Random Selection with an Adversarial Majority*

Ronen Gradwohl[1**], Salil Vadhan[2***], and David Zuckerman[3†]

[1] Department of Computer Science and Applied Math, Weizmann Institute of Science
ronen.gradwohl@weizmann.ac.il
[2] Division of Engineering & Applied Sciences, Harvard University
salil@eecs.harvard.edu
[3] Department of Computer Science, University of Texas at Austin
diz@cs.utexas.edu

Abstract. We consider the problem of random selection, where p players follow a protocol to jointly select a random element of a universe of size n. However, some of the players may be adversarial and collude to force the output to lie in a small subset of the universe. We describe essentially the first protocols that solve this problem in the presence of a dishonest majority in the full-information model (where the adversary is computationally unbounded and all communication is via non-simultaneous broadcast). Our protocols are nearly optimal in several parameters, including the round complexity (as a function of n), the randomness complexity, the communication complexity, and the tradeoffs between the fraction of honest players, the probability that the output lies in a small subset of the universe, and the density of this subset.

1 Introduction

Suppose p players wish to jointly make a random choice from a universe of size n. They follow some protocol, and if all parties play honestly, the output is indeed a uniformly random one. However, some of the players may form a coalition and deviate arbitrarily from the protocol, in an attempt to force some output. The problem of random selection is that of designing a protocol in which the influence of coalitions of dishonest players is somehow limited.

Random selection is a very useful building block for distributed algorithms and cryptographic protocols, because it allows one to first design protocols assuming a public source of randomness, which is often an easier task, and then

* The full version of this paper is available on *ECCC* [22].

** Research supported by US-Israel Binational Science Foundation Grant 2002246.
*** Supported by US-Israel BSF grant 2002246, NSF grants CNS-0430336 and CCR-0133096, and ONR Grant N00014-04-1-0478.
† Most of this work was done while visiting Harvard University, and was supported in part by a Radcliffe Institute for Advanced Study Fellowship, a John Simon Guggenheim Memorial Foundation Fellowship, a David and Lucile Packard Fellowship for Science and Engineering, and NSF Grant CCR-0310960.

replace public randomness with the output of a random selection protocol. Of course, for this to work, there must be a good match between the guarantees of the random selection protocol and the requirements of the application at hand. Nevertheless, this general paradigm has been applied successfully numerous times in the past in various settings, e.g., [42, 20, 17, 32, 12, 13, 30, 21, 25, 3, 23]. This motivates a systematic study of random selection in its own right, like the one we undertake in this paper.

The Setting. The problem of random selection has been widely studied in a variety of settings, which differ in the following respects:

Adversary's Computational Power. In some work on random selection, such as Blum's 'coin-tossing over the telephone' [4], the adversary is assumed to be computationally bounded (e.g., probabilistic polynomial time). Generally, in this setting one utilizes one-way functions and other cryptographic primitives to limit the adversary's ability to cheat, and thus the resulting protocols rely on complexity assumptions. In this paper, we study the *information-theoretic* setting, where the adversary is computationally unbounded (so complexity assumptions are useless).

Communication Model and the Adversary's Information. There is a choice between having point-to-point communication channels, a broadcast channel, or both. In the case of point-to-point communication, one can either assume private channels, as in [6, 10], or allow the adversary full access to all communication, as in the *full-information model* of Ben-Or and Linial [7]. We allow a broadcast channel and work in the full-information model (so there is no benefit from point-to-point channels). We do not assume simultaneous communication, and thus consider a 'rushing' adversary, which can send its messages in a given round after receiving those of the honest players.

Number of Players. There has been work specifically studying two-party protocols where one of the players is adversarial; examples in the full-information model include [17, 39]. Other works study p-player protocols for large p, such as the large body of work on collective coin-flipping (random selection where the universe is of size $n = 2$) and leader election [7, 38, 1, 11, 31, 8, 43, 37, 15]. In this paper, we focus on the latter setting of *p-player protocols*, but some of our results are significant even for $p = 2$.

To summarize, here we study general multiparty protocols for random selection in the full-information model (with a broadcast channel). This is the first work in this setting to focus on the case that a majority of the players may be dishonest.[4] It may be surprising that protocols exist for this case, as the other two other well-studied problems in this setting, leader election and collective coin-flipping, are provably impossible to solve with an adversarial majority [38].

[4] We note that dishonest majorities have been studied extensively in the settings of computationally bounded parties and private channels, both for Byzantine agreement and secure computation, e.g., [20, 26, 19].

The Goal: Construct p-player protocols for selecting an element of $[n]$ such that even if a β fraction of players are cheating, the probability that the output lands in any small subset of $[n]$, of density μ, is at most ε.

Particular applications of random selection protocols often have special additional requirements, such as "simulatability." However, all of the existing work on random selection with information-theoretic security, such as [7, 38, 17, 28, 1, 12, 13, 21, 37, 15, 14, 39], seem to include at least some variant of our requirement above. Thus it is of interest to understand this requirement on its own, in particular the tradeoffs between the parameters p, n, β, μ, and ε, as well as the efficiency of protocols meeting the requirement.

As these five parameters vary, we have a very general class of problems, which includes many previously studied problems as special cases (See Section 2.2.). Some natural settings of parameters are n being exponentially large in the security parameter (e.g., choosing a random k-bit string), p being constant or polynomial, β being a constant in $(0, 1)$ (we are particularly interested in $\beta \geq 1/2$), and μ, ε either being constants in $(0, 1)$ or tending to zero.

Regarding protocol efficiency, we focus primarily on information-theoretic measures, such as the communication and round complexities, but we also provide some computationally efficient versions of our protocols.

Our Results. In this paper, we give several protocols for random selection that tolerate an arbitrarily large fraction of cheating players $\beta < 1$. The protocols are nearly optimal in many of the parameters, for example:

- One of our protocols achieves an error probability of $\varepsilon = \tilde{O}(\mu^{1-\beta})$, when the number of players is constant and the density μ of bad outcomes is arbitrary. This comes close to the lower bound of $\varepsilon \geq \mu^{1-\beta}$ proven by Goldreich, Goldwasser, and Linial [17]. For a nonconstant number of players, we can come polynomially close to the lower bound, achieving $\varepsilon = \mu^{\Omega(1-\beta)}$, provided that the fraction β of cheating players is bounded away from 1.
- One of our protocols can handle any density μ of bad outcomes that is smaller than the fraction $\alpha = 1 - \beta$ of honest players while achieving an error probability ε that is bounded away from 1. More generally, we can handle any constants α, μ such that $\lfloor 1/\alpha \rfloor \leq \lceil 1/\mu \rceil - 1$, which is a tight tradeoff by a lower bound of Feige [15].
- In our protocols, the total number of coins tossed by the honest parties is $\log n + o(\log n)$ (when the other parameters are constant), which almost equals the lower bound of $\log n - O(1)$. As the only bits communicated in our protocols are the random coin tosses, the communication complexity is also nearly optimal.
- As a function of n, the round complexity of our protocols is $\log^* n + O(1)$ (when the other parameters are constant). This is within a factor of essentially 2 of the $(1/2 - o(1)) \log^* n$ lower bound proven by Sanghvi and Vadhan [39], which applies whenever $\beta \geq 1/2$, and $\mu > 0$ and $\varepsilon < 1$ are constants.

