Lower Bounds for Leader Election and Collective Coin Flipping, Revisited

Eshan Chattopadhyay*
Cornell University
eshan@cs.cornell.edu

Mohit Gurumukhani*
Cornell University
mgurumuk@cs.cornell.edu

Noam Ringach †
Cornell University
nomir@cs.cornell.edu

Rocco Servedio‡
Columbia University
rocco@cs.columbia.edu

Abstract

In fault-tolerant distributed computing, two fundamental tasks are collective coin flipping—where processors agree on a common random bit—and leader election—where they designate a leader among themselves. We study these problems in the full-information model, where processors communicate via a single broadcast channel, have access to private randomness, and face a computationally unbounded adversary that controls some of the processors. Despite decades of study, key gaps remain in our understanding of the trade-offs between round complexity, communication per player in each round, and adversarial resilience.

We make progress by proving new lower bounds for coin flipping protocols, which also imply lower bounds for leader election protocols. Specifically, we show that any k-round coin flipping protocol, where each of ℓ players sends 1 bit per round, can be biased by $O(\ell/\log^{(k)}(\ell))$ bad players. For all $k > 1$ this strengthens the previous best lower bounds [RSZ, SICOMP 2002], which ruled out protocols resilient to adversaries controlling $O(\ell/\log^{(2k-1)}(\ell))$ players. As a consequence, we establish that any protocol tolerating a linear fraction of corrupt players, while restricting player messages to 1 bit per round, must run for at least $\log^* \ell - O(1)$ rounds, improving on the prior best lower bound of $\frac{1}{2}\log^* \ell - \log^* \log^* \ell$. This lower bound also matches the number of rounds, $\log^* \ell$, taken by the current best coin flipping protocols from [RZ, JCSS 2001], [F, FOCS 1999] that can handle a linear sized coalition of bad players, given the additional freedom that players can send unlimited bits per round. We also extend our techniques to derive lower bounds for protocols allowing multi-bit messages per round. Our results show that the protocols from [RZ, JCSS 2001], [F, FOCS 1999] that handle a linear number of corrupt players are almost optimal in terms of round complexity and communication per player in a round.

A key technical ingredient in proving our lower bounds is a new result regarding biasing most functions from a family of functions using a common set of bad players and a small specialized set of bad players specific to each function that is biased.

Complementing our lower bound results, we give improved constant-round coin flipping protocols in the setting that each player can send 1 bit per round. For example, in the case of two rounds, our protocol can handle $O(\ell/(\log \ell)(\log\log \ell)^2)$ sized coalition of bad players; this is better than the best one-round protocol (also called a resilient function) by [AL, Combinatorica 1993] in this setting, which can only handle $O(\ell/(\log \ell)^2)$ sized coalition of bad players.

*Supported by a Sloan Research Fellowship and NSF CAREER Award 2045576.
†Supported by NSF GRFP grant DGE – 2139899, NSF CAREER Award 2045576 and a Sloan Research Fellowship.
‡Supported by NSF Award CCF-2106429 and NSF Award CCF-2211238.

Contents

1 Introduction

Two related and fundamental tasks in fault tolerant distributed computing are for processors to (1) collectively *flip a coin*, i.e. agree on a common random bit, and to (2) *select a 'leader'* among them, in the presence of adversarial faults. These tasks have been widely studied under various assumptions about the communication channels between the processors, the kind of faults allowed, the power of the adversary, and the kinds of randomness each processor has access to. We study this problem under the assumption that there is only a single common broadcast channel, that adversarial processors are computationally unbounded, and that each processor has access to private randomness. This widely studied model was introduced by Ben-Or and Linial [BL85], and is known as the *full information model* since all processors have access to the same information.

The standard way of modeling how processors coordinate their efforts is through the notion of a *protocol*. All players (processors henceforth will be called players) agree on a fixed protocol π beforehand and execute it, at the end of which they all agree on a common random bit or a leader. We assume that there is an adversary $\mathcal{A}$ that selects a subset of players before the protocol begins and continues controlling them throughout the execution of the protocol. We refer to the coalition of controlled players as 'bad' and the remaining players as 'good'. A protocol consists of one or more rounds. In each round, all players should flip r private random coins and broadcast those r bits to everyone else. The identity of the sender of the bits is always known. This continues for k rounds, after which the protocol determines the common random bit or selects the leader. We assume that each round of the protocol is asynchronous, and we always consider the worst-case scenario in which all good players broadcast their bits at the beginning of the round. Then, based on their outputs, the adversary $\mathcal{A}$ determines the r bits output by each of the bad players. As a result, the outputs of the bad players are coordinated and depend on the outputs of the good players and the outputs of previous rounds. We note that the players are synchronized in between rounds, i.e., a round ends only when all outputs of all players are received. For formal definitions of these protocols, we refer the reader to Section 3.3.1.

Since bad players in this model are computationally unbounded, cryptography-based protocols, which are standard for the Byzantine generals problem and other related models where the adversary is (say) polynomially bounded, are not useful in our setting. Nevertheless, remarkable protocols do exist in this model that can guarantee that a good leader is always chosen with non-trivial probability, or that the outcome of the coin flip is not too biased towards any particular result.

1.1 Related Work

We first survey the best protocol constructions that are currently known. We mention that given a leader election protocol, one can turn it into a collective coin flipping protocol at the expense of one extra round by making the elected leader flip a coin and output the result (see Claim 3.10 for a proof). Hence, most constructions have focused on leader election protocols. In all these protocols, the goal is to minimize the number of rounds, minimize the number of random bits each player can send per round, and maximize the number of bad players the protocol can handle while electing a good leader with non-trivial probability / flipping a coin where each output in $\{0, 1\}$ has at least constant probability of appearing.

Starting with Ajtai and Linial's non-explicit function [AL93], later made explicit by [CZ19, Mek17, IMV23, IV24], one can construct a one-round collective coin flipping protocol, where each player sends 1 random bit and the output remains close to uniform even in the presence of $O(\ell/(\log \ell)^2)$ bad players. Leader election protocols that can guarantee a constant probability of electing a good leader in the presence of $O(\ell)$ bad players have been widely studied: when the number of random bits per round is restricted to be 1, a long line of works [BL85, Sak89, AN93, CL95, ORV94, Zuc97, Fei99, Ant06] constructed various explicit protocols that run in $O(\log \ell)$ rounds. Moreover, when the number of random bits per round is unrestricted, [RZ01, Fei99] constructed explicit leader election protocols that run in $\log^*(\ell) + O(1)$ rounds.[1]

On the negative side, very few results are known regarding the non-existence of protocols with various parameters. Before stating the known results, we observe that since leader election protocols imply coin flipping protocols, lower bounds for coin flipping protocols imply lower bounds for leader election protocols with one less round (see Corollary 3.11 for formal claim). Hence, all efforts to date have focused on providing lower bounds for coin flipping protocols. First, it is well known that no protocol can handle $\ell/2$ or more bad players [Sak89], regardless of the number of rounds and the number of random bits allowed per round. For 1 round coin flipping protocols where players send 1 bit per round, [KKL88] showed every protocol can be biased towards some outcome by some set of $O(\ell/\log \ell)$ bad players, almost matching the construction of [AL93]. For multiple rounds, [RSZ02] showed that any k-round coin flipping protocol where each player sends 1 bit per round can be biased towards some outcome by $O(\ell/\log^{(2k-1)} \ell)$ bad players.[2] This also implies that in order to handle $\Theta(\ell)$ bad players when the number of bits per round is 1, the number of rounds required is $\frac{1}{2}\log^*(\ell) - \log^* \log^* \ell$. For coin flipping protocols where players can send $\left(\log^{(2i-1)} \ell\right)^{1-o(1)}$ random bits in round i, [RSZ02] showed that there always exists a set of $o(\ell)$ bad players that can bias them towards some outcome.

The work of Filmus et al. [FHHHZ19] gave lower bounds on collective coin flipping protocols under arbitrary product distributions, rather than the uniform distribution, on the Boolean cube. For further details, we refer the reader to the excellent survey of Dodis [Dod06] on protocols and lower bounds in the full information model.

1.2 Our Results

Our main results are various improved lower bounds for coin flipping protocols. As noted above, since a leader election protocol can be turned into a coin flipping protocol (with one extra round), our lower bounds for coin flipping protocols immediately imply lower bounds for leader election protocols as well.

Our lower bounds improve and subsume all previous multi-round lower bounds (established by [RSZ02]). As discussed below, in some sense our lower bounds show that the protocols of [RSZ02] and [Fei99] are optimal.

Theorem 1 (Informal version of Theorem 5.1)**.** *For any k-round coin flipping protocol π over ℓ players where players send 1 bit per round, there exists a set of bad players $B \subset [\ell]$ and an outcome $o \in \{0,1\}$, with $|B| \le O\left(\frac{\ell}{\log^{(k)}(\ell)}\right)$, such that the players in B can bias π to output o with probability 0.99.*

[1] Recall that $\log^*(\ell)$ is the minimum number of logs that need to be applied to ℓ until it attains value at most 1.

[2] We use the notation $\log^{(i)} \ell$ to denote i-times-iterated logarithm, i.e. $\overbrace{\log \cdots \log}^{i \text{ times}} \ell$.

This directly improves upon the previous best lower bounds, due to [RSZ02], which showed that $O\left(\frac{\ell}{\log^{(2k-1)}(\ell)}\right)$ players can corrupt such protocols, achieving a strictly better bound for all $k > 1$.

We also obtain a lower bound on the number of rounds required for protocols that can handle a linear number of bad players:

Corollary 1 (Informal version of Corollary 5.2)**.** *For any k-round coin flipping protocol π over ℓ players where players send 1 bit per round and where $k \leq \log^* \ell - O(1)$, there exists a set of bad players $B \subset [\ell]$ with $|B| \leq 0.01\ell$ and an outcome $o \in \{0, 1\}$ such that the players in B can bias π to output o with probability ≥ 0.99.*

This result shows that any coin flipping protocol, where players send 1 bit per round, that can handle linear sized coalitions requires $\log^* \ell - O(1)$ rounds. This improves upon the best previous lower bound of $\frac{1}{2} \log^* \ell - \log^* \log^* \ell$ rounds, due to [RSZ02] . This lower bound also essentially matches the number of rounds, $\log^* \ell$, used by the current best coin flipping protocols from [RZ01, Fei99] that can handle linear sized coalition of bad players, given the additional freedom that players can send unlimited bits per round.

We remark that a key ingredient towards proving Theorem 1 is a new lemma regarding biasing most functions from a family of functions using a common set of bad players and a small specialized set of bad players specific to each function being biased. For a formal statement, refer to Theorem 4.1

We also obtain lower bounds against protocols where players can send multiple bits per round:

Theorem 2 (Informal version of Theorem 5.5)**.** *Let π be a k-round coin flipping protocol over ℓ players where each player can send $(\log^{(i)} \ell)^{0.99}$ many bits in round i, and $k \leq \log^* \ell - O(1)$. Then, there exists a set of bad players $B \subset [\ell]$ and an outcome $o \in \{0, 1\}$, with $|\mathbf{B}| \leq 0.01\ell$, such that the players in B can bias π to output o with probability ≥ 0.99.*

The coin flipping protocols from [RZ01, Fei99] that can handle a linear sized coalition of bad players take $\log^* \ell + O(1)$ rounds, and in round i, only require that all players send $O(\log^{(i)} \ell)$ many bits. Hence, Theorem 2 shows that the protocols from [RZ01, Fei99] are essentially optimal, in the sense that if the number of rounds and the number of bits each player can send in round i are even slightly lowered, then no protocol can handle a linear sized coalition of bad players. This result also directly improves on a bound in this setting obtained by [RSZ02]; their result showed that if each player can send at most $(\log^{(2i-1)} \ell)^{0.99}$ bits in round i, then any such protocol cannot handle a linear sized coalition of bad players. This resolves an open problem raised in [RSZ02], where they exactly asked for a result along the lines of Theorem 2.

We also complement the lower bound result described above, by constructing improved constant-round protocols in which each player can send one bit per round. Our goal is to maximize the number of bad players the protocol can tolerate while still ensuring that the output coin flip remains nearly balanced. We provide constructions of such protocols.

Theorem 3 (Informal version of Theorem 6.1)**.** *For any $k \geq 2$, there exists an explicit k-round coin flipping protocol over ℓ players where each player sends 1 bit per round such that when the number of bad players is at most $O\left(\frac{\ell}{(\log \ell)(\log^{(k)} \ell)^2}\right)$, the output bit is 0.01-close to uniform.*

When $k = 2$, this protocol can handle any $O\left(\frac{\ell}{(\log \ell)(\log \log \ell)^2}\right)$-sized coalition of bad players, which is better than the best 1 round protocol by [AL93] in this setting which can only handle an

$O(\ell/(\log \ell)^2)$-sized coalition of bad players. To compare with our lower bound, when $k = 2$ our lower bound from Theorem 1 shows any such protocol can be biased by a $O(\ell/(\log\log \ell))$ sized coalition of bad players. It is a very interesting open problem to settle the exact number of bad players that the best two-round coin flipping protocol can handle.

