Game-Theoretic Fairness Meets Multi-Party Protocols: The Case of Leader Election

Kai-Min Chung[1], T-H. Hubert Chan[2], Ting Wen[2], and Elaine Shi[3]*

[1] Academia Sinica kmchung@iis.sinica.edu.tw
[2] The University of Hong Kong hubert@cs.hku.hk,twen.hku@gmail.com
[3] Carnegie Mellon University runting@cs.cmu.edu

Abstract. Suppose that n players want to elect a random leader and they communicate by posting messages to a common broadcast channel. This problem is called leader election, and it is fundamental to the distributed systems and cryptography literature. Recently, it has attracted renewed interests due to its promised applications in decentralized environments. In a game theoretically fair leader election protocol, roughly speaking, we want that even a majority coalition cannot increase its own chance of getting elected, nor hurt the chance of any honest individual. The folklore tournament-tree protocol, which completes in logarithmically many rounds, can easily be shown to satisfy game theoretic security. To the best of our knowledge, no sub-logarithmic round protocol was known in the setting that we consider.

We show that by adopting an appropriate notion of approximate game-theoretic fairness, and under standard cryptographic assumption, we can achieve $(1-1/2^{\Theta(r)})$-fairness in r rounds for $\Theta(\log\log n) \leq r \leq \Theta(\log n)$, where n denotes the number of players. In particular, this means that we can approximately match the fairness of the tournament tree protocol using as few as $O(\log\log n)$ rounds. We also prove a lower bound showing that logarithmically many rounds are necessary if we restrict ourselves to "perfect" game-theoretic fairness and protocols that are "very similar in structure" to the tournament-tree protocol.

Although leader election is a well-studied problem in other contexts in distributed computing, our work is the first exploration of the round complexity of *game-theoretically fair* leader election in the presence of a possibly majority coalition. As a by-product of our exploration, we suggest a new, approximate game-theoretic fairness notion, called "approximate sequential fairness", which provides a more desirable solution concept than some previously studied approximate fairness notions.

1 Introduction

Suppose that Murphy and Moody simultaneously solve a long-standing open problem in cryptography and they each submit a paper with identical result

[*] Author ordering is randomized. See our online full version [15] for full details and proofs.

to CRYPTO'21. The amazing CRYPTO'21 program committee recommends a hard merge of the two papers. Murphy and Moody decide to flip a random coin over the Internet to decide who gets to present the result at the prestigious CRYPTO'21 conference, to be held on the beautiful virtual beaches of Santa Barbara. Murphy and Moody both want to make sure that the outcome of the coin toss is fair, even when the other player may be behaving selfishly. There is good news and bad news. The bad news is that a famous lower bound by Cleve [16] proved that a strong notion of fairness, henceforth called *unbiasability*, is impossible in any n-player coin toss protocol in the presence of corrupt majority. Specifically, for any r-round protocol, a coalition controlling half or more of the players can implement an efficient attack that biases the outcome by $\Omega(\frac{1}{r})$. This impossibility result also holds in the two-party setting where one of the parties can be corrupt. This strong unbiasability notion is also the *de facto* notion in the long line of work on multi-party computation [8, 13, 26]. The good news is that Cleve's lower bound is not a deal-breaker for Murphy and Moody. In fact, they can simply run Blum's celebrated coin toss protocol [10]: each player picks a random bit and posts a commitment of the bit to a public bulletin board (e.g., a broadcast channel, a blockchain); then both parties open their committed bits and the XOR of the two bits is used to decide the winner. If either player ever aborts from the protocol or opens the commitment wrongly, it automatically forfeits and the other is declared the winner. Blum's protocol is *not* unbiasable, i.e., a player can indeed misbehave and bias the coin — however, the bias will simply benefit the other player and hurt itself. Although not explicitly stated in Blum's original paper, in fact, his celebrated protocol achieves a *game-theoretic* notion of fairness which is strictly weaker than the de facto unbiasability notion. Specifically, no player can benefit itself or hurt the other by deviating from the protocol, and thus the honest protocol is a Nash equilibrium in which no player would be incentivized to deviate.

The above example shows that in the two-party setting, adopting a game theoretic notion of fairness allows us to circumvent the impossibility of fairness in the corrupt majority setting [16]. Therefore, a natural question is whether such game theoretic notions can also help us in the multi-party setting. Surprisingly, this very natural question has traditionally been overlooked in the long line of work on multi-party protocols. Only very recently, an elegant work by Chung et al. [14] initiated the study of game-theoretic fairness in a multi-party setting. Unfortunately, Chung et al. [14] proved broad impossibility results (in the corrupt majority setting) for a particular formulation of the multi-party coin toss problem for natural game-theoretic fairness notions. Specifically, suppose that n parties want to toss a *binary* coin, and each player has preference for either the bit 0 or 1. If the outcome agrees with a player's preference, it obtains a utility 1; otherwise, it obtains a utility of 0. Chung et al. [14] showed that roughly speaking, unless all players but one prefer the same coin, the following natural fairness notions can be ruled out in the corrupt majority setting: 1) *maximin fairness*, which requires that no coalition can harm any honest individual; and 2) *cooperative strategy proofness* (also called *CSP-fairness* for short), which requires that no coalition can benefit itself.

Philosophically, if a protocol satisfies maximin fairness and CSP fairness, then no individual should be incentivized to deviate from this equilibrium, no matter whether the coalition/individual is greedy and profit-seeking, malicious and aiming to harm others, or paranoid and aiming to defend itself in the worst-possible scenario. Such protocols are also said to be *incentive compatible*.

1.1 Leader Election: Another Formulation of Multi-Party Coin Toss

In this paper, we revisit the question of game-theoretically fair multi-party coin toss. Specifically, we consider an alternative formulation. Instead of tossing a binary coin, we consider the problem of *leader election* which can be viewed as tossing an n-way coin among n parties. Suppose that all parties prefer to be elected: the elected leader gains a utility of 1 (or equivalently, a utility of an arbitrary positive value), whereas everyone else gains a utility of 0. This natural utility notion is often encountered in practical applications as we mention in Section 1.3. Intriguingly, for this formulation, the theoretical landscape appears starkly different from the binary-coin case[4]. The broad impossiblity results of Chung et al. [14] for the binary case no longer apply. A folklore approach henceforth called the tournament-tree protocol [6, 31] establishes the feasilibity of a logarithmic round, game-theoretically fair leader election protocol, even in the presence of majority coalitions:

- Each pair of players duels with each other to select a winner using Blum's coin toss [10]; again, aborting is treated as forfeiting.
- Now the $\frac{n}{2}$ winners of the previous iteration form pairs and run the same protocol to elect $\frac{n}{4}$ winners.
- After logarithmically many rounds, the final winner is called the leader.

Like Blum's protocol, the tournament-tree protocol also does not satisfy unbiasability, since anyone can abort and bias the outcome in a direction that harms itself. However, one can show that it indeed satisfies the aforemnetioned maximin fairness and CSP fairness notions, i.e., no coalition can harm an honest individual or benefit itself. In light of this folklore protocol, one important and natural open question is to understand the *round complexity* of game-theoretically fair, multi-party leader election in the corrupt majority setting. Specifically, *can we have an n-party, game-theoretically fair leader election protocol that tolerates majority coalitions, and completes in $o(\log n)$ number of rounds?* A naïve idea is to directly collapse the tournament-tree protocol to two rounds — in the first round, all players commit all random coins they ever need to use in the protocol; and in the second round, they open all random coins. It turns out that this naïve approach completely fails in the sense that a majority coalition can have a definitive winning strategy (see the online full version [15]).

[4] Game theoretically fair leader election and binary coin toss are different in nature partly due to the different utility functions.

Throughout this paper, we shall consider the *plain setting without trusted setup*, and allowing *standard cryptographic assumptions*. This rules out naïve solutions such as having the trusted setup choose the coin toss outcome, or using Verifiable Delay Functions [11, 12]. Also, recall that in the honest majority setting, the standard multi-party computation literature gives us constant-round solutions [7, 18] that achieves the stronger notion of unbiasability. Therefore, we will focus on the corrupt majority setting. We also stress that the game-theoretic fairness notions we consider are stronger than in some previous contexts. For example, a strictly weaker notion is called *resilience*, which requires that an honest player is elected with constant probability [19, 20, 35, 36]. The resilience notion may be sufficient in certain contexts, however, it does not provide *incentive compatibility* like our notions.

1.2 Our Results and Contributions

We initiate the study of the round complexity of game-theoretically fair, multi-party leader election. Below, we first describe our new upper bound result and techniques informally, and then we will discuss the interesting definitional subtleties we encountered and our definitional contributions — it turns out that even defining an *approximate* notion of (game-theoretic) fairness is rather non-trivial, and the notions that existed in the literature appear somewhat lacking.