Techniques. Our protocols build upon recent work on round-efficient leader election [37, 15] and round-efficient two-party random selection [39]. Specifically, the leader election protocols of Russell and Zuckerman [37] and Feige [15] work by iterating a one-round protocol that reduces the task of electing a leader from p players to that of electing from polylog(p) players. Similarly, the two-party random selection protocol of Sanghvi and Vadhan [39] utilizes a one-round protocol that reduces selecting from a universe of size n to selecting from one of size polylog(n). We combine these approaches, iteratively reducing both the number of players and the universe size in parallel. To do this, we construct new one-round universe reduction protocols that work for many parties (instead of just two, as in [39]). We obtain these by establishing a connection between randomness extractors [29] (or, equivalently, randomness-efficient samplers) and universe reduction protocols. Optimizing parameters of the underlying extractors then translates to optimizing parameters of the universe reduction protocols, resulting in the near-optimal bounds we achieve in our final protocols. Our main results, as outlined above, refer to protocols that use optimal extractors, as proven to exist via the probabilistic method, and thus are not explicit or computationally efficient. In the full version of this work [22], we also give computationally efficient versions of our protocols, using some of the best known explicit constructions of extractors. Any additional deficiencies in these protocols are due to limitations in the state-of-the-art in constructing extractors, which we view as orthogonal to the issues we study here. Indeed, if the loss turns out to be too much for some application, then that would provide motivation for further research on explicit constructions of extractors.

Organization. Section 2 includes definitions, a more detailed description of previous work and how it relates to this paper, and our results. Section 3 contains the one-round selection protocols that are the final ingredient in our protocols, and in Section 4 we give protocols that reduce the number of players and the size of the universe. In Section 5 we informally describe how the different pieces fit together to form our final protocols, and defer details and formal proofs from this section to the full version [22]. Finally, in Section 6, we state known and new lower bounds on the various parameters of random selection.

2 Definitions and Results

2.1 Random Selection Protocols

We now define random selection protocols, the model, and the complexity measures in which we are interested.

A (p, n)-*selection protocol* is a p-player protocol for selecting an element of $[n]$. In each round of the protocol, the players broadcast messages that they may base on the messages sent by all players in previous rounds, as well as their own internal coin tosses. The players may not legally base their outputs in round i on the outputs of other players in round i. However, since we can not guarantee simultaneity within a round, we allow the dishonest players to base

their outputs on the outputs of other players from the same round (but not from later rounds). This is known as rushing. At the end, a predetermined function of all sent messages is computed, outputting an element of $[n]$.

Given this definition, we have the following notion of security.

Definition 1. *A* (p, n)-selection protocol *is called* $(\beta, \mu, \varepsilon)$-resilient *if when at most a* β *fraction of players are cheating and* S *is any subset of* $[n]$ *of density at most* μ*, the probability that the output lands in* S *is at most* ε*. We refer to* S *as a* bad set*.*

We will be interested in the asymptotic behavior of protocols, so when we discuss (p, n)-selection protocols, we are implicitly referring to a family of protocols, one for each value of p and n (or some infinite set of pairs (p, n)). We are then interested in optimizing a variety of complexity measures:

The *computation time* of a (p, n)-selection protocol is the maximum total time spent by all (honest) players (to compute their messages and the final function) in an execution of the protocol. We call a protocol *explicit* if its computation time is $\text{poly}(\log n, p)$. The *round complexity* is the total number of rounds of the protocol. The *randomness complexity* of a protocol is the (maximum possible) total number of random bits used by the honest players. [5] (Typically this maximum is achieved when all players are honest.) The *communication complexity* of a protocol is the total number of bits communicated by the honest players.[6]

All our protocols are public-coin, in the sense that the honest players flip their random coins and broadcast the results. Thus, the communication complexity is equal to the randomness complexity. By convention, we assume that if a player sends a message that deviates from the protocol in some syntactically obvious way (e.g. the player outputs more bits than requested), then its message is replaced with some canonical string of the correct form (e.g. the all-zeroes string).

2.2 Previous Work

We now discuss the relationship of the above definitions, specifically of $(\beta, \mu, \varepsilon)$-resilient (p, n)-selection protocols, to existing notions and results in the literature.

Two-Party Random Selection. This is the special case where $p = 2$ and $\beta = 1/2$, and attention in previous work has focused on the tradeoff between μ and ε as well as the round complexity. Specifically,

- Goldreich, Goldwasser, and Linial [17] constructed, for every $n = 2^i$, an explicit $(2, n)$-selection protocol that is $(1/2, \mu, O(\sqrt{\mu}))$-resilient for every

[5] Actually, it will be convenient to allow the players to pick elements uniformly at random from $\{1, \ldots, m\}$ where m is determined during the protocol and may not be a power of 2, and in such a case we view this as costing $\log_2 m$ random bits.

[6] As with randomness complexity, it will be convenient to allow players to send elements of $\{1, \ldots, m\}$, in which case we charge $\log_2 m$ bits of communication.

$\mu > 0$. The protocol takes $2 \log n$ rounds. They also prove that the bound of $\varepsilon = O(\sqrt{\mu})$ is tight (as a special case of a more general result mentioned later).
- Sanghvi and Vadhan [39] constructed, for every constant $\delta > 0$ and every n, an explicit $(2, n)$-selection protocol that is $(1/2, \mu, O(\sqrt{\mu + \delta}))$-resilient for every $\mu > 0$. Their protocol takes $\log^* n + O(1)$ rounds. They also prove that $(\log^* n - \log^* \log^* n - O(1))/2$ rounds are necessary for any $(2, n)$-selection protocol that is $(1/2, \mu, \varepsilon)$-resilient for constants $\mu > 0$ and $\varepsilon < 1$.