2 Proof Overview

We give a brief overview of our strategy for our new lower bound and upper bounds for k-round coin flipping protocols. In Section 2.1, we sketch our proof of lower bounds for coin flipping protocols. In Section 2.2, we prove lower bounds against protocols where players can send many bits per round. In Section 2.3, we prove a useful helper theorem that shows how any family of functions can be biased by a random set of variables and a small set of variables, specific to each function. Finally, in Section 2.5, we construct new improved constant-round coin flipping protocols.

2.1 Biasing Coin Flipping Protocols

We focus on the two-round setting which contains most of the core ideas behind our proof and allows us to simplify our analysis. For any two-round coin flipping protocol $\pi : (\{0,1\}^\ell)^2 \to \{0,1\}$ we show:

Theorem 2.1 (Two-round version of Theorem 5.1)**.** *For any 2-round coin flipping protocol π over ℓ players where $\Pr[\pi = 1] \geq 0.01$, there exists a set of bad players $B \subseteq [\ell]$ with $|B| = O\left(\frac{\ell}{\log\log \ell}\right)$ so that $\Pr[\pi|_B = 1] \geq 0.99$.*

The main tool that we will use to prove this theorem is that for any family $\mathcal{F}$ of functions from $\{0,1\}^\ell \to \{0,1\}$, there exists some common set of bad players $B_R \subset [\ell]$ such that for almost every function $f \in \mathcal{F}$, there exist a small set of "heavy" bad players B_H, depending on f, such that $B_R \cup B_H$ can bias f. Formally:

Theorem 2.2 (Simplified version of Theorem 4.1)**.** *Let $\mathcal{F}$ be any family of functions from $\{0,1\}^\ell \to \{0,1\}$ where $\Pr[f = 1] \geq 0.001$ for each $f \in \mathcal{F}$. Then, there exists a common bad set $B_R \subseteq [\ell]$ with $|B_R| = O\left(\frac{\ell}{\log\log \ell}\right)$ such that for 0.999 fraction of functions $f \in \mathcal{F}$, there exists a heavy bad set $B_H = B_H(f) \subseteq [\ell]$ with $|B_H| = (\log \ell)^{0.001}$ such that $B_R \cup B_H$ can 0.999-bias f towards 1.*

We will also require a simpler version of the above theorem that follows by inductively applying the KKL theorem [KKL88]:

Theorem 2.3 (KKL)**.** *Let $f : \{0,1\}^\ell \to \{0,1\}$ be such that $\Pr[f = 1] \geq 0.001$. Then, there exists a set of bad players B with $|B| \leq O\left(\frac{\ell}{\log \ell}\right)$ such that $\Pr[f|_B = 1] \geq 0.999$.*

We will sketch a proof of Theorem 2.2 in Section 2.3. Let us see how Theorem 2.1 follows from Theorem 2.2:

Proof Sketch for Theorem 2.1. In our proof we will separately find sets of bad players to corrupt in the first round and the second round, and at the end we will take the union of the two sets. In particular, we let B_I (initially empty) be the set of bad players that we will corrupt from the first round.

For $\alpha \in \{0,1\}^\ell$, let $\pi_\alpha : \{0,1\}^\ell \to \{0,1\}$ be the induced second round protocol when the players output α in the first round. Since $\Pr[\pi = 1] \geq 0.01$, by a reverse Markov argument (see Claim 3.3), we have that $\Pr[\pi_\alpha = 1] \geq 0.001$ for at least 0.009 fraction of α. Let $\mathcal{F}$ be the family of functions consisting of these functions π_α. We apply Theorem 2.2 to find a common set of bad players B_R such that for 99% of $\pi_\alpha \in \mathcal{F}$, there exists a set of bad players $g(\alpha)$ such that $B_R \cup g(\alpha)$ can 0.999-bias π_α towards 1. Moreover, $|B_R| \leq O\left(\frac{\ell}{\log\log\ell}\right)$ and $|g(\alpha)| = h = (\log \ell)^{0.001}$.

For each α for which there does not exist such a set of bad players that along with B_R can corrupt π_α, we let $g(\alpha) = \perp$. Then, we know that $\Pr[g \neq \perp] \geq 0.008$. We use Theorem 2.3 to find a set of bad players $B^{(1)}_{temp}$ with $\left|B^{(1)}_{temp}\right| \leq O\left(\frac{\ell}{\log \ell}\right)$ to bias the first round protocol so that $\Pr[g|_{B^{(1)}_{temp}} \neq \perp] \geq 0.999$. We add all players from $B^{(1)}_{temp}$ to B_I.

We view $g = (g_1, \ldots, g_h)$ where each $g_i : \{0,1\}^\ell \to [\ell] \cup \perp$ where g_i outputs the i-th largest element from $g(\alpha)$ and $g_i(\alpha) = \perp$ if and only if $g(\alpha) = \perp$. We will repeatedly use Theorem 2.3 to find bad players that can bias each of $g_1, \ldots, g_h$ so that over most inputs, each of their images lies in some set of size $c = \frac{\ell}{(\log \ell)^{1000}}$. Towards this goal, we maintain sets $C_1, \ldots, C_h$ that will satisfy the following invariant (initially each of these sets equals $[\ell]$):

> $(*)$ For every α, if there exists $i \in [h]$ such that $g_i(\alpha) \notin C_i$, then $g(\alpha) = \perp$. (Equivalently, if $g(\alpha) \neq \perp$, then $g_i(\alpha) \in C_i$ for all $i \in [h]$.)

We now find bad players to bias each of these g_i and while doing so, maintain that at the beginning of each iteration of the loop, $\Pr[g|_{B_I} \neq \perp] \geq 0.999$. Formally, we proceed as follows:

- While there exists $i \in [h]$ such that $|C_i| > c$ $\left(\text{Recall that } c = \frac{\ell}{(\log \ell)^{1000}}, h = (\log \ell)^{0.001}\right)$:
  1. Let $\mathbf{X}_i$ be a random subset of C_i with $|\mathbf{X}_i| = |C_i|/2$. Since $\Pr[g_i|_{B_I} \neq \perp] \geq 0.999$, by our invariant property, $\Pr[g_i|_{B_I} \in C_i] \geq 0.999$. Hence, $\mathbb{E}[\Pr[g_i|_{B_I} \in \mathbf{X}_i]] \geq 0.49$. In particular, there exists $C_i' \subset C_i$ with $|C_i'| = |C_i'|/2$ such that $\Pr[g_i|_{B_I} \in C_i'] \geq 0.49$. We set C_i to equal C_i'.
  2. To maintain our invariant, we can now only guarantee that $\Pr[g|_{B_I} \neq \perp] \geq 0.49$. To increase this, we apply Theorem 2.3 to the function $g|_{B_I}$ to find a set of bad players $B^{(2)}_{temp}$ with $\left|B^{(2)}_{temp}\right| \leq O\left(\frac{\ell}{\log \ell}\right)$ so that $\Pr[g|_{B_I \cup B^{(2)}_{temp}} \neq \perp] \geq 0.999$. We add all players from $B^{(2)}_{temp}$ to B_I.

Finally, we let $B_H = \bigcup_{i=1}^h C_i$. Our final set of bad players will be $B_R \cup B_H \cup B_I$. We see that $\Pr[g|_{B_I} \neq \perp] \geq 0.999$. Whenever $g|_{B_I}(\alpha) \neq \perp$, $\Pr[\pi_\alpha|_{B_R \cup B_H} = 1] \geq 0.999$. Hence, $\Pr[\pi_{B_R \cup B_H \cup B_I} = 1] \geq 0.999 \cdot 0.999 \geq 0.99$ as desired.

We finally bound the number of bad players we control. We know that $|B_R| \leq O\left(\frac{\ell}{\log\log\ell}\right)$. We also have that $|B_H| \leq c \cdot h = \frac{\ell}{(\log \ell)^{1000}} \cdot (\log \ell)^{0.001} < \frac{\ell}{(\log \ell)^{999}}$. Since we decrease the size of each C_i by a factor of 2 each time we execute the loop, and we stop once every $i \in [h]$ has $|C_i| \leq c$, the total number of iterations of the loop is $h \cdot \log(\ell/c) = (\log \ell)^{0.001} \cdot (1000 \log\log \ell) \leq (\log \ell)^{0.002}$. Since each time through the loop we add $O\left(\frac{\ell}{\log \ell}\right)$ players to B_I, we bound the size of B_I as $|B_I| \leq O\left(\frac{\ell}{\log \ell} + (\log \ell)^{0.002} \cdot \frac{\ell}{\log \ell}\right) \leq \frac{\ell}{(\log \ell)^{0.99}}$. Hence, the total number of bad players we control is $|B_R| + |B_H| + |B_I| \leq O\left(\frac{\ell}{\log\log\ell}\right)$. $\square$

We briefly mention that the above ideas more or less work to extend our two-round lower bound to k rounds; the main difference is that we inductively use our bounds to bias $(k-1)$-round protocols in place of the second use case of the KKL theorem above (see Item 2(a)iii in the "Formal description of algorithm" in Section 5).

2.2 Biasing Coin Flipping Protocols with Longer Messages

We now briefly sketch a proof for lower bounds for protocols where players can send more than one bit per round:

Theorem 2.4 (Simplified version of Theorem 5.5)**.** *Let π be a k-round coin flipping protocol over ℓ players where in round i, each player can send $r_i = (\log^{(i)} \ell)^{0.99}$ many bits. Then, there exists a set of bad players $B \subset [\ell]$ with $|B| \leq 0.01\ell$ so that $\Pr[\pi|_B = 1] \geq 0.99$.*

Proof sketch. We essentially follow the same proof strategy as in the k-round version of Theorem 2.1. To do this, we treat each round i of the protocol as being over $r_i \cdot \ell$ bits and whenever our lower bound asks us to corrupt a bit, we corrupt the corresponding player that controls that bit. For instance, in the one-round version of this theorem, we treat the input as being over $\ell \cdot r_1$ bits. Then, we use our lower bounds to find a set of $\frac{\ell \cdot r_1}{\log(\ell \cdot r_1)}$ bits such that if bad players control them, then they can bias the protocol. So, we let B be the set of all players that control each such bit. It must be the case that $|B| \leq \frac{\ell \cdot r_1}{\log(\ell \cdot r_1)}$. Since $|r_1| \leq (\log \ell)^{0.99}$, we end up controlling $o(\ell)$ players overall to corrupt π. This idea, with the right setting of parameters, generalizes to k-round protocols. □

2.3 Biasing a Family of Functions

We now sketch a proof of Theorem 2.2. Towards proving it, we will require the following lemma which is about biasing individual boolean functions:

Lemma 2.5 (Simplified version of Lemma 4.2)**.** *Fix $1 \leq h \leq \ell^{0.99}$ and let $f : \{0,1\}^\ell \to \{0,1\}$ be such that $\Pr[f = 1] \geq 0.01$. Then, for 0.99 fraction of $B_R \subset [\ell]$ with $|B_R| = O\left(\frac{\ell}{\log(h)}\right)$, there exists $B_H = B_H(B_R) \subset [\ell]$ with $|B_H| \leq h$ such that $B_R \cup B_H$ can 0.99-bias f towards 1.*

We prove this lemma later. Now we show how, using Lemma 2.5, we can prove Theorem 2.2.

Proof sketch of Theorem 2.2. We set $h = (\log \ell)^{0.001}$. Let $G = (U, V)$ be a bipartite graph where the left part U equals $\mathcal{F}$, and the elements of the right part V are all size-$\frac{\ell}{\log(h)}$ subsets of $[\ell]$. We add an edge between $f \in U$ and a set $R \in V$ if there exists $H \subset [\ell]$ with $|H| \leq h$ such that $R \cup H$ can 0.99-bias f towards 1. By Lemma 2.5, the degree of each $f \in U$ is at least $0.99 \cdot |V|$. Therefore, there exists $B_R \in V$ such that B_R has degree at least $0.99 \cdot |U|$. Such a B_R satisfies the conditions of the theorem. □

We now focus on proving Lemma 2.5. Towards proving this, we will need the following result regarding influence of boolean functions that was proven in [RSZ02] by slightly building up on the result of [KKL88]:

Lemma 2.6 (Simplified version of Lemma 4.4)**.** *Fix $1 \leq h \leq \ell^{0.99}$. Let $f : \{0,1\}^\ell \to \{0,1\}$ be such that $0.01 \leq \Pr[f = 1] \leq 0.99$. If $I_f(i) \leq \frac{1}{h}$ for all $i \in [\ell]$, then $\sum_{i=1}^{\ell} I_f(i) \geq \frac{\log(h)}{2000}$.*

We use Lemma 2.6 to establish Lemma 2.5 by analyzing a semi-random process that was first studied in [RSZ02]. Our analysis closely follows that of [RSZ02], and only differs in the final conclusion.