New upper bounds and techniques. Roughly speaking, we prove that one can *approximately* match the fairness of the tournament-tree protocol, in as small as $O(\log\log n)$ rounds. Specifically, we give the following parametrized result that allows one to trade off the round complexity and approximation factor.

Theorem 1 (Informal: round-efficient, game theoretically fair leader election). *For $r \in [C_0 \log\log n, C_1 \log n]$ where C_0 and C_1 are suitable constants, r-round protocols exist that achieve $\left(1 - \frac{1}{2^{\Theta(r)}}\right)$-approximate fairness in the presence of a coalition of size at most $\left(1 - \frac{1}{2^{\Theta(r)}}\right) \cdot n$.*

In the above, roughly speaking, 1-fairness means perfect fairness and 0-fairness means no fairness. Observe that if we plug in $r = \Theta(\log\log n)$, we can achieve $(1 - o(1))$-fairness against coalitions of size $n - o(n)$. It is also interesting to contrast our result with the classical notion of approximate unbiasability — it is well-known that r-round protocols cannot achieve better than $O(1/r)$-unbiasability in the presence of a majority coalition [16]. In contrast, our approximation factor, i.e., $\frac{1}{2^{\Theta(r)}}$, is exponentially sharper than the case of approximate unbiasability. We review more related work on ϵ-unbiasability in the online full version [15].

The techniques for achieving our upper bound are intriguing and somewhat surprising at first sight. We describe a novel approach that combines combinatorial techniques such as extractors, as well as cryptographic multiparty computation (MPC). Intriguingly, for designing game theoretically secure protocols, some of our classical insights in the standard MPC literature do not apply.

Several aspects of our protocol design are counter-intuitive at first sight. For example, jumping ahead, we defend against "a *large* coalition benefitting itself" using (non-trivial) combinatorial techniques; but these combinatorial techniques provide no meaningful defense against a *small* coalition benefitting itself — it is initially surprising that small coalitions turn out to be more challenging to defend against. To defend against a small coalition, we employ a special *honest-majority* MPC protocol as part of our final construction. The fact that an honest-majority MPC can provide meaningful guarantees in a corrupt majority setting is initially surprising too. Of course, weaving together the combinatorial and the cryptographic techniques also has various subtleties as we elaborate on in subsequent sections. We believe our design paradigm can potentially lend to the design of other game-theoretically fair protocols.

New definition of approximate fairness. It turns out that how to define a good *approximate* fairness notion requires careful thought. The most natural (but somewhat flawed) way to define $(1-\epsilon)$-fairness is to require that even a majority coalition cannot increase its own chances by more than an ϵ factor, or reduce an honest individual's chance by more than ϵ. Throughout the paper, we allow the coalition's *action space* to include *arbitrary deviations from the prescribed protocol*, as long as the coalition is subject to probabilistic polynomial-time (p.p.t.) computations. We consider a multiplicative notion of error, i.e., we want that a coalition A's expected utility is at most $\frac{|A|}{(1-\epsilon)\cdot n}$ where $\frac{|A|}{n}$ is the coalition's fair share had it played honestly; moreover, we want that any honest individual's expected utility is at least $(1-\epsilon)/n$ where $1/n$ is its utility if everyone participated honestly. We prefer a multiplicative notion to an additive notion, because in practical settings, the game may be repeated many times and the absolute value of the utility may not be as informative or meaningful. The relative gain or loss often matters more.

Indeed, some earlier works considered such an approximate fairness notion — for example, Pass and Shi [33] considered such a notion in the context of consensus protocols; they want that a (minority) coalition cannot act selfishly to increase its own gains by more than ϵ[5]. We realize, however, that such an approximate notion is somewhat flawed and may fail to rule out some undesirable protocols. Specifically, consider a protocol in which some bad event happens with small but non-negligible probability, and if the bad event happens, it makes sense for the coalition to deviate. For example, consider a contrived example.

> *Example.* Suppose that Alice and Bob run Blum's coin toss except that with ϵ probability, Bob sends all his random coins for the commitment to Alice in the first round. If this small-probability bad event happens, Alice should choose a coin that lets her win. This is not a desirable protocol because with small but non-negligible probability, it strongly incentivizes Alice to deviate.

However, the above protocol is not ruled out by the aforementioned notion of approximate fairness: since the probability of the bad event is small, the a-

[5] Pass and Shi [33] do not consider the threat of a coalition targeting an individual.

priori motivation for Alice or Bob to deviate is indeed small. In the online full version [15], we give another (arguably less contrived) counter-example that also violates sequential fairness.

We propose a new approximate fairness notion called *sequential approximate fairness* that avoids this drawback, and characterizes a more desirable space of solution concepts. At a very high level, our new notion says, it is not enough for a coalition to not have *a-priori* noticeable incentives to deviate, rather, we want the following stronger guarantee: *except with negligible probability, at no point during the protocol execution should a coalition have noticeable (i.e., ϵ) incentive to deviate, even after having observed the history of the execution so far.*

Remark 1. In the online full version [15], we show that the non-sequential approximate fairness notion is in fact equivalent to a multiplicative approximate variant of the Rational Protocol Design (RPD) notion proposed by Garay et al. [22–24]. However, as mentioned, we believe that our new *sequential* approximate notion provides a better solution concept.

Lower bound. The tournament-tree protocol achieves perfect fairness (i.e., $\epsilon = 0$) in an ideal "commit-and-immediately-open" model. That is, the protocol proceeds in $\log n$ iterations where each iteration consists of a commitment and a subsequent opening for every player. In the online full version [15], we prove a lower bound showing that in the operational model of the tournament-tree protocol, i.e., if we insist on perfect fairness (assuming idealized commitments) as well as immediate opening of committed values, unfortunately $\Theta(\log n)$ rounds is optimal. This lower bound provides a useful sanity check and guideline for protocol design. In comparison, our protocol achieves sub-logarithmic round complexity by introducing the approximate fairness relaxation and general cryptographic techniques. It is an open direction to precisely characterize the minimal conditions/assumptions under which sub-logarithmic rounds become possible.

Theorem 2 (Informal: some relaxations in our design are necessary). *Assume the ideal commitment model. If commitments must be opened immediately in the next round and perfect fairness is required, then $\Omega(\log n)$ rounds is necessary.*

Our work complements the recent prior work of Chung et al. [14] and makes a new step forward at understanding the mathematical landscape of game-theoretically fair, multi-party coin toss. Unlike the *de facto* unbiasability notion, however, our understanding of game-theoretic fairness in multi-party protocols is only just beginning, and there are numerous open questions. We describe some open questions in the online full version [15].

1.3 Motivating Applications and Scope of Our Work

Our work should be viewed as an *initial theoretical exploration* of the round complexity of game-theoretically fair leader-election. We do not claim practicality;

however, it is indeed an exciting future direction to design practical variants of our ideas.

Having said this, interestingly, the original inspiration that led the formulation of this problem as well as our game theoretic notions comes from emerging decentralized applications [5, 6, 9, 31]. In a decentralized environment, often pseudonyms or public keys are cheap to create, and thus it may well be that many pseudonyms are controlled by the same entity, i.e., *the classical honest majority assumption is not reasonable*. Some works orthogonal and complementary to our paper [30] aim to make it more costly to establish identities in decentralized applications, nonetheless, even with such DoS-defense mechanisms, honest majority may not be a reasonable assumption.

A line of work [5, 9] considered how to achieve a "financially fair" n-party lottery over cryptocurrencies such as Bitcoin and Ethereum. These works adopt game-theoretic fairness notions similar in spirit to ours, but they rely on collateral and penalty mechanisms to achieve fairness. In comparison, in our model, we aim to achieve fairness *without having to rely on additional assumptions such as collateral and penalty*. A couple recent works [6, 31] also pointed out that collateral and penalty mechanism can be undesirable and should be minimized in mechanism design in decentralized blockchain environments.

Leader election is also needed in decentralized smart contracts where one may want to select a service provider among a pool to provide some service, e.g., act as the block proposer, generate a verifiable random beacon, or verifiably perform some computational task, in exchange for rewards. In this case, providers may wish to get elected to earn a profit. A coalition may also wish to monopolize the eco-system by harming and driving away smaller players (potentially even at the cost of near-term loss). Conversely, a small player may be concerned about protecting itself in worst-possible scenarios. Our game-theoretic notion guarantees that no matter which of objectives a player or coalition has, it has no noticeable incentive to deviate from the honest protocol. In such blockchain settings, typically the blockchain itself can serve as a broadcast channel, and a round can be a confirmation delay of the blockchain[6].

2 Technical Overview

In this section, we will go through a few stepping stones to derive an $O(\log \log n)$-round protocol achieving $(1 - o(1))$-approximate fairness. We defer the fully parametrized version to the subsequent formal sections.

