant number of rounds, e.g. [Cle86, NOVY98, GGL98, RZ98]. Lower bounds on round complexity, however, have proven much more difficult to obtain. In the computational setting, Cleve [Cle86] proved that for 2-party random selection protocols, the number of rounds must grow linearly as the bias of the output tends to zero. (See also [CI93].) Ambainis [Amb04] gave a similar kind of result for 2-party quantum protocols for leader election, a.k.a. weak coin-flipping. As far as we know, all other previous round complexity lower bounds impose additional constraints on the protocol (beyond the basic security guarantee of bounded bias). For example, in the computational setting, it has been recently shown that 5 rounds are necessary and sufficient for random selection protocols satisfying a certain "black-box simulation" condition [KO04]. In the information-theoretic setting, a long line of work on the collective coin-flipping problem has culminated in the $(\log^* n + O(1))$-round protocol of Russell and Zuckerman [RZ98] (see also Feige [Fei99]), but the only known lower bound (of $\Omega(\log^* n)$ rounds), due to Russell, Saks, and Zuckerman [RSZ99], is restricted to protocols where each party can only communicate a small number of bits per round. Without this restriction, it is not even known how to prove that 1 round is impossible.

The problem and main results. As mentioned above, previous works on random selection have considered a number of different measures of the bias of the output, typically motivated by particular applications. Here we focus on what we consider to be the most natural measure — the statistical difference (i.e., total variation distance) of the output from the uniform distribution. The *statistical difference* between two random variables X and Y taking values in a universe $\mathcal{U}$ is defined to be $\max_{S\subseteq\mathcal{U}} |\Pr[X \in S] - \Pr[Y \in S]|$. We call the maximum statistical difference of the output from uniform when a player is honest (but when the other may deviate arbitrarily from the protocol) that player's *statistical guarantee*. We seek a two-party protocol that produces an output in $\{0,1\}^n$ such that both players' statistical guarantees are constant (i.e., bounded away from 1). Equivalently (see Lemma 2.4), we want to satisfy the following criterion.

Statistical Criterion: There are fixed constants $\mu > 0$ and $\epsilon > 0$ such that for every n and every subset $T \subseteq \{0,1\}^n$ of density at most μ, the probability that the output lands in T is at most $1-\epsilon$, even if one party deviates arbitrarily from the specified protocol.

In addition to being a natural choice, this criterion is closely related to others considered in the literature. In particular, the standard criterion for the "collective coin-flipping" problem is that the output bit $B \in \{0,1\}$ satisfies $\max\{\Pr[B=0], \Pr[B=1]\} < p$, where p is a constant less than 1; this is equivalent to B's statistical difference from uniform being bounded away from 1. (Here we see that the problem we consider is in some sense "dual" to collective coin-flipping — we restrict to two players but the output comes from a large set, whereas in collective coin-flipping there are many players but the output has only two possibilities.)

Of course, the first question is whether or not the Statistical Criterion can be met at all, regardless of round complexity. Indeed, being able to tolerate computationally unbounded cheating strategies is a strong requirement. In fact, when $n = 1$ (i.e. the output is a single bit), it turns out that one of the two parties can always force the outcome to be constant [Sak89]. This implies that the Statistical Criterion is impossible to meet for $\mu = 1/2$. Surprisingly, the criterion is achievable, however, for some smaller constant $\mu > 0$. This is implied by the following result of Goldreich, Goldwasser, and Linial [GGL98].

Theorem 1.1 ([GGL98]) *For every n, there is a two-party protocol producing output in $\{0,1\}^n$ such that, as long as one party plays honestly, the probability that the output lands in any set $T \subseteq \{0,1\}^n$ of density μ is at most $p = O(\sqrt{\mu})$. The protocol has $2n$ rounds.*

Notice that for sufficiently small μ, the probability p is indeed a constant less than 1. This implies that the Statistical Criterion is achievable with a *linear* number of rounds. Our goal is to determine the minimal round complexity of this problem.

First, we give a protocol achieving the Statistical Criterion with substantially fewer rounds than the above.

Theorem 1.2 *For every constant $\delta > 0$, there is an efficient two-party protocol producing output in $\{0,1\}^n$ with $2\log^* n + O(1)$ rounds such that, as long as one party plays honestly, the probability that the output lands in any set T of density μ is at most $p = O(\sqrt{\mu + \delta})$.*

Our protocol is inspired by the $\log^* n$-round protocols for leader election [RZ98, Fei99] and Lautemann's proof that **BPP** is contained in the polynomial hierarchy [Lau83]. Specifically, we exhibit a 2-round protocol that reduces the universe of size $N = 2^n$ to a universe of size $\text{polylog}(N)$, while approximately preserving the density of the set T with high probability. Repeating this protocol $\log^* n$ times reduces the universe size to a constant, after which point we apply the GGL protocol.

Second, we prove a lower bound that matches the above up to a factor of $2 + o(1)$.

Theorem 1.3 *Any two-party protocol producing output in $\{0,1\}^n$ that satisfies the Statistical Criterion must have at least $\log^* n - \log^* \log^* n - O(1)$ rounds.*

Our proof of this theorem is a technically intricate induction on the game tree of the protocol. Roughly speaking, we associate to each node z of the game tree a collection $\mathcal{H}$ of very small sets such that if the protocol is started at z and R is a random subset of the universe of density $o(1)$, one of the players X can force the outcome of the protocol to land in $R \cup S$ with probability $1 - o(1)$ for any $S \in \mathcal{H}$. The challenge is to keep the size of the sets in the collections $\mathcal{H}$ small as we induct up the game tree (so that they remain of density $o(1)$ when z is the root, which yields the desired lower bound). In particular, a node can have an arbitrary number of children, so we cannot afford to take unions of sets S occurring across all children. The key idea that allows us to keep the sets small is the following: We consider two cases: If we have a collection of sets that contains a large disjoint subcollection, then the random set R will contain one of the sets with high probability and so we do not need to carry the set through the recursion. On the other hand, if the collection of sets has no large disjoint subcollection, then we show how we can use this fact to construct a successful strategy for the *other* player (based on how we inductively construct the collections $\mathcal{H}$).

We stress that our lower bound does not impose any additional constraint on the protocol, such as the number of bits sent per round. Thus, we hope that our techniques can help in establishing unrestricted lower bounds on round complexity for other problems, in particular for the collective coin-flipping (and leader election) problem.

Results on multiplicative guarantees. A different measure of the quality of random selection protocol is a *multiplicative guarantee* ρ, whereby we require that even if one player cheats, the probability that the outcome lands in any set T of density μ is at most $\rho \cdot \mu$. The goal, naturally, is for ρ to be as small as possible (ideally a constant independent of n). Previous protocols, e.g. the one in [DGW94], have given a multiplicative guarantee to one player while the other has a statistical guarantee (i.e. a bound on the output's statistical difference from uniform if the other cheats). Our observations and results on multiplicative guarantees are the following:

- If both parties have multiplicative guarantees ρ_A and ρ_B, then an argument of [GGL98] implies $\rho_A \cdot \rho_B \geq 2^n$, regardless of the number of rounds. On the other hand, for any desired ρ_A, there is a simple two-round protocol with multiplicative guarantees of ρ_A and $2^n/\rho_A$ for the two players.

- If one party has a multiplicative guarantee ρ and the other has a statistical guarantee ε, then $\varepsilon \geq 1/\rho - 1/2^n$. This explains inverse relationships in existing protocols of [DGW94] (where $\varepsilon = 1/\text{poly}(n)$ and $\rho = \text{poly}(n)$) and [GSV98] (where $\varepsilon = \text{poly}(n) \cdot 2^{-k}$ and $\rho = 2^k$ for any k).[1]

[1] Actually, the protocol of [GSV98] does not provide a multiplicative guarantee of 2^k, but rather ensures that the

- There is a protocol with $2\log^* n + O(1)$ rounds that provides a constant statistical guarantee to one player and a $1+\delta$ multiplicative guarantee to the other, for an arbitrarily small constant δ. Theorem 1.3 implies that this round complexity is tight up to a constant factor, because a constant multiplicative guarantee implies a constant statistical guarantee.

Notation for Logarithms. As in other work [RZ98], for the purposes of this paper, we define $\log_b^{(k)} n$ to be k base-b iterated logarithms of n, with 1 being a minimum value:

$$\log_b^{(k)} n = \begin{cases} 1 & : \text{ if } \log_b^{(k-1)} n < b \\ \log_b \left(\log_b^{(k-1)} n\right) & : \text{ otherwise} \end{cases}$$

with $\log_b^{(0)} n = n$. Moreover, for $n \geq 1$, we define $\log_b^* n$ to be the least natural number k such that $\log_b^{(k)} n = 1$. Throughout the paper we take the base of the logarithms to be $b = 2$ unless otherwise specified.

2 Defining Random Selection Protocols

Although we introduced the problem for a universe $\{0,1\}^n$, for the rest of the paper we assume we have an arbitrary universe $\mathcal{U}$. We can formally characterize a random selection protocol as follows:

Definition 2.1 *A* random selection protocol $\Pi = (A, B, f)$ *over a universe* $\mathcal{U}$ *consists of a pair of programs A and B and a function f such that:*

- *Both A (Alice) and B (Bob) alternately output strings ("messages") m_i of arbitrary length that are a function of the conversation thus far and their sequences of random coin tosses r_A and r_B, respectively. That is, $m_1 = A(r_A)$, $m_2 = B(r_B, m_1)$, $m_3 = A(r_A, m_1 m_2)$, etc.*
- *The* conversation *between Alice and Bob is the transcript $(A, B) = (m_1, m_2, \ldots, m_r)$, where r is a parameter defining the number of messages (also called the number of "rounds" or "turns") of the protocol.*
- *The* output *of the protocol is $f((m_1, m_2, \ldots, m_r))$, which is some element of $\mathcal{U}$.*

We are interested in the behavior of the protocol when one of these programs is replaced with an arbitrary "cheating" program A^* or B^*, which may send its messages as an arbitrary function of the conversation and input length. If the cheating program "aborts" or sends an irregular message (too long, ill-formed, etc.) the protocol can assume it has sent the empty string.

Although the formulation we have provided assumes a protocol operates over a single fixed universe, in general we will be interested in studying asymptotic behavior of protocols as the universe size increases. Thus, we define a *random selection protocol ensemble* to be a sequence $(\Pi^{(1)}, \Pi^{(2)}, \ldots)$ where each $\Pi^{(N)}$ is a protocol over $\mathcal{U} = \{1, \ldots, N\}$. From now on, we blur the distinction between random selection protocols over a fixed universe and random selection protocol ensembles.

probability that the output lands in any set T of density μ is at most $2^k \cdot \mu + o(1)$. Our lower bound also applies to this more general type of guarantee.

Two additional desirable properties of random selection protocols are the following: (a) the output is uniformly distributed in $\mathcal{U}$ assuming honest players; (b) in a protocol ensemble, honest strategies can be computed in time polynomial in the output length, $\log N$. Our protocols will have these properties, but our lower bounds will apply even to protocols without them.

We now introduce a formalism that will be essential in the proofs of this paper.

Definition 2.2 *Given a protocol* Π *over universe* $\mathcal{U}$*, define the* game tree T *as follows:*

- *A set of nodes* V*, each representing a partial transcript of messages,* $(m_1, \ldots, m_i)$.
- *A set of edges* E*, defined by* $(u, v) \in E$ *if and only if* $u = (m_1, \ldots, m_i)$ *and (abusing notation)* $v = (u, m_{i+1})$*, for some message* m_{i+1}*. That is,* u *has* v *as a* child *if* v *is a potential protocol state one message after state* u*. Note that this makes* T *a tree, rooted at the empty transcript.*
- *For each node* z*, a distribution* $\mathcal{D}_z$ *over the children* z_i *whereby* A *or* B *chooses the next message. (where the children are all nodes* z_i *such that* $(z, z_i) \in E$*).*
- *For every leaf* $z = (m_1, \ldots, m_r)$*, a label equal to* $f((m_1, \ldots, m_r))$*, the output of the protocol ending at node* z*.*

One can verify that this formalism produces an equivalent specification to Definition 2.1 of a random selection protocol.