Proof sketch of Lemma 2.5. Let $r = \frac{C\ell}{\log(h)}$ where C is a very large universal constant. Consider the following semi-random process:

1. Initialize $B_R = B_H = \emptyset$. Repeat the following for r steps or until $\Pr[f|_{B_R \cup B_H} = 1] \geq 0.99$:
   (a) (Heavy Case) If there exists $i \in [\ell] \setminus (B_R \cup B_H)$ such that the influence of i on $f|_{B_R \cup B_H}$ is at least $\frac{2}{h}$, then add i to B_H.
   (b) (Random Case) Otherwise, pick a random $i \in [\ell] \setminus (B_R \cup B_H)$ and add i to B_R.

We say that the above process *succeeds* if at the end, $\Pr[f|_{B_R \cup B_H} = 1] \geq 0.99$. We will establish the following:

Claim 2.7 (Simplified version of Claim 4.5)**.** *The above process succeeds with probability at least* 0.999.

We will prove this claim later. For now we show how Lemma 2.5 follows from it. First, we observe that $|B_H| \leq h$ (since each player added to B_H 'pushes' f towards 1 by a factor of $1/h$, this can happen at most h times). Next, consider a random set $\mathbf{R} \subset [\ell]$ with $|\mathbf{R}| = \ell$ and a random permutation $\boldsymbol{\pi}$ of $[r]$ so that $(\mathbf{R}, \boldsymbol{\pi})$ fixes an ordering of the elements of $\mathbf{R}$. Consider the modified semi-random process where an initially randomly choice of $(\mathbf{R}, \boldsymbol{\pi})$ is made, and then in the random case of the process, the earliest element from $\mathbf{R}$ that is not in $B_R \cup B_H$ is chosen. At the end of the process, $B_R \subset \mathbf{R}$ and if the process succeeds, $B_H \cup \mathbf{R}$ can indeed bias f as desired. By Claim 2.7, the process succeeds with probability at least 0.999 and so for 0.999 fraction of choices of (R, π), the modified process succeeds. So, there exists π^* such that for 0.999 fraction of (R, π^*), the modified random process succeeds. Thus, for 0.999 fraction of sets R, there exists a set $B_H = B_H(R)$ (given by the process) such that $R \cup B_H$ can 0.99 bias f as desired.

We finally prove our claim that the process indeed succeeds with high probability:

Proof sketch of Claim 2.7. For $j \in [r]$, let v_j represent the variable chosen at step j of the semi-random process. Let $\mathbf{X}_j$ equal the influence of v_j on $f|_{\{v_1,\ldots,v_{j-1}\}}$ if the process hasn't stopped before step j and let $\mathbf{X}_j$ equal 1 otherwise. Let $j \in [r]$ be a step such that $\Pr[f|_{\{v_1,\ldots,v_{j-1}\}} = 1] \leq 0.99$. Then, by Lemma 2.6, in step j, either there exists a variable with influence $\frac{2}{h}$ or else a random variable will have expected influence $\geq \frac{\log(h)}{2000\ell}$. Since $\frac{1}{h} \geq \frac{\log(h)}{2000\ell}$ (recall that we assumed $h \leq \ell^{0.99}$), we always have that $\mathbb{E}[\mathbf{X}_j | \mathbf{X}_1, \ldots, \mathbf{X}_{j-1}] \geq \frac{\log(h)}{2000\ell}$.

For $j \in [r]$, let $\mathbf{Z}_j = \sum_{k \leq j} \mathbf{X}_k$. Then, $\mathbf{Z}_1, \ldots, \mathbf{Z}_r$ forms a submartingale with

$$\mathbb{E}[\mathbf{Z}_r | \mathbf{Z}_1, \ldots, \mathbf{Z}_{r-1}] = r \cdot \frac{\log(h)}{2000\ell} = C/2000 \geq 10^6.$$

We then apply Azuma's inequality (see Lemma 4.3) to infer that $\mathbf{Z}_r \geq 2$ with probability at least 0.999. Since $\mathbf{Z}_r$ represents the sum total of 'contributions' of influences of variables towards 1, whenever $\mathbf{Z}_r \geq 2$ it must be that f has been 'pushed' towards 1 by a 'total amount' of 1. In all such cases, $\Pr[f|_{B_R \cup B_H} = 1] \geq 0.999$ and the process must succeed. Since this happens with probability at least 0.999, the process succeeds with at least this probability. ☐

This concludes the proof sketch of Lemma 2.5. ☐

2.4 Comparison to Proof Strategy of [RSZ02]

Since our results are inspired by the work of [RSZ02] and our proofs build on their arguments, we here briefly sketch their strategy for 2 rounds and how we improve upon it.

Let $\mathcal{F} = \{\pi_\alpha\}$ where $\alpha \in \{0,1\}^\ell$ and let π_α be the induced second round protocol when in the first round the players output α. [RSZ02] proceeds by finding a set of bad players in the second round to bias the functions in $\mathcal{F}$. To do this, as hinted at earlier, [RSZ02] analyzes the same semi-random process as in Lemma 2.5. At the end of the process, they also obtain a common random set B_R and a specialized heavy set $B_H = B_H(f)$ such that for most functions $f \in \mathcal{F}$, $B_R \cup B_H(f)$ can bias the function f. They then pick another random set $B_{R'}$ and bound the probability that $B_H(f) \subset B_{R'}$ for fixed f, and argue that in expectation, $\delta \approx \left(\frac{r}{\ell}\right)^{2^{\ell/r}}$ fraction of functions $f \in \mathcal{F}$ will be such that $B_H(f) \subset B_{R'}$ where $|B_{R'}| = r$. They let $B_R \cup B_{R'}$ be the set of bad players that will bias the second round protocols.

Now, let $S \subset \{0,1\}^\ell$ be the set of $\alpha \in \{0,1\}^\ell$ such that π_α can be biased by $B_R \cup B_{R'}$, and let $g : \{0,1\}^\ell \to \{0,1\}$ be such that $g(\alpha) = 1 \iff \alpha \in S$. We know from above that $\Pr[g = 1] \geq \delta$. Now, the goal for the first round is to find bad players to bias g towards 1 so that with high probability, the resulting second round protocol can be biased using $B_R \cup B_{R'}$. To do this, [RSZ02] use the KKL theorem [KKL88] which lets them do this by finding $O(\ell/(\delta \log \ell))$ bad players. To minimize the total number of bad players, they end up using $O(\ell/\log^{(3)} \ell)$ bad players. For multiple rounds, their result follows by an inductive argument, showing that $O\left(\ell/\log^{(2k-1)} \ell\right)$ bad players suffice to bias k round protocols.

Our key insight is that instead of first committing to the set $B_{R'}$ that covers a small non-trivial fraction of $\mathcal{F}$ and then biasing the first round to ensure that $B_R \cup B_{R'}$ can bias most of the resulting second round protocols, we can delay committing and instead *commit incrementally*. We do this by initially setting $B_{R'} = [\ell]$, so that it can bias most of $\mathcal{F}$; we then "chisel away" half of $B_{R'}$, so that we can only bias half of the functions from $\mathcal{F}$. To fix this, we use the KKL theorem so that out of the set of π_α functions that are being considered, we again can bias most of them. We repeatedly do this until $B_{R'}$ becomes very small, while always maintaining that we can bias most of the resultant second round protocols. Doing this gives us an exponential improvement, with $O\left(\ell/\log^{(2)} \ell\right)$ bad players sufficing to bias two-round protocols.

To illustrate in a simple setting why delaying committing can give an exponential improvement, we consider the following scenario: Suppose that we are given a function $f : \{0,1\}^\ell \to \{0,1\}^t$ and our goal is to bias f so that it outputs some single element y from $\{0,1\}^t$ with high probability. One way to do this is to fix an element $y \in \{0,1\}^t$ that appears with probability at least $\frac{1}{2^t}$, then use the KKL theorem to obtain a set of bad players that can bias f to ensure that y is output with probability 0.99. While this works, the upper bound on the size of the set of bad players resulting from this approach is $O\left(2^t \cdot \frac{\ell}{\log \ell}\right)$ (see Theorem 5.3). In contrast, a "delayed commitment" approach is to maintain a set S (initially setting $S = \{0,1\}^t$); cut S in half; use KKL to find $O\left(\frac{\ell}{\log \ell}\right)$ bad players that will ensure that at least 0.99 fraction of the inputs lie in S; and repeat. We do this cutting process t times, so that at the end of it $|S| = 1$, and we let its sole element be the final output string y. With this approach, the number of bad players required used will be only $O\left(t \cdot \frac{\ell}{\log \ell}\right)$, which is an exponential improvement over $O\left(2^t \cdot \frac{\ell}{\log \ell}\right)$ in terms of the dependence on t.

2.5 Constructing Coin Flipping Protocols

We now move on to sketching the main ideas involved in constructing better coin-flipping protocols in the setting that each player sends one bit per round. For simplicity, we focus on the setting of two-round protocols, which will give a sense of how our general construction (for k rounds) works.

Theorem 2.8 (Two-round version of Theorem 6.1)**.** *For any $\gamma > 0$ there exists an explicit two-round coin flipping protocol over ℓ players such that when the number of bad players is at most $\frac{\gamma\ell}{(\log\ell)(\log\log\ell)^2}$, the output coin flip is ε-close to uniform where $\varepsilon = O\left(\gamma + (\log\ell)^{-0.2}\right)$.*

To help construct such protocols, we use the explicit resilient functions of [IV24]:

Theorem 2.9 (Explicit resilient function from [IV24])**.** *There exists an explicit 1 round protocol $f : \{0,1\}^n \to \{0,1\}$ such that when b out of the n players are bad, the resultant output bit is $O\left(\frac{b(\log n)^2}{n} + n^{-0.99}\right)$ close to uniform.*

Using this, our two-round protocol proceeds as follows:

Proof sketch of Theorem 2.8. We partition the ℓ players arbitrarily into parts $P_1, \ldots, P_{\ell/\log\ell}$ where $|P_i| = \log\ell$. Players in the same part will act as a single entity that can make $\log\ell$ many coin flips. We say entity P_i is *good* if all players that are part of the entity are good. We say the entity is *bad* otherwise. Since at most $\frac{\gamma\ell}{(\log\ell)(\log\log\ell)^2}$ players are bad, the number of bad entities is at most $\frac{\gamma\ell}{(\log\ell)(\log\log\ell)^2}$.

We now deploy the lightest bin protocol of [Fei99] on the entities in the first round. In particular, we introduce 'bins' $B_1, \ldots, B_{\ell/(\log\ell)^3}$. We ask each entity to 'vote' for a bin by outputting a random number between 1 and $\ell/(\log\ell)^3$ (this can be done since each entity has access to $\log\ell$ random bits). Let $i^* \in [\ell/(\log\ell)^3]$ be such that among all the bins, bin B_{i^*} "is the lightest" (has the smallest number of entities voting for it). Let $S \subset [\ell/\log\ell]$ be the set of entities that voted for B_{i^*}.

In the second round, we apply the explicit resilient function from Theorem 2.9 to the entities in S (asking them to output 1 bit) and let that be the output of the protocol.

We now analyze this protocol. First, since there are $\ell/(\log\ell)^3$ bins and $\ell/\log\ell$ entities, the lightest bin B_{i^*} will have at most $(\log\ell)^2$ players voting for it, i.e., $|S| \le (\log\ell)^2$. Next, since the number of good entities is at least $\frac{\ell}{\log\ell} - \frac{\gamma\ell}{(\log\ell)(\log\log\ell)^2}$, by a Chernoff bound (see Claim 3.4) and a union bound, with high probability, every bin will have roughly $\frac{\frac{\ell}{\log\ell} - \frac{\gamma\ell}{(\log\ell)(\log\log\ell)^2}}{\frac{\ell}{(\log\ell)^3}} = (\log\ell)^2 - \frac{\gamma(\log\ell)^2}{(\log\log\ell)^2}$ good entities. Equivalently, the number of bad entities in bin B_{i^*} is at most $\frac{\gamma(\log\ell)^2}{(\log\log\ell)^2}$ (out of roughly $(\log\ell)^2$ entities in it). Hence, when we apply Theorem 2.9 with $n \approx (\log\ell)^2$ players, the output coin flip will be $O(\gamma)$ close to the uniform distribution (the extra $(\log\ell)^{-0.2}$ error term is because of the error in the Chernoff bound). $\square$

Organization We introduce necessary preliminaries in Section 3. In Section 4, we prove our main technical result on jointly biasing most functions in any family of functions. We prove our main lower bound result on the resilience of coin flipping protocols in Section 5. In Section 6, we present new coin flipping protocols. We conclude with some discussion and open problems in Section 7.