[6] Why and how blockchain can formally realize/approximate a broadcast channel is outside the scope of our paper, and has been extensively studied in a line of works on distributed consensus. We simply assume broadcast as given, a modeling approach that has been adopted in the long line of work on multi-party computation. In fact, our protocol execution model is no different from the standard literature on multi-party computation — see Section 2.1.

2.1 Leader Election Protocol

A leader election protocol (also called lottery) involves n players which exchange messages over *pairwise private channels* as well as a *common broadcast channel*. The protocol execution proceeds in synchronous rounds: in every round, players first receive new messages, then they perform some local computation, and send new messages. We assume a *synchronous network* where messages posted by honest players can be received by honest recipients in the immediate next round. At the end of the final round, everyone can apply an a-priori fixed function f over all messages on the broadcast channel to determine a unique leader from $[n]$, i.e., the result is *publicly verifiable*. For *correctness*, we require that in an honest execution where all players faithfully follow the protocol, the elected leader be chosen uniformly at random from $[n]$.

A subset of the players (often called a coalition) may decide to deviate from the honest strategy. Such a coalition can perform a *rushing* attack: during a round, players in the coalition (also called corrupt players) can wait to read all messages sent by honest players in this round, then decide what messages they should send in the same round.

Throughout the paper, we assume that an execution of the protocol is parametrized with a security parameter κ, since the protocol may adopt cryptographic primitives. We assume that the number of players n is a polynomially bounded function in κ; without loss of generality we assume that $n \geq \kappa$.

2.2 Non-Sequential Approximate Fairness

For simplicity, we first present an overview of our upper bound using the *non-sequential* notion of approximate fairness. However, in subsequent formal sections, we will actually define a better solution concept called *sequential approximate fairness*, and prove our protocols secure under this better solution concept.

Chung et al. [14] considered game theoretic fairness in a setting where n parties wish to toss a binary coin. They considered *perfect* fairness notions and coined them cooperative-strategy-proofness and maximin fairness, respectively. Below we give the natural approximate versions of these notions:

- *CSP-fairness:* we say that a leader election protocol achieves $(1-\epsilon)$-cooperative-strategy-proofness against a (non-uniform p.p.t.) coalition $A \subset [n]$, iff no matter what (non-uniform p.p.t.) strategy A adopts, its expected utility is at most $\frac{|A|}{(1-\epsilon)n}$. We often write CSP-fairness in place of "cooperative strategy proofness" for short.
- *Maximin fairness:* we say that a leader election protocol achieves $(1-\epsilon)$-maximin-fairness against a (non-uniform p.p.t.) coalition $A \subset [n]$, iff no matter what (non-uniform p.p.t.) strategy A adopts, any honest individual's expected utility is at least $(1-\epsilon)/n$.

Approximate maximin-fairness and approximate CSP-fairness are not equivalent — we give more explanations in the online full version [15].

Remark 2 (Coalition-resistant notions of equilibrium). In our definitions, we consider the deviation of a single coalition. This definitional approach is standard in game theory [1–4, 19–21, 25, 29, 36, 38], since the philosophy is to capture the notion of an approximate equilibrium in the sense that no coalition has noticeable incentives to deviate. Our equilibrium notion is coalition-resistant. In comparison, the standard notion of (approximate) Nash equilibrium typically considers deviation of a single player, and therefore is weaker than our notions in this sense.

Remark 3 (Choice of ϵ). In our formal results later, we will use $\epsilon = o(1)$ — in fact, our result will be parametrized. For simplicity, in the informal roadmap, it helps to think of $\epsilon = 1\%$.

2.3 A Strawman Scheme

Although *in our final scheme we do NOT use random oracles* (RO), it is instructive to think about a strawman scheme with an RO. Interestingly, this approach is inspired by recent proof-of-stake consensus protocols [17, 28].

Strawman: RO-based committee election + tournament tree

1. Every player $i \in [n]$ broadcasts a bit $x_i \in \{0, 1\}$, and we use $\mathsf{RO}(x_1, \ldots, x_n)$ to elect committee of size $\log^9 n$. If a player i fails to post a bit, we treat $x_i := 0$.
2. The committee runs the tournament-tree protocol to elect a final leader.

One can easily show that this approach achieves $(1 - \epsilon)$-CSP-fairness against any coalition A containing *at least $\epsilon/2$ fraction of the players* — we call a coalition at least $\epsilon/2$ fraction in size a *large coalition*. The argument is as follows. Since the second step, i.e., tournament tree, is in some sense "ideal", to increase its expected utility, the coalition $A \subset [n]$ must include as many of its own members in the committee as possible. Suppose that $\epsilon = 1\%$. For a fixed RO query, the probability that it selects a *bad* committee, i.e., one with more than $\frac{|A|}{(1-\epsilon)\cdot n}$ fraction of coalition players, is negligibly small by the Chernoff bound. Since the coalition is computationally bounded and can make at most polynomially many queries to RO, by the union bound, except with negligible probability, all of its RO queries select a good committee.

Unfortunately, this scheme suffers from a couple serious flaws:

- *Drawback 1: NOT approximately maximin-fair*: a coalition A can harm an individual $i \notin A$ as follows: wait till everyone not in A broadcasts their bits, and then try different combinations of bits for those in A to find a combination that excludes the player i from the committee. This attack can succeed with $1 - o(1)$ probability if $|A| = \Theta(\log n)$.
- *Drawback 2: NOT approximately CSP-fair against a small coalition*: a profit-seeking individual i is incentivized to deviate in the following manner: i can

wait for everyone else to post bits before posting its own bit denoted x_i. In this way it can increase its advantage roughly by a factor of 2 since it can try two choices of x_i. This attack can be extended to work for small coalitions too.

The second drawback is somewhat surprising at first sight, since we proved the strawman scheme to be CSP-fair against large coalitions (i.e., at least $\epsilon/2$ fraction in size). The reason is because the Chernoff bound proof gives only statistical guarantees about a population, but does not give meaningful guarantees about an individual or a very small group of players.

Remark 4. In the above strawman, one can also replace the committee election with a single iteration of Feige's lightest bin protocol [20]. The resulting protocol would still be $(1-\epsilon)$-CSP-fair, although it suffers from exactly the same drawbacks as the RO-based strawman. The upgrade techniques described in Section 2.4, however, is compatible only with the RO-based approach — and this is why we start with the RO-based approach. However, intriguingly, we will indeed make use of the lightest bin protocol later in Section 2.5 where we show how to get rid of the RO.

2.4 Warmup: A Game Theoretically Fair, RO-Based Protocol

We now discuss how to fix the two drawbacks in the previous strawman scheme. We will still have an RO in the resulting warmup scheme; however, in the immediate next subsection, we will discuss techniques for removing the RO, and obtain our final construction.

The first drawback is due to a potentially large coalition A choosing its coins (after examining honest coins) to exclude some individual $i \notin A$ from the committee. The second drawback is due to a small coalition A containing less than ϵ fraction of the players choosing its coins to help its members get included. To tackle these drawbacks, our idea is to introduce virtual identities henceforth called v-ids for short. Basically, we will use the RO to select a committee consisting of v-ids. When the RO's inputs are being jointly selected, we make sure that 1) a potentially large coalition A has no idea what each honest individual's v-id is and thus A has no idea which v-id to target; and 2) a small coalition has no idea what its own v-ids are, and thus it has no idea which v-ids to help.

To achieve this, each player i's final v-id will be the xor of two shares: a share chosen by the player itself henceforth called the *unmasked* v-id, and a share jointly chosen by a special, honest-majority protocol, henceforth called the *mask*. In the beginning, the player itself commits to its own unmasked v-id, and the MPC protocol jointly commits to each player's mask. Next, the players jointly choose the inputs to the RO. Finally, each player reveals its own unmasked v-id, and then the MPC protocol reconstructs all players' masks.

Special honest-majority MPC. Instantiating these ideas correctly, however, turns out to be rather subtle. A generic honest-majority MPC protocol does not guarantee anything when there is a large coalition. In our case, when the coalition is large, it can fully control the mask value. However, we do need that even with $(1-\epsilon)n$-sized coalitions, the mask value must be uniquely determined at the end of the sharing phase, and reconstruction is guaranteed. More specifically, we want our special, honest-majority MPC to satisfy the following properties for some small $\eta \in (0,1)$ (think of $\eta = \epsilon/2$):

- If $|A| \leq \eta n$, we want that at the end of this sharing phase, A has no idea what its own masks are;
- As long as $|A| < (1-2\eta)n$, at the end of the sharing phase, the mask value to be reconstructed is uniquely determined, and moreover, reconstruction is guaranteed to be successful.