Collective Coin-Flipping [7]. This is the special case when $n = 2$ and $\mu = 1/2$. Attention in the literature has focused on constructing efficient protocols that are $(\beta, 1/2, \varepsilon)$-resilient where β and ε are constants (independent of p), β is as large as possible, and $\varepsilon < 1$. Such a protocol exists for every constant $\beta < 1/2$ [8] and can be made explicit [43]. Conversely, it is impossible to achieve $\beta = 1/2$ and $\varepsilon < 1$ [38]. Efficient constructions of such protocols have been based on leader election (described below).

Leader Election.

Definition 2. *A* p-player leader election protocol *is a* (p, p)*-selection protocol. It is* (β, ε)-resilient *if when at most a* β *fraction of players are cheating, the probability that the output is the index of a cheating player is at most* ε*.*

- Every $(\beta, \beta, \varepsilon)$-resilient (p, p)-selection protocol is a (β, ε)-resilient p-player leader election protocol. The converse does not hold because the the former considers *each* subset $S \subset [p]$ of density at most β as a potential bad set of outcomes, but the latter only considers the subset consisting of the cheating players.
- Nevertheless, a p-player leader election protocol can be used to construct a (p, n)-selection protocol for any n by having the elected leader choose a uniform, random element of $[n]$ as the output. If the election protocol is (β, ε)-resilient, then the resulting selection protocol will be $(\beta, \mu, \varepsilon + (1 - \varepsilon) \cdot \mu)$-resilient for every $\mu \geq 0$.
- By the impossibility result for collective coin-flipping mentioned above [38] and the previous bullet, it is impossible to have an election protocol that is (β, ε)-resilient for $\beta = 1/2$ and $\varepsilon < 1$.
- A long line of work [1, 11, 31, 43, 37, 15] on optimizing the resilience and round complexity for leader election has culminated in the following result of Russell and Zuckerman [37].[7] For every constant $\beta < 1/2$, there exists an $\varepsilon < 1$ such that for all p, there is an explicit (β, ε)-resilient p-player leader election protocol of round complexity $\log^* p + O(1)$. Consequently, for all constants $\beta < 1/2$ and $\mu > 0$, there is a constant $\varepsilon < 1$ such that for all p and n, there is an explicit $(\beta, \mu, \varepsilon)$-resilient (p, n)-selection protocol.

[7] A very recent paper [2] aims to optimize ε as a function of β, obtaining efficient leader election protocols with $\varepsilon = O(\beta)$.

Multi-Party Random Selection. This is the general problem that encompasses the previous special cases.

- Goldreich, Goldwasser, and Linial [17] constructed, for every $n = 2^i$ and every p, an explicit (p, n)-selection protocol that is $(\beta, \mu, \mu^{1-O(\beta)})$-resilient for all sufficiently small β and every $\mu > 0$. The protocol runs in polylog(n) rounds. They also showed that any $(\beta, \mu, \varepsilon)$-resilient protocol must satisfy $\varepsilon \geq \mu^{1-\beta}$.
- Russell and Zuckerman [37] constructed, for every n and p such that $n \geq p^c$ for a constant c, an explicit one-round (p, n)-selection protocol that is $(\beta, \mu, \mu \cdot n/n^{\Omega(1-\beta)})$-resilient for every $\mu > 0$ and $1 > \beta > 0$.

Notice that all but the last of the above results require that the fraction β of bad players satisfies $\beta \leq 1/2$.[8] For collective coin-flipping and leader election, this is supported by impossibility results showing that $\beta \geq 1/2$ is impossible. For 2-party random selection, it does not make sense to discuss $\beta > 1/2$. The only result which applies to $\beta \geq 1/2$ is the last one (of [37]). However, the resilience $\mu \cdot n/n^{\Omega(1-\beta)}$ is quite weak and only interesting when the density μ of the bad set is close to $1/n$.[9] Our work is the first to show strong results for the case $\beta > 1/2$.

2.3 Our Results

In this section, we present our main results. All of our protocols utilize certain kinds of randomness-efficient samplers (equivalently, randomness extractors). Here we present the versions of our results obtained by using optimal samplers, proven to exist via the probabilistic method. We also have explicit (i.e., computationally efficient) versions of our protocols, obtained by using best known explicit constructions of samplers. One such protocol is given by Theorem 7, and the rest are deferred to the full version of this work [22].

The first main result of this paper is the following:

Theorem 3. *For all constants $k \in \mathbb{N}$, $k > 0$ and $\delta > 0$, there exists a constant $\varepsilon < 1$ and a (p, n)-selection protocol with the following properties:*

(i) The protocol has $\max(\log^* p, \log^* n) + O(1)$ *rounds.*
(ii) The protocol is $(1 - \alpha, \mu, \varepsilon)$*-resilient for* $\alpha = 1/(k+1) + \delta$ *and* $\mu = 1/k - \delta$.
(iii) The randomness complexity of the protocol is $(\log n)/\alpha + o(\log n) + O(p \log p)$.

The tradeoff between α and μ in the above theorem is optimal up to the slackness parameter δ. This is shown in Corollary 27, as a consequence of a lower bound of Feige [15]. Furthermore, the round and randomness complexity are nearly optimal as functions of n, as shown by Corollary 25 and Theorem 28.

Setting $p = 2$ and $\alpha = 1/2$, we obtain the following two-party protocol:

[8] The hidden constant in the protocol of [17] is larger than 2.

[9] The significance of the [37] protocol is that it is one round and only requires n polynomial in p; in fact, there is a trivial protocol with somewhat better parameters when n is exponential in p (Lemma 10).

Corollary 4. *For every constant $\delta > 0$, there exists a constant $\varepsilon < 1$ and a $(2, n)$-selection protocol with the following properties:*

(i) The protocol has $\log^* n + O(1)$ rounds.
(ii) The protocol is $(1/2, 1/2 - \delta, \varepsilon)$-resilient.
(iii) The randomness complexity of the protocol is $2\log n + o(\log n)$.

This protocol improves upon the two-party protocol of [39][10] in two ways: first, the randomness complexity is a nearly optimal $2\log n + o(\log n)$, and not polylog(n). Second, their protocol is $(1/2, \nu, \varepsilon')$-resilient for some small constant ν, and not for the nearly optimal $\frac{1}{2} - \delta$. In other words, their resilience is not optimal in the density of the bad set. On the other hand, the error probability ε' of their protocol is smaller than that of ours. However, a special case of our second theorem below gives the parameters of [39] with the added benefit of optimal randomness complexity.

Our next two results optimize the error probability ε as a function of the density μ of the bad set and fraction β of cheating players. The first achieves a near-optimal tradeoff when the number of players is small (e.g., constant).