3 Preliminaries

3.1 Notation and Terminology

For $0 < \varepsilon < 1/2$ and $o \in \{0,1\}$, we say that $S \subset [\ell]$ can $(1-\varepsilon)$-*bias a function* $f : \{0,1\}^\ell \to \{0,1\}$ *towards* o if there exists an adversary $\mathcal{A}_o : \{0,1\}^{\ell-|S|} \to \{0,1\}^{|S|}$ such that the function $g_o(y) := f(\mathcal{A}_o(y), y)$ has $\Pr_{\mathbf{y}}[g_o(\mathbf{y}) = o] \geq 1-\varepsilon$. We will similarly use the notion that a set $S \subseteq [\ell]$ can $(1-\varepsilon)$-bias a protocol π towards o.

For a function $f : \{0,1\}^\ell \to \{0,1\}, S \subset [\ell]$, and $\mathcal{A} : \{0,1\}^{\ell-|S|} \to \{0,1\}^{|S|}$, we define $f|_S : \{0,1\}^{\ell-|S|} \to \{0,1\}$ to be $f|_S(x) = f(\mathcal{A}(x), x)$. Often when we do this, $\mathcal{A}$ will be implicitly defined by some other claim, and hence the notation reflects this by only mentioning f and S but not $\mathcal{A}$. We similarly extend this notation to protocols so that for $S \subset [\ell]$, a protocol π, and some adversarial function $\mathcal{A}$, the protocol π_S is well defined.

For a function $f : \{0,1\}^\ell \to \{0,1\}$, we sometimes write "$\Pr[f = o]$" to denote $\Pr_{\mathbf{y}\sim\{0,1\}^\ell}[f(\mathbf{y}) = o]$. Similarly, for π a k-round protocol over ℓ players in which each player outputs one bit per round, we sometimes write "$\Pr[\pi = o]$" to denote $\Pr_{\mathbf{x}\sim(\{0,1\}^\ell)^k}[\pi(\mathbf{x}) = o]$.

3.2 Probability

3.2.1 Useful Definitions

Definition 3.1. *A* submartingale *is a sequence of real valued random variables* $\mathbf{Z}_0, \mathbf{Z}_1, \ldots,$ *for which* $\mathbb{E}[\mathbf{Z}_i | \mathbf{Z}_{i-1}] \geq \mathbf{Z}_{i-1}$.

We will also need the following definition of influence of a variable with respect to a boolean function:

Definition 3.2 (Influence)**.** *For* $f : \{0,1\}^n \to \{0,1\}$ *and* $i \in [n]$*, we say that the* influence *of coordinate* i *on* f*, denoted* $I_i(f)$*, is*

$$\Pr_{\mathbf{x}\sim\mathbf{U}_{i-1},\mathbf{y}\sim\mathbf{U}_{n-i}}[|f(\mathbf{x},1,\mathbf{y}) - f(\mathbf{x},0,\mathbf{y})|], \quad \textit{or equivalently,} \quad \Pr_{\mathbf{z}\sim\mathbf{U}_n}[f(\mathbf{z}) \neq f(\mathbf{z}^{\oplus i})].$$

3.2.2 Helpful Claims

We will make use of the following reverse Markov style inequality:

Claim 3.3 (Reverse Markov)**.** *Let* $\mathbf{X}$ *be a random variable taking values in* $[0,1]$*. Then, for* $0 \leq p < \mathbb{E}[X]$*, it holds that*

$$\Pr[\mathbf{X} > p] \geq \frac{\mathbb{E}[\mathbf{X}] - p}{1-p}.$$

We will also use the following lower tail Chernoff bound:

Claim 3.4 (Lower tail Chernoff bound)**.** *For* $n \geq 1$ *and* $p \in [0,1]$*, let* $\mathbf{X}_1, \ldots, \mathbf{X}_n$ *be independent random variables such that for each* $i \in [n]$*,* $\Pr[\mathbf{X}_i = 1] = p$ *and* $\Pr[\mathbf{X}_i = 0] = 1-p$*. Let* $\mathbf{X} = \sum_{i=1}^n \mathbf{X}_i$ *and let* $\mu = \mathbb{E}[\mathbf{X}] = pn$*. Then, for all* $\delta \in (0,1)$*:*

$$\Pr[\mathbf{X} \leq (1-\delta)\mu] \leq \exp(-\delta^2\mu/2).$$

3.3 Leader Election and Collective Coin Flipping Protocols

3.3.1 Protocols in the full information model

We formalize the definition of protocols in the full information model. Collective coin flipping protocols and leader election protocols are special cases of such protocols where the output domain is $\{0,1\}$ and $[\ell]$ respectively.

Definition 3.5 (Protocol in the full information model)**.** *A* k-round protocol with output domain Y over ℓ players where each player sends r random bits per round *is a function*

$$\pi : \left((\{0,1\}^r)^\ell \right)^k \to Y$$

that takes in the input of each of the players during each round and outputs an element from set Y *which is the outcome of the protocol.*

Here is how the protocol operates in the presence of a set $B \subset [\ell]$ *of bad players: In round* i*, each of the players from* $[\ell] \setminus B$ *independently outputs a uniformly random element from* $\{0,1\}^r$*. Let their collective outputs be* $\alpha_i \in (\{0,1\}^r)^{[\ell]\setminus B}$*. Then, depending on* $\alpha_1, \dots, \alpha_i$*, the players in* B *together output an element of* $(\{0,1\}^r)^B$*. Hence, we model the strategy of the bad players as a sequence of functions* $\sigma = (\sigma_1, \dots, \sigma_k)$*, where*

$$\sigma_i : \left((\{0,1\}^r)^{[\ell]\setminus B} \right)^i \to (\{0,1\}^r)^B ,$$

where σ_i *takes in the inputs of the good players from the first* i *rounds and maps it to the output of the bad players for round* i*. For a fixed strategy* σ*, the outcome of the protocol can be modeled as follows: uniform random strings* $\alpha_1, \dots, \alpha_k \in (\{0,1\}^r)^{[\ell]\setminus B}$ *are chosen, and the outcome of the protocol is*

$$\pi(\alpha_1 : \sigma_1(\alpha_1), \alpha_2 : \sigma_2(\alpha_1, \alpha_2), \dots, \alpha_k : \sigma_k(\alpha_1, \dots, \alpha_k)).$$

We now specialize this definition to define collective coin flipping protocols

Definition 3.6 (Collective coin flipping protocol)**.** *A* collective coin flipping protocol π *is a protocol in the full information model with output domain* $Y = \{0,1\}$*. Furthermore, we say* π *is* (b,γ)-resilient *if in the presence of any set* B *of bad players with* $|B| \leq b$*, we have that* $\max_{o\in\{0,1\}} \Pr[\pi|_B = o] \leq 1-\gamma$*.*

Remark 3.7. *Typically in the pseudorandomness literature, the quality of a coin flip is measured by its distance to the uniform distribution. The definition of resilience that we use, which is standard in the leader election and collective coin flipping literature, has a weaker requirement that each outcome has probability at most* $1-\gamma$*. Our lower bounds results rule out, for any small* γ*, that there exist coin flipping protocols that are* (b,γ)*-resilient, with tradeoffs between the number of "bad" players* b *and the number of rounds* k*. On the other hand, our positive results (collective coin flipping protocols) satisfy the stronger measure of quality that is standard in the pseudorandomness literature: their output is* ε*-close to the uniform distribution over* $\{0,1\}$*, for small* ε*, even in the presence of bad players who are colluding using any strategies (with tradeoffs between the number of rounds, the number of bad players, and the closeness to the uniform distribution).*

In this paper we will mostly consider the case in which $r = 1$ and this will be the default assumption unless stated otherwise. Note that when $k = 1$, the protocol π just becomes a function over $\{0,1\}^\ell$; such 1-round coin flipping protocols which cannot be biased by any small set of bad players are also known as *resilient functions*.

Remark 3.8. *In this paper, when showing that some subset of bad players can corrupt any coin flipping protocol, we will often construct strategies in which some subset of bad players output a random string in some rounds. By an averaging argument we can always turn such a strategy into one in which each bad player outputs a deterministic string in every round, while maintaining the bias of the protocol.*

We also specialize the definition of protocols to define leader election protocols:

Definition 3.9 (Leader election protocol)**.** *A* leader election protocol π *is a protocol in the full information model with output domain* $Y = [\ell]$*, the number of players the protocol is operating on. Furthermore, we say* π *is* (b,γ)-resilient *if in the presence of any set* B *of bad players with* $|B| \le b$*, we have that* $\Pr[\pi|_B \in B] \le 1-\gamma$*.*

3.3.2 A Useful Claim

We record the following well-known claim which states that any leader election protocol implies a collective coin flipping protocol. For completeness, we supply a proof.

Claim 3.10. *Fix* $\gamma \in (0,1/2]$*. Let* π *be a* k*-round leader election protocol where each player sends* r *bits per round and where in presence of any* b *bad players and their colluding strategy,* π *guarantees that a good leader is chosen with probability at least* γ*. Then, there exists a* $(k+1)$*-round collective coin flipping protocol* π' *that is* $(b,\gamma/2)$*-resilient in which each player sends* r *bits per round.*

Proof. Let π' be the protocol that executes the protocol π in the first k rounds and in round $k+1$ asks the elected leader to flip a coin. Formally, say the messages sent in the first k rounds in π' are α. Let i be the index of the player that is chosen as the leader by π on input α. Then, in round $k+1$, π' outputs the first bit of the message sent by player i.

We see that whenever π selects a good leader, π' outputs a truly random coin toss, giving us the desired resilience parameter for π'. □

An immediate consequence is that proving lower bounds for con flipping protocols also gives lower bounds for leader election protocols.

Corollary 3.11. *Fix* $\gamma \in (0,1/2]$*. Suppose that for every* k *round coin flipping protocol* π' *over* ℓ *players, there exist* b *bad players that can* $(1-\gamma)$*-bias* π *towards a particular outcome* $o \in \{0,1\}$*. Then, for every* $(k-1)$*-round leader election protocol* π*, there exist* b *bad players such that that a good player is elected as a leader with probability at most* 2γ*.*

4 Biasing a Family of Functions

In this section, we show that for any family of functions mapping $\{0,1\}^\ell \to \{0,1\}$, there exists a "common set" B_R of coordinates in $[\ell]$ such that for almost every function f in the family, B_R along with a small "heavy set" of coordinates (which may depend on f) can together bias f. Moreover, crucially, neither B_R nor the "heavy set" is too large. This result plays an essential role in our main lower bound for k-round coin flipping protocols, just as its simplified version, Theorem 2.2, played a crucial role in the sketch of our two-round lower bound given in Section 2.

Formally, we show the following:

Theorem 4.1. *Let $0 < \gamma, \delta < 1/2, h \in \mathbb{N}, \ell \in \mathbb{N}$ be such that $h \log(h/2) < 40\ell/\gamma$ and $8 \leq h \leq \ell$. Let $\mathcal{F}$ be any family of functions from $\{0,1\}^\ell$ to $\{0,1\}$, and for each $f \in \mathcal{F}$ let $o_f \in \{0,1\}$ be such that $\Pr[f = o_f] \geq \gamma$. Then, there exists $B_R \subset [\ell]$ with $|B_R| \leq \frac{100\ell \log(1/\delta)}{\gamma \log(h)}$ such that for $(1-\delta)$ fraction of functions $f \in \mathcal{F}$, there exists $B_H = B_H(f) \subset [\ell]$ with $|B_H| \leq h$ such that $B_R \cup B_H$ can $(1-\gamma)$-bias f towards o_f.*

We will use the following main lemma to prove Theorem 4.1:

Lemma 4.2. *As in Theorem 4.1 let $0 < \gamma, \delta < 1/2, h \in \mathbb{N}, \ell \in \mathbb{N}$ be such that $h \log(h/2) < 40\ell/\gamma$ and $8 \leq h \leq \ell$. Let $f : \{0,1\}^\ell \to \{0,1\}$ and $o \in \{0,1\}$ be such that $\Pr[f = o] \geq \gamma$. Then, for $(1-\delta)$ fraction of $B_R \subset [\ell]$ with $|B_R| = \frac{100\ell \log(1/\delta)}{\gamma \log(h)}$, there exists $B_H = B_H(B_R) \subset [\ell]$ with $|B_H| \leq h$ such that $B_R \cup B_H$ can $(1-\gamma)$-bias f towards o.*

Let us see how Lemma 4.2 yields Theorem 4.1:

Proof of Theorem 4.1. Let G be the bipartite graph with left vertex set U given by $\mathcal{F}$ and right vertex set V consisting of all size-r subsets of $[\ell]$, where $r = \frac{100\ell \log(1/\delta)}{\gamma \log(h)}$. G contains an edge between $f \in U$ and $B_R \in V$ if there exists $B_H = B_H(f, B_R) \subset [\ell]$ with $|B_H| \leq h$ such that $B_R \cup B_H$ can $(1-\gamma)$-bias f towards o_f. By Lemma 4.2, the minimum left degree of this graph is at least $(1-\delta)\binom{\ell}{r}$. This implies that the average right degree of this graph is at least $(1-\delta)|\mathcal{F}|$, and hence, there exists $B_R^* \in V$ with degree at least $(1-\delta)|\mathcal{F}|$. Thus, there must exist some fixed B_R^* such that for $(1-\delta)$ fraction of $f \in \mathcal{F}$, there exists $B_H = B_H(f)$ with $|B_H| \leq h$ such that $B_R^* \cup B_H$ can $(1-\gamma)$-bias f towards o_f. ☐

4.1 Proving the Main Lemma

In this subsection we will prove Lemma 4.2.