The following $\mathcal{F}^{\eta}_{\mathrm{mpc}}$ ideal functionality describes what we need from the honest-majority MPC. For simplicity, in our informal overview, we will describe our protocols assuming the existence of this $\mathcal{F}^{\eta}_{\mathrm{mpc}}$ ideal functionality. Later in Section 4.2, we will instantiate it with an actual, constant-round cryptographic protocol using bounded concurrent MPC techniques [32]. Technically, the real-world cryptographic instantiation does not securely emulate $\mathcal{F}_{\mathrm{mpc}}$ by a standard simulation-based notion; nonetheless, we prove in the online full version [15] that the fairness properties we care about in the ideal-world protocol (using idealized cryptography) extend to the real-world protocol (using actual cryptography).

$\mathcal{F}^{\eta}_{\mathrm{mpc}}$: special, honest-majority MPC functionality

Sharing phase. Upon receiving `share` from all honest players, choose a random string coins. If the coalition size $|A| \geq \eta n$, the adversary is asked to overwrite the the variable coins to any value of its choice. Send `ok` to all honest players.

Reconstruction phase. Upon receiving `recons` from all honest players: if $|A| \geq (1-2\eta)n$, the adversary may, at this point, overwrite the string coins to its choice. Afterwards, in any case, send coins to all honest players.

Our warmup RO-based protocol. Now, it helps to describe our protocol first, then we explain the additional subtleties. We describe our warmup protocol using an idealized commitment scheme, as well as the $\mathcal{F}_{\mathrm{mpc}}$ functionality described earlier.

Our warmup RO-based protocol

1. Every player $i \in [n]$ commits to a randomly selected unmasked v-id $y_i \in \{0,1\}^v$ where $2^v = n \cdot \mathsf{poly}\log n$.
2. Send `share` to $\mathcal{F}^{\epsilon/2}_{\mathrm{mpc}}$ and receive `ok` from $\mathcal{F}_{\mathrm{mpc}}$.
3. Every player $i \in [n]$ broadcasts a bit x_i. Let x be the concatenation of all of $\{x_i\}_{i\in[n]}$ in increasing order of the players' indices — here for any player j who has aborted, its x_j is treated as 0.
4. Every player $i \in [n]$ now opens its committed unmasked v-id $y_i \in \{0,1\}^v$.

5. All honest players send recons to $\mathcal{F}_{\text{mpc}}^{\epsilon/2}$, and they each receive a mask vector z from $\mathcal{F}_{\text{mpc}}^{\epsilon/2}$.
6. Parse $z := (z_1, \ldots, z_n)$ where each $z_j \in \{0,1\}^v$ for $j \in [n]$. We now view $y_i \oplus z_i$ player i's final v-id. A player i is a member of the committee $\mathcal{C}$ iff 1) it correctly committed and opened its unmasked v-id y_i; 2) its final v-id $y_i \oplus z_i$ is chosen by $\mathsf{RO}(\mathbf{x})$; and 3) its final v-id $y_i \oplus z_i$ does not collide with anyone else's final v-id— we may assume that anyone who aborted has the final v-id $\perp$.
7. The committe $\mathcal{C}$ runs the tournament-tree protocol to elect a leader.

Additional subtleties. At this moment, it helps to point out a few additional subtleties.

1. *Unique reconstruction even under a majority coalition.* First, recall that even in the presence of a $(1-\epsilon)$-coalition, we wanted our $\mathcal{F}_{\text{mpc}}$ to guarantee uniqueness of the reconstructed mask z at the end of the sharing phase. This is important because we do not want the coalition to see the RO's outputs and then choose the mask vector z a-posteriori to exclude some honest individual from the final committee or to include all of the coalition members.
2. *The need for collision detection.* Second, notice that the protocol prevents colliding final v-ids from being elected into the final committee. Such a collision detection mechanism is necessary since otherwise, the following attack would be possible[7]: a 99% coalition can make all of its members choose the same final v-id— it can do that because it controls its members' unmasked v-ids as well as the mask value. Now, the 99% coalition can choose its input bits to the RO to help this particular final v-id. In this way, with high probability, all coalition members can be elected into the final committee.
3. *Proving sequential approximate fairness.* Last but not the least, so far we have only focused on the non-sequential notion of fairness, and it turns out that proving the sequential notion is much more subtle. In our formal proofs later (see Sections 5 and the online full version [15]), we will do a round-by-round argument to show that except with negligible probability, in no round of the protocol would the coalition have noticeable incentive to deviate.

Since this warmup construction is not our final scheme, we will not formally prove the warmup construction. Instead, we now explain how to get rid of the RO to get our final scheme.

2.5 Final Construction: Removing the Random Oracle

To remove the RO, our idea is to replace the committee election with a two-phase approach, where the first phase uses a single iteration of Feige's lightest-bin protocol [20] and the second phase uses a combinatorial object called a

[7] We describe this attack for illustration purposes to help understanding. Of course, we will later prove our final construction secure against all possible p.p.t. coalition strategies.

sampler [37] in place of the RO. We briefly describe the intuition below. The actual scheme, calculations, and proofs are more involved especially for getting the more general, parametrized result, and we defer the full description to the subsequent formal sections.

Background. We will rely on a combinatorial object called a *sampler* which is known to be equivalent to a seeded extractor [37][8]. A sampler, denoted as Samp, is a combinatorial object with the following syntax and properties: given an input $x \in \{0,1\}^u$, $\mathsf{Samp}(x)$ returns d sample points $z_1, \ldots, z_d \in \{0,1\}^v$ from its output space. A sampler is supposed to have good, random-sampling-like properties. Consider a predicate function $f : \{0,1\}^v \rightarrow \{0,1\}$. The *population mean* of f over its inputs is defined as is $\frac{1}{2^v} \sum_{z \in \{0,1\}^v} f(z)$. The d sample points define a *sample mean* $\frac{1}{d} \sum_{j=1}^{d} f(z_j)$, which ideally should be close to the population mean. An $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-averaging sampler Samp guarantees that for any f, at least a $1 - \delta_{\mathrm{s}}$ fraction of the inputs will lead to a sample mean that differs from the population mean by at most ϵ_{s} additively.

Intuition. A flawed idea is to directly replace the RO in the warmup scheme with a sampler. To do so, the nature of our proof for this specific step will have to change: in the warmup scheme, we relied on the fact that the coalition can make only polynomially many queries to RO in our fairness proof. With a sampler, however, we must make a combinatorial argument here that does not depend on the adversary's computional bounds (although to reason about other parts of the scheme involving the commitment and the MPC, we still need to make computational assumptions on the adversarial coalition). Specifically, we want to argue that no matter which subset of players form a coalition, as long as the coalition's size is, say, between $0.01n$ and $0.99n$, then almost all honest inputs x_H *resist even the worst-case attack*, in the sense that there does not exist a x_A such that $x = (x_H, x_A)$ would form a bad input to Samp[9]. Here x is said to be a bad input to Samp if $\mathsf{Samp}(x)$ selects a committee in which the fraction of coalition players is noticeably higher than $|A|/n$.

Suppose that we want to select a $\log^9 n$-sized committee, and the final v-id space is of size $n \log^3 n$. In this case, we would need the sampler to select roughly $d = \log^{12} n$ output points. A calculation using the probabilistic method suggests that in this case, we cannot start with n players who jointly select the input to the sampler — if so, there would simply be too many combinations the adversarial coalition could try for its own input bits; and the number of bad inputs to the sampler simply is not sparse enough to defeat so many adversarial combinations.

The parameters would work out, however, if we start out with, say, $\log^3 n$ players who jointly choose the input to the sampler. In our subsequent formal sections, we will select parameters that work with the best known explicit sampler construction [27, 34, 37].

[8] We stress that our construction does not need a common reference string as the seed.

[9] Throughout the paper, for $S \subseteq [n]$, we use $x_S := \{x_i\}_{i \in S}$ to denote the coordinates of the vector x corresponding to all players in S.

Our idea. Given the above intuition, our idea is to adopt a *two-phase committee election* approach. We first down-select to a preliminary committee of size $\log^3 n$, and then the preliminary committee jointly choose input bits to a sampler to select a *final committee* among *all* players, and the final committee runs the tournament tree protocol to elect a leader among the final committee. We sketch the protocol below while deferring a more formal description to Section 4:

- *Commitment phase.* As before, players commit to their unmasked v-ids and use an honest-majority MPC to jointly commit to a mask first.
- *Preliminary committee election.* First, we elect a $\log^3 n$-sized *preliminary committee* such that the fraction of honest players on the preliminary committee approximately matches the fraction of honest players in the overall population. Here we do not care about the threat where a potentially large coalition seek to exclude a specific individual or a small coalition or individual try to include itself. It turns out that this can be accomplished by running a single iteration of Feige's elegant lightest bin protocol [20] in the plain model.
- *Final committee election.* Next, the preliminary committee jointly selects an input to the sampler, which is used to select $\log^9 n$ final v-ids among the space of all possible v-ids— these final v-ids would form the *final* committee. At this moment, the players open their unmasked v-ids, and reconstruct the mask that was secret shared earlier by the MPC. The players' final v-ids are now revealed, and the final committee determined.
- *Leader election.* Finally, the elected, poly-logarithmically sized final committee runs the tournament-tree protocol to elect a final leader.