Theorem 5. *For all $\mu, \alpha > 0$ and $p, n \in \mathbb{N}$, there exists a (p, n)-selection protocol with the following properties:*

(i) The protocol has $\log^* n - \log^*(1/\mu) + O(1)$ rounds.
(ii) The protocol is $(1 - \alpha, \mu, \varepsilon)$-resilient for

$$\varepsilon = \mu^{\alpha} \cdot O\left(\frac{1}{\alpha} \cdot \log \frac{1}{\mu} + (1-\alpha)p\right)^{1-\alpha} \cdot 2^{(1-\alpha)p}.$$

(iii) The randomness complexity is $[\log n + o(\log n) + O(p + \log(1/\mu))]/\alpha + \log(1/(1-\alpha)) + O(p \log p)$.

Note that when the number p of players and the fraction α of honest players are constants, the bound becomes $\varepsilon = \tilde{O}(\mu^{\alpha})$, which nearly matches the lower bound of $\varepsilon \geq \mu^{\alpha}$ proven in [17] (see Theorem 26). However, the bound on ε grows exponentially with p. This is removed in the following theorem, albeit at the price of achieving a slightly worse error probability of $\mu^{\Omega(\alpha)}$ (for constrained values of α).

Theorem 6. *There is a universal constant c such that for all $p, n \in \mathbb{N}$, $\mu, \alpha > 0$ satisfying $\alpha \geq \sqrt{c \log\log(1/\mu)/\log(1/\mu)}$, there exists a (p, n)-selection protocol with the following properties:*

(i) The protocol has $\max\{\log^* p, \log^* n\} - \log^*(1/\mu) + O(1)$ rounds.
(ii) The protocol is $(1 - \alpha, \mu, \mu^{\Omega(\alpha)})$-resilient.

[10] Note that in [39], the claimed round complexity is $2\log^* n + O(1)$, but this difference from our claim is only a difference of convention: in their model, only one player may communicate in each round, whereas we use the convention of multi-party protocols, in which all players may communicate simultaneously in one round.

(iii) The randomness complexity is $[\log n + o(\log n) + O(p)]/\alpha + O(p \log p) + \mathrm{poly}(1/\alpha, \log(1/\mu))$.

One disadvantage of the above two theorems (as compared to, say, the honest-majority protocols of [17]) is that the protocols require an a priori upper-bound μ on the density of the bad set. However, we also benefit from this, in that the round complexity improves as μ tends to zero. In particular, if $\mu \leq 1/\log^{(k)} n$ for some constant k, where $\log^{(k)}$ denotes k iterated logarithms, then the round complexity is *constant*.

An explicit version of the protocol of Theorem 3 is the following theorem:

Theorem 7. *For all constants* $k \in \mathbb{N}$, $k > 0$,$\gamma > 0$ *and* $\delta > 0$, *there exists a constant* $\varepsilon < 1$ *and an explicit* (p, n)*-selection protocol with the following properties:*

(i) The protocol has $\max(\log^* p, \log^* n) + O(1)$ *rounds.*
(ii) The protocol is $(1 - \alpha, \mu, \varepsilon)$*-resilient for* $\alpha = 1/(k + 1) + \delta$ *and* $\mu = 1/k - \delta$.
(iii) The randomness complexity of the protocol is $(\log n)^{1+\gamma} + O(p \log p)$.

Apart from its explicitness, note that the randomness complexity of the above theorem is now $(\log n)^{1+\gamma}$ for an arbitrarily small constant γ, rather than $(1 + o(1))(\log n)/\alpha$. Intuitively, this occurs because the explicit sampler we use (based on an extractor of [34]) only has randomness complexity polynomially close to optimal. It is possible to remedy this and obtain a randomness complexity of $(1 + o(1)) \log n$ by using other samplers (e.g. based on the extractors of [33]) for the first few rounds of universe reduction, but this creates some messy constraints on the other parameters, so we omit a formal statement.

We also have explicit versions of Theorem 5 and Theorem 6.

Theorem 8. *For every constant* $\gamma > 0$, *and every* $p, n \in \mathbb{N}$, $\mu, \alpha > 0$, *there exists a* (p, n)*-selection protocol with the following properties:*

(i) The protocol has $\log^* n - \log^*(1/\mu) + O(1)$ *rounds.*
(ii) The protocol is $(1 - \alpha, \mu, \varepsilon)$*-resilient for*

$$\varepsilon = \mu^{\alpha} \cdot O\left(\frac{1}{\alpha} \cdot \log \frac{1}{\mu} + (1 - \alpha)\, p\right)^{1-\alpha} \cdot 2^{(1-\alpha)p}.$$

(iii) The randomness complexity is $[(\log n)^{1+\gamma} + O(p + \log(1/\mu))]/\alpha + O(\log(1/(1 - \alpha))) + O(p \log p)$.
(iv) The protocol is explicit given appropriate samplers of size

$$s = \mathrm{poly}(2^p, 1/\mu, \log^{(3)} n)^{1/(\alpha)}.$$

which can be obtained probabilistically in time $O(s)$ *and deterministically in time* $2^{O(s)}$.

Theorem 9. *There is a universal constant* c *such that for every constant* $\gamma > 0$ *and every* $p, n \in \mathbb{N}$, $\mu, \alpha > 0$ *satisfying* $\alpha \geq \sqrt{c \log\log(1/\mu)/\log(1/\mu)}$, *there exists a* (p, n)*-selection protocol with the following properties:*

(i) The protocol has $\max\{\log^* p, \log^* n\} - \log^*(1/\mu) + O(1)$ *rounds.*
(ii) The protocol is $(1-\alpha, \mu, \mu^{\Omega(\alpha)})$*-resilient.*
(iii) The randomness complexity is $[(\log n)^{1+\gamma}+O(p)]/\alpha+O(p \log p)+\mathrm{poly}(1/\alpha, \log(1/\mu))$.
(iv) The protocol is explicit given appropriate samplers of size

$$s = \mathrm{poly}(1/\mu, 1/\alpha, \log^{(3)} n)^{1/\alpha}.$$

which can be obtained probabilistically in time $O(s)$ *and deterministically in time* $2^{O(s)}$.

Note that the protocols are explicit whenever $s = O(\log \log n)$ (in particular, when μ and α are constants). Due to space constraints, details of these explicit constructions are deferred to the full version [22].

3 One-round Protocols

We start with some simple one-round protocols that will play a role in our later constructions.

Lemma 10. *For every* $p, \ell \in \mathbb{N}$ *and* $n = \ell^p$*, there is an explicit* (p, n)*-selection protocol that is* $(\beta, \mu, n^\beta \cdot \mu)$*-resilient for every* $\beta, \mu > 0$.