We will use the following concentration inequality regarding submartingales from [RSZ02]:

Lemma 4.3 (Lemma 9 of [RSZ02])**.** *Let $0 < \mu < 1, 0 < \eta < 1, \ell \in \mathbb{N}$ be arbitrary. Let $\mathbf{Z}_0, \mathbf{Z}_1, \ldots, \mathbf{Z}_\ell$ form a submartingale with $\mathbf{Z}_0 = 0$, and suppose that for $i \in [\ell]$ we have $\mathbf{Z}_i - \mathbf{Z}_{i-1} \in [0,1]$ and $\mathbb{E}[\mathbf{Z}_i - \mathbf{Z}_{i-1}] \geq \mu$. Then,*

$$\Pr[\mathbf{Z}_\ell < (1-\eta)\ell\mu] < e^{-\eta^2 \mu \ell / 2}.$$

We will also use the following result from [RSZ02], which follows as a slight extension of [KKL88]:

Lemma 4.4. *Let $0 < \gamma < \frac{1}{2}, 0 < \theta < \frac{1}{8}, \ell \in \mathbb{N}$. Let $f : \{0,1\}^\ell \to \{0,1\}$ be such that $\gamma \leq \Pr[f = 1] \leq 1 - \gamma$. If $I_i(f) \leq \theta$ for all $i \in [\ell]$, then*

$$\sum_{i=1}^{\ell} I_i(f) \geq \frac{\gamma \log(1/\theta)}{20}.$$

Let us see how Lemma 4.2 follows using these results:

Proof of Lemma 4.2. Without loss of generality, assume $o = 1$. Let $r = \frac{100\ell \log(1/\delta)}{\gamma \log(h/2)}$. Consider the following semi-random procedure.

1. Initialize $B_R \leftarrow \emptyset, B_H \leftarrow \emptyset$. Do the following for r steps or until $\Pr[f|_{B_R \cup B_H} = 1] \geq 1 - \gamma$:
   (a) (Heavy Case) If there exists $i \in [\ell] \setminus (B_R \cup B_H)$ with $I_i(f|_{B_R \cup B_H}) \geq \frac{2}{h}$, then add i to B_H.
   (b) (Random Case) Otherwise, pick a random $i \in [\ell] \setminus (B_R \cup B_H)$ and add i to B_R.

We say that this procedure *succeeds* if at the end of this process, $B_R \cup B_H$ can $(1-\gamma)$-bias f. We will show that this procedure succeeds with high probability:

Claim 4.5. *With probability $\geq 1 - \delta$ over the above process, $B_R \cup B_H$ can $(1-\gamma)$-bias f towards 1.*

We will prove this claim later. Using this, we now show that for $1 - \delta$ fraction of B_R, there exists B_H with $|B_H| \leq h$ such that $B_R \cup B_H$ can $(1-\gamma)$-bias f towards 1.

Firstly, since every element that is added to B_H increases the probability of outputting 1 by $\frac{1}{h}$, it is always the case that $|B_H| \leq h$. We now prove the remaining statement.

Consider a random set $\mathbf{R} \subset [\ell]$ with $|\mathbf{R}| = r$ and a random permutation $\boldsymbol{\pi}$ of $[r]$ so that $(\mathbf{R}, \boldsymbol{\pi})$ determines an ordering of the elements from $\mathbf{R}$. For any fixed (R, π), consider the modified random process from above where at each step in the Random case, instead of picking a truly random element, we pick the earliest element from R that has not yet been added to $B_R \cup B_H$. Since the process always ends with $|B_R| \leq r$, this is a well defined operation and we will always have that $B_R \subset R$. Hence if the process succeeds, we have that $R \cup B_H$ can $(1-\gamma)$-bias f towards 1.

By Claim 4.5, we know that for $(1-\delta)$ fraction of choices of (R, π), the above process succeeds. By an averaging argument, this implies there exists a fixed permutation π^* such that for $(1-\delta)$ fraction of choices of R, it holds that $R \cup B_H$ (where $B_H = B_H(R, \pi^*)$) $(1-\gamma)$-biases f towards 1, as desired.

We finally prove the remaining claim:

Proof of Claim 4.5. For $j \in [r]$, let the variable chosen at each step j of the above semi-random process be v_j. For $j \in [r]$, let:

$$\mathbf{X}_j = \begin{cases} I_{v_j}(f|_{v_1,\ldots,v_{j-1}}) & \text{if } \Pr[f|_{v_1,\ldots,v_{j-1}} = 1] < 1 - \gamma \\ 1 & \text{otherwise} \end{cases}$$

At every step $j \in [r]$ where $\Pr[f|_{v_1,\ldots,v_{j-1}} = 1] < 1 - \gamma$ (note that since initially $\Pr[f = 1] \geq \gamma$, this also holds at all steps j as well), by Lemma 4.4, either there exists a variable with influence $\frac{2}{h}$ or the sum of influences of all the variables is at least $\frac{\gamma \log(h/2)}{20}$. By our choice of h we have that $\frac{2}{h} \geq \frac{\gamma \log(h/2)}{20\ell}$, and so for all $j \in [r]$, $\mathbb{E}[\mathbf{X}_j | \mathbf{X}_1, \ldots, \mathbf{X}_{j-1}] \geq \frac{\gamma \log(h/2)}{20\ell}$.

For $j \in [r]$, let $\mathbf{Z}_j = \sum_{k \leq j} \mathbf{X}_k$. We then see that $\mathbf{Z}_1, \ldots, \mathbf{Z}_r$ form a submartingale with $\mathbf{Z}_0 = 0$. We also have that

$$\mathbb{E}[\mathbf{Z}_r | \mathbf{Z}_1, \ldots, \mathbf{Z}_{r-1}] = \frac{r\gamma \log(h/2)}{20\ell} \geq 4 \log(1/\delta),$$

where in the last inequality, we used the fact that $r = \frac{100\ell \log(1/\delta)}{\gamma \log(h/2)}$. Applying Lemma 4.3 (with $\eta = \frac{1}{2}$), we infer that

$$\Pr[\mathbf{Z}_r < 2\log(1/\delta)] \leq e^{-1/4(4\log(1/\delta))} < (1/e)^{\log(1/\delta)} < \delta.$$

Since $\delta < \frac{1}{2}, 2\log(1/\delta) > 2$. Hence, with probability at least $1-\delta$, we have $\mathbf{Z}_r > 2$.

We claim that whenever this happens, we must be in the case that $\Pr[f|_{B_R \cup B_H} = 1] \geq 1-\gamma$. For $1 \leq j \leq r$, let the random variables $\mathbf{X}_j, \mathbf{Z}_j$ take on values x_j, z_j respectively. Let $j^* \leq r$ be such that at the end of that step, $\Pr[f|v_1, \ldots, v_{j^*} = 1] \geq 1-\gamma$; if this doesn't happen by then, let $j^* = r+1$. Then, since we picked an influential variable and bias it towards 1, for $1 \leq j \leq j^*$ we have that

$$\Pr[f|v_1, \ldots, v_{j-1}, v_j = 1] \geq \Pr[f|v_1, \ldots, v_{j-1} = 1] + x_j/2.$$

So, for all $1 \leq j \leq j^*$,

$$\Pr[f|v_1, \ldots, v_{j-1}, v_j = 1] \geq \gamma + z_j/2$$

Since the probability of any event is always less than 1, and we know that $z_r \geq 2$, it must be the case that $j^* \leq r$. So, we have that at the end of step j^*, $\Pr[f|v_1, \ldots, v_{j^*} = 1] \geq 1-\gamma$ and the process succeeded. Hence, the process indeed succeeds with probability $1-\delta$ as desired. This concludes the proof of Claim 4.5. ☐

This concludes the proof of Lemma 4.2. ☐

5 Biasing Coin Flipping Protocols

We will prove the following lower bound regarding coin flipping protocols:

Theorem 5.1. *There exist universal constants $C = 10^7, \ell_0 \in \mathbb{N}$ such that for all $k \in \mathbb{N}, \ell \in \mathbb{N}, 0 < \gamma < 1/4$, the following holds: For any k-round coin flipping protocol π over $\ell \geq \ell_0$ players where $\Pr[\pi = 1] \geq \gamma$, there exists a set of bad players $B \subset [\ell]$ with $|B| \leq \frac{C\ell}{\gamma \log^{(k)}(\ell)}$ such that $\Pr[\pi|_B = 1] \geq 1-\gamma$.*

From this, we obtain the following lower bound on the number of rounds required for coin flipping in the presence of a linear sized coalition of bad players:

Corollary 5.2. *There exist universal constants ℓ_0, C such that for all $\ell \in \mathbb{N}$ where $\ell \geq \ell_0$, the following holds: Let π be a k-round coin flipping protocol where $k \leq \log^*(\ell) - C$. Then, there exists a a set of bad players $B \subset [\ell]$ with $|B| \leq 0.01\ell$ and $o \in \{0,1\}$ such that $\Pr[\pi|_B = o] \geq 0.999$.*

To prove Theorem 5.1, we will primarily utilize Theorem 4.1. We will also use a simpler result that can be established by inductively applying the KKL theorem [KKL88]:

Theorem 5.3 (KKL)**.** *There exists a universal constant $C \leq 10^7$ such that for all $\ell \in \mathbb{N}, \gamma > 0$ the following holds: Let $f : \{0,1\}^\ell \to \{0,1\}$ be such that $\Pr[f = 1] \geq \gamma$. Then, there exists a set of bad players B with $|B| \leq \frac{C\ell}{\gamma \log(\ell)}$ such that $\Pr[f|_B = 1] \geq 1-\gamma$.*

Using this, we finally prove our main theorem:

Proof of Theorem 5.1. We apply induction on $k \geq 1$. For $k = 1$, the result follows by Theorem 5.3 For $k \geq 2$, we present an 'algorithm' which, given as input the protocol π, produces a small coalition of bad players that $(1-\gamma)$-biases π towards 1.

Setup. We will maintain two kinds of bad players: First, let $B_R, B_H \subset [\ell]$ be the sets of bad players that we will use to bias round k. Second, we let $B_I \subset [\ell]$ be the set of bad players that we will use to bias rounds 1 to $k-1$ (these will be provided to us by induction). At the end, we will take the union of all these sets, and that will be the final set of bad players that we will use to bias π. Let $h, c \in [\ell]$ be parameters that we set later (for intuition, h will be set to be very small and c will be set to be very large).

Given $\alpha \in (\{0,1\}^\ell)^{(k-1)}$, we write $\pi_\alpha(x)$ to denote the function from $\{0,1\}^\ell$ to $\{0,1\}$ that corresponds to the output of the protocol when the bits that the ℓ players output in rounds $1, \ldots, k-1$ are given by α and the bits that the players output in round k are given by x.

Informal description of algorithm. Since $\Pr[\pi = 1] \geq \gamma$, for a non-trivial fraction of $\alpha \in (\{0,1\}^\ell)^{k-1}$, the induced function $\pi_\alpha : \{0,1\}^\ell \to \{0,1\}$ has a non-trivial probability of outputting 1. For the family of functions consisting of all such π_α, we use Theorem 4.1 to find a set B_R such that for most such α, there exists a set $H = g(\alpha)$ with $|H| = h$ such that $B_R \cup g(\alpha)$ can bias π_α towards 1. We then use induction for the first time, biasing players from the first $k-1$ rounds, so that the overall fraction of α for which $B_R \cup g(\alpha)$ can bias π_α towards 1 becomes at least $3/4$.

Note that the function $g : (\{0,1\}^\ell)^{k-1} \to 2^{[\ell]}$ can be viewed as $g(\alpha) = (g_1(\alpha), \ldots, g_h(\alpha))$ where each g_j maps $(\{0,1\}^\ell)^{k-1}$ to $[\ell]$ and the value $g_j(\alpha)$ is simply the j-th largest element of the set $g(\alpha)$. We will inductively find bad players that can bias each of $g_1, \ldots, g_h$ so that over most inputs, their image lies in a set of size c.

We maintain sets $C_1, \ldots, C_h$ so that each of $g_1, \ldots, g_h$ outputs an element from these respective sets with high probability. Initially all these sets are set to equal $[\ell]$. We then iteratively find some C_j such that $|C_j| \geq c$ and cut C_j in half. Doing this decreases by a factor of 2 the fraction of inputs α whose outputs fall into the sets $C_1, \ldots, C_h$; so the fraction of inputs α for which $B_R \cup g(\alpha)$ can corrupt π_α is halved. To fix this, we use induction for a second time and find bad players that can corrupt the first $k-1$ rounds so that this fraction again becomes at least $3/4$.