3 Defining Sequential Approximate Fairness

3.1 Sequential Approximate Fairness

The non-sequential fairness notions mentioned in Section 2.2 does not rule out some undesirable protocols that may offer incentives for a coalition to deviate with non-negligible probability. Recall the example given in Section 1 where two parties run Blum's coin toss except that with some small ϵ probability, Bob broadcasts all its private coins in the first round. If the small (but non-negligible) probability bad event happens, Alice should deviate and choose her coins to definitively win. However, *a-priori* Alice does not have much incentive to deviate: since the bad event happens with only ϵ probability, her a-priori probability if winning is at most $\epsilon \cdot 1 + (1-\epsilon) \cdot \frac{1}{2} = (1+\epsilon) \cdot \frac{1}{2}$, and this is only an ϵ fraction more than her fair share. Nonetheless, we do want to rule out such bad protocols since such a protocol has a non-negligible probability ϵ of creating incentives for Alice to deviate.

We propose a better solution concept called sequential approximate fairness. Roughly speaking, we require that even if the coalition is allowed to re-evaluate whether to deviate at the beginning of every round in the protocol, except with negligible probability, no p.p.t. coalition (of size at most $(1-\epsilon)n$) should have ϵ incentive to deviate at any time.

When we try to formalize this notion of sequential rationality, we encounter another subtlety: since our protocols will rely on cryptographic commitment schemes, our definitions should capture the fact that the coalition is polynomially bounded. For example, it could be that there *exists* a set of execution prefixes that account for non-negligible probability mass, such that if A deviated conditioned on having observed those prefixes, it would have gained noticeably. However, it might be that these prefixes are computationally infeasible to recognize, since recognizing them might involve, say, breaking cryptographic commitments. As a result, our definitions actually stipulate that, for any *polynomially bounded* coalition strategy that *wants to deviate with non-negligible probability* at some point in the execution, deviating will not *conditionally* improve the coalition's utility by more than a noticeable amount.

To formally define our sequentially approximately fair notions, we first introduce some probability notations.

Probability notation. In this paper, we use the acronym p.p.t. to mean expected probabilistic polynomial-time. Let Π denote the original honest protocol. However, a non-uniform p.p.t. coalition $A \subset [n]$ might deviate from the original protocol and we use S to denote the strategy of A. As a special case, we use the notation $A(\Pi)$ to mean that the coalition A simply follows the honest protocol and does not deviate. Let κ be the security parameter. We use the notation $tr \leftarrow \mathsf{Exec}^{A(S)}$ to denote a random sample of the protocol execution, where the honest players $[n]\backslash A$, interact with the coalition A which adopts the strategy S. The random experiment $\mathsf{Exec}^{A(S)}$ produces an *execution trace* tr (also called a *trace* for short), which consists of all the messages and the internal states of all players throughout the entire execution. Once the coalition A's strategy S is fixed, all players' internal states and messages in all rounds would be uniquely determined by all players' randomness in all rounds — thus one can also equivalently think of tr as the sequence of *all* players' random coins in all rounds.

An event $\mathsf{Evt}(tr)$ is identified with its indicator function that takes a trace tr and returns either 1 (meaning the event happens) or 0. For example, we use $W^A(tr) = 1$ to indicate that one player in A is elected as the leader in the end.

We use $\Pr[\mathsf{Exec}^{A(S)}(1^\kappa) : \mathsf{Evt}] := \Pr[tr \leftarrow \mathsf{Exec}^{\Pi,A(S)}(1^\kappa) : \mathsf{Evt}(tr)]$ to denote the probability that when the coalition A adopts strategy S, the event Evt happens. Similarly, given events Evt_1 and Evt_2, we use $\Pr[\mathsf{Exec}^{A(S)}(1^\kappa) : \mathsf{Evt}_1 \mid \mathsf{Evt}_2]$ to denote the conditional probability that when the coalition A adopts strategy S and conditioning on the event Evt_2, event Evt_1 also happens. The same notation extends to expectation $\mathbf{E}[\cdot]$.

Deviation event. Given a strategy S of the coalition A, we define the deviation event $\mathsf{Dev}^{A(S)}(tr)$ as follows:

- for each round $r = 1, 2, \ldots$: replay the trace tr (which contains all players' random coins) till the beginning of round r, immedately after the coalition A has observed all honest nodes' round-r messages; at this moment, check whether the strategy S adopted by A would deviate from the honest protocol

Π in round r (i.e., whether S would send a message that differs from what the honest strategy would have sent, suppose that the random coins of S have been fixed by the trace tr); if yes, return 1;
– return 0 if the strategy S adopted by A does not actually deviate from Π till the end.

Intuitively, we say that a protocol satisfies sequential CSP-fairness against the coalition A iff either A never wants to deviate except with negligible probability (condition 1 in Definition 1); or conditioned on deviating, A does not do noticeably better (condition 2 in Definition 1).

Definition 1 (Sequential CSP-fairness). *Let $\epsilon \in (0,1)$. We say that a leader election protocol Π achieves $(1-\epsilon)$-sequential-CSP-fairness against a (non-uniform p.p.t.) coalition $A \subseteq [n]$ iff for any strategy S by A, there exist a negligible function* $\mathsf{negl}(\cdot)$*, such that and for all κ, at least one of the following holds — recall that W^A is the event that one of the coalition members in A is elected leader:*

1. $\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : \mathsf{Dev}^{A(S)}\right] \leq \mathsf{negl}(\kappa)$,
2. $\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : W^A \;\middle|\; \mathsf{Dev}^{A(S)}\right] \leq \frac{1}{1-\epsilon} \cdot \Pr\left[\mathsf{Exec}^{A(\Pi)}(1^\kappa) : W^A \;\middle|\; \mathsf{Dev}^{A(S)}\right] + \mathsf{negl}(\kappa)$.

In the above, the left-hand-side $\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : W^A \;\middle|\; \mathsf{Dev}^{A(S)}\right]$ means the conditional probability that $A(S)$, i.e., a coalition A adopting strategy S, is elected leader, conditioned on $\mathsf{Dev}^{A(S)}$, i.e., that $A(S)$ decided to deviate from honest behavior. The right-hand-side $\Pr\left[\mathsf{Exec}^{A(\Pi)}(1^\kappa) : W^A \;\middle|\; \mathsf{Dev}^{A(S)}\right]$ means *the conditional probability for A to win, had A continued to adopt the honest strategy throughout, even though $A(S)$ had wanted to deviate at some point in the protocol* — the conditional probability is calculated when conditioning on traces where $A(S)$ would have deviated[10]. Intuitively, Condition 2 above says that conditioned on the strategy S deciding to deviate, the coalition A cannot benefit itself noticeably in comparison with just executing honestly to the end.

We can similarly define the sequential approximate maximin fairness.

Definition 2 (Sequential maximin fairness). *Let $\epsilon \in (0,1)$. We say that a leader election protocol Π achieves $(1-\epsilon)$-sequential-maximin-fairness against a (non-uniform p.p.t.) coalition $A \subseteq [n]$ iff for any strategy S by A, there exist a negligible function* $\mathsf{negl}(\cdot)$*, such that for all κ, at least one of the following holds:*

[10] Note that the event $\mathsf{Dev}^{A(S)}(tr)$ is well-defined, even if tr is sampled from $\mathsf{Exec}^{A(\Pi)}$, i.e., an execution in which A adopts the honest strategy. In this case, $\mathsf{Dev}^{A(S)}(tr)$ means the following: had A instead adopted the strategy S rather than the honest strategy Π, *is there a round in which S would have started to deviate from the honest protocol, given that all players' randomness in all rounds is fixed by tr.*

1. $\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : \mathsf{Dev}^{A(S)}\right] \leq \mathsf{negl}(\kappa)$,

2. for any $i \notin A$, let W^i be the event that player i is elected as the leader, it holds that

$$\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : W^i \;\middle|\; \mathsf{Dev}^{A(S)}\right] \geq (1-\epsilon)\cdot\Pr\left[\mathsf{Exec}^{A(\Pi)}(1^\kappa) : W^i \;\middle|\; \mathsf{Dev}^{A(S)}\right] - \mathsf{negl}(\kappa).$$

The following fact says that the sequentially rational notions implies the corresponding non-sequential counterparts defined earlier in Section 2.2.