Just as any node of a tree can be viewed as the root of another tree, any node of a protocol's game tree induces its own random selection protocol starting from that state. We simply fix the messages leading to that node, and have the players choose the remaining messages as in the original protocol. This observation is one of the main reasons that the abstraction of a random selection protocol as a tree will prove useful.

Evaluating a Random Selection Protocol. We evaluate random selection protocols with metrics measuring how "close" the output is to the uniform distribution on $\mathcal{U}$. The primary metric we use is the following:

Definition 2.3 *The* statistical difference from uniform *of a distribution* X *over universe* $\mathcal{U}$ *is defined to be*

$$\max_T \left| \Pr_{x \leftarrow X}[x \in T] - \mu(T) \right| = \max_T \left(\Pr_{x \leftarrow X}[x \in T] - \mu(T) \right),$$

where $T \subseteq \mathcal{U}$ *and* $\mu(T)$ *is the density of* T *in* $\mathcal{U}$*,* $|T|/|\mathcal{U}|$.

It can be verified that this distance is in the interval $[0, 1-1/N]$, where N is the size of the universe $\mathcal{U}$. A statistical difference of 0 implies that X is uniform, and $1 - 1/N$ implies X is concentrated on a single point. It is equal to one-half of the ℓ_1 distance between X and the uniform distribution, where we view each as a vector in $[0, 1]^N$.

We will want to avoid output distributions X whose statistical difference from uniform is very close to 1. The following lemma demonstrates that this (undesirable) property is equivalent to X landing in a small set with high probability.

Lemma 2.4 *If* X *has statistical difference at least* $1 - \epsilon$ *from uniform, then there exists a set* T *such that* $\mu(T) \leq \epsilon$ *and* $\Pr_{x \leftarrow X}[x \in T] \geq 1 - \epsilon$*. Conversely, if there exists such a set* T*, then* X *has statistical difference at least* $1 - 2\epsilon$ *from uniform.*

Proof: If X has statistical difference at least $1-\epsilon$ from uniform, then there exists a set T such that $\Pr[x \in T] - \mu(T) \geq 1-\epsilon$. Since $\Pr[x \in T] \leq 1$, we can conclude $\mu(T) \leq \epsilon$, and since $\mu(T) \geq 0$, we can conclude $\Pr[x \in T] \geq 1-\epsilon$.

The second statement follows directly from the definition of statistical difference. ■

Given these metrics, we can define the following:

Definition 2.5 *Let $\Pi = (A, B, f)$ be a random selection protocol. The* statistical guarantee *for Alice playing honest strategy A in Π, denoted ϵ_A, is the maximum over all B^* of the statistical difference between the distribution of $f((A, B^*))$ and the uniform distribution over $\mathcal{U}$. The guarantee for Bob is defined analogously.*

Intuitively, the guarantee of a protocol for a player bounds the damage that the opponent can effect on the distribution by deviating from the protocol. Unfortunately, the terminology here is a bit counterintuitive — the lower the number, the better the guarantee. We will try to avoid confusion by saying a guarantee is "at best x", rather than "at least x".

Armed with this notion of a guarantee, we can state the following important equivalence, following directly from Lemma 2.4:

Proposition 2.6 *The Statistical Criterion is equivalent to both of the statistical guarantees of a protocol being bounded away from 1.*

Later on, we will prove the following proposition, which lower bounds the ability of any protocol to provide strong statistical guarantees to both players simultaneously:

Proposition 2.7 *In any random selection protocol Π over universe $\mathcal{U}$ achieving statistical guarantees ϵ_A and ϵ_B, $\epsilon_A + \epsilon_B \geq 1 - 1/N$, where $N = |\mathcal{U}|$.*

In addition to statistical guarantees, we also consider multiplicative guarantees, which come from bounds on multiplicative difference:

Definition 2.8 *The* multiplicative ratio *of a distribution X is*

$$\max_{T} \Pr_{x \leftarrow X}[x \in T]/\mu(T),$$

where T ranges over all nonempty subsets of $\mathcal{U}$. Similarly to Definition 2.5, for a random selection protocol $\Pi = (A, B, f)$, we define the multiplicative guarantee *ρ_A (respectively, ρ_B) for Alice (resp., Bob), by taking the maximum multiplicative ratio over all cheating strategies B^* (resp., A^*).*

The multiplicative ratio is always a rational in $[1, N]$, where 1 implies the uniform distribution, and N implies one element is chosen with probability 1. The multiplicative ratio of a distribution is actually equal to the factor by which a single element's probability of being the protocol's output can be increased from uniform. Formally:

Lemma 2.9 *The multiplicative ratio of a distribution X is equal to*

$$\max_{s \in \mathcal{U}} \left(N \cdot \Pr_{x \leftarrow X}[x = s] \right).$$

□

Later on, we will prove the following result from [GGL98]:

Theorem 2.10 [GGL98] *For any protocol* Π*, we have* $\rho_A \cdot \rho_B \geq N$*.*

When the universe is $\{0,1\}^n$ (i.e., $N = 2^n$), this theorem implies that one player can always increase the probability a single element is chosen by an exponential factor — namely, $2^{n/2}$.

While both measures bound the deviation of a distribution from uniform, multiplicative ratio tends to focus on the concentration of probability into small sets (indeed, by Lemma 2.9, sets of size 1), while statistical difference will prove more useful when considering larger subsets (e.g., a constant fraction of the universe).

This said, we can prove some basic relationships between the two metrics that will prove useful:

Lemma 2.11 *Let* X *be an arbitrary distribution over universe* $\mathcal{U}$*, with* $N = |\mathcal{U}|$*. Denote by* ϵ *the statistical difference of* X *from uniform and by* ρ *the multiplicative difference from uniform. Then:*

1. $\rho \leq N\epsilon + 1$
2. $\epsilon \leq 1 - 1/\rho$.

Part 2 of Lemma 2.11 implies that a distribution with a constant multiplicative ratio will have a constant statistical difference, though the converse is not necessarily true. Put another way, a strong multiplicative guarantee is harder to achieve than a strong statistical guarantee.

Proof of Lemma 2.11: By Lemma 2.9, we have $\rho = N \cdot \max_{s \in \mathcal{U}} \Pr_{x \leftarrow X}[x = s]$. But by the definition of statistical difference, we have for any $x \in \mathcal{U}$, $\Pr_{x \leftarrow X}[x = s] \leq \epsilon + 1/N$, by setting $T = \{s\}$. Part (1) of the lemma follows.

For Part (2), it suffices to show that for all T, $\Pr_{x \leftarrow X}[x \in T] - \mu(T) \leq 1 - 1/\rho$. If $\mu(T) \geq 1/\rho$, then this certainly holds. Otherwise, $\Pr_{x \leftarrow X}[x \in T] \leq |T| \cdot (\rho/N) = \mu(T)\rho$, which implies that $\Pr_{x \leftarrow X}[x \in T] - \mu(T) \leq Pr_{x \leftarrow X}[x \in T](1 - 1/\rho) \leq 1 - 1/\rho$. ■

3 The Iterated Random Shift Protocol

In this section we describe the main protocol of this paper, the Iterated Random Shift Protocol, and prove its main properties. That is, we show that for any constant δ, Iterated Random Shift is a $(2 \log^* N + O(1))$-round protocol where the probability the output falls in a set of density μ is at most $O(\sqrt{\mu + \delta})$. It follows that the protocol satisfies the Statistical Criterion given above.

3.1 The GGL Protocol

We begin by briefly describing the $2 \log N$ round protocol satisfying the Statistical Criterion given by Goldreich, Goldwasser, and Linial [GGL98], which we will use in the construction of our protocol:

The GGL Protocol (Idealized): Let $\mathcal{U}$ be a universe of size N, where N is a power of 2.

1. Alice randomly divides $\mathcal{U}$ into two equal-sized halves, and sends this partition to Bob.

2. Bob randomly selects one of the halves, randomly divides it, and sends the resulting partition to Alice.

3. This process continues until one element remains: this element is the output of the protocol.

Actually, to obtain a protocol with computation time polylogarthmic in N, the authors use pairwise independent partitions of the universe. They first prove that this protocol achieves that, as long as one party plays honestly, the probability that the output lands in any set $T \subseteq \mathcal{U}$ of density μ is at most $p = O(\mu^{1/4})$.

They go on to improve this result by using a slightly different protocol to achieve Theorem 1.1 (Theorem 23 in [GGL98]), which improves the bound to $p = O(\sqrt{\mu})$.

3.2 The Random Shift Protocol

The Iterated Random Shift Protocol is inspired by the $\log^* n$-round protocols for leader election [RZ98, Fei99] and Lautemann's proof that **BPP** is contained in the polynomial hierarchy [Lau83]. It is built by iteration of the following 2-round protocol, which we will call the Random Shift Protocol:

The Random Shift Protocol: Given a universe $\mathcal{U}$ of size N and $m \in \mathbb{N}$:

1. Alice uniformly randomly selects and sends a sequence of elements $a_1, \ldots, a_m \in \mathcal{U}$.

2. Bob uniformly randomly selects and sends a sequence of elements $b_1, \ldots, b_m \in \mathcal{U}$.

3. Output the sequence $(a_i + b_j)_{1 \leq i,j \leq m}$, where $+$ is a group operation over $\mathcal{U}$.

Note that the Random Shift Protocol is not, strictly speaking, a random selection protocol over $\mathcal{U}$: its output is a *sequence* of elements from the universe. In using it, we will typically choose the parameter m so that the number of output elements, m^2, is much smaller than N (e.g. $m = \text{polylog}(N)$), and recursively use our random selection protocol to select one of the m^2 output elements. To show that this approach yields a protocol with bounded statistical guarantees, we argue that even if one of the players cheats, any subset T of the universe is unlikely to appear in much more than a $\mu(T)$ fraction of the outputs of the Random Shift Protocol. This is formalized by the following lemma.

Lemma 3.1 *Let T be an arbitrary subset of $\mathcal{U}$. Let $\mu(T) = |T|/N$, and let $\mu'(T)$ denote the density of T in the sequence output by the Random Shift Protocol: $\mu'(T) = \#\{(i,j) : a_i + b_j \in T\}/m^2$. Then as long as one player plays honestly (i.e., chooses elements uniformly at random) and $m \geq (1/2\delta^2) \cdot \log(N/\epsilon)$, we have*

$$\Pr[\mu'(T) \geq \mu(T) + \delta] \leq \epsilon.$$

That is, when one player is honest, the sequence $(a_i + b_j)_{ij}$ will be sufficiently random so that it is very unlikely that the density of T in the output sequence will increase substantially.

Proof: Suppose Alice plays honestly and chooses her elements $a_1, \ldots, a_m$ uniformly at random from $\mathcal{U}$. The lemma certainly holds a fortiori for an honest Alice, as a cheating Bob can see what elements Alice has selected.

For each element $b \in \mathcal{U}$, define the random variables

$$X_i^{(b)} = \begin{cases} 1 & \text{if } (a_i + b) \in T \\ 0 & \text{otherwise} \end{cases}$$

Then define $X^{(b)} = (1/m) \sum_{i=1}^m X_i^{(b)}$. Notice that $\mathrm{E}[X_i^{(b)}] = \mu(T)$ and for each b, the random variables $X_1^{(b)}, \ldots, X_m^{(b)}$ are mutually independent.