At the end of this process, we are guaranteed that $|C_1| \leq c, \ldots, |C_h| \leq c$. We then once again for a third time use induction and find bad players to bias the first $k-1$ rounds so that the fraction of α for which $B_R \cup g(\alpha)$ can corrupt π_α is as large as desired. We collect the union of the set of bad players outputted by g, i.e. the union of sets $C_1, \ldots, C_h$, along with B_R and the bad players from all our inductive calls and let that be our final set of bad players.

We now briefly describe the parameters. Since we cut each of $C_1, \ldots, C_h$ in half until its size is less than c, the second time we use induction we call the inductive hypothesis $h \log(\ell/c) + O(1)$ times and this will dominate the cost of our inductive calls. Each call adds $\ell/\log^{(k-1)}(\ell)$ bad players. We also get $hc + |B_r| = hc + O(\ell/\log(h))$ bad players that will bias the last round. Setting $h = \left(\log^{(k-1)}(\ell)\right)^{1/C_0}$ and $c = \ell/\left(\log^{(k-1)}(\ell)\right)^{C_0}$ for some large constant C_0 yields the desired bound.

Formal description of algorithm. Formally, our algorithm proceeds as follows:

1. Let $B_R, B_H, B_I \leftarrow \emptyset$. We begin by biasing the last round.

   (a) For each $\alpha \in \left(\{0,1\}^\ell\right)^{k-1}$, let $\pi_\alpha : \{0,1\}^\ell \to \{0,1\}$ be the induced round k protocol.

   (b) Set $\mathcal{F} = \{\pi_\alpha\}_{\mathbb{E}[\pi_\alpha] \geq \gamma/2}$ and apply Theorem 4.1 to $\mathcal{F}$ with parameters $h, \gamma/2$ and $\delta = 1/3$. Let B_R be the set given by Theorem 4.1, and for $\alpha \in \left(\{0,1\}^\ell\right)^{k-1}$ let $g : \left(\{0,1\}^\ell\right)^{k-1} \to$

$([\ell])^h \cup \perp$ be such that (viewing $g(\alpha)$ as a set in the obvious way) $B_R \cup g(\alpha)$ can $(1-\gamma/2)$-bias π_α towards 1 if possible, otherwise $g(\alpha) = \perp$.

We assert that $\Pr[g \neq \perp] \geq \gamma/6$ (this will be proven in Claim 5.4).

(c) Let $\pi^{(1)}_{temp}$ be the $(k-1)$-round protocol such that $\pi^{(1)}_{temp}(\alpha) = 1$ iff $g(\alpha) \neq \perp$. We have that $\Pr[\pi^{(1)}_{temp} = 1] \geq \gamma/6$. By induction (the first use of induction mentioned in the informal overview) we can find a set of bad players, which we denote $B^{(1)}_{temp}$, so that $\Pr[\pi^{(1)}_{temp}|_{B^{(1)}_{temp}} = 1] \geq (1-\gamma/6) \geq 3/4$. Add all the bad players from $B^{(1)}_{temp}$ to B_I.

2. We now find bad players that can bias g. Initialize sets $C_1 = \cdots = C_h = [\ell]$. We will change these sets below while maintaining the following invariant:

   $(*)$ For every α, if there exists $j \in [h]$ such that $g_j(\alpha) \notin C_j$, then $g(\alpha) = \perp$. (Equivalently, if $g(\alpha) \neq \perp$, then $g_j(\alpha) \in C_j$ for all $j \in [h]$.)

   We iteratively bias players from rounds 1 to $k-1$ so that for all $j \in [h]$, C_j will have $|C_j| \leq c$. At the end of each iteration, we maintain that $\Pr[g \neq \perp] \geq 3/4$ (this is equivalent to having $\Pr_\alpha[g_1(\alpha) \in C_1, \ldots, g_h(\alpha) \in C_h] \geq 3/4$). This is done as follows:

   (a) While there exists $j \in [h]$ with $|C_j| > c$, do the following:
      i. Let $\mathbf{X}$ be a random subset of C_j of size $|C_j|/2$. Since $\Pr_\alpha[g_j(\alpha) \in C_j] \geq 3/4$, we have that $\mathbb{E}_{\mathbf{X}}[\Pr[g_j(\alpha) \in \mathbf{X}]] \geq 3/8$. So, there exists $C'_j \subset [\ell]$ with $\left|C'_j\right| = |C_j|/2$ such that $\Pr_\alpha[g_j(\alpha) \in C'_j] \geq 3/8$.
      ii. Update $C_j \leftarrow C'_j$ and let $g_j(\alpha) = \perp$ if $g_j(\alpha) \notin C_j$. We now have $\Pr_\alpha[g(\alpha) \neq \perp] \geq 3/8 \geq 1/4$.
      iii. Let $\pi^{(2)}_{temp}$ be the $(k-1)$-round protocol such that $\pi^{(2)}_{temp}(\alpha) = 1$ iff $g(\alpha) \neq \perp$. We have that $\Pr[\pi^{(2)}_{temp} = 1] \geq 1/4$. By induction (the second use of induction mentioned in the informal overview), we can find a set of bad players, which we denote $B^{(2)}_{temp}$, so that $\Pr[\pi^{(2)}_{temp}|_{B^{(2)}_{temp}} = 1] \geq 3/4$. Add all the bad players from $B^{(2)}_{temp}$ to B_I.

3. Let $\pi^{(3)}_{temp}$ be the $(k-1)$-round protocol such that $\pi^{(3)}_{temp}(\alpha) = 1$ iff $g(\alpha) \neq \perp$. We have that $\Pr[\pi^{(3)}_{temp} = 1] \geq 3/4 \geq \gamma/2$. By induction (the third use of induction mentioned in the informal overview), we can find a set of bad players, which we denote $B^{(3)}_{temp}$, so that $\Pr[\pi^{(3)}_{temp}|_{B^{(3)}_{temp}} = 1] \geq 1-\gamma/2$. Add all the bad players from $B^{(3)}_{temp}$ to B_I.

4. Finally, let $B_H = \bigcup_{j \in [h]} C_j$.

Correctness of the algorithm. We now prove the correctness of our procedure above. We first prove the following claim that we informally asserted at end of step 1(b) above:

Claim 5.4. *At the end of step* 1(*b*), $\Pr[g \neq \perp] \geq \gamma/6$.

Proof. Using the reverse Markov argument from Claim 3.3, since $\Pr[\pi = 1] \geq \gamma$, we have that $\frac{|\mathcal{F}|}{2^{(k-1)\ell}} = \frac{\left|\alpha \in (\{0,1\}^\ell)^{(k-1)} : \mathbb{E}[\pi_\alpha] \geq \gamma/2\right|}{2^{(k-1)\ell}} \geq \gamma/4$. Since we set $\delta = 2/3$ while using Theorem 4.1, we indeed infer that $\Pr[g \neq \perp] \geq \gamma/6$. $\square$

Lastly, we show that the protocol is indeed $(1-\gamma)$ biased by our set of bad players. At the end of the last step, we have that $\Pr[g \neq \perp] \geq (1 - \gamma/2)$. Moreover, by the choice of B_R and g in step 1, we know that if for $\alpha \in \left(\{0,1\}^\ell\right)^{k-1}$ we have that $g(\alpha) \neq \perp$, then $B_R \cup g(\alpha)$ can bias π_α so that $\Pr[\pi_\alpha|_{B_R \cup B_H} = 1] \geq 1 - \gamma/2$. Combining these, we conclude

$$\Pr[\pi|_{B_R \cup B_H \cup B_I} = 1] \geq \Pr[g \neq \perp] \cdot (1 - \gamma/2) \geq (1 - \gamma/2)^2 \geq 1 - \gamma$$

as desired.

Setting parameters. We finally set parameters. We set $h = \left(\log^{(k-1)}(\ell)\right)^{1/10^4}, c = \frac{\ell}{\left(\log^{(k-1)}(\ell)\right)^{10^4}}$. Recalling Theorem 4.1, we see that

$$|B_R| \leq \frac{100\ell \log(3/2)}{\gamma \log(h)} \leq \frac{10^6 \ell}{\gamma \log^{(k)}(\ell)}. \tag{1}$$

We also have that

$$|B_H| \leq c \cdot h \leq \frac{\ell}{\left(\log^{(k-1)}(\ell)\right)^{10^4}} \cdot \left(\log^{(k-1)}(\ell)\right)^{1/10^4} \leq \frac{\ell}{\left(\log^{(k-1)}(\ell)\right)^{10^4 - 1}} \tag{2}$$

We see that the while loop of step 2(b) iterates for at most $h \cdot \log(\ell/c) \leq 10^4 \log^{(k)}(\ell) \cdot \left(\log^{(k-1)}(\ell)\right)^{1/10^4}$ steps. In each iteration, we inductively have $\left|B^{(2)}_{temp}\right| \leq \frac{10^7 \ell}{(1/4)\log^{(k-1)}(\ell)}$. We also call our inductive bound twice, once before the loop and once after, and for each of these calls the corresponding sets inductively satisfy $\left|B^{(1)}_{temp}\right|, \left|B^{(3)}_{temp}\right| \leq \frac{10^7 \ell}{(\gamma/6)\log^{(k-1)}(\ell)}$. This gives

$$|B_I| \leq h \cdot \log(\ell/c) \cdot \frac{10^7 \ell}{(1/4)\log^{(k-1)}(\ell)} + 2 \cdot \frac{10^7 \ell}{(\gamma/6)\log^{(k-1)}(\ell)} \leq \frac{10^{12} \log^{(k)}(\ell)}{\left(\log^{(k-1)}(\ell)\right)^{1 - 1/10^4}} + \frac{10^9 \ell}{\gamma \log^{(k-1)}(\ell)}. \tag{3}$$

Hence, the total number of bad players we used is

$$\begin{aligned}
|B_R| + |B_H| + |B_I| &\leq \frac{10^6 \ell}{\gamma \log^{(k)}(\ell)} + \frac{\ell}{\left(\log^{(k-1)}(\ell)\right)^{10^4 - 1}} + \frac{10^{12} \log^{(k)}(\ell)}{\left(\log^{(k-1)}(\ell)\right)^{1 - 1/10^4}} + \frac{10^9 \ell}{\gamma \log^{(k-1)}(\ell)} \\
&\leq \frac{10^7 \ell}{\gamma \log^{(k)}(\ell)}
\end{aligned}$$

wherein for the last inequality, we set ℓ_0 large enough and use the fact that $\ell \geq \ell_0$. $\square$

5.1 Protocols with Longer Messages

In this subsection, we show that coin flipping protocols in which each player is allowed to send more than one bit per round can also be biased by relatively few bad players. In particular, we show:

Theorem 5.5. *There exist universal constants $C = 10^7, \ell_0 \in \mathbb{N}$ such that for all $k \in \mathbb{N}, \ell \in \mathbb{N}, 0 < \gamma < 1/4, 0 < \delta < 1/10^6, \varepsilon, r_1, \ldots, r_k \in \mathbb{N}$ where $\varepsilon \geq \frac{C\log(1/\gamma\delta)\log^{(k+2)}(\ell)}{\log^{(k+1)}(\ell)}$ and for $i \in [k]$, $1 \leq r_i \leq (\log^{(i)}(\ell))^{1-\varepsilon}$, the following holds: Let π be a k-round coin flipping protocol over $\ell \geq \ell_0$ players where $\Pr[\pi = 1] \geq \gamma$ and in round i, each player is allowed to send r_i bits. Then, there exists a set of bad players $B \subset [\ell]$ with $|B| \leq \delta\ell$ such that $\Pr[\pi|_B = 1] \geq 1 - \gamma$.*

Setting parameters, we obtain the following corollaries regarding corrupting protocols using 0.001ℓ many bad players:

Corollary 5.6. *Let $\Delta : \mathbb{N} \to \mathbb{N}$ be such that $\Delta(\ell) \geq \omega(1)$. Then, there exists $\varepsilon : \mathbb{N} \to [0,1]$ where $\varepsilon(\ell) \leq o(1)$ and $\ell_0 \in \mathbb{N}$ such that for all $k, \ell \in \mathbb{N}$ with $\ell \geq \ell_0$, and $k \leq \log^*(\ell) - \Delta(\ell)$, the following holds: Let π be a k round protocol over ℓ players where in round $i \in [k]$, the number of bits each player can send is at most $r_i = (\log^{(i)}(\ell))^{1-\varepsilon(\ell)} = (\log^{(i)}(\ell))^{1-o(1)}$. Then, there exists a set of bad players $B \subset [\ell]$ and an outcome $o \in \{0,1\}$ with $|B| \leq 0.001\ell$ such that $\Pr[\pi|_B = o] \geq 0.999$.*