Fact 1 (Sequential notions are stronger) *Let $\epsilon(n, \kappa) \in (0, 1)$ be a non-negligible function. If a leader election protocol satisfies $(1-\epsilon)$-sequential-CSP-fairness (or $(1-\epsilon)$-sequential-maximin-fairness resp.) against the coalition $A \subseteq [n]$, then for $\epsilon'(n, \kappa) = \epsilon(n, \kappa) + \mathsf{negl}(\kappa)$ where $\mathsf{negl}(\cdot)$ is some negligible function, then, the same protocol also satisfies non-sequential $(1-\epsilon')$-CSP-fairness (or non-sequential $(1-\epsilon')$-maximin-fairness resp.) against A.*

Proof. Deferred to the online full version [15].

We show that if the slack ϵ is constrained to being negligibly small, then in fact the non-sequential notions imply the sequential notions too. However, this direction is not true when the slack ϵ may be non-negligible.

Fact 2 *If a protocol Π satisfies $(1-\mathsf{negl}(\kappa))$-CSP-fairness (or $(1-\mathsf{negl}(\kappa))$-maximin-fairness resp.) against the coalition $A \subset [n]$ for some negligible function $\mathsf{negl}(\cdot)$, then Π satisfies $(1-\mathsf{negl}'(\kappa))$-sequential-CSP-fairness (or $(1-\mathsf{negl}(\kappa))$-sequential-maximin-fairness resp.) against A for some negligible function $\mathsf{negl}'(\cdot)$.*

Proof. Deferred to the online full version [15].

3.2 Fairness of the Tournament Tree Protocol

Instantiated with a suitable cryptographic commitment protocol (described in the online full version [15]), the folklore tournament-tree protocol satisfies $(1-\mathsf{negl}(\kappa))$-sequential-CSP-fairness and $(1-\mathsf{negl}(\kappa))$-sequential-maximin-fairness against coalitions of arbitrarily sizes, as stated below:

Theorem 3 (Tournament-tree protocol). *Suppose that n is the number of players and κ is the security parameter. Then, the tournament-tree protocol, when instantiated with a suitable publicly verifiable, non-malleable commitment scheme as defined in the online full version [15], satisfies $(1-\mathsf{negl}(\kappa))$-sequential-CSP-fairness and $(1-\mathsf{negl}(\kappa))$-sequential-maximin-fairness against coalitions of arbitrarily sizes. Moreover, the number of rounds is $O(\log n)$.*

Proof. Deferred to the online full version [15].

4 Formal Description of Our Scheme

4.1 Description of Our Scheme Assuming Idealized Cryptography

Our scheme makes use of an $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-averaging sampler which we define in the online full version [15]. We will first describe our scheme assuming idealized commitments $\mathcal{F}_{\mathrm{comm}}$ and an ideal MPC functionality $\mathcal{F}_{\mathrm{mpc}}$ described earlier in Section 2.4. Later in Section 4.2, we will instantiate the ideal cryptographic primitives with actual cryptography. In the scheme below, committing to a value is performed by sending it to $\mathcal{F}_{\mathrm{comm}}$, and opening is performed by instructing $\mathcal{F}_{\mathrm{comm}}$ to send the opening to everyone.

Our leader election protocol (assuming idealized cryptography)

Parameters. For some $r := r(n)$, suppose that we would like to achieve round complexity $O(r)$ satisfying $C_0 \log\log n < r(n) < C_1 \log n$, where C_0 and C_1 are suitable constants. We set the parameters as follows:

- Let $B := \frac{n}{2^{9r}}$ such that the expected number of players in a bin (assuming honest behavior) is $\frac{n}{B} = 2^{9r}$ in the preliminary committee election.
- The parameters of the sampler are chosen as below: v is chosen such that $\frac{2^v}{n} = 2^{0.5r}$. Let $\epsilon_{\mathrm{s}} := 2^{-6r}$, and $\delta_{\mathrm{s}} := 2^{-(1-\frac{\psi}{2})|\mathcal{U}|}$, where ψ denotes a lower bound on the fraction of honest players, we shall assume $\psi \geq \frac{1}{2^{\Theta(r)}}$, which means that $|A| \leq (1 - \frac{1}{2^{\Theta(r)}})n$. Let $d = (|\mathcal{U}|/\epsilon_{\mathrm{s}})^{\widetilde{c}}$, where $\widetilde{c}$ is the universal constant specificied in the online full version [15].
- Let $\eta := 1/2^{0.2r}$.

Our protocol.

1. *Elect the preliminary committee $\mathcal{U}$ using lightest bin.* Everyone $i \in [n]$ broadcasts a random index $\beta_i \in [B]$ indicating its choice of bin where B denotes the number of bins. The bin with the lightest load is selected as the preliminary committee $\mathcal{U}$. Break ties with lexicographically the smallest bin.
2. *Elect the final committee $\mathcal{C}$.* Let $\mathsf{Samp} : \{0,1\}^{|\mathcal{U}|} \to \{\{0,1\}^v\}^d$ denote an explicit $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-averaging sampler. If it is not the case that $|\mathcal{U}| \geq \log \frac{1}{\delta_{\mathrm{s}}} + c \cdot v$ (see the online full version [15]), simply abort with the exception param_error and output player 1 as the leader.
   (a) Every player sends `share` to $\mathcal{F}^{\eta}_{\mathrm{mpc}}$, and receives `ok` from $\mathcal{F}^{\eta}_{\mathrm{mpc}}$.
   (b) Every player $i \in [n]$ commits to a randomly selected unmasked v-id henceforth denoted $y_i \in \{0,1\}^v$.
   (c) Every player in the preliminary committee $i \in \mathcal{U}$ broadcasts a bit x_i. Let x be the concatenation of all of $\{x_i\}_{i \in \mathcal{U}}$ in increasing order of the players' indices — here for any player j who has aborted, its x_j is treated as 0.
   (d) Every player $i \in [n]$ now opens the committed string $y_i \in \{0,1\}^v$.

(e) Input `recons` to $\mathcal{F}^{\eta}_{\text{mpc}}$, and receive a mask vector z from $\mathcal{F}^{\eta}_{\text{mpc}}$.

(f) Parse $z := (z_1, \ldots, z_n)$ where each $z_j \in \{0,1\}^v$ for $j \in [n]$. We now view $y_i \oplus z_i$ as player i's finalized v-id, which corresponds to a point in the output range of the sampler Samp. The final committee $\mathcal{C}$ is defined as a *multiset* constructed as follows: for $j \in [d]$, if there is exactly one player $i \in [n]$ who opened y_i and whose final v-id $y_i \oplus z_i = \mathsf{Samp}_j(x)$, then add i to $\mathcal{C}$.

3. *Elect leader among final committee.* The final committee run the tournament-tree protocol to elect a final leader.[a] In case the final committee is empty, simply output player 1 as the leader.

[a] When the ideal $\mathcal{F}_{\text{comm}}$ and $\mathcal{F}^{\eta}_{\text{mpc}}$ are instantiated with actual cryptography later in Section 4.2, the opening/reconstruction messages will be posted to the broadcast channel such that the elected leader can be determined from the collection of messages posted to the broadcast channel.

4.2 Instantiating the Scheme with Real-World Cryptography

Our final protocol replaces the ideal commitment and $\mathcal{F}_{\text{mpc}}$ with actual cryptography. To achieve this, we take an intermediate step and consider an **IdealZK**-hybrid protocol where **IdealZK** is an idealized zero-knowledge proof functionality which we formally define in the online full version [15]. We first instantiate the ideal commitment and $\mathcal{F}_{\text{mpc}}$ using a protocol in the **IdealZK**-hybrid world, and then we use the elegant techniques of Pass [32] to instantiate the protocol with actual cryptography with only $O(1)$ round blowup, while allowing bounded concurrent composition *without any common reference string or trusted setup.* In our case, the total number of concurrent sessions of the cryptographic protocols is a-priori known given n.

Instantiating the ideal commitments with non-malleable commitments. We will instantiate the ideal commitments using a publicly verifiable, non-malleable commitment (NMC) scheme which is defined in the online full version [15]. Basically, to commit to a string, a player invokes n instances of NMC, one for each of the n recipients. To open a previously committed string, post the openings corresponding to all n instances, and the opening is successful iff all n instances open to the same string. We may assume that messages are posted to the broadcast channel and it can be publicly checked what a commitment opens to. An honest committer's commitment will always successfully open even when the receiver is malicious.

Instantiating the $\mathcal{F}_{\text{mpc}}$ with bounded concurrent zero-knowledge proofs. To instantiate $\mathcal{F}_{\text{mpc}}$ with actual cryptography, we first instantiate it in **IdealZK**-hybrid world. Then, we use the bounded concurrent zero-knowledge proofs of Pass [32] to replace the **IdealZK** instances with actual zero-knowledge proofs.