By a Chernoff bound, we may conclude the following for any δ:

$$\Pr[X^{(b)} \geq \mu(T) + \delta] \leq e^{-2\delta^2 m} \leq 2^{-2\delta^2 m} \leq \epsilon/N.$$

Using a union bound, we conclude:

$$\Pr[\exists b \in \mathcal{U} \text{ such that } X^{(b)} \geq \mu(T) + \delta] \leq N \cdot \epsilon/N = \epsilon.$$

But if for all b, we have $X^{(b)} < \mu(T) + \delta$, then no matter what elements $b_1, \ldots, b_m$ Bob chooses, we have

$$\mu'(T) = (1/m) \sum_{j=1}^m X^{(b_j)} < \mu(T) + \delta.$$

It follows that $\Pr[\mu'(T) \geq \mu(T) + \delta] \leq \epsilon$ as desired. ■

Remark 3.2 We note that the number of elements sent by Bob need only be $(1/2\delta^2) \cdot \log(1/\epsilon)$ (i.e. the $\log N$ factor can be eliminated), since there is no need to do a union bound as in the above proof when proving Bob's guarantee. However, the symmetry of the protocol as presented above has the advantage that it can actually be implemented in *one* round in a model of simultaneous communication (where honest parties can send messages at the same time, but a cheating party may wait to see the other party's message before sending its own message), as is typically used in many-party protocols (e.g. leader election and collective coin-flipping). This reduces the round complexity of our Iterated Random Shift Protocol below to $\log^* N + O(1)$ in the simultaneous communication model. It is interesting to know whether our lower bound of $\log^* N - \log^* \log^* N - O(1)$ rounds (in Section 4.2) can be extended to the simultaneous communication model (without paying the factor of 2 in the trivial reduction to our non-simultaneous model), since we would then have bounds in that model that are tight up to a factor of $1 + o(1)$.

3.3 The Iterated Random Shift Protocol

We now describe our Iterated Random Shift Protocol satisfying Theorem 1.2, which consists of recursively applying the Random Shift Protocol until the universe size is small (say, less than a fixed constant), after which we apply the GGL Protocol from [GGL98] discussed in Section 3.1. Formally:

The Iterated Random Shift Protocol: Given a universe $\mathcal{U}$ of size N and $M \in \mathbb{N}$ being a sufficiently large "cutoff parameter" that is a power of 2. Three cases:

1. If $N > M^2$, then letting $m = \max\{M, \lceil \log^3 N \rceil\}$, execute the following:
   (a) Alice uniformly randomly selects and sends a sequence of elements $a_1, \ldots, a_m \in \mathcal{U}$.
   (b) Bob uniformly randomly selects and sends a sequence of elements $b_1, \ldots, b_m \in \mathcal{U}$.
   (c) Recursively execute the protocol on universe $\mathcal{U}' = [m] \times [m]$ to obtain result (i, j) and output $a_i + b_j$.
2. If $N = M^2$, run the GGL Protocol on $\mathcal{U}$ and output its result.
3. If $N < M^2$, recursively use the protocol on universe $\mathcal{U}' = \mathcal{U} \times [M^2]$ to obtain result (x, y) and output x.

First, we observe that, assuming M is chosen sufficiently large, the recursion will always terminate with an application of Case 2 (the GGL Protocol): When we run Case 1, we have $|\mathcal{U}'| = m^2 = \max\{M^2, \lceil \log^3 N \rceil\} \leq \max\{M^2, N-1\}$ for sufficiently large N, so eventually the universe size will equal M^2, provided M is large enough. Case 3 will be executed at most once, and is present only to avoid running the GGL Protocol immediately if the universe size is not a power of 2.

Observe that the output of this protocol is uniform if both players are playing honestly — by symmetry, all elements of the universe are equally likely to be selected. We note also that, by inspection, the honest players' strategies and the output of the protocol can certainly be computed in polynomial time.

We now analyze the round complexity of the Iterated Random Shift Protocol:

Proposition 3.3 *For all sufficiently large M and all N, the Iterated Random Shift Protocol over a universe $\mathcal{U}$ of size N with parameter M takes $2\log^* N + O(\log M)$ rounds. Moreover, the strategies of the players are computable in time $\mathrm{poly}(\log N, \log M)$.*

Proof: Each application of the Random Shift Protocol (except for the last) reduces the universe size from N to $\lceil \log^3 N \rceil^2 < \log^7 N$ for sufficiently large N, and takes two rounds. A lemma proven in [RZ98] (Lemma 5) states that if $f(n) \leq \log^a n$ for some constant a, then $f^{(\log^* n)}(n) \leq b$ for some constant b depending only on a, where $f^{(k)}$ represents k repeated applications of f. This implies that, if M is sufficiently large and the initial universe size is $N \geq M^2$, the Random Shift Protocol is applied at most $\log^* N$ times. (If $N < M^2$, then we apply the Random Shift Protocol at most $\log^*(NM^2) = \log^* N + O(\log M)$ times.) By Theorem 1.1, the GGL Protocol on a universe of size at most M^2 takes at most $4 \log M$ rounds.

As for the efficiency of the protocol, note that the players need only generate in each round a number of random bits that is polylogarithmic in the size of the universe. ■

3.4 The Statistical Guarantees of the Iterated Random Shift Protocol

Theorem 3.4 *If $M \geq 1/\delta^3$, then for any set $T \subseteq \mathcal{U}$, the probability that the output of the Iterated Random Shift Protocol lands in T is $O(\sqrt{\mu + \delta})$, where $\mu = \mu(T)$, assuming at least one player plays honestly.*

Corollary 3.5 *For a sufficiently large constant M, the Iterated Random Shift Protocol satisfies the Statistical Criterion. Equivalently, there exists a constant $\epsilon > 0$ such that the Iterated Random Shift Protocol achieves* $\max\{\epsilon_A, \epsilon_B\} \leq 1 - \epsilon$.

Observe that Theorem 3.4 is much stronger than what we need to show Corollary 3.5. Using Theorem 3.4, we know that when one player is honest, for any "small" set T, the probability the output falls in T is close to zero. The Statistical Criterion requires only that this probability is not arbitrarily close to 1.

Proof of Theorem 3.4: The key idea is that in the ith application of the Random Shift Protocol, we can bound the increase in density of any particular set T by at most some small δ_i (with high probability) and these δ_i's can be chosen so that $\sum_i \delta_i \leq \delta$. The Iterated Random Shift Protocol concludes by applying the GGL Protocol to this small universe, and then Theorem 1.1 gives us the result.

We first note that the modification of the protocol in case $N < M^2$, selecting from $\mathcal{U} \times [M^2]$ and taking the first component, does not affect the property claimed in the theorem (because the density of $T \times [M^2]$ in $\mathcal{U} \times [M^2]$ equals the density of T in $\mathcal{U}$). Thus we assume that $N \geq M^2$, and let $N_0, N_1, \ldots, N_{k^*}$ be the universe sizes in an execution of the Iterated Random Shift Protocol, where k^* is the first value of k such that $N_k = M^2$. That is,

$$\begin{aligned} N_0 &= N, \\ N_k &= m_k^2, \text{ for } m_k = \max\{M, \lceil \log^3 N_{k-1} \rceil\}. \end{aligned}$$

Note that for sufficiently large M, the sequence of N_i's is strictly decreasing and there exists a finite k^* such that $N_{k^*} = M^2$.

Given a subset $T \subseteq \mathcal{U}$, we track how T evolves through an execution of the Iterated Random Shift Protocol using the following notation for $k = 0, \ldots, k^*$:

$$\begin{array}{ll} \mathcal{U}_0 = \mathcal{U}, & \mathcal{U}_k = [m_k] \times [m_k], \\ T_0 = T, & T_k = \{(i,j) \in \mathcal{U}_k : (a_i + b_j) \in T_{k-1}, 1 \leq i,j \leq m_k\}, \\ \mu(T_k) = |T_k|/|\mathcal{U}_k|, & \end{array}$$

where in the definition of T_k, (a_i) and (b_j) are the sequences of elements of $\mathcal{U}_{k-1}$ chosen by Alice and Bob in the kth application of the Random Shift Protocol, and $+$ is the group operation over $\mathcal{U}_{k-1}$ used in the protocol.

Intuitively, $\mathcal{U}_k$ is the remaining universe (of size N_k) after k iterations and T_k represents the portion of the remaining universe such that choosing $(i,j) \in T_k$ will lead to an element of T being the output of the whole protocol. We call $\mu(T_k)$ the "effective density" of T in the kth iteration.

Claim 3.6 *There is a finite constant C independent of N and M such that, we have*

$$\Pr\left[\mu(T_{k^*}) \geq \mu(T) + C \cdot M^{-1/3}\right] \leq C \cdot 2^{-M^{1/3}},$$

provided at least one party plays honestly.

Proof: Recall that in the kth iteration, we are applying the Random Shift Protocol with parameter $m = m_k = \max\{M, \lceil \log^3 N_{k-1} \rceil\}$. Define $\epsilon_k = 2^{-m_k^{1/3}}$, and $\delta_k = 1/m_k^{1/3}$. Notice that $m_k \geq (1/2\delta_k^2) \cdot \log(N_{k-1}/\epsilon_k)$.

Using Lemma 3.1 repeatedly in an induction and using a union bound, we have that for any k,

$$\Pr\left[\mu(T_k) \geq \mu(T) + \sum_{i=1}^{k} \delta_i\right] \leq \sum_{i=1}^{k} \epsilon_i.$$

Since the N_k's are decreasing exponentially fast, we have

$$\begin{aligned} \sum_{i=0}^{k^*} \delta_i &= O(\delta_{k^*}) \\ &= O(1/m_{k^*}^{1/3}) \\ &= O(1/M^{1/3}). \end{aligned}$$

Similarly,

$$\sum_{i=1}^{k^*} \epsilon_i = O(2^{-m_{k^*}^{1/3}}) = O(2^{-M^{1/3}}).$$

This completes the proof. ■

Applying Claim 3.6 and using Theorem 1.1, we deduce that the probability that the output lands in T is at most

$$O\left(\sqrt{\mu(T) + C \cdot M^{-1/3}}\right) + C \cdot 2^{-M^{1/3}} = O\left(\sqrt{\mu(T) + M^{-1/3}}\right) = O\left(\sqrt{\mu(T) + \delta}\right),$$

using $M \geq 1/\delta^3$. Theorem 3.4 is proven. ■

Recalling Proposition 3.3, and taking M to be a sufficiently large constant, we obtain a protocol with $2\log^* N + O(1)$ rounds satisfying the Statistical Criterion, thereby proving Theorem 1.2. More generally, we obtain a protocol of $2\log^* N + O(\log(1/\delta))$ rounds such that the output lands in a sets of density μ with probability at most $O(\sqrt{\mu + \delta})$. Note that we can take δ to be a slowly vanishing function of N and still have $O(\log^* N)$ rounds.

3.5 The Multiplicative Guarantees of the Iterated Random Shift Protocol

In this section we discuss the multiplicative guarantees provided by the Iterated Random Shift Protocol. Later we will see how lower bounds require that one of the players (in this case, Alice) receives a very poor multiplicative guarantee; however, we will see that Bob receives a very strong guarantee. In this way, we can say something about the ability of a protocol to provide a strong multiplicative guarantee to one player, while providing a strong statistical guarantee to the other.

Specifically, we show:

Theorem 3.7 *There exist constants $\epsilon < 1$ and ρ such that the Iterated Random Shift Protocol with the cutoff parameter M taken to be a sufficiently large constant achieves guarantees $\rho_B \leq \rho$ and $\epsilon_A \leq \epsilon$.*

This is the first protocol achieving constant statistical and multiplicative guarantees that we know of, and later we will prove Theorem 1.3, which, together with Lemma 2.11, implies that it has optimal round complexity (up to a factor of $2 + o(1)$). (See Corollary 5.2.)

Given Corollary 3.5, to prove Theorem 3.7, it suffices to show the following:

Proposition 3.8 *Let Π be a Iterated Random Shift Protocol defined with constant cutoff parameter M. Then Π provides a constant multiplicative guarantee to Bob: there exists constant ρ such that, as long as Bob plays honestly, the output of the Iterated Random Shift Protocol will fall in a set T with probability at most $M^2 \cdot \mu(T)$, for any set T.*

Proof of Proposition 3.8: Fix an arbitrary set $T \subseteq \mathcal{U}$. We use the notation from the proof of Theorem 1.2, in particular $\mathcal{U}_k$ is the remaining universe after k iterations and T_k is the set of elements of $\mathcal{U}_k$ corresponding to elements of T. The following is the key lemma:

Lemma 3.9 *Assuming Bob plays honestly,* $\mathrm{E}[\mu(T_k)] = \mathrm{E}[\mu(T_{k-1})]$ *for all $k = 1, \ldots, k^*$.*

Proof: Consider the Random Shift Protocol. Let $a_1, \ldots, a_m$ be given. Then if $b_1, \ldots, b_m$ are chosen uniformly at random, it follows that for each i, j, the element $a_i + b_j$ is uniform over $\mathcal{U}$ (since $+$ is a group operation), and thus $\Pr[a_i + b_j \in T] = \mu(T)$. By linearity of expectations, we can conclude that $\mathrm{E}[\#(a_i + b_j) \in T] = \mu(T) \cdot m^2$ (where m^2 is the size of the new universe), and thus $\mathrm{E}[\mu(T')] = \mu(T)$, where $\mu(T')$ is the residual density of T in the resulting universe.