Proof sketch. Each player outputs a random element of $[\ell]$, and we take the concatenation of the player's outputs. □

The above protocol has two main disadvantages. First, the size of the universe $n = \ell^p$ must be at least exponential in the number of players. (We note that Russell and Zuckerman [37] showed how to reduce this requirement to be only polynomial, at the price of a somewhat worse resilience. We will avoid this difficulty in a different manner, by first reducing the number of players.) Second, in terms of resilience, a bad set of density μ gets multiplied by a factor that grows polynomially with the universe size (namely, n^β). However, when the number of players is small (e.g. a fixed constant) and the universe is small (e.g. $n = O(1/\mu)$), it can achieve a nearly optimal bound on ε as a function of β and μ (cf. Theorem 26).

Lemma 11 ([15], Cor. 5). *For every* $p, n \in \mathbb{N}$ *and* $\alpha, \mu \in [0, 1]$ *such that* $\lfloor 1/\alpha \rfloor \leq \lceil 1/\mu \rceil - 1$*, there exists an* $\varepsilon < 1$ *and a* (p, n)*-selection protocol that is* $(1-\alpha, \mu, \varepsilon)$*-resilient. Specifically, one can take* $\varepsilon = 1 - \exp(-\Omega(\alpha \cdot (1-\mu) \cdot np))$.

Proof sketch. Every player chooses a random subset of $[n]$ of density at least $1 - \mu$, and the output is the first element of $[n]$ that is contained in every set S that was picked by at least an α fraction of players. Such an element exists because there exist at most $\lfloor 1/\alpha \rfloor \leq \lceil 1/\mu \rceil - 1$ such sets S, but any $\lceil 1/\mu \rceil - 1$ sets of density at least $1 - \mu$ must have a common intersection. □

The advantage of the above protocol is that it achieves an optimal tradeoff between α and μ (cf. Theorem 27). The main disadvantage is that ε can depend on p and n (this time with exponentially bad dependence), and that it is not sufficiently explicit — even the communication is of length $\Theta(n)$ (rather than $\mathrm{polylog}(n)$).

4 Universe and Player Reduction

The simple 1-round protocols of the previous section behave well when the number of players p and universe size n are small. Thus, as in previous work, our main efforts will be aimed at giving protocols to reduce p and n while approximately preserving the fraction β of bad players and the density μ of the bad set. Roughly speaking, in one round we will reduce p and n to polylog(p) and polylog(n), respectively. For this, we consider protocols that select a subset of the universe (or a subset of the players).

4.1 Definitions

Definition 12. *A* $[(p, n) \mapsto n']$-universe reduction protocol *is a p-player protocol whose output is a sequence $(s_1, \ldots, s_{n'})$ of elements of $[n]$. Such a protocol is* $[(\beta, \mu) \overset{\gamma}{\mapsto} \mu']$-resilient *if when at most a β fraction of players are cheating and S is any subset of $[n]$ of density at most μ, the probability that at most a μ' fraction of the output sequence is in S is at least γ. It is* explicit *if the players' strategies are computable in time* $\mathrm{poly}(\log n, p)$*, and given the protocol transcript and $i \in [n']$, the i'th element of the output sequence is computable in time* $\mathrm{poly}(\log n, p)$.

Notice that a (p, n)-selection protocol is equivalent to a $[(p, n) \mapsto n']$-universe reduction protocol with $n' = 1$, and the former is $(\beta, \mu, \varepsilon)$-resilient if and only if the latter is $[(\beta, \mu) \overset{1-\varepsilon}{\mapsto} 0]$-resilient.

Definition 13. *A* $[p \mapsto p']$-player reduction protocol *is a* $[(p, p) \mapsto p']$*-universe reduction protocol. It is* $[\beta \overset{\gamma}{\mapsto} \beta']$-resilient *if when at most a β fraction of players are cheating, the probability that at most a β' fraction of the output sequence are indices of cheating players is at least γ.*

Definition 14. *A* $[(p, n) \mapsto (p', n')]$-universe+player reduction protocol *is a p-player protocol whose output is a sequence $(s_1, \ldots, s_{n'})$ of elements of $[n]$ and a sequence $(t_1, \ldots, t_{p'})$ of elements of $[p]$. Such a protocol is* $[(\beta, \mu) \overset{\gamma}{\mapsto} (\beta', \mu')]$-resilient *if when at most a β fraction of players are cheating and S is any subset of $[n]$ of density at most μ, the probability that at most a β' fraction of the first output sequence are cheating players and at most a μ' fraction of the second output sequence is in S is at least γ. It is* explicit *if the players' strategies are computable in time* $\mathrm{poly}(\log n, p)$*, and given the protocol transcript and $i \in [n']$ (resp., $j \in [p']$), the i'th (resp., j'th) element of the first (resp., second) output sequence is computable in time* $\mathrm{poly}(\log n, p)$.

4.2 One-Round Reduction Protocols

In the following one-round protocols, think of $\theta = 1/\mathrm{polylog}(n)$ and $\varepsilon = 1/\mathrm{poly}(n)$. We use both a one-round player-reduction protocol and a one-round universe-reduction protocol.

Theorem 15 ([37, 15]). *For every $p \in \mathbb{N}$, $\varepsilon > 0$, and $\theta > 0$, there is an explicit, one-round $[p \mapsto p']$-player reduction protocol with*

$$p' = O\left(\frac{1-\beta}{\theta^2} \cdot \log \frac{p}{\varepsilon}\right),$$

that is $[\beta \overset{1-\varepsilon}{\mapsto} \beta+\theta]$-resilient for all $\beta > 0$. Moreover, the randomness complexity is $p \cdot \log(p/p')$.

The starting point for our universe reduction protocol is the simple protocol of Lemma 10. That protocol has the property that a β fraction of cheating players cannot make any outcome in $[n]$ appear with probability more than $1/n^{1-\beta}$. (This can be seen by taking $\mu = 1/n$.) Thus the output can be viewed as a source with "min-entropy rate" at least $1 - \beta$.[11] To get a higher quality output, it is natural to try applying a *randomness extractor*, a function that extracts almost-uniform bits from sources with sufficient min-entropy. However, randomness extractors require an additional random *seed* to do such extraction. Thus we will enumerate over all seeds of the extractor, and the resulting sequence will be the output of our universe reduction protocol. Fortunately, there exist extractors where the number of seeds is only polylogarithmic in n, the domain of the source.

Actually, it is more convenient for us to work with an object that is essentially equivalent to extractors, namely (averaging) samplers (cf., [9, 43, 16]). Samplers are functions that output sample points of a given universe, with the property that the fraction of samples from any particular subset of the universe is roughly equal to the density of that subset. In the following definition, U_r denotes an element of $[r]$ chosen uniformly at random.