Therefore, it suffices to describe how to replace $\mathcal{F}_{\text{mpc}}$ with a protocol Π_{mpc} in the **IdealZK**-hybrid world. This protocol actually does not realize $\mathcal{F}_{\text{mpc}}$ with full simulation security[11]. Yet, we can later prove that when we replace $\mathcal{F}_{\text{mpc}}$ with this protocol, the game theoretic fairness properties we care about extend to the real-world protocol.

Π_{mpc}: instantiating $\mathcal{F}^{\eta}_{\text{mpc}}$ in the IdealZK-hybrid world

Let comm be a perfectly binding and computationally hiding (non-interactive) commitment scheme. We assume that committing to a string is accomplished by committing to each individual bit. Let $\eta \in (0, 1)$ be a parameter.

Sharing phase.

1. Every player i chooses a random string $\mathsf{coins}_i \in \{0,1\}^{vn}$. It splits coins_i into a $\lceil \eta \cdot n \rceil$-out-of-$n$ Shamir secret shares, and let $\mathsf{coins}_{i,j}$ be the j-th share. Next, for each $j \in [n]$, player i computes the commitment $\overline{\mathsf{coins}}_{i,j} := \mathsf{comm}(\mathsf{coins}_{i,j}, \rho_{i,j})$ where $\rho_{i,j}$ denotes some fresh randomess consumed by the commitment scheme, and it posts the commitment message $\{\overline{\mathsf{coins}}_{i,j}\}_{j\in[n]}$ to the broadcast channel.
2. Player i does the following for each $j \in [n]$:
   - invokes an **IdealZK** instance denoted $\mathbf{IdealZK}_{i,j}$ to prove that the commitment message $\{\overline{\mathsf{coins}}_{i,k}\}_{k\in[n]}$ it has posted is computed correctly, by supplying to $\mathbf{IdealZK}_{i,j}$ 1) the statement $\{\overline{\mathsf{coins}}_{i,k}\}_{k\in[n]}$ and 2) all the random coins used in computing the commitment message. $\mathbf{IdealZK}_{i,j}$ checks the following NP relation: all the commitments are computed correctly, and moreover, the openings form a valid $\lceil \eta n \rceil$-out-of-n secret sharing.
   - gives player j the opening $(\mathsf{coins}_{i,j}, \rho_{i,j})$.
3. A player $i \in [n]$ does the following: for every $j \in [n]$, if player i
   - has seen a message $\{\overline{\mathsf{coins}}_{j,k}\}_{k\in[n]}$ posted by j;
   - has received the message $(\{\overline{\mathsf{coins}}_{j,k}\}_{k\in[n]}, 1)$ from $\mathbf{IdealZK}_{j,i}$ where the statement must match the message posted by j; and
   - has received a correct opening $(\mathsf{coins}_{j,i}, \rho_{j,i})$ w.r.t. the i-th coordinate of j's posted message $\{\overline{\mathsf{coins}}_{j,k}\}_{k\in[n]}$, that is, $\overline{\mathsf{coins}}_{j,i}$.

   then, it posts the tuple $(\mathtt{ok}, j)$ to the broadcast channel.
4. Every player i does the following: for every $j \in [n]$ who has obtained an approval message $\mathtt{ok}$ from at least $(1-\eta)n$ players, add j to the set S. If $|S| \geq \eta n$, then let $\mathsf{succ} := 1$; else let $\mathsf{succ} := 0$. Output $\mathtt{ok}$.

Reconstruction phase. If $\mathsf{succ} = 0$, simply output the $\mathbf{0}$ vector. Else continue with the following.

[11] The reason we do not fully simulate $\mathcal{F}_{\text{mpc}}$ is due to technicalities arising from the requirement that the outcome of the leader election be publicly computable from all the messages posted to the broadcast channel.

1. For every player $j \in S$, if the current player i posted $(\texttt{ok},\ j)$ during the sharing phase, then let $(\mathsf{coins}_{j,i}, \rho_{j,i})$ be the correct opening received from j during the sharing phase, post $(j, \mathsf{coins}_{j,i}, \rho_{j,i})$ to the broadcast channel.
2. For every tuple $(j, \mathsf{coins}_{j,k}, \rho_{j,k})$ received from some player $k \in [n]$, if $j \in S$ and $(\mathsf{coins}_{j,k}, \rho_{j,k})$ is a valid opening w.r.t. the k-th coordinate of j's commitment message posted during the sharing phase, then accept this share $(k, \mathsf{coins}_{j,k})$ of coins_j.

   For every $j \in S$, use all accepted shares to reconstruct coins_j. Output $z := \oplus_{j \in S} \mathsf{coins}_j$ if the reconstruction of every coins_j for $j \in S$ is successful; else output the vector $\mathbf{0}$.

Theorem 4 (Main theorem). *Assume the existence of enhanced trapdoor permutations and collision resistant hash functions. Then, there exists an* $O(r)$*-round leader election protocol that achieves* $(1 - 2^{-\Theta(r)})$*-sequential-maximin-fairness against a non-uniform p.p.t. coalition of size at most* $(1 - 2^{-\Theta(r)}) \cdot n$*, and* $(1 - 2^{-\Theta(r)})$*-sequential-CSP-fairness against a non-uniform p.p.t. coalition of arbitrary size.*

Proof. The theorem results from the construction presented in this section. The detailed proofs are presented in Sections 5 and the online full version [15].

5 Proofs for the Ideal-World Protocol

5.1 Bounding the Preliminary Committee's Size

Since the preliminary committee $\mathcal{U}$ is chosen from a lightest bin, it is immediate that $|\mathcal{U}| \leq \left\lfloor \frac{n}{B} \right\rfloor$. The next lemma states that there is a sufficient number of honest players in $\mathcal{U}$ with high probability.

Lemma 1 (Sufficient honest players in the preliminary committee). *Suppose for some* $\psi \in (0, 0.5)$*, there are at least* $\psi \cdot n$ *honest players. Let* $|\mathcal{U}_H|$ *denote the number of honest players in the preliminary committee* $\mathcal{U}$*. Then, for* $\gamma \in (0, 1)$*, the following holds:*

$$\Pr\left[|\mathcal{U}_H| \leq (1 - \gamma) \cdot \frac{\psi n}{B}\right] \leq B \cdot \exp\left(-\gamma^2 \cdot \frac{\psi n}{2B}\right).$$

In particular, if $\frac{n}{B} = 2^{9r}$ *and* $C_0 \log\log n \leq r \leq C_1 \log\log n$ *for appropriate constants* C_0 *and* C_1*, and* $\psi \geq 2^{-r}$*, then the number of honest players in the preliminary committee is at least* $0.9\psi n/B$*, except with* $\exp(-2^{7r})$ *probability.*

Proof. By the Chernoff bound, except with probability $\exp\left(-\gamma^2 \cdot \frac{\psi n}{2B}\right)$, the number of honest players in any particular bin is greater than $(1 - \gamma) \cdot \frac{\psi n}{B}$. The union bound over all the B bins gives the required result.

The following fact makes sure that the sampler needed by our protocol exists except with doubly-exponentially small in r probability as long as at least a $\psi(n) \geq 1/2^r$ fraction of the players are honest.

Fact 3 *Suppose that the honest fraction $\psi \geq \frac{1}{2^r}$ and that our protocol uses the aforementioned parameters. We have that $|\mathcal{U}| \geq \log(1/\delta_{\mathrm{s}}) + c \cdot v$ except with $\exp(-\Omega(2^{7r}))$ probability.*

Proof. Since we choose $\delta_{\mathrm{s}} := 2^{-(1-\frac{\psi}{2})|\mathcal{U}|}$, the expression to verify can be rewritten as $|\mathcal{U}| \geq (1-\psi/2)|\mathcal{U}| + c \cdot v$, which is equivalent to:

$$0.5\psi \cdot |\mathcal{U}| \geq c \cdot v = c \cdot (\log n + 0.5r).$$

Due to Lemma 1, the size of the preliminary committee is at least $\frac{0.9\psi n}{B}$, except $\exp(-\Omega(2^{7r}))$ probability. Therefore, it suffices to show that

$$0.5\psi \cdot 0.9\psi n/B \geq 0.45 \cdot 2^{-2r} \cdot 2^{9r} \geq c \cdot (\log n + 0.5r),$$

where the last inequality holds as long as $r \geq C_0 \log\log n$ for a sufficiently large constant C_0.

5.2 Terminology and Notations

We first present proofs for our protocol in Section 4 assuming idealized $\mathcal{F}_{\mathrm{comm}}$ and $\mathcal{F}_{\mathrm{mpc}}$. However, we shall assume that the tournament-tree protocol is instantiated with real cryptography as explained in the online full version [15], since we will use the tournament-tree protocol's fairness properties as a blackbox in our proofs. In the online full version [15], we prove that the relevant security properties extend to the real-world protocol when the idealized cryptographic primitives are instantiated with actual cryptography.