Applying this logic within the Iterated Random Shift Protocol, the lemma is proven (since for given $\mu(T_{k-1})$, we know $\mathrm{E}[\mu(T_k)] = \mu(T_{k-1})$).

■

By induction, we then have that for all k, $\mathrm{E}[\mu(T_k)] = \mu(T)$. In particular, this is true for $k = k^*$. We can then derive:

$$
\begin{aligned}
\mu(T) &= \mathrm{E}[\mu(T_{k^*})] \\
&\geq (1/M^2) \cdot \mathrm{E}[|T_{k^*}|] \\
&\geq (1/M^2) \cdot \Pr[|T_{k^*}| > 0].
\end{aligned}
$$

Since if $|T_{k^*}| = 0$, the protocol's output cannot possibly fall in T, we conclude that the probability the output falls in T is at most $M^2 \cdot \mu(T)$. This proves the proposition, and thus Theorem 3.7. ■

As an aside, notice that by using Lemma 2.11, Proposition 3.8 allows us to conclude one half of Theorem 1.2: the Iterated Random Shift Protocol provides a constant statistical guarantee to Bob.

We can conclude that the Iterated Random Shift Protocol has the following properties:

- It has only $2 \log^* N + O(1)$ rounds.
- It provides both Alice and Bob with constant statistical guarantees (equivalently, it satisfies the Statistical Criterion).
- It provides Bob with a constant multiplicative guarantee.

Notice that in the above proof, we never used the multiplicative guarantee properties of the GGL Protocol–we simply relied on the initial recursions of Random Shift to provide the strong guarantee to Bob.

In fact, by changing the protocol used when the universe size becomes of size M^2 in the definition of the Iterated Random Shift Protocol, we can improve even further the multiplicative guarantee given to Bob. The current protocol only implies that Bob gets *some* constant multiplicative guarantee. Specifically, consider the following simple 2-round protocol:

The Random Set Protocol: Given universe $\mathcal{U}$ of size N and parameter K,

1. Alice selects a subset S of $\mathcal{U}$ of size K, uniformly at random, and sends S to Bob.
2. Bob selects an element $x \in S$, uniformly at random.
3. The output is x.

It is straightforward to prove the following:

Proposition 3.10 *For all positive integers $N \geq K$, the Random Set Protocol provides multiplicative guarantees $\rho_A = K$ and $\rho_B = N/K$.*

□

Thus, by using the Random Set Protocol on the universe of size M^2 with parameter $K = M^2/(1+\gamma)$ instead of GGL, Bob can achieve a multiplicative guarantee $1+\gamma$, while still keeping Alice's statistical guarantee constant (when γ is constant — if the residual density of Bob's target set T is smaller than $1 - 1/(1+\gamma)$, there is a nonzero probability that Alice will choose a set S disjoint from T).

In the next section, we will prove that the Iterated Random Shift Protocol has optimal round complexity, up to a factor of $2 + o(1)$, among protocols achieving the Statistical Criterion.

4 Lower Bounds on Statistical Guarantees

4.1 Tradeoffs between Statistical Guarantees

As a warmup to our main lower bound, in this section we present a tradeoff between the statistical guarantees ϵ_A and ϵ_B of Alice and Bob, respectively:

Proposition 4.1 (Proposition 2.7, restated) *In any random selection protocol Π over universe $\mathcal{U}$ achieving statistical guarantees ϵ_A and ϵ_B, $\epsilon_A + \epsilon_B \geq 1 - 1/N$, where $N = |\mathcal{U}|$.*

Corollary 4.2 *In any random selection protocol Π, $\max\{\epsilon_A, \epsilon_B\} \geq 1/2 - 1/(2N)$.*

Proof: Suppose we have a protocol (A, B, f) where $\epsilon_A + \epsilon_B < 1 - 1/N$. Then we can partition the universe into two sets, S and $\mathcal{U} \setminus S$, where $|S| > \epsilon_A N$ and $|\mathcal{U} \setminus S| > \epsilon_B N$. Now the argument follows logic similar to the impossibility result of Saks [Sak89] for collective coin-flipping when at least half of the players are dishonest (where we think of outcomes in S as "heads" and $\mathcal{U} \setminus S$ as "tails").

Specifically, we view the protocol as a game where Alice wins if the output lands in S and Bob wins if the output lands in $\mathcal{U} \setminus S$. A well-known result in game theory is Zermelo's theorem: it implies that one of the players will have a winning strategy (one that wins regardless of how the other player plays). The basic reasoning is *backwards induction* on the game tree: every leaf node can be labeled A-WIN or B-WIN, depending on whether the output is in S or $\mathcal{U} \setminus S$ respectively, and then we can inductively label the remaining nodes depending on whether there exists a winning child for the current player to select. If there is, the current player can choose that child and will thus have a winning strategy from the current node. If there is not, then the opposing player can certainly win from the current node, as he or she has a winning strategy from all the children of the node.

This result implies one of the following:

- There exists strategy A^* where $\Pr[f((A^*, B^*)) \in S] = 1$, for any B^*. Taking $B^* = B$ (Bob's honest strategy), we have $\Pr[f((A^*, B)) \in S] - \mu(S) = 1 - \mu(S) > \epsilon_B$. This contradicts the guarantee of ϵ_B for Bob.
- There exists strategy B^* where $\Pr[f((A^*, B^*)) \in \mathcal{U} \setminus S] = 1$, for any A^*. This similarly contradicts guarantee ϵ_A.

■

The main intuition behind the above proof is that, at every stage, either there exists a move that is good for the current player or all moves are good for the other player. In either case, the result is good for one of the two players. All that is needed is a way to make sure that every node on the bottom level can be defined as "winning" for someone, and that this notion can propagate up the tree. As we will see, this type of reasoning will figure strongly in the proof of our main lower bound. There, the primary challenge will be to handle the cases when some nodes do not appear to be "winning" for either player.

4.2 The Main Lower Bound

In this section, we prove Theorem 1.3, giving a lower bound on round complexity matching the Iterated Random Shift Protocol up to a factor of $2 + o(1)$.

Theorem 4.3 (Theorem 1.3, strengthened) *For any $\epsilon, \mu > 0$ and $N \in \mathbb{N}$, any random selection protocol on a universe of size N satisfying the Statistical Criterion with parameters ϵ and μ requires at least $\log^* N - \log^*(\max\{\log^* N, 1/\epsilon, 1/\mu\}) - O(1)$ rounds.*

Corollary 4.4 *For every constant $\delta > 0$, there exists a constant C such that if a protocol Π achieves ϵ_A, $\epsilon_B \leq 1 - \delta$, then Π has at least $\log^* N - \log^* \log^* N - C$ rounds.*

To prove this theorem, we must show that in a protocol with "few" rounds, one of the two players will be able to find a set of small size that will contain the output with high probability. We will refer to such a set (into which the cheating player is trying to make the output fall) as a *cheating set*. The proof will rely to some degree on the probabilistic method: we will show the existence of such a cheating set by choosing it *randomly*, at least in part. The distribution on sets we will use is the following:

Definition 4.5 *For a universe $\mathcal{U}$ and a parameter $\mu \in [0, 1]$, we define* "a random subset of $\mathcal{U}$ of expected density μ" *to be a set $S \subseteq \mathcal{U}$ obtained by including each element $x \in \mathcal{U}$ in S independently with probability μ.*

Notice that the expected density of sets S chosen in this way is μ (and with high probability the density will not deviate significantly from μ).

We now can state the main helper theorem that will allow us to prove Theorem 4.3:

Theorem 4.6 *There exists a function h such that for any $\mu, \epsilon > 0, r \in \mathbb{N}$ and protocol Π with r rounds, one of the following three cases holds:*

1. *(A-EASY-WIN) When R is a randomly chosen set of expected density μ, and Alice plays a strategy maximizing the probability that the output of the protocol falls in R assuming that Bob plays honestly, she will succeed with probability $1-\epsilon$, on average over all possible R. Formally, $\mathrm{E}_R\left[\max_{A^*}\left\{\Pr_B\left[\Pi(A^*, B) \in R\right]\right\}\right] \geq 1-\epsilon$.*

2. *(B-EASY-WIN) $\mathrm{E}_R[\max_{B^*}\{\Pr_A[\Pi(A, B^*) \in R]\}] \geq 1-\epsilon$.*

3. *(DIFFICULT-WIN-WIN) When R is a randomly chosen set of expected density μ, both Alice and Bob can force the output into R plus an additional $h(r, \epsilon, \mu)$ elements with high probability. That is, the following two conditions hold:*

   (a) $\exists T$, $|T| \leq h(r, \epsilon, \mu)$, such that

   $$\mathop{\mathrm{E}}_R[\max_{A^*}\{\Pr_B[\Pi(A^*, B) \in R \cup T]\}] \geq 1-\epsilon,$$

   (b) $\exists S$, $|S| \leq h(r, \epsilon, \mu)$, such that

   $$\mathop{\mathrm{E}}_R[\max_{B^*}\{\Pr_A[\Pi(A, B^*) \in R \cup S]\}] \geq 1-\epsilon.$$

Moreover, h does not grow too fast in r. Specifically, there is a constant C such that $h(r, \epsilon, \mu) \leq \mu N$ for all $r \leq \log^* N - \log^*(\max\{\log^* N, 1/\epsilon, 1/\mu\}) - C$.

Putting the three conditions together, this theorem says that either one player can make the output fall into a random set of a certain expected density with high probability (EASY-WIN), or *both* players can make the output fall into a set consisting of a randomly chosen set of a certain expected density and a certain bounded number of (non-random) elements (DIFFICULT-WIN-WIN).

Proof of Theorem 4.3:

Let μ and ϵ be given and set $\mu' = \mu/4$, $\epsilon' = \epsilon/4$. By Theorem 4.6, we know that one of the players can force the output into a set $R \cup X$, where $|X| = h(r, \epsilon', \mu')$, with probability $1-\epsilon'$ in expectation over selecting R of expected density μ', for any protocol using r rounds. Suppose without loss of generality that the cheating player is Alice, playing with strategy A^*. Then we have:

$$\mathop{\mathrm{E}}_R[\max_{A^*}\{\Pr_B[\Pi(A^*, B) \in R \cup X]\}] \geq 1-\epsilon',$$

where R is a random set of expected density μ'. By a Chernoff bound, we have that,

$$\Pr_R[\mu(R) \geq 2\mu'] \leq e^{-2(\mu')^2 N} = e^{-\mu^2 N/8} \leq \epsilon',$$

where we may assume the last inequality holds because otherwise $\log^* N - \log^*(\max\{\log^* N, 1/\epsilon, 1/\mu\}) - C \leq 0$ for a constant C and the lower bound to be proven is trivial. But then it follows that

$$\mathop{\mathrm{E}}_R[\max_{A^*}\{\Pr_B[\Pi(A^*, B) \in R \cup X]\} \cdot I(\mu(R) < 2\mu')] \geq 1-2\epsilon',$$

where $I(\mu(R) < 2\mu')$ is the indicator random variable for the event $\mu(R) < 2\mu'$. By averaging, we can find a particular set R^*, $\mu(R^*) < 2\mu'$ such that

$$\max_{A^*}\{\Pr_B[\Pi(A^*, B) \in R^* \cup X]\}] \geq 1 - 2\epsilon' > 1 - \epsilon.$$

Assuming for contradiction that

$$r \leq \log^* N - \log^*(\max\{\log^* N, 1/\epsilon', 1/\mu'\}) - C = \log^* N - \log^*(\max\{\log^* N, 1/\epsilon, 1/\mu\}) - O(1),$$

we have $h(r, \epsilon', \mu')/N < \mu' N$ and so $|R^* \cup X| \leq 3\mu' N < \mu N$, violating the Statistical Criterion.