Definition 16. *A function* $\mathrm{Samp} : [r] \to [n]^t$ *is a* (θ, ε) sampler *if for every set* $S \subseteq [n]$,

$$\Pr_{(i_1,\ldots,i_t) \leftarrow \mathrm{Samp}(U_r)} \left[\frac{\#\{j : i_j \in S\}}{t} > \frac{|S|}{n} + \theta\right] \leq \varepsilon.$$

We say that Samp *is* explicit *if for every $x \in [r]$ and every $j \in [t]$, the j'th component of* $\mathrm{Samp}(x)$ *can be computed in time* $\mathrm{poly}(\log r, \log n)$.[12]

Zuckerman [43] showed that samplers (as defined above) are essentially equivalent to randomness extractors.

Given $p, \ell \in \mathbb{N}$ and a sampler $\mathrm{Samp} : [r] \to [n]^{n'}$ with $r = \ell^p$, we obtain a $[(p, n) \mapsto n']$-universe reduction protocol Π_{Samp} as follows: the players use

[11] The *min-entropy* of a random variable X is defined as $H_\infty(X) = \max_x \Pr[X = x]$. If X takes values in a universe U, then its *min-entropy rate* is defined to be $H_\infty(X)/\log|U|$.

[12] Often the definition of samplers also requires that the fraction of samples that lie in S is also not much larger than the density of S. However, this follows from our definition (paying a factor of 2 in ε) by considering $\overline{S}$. Moreover, we will only need the one-sided version, and below, in Definition 20 we will consider a variant which is not symmetric with respect to approximation from above and below.

the protocol of Lemma 10 to select an element $x \in [\ell^p]$, and then output the sequence $\mathrm{Samp}(x)$.

Lemma 17. *If* Samp *is a* (θ, ε) *averaging sampler, then for every* $\mu, \beta > 0$, Π_{Samp} *is* $[(\beta, \mu) \overset{\gamma}{\mapsto} \mu + \theta]$*-resilient for* $\gamma = 1 - r^{\beta} \cdot \varepsilon$*. Moreover, the randomness complexity is* $\log r$*.*

Proof. Call $x \in [r]$ "bad" if $\#\{j : i_j \in S\}/t > |S|/n + \theta$ when $(i_1, \ldots, i_t) \leftarrow \mathrm{Samp}(x)$, and note that the number of bad x's is at most $\varepsilon \cdot r$ by the properties of the sampler. The players use the protocol of Lemma 10 to select an element x from a universe of size r, where the fraction of bad elements is ε. This is a (p, r)-selection protocol that is $(\beta, \varepsilon, r^{\beta} \cdot \varepsilon)$-resilient, and so the probability of selecting a good x is at least $\gamma = 1 - r^{\beta} \cdot \varepsilon$. If a good x is selected, then the fraction of bad elements is increased by at most θ. □

Notice that for this to be useful, we need the error probability ε of the sampler to be smaller than $r^{-\beta}$, and in fact we will be interested in β that are arbitrarily close to 1. Fortunately, we have samplers that achieve this. (This is equivalent to the fact that we have extractors that work for min-entropy rate arbitrarily close to 0.)

Lemma 18 (nonconstructive samplers [36, 43]). *There is a universal constant* c *such that for every* $n \in \mathbb{N}, \theta > 0, \varepsilon > 0$ *and* $r \geq c \cdot n/(\varepsilon\theta^2)$*, there exists a* (θ, ε) *sampler* $\mathrm{Samp} : [r] \to [n]^t$ *with* $t = O(\log(1/\varepsilon)/\theta^2)$*.*

It is important to note that the lower bound on r depends linearly on $1/\varepsilon$; this means that we can make the error $\varepsilon \leq r^{-\beta}$ for any $\beta < 1$. Combining the above two lemmas, we have:

Theorem 19 (nonconstructive 1-round universe reduction). *For every* $p, n \in \mathbb{N}$*,* $\beta, \varepsilon, \theta > 0$*, there exists a 1-round* $[(p, n) \mapsto n']$*-universe reduction protocol that is* $[(\beta, \mu) \overset{1-\varepsilon}{\mapsto} \mu + \theta]$*-resilient for every* $\mu > 0$*, with*

$$n' = O\left(\frac{\log(1/\varepsilon) + (\beta/(1-\beta)) \cdot \log n + \beta \cdot p}{\theta^2}\right).$$

Moreover, the randomness complexity is $p + (\log n + \log(1/\varepsilon) + 2\log(1/\theta))/(1 - \beta) + O(1)$*.*

Proof. First note that without loss of generality, $\theta \geq 1/n$, otherwise we can use the trivial protocol that outputs the entire universe. So now choose $r \in [(cn/(\varepsilon\theta^2))^{1/(1-\beta)}, 2^p \cdot (cn/(\varepsilon\theta^2))^{1/(1-\beta)}]$ such that r is the p'th power of some natural number, and apply Lemma 17 with $\varepsilon' = \varepsilon/r^{\beta}$. □

Thus, for $p = \mathrm{polylog}(n)$, $\theta = 1/\mathrm{polylog}(n)$, $\varepsilon = 1/\mathrm{poly}(n)$, and $\beta = 1 - 1/\mathrm{polylog}(n)$, we can reduce the universe size from n to $\mathrm{polylog}(n)$. If the number of players is constant, then we can iterate this $\log^* n$ times to reduce the universe size to a constant. However, if the number of players p is large, then the above

will not reduce the universe size below βp. Therefore, we will combine this with the player reduction of Theorem 15.

Notice that if we want to preserve the density μ of the bad set up to a constant factor, then we can set $\theta = 1/\mu$ and the above protocol will reduce the universe size to n' depending polynomially on $1/\mu$. However, to obtain some of our results (namely, Theorems 6, 5, and their explicit versions), it will be beneficial to reduce to a universe size that depends almost-linearly on $1/\mu$. To achieve this, we use a variant of our sampler-based protocol that is tailored to a particular value of μ.

Definition 20. *A function* $\mathrm{Samp} : [r] \to [n]^t$ *is a* $(\mu, \theta, \varepsilon)$ density-tailored sampler *if for every set* $S \subseteq [n]$ *with* $|S| \leq \mu \cdot n$,

$$\Pr_{(i_1,\ldots,i_t)\leftarrow \mathrm{Samp}(U_r)} \left[\frac{\#\{j : i_j \in S\}}{t} > \mu + \theta \right] \leq 1 - \varepsilon.$$

We say that Samp *is* explicit *if for every* $x \in [r]$ *and every* $j \in [t]$*, the* j*'th component of* $\mathrm{Samp}(x)$ *can be computed in time* $\mathrm{poly}(\log r, \log n)$.