Recall that A denotes the coalition; we often refer to players in A as corrupt and players outside A as honest. Further, we often use the notation $H := [n] \backslash A$ to denote the set of honest players. For $S \subseteq [n]$, we use the notation $x_S := \{x_i\}_{i \in S}$ and y_S is also similarly defined.

5.3 Composition of the Final Committee

Lemma 2 (Final committee composition). *Suppose that the honest fraction $\psi \geq 2\eta = 2 \cdot \frac{1}{2^{0.2r}}$ and that our protocol uses the aforementioned parameters. Fix $\mathcal{N}$ to be an arbitrary set of (distinct) final* v-ids *in the sampler's output range $\{0,1\}^v$ where $|\mathcal{N}| \leq n$. Let $\mathcal{C}_{\mathcal{N}}$ be the (multi-)set of final v-ids in $\mathcal{N}$ chosen by $\mathsf{Samp}(x)$. Let[12] $\epsilon_0 = \epsilon_{\mathrm{s}} \cdot \frac{2^v}{|\mathcal{N}|}$. Then, conditioned on no $\mathsf{param\_error}$ and $|\mathcal{U}_H| \geq 0.9\psi \cdot n/B$, with probability at least $1 - \exp(-\Omega(2^{7r}))$ over the choice of x_H, $\mathcal{C}_{\mathcal{N}}$ has size in the range $[1-\epsilon_0, 1+\epsilon_0] \cdot d \cdot \frac{|\mathcal{N}|}{2^v}$.*

[12] Note that ϵ_0 would be very large if $\mathcal{N}$ is too tiny, but our usage later will guarantee that $\mathcal{N}$ is not too tiny.

Alternatively, suppose there is some upper bound $|\mathcal{N}| \leq N$, and we set $\epsilon_0 = \epsilon_{\mathrm{s}} \cdot \frac{2^v}{N}$. Then, with conditional probability at least $1 - \exp(-\Omega(2^{7r}))$ under the events, $\mathcal{C}_{\mathcal{N}}$ has size at most $(1+\epsilon_0) \cdot d \cdot \frac{N}{2^v}$.

Proof. Let the final committee $\mathcal{C}_{\mathcal{N}}$ be the multi-set of v-ids in $\mathcal{N}$ chosen by the $\mathsf{Samp}(x)$. We shall show that, using the sampler theorem in the online full version [15], except with probability $p := \exp(-\Omega(2^{6r}))$ over the choice of x_H,

$$|\mathcal{C}_{\mathcal{N}}| \in [1-\epsilon_0, 1+\epsilon_0] \cdot d \cdot \frac{|\mathcal{N}|}{2^v}. \tag{1}$$

Observing that $\epsilon_{\mathrm{s}} = \epsilon_0 \cdot \frac{|\mathcal{N}|}{2^v}$, by the property of the $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-averaging sampler, except for at most $2^{|\mathcal{U}|} \cdot \delta_{\mathrm{s}} = 2^{0.5\psi|\mathcal{U}|}$ number of *bad* inputs to the sampler, the size of $\mathcal{C}_{\mathcal{N}}$ satisfies (1).

We say that some choice of $x_{H \cap \mathcal{U}}$ is *bad* if there exists a corrupt choice of $x_{A \cap \mathcal{U}}$ such that the combination of $x_{H \cap \mathcal{U}}$ and $x_{A \cap \mathcal{U}}$ (arranged in the right order) will lead to $\mathcal{C}_{\mathcal{N}}$ such that (1) is violated. Otherwise, we say that $x_{H \cap \mathcal{U}}$ is good. Note that if $x_{H \cap \mathcal{U}}$ is good, it means that no matter how the adversary chooses $x_{A \cap \mathcal{U}}$, it cannot make $\mathcal{C}_{\mathcal{N}}$ violate (1).

Since honest players choose their $x_{H \cap \mathcal{U}}$ at random, we next claim that the fraction of bad $x_{H \cap \mathcal{U}}$ is bounded by $2^{-0.3\psi|\mathcal{U}|} \leq 2^{-0.27\psi^2 \cdot n/B} \leq 2^{-\Omega(2^{7r})}$. The claim is true; otherwise, the number of bad inputs to the sampler is at least $2^{-0.3\psi|\mathcal{U}|} \cdot 2^{0.9\psi|\mathcal{U}|} = 2^{0.6\psi|\mathcal{U}|}$ and thus we have reached a contradiction. Finally, a union bound over all the above bad events shows that except with probability at most $\exp(-\Omega(2^{7r}))$, $\mathcal{C}_{\mathcal{N}}$ respects the range in (1).

The alternative case when there is an upper bound $|\mathcal{N}| \leq N$ uses the same argument, but we just need one direction of the inequality from the sampler.

The above Lemma 2 immediately implies the following bound on the final committee size.

Lemma 3 (Final committee not too large). *Suppose that the honest fraction $\psi > 2\eta = 2 \cdot \frac{1}{2^{0.2r}}$ and that our protocol uses the aforementioned parameters. Let $\epsilon_0 = \epsilon_{\mathrm{s}} \cdot \frac{2^v}{n} = 2^{-5.5r}$. Then, with probability at least $1 - \exp(-\Omega(2^{6r}))$, the final committee $\mathcal{C}$ has size at most $(1+\epsilon_0) \cdot d \cdot \frac{n}{2^v} \leq 2^{O(r)}$, and the protocol does not throw* $\mathsf{param\_error}$. *In particular, with probability at least $1 - \exp(-\Omega(2^{6r}))$, the protocol has round complexity at most $O(r)$.*

Proof. Due to Lemma 1, except with $\exp(-\Omega(2^{7r}))$ probability, $|\mathcal{U}_H| \geq 0.9\psi \cdot n/B \geq 0.9\psi \cdot |\mathcal{U}|$. Further, due to Fact 3, $\mathsf{param\_error}$ does not happen except with $\exp(-\Omega(2^{7r}))$ probability. Conditioned on these bad events not happening, we now use Lemma 2. In this case, the n players can choose at most n final v-ids, i.e., $|\mathcal{N}| \leq n$. The range in (1) implies that except with $\exp(-\Omega(2^{6r}))$ over the choice of x_H, the final committee $\mathcal{C}$ has size at most:

$$d(\frac{n}{2^v}+\epsilon_{\mathrm{s}}) = (1+\epsilon_0) \cdot d \cdot \frac{n}{2^v} \leq d \cdot (2^{-0.5r}+2^{-6r}) = (1+2^{-5.5r}) \cdot (|\mathcal{U}|/\epsilon_{\mathrm{s}})^{\widetilde{c}} \cdot 2^{-0.5r} \leq (1+2^{-5.5r}) \cdot 2^{15r\widetilde{c}} \cdot 2^{-0.5r}.$$

We shall consider the following bad events in our proofs. Recall that conditioned on any coin used in the lightest-bin protocol for the preliminary committee election, the protocol still has independent randomness x chosen by the preliminary committee as input for the averaging sampler, the unmasked v-ids y chosen by all players, as well as the mask vector z.

- Event $\mathsf{param\_error}$. Recall that this happens when the preliminary comittee selected does not have the desirable properties; by Lemma 1 and Fact 3, this bad event happens with probability at most $\exp(-\Omega(2^{7r}))$.
- Event bad_1: out of the d samples from the $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-sampler, at least $(1 + \epsilon_0) \cdot d \cdot \frac{n}{2^v}$ number of them correspond to corrupt players' final v-ids, where $\epsilon_0 := 2^{-6r} \cdot 2^{0.5r}$ is defined as in Lemma 3. Assuming the honest fraction $\psi \geq 2\eta$, by Lemma 3, $\Pr[\mathsf{bad}_1] \leq \exp(-\Omega(2^{6r}))$. Moreover, observe that bad_1 is determined by x, y_A, and z_A, and is independent of y_H and z_H.
- Event bad_2: the final committee $\mathcal{C}$ has size greater than $(1+\epsilon_0) \cdot d \cdot \frac{n}{2^v}$. Again assuming $\psi \geq 2\eta$, Lemma 3 implies that $\Pr[\mathsf{bad}_2] \leq \exp(-\Omega(2^{6r}))$. Observe that bad_2 depends on x, y, and z.

Lemma 4 (Influence of an honest player in the final committee). *Suppose that $|A| < (1 - 2\eta)n$, i.e., $\frac{h}{n} = \psi > 2\eta \geq \frac{1}{2^r}$. For an honest player $i \notin A$, let M_i be its multiplicity in the final committee $\mathcal{C}$. Define a random variable Υ_i that equals $\frac{M_i}{|\mathcal