In this first corollary, we let the number of round be $\log^*(\ell) - \Delta(\ell)$ where $\Delta(\ell) \geq \omega(1)$ and from that obtained a constraint on number of bits per round. We also note down the following corollary which follows from setting Δ to be a very large fixed constant:

Corollary 5.7. *There exists universal constants $\ell_0, \Delta \in \mathbb{N}$ such that for all $k, \ell \in \mathbb{N}$ with $\ell \geq \ell_0$, and $k \leq \log^*(\ell) - \Delta$, the following holds: Let π be a k round protocol over ℓ players where in round $i \in [k]$, the number of bits each player can send is at most $r_i = (\log^{(i)}(\ell))^{0.999}$. Then, there exists a set of bad players $B \subset [\ell]$ and an outcome $o \in \{0,1\}$ with $|B| \leq 0.001\ell$ such that $\Pr[\pi|_B = o] \geq 0.999$.*

To prove Theorem 5.5, we will slightly modify the proof of Theorem 5.1 and set parameters differently. Here's how we proceed:

Proof of Theorem 5.5. We proceed by induction on k. For $k = 1$, we view the protocol $\pi : (\{0,1\}^{r_1})^\ell$ as a function over $\ell \cdot r_1$ bits and apply Theorem 5.3 to find a set $B \subset [\ell r]$ with $|B| \leq \frac{C\ell r_1}{\gamma \log(\ell r_1)}$ such that if all bits in B are controlled by bad players, then they can $(1-\gamma)$-bias π towards 1. We let $B' \subset [\ell]$ denote the set of players that control all of the bits in B. We see that

$$|B'| \leq |B| \leq \frac{C\ell r_1}{\gamma \log(\ell r_1)} \leq \frac{C\ell(\log(\ell))^{1-\varepsilon}}{\gamma \log(\ell)} = \frac{C\ell}{\gamma(\log(\ell))^{\varepsilon}}$$

We need to show $|B'| \leq \delta\ell$, or equivalently that $\frac{C}{\gamma(\log \ell)^\varepsilon} \leq \delta$. We rewrite this inequality as $\varepsilon \geq \frac{\log(C/\gamma\delta)}{\log^{(2)}(\ell)}$ which holds by choice of ε.

For $k \geq 2$, we proceed by using the same exact algorithm as in Theorem 5.1 - for $\alpha \in \{0,1\}^{r_1} \times \cdots \times \{0,1\}^{r_{k-1}}$, we view $\pi_\alpha : (\{0,1\}^{r_k})^\ell$ as a function over ℓr_k bits. We set some parameters c, h and find sets B_R, B_H, B_I where B_R, B_H will find bad bits for the last round and B_I will apply induction to find bad players that can corrupt the first $i-1$ rounds. When applying induction for $(k-1)$-round protocols, we call it with parameter δ_I that we set later. So, our inductive calls to

$k-1$ round protocol will be such that it corrupts at most $\delta_I \ell$ bad players. We will then choose bad players for the last round such that they will control the bits that are in $B_R \cup B_H$. We see that in that case, the number of bad players we control over last round is at most $|B_R| + |B_H|$.

Before setting parameters, we first recall from Equation (1) and Equation (2) that $|B_R| \leq \frac{100\ell r_k \log(3/2)}{\gamma \log(h)}$ and $|B_H| \leq c \cdot h$. We also see, from Equation (3) and the discussion preceding it, that $|B_I| \leq (h \cdot \log(\ell/c) + 2) \cdot \delta_I \ell$. Also to enforce that we overall have $\delta\ell$ bad players, we will ensure $|B_R| \leq \frac{\delta\ell}{4}, |B_H| \leq \frac{\delta\ell}{4}$ and $|B_I| \leq \frac{\delta\ell}{4}$.

We finally set parameters. We let h be such that $\log h = \frac{400\log(3/2)(\log^{(k)}(\ell))^{1-\varepsilon}}{\gamma\delta}$ and let $c = \frac{\delta\ell}{4h}$. We set $\delta_I = \delta/h^2$. We now bound each of B_R, B_H, B_I and also, we show that we can indeed set δ_I to the prescribed value for induction, checking that it satisfies the restriction on ε. First, we bound $|B_R|$:

$$|B_R| \leq \frac{100\ell r_k \log(3/2)}{\gamma \log(h)} = \frac{\delta\ell r_k}{4(\log^{(k)}(\ell))^{1-\varepsilon}} \leq \frac{\delta\ell}{4}.$$

Second, we bound $|B_H|$:

$$|B_H| \leq c \cdot h \leq \frac{\delta\ell}{4}.$$

Finally, we bound $|B_I|$. To do that, we first observe that by choice of h, we have that $h > \frac{\log(h)}{10^6}$ and $h > \frac{\log(1/\delta)}{10^6}$. So, $2\log(4h/\delta) < h/4$. Using this, we see:

$$\begin{aligned}
|B_I| &= (h\log(\ell/c) + 2)\delta_I \ell \\
&\leq 2(h\log(\ell/c))\frac{\delta}{h^2}\ell \\
&= \frac{2\log(4h/\delta)}{h} \cdot \delta\ell \\
&\leq \frac{\delta\ell}{4},
\end{aligned}$$

where in the last inequality we used the bound $2\log(4h/\delta) < h/4$. Lastly, we show that we can indeed set δ_I to the desired value, i.e., that $\varepsilon \geq \frac{C\log(1/\gamma\delta_I)\log^{(k+1)}(\ell)}{\log^{(k)}(\ell)}$. To show this inequality holds, we will make use of the fact that $\varepsilon \geq \frac{C\log(1/\gamma\delta)\log^{(k+2)}(\ell)}{\log^{(k+1)}(\ell)}$. We see that

$$\begin{aligned}
\frac{C\log(1/\gamma\delta_I)\log^{(k+1)}(\ell)}{\log^{(k)}(\ell)} &= \frac{C\log(h^2/\gamma\delta)\log^{(k+1)}(\ell)}{\log^{(k)}(\ell)} \\
&\leq \frac{C\log(h^3)\log^{(k+1)}(\ell)}{\log^{(k)}(\ell)} && \text{(since } h \geq (1/\gamma\delta)\text{)} \\
&= \frac{3C\log(h)\log^{(k+1)}(\ell)}{\log^{(k)}(\ell)} \\
&= \frac{1200\log(3/2)C(\log^{(k)}(\ell))^{1-\varepsilon}\log^{(k+1)}(\ell)}{\log^{(k)}(\ell)} \\
&\leq \frac{1200C\log^{(k+1)}(\ell)}{\left(\log^{(k)}(\ell)\right)^{\varepsilon}}
\end{aligned}$$

Hence, it suffices to show that $1200C \log^{(k+1)}(\ell) \leq \varepsilon(\log^{(k)}(\ell))^{\varepsilon}$. Equivalently, we want to show that $\log(1200C) + \log^{(k+2)}(\ell) \leq \log(\varepsilon) + \varepsilon(\log^{(k+1)}(\ell))$. This last inequality indeed holds since by assumption we have that $\varepsilon \geq \frac{C \log^{(k+2)}(\ell)}{\log^{(k+1)}(\ell)}$. $\square$

6 Constructing Improved Constant-Round Coin Flipping Protocols

Our main results in this section are improved explicit constant-round coin flipping protocols, where each player is allowed to send one bit per round. Formally, we prove:

Theorem 6.1. *For any $k, \ell \in \mathbb{N}, 0 < \gamma < 1/2$ with $k \geq 2$, there exists a k-round coin flipping protocol over ℓ players with each player sending one bit per round such that when the number of bad players is at most $\frac{\gamma\ell}{\log(\ell)(\log^{(k)}(\ell))^2}$, the output coin flip is ε-close to uniform where $\varepsilon = O\left(\gamma + \left(\log^{(k-1)}(\ell)\right)^{-0.2}\right)$.*

To help construct our protocol, we need to introduce the following notion of an *assembly*:

Definition 6.2 (Assembly)**.** *A* (b, n, s)-assembly over ℓ players, *where each player is labeled either 'good' or 'bad,' is a collection of n disjoint subsets $S_1, \ldots, S_n \subset [\ell]$, each of size s, with the following property:*

- *There exists some fixed $B \subset [n]$ with $|B| \leq b$ such that if $i \notin B$, then the set S_i and all players within it are labeled as good. We say that such sets are 'good' and that the remaining sets are 'bad.'*

We will use the following ways of transforming a given assembly into an assembly with different parameters (the proof will be given later):

Lemma 6.3 (Transforming assemblies)**.** *Let b, n, s be arbitrary and let A be a (b, n, s)-assembly. Then we can explicitly transform A, oblivious to which players are labelled good or bad, in any of the following ways (without changing the underlying labels of players being good or bad):*

1. *(Grouping) For any $t \geq 1$, in 0 rounds, A can be transformed into a $(b, n/t, st)$-assembly.*
2. *(Splitting) For any $t \geq 1$, in 0 rounds, A can be transformed into a $(bt, nt, s/t)$-assembly.*
3. *(Feige's Lightest Bin Protocol [Fei99]) For any $\beta \leq 2^s, 0 < \delta < 1$, in 1 round, A can be transformed into a $\left(\frac{b}{\beta} + \frac{\delta(n-b)}{\beta}, \frac{n}{\beta}, s\right)$-assembly, with success probability at least $1 - \beta \cdot \exp\left(\frac{-\delta^2(n-b)}{2\beta}\right)$*

We will also use the following explicit resilient function. This can be viewed as a one-round operation which, given a suitable assembly, generates an almost-fair coin toss. Formally:

Lemma 6.4 (Resilient function from [AL93, IV24])**.** *In one round, a (b, n, s)-assembly can be used to output a single bit $\in \{0, 1\}$ that is $O\left(\frac{b(\log(n))^2}{n} + n^{-0.99}\right)$ close to uniform.*

Let us show that using Lemmas 6.3 and 6.4 we can obtain our desired k-round protocol.

Proof of Theorem 6.1 using Lemmas 6.3 and 6.4. We first informally describe our protocol and then we will formally describe it and analyze it.

Informal description First, we arbitrarily group players together and treat each group as a single entity (each group will correspond to a set of the assembly) which has access to multiple random bits. We will have many such entities, and most of them will contain solely of good players; call those 'good' entities and the remaining entities 'bad' entities. Then, we make these entities participate in one round of the lightest bin protocol with many bins. This vastly reduces the number of entities while roughly maintaining the absolute fraction of entities that are bad. We then again split each of the entities into many entities with fewer players in each one of them. This still maintains the absolute fraction of entities that are bad. We then repeat the lightest bin protocol and continue doing this until it is time for the last round. In the last round, we take one player from each entity and apply a resilient function to obtain our output coin flip.

Formal protocol We formally proceed as follows. Since we know that at most $\frac{\gamma\ell}{\log(\ell)\log^{(k)}(\ell)}$ players are bad, our input set of players form a $\left(\frac{\gamma\ell}{\log(\ell)\log^{(k)}(\ell)}, \ell, 1\right)$-assembly. Our protocol will repeatedly perform transformations on this assembly as described below. (We remark that the warmup case discussed in Section 2.5 (the case of $k = 2$) corresponds to performing only steps 1 and 3 below.)

1. In the first round we proceed as follows:
   (a) We use transformation 1 from Lemma 6.3 (grouping), setting $t = 3\log(\ell)$, to transform our $\left(\frac{\gamma\ell}{\log(\ell)\log^{(k)}(\ell)}, \ell, 1\right)$-assembly into a $\left(\frac{\gamma\ell}{\log(\ell)\log^{(k)}(\ell)}, \frac{\ell}{3\log(\ell)}, 3\log(\ell)\right)$-assembly.
   (b) We use transformation 3 from Lemma 6.3 (Feige's lightest bin protocol) with $\beta = \frac{\ell}{(\log(\ell))^3}, \delta = (\log(\ell))^{-0.25}$ to transform the $\left(\frac{\gamma\ell}{\log(\ell)\log^{(k)}(\ell)}, \frac{\ell}{3\log(\ell)}, 3\log(\ell)\right)$-assembly into a $\left(\frac{\gamma(\log(\ell))^2}{(\log^{(k)}(\ell))^2} + (\log(\ell))^{1.76}, \frac{(\log(\ell))^2}{3}, 3\log(\ell)\right)$-assembly, succeeding with probability $\geq 1 - \ell\exp\left(-(\log(\ell))^{1.49}\right)$.