Density-tailored samplers are essentially equivalent to 'slice extractors', defined in [36]. As in Lemma 17, these density-tailored samplers also induce selection protocols.

Lemma 21. *If* Samp *is a* $(\mu, \theta, \varepsilon)$ *density-tailored sampler, then for every* $\beta > 0$*,* Π_{Samp} *is* $[(\beta, \mu) \overset{\gamma}{\mapsto} \mu + \theta]$*-resilient for* $\gamma = 1 - r^{\beta} \cdot \varepsilon$*. Moreover the randomness complexity is* $\log r$.

The reason we are interested in these density-tailored samplers is that they exist with slightly better parameters for certain values of μ.

Lemma 22 (nonconstructive density-tailored samplers [40]). *There is a universal constant* c *such that for every* $n \in \mathbb{N}, \mu > 0, \theta > 0, \varepsilon > 0$*,* $t \geq c \cdot \log(1/\varepsilon) \cdot \max\{1/\mu, \mu/\theta^2\}$*, and* $r \geq c \cdot n \cdot (\mu \log(1/\mu))/(\varepsilon \log(1/\varepsilon))$*, there exists a* $(\mu, \theta, \varepsilon)$ *density-tailored sampler* $\mathrm{Samp} : [r] \to [n]^t$.

Note that the number of samples t in these samplers depends linearly on $1/\mu$ (if $\theta = \Omega(\mu)$) and not polynomially as in Lemma 18. Combining the above lemma with Lemma 17 we get a nonconstructive 1-round universe reduction protocol with different parameters from those of Theorem 19:

Theorem 23 (nonconstructive, density-tailored 1-round universe reduction). *There is a universal constant* c *such that for every* $p, n \in \mathbb{N}$*,* $\beta, \mu, \varepsilon, \theta > 0$ *and every*

$$n' \geq c \cdot \max\left\{\frac{\mu}{\theta^2}, \frac{1}{\theta}\right\} \cdot \left(\log\frac{1}{\varepsilon} + \frac{\beta}{1-\beta} \cdot \left(\log n + \log\frac{1}{\beta}\right) + \beta \cdot p\right),$$

there exists a 1-round $[(p, n) \mapsto n']$*-universe reduction protocol that is* $[(\beta, \mu) \overset{1-\varepsilon}{\mapsto} \mu + \theta]$*-resilient. Moreover, the randomness complexity is* $p + [\log n + \log(1/\varepsilon) - \log\log(1/\varepsilon) - \log(1/\mu) + \log\log(1/\mu) + \log(1/\beta)]/(1-\beta) + O(1)$.

Proof. We can choose $r \in [r', 2^p \cdot r']$ such that r is the p'th power of some natural number and

$$r' = \left(\frac{cn \cdot \mu \log \frac{1}{\mu}}{\beta \cdot \varepsilon \log \frac{1}{\varepsilon}} \right)^{\frac{1}{1-\beta}}.$$

We then apply Lemmas 21 and 22 with $\varepsilon' = \varepsilon / r^\beta$. □

5 Putting It Together

In this section we give a sketch of how our main results are obtained by combining the various building blocks described above. Due to space constraints, details and formal proofs are deferred to the full version of this work [22].

First we construct a sub-protocol that reduces the size of the universe and the number of players to $n', p' = \text{poly}(\log(1/\varepsilon), 1/\theta)$ (Theorem 24). This is accomplished by iterating the 1-round player reduction protocol of Theorem 15 to reduce the number of players, and iterating the 1-round universe reduction protocol of Theorem 19 to reduce the size of the universe. To save on the round complexity, the player reduction and universe reduction can be done in parallel. This yields the following theorem, which is the main component of our protocols:

Theorem 24 (many-round universe+player reduction). *For every $n, p \in \mathbb{N}$ and every $\beta, \theta, \varepsilon > 0$, there exists a $[(p, n) \mapsto (p', n')]$-universe+player reduction protocol that is $[(\beta, \mu) \overset{1-\varepsilon}{\mapsto} (\beta + \theta, \mu + \theta)]$-resilient for every $\mu > 0$, with*

$$n' = \text{poly}(\log(1/\varepsilon), 1/\theta)$$
$$p' = \text{poly}(\log(1/\varepsilon), 1/\theta).$$

Moreover, the number of rounds is $t = \max\{\log^* n, \log^* p\} - \log^* n' + O(1)$ and the randomness complexity is $[\log n + o(\log n) + O(p + t \cdot \log(1/\varepsilon) + t \cdot \log(1/\theta))]/(1-\beta) + O(p \log p)$.

For Theorem 3, the size of the universe and the number of players are both reduced to constant ($\text{poly}(1/\delta)$, where δ is the constant in the statement of the theorem). We can then run the protocol of Lemma 11, and since all parameters are constant, the error probability of the final output is also bounded by a constant less than 1.

For Theorem 5, we wish to use the protocol of Lemma 10 as the final protocol, and thus obtain a near optimal probability of error. For this to work well, however, we require that n' be exponential in p, and that $n' = \tilde{O}(1/\mu)$. This requires a more delicate reduction than the one above.

First, we reduce the size of the universe to $\text{poly}(1/\mu)$ in the same way as above (Theorem 24, ignoring the player reduction). Next we further reduce the size of the universe to $n' = \tilde{O}(1/\mu)$, using the density-tailored 1-round universe reduction protocol of Theorem 23. This use of a density-tailored reduction is critical, as it allows the reduction of n' to $1/\mu$. Once this is accomplished, we can use the protocol of Lemma 10.

The protocol of Theorem 6 is the same as that of Theorem 5, except that it is combined with an appropriate player-reduction protocol (which can be run in parallel).

6 Lower Bounds

The are several known and new lower bounds for the different parameters of random selection, and here we state the relevant theorems. A lower bound on the round complexity is a corollary of a theorem of [39]:

Corollary 25. *For any (p, n)-selection protocol that is $(\beta, \mu, \varepsilon)$-resilient with $\beta \geq 1/2$ and for constants $\mu > 0$ and $\varepsilon < 1$, the round complexity is at least $(\log^* n - \log^* \log^* n - O(1))/2$.*

The lower bound on the error probability ε is given by a theorem of [17]:

Theorem 26 ([17]). *For any (p, n)-selection protocol that is $(1-\alpha, \mu, \varepsilon)$-resilient, $\varepsilon \geq \mu^{\alpha}$.*

The optimal tradeoff between μ and α is given by the following corollary of [15, Thm. 4]:

Corollary 27. *For any (p, n)-selection protocol that is $(1-\alpha, \mu, \varepsilon)$-resilient with $\varepsilon < 1$, $\lfloor 1/\alpha \rfloor \leq \lceil 1/\mu \rceil - 1$.*

Finally, in the full version of this work [22], we prove a lower bound on the randomness and communication complexities.