■

4.2.1 Proof Outline

In this section, we give an overview of the proof of Theorem 4.6. A detailed implementation is contained in Section 4.2.2. A pictorial depiction of the proof for protocols with up to 3 rounds can be found at [San05].

Proving Theorem 4.6 will require an intricate analysis of the game tree using backwards induction. Like the proof of Proposition 4.1, we will show how to "label" the nodes of the game tree, where each label corresponds to a power of a player to force a particular outcome.

The labels. We will use three labels, corresponding precisely to the three cases of a protocol in Theorem 4.6. Specifically, the labels for a node on level k of the game tree will correspond to the following (where the leaves are at level 0):

- A-EASY-WIN: Alice could from that point choose a cheating set of small (say, constant) density at random and "win" — that is, make the output fall in that set with high probability.
- B-EASY-WIN: Bob could choose a cheating set at random and win.
- DIFFICULT-WIN-WIN: Neither player can win easily by choosing a totally random set, but *both* can win by choosing a set partly at random, but also including a small (e.g. constant or very slowly growing) number of non-random elements (what we call a "helper set").

A node z labeled as DIFFICULT-WIN-WIN will have two collections of sets associated with it: $A\text{-}\mathcal{H}_z$ and $B\text{-}\mathcal{H}_z$. If the node is on level k of the game tree and it is Alice's turn to act, then $A\text{-}\mathcal{H}_z$ consists of sets of size s_{k-1} and $B\text{-}\mathcal{H}_z$ consists of sets of size s_k, where $s_1, \ldots, s_r$ will be an ascending sequence of appropriately defined constants (where r is the number of rounds of the protocol).[2] Each set $H \in A\text{-}\mathcal{H}_z$ is a set Alice can use as a "helper" — after choosing a cheating set at random and then adding the helper set, Alice can win from the given node with high probability. Similarly every set in $B\text{-}\mathcal{H}_z$ can be used by Bob as a "helper."

Our main challenge is to show that every node can be given a label as above. Once we do that, Theorem 4.6 and thus Theorem 4.3 would follow readily. As a start, notice that the leaves of the tree (the base case of our induction) can certainly be labeled DIFFICULT-WIN-WIN, simply by setting $A\text{-}\mathcal{H}_z = B\text{-}\mathcal{H}_z = \{\{x\}\}$, where x is the output of the protocol at leaf z.

[2] In the actual proof, they will be very slowly growing functions of N, but for this outline one may think of them as constants.

Piths. Before demonstrating how the internal nodes will be labeled, we define the following key concept:

Definition 4.7 *Given a collection $\mathcal{H}$ of nonempty subsets of a universe $\mathcal{U}$, the s-*pith *of $\mathcal{H}$ is the collection of all sets $S \subseteq \mathcal{U}$, $|S| \leq s$ that intersect every $H \in \mathcal{H}$ (that is, $S \cap H \neq \emptyset$). We call each such S in the pith an* intersect-set *of $\mathcal{H}$.*

Three combinatorial facts about piths and collections of disjoint sets will prove useful:

Fact 4.8 *Suppose $\mathcal{H}$ consists of nonempty sets of size at most s. Then for any s', either $\mathcal{H}$ has a disjoint subcollection of size at least s'/s or it has a nonempty s'-pith.*

Proof: Take a maximal disjoint subcollection $\mathcal{P}$ of $\mathcal{H}$. If it is not of size at least s'/s, then the union of all sets in $\mathcal{P}$ will be a set of size at most s' intersecting every set in $\mathcal{H}$ (because $\mathcal{P}$ is maximal). ■

Fact 4.9 *Suppose $\mathcal{H}$ consists of m disjoint sets of size at most s and $m \geq (1/\mu)^s \cdot \ln(1/\epsilon)$. Then the probability a random set R of expected density μ will encompass a set in $\mathcal{H}$ (i.e., there exists $H_i \in \mathcal{H}$ with $H_i \subseteq R$) is at least $1 - \epsilon$.*

Proof: The probability of failure is at most $(1 - \mu^s)^m \leq e^{-\mu^s m}$. The result follows. ■

Putting these two facts together, note that either a set in $\mathcal{H}$ is encompassed by a random set with probability $1 - \epsilon$, or $\mathcal{H}$ has a nonempty s'-pith, as long as $s' \geq (1/\mu)^s \cdot \ln(1/\epsilon) \cdot s$. Finally:

Fact 4.10 *Suppose $\mathcal{H}$ consists of sets of size at most s, and a set S intersects every set in the s'-pith of $\mathcal{H}$. Then either S encompasses a set in $\mathcal{H}$ or $\mathcal{H}' = \{H \backslash S : H \in \mathcal{H}\}$ has a disjoint subcollection of size at least s'/s.*

Proof: Say S does not encompass a set in $\mathcal{H}$. Then every set in $\mathcal{H}'$ is nonempty. If $\mathcal{H}'$ does not have a disjoint subcollection of size s'/s, then by Fact 4.8 $\mathcal{H}'$ has a nonempty s'-pith. But if a set T is in the s'-pith of $\mathcal{H}'$ then $T \backslash S$ is in the s'-pith of $\mathcal{H}$, contradicting the definition of S. ■

This strange last fact is actually an important key to the whole proof.

Labeling the Nodes. We now can describe how we will inductively label a node z on level k of the game tree, assuming it is Alice's turn at that node. First, we define the constants s_k to obey $s_k \geq (1/\mu)^{s_{k-1}} \cdot \ln(1/\epsilon) \cdot s_{k-1}$. Next:

- If all children of z are labeled B-EASY-WIN, then certainly we can give z the label B-EASY-WIN.
- If there exists a child of z that is labeled A-EASY-WIN, then we can give z the label A-EASY-WIN (Alice would just choose that child on her turn).

If neither of these cases occur, then we know that all of the children of z are either labeled DIFFICULT-WIN-WIN or B-EASY-WIN (with at least one labeled DIFFICULT-WIN-WIN). Let $\mathcal{X}$ be the union of all the collections A-$\mathcal{H}_{z_i}$, over all children z_i labeled DIFFICULT-WIN-WIN. $\mathcal{X}$ contains the helper sets that we know Alice can use to win from any such child. There are two cases:

- Suppose $\mathcal{X}$ has a large disjoint subcollection (specifically, of size at least s_k/s_{k-1}). Then label z as A-EASY-WIN.
- Else, label z as DIFFICULT-WIN-WIN, letting $A\text{-}\mathcal{H}_z$ equal $\mathcal{X}$ and letting $B\text{-}\mathcal{H}_z$ equal the s_k-pith of $\mathcal{X}$.

The correctness of the first bullet follows quickly from Fact 4.9 — with high probability a random cheating set R chosen by Alice will encompass one of the sets $X \in \mathcal{X}$ and Alice can then choose the child z_i associated with X (i.e., $X \in A\text{-}\mathcal{H}_{z_i}$) and force the output into $X \cup R = R$ with high probability.

As for the second bullet, Alice certainly has a difficult win (that is, she can win with any helper set $X \in \mathcal{X}$) because she can choose the child associated with X.

The crux of the proof is showing that Bob has a difficult win from this point, using any helper set S from the pith of $\mathcal{X}$. Note first that by Fact 4.8 and the fact that we didn't fall into the first bullet, we know this pith is nonempty.

It suffices to show that no matter what child Alice chooses, Bob can win using S — that is, force the output into $R \cup S$ with high probability where R is a random set. Certainly, if she chooses a B-EASY-WIN node, Bob can win, even without S. So suppose she chooses a child node z_i labeled DIFFICULT-WIN-WIN. Inductively, we know Bob could win from this node z_i using any of the helper sets in $B\text{-}\mathcal{H}_{z_i}$, and that the pith of $B\text{-}\mathcal{H}_{z_i}$ is $A\text{-}\mathcal{H}_{z_i} \subseteq \mathcal{X}$. Since S is in the pith of $\mathcal{X}$, we know in particular that it intersects every set in $A\text{-}\mathcal{H}_{z_i}$.

Applying Fact 4.10, we then have two cases, both of which ensure Bob has some $T \in B\text{-}\mathcal{H}_{z_i}$ encompassed by his cheating set $R \cup S$, allowing him to win:

- S encompasses a set T in $B\text{-}\mathcal{H}_{z_i}$.
- $T' = \{T \backslash S : T \in B\text{-}\mathcal{H}_{z_i}\}$ has a large disjoint subcollection (i.e., of size s_{k-1}/s_{k-2}), in which case the random set will encompass one of these sets $T \backslash S$ with high probability (by Fact 4.9), and thus $T \subseteq R \cup S$.

This completes the induction and the proof sketch. The bulk of the ideas in the main proof were demonstrated above. What remains to flesh out is proper bookkeeping of parameters, namely the randomly chosen set and the sizes of the helper sets $A\text{-}\mathcal{H}_z$ and $B\text{-}\mathcal{H}_z$ to derive the precise round complexity bound. Jumping ahead, notice that the induction will stop at $\log^* N$ rounds because the sizes of these helper sets grow as tower in the number of rounds — each s_{k+1} is exponential in s_k. Thus, if there are more than $\log^* N$ rounds, the helper sets could contain all of the elements of the universe, rendering them useless for violating the Statistical Criterion.

4.2.2 Proof of Theorem 4.6

We proceed by backwards induction on the game tree of the protocol.

Definition 4.11 *Given a protocol* Π *with* r *rounds and constants* ϵ *and* μ*, let* $h(r,\epsilon,\mu) = g(r,r,\epsilon,\mu)$*, where*

$$\begin{aligned} g(0,r,\epsilon,\mu) &= 1, \\ g(k,r,\epsilon,\mu) &= \ln(r/\epsilon) \cdot (r/\mu)^{g(k-1,r,\epsilon,\mu)} \cdot g(k-1,r,\epsilon,\mu). \end{aligned}$$

For readability, we write s_k for $g(k, r, \epsilon, \mu)$, as r, ϵ, and μ will remain fixed throughout the proof. So

$$
\begin{aligned}
s_0 &= 1, \\
s_k &= \ln(r/\epsilon) \cdot (r/\mu)^{s_{k-1}} \cdot s_{k-1}.
\end{aligned}
$$

Now, fix a protocol Π with r rounds, and consider the game tree T it induces (see Definition 2.2).

We inductively label the nodes of the tree as either A-EASY-WIN, B-EASY-WIN, or DIFFICULT-WIN-WIN. For each of the DIFFICULT-WIN-WIN nodes, we will also associate two collections of sets (which are subsets of $\mathcal{U}$): A-$\mathcal{H}_z$ and B-$\mathcal{H}_z$, as defined below. The sets in these collections will correspond to the sets S and T of case (3) of Theorem 4.6. As we will see, the labels have been chosen to indicate the power of one or both of the players to manipulate the output effectively from that point in the tree.

Definition 4.12 *Fix a protocol Π. Let z be a node on its game tree at level k (where leaves are at level 0). Assume it is Alice's turn at this node. (If it is Bob's turn, swap "A"/"A" with "B"/"B" everywhere in the description below).*

If $k = 0$ (i.e., z is a leaf of the tree) then label z as DIFFICULT-WIN-WIN. *Moreover, let B-$\mathcal{H}_z = A$-$\mathcal{H}_z = \{\{x\}\}$, where x is the output of the protocol ending at node z.*

If $k > 0$, consider the children $z_1, \ldots, z_\ell$ of z. Use the following rules to label the nodes:

1. *If there exists $1 \le i \le \ell$ such that z_i is in case* A-EASY-WIN, *then label z as* A-EASY-WIN.
2. *If, for all $1 \le i \le \ell$, z_i is in case* B-EASY-WIN, *then label z as* B-EASY-WIN.
3. *Otherwise, denote $\bigcup_{z_i} A$-$\mathcal{H}_{z_i} = \{S : z_i$ is* DIFFICULT-WIN-WIN *and $S \in A$-$\mathcal{H}_{z_i}\}$. That is, $\bigcup_{z_i} A$-$\mathcal{H}_{z_i}$ is the union of the collections of sets associated with all children of z that are labeled* DIFFICULT-WIN-WIN. *Now, let $\mathcal{P}$ denote the largest disjoint subcollection of $\bigcup_{z_i} A$-$\mathcal{H}_{z_i}$ (break ties arbitrarily) and let s_k, s_{k-1} be defined as in Definition 4.11.*

   Two cases:

   (a) *$|\mathcal{P}| \ge s_k/s_{k-1} \Rightarrow$ label z as* A-EASY-WIN.