2. For rounds 2 to $k-1$, we proceed as follows: In round i, we are given a $\left(\frac{\gamma(\log^{(i-1)}(\ell))^2}{(\log^{(k)}(\ell))^2} + (\log^{(i-1)}(\ell))^{1.76}, \frac{(\log^{(i-1)}(\ell))^2}{3}, 3\log^{(i-1)}(\ell)\right)$-assembly.
   (a) We use transformation 2 from Lemma 6.3 (splitting), setting $t = \frac{\log^{(i-1)}(\ell)}{\log^{(i)}(\ell)}$, to transform our assembly above into a $\left(\frac{\gamma(\log^{(i-1)}(\ell))^3}{\log^{(i)}(\ell)(\log^{(k)}(\ell))^2} + (\log^{(i-1)}(\ell))^{2.76}, \frac{(\log^{(i-1)}(\ell))^3}{3\log^{(i)}(\ell)}, 3\log^{(i)}(\ell)\right)$-assembly.
   (b) We use transformation 3 from Lemma 6.3 (Feige's lightest bin protocol) with $\beta = \frac{(\log^{(i-1)}(\ell))^3}{(\log^{(i)}(\ell))^3}, \delta = (\log^{(i)}(\ell))^{-0.25}$ to transform the $\left(\frac{\gamma(\log^{(i-1)}(\ell))^3}{(\log^{(i)}(\ell))(\log^{(k)}(\ell))^2} + (\log^{(i-1)}(\ell))^{2.76}, \frac{(\log^{(i-1)}(\ell))^3}{3\log^{(i)}(\ell)}, 3\log^{(i)}(\ell)\right)$-assembly into a $\left(\frac{\gamma(\log^{(i)}(\ell))^2}{(\log^{(k)}(\ell))^2} + (\log^{(i)}(\ell))^{1.76}, \frac{(\log^{(i)}(\ell))^2}{3}, 3\log^{(i)}(\ell)\right)$-assembly, succeeding with probability $\geq 1 - (\log^{(i-1)}(\ell))^3\exp\left(-(\log^{(i)}(\ell))^{1.49}\right)$.

3. In round k, we are given a $\left(\frac{\gamma(\log^{(k-1)}(\ell))^2}{(\log^{(k)}(\ell))^2} + (\log^{(k-1)}(\ell))^{1.76}, \frac{(\log^{(k-1)}(\ell))^2}{3}, 3\log^{(k-1)}(\ell)\right)$-assembly. We use the resilient function from Lemma 6.4 to obtain a single output bit $\in \{0,1\}$ that is $O\left(\gamma + (\log^{(k-1)}(\ell))^{-0.23}\right)$ close to uniform.

By a union bound over all k rounds (the failure probability is dominated by the $i = k-1$ iteration of step 2(b) above), we get that all our transformations succeed with probability $1 - \exp(-(\log^{(k-1)}(\ell))^{1.48})$. In such an event, the error from our final round - round k from applying the resilient function is $O\left(\gamma + (\log^{(k-1)}(\ell))^{-0.23}\right)$. Since $(\log^{(k-1)}(\ell))^{-0.23} > \exp(-(\log^{(k-1)}(\ell))^{1.48})$, we get that the final error of our coin flip is at most $O\left(\gamma + (\log^{(k-1)}(\ell))^{-0.2}\right)$ as desired. $\square$

It remains to prove that we can indeed obtain each of the transformations described earlier.

Proof of Lemma 6.3. We present proofs of each of the claimed transformations.

1. To do this, we arbitrarily partition $[n]$ into parts of size t and take the union of all sets that are in the same part. The number of sets indeed decreases to n/t and the size of each set increases to st. Also, since b sets were bad, there are at most b parts which have at least one bad set. For the rest of the parts, all the players in each of the sets comprising the part are good, and hence after taking the union, the resulting set will also only consist of good players.

2. To do this, we split each set S_i arbitrarily into t parts. We indeed obtain tn sets of size s/t each. Moreover, whenever we split a set consisting of only good players, all of the resulting sets only consist of good players. Hence, the number of such sets in the resultant assembly is $(n-b)t = nt - bt$ as desired.

3. To do this, we use Feige's lightest bin protocol from [Fei99]. In particular, we set up β bins $B_1, \ldots, B_\beta$ which receive 'votes' as follows. In one round, the players in each set $S_1, \ldots, S_n$ each flip a random coin. For each set S_i, we interpret their s coin flips together as generating a random number between 1 and β (since $\beta \le 2^s$) and that number is the 'vote' for which bin the set should go to. After this voting process, we pick the lightest bin (the one that has received the fewest votes), say B_{i^*}, and let the sets in that bin (along with arbitrary other sets, chosen in some canonical way so that number of sets in the output assembly is n/β) form the output assembly .

   We analyze the process described above. We treat all players in any bad set as bad, i.e., they are allowed to cast their votes after seeing the votes of all the good sets. We first note that the number of sets in the lightest bin B_{i^*} is at most n/β (since there are n sets and β bins). Hence, the above process always outputs an assembly consisting of n/β many sets of size s each. We now lower bound the number of good players in each of the bins. Fix any bin $i \in [\beta]$ and let $\mathbf{X}_i$ be the random variable representing the number of good sets that voted i. We think of the vote of each of the good sets as independently choosing whether to vote for bin i or not, where the vote is cast for bin i with probability $1/\beta$. Hence $\mathbb{E}[\mathbf{X}_i] = (n-b)/\beta$, and by the Chernoff bound Claim 3.4, we have that

$$\Pr[g_i \le (1-\delta)(n-b)/\beta] \le \exp\left(\frac{-\delta^2(n-b)}{2\beta}\right).$$

   We take union bound over all β bins to infer that with probability at least $1-\beta\exp\left(\frac{-\delta^2(n-b)}{2\beta}\right)$, every bin will have at least $\frac{(1-\delta)(n-b)}{\beta}$ good sets. As this is true for all bins, this is also true for the lightest bin i^*. Hence the number of bad sets in bin i^* is at most

$$\frac{n}{\beta} - \frac{(1-\delta)(n-b)}{\beta} = \frac{b}{\beta} + \frac{\delta(n-b)}{\beta}$$

as desired.

$\square$

7 Conclusion and Open Problems

There still remains a gap between the best known upper bounds and lower bounds for coin flipping protocols. For instance, even when restricted to two-round protocols, our protocols from Theorem 6.1 can handle $\frac{\ell}{(\log \ell)\operatorname{poly}(\log^{(2)} \ell)}$ bad players while our lower bound from Theorem 5.1 requires $\frac{\ell}{\log^{(2)} \ell}$ bad players. Towards proving stronger lower bounds to bridge this gap, we present the following question regarding simultaneously biasing functions.

Question 7.1 (Simultaneous biasing)**.** *Prove or disprove the following: Let $f = (f_1, \dots, f_m) : \{0,1\}^\ell \to \{0,1\}^m$ where $m = \ell^{0.01}$ be an arbitrary map. Then, there exist $B \subset [\ell], P \subset [m], o \in \{0,1\}^{|P|}$ with $|B| \le O(\ell/\log \ell)$, $|P| = 0.01m$ such that for $g = f|_B$, and all $i \in [P]$, $\Pr[g_i = o_i] \ge 0.99$.*

Note that a much weaker version of Question 7.1 indeed holds by Theorem 4.1. The difference is that in Theorem 4.1, the (slightly different) sets of bad players can corrupt each f_i differently while here we want each of the functions indexed by P to be simultaneously corrupted not only by the same *set* of bad players B, but also by the same *behavior* of the bad players in that set.

We note that when $m = O(1)$, by repeatedly applying the result of [KKL88], Question 7.1 does hold. However, that strategy requires more than n players once $m \ge \log \ell$, which is trivial. Answering the above question in the positive will lead to almost matching lower bounds for coin flipping - in particular, this would imply that $\frac{\ell \operatorname{poly}(\log^{(2)} \ell)}{\log \ell}$ bad players suffice to bias any two-round protocol. This bound would follow by using the same strategy as the proof of Theorem 5.1, with the 'second induction step' replaced by the simultaneous biasing conjecture. For multi-round lower bounds nearly matching our constructions from Theorem 6.1, we would need a stronger version of Question 7.1 where we would want to simultaneously bias not just a function map but a k-round protocol map.

Acknowledgements

We thank the organizers of the Dagstuhl Seminar on Algebraic and Analytic Methods in Computational Complexity and Schloss Dagstuhl for providing a stimulating research environment, where discussions between R.S. and E.C. initiated this collaboration.

References

[AL93] Miklós Ajtai and Nathan Linial. "The influence of large coalitions". In: *Comb.* 13.2 (1993), pp. 129–145. DOI: 10.1007/BF01303199 (cit. on pp. 1–3, 21).

[AN93] Noga Alon and Moni Naor. "Coin-Flipping Games Immune Against Linear-Sized Coalitions". In: *SIAM J. Comput.* 22.2 (1993), pp. 403–417. DOI: 10.1137/0222030 (cit. on p. 2).

[Ant06] Spyridon Antonakopoulos. "Fast leader-election protocols with bounded cheaters' edge". In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*. Ed. by Jon M. Kleinberg. ACM, 2006, pp. 187–196. DOI: 10.1145/1132516.1132544 (cit. on p. 2).

[BL85] Michael Ben-Or and Nathan Linial. "Collective Coin Flipping, Robust Voting Schemes and Minima of Banzhaf Values". In: *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*. IEEE Computer Society, 1985, pp. 408–416. DOI: 10.1109/SFCS.1985.15 (cit. on pp. 1, 2).

[CZ19] Eshan Chattopadhyay and David Zuckerman. "Explicit two-source extractors and resilient functions". In: *Annals of Mathematics* 189.3 (2019), pp. 653–705. DOI: 10.4007/annals.2019.189.3.1 (cit. on p. 1).

[CL95] Jason Cooper and Nathan Linial. "Fast Perfect-Information Leader-Election Protocols with Linear Immunity". In: *Comb.* 15.3 (1995), pp. 319–332. DOI: 10.1007/BF01299739 (cit. on p. 2).

[Dod06] Yevgeniy Dodis. *Fault-tolerant leader election and collective coin-flipping in the full information model*. 2006 (cit. on p. 2).

[Fei99] Uriel Feige. "Noncryptographic Selection Protocols". In: *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*. IEEE Computer Society, 1999, pp. 142–153. DOI: 10.1109/SFFCS.1999.814586 (cit. on pp. 2, 3, 9, 21, 23).

[FHHHZ19] Yuval Filmus, Lianna Hambardzumyan, Hamed Hatami, Pooya Hatami, and David Zuckerman. "Biasing Boolean Functions and Collective Coin-Flipping Protocols over Arbitrary Product Distributions". In: *46th International Colloquium on Automata, Languages, and Programming (ICALP)*. Vol. 132. LIPIcs. 2019, 58:1–58:13 (cit. on p. 2).

[IMV23] Peter Ivanov, Raghu Meka, and Emanuele Viola. "Efficient resilient functions". In: *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*. Ed. by Nikhil Bansal and Viswanath Nagarajan. SIAM, 2023, pp. 2867–2874. DOI: 10.1137/1.9781611977554.CH108 (cit. on p. 1).

[IV24] Peter Ivanov and Emanuele Viola. "Resilient functions: Optimized, simplified, and generalized". In: *CoRR* abs/2406.19467 (2024). DOI: 10.48550/ARXIV.2406.19467. arXiv: 2406.19467 (cit. on pp. 1, 9, 21).

[KKL88] Jeff Kahn, Gil Kalai, and Nathan Linial. "The Influence of Variables on Boolean Functions (Extended Abstract)". In: *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*. IEEE Computer Society, 1988, pp. 68–80. DOI: 10.1109/SFCS.1988.21923 (cit. on pp. 2, 4, 6, 8, 13, 15, 24).

[Mek17] Raghu Meka. “Explicit Resilient Functions Matching Ajtai-Linial”. In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*. Ed. by Philip N. Klein. SIAM, 2017, pp. 1132–1148. DOI: 10.1137/1.9781611974782.73 (cit. on p. 1).

[ORV94] Rafail Ostrovsky, Sridhar Rajagopalan, and Umesh V. Vazirani. “Simple and efficient leader election in the full information model”. In: *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*. Ed. by Frank Thomson Leighton and Michael T. Goodrich. ACM, 1994, pp. 234–242. DOI: 10.1145/195058.195141 (cit. on p. 2).

[RSZ02] Alexander Russell, Michael E. Saks, and David Zuckerman. “Lower Bounds for Leader Election and Collective Coin-Flipping in the Perfect Information Model”. In: *SIAM J. Comput.* 31.6 (2002), pp. 1645–1662. DOI: 10.1137/S0097539700376007 (cit. on pp. 2, 3, 6–8, 13).

[RZ01] Alexander Russell and David Zuckerman. “Perfect Information Leader Election in log* n+O (1) Rounds”. In: *J. Comput. Syst. Sci.* 63.4 (2001), pp. 612–626. DOI: 10.1006/JCSS.2001.1776 (cit. on pp. 2, 3).

[Sak89] Michael E. Saks. “A Robust Noncryptographic Protocol for Collective Coin Flipping”. In: *SIAM J. Discret. Math.* 2.2 (1989), pp. 240–244. DOI: 10.1137/0402020 (cit. on p. 2).

[Zuc97] David Zuckerman. “Randomness-optimal oblivious sampling”. In: *Random Structures & Algorithms* 11.4 (1997), pp. 345–367 (cit. on p. 2).