Theorem 28. *For any (p, n)-selection protocol that is $(1 - \alpha', \mu, \varepsilon)$-resilient for $\varepsilon < 1$, the randomness and communication complexities are at least*

$$\max\left\{(1-\alpha')p, \frac{1-\varepsilon}{\alpha}\log\frac{\mu n}{\varepsilon}\right\},$$

where $\alpha = \lfloor \alpha' p \rfloor / p$.

Acknowledgements

We thank Omer Reingold for many useful discussions, and the anonymous referees for helpful comments.

References

1. N. Alon and M. Naor. Coin-flipping games immune against linear-sized coalitions. *SIAM J. Computing*, 22(2):403–417, April 1993.
2. S. Antonakopoulos. Fast leader-election protocols with bounded cheaters' edge. In *Proc. 38th STOC*, 187–196, 2006.

3. B. Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model . In *43rd FOCS*, 2002.
4. M. Blum. Coin flipping by telephone. In *IEEE Spring COMPCOM*, 1982.
5. D. Beaver and S. Goldwasser. Multiparty computation with faulty majority. In *CRYPTO 89*, Springer LNCS 435, 589–590, 1989.
6. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *20th STOC*, 1–10, 1988.
7. M. Ben-Or and N. Linial. Collective coin fliping. In *Advances in Computing Research*, volume 5: Randomness and Computation, JAI Press, 1989, 91–115.
8. R. Boppana and B. Narayanan. Perfect-information leader election with optimal resilience. *SIAM J. Computing*. 29(4): 1304-1320 (2000).
9. M. Bellare and J. Rompel. Randomness-efficient oblivious sampling. In *35th FOCS*, 1994.
10. D. Chaum, C. Crépeau, and I. Damgård. Multiparty unconditionally secure protocols. In *20th STOC*, 11–19, 1988.
11. J. Cooper and N. Linial. Fast perfect-information leader-election protocols with linear immunity. *Combinatorica*, 15:319-332, 1995.
12. I. Damgård. Interactive hashing can simplify zero-knowledge protocol design without computational assumptions (extended abstract). In *CRYPTO*, 1993.
13. I. Damgård, Oded Goldreich, and Ave Wigderson. Hashing functions can simplify zero-knowledge protocol design (too). TR RS-94-39. BRICS, 1994.
14. Y. Z. Ding, D. Harnik, A. Rosen, and R. Shaltiel. Constant-round oblivious transfer in the bounded storage model. In *1st TCC*, 2004.
15. U. Feige. Noncryptographic selection protocols. In *40th FOCS*, 142–152, 1999.
16. O. Goldreich. A sample of samplers - a computational perspective on sampling (survey). Report 97-020, *Electronic Colloquium on Computational Complexity*, 1997.
17. O. Goldreich, S. Goldwasser, N. Linial. Fault-tolerant computation in the full information model. *SIAM J. Computing* 27(2), 1998.
18. S. Goldwasser and L. Levin. Fair computation of general functions in presence of immoral majority. In *CRYPTO 90*, Springer LNCS 537, 77–93, 1990.
19. S. Goldwasser and Y. Lindell. Secure computation without agreement. In *DISC* 2002: 17-32
20. O. Goldreich, S. Micali, and A. Wigderson. How to play ANY mental game. In *19th STOC*, 218–229, 1987.
21. O. Goldreich, A. Sahai, and S. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *30th STOC*, 1998.
22. R. Gradwohl, S. Vadhan, and D. Zuckerman. Random Selection with an Adversarial Majority. Report TR06-26, *Electronic Colloquium on Computational Complexity*, Feb. 2006.
23. J. Katz and R. Ostrovsky. Round-optimal secure two-party computation. In *CRYPTO*, 2004.
24. J. Katz, R. Ostrovsky, and A. Smith: Round efficiency of multi-party computation with a dishonest majority. In *EUROCRYPT 2003*: 578-595.
25. Y. Lindell. Parallel coin-tossing and constant-round secure two-party computation. In *CRYPTO*, 2001.
26. M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
27. C. Lu, O. Reingold, S. Vadhan, and A. Wigderson. Extractors: Optimal up to constant factors. In *35th STOC*, 2003.

28. 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.
29. N. Nisan and D. Zuckerman. Randomness is linear in space. *J. Computer and System Sci.*, 52(1): 43–52, 1996.
30. T. Okamoto. On relationships between statistical zero-knowledge proofs. *J. Computer and System Sci.*, 60(1): 47–108, 2000.
31. R. Ostrovsky, S. Rajagopalan, and U. Vazirani. Simple and efficient leader election in the full information model. In *Proc. 26th STOC*, 234–242, 1994.
32. R. Ostrovsky, R. Venkatesan. and M. Yung. Interactive hashing simplifies zero-knowledge protocol design. In *EUROCRYPT*, 1993.
33. R. Raz, O. Reingold, and S. Vadhan. Extracting all the randomness and reducing the error in Trevisan's extractors. *J. Computer and System Sc.*, 65(1):97–128, 2002.
34. R. Raz, O. Reingold, and S. Vadhan. Error Reduction for Extractors. In *40th FOCS*, 1999.
35. Alexander Russell, Michael Saks and David Zuckerman. Lower bounds for leader election and collective coin- flipping in the perfect information model. *SIAM J. Computing*, 31:1645–1662, 2002.
36. J. Radhakrishnan and A. Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM J. Discrete Math.*, 13(1):2–24, 2000.
37. A. Russell and D. Zuckerman. Perfect-information leader election in $\log^* n + O(1)$ rounds. *J. Computer and System Sci.*, 63:612–626, 2001.
38. M. Saks. A robust noncryptographic protocol for collective coin flipping. *SIAM J. Discrete Math.,* 2(2):240-244, 1989.
39. S. Sanghvi and S. Vadhan. The round complexity of two-party random selection. In *37th STOC*, 2005.
40. S. Vadhan. Constructing locally computable extractors and cryptosystems in the bounded-storage model. *J. Cryptology* 17(1), 43–77, 2004.
41. A. Wigderson and D. Zuckerman, Expanders that beat the eigenvalue bound: explicit construction and applications. *Combinatorica*, 19 (1999) 125–138. 17
42. A. Yao. How to generate and exchange secrets. In *Proc. 27th FOCS*, 1986.
43. D. Zuckerman. Randomness-optimal oblivious sampling. *Random Structures and Algorithms,* 11(4): 345-367 (1997).