   (b) *$|\mathcal{P}| < s_k/s_{k-1} \Rightarrow$ label z as* DIFFICULT-WIN-WIN, *and define A-$\mathcal{H}_z$ to be $\bigcup_{z_i} A$-$\mathcal{H}_{z_i}$, and B-$\mathcal{H}_z$ to be the s_k-pith of A-$\mathcal{H}_z$ (i.e., all sets of size at most s_k intersecting all sets in A-$\mathcal{H}_z$).*

Intuitively, this structure defines the power of the players at various stages of the protocol. The A-EASY-WIN, B-EASY-WIN, and DIFFICULT-WIN-WIN nodes refer to cases (1), (2), and (3) of Theorem 4.6, respectively. Moreover, the sets in the collections B-$\mathcal{H}_z$ and A-$\mathcal{H}_z$ will correspond to S and T in Case (3) of Theorem 4.6.

We will codify this power below in Lemma 4.14. Before stating it, it will help to define the following:

Definition 4.13 *Let $\Pi = (A, B, f)$ be a protocol, and let T be its equivalent game tree (see Definition 2.2). For any node $z = (m_1, \ldots, m_{r-k})$ on level k of the tree T, let $\Pi_z = (A_z, B_z, f_z)$ be a* *protocol of k rounds, where $f_z((m'_1, \ldots, m'_k)) = f((m_1, \ldots, m_{r-k}, m'_1, \ldots, m'_k))$, and where A_z and B_z denote the strategies of A and B conditioned on history z (i.e. we choose their coin tosses r_A and r_B uniformly from those consistent with the history).*

Intuitively, Π_z is the protocol induced by starting the protocol at node z (i.e., assuming all messages leading to z are fixed in advance).

Lemma 4.14 *Fix ϵ and μ, and suppose the protocol has r turns. Let z be some node on the tree at level k, at which it is Alice's turn to play. Throughout, let R be a random subset of $\mathcal{U}$ of expected density $k\mu/r$.*

1. *If z is in case* A-EASY-WIN*, then*

$$\mathop{\mathrm{E}}_{R}\left[\max_{A^*}\left\{\Pr_{B}\left[\Pi_z(A^*, B) \in R\right]\right\}\right] \geq 1 - k\epsilon/r,$$

   where Π_z is the protocol induced by beginning at node z, as defined in Definition 4.13. (We say Alice can "win" from node z).

2. *If z is in case* B-EASY-WIN*, then similarly,*

$$\mathop{\mathrm{E}}_{R}\left[\max_{B^*}\left\{\Pr_{A}\left[\Pi_z(A, B^*) \in R\right]\right\}\right] \geq 1 - k\epsilon/r.$$

   (We say Bob can "win" from node z).

3. *If z is in case* DIFFICULT-WIN-WIN*, then:*

   (a) B-$\mathcal{H}_z$ and A-$\mathcal{H}_z$ are nonempty.

   (b) For any $T \in A$-$\mathcal{H}_z$,

$$\mathop{\mathrm{E}}_{R}\left[\max_{A^*}\left\{\Pr_{B}\left[\Pi_z(A^*, B) \in R \cup T\right]\right\}\right] \geq 1 - k\epsilon/r.$$

   (c) For any $S \in B$-$\mathcal{H}_z$,

$$\mathop{\mathrm{E}}_{R}\left[\max_{B^*}\left\{\Pr_{A}\left[\Pi_z(A, B^*) \in R \cup S\right]\right\}\right] \geq 1 - k\epsilon/r.$$

   (We say both Alice and Bob "win" from node z, with "helper sets" T and S, respectively).

*Moreover, the same (with "Alice"/"A"/"*A*" exchanged for "Bob"/"B"/"*B*" respectively) holds for all nodes for which it is Bob's turn.*

Lemma 4.14 more precisely asserts Theorem 4.6 at each level of the game tree. To use this lemma to prove Theorem 4.6, we simply need to apply it with $k = r$ and z being the root of the game tree. Certainly, if z_r is in Case (1) or (2) of Lemma 4.14, it is in Case (1) or (2) of Theorem 4.6, respectively. If z_r is in Case (3) of Lemma 4.14, then subcases (3.b) and (3.c) directly prove subcases (3.a) and (3.b) respectively, where the sets in A-$\mathcal{H}_z$ and B-$\mathcal{H}_z$ of (3.a) and (3.b) correspond precisely to the sets T and S we need in those subcases of the theorem. The existence of such sets is guaranteed by subcase $(3.a)$ of the lemma.

Thus, after proving Lemma 4.14, all that will remain will be to bound the function $h(r, \epsilon, \mu)$ to prove Theorem 4.6, which we will do in Lemma 4.20.

We prove Lemma 4.14 by induction on the levels of the tree.

Base Case:

$\mathbf{k = 0}$: So z is a leaf node, and the output of Π_z is just deterministically fixed at, say, x. According to Definition 4.12, $A\text{-}\mathcal{H}_z = B\text{-}\mathcal{H}_z = \{\{x\}\}$, and we are in case DIFFICULT-WIN-WIN. Since the density of R is chosen to be zero (it is $k\mu/r$), $R = \emptyset$ and so we need to show that

$$\max_{B^*} \left\{ \Pr_A \left[\Pi_z(A, B^*) \in \{x\} \right] \right\} = 1,$$

and similarly with A and B swapped. This of course holds because the output is fixed at x.

Inductive Step. Suppose Lemma 4.14 holds for nodes on all levels up to level $k-1$. We will show that it holds for an arbitrary node z on level k. Assume it is Alice's turn at z. There are several possibilities:

Claim 4.15 *If z is in case* B-EASY-WIN, *then*

$$\mathop{\mathrm{E}}_{R} \left[\max_{B^*} \left\{ \Pr_A \left[\Pi_z(A, B^*) \in R \right] \right\} \right] \geq 1 - k\epsilon/r,$$

where R is a random subset of expected density $k\mu/r$.

Proof: We will use Definition 4.12 and the inductive hypothesis to show that every child node z_i is "good" for Bob—that is, on average over R, B^* can make the outcome land in R with probability at least $1 - (k-1)\epsilon/r$. Then certainly the same holds for node z, since Alice cannot help but move to such a node.

Formally, we first notice that it suffices to show:

$$\mathop{\mathrm{E}}_{R'} \left[\max_{B^*} \left\{ \Pr_A \left[\Pi_z(A, B^*) \in R' \right] \right\} \right] \geq 1 - (k-1)\epsilon/r,$$

where R' is a random subset of expected density $\mu' = (k-1)\mu/r$, since a random set R of expected density $k\mu/r$ "contains" such an R'. (Formally, R can be obtained by first picking R' and then adding each element $x \notin R'$ to R with probability $(\mu/r)/(1-(k-1)\mu/r)$).

Now, for z to be labeled B-EASY-WIN, we must have used Rule 2 of Definition 4.12. Thus, all of the children of z are in case B-EASY-WIN. By the inductive hypothesis:

$$\mathop{\mathrm{E}}_{R'} \left[\max_{B^*} \left\{ \Pr_A \left[\Pi_{z_i}(A, B^*) \in R' \right] \right\} \right] \geq 1 - (k-1)\epsilon/r, \tag{1}$$

for each child z_i of z. Since at node z it is Alice's turn, we have

$$\begin{aligned} \mathop{\mathrm{E}}_{R'} \left[\max_{B^*} \left\{ \Pr_A \left[\Pi_z(A, B^*) \in R' \right] \right\} \right] &= \mathop{\mathrm{E}}_{R', z_i \leftarrow \mathcal{Z}} \left[\max_{B^*} \left\{ \Pr_A \left[\Pi_{z_i}(A, B^*) \in R' \right] \right\} \right] \\ &\geq 1 - (k-1)\epsilon/r, \end{aligned}$$

where $\mathcal{Z}$ is the distribution according to which Alice chooses child z_i of z when playing honestly, and the last inequality is by (1). ■

Claim 4.16 *If z is in case* A-EASY-WIN*, then*

$$\mathop{\mathrm{E}}_{R}\left[\max_{A^*}\left\{\Pr_{B}\left[\Pi_z\left(A^*,B\right)\in R\right]\right\}\right]\geq 1-k\epsilon/r,$$

where R is a random subset of expected density $k\mu/r$.

Proof: By Definition 4.12, z could have been labeled A-EASY-WIN either by Rule 1 or Rule 3a.

In Rule 1, z has a child z_j that is in case A-EASY-WIN. Since it is Alice's turn at node z, if she can choose a node z_j "good" for her then node z will be "good" for her too. Formally, by the inductive hypothesis applied to z_j, and again noting that a random set of expected density $k\mu/r$ "contains" a random set of expected density $(k-1)\mu/r$, we have that

$$\mathop{\mathrm{E}}_{R}\left[\max_{A^*}\left\{\Pr_{B}\left[\Pi_{z_j}(A^*,B)\in R\right]\right\}\right]\geq 1-(k-1)\epsilon/r.$$

But $\max_{A^*}\left\{\Pr_B\left[\Pi_z(A^*,B)\in R\right]\right\}$ is at least $\max_{A^*}\left\{\Pr_B\left[\Pi_{z_j}(A^*,B)\in R\right]\right\}$, since node z is Alice's turn and she can always choose z_j. Taking expectations of both sides, the claim is proven for this case.

The alternative possibility is that z is in A-EASY-WIN because of Rule 3a. So among the sets $\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i}$ (for all children z_i in DIFFICULT-WIN-WIN), we can find a disjoint subcollection $\mathcal{P}$, where $|\mathcal{P}|\geq s_k/s_{k-1}$.

Intuitively, since no A-EASY-WIN nodes are available among the children of z, Alice cannot simply choose such a branch as above. However, we know that from the DIFFICULT-WIN-WIN nodes, Alice can ensure the output lands in $S\cup R$ for any $S\in A\text{-}\mathcal{H}_{z_i}$, with high probability over the choice of a random set R. But this is true for many possible sets S—not only at a given child, but also across all the potential children that are in case DIFFICULT-WIN-WIN (i.e., any $S\in\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i}$). Thus, we can expect that with sufficiently many disjoint sets in $\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i}$, the random set R will encompass some $S\in\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i}$ with high probability. The inductive hypothesis will then give the desired result.

Formally, since $\mathcal{P}\subseteq\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i}$ consists of at least $s_k/s_{k-1}=\ln(r/\epsilon)(r/\mu)^{s_{k-1}}$ (disjoint) sets of size at most s_{k-1}, Fact 4.9 tells us that:

$$\mathop{\mathrm{E}}_{R_1}\left[\exists S\in\bigcup_{z_i} A\text{-}H_{z_i}, S\subseteq R_1\right]\geq 1-\epsilon/r, \tag{2}$$

where R_1 is a random subset of expected density μ/r.

For any $S\in\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i}$, we can then assert:

$$\mathop{\mathrm{E}}_{R_2}\left[\max_{A^*}\left\{\Pr_{B}\left[\Pi_z(A^*,B)\in R_2\cup S\right]\right\}\right]\geq 1-(k-1)\epsilon/r, \tag{3}$$

where R_2 is a random subset of expected density $(k-1)\mu/r$. This comes from applying the inductive hypothesis to the child z_j such that $S\in A\text{-}\mathcal{H}_{z_j}$, and noting that $\max_{A^*}\left\{\Pr_B\left[\Pi_z(A^*,B)\in R_2\cup S\right]\right\}$ is at least $\max_{A^*}\left\{\Pr_B\left[\Pi_{z_j}(A^*,B)\in R_2\cup S\right]\right\}$ (because at node z it is Alice's turn).

Now, since a random subset R of expected density $k\mu/r$ "contains" $R_1 \cup R_2$ where R_1 and R_2 are independent random subsets of expected densities μ/r and $(k-1)\mu/r$ respectively,[3] we can combine (2) and (3) to derive

$$\begin{aligned} \mathop{\mathrm{E}}_{R}\left[\max_{A^*}\left\{\Pr_{B}\left[\Pi_z(A^*,B)\in R\right]\right\}\right] &\geq (1-\epsilon/r)\cdot(1-(k-1)\epsilon/r) \\ &\geq 1-k\epsilon/r. \end{aligned}$$

The claim follows. ■

The final possibility is that z is in case DIFFICULT-WIN-WIN. Since z is not a leaf, this can only come about by Rule 3b from Definition 4.12. That is, no children of z are in case A-EASY-WIN, and at least some are in DIFFICULT-WIN-WIN. Moreover, among $\bigcup_{z_i} A\text{-}\mathcal{H}_{z_i} = A\text{-}\mathcal{H}_z$ the largest (maximal) disjoint subcollection $\mathcal{P}$ has less than s_k/s_{k-1} elements.

We must prove the following: $B\text{-}\mathcal{H}_z$ is nonempty, $A\text{-}\mathcal{H}_z$ is nonempty, Alice can win from this node with a helper set from $A\text{-}\mathcal{H}_z$, and Bob can win from this node with a helper set from $B\text{-}\mathcal{H}_z$ (see Lemma 4.14).

Claim 4.17 *$B\text{-}\mathcal{H}_z \neq \emptyset$ and $A\text{-}\mathcal{H}_z \neq \emptyset$.*

Proof: We have already established that z has children in case DIFFICULT-WIN-WIN (this follows from Definition 4.12 and from our assumption that $z \in$ DIFFICULT-WIN-WIN). By the inductive hypothesis on such a child z_i, $A\text{-}\mathcal{H}_{z_i}$, and thus $A\text{-}\mathcal{H}_z$ is nonempty. Since the largest disjoint subcollection $\mathcal{P}$ of $A\text{-}\mathcal{H}_z$ has size less than s_k/s_{k-1} and since all $S \in \mathcal{P}$ have size at most s_{k-1}, Fact 4.8 tells us that the s_k-pith of $A\text{-}\mathcal{H}_z$ — namely, $B\text{-}\mathcal{H}_z$ — is nonempty. ■

Claim 4.18 *For any $S \in A\text{-}\mathcal{H}_z$, $\mathrm{E}_R[\max_{A^*}\{\Pr_B[\Pi_z(A^*,B) \in S \cup R]\}] \geq 1 - k\epsilon/r$ where R is a random subset of expected density $k\mu/r$.*

The proof of this claim is identical to the proof of equation (3) in the proof of Claim 4.16, noting also that a random set R of expected density $k\mu/r$ contains a random set R_2 of expected density $(k-1)\mu/r$.

The final claim required to prove Lemma 4.14 is the following:

Claim 4.19 *For any $S \in B\text{-}\mathcal{H}_z$, $\mathrm{E}_R[\max_{B^*}\{\Pr_A[\Pi_z(A,B^*) \in S \cup R]\}] \geq 1 - k\epsilon/r$ where R is a random subset of expected density $k\mu/r$.*

This claim is the heart of the entire proof. All we know now is that there is at least one DIFFICULT-WIN-WIN node that is a child of the current node z, and that among the corresponding sets in $A\text{-}\mathcal{H}_z$, the largest disjoint subcollection $\mathcal{P} \subseteq A\text{-}\mathcal{H}_z$ contains fewer than s_k/s_{k-1} sets. That $\mathcal{P}$ is so small is a limitation on the power of Alice, who would like there to be enough such disjoint sets in $\mathcal{P}$ that she could choose randomly and encompass a set in $\mathcal{P}$ with high probability. The key to this proof is converting this *limitation* on Alice into an *ability* for Bob to cheat.

[3]Formally, R can be obtained by first picking R_1 and R_2 and adding each element $x \notin R_1 \cup R_2$ to $R_1 \cup R_2$ with probability $(\mu/r) \cdot ((k-1)\mu/r)/[(1-\mu/r)\cdot(1-(k-1)\mu)/r]$.

Proof: Fix a set $S \in B\text{-}\mathcal{H}_z$, which recall is the s_k-pith of $A\text{-}\mathcal{H}_z$. Since an honest Alice will choose a child z_i at random, it suffices to prove the following for each child z_i:

$$\mathop{\mathrm{E}}_{R}\left[\max_{B^*}\left\{\Pr_{A}\left[\Pi_{z_i}(A, B^*) \in S \cup R\right]\right\}\right] \geq 1 - k\epsilon/r, \tag{4}$$

where R is a random subset of expected density $k\mu/r$. So fix an arbitrary child z_i. Looking to Definition 4.12, the only way we could have defined z to be in case DIFFICULT-WIN-WIN is if all children z_i are either in case B-EASY-WIN or case DIFFICULT-WIN-WIN. So z_i is in one of these two cases.

If z_i is in case B-EASY-WIN, then we are done by the inductive hypothesis. So suppose z_i is in case DIFFICULT-WIN-WIN. Applying the inductive hypothesis to z_i, we know that $B\text{-}\mathcal{H}_{z_i}$ is nonempty. Moreover, for any $T \in B\text{-}\mathcal{H}_{z_i}$:

$$\mathop{\mathrm{E}}_{R_1}\left[\max_{B^*}\left\{\Pr_{A}\left[\Pi_{z_i}(A, B^*) \in T \cup R_1\right]\right\}\right] \geq 1 - (k-1)\epsilon/r, \tag{5}$$

where R_1 is a random subset of expected density $(k-1)\mu/r$.

We have that S is in the s_k-pith of $A\text{-}\mathcal{H}_z$, which means in particular it intersects every set in $A\text{-}\mathcal{H}_{z_i}$, which in turn is the s_{k-1}-pith of $B\text{-}\mathcal{H}_{z_i}$ (whose sets are of size s_{k-2}, when $k > 1$). By Fact 4.10, either there exists a set T in $B\text{-}\mathcal{H}_{z_i}$ such that $T \subseteq S$ (in which case (4) follows immediately from (5)), or else $\mathcal{T} = \{T \backslash S : T \in B\text{-}\mathcal{H}_{z_i}\}$ has a disjoint subcollection of size s_{k-1}/s_{k-2}. (When $k = 1$, $B\text{-}\mathcal{H}_{z_i}$ contains only the set $T = \{x\}$, where x is the output of the protocol at leaf z_i, and we also have $x \in S$ because S intersects every set in $A\text{-}\mathcal{H}_{z_i} = \{\{x\}\}$. So we have $T \subseteq S$, and (4) follows immediately from (5).)

Informally, there aren't many disjoint sets in $B\text{-}\mathcal{H}_{z_i}$—if there were, we would have labeled z_i as a case B-EASY-WIN node for Bob. That said, by intersecting every (small) set that intersected every set in $B\text{-}\mathcal{H}_{z_i}$, S captures the lack of disjointness of $B\text{-}\mathcal{H}_{z_i}$ in the first place. Once the elements of S are removed from consideration, the result has a large number of disjoint sets.

Returning to the proof of Claim 4.19, by Fact 4.9 we may conclude the following:

$$\Pr_{R_2}\left[\exists T' \in \mathcal{T}, T' \subseteq R_2\right] \geq 1 - \epsilon/r,$$

where R_2 is a random subset of expected density μ/r. By the definition of $\mathcal{T}$, this in turn implies:

$$\Pr_{R_2}\left[\exists T \in B\text{-}H_{z_i}, T \subseteq S \cup R_2\right] \geq 1 - \epsilon/r.$$

Using (5) and choosing R through independent choices of R_1 and R_2 as in the proof of Claim 4.16, we are done. ■

Taking together Claims 4.15, 4.16, 4.17, 4.18, and 4.19, the proof of Lemma 4.14 is complete. ■

To conclude Theorem 4.6, it remains to prove the function h defining the set sizes s_k does not grow too fast in the number of rounds. Intuitively, the reason the lower bound only holds for protocols with fewer than $\log^* N - \log^* \log^* N - O(1)$ rounds is that these "helper sets" must have much fewer than N elements to be useful, but this function h grows as a tower—where the height (and base) of the tower grow with the number of rounds. Our challenge is to lower bound the number of rounds that keep this tower of size $o(N)$.

Lemma 4.20 *Recall the definition* $h(r, \epsilon, \mu) = g(r, r, \epsilon, \mu)$*, where we define*

$$\begin{aligned} g(0, r, \epsilon, \mu) &= 1 \\ g(k, r, \epsilon, \mu) &= \ln(r/\epsilon) \cdot (r/\mu)^{g(k-1,r,\epsilon,\mu)} \cdot g(k-1, r, \epsilon, \mu). \end{aligned}$$

There exists a constant C *such that when* $r < \log^* N - \log^*(\max\{\log^* N, 1/\epsilon, 1/\mu\}) - C$*, we have* $h(r, \epsilon, \mu) \leq \mu N$*.*

Proof: First, bound r by $\log^* N$. Again, for shorthand we will write s_k for $g(k, r, \epsilon, \mu)$. Thus, we have that

$$s_k = \ln(r/\epsilon) \cdot (r/\mu)^{s_{k-1}} \cdot s_{k-1}.$$

Notice that this is no more than $(r \ln(r/\epsilon)/\mu)^{s_{k-1}}$. ($xy \leq x^y$ if $x \geq 2$ and $y \geq 1$.) Letting $d = (r \ln(r/\epsilon)/\mu)$, we can then bound s_k by t_k, where t_k is defined by $t_0 = 1$ and $t_k = d^{t_{k-1}}$.

This means that we can set $k = \log_d^* N - 1$ (recall that by our definition, $\log_b^* N$ is always an integer, for any b or N) and still have $s_k \leq t_k \leq \log N \leq \mu N$, where we may assume that the last inequality holds, because otherwise $\log^* N - \log^*(1/\mu) - C < 0$ and the lemma is vacuously true. It only remains to relate this to a base 2 logarithm:

Claim 4.21 *If* $d \geq 4$*,* $\log_d^* N \geq \log^* N - \log^*(2 \log d)$*.*

Proof: Recall $\log^{(k)} N$ is k iterated logarithms of N. We claim the following:

Claim 4.22 *For* $k \leq \log_d^* N$*,* $d \geq 4$*,* $\log^{(k)} N \leq (2 \log d) \log_d^{(k)} N$*.*

Proof: The base case $k = 0$ is clear. Assume, then, that $\log^{(k-1)} N \leq (2 \log d) \log_d^{(k-1)} N$. Applying log to both sides, we have that:

$$\begin{aligned} \log^{(k)} N &\leq \log(2 \log d) + \log(\log_d^{(k-1)} N) \\ &\leq \log(2 \log d) + (\log_d^{(k)} N)(\log d) \\ &\leq (2 \log d)(\log_d^{(k)} N), \end{aligned}$$

where the last line follows because for $d \geq 4$, $d \geq 2 \log d$ and for $k \leq \log_d^* N$, $\log_d^{(k)} N \geq 1$. ■

Plugging in $k = \log_d^* N$, then we have that $\log^{(\log_d^* N)} N \leq 2 \log d$. Applying $\log^*(2 \log d)$ logarithms to both sides, we have $\log^{(\log_d^* N + \log^*(2 \log d))} N \leq 1$. Since $\log^* N$ is defined to be the least k such that $\log^{(k)} N \leq 1$, it follows that $\log^* N \leq \log_d^* N + \log^*(2 \log d)$. ■

Thus we have that we can set k to be at least $\log_d^* N - 1$ and s_k will be no more than $\log N$. Moreover,

$$\begin{aligned} \log_d^* N - 1 &\geq \log^* N - \log^*(2 \log d) - 1 \\ &= \log^* N - \log^*(2 \log((r \ln(r/\epsilon))/\mu)) - 1 \\ &\geq \log^* N - \log^*(\max\{\log^* N, 1/\epsilon, 1/\mu\}) - O(1). \end{aligned}$$

■

By applying Lemma 4.14 to the root of the tree and using Lemma 4.20, we prove Theorem 4.6 and thus Theorem 4.3.

5 Multiplicative Lower Bounds

In this section, we concentrate on lower bounds regarding multiplicative guarantees–and indeed show that no protocol exists that provides constant multiplicative guarantees to both players. This is a very strong limitation on the ability of protocols to limit cheating player's power in this regard.

An Initial Lower Bound. Proposition 4.1 can be adapted to provide a quick lower bound for multiplicative guarantees. Specifically:

Proposition 5.1 *In any random selection protocol,* $(\rho_A - 1)/\rho_A + (\rho_B - 1)/\rho_B \geq 1 - 1/N$. *Moreover,* $\epsilon_A + (\rho_B - 1)/\rho_B \geq 1 - 1/N$ *(or equivalently,* $\epsilon_A \geq 1/\rho_B - 1/N$*).*

Proof: The results follow immediately from Theorem 4.1 and from the second part of Lemma 2.11. ■

This lower bound for multiplicative guarantees is not very strong—ρ is a number from 1 to N, but this lower bound is satisfied (for instance) as long as both ρ_A and ρ_B are at least 2. In Theorem 5.3 we will prove that $\rho_A \rho_B \geq N$, which is a substantially stronger result.

On the other hand, when looking at one player getting a statistical guarantee and the other player getting a multiplicative guarantee, Proposition 5.1 does provide some useful information. Specifically, it tells us that (minus a small $1/N$ term) we can always expect the statistical guarantee for one player to be worse than the reciprocal of the multiplicative guarantee to the other player. This explains inverse relationships in existing protocols of [DGW94] (where $\varepsilon = 1/\mathrm{poly}(n)$ and $\rho = \mathrm{poly}(n)$) and [GSV98] (where $\varepsilon = \mathrm{poly}(n) \cdot 2^{-k}$ and $\rho = 2^k$ for any k).[4] Notice that these earlier works focus on the case of nonconstant guarantees ($\varepsilon \to 0$ and $\rho \to \infty$). Earlier, we showed that the Iterated Random Shift Protocol achieves simultaneous *constant* statistical and multiplicative guarantees. From Theorem 4.3 and Lemma 2.11, it follows that our protocol has optimal round complexity up to a factor of $2 + o(1)$ among those achieving simultaneous constant statistical and multiplicative guarantees:

Corollary 5.2 *For every two constants* $\epsilon_A < 1$ *and* ρ_B*, there exists a constant* C *such that any protocol* Π *selecting from a universe of size* N *and achieving statistical guarantee* ϵ_A *and multiplicative guarantee* ρ_B *will have at least* $\log^* N - \log^* \log^* N - C$ *rounds. Similarly for* ρ_A *and* ϵ_B*.*

A Tight Lower Bound. Despite the inability of the above to give a strong lower bound for simultaneous multiplicative guarantees, in this section we present a tight lower bound in this setting, which follows from the work of Goldreich, Goldwasser, and Linial [GGL98].

Theorem 5.3 (Theorem 2.10, restated) [GGL98] *For any protocol* Π, $\rho_A \cdot \rho_B \geq N$.

Corollary 5.4 *In any protocol* Π, $\max\{\rho_A, \rho_B\} \geq \sqrt{N}$.

[4] Actually, the protocol of [GSV98] does not provide a multiplicative guarantee of 2^k, but rather ensures that the probability that the output lands in any set T of density μ is at most $2^k \cdot \mu + o(1)$. Our lower bound also applies to this more general type of guarantee.

Recalling the multiplicative guarantee ρ_A is the greatest factor by which Bob can improve the probability that a single element is chosen over uniform, we conclude:

In any random selection protocol, at least one of the players can improve the probability that a single element is chosen by a factor *exponential* in the length of the output (which equals $\log N$).

Goldreich et. al. [GGL98] showed a more general result than Theorem 5.3 (for multiparty protocols) using different language and a moderately involved proof. We include below the simplification of their proof for the two-party case.

Proof of Theorem 5.3: Fix some element v of the universe. Now, consider the game tree T of the protocol (see Definition 2.2). At each node z of the tree, denote the protocol induced by beginning at node z to be Π_z. Then define:

$$\begin{aligned}\phi_A^z &= \max_{A^*} \Pr_B[\Pi_z((A^*, B)) = v], \\ \phi_B^z &= \max_{B^*} \Pr_A[\Pi_z((A, B^*)) = v], \\ p_z &= \Pr_{A,B}[\Pi_z((A, B)) = v].\end{aligned}$$

That is, ϕ_A^z (resp. ϕ_B^z) is the highest probability Alice (resp. Bob) can make the the output to be v, given that the protocol is now at node z and that Bob (resp. Alice) is playing honestly. p_z is the probability that v is chosen starting from z and assuming both player play honestly.

To prove the theorem, we will show that for every node z on T, $\phi_A^z \cdot \phi_B^z \geq p_z$. Applying this fact to the root r of the tree r and noting that we can choose v so that $p_r \geq 1/N$, the theorem follows easily.

We will prove $\phi_A^z \cdot \phi_B^z \geq p_z$ by backwards induction on the levels of the tree.

When z is a leaf, the protocol is complete. If v is the output of the protocol at leaf z, then $\phi_A^z = \phi_B^z = p_z = 1$. Otherwise, $\phi_A^z = \phi_B^z = p_z = 0$.

Now, suppose that the lemma holds for all children of z—denote them $z_1, \dots z_m$. Thus, we know $\phi_A^{z_i} \phi_B^{z_i} \geq p_{z_i}$ for all children z_i. Suppose also, without loss of generality, that at node z it is Alice's turn.

Suppose an honest Alice chooses child node z_i with probability λ_i. Then $p_z = \sum \lambda_i p_{z_i}$, and $\phi_B^z = \sum \lambda_i \phi_B^{z_i}$. When considering ϕ_A^z, however, Alice will cheat and choose the best child available. Thus, $\phi_A^z = \max_{z_i} \phi_A^{z_i}$, and so in particular for all i, $\phi_A^z \geq \phi_A^{z_i}$.

Now just compute:

$$\phi_A^z \phi_B^z = \phi_A^z \sum \lambda_i \phi_B^{z_i} \geq \sum \lambda_i \phi_A^{z_i} \phi_B^{z_i} \geq \sum \lambda_i p_{z_i} = p_z.$$

■

To understand this result intuitively, suppose that there were only one path down the tree that led to v being chosen as the output. At each node along that path, starting from the root, there is a certain probability that an honest player will choose the (unique) next node in the path. So the probability that v is chosen is the product of these probabilities when both players play honestly. If Alice (resp. Bob) is cheating, then the probability v will be chosen is the product of the probabilities at nodes where it's Bob's (resp. Alice's) turn. In this case, $\phi_A^z \cdot \phi_B^z = p_z$. More paths yielding v merely provide more options to the cheating player, and so $\phi_A^z \cdot \phi_B^z \geq p_z$.

Note that, unlike Theorem 4.1, this result relies centrally on the assumption that, when one player is cheating, the other player is playing *honestly*.

Theorem 5.3 is, in fact, tight. In Proposition 3.10 we noted that the Random Set Protocol — in which Alice chooses a uniformly random subset of fixed size K and Bob chooses a random element of this set — achieves multiplicative guarantees of K and N/K for the two players.

Note that one negative aspect of the Random Set Protocol is that it is not efficient—sending a description of the random subset requires communication linear in N (rather than polylog(N)). This is certainly not necessary to achieve $\rho_A \rho_B = N$, however: other very simple and efficient protocols achieve this tradeoff. Specifically, instead of using all sets of size K, we can use any subcollection such that every element of $[N]$ is contained in the same number of sets. For example, if $N = K \cdot L$ for an integer L, then we can view the universe as $[K] \times [L]$ and use only the sets of the form $S_a = [K] \times \{a\}$ for each $a \in [L]$, so the communication becomes $\log L + \log K = \log N$. The optimality of such a trivial protocol suggests that, ultimately, multiplicative guarantees are not by themselves likely to be sufficient metrics of study for two-party random selection protocols.

Acknowledgements

We thank Oded Goldreich, Grant Schoenebeck, and Alexander Healy for several helpful discussions. We are also grateful to anonymous reviewers for suggestions that greatly improved the presentation.

References

[AN93] N. Alon and M. Naor. Coin-flipping games immune against linear-sized coalitions. *SIAM J. Computing* 22(2), 1993.

[Amb04] A. Ambainis. A new protocol and lower bounds for quantum coin flipping. *Journal of Computer and System Sciences* 68 (2), 398–416, 2004.

[BL89] M. Ben-Or and N. Linial. Collective coin-flipping. *Randomness and Computation*, S. Micali ed., Academic Press, New York, 1989.

[Blu82] M. Blum. Coin flipping by telephone. In *Proc. 24th COMPCOM*, 1982.

[BN98] R. Boppana and B. Narayanan. Perfect information leader election with optimal resilience. *SIAM J. Computing* 29(4), 1998.

[CCM98] C. Cachin, C. Crépeau, and J. Marcil. Oblivious transfer with a memory-bounded receiver. In *Proc. 39th FOCS*, 1998.

[Cle86] R. Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proc. 18th STOC*, pp. 364–369, 1986.

[CI93] R. Cleve and R. Impagliazzo. Martingales, collective coin flipping and discrete control processes. Unpublished manuscript, 1993.

[Dam94] I. Damgård. Interactive hashing can simplify zero-knowledge protocol design without computational assumptions. In *Proc. CRYPTO '93*, Springer LNCS 403, 1994.

[DGW94] I. Damgård, O. Goldreich, and A. Wigderson. Hashing functions can simplify zero-knowledge protocol design (too). TR RS-94-39. BRICS, 1994.

[DHRS04] Y. Ding, D. Harnik, A. Rosen, and R. Shaltiel. Constant-round oblivious transfer in the bounded storage model. In *Proc. 1st TCC*, Springer LNCS 2951, 2004.

[Fei99] U. Feige. Noncryptographic selection protocols. In *Proc. 40th FOCS*, 1999.

[GGL98] O. Goldreich and S. Goldwasser and N. Linial. Fault-tolerant computation in the full information model. *SIAM J. Computing* 27(2), 1998.

[GMW91] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity or all languages in **NP** have zero-knowledge proof systems. *J. ACM* 38(1), 1991.

[GSV98] O. Goldreich, A. Sahai, and S. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *Proc. 30th STOC*, 1998.

[KO04] J. Katz and R. Ostrovsky. Round-optimal secure two-party computation. In *Proc. CRYPTO '04*. Springer LNCS 3152, 2004.

[Lau83] C. Lautemann. **BPP** and the polynomial hierarchy. *Information Processing Letters* 14, 1983.

[Lin01] Y. Lindell. Parallel coin-tossing and constant-round secure two-party computation. In *Proc. CRYPTO '01*. Springer LNCS, 2001.

[NOVY98] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect zero-knowledge arguments for **NP** can be based on general complexity assumptions. *J. Cryptology* 11, 1998.

[ORV94] R. Ostrovsky, S. Rajagopalan, U. Vazirani. Simple and efficient leader election in the full information model. In *Proc. 26th STOC*, 1994.

[RSZ99] A. Russell, M. Saks, and D. Zuckerman. Lower bounds for leader election and collective coin-flipping in the perfect information model. In *Proc. 31st STOC*, 1999.

[RZ98] A. Russell and D. Zuckerman. Perfect information leader election in $\log^* n + O(1)$ rounds. In *Proc. 39th FOCS*, 1998.

[Sak89] M. Saks. A robust noncryptographic protocol for collective coin-flipping. *SIAM J. Discrete Math*, 1989.

[San04] S. Sanghvi. A study of two-party random selection protocols. Undergraduate thesis. Harvard University, 2004.

[San05] S. Sanghvi. The round complexity of two-party random selection. Slides of presentation given at STOC 2005. Available at http://eecs.harvard.edu/˜salil/papers/randsel-abs.html.

[SV05] S. Sanghvi and S. Vadhan. The round complexity of two-party random selection. In *Proc. 37th STOC*, 2005.

[Yao86] A. Yao. How to generate and exchange secrets. In *Proc. 27th FOCS*, 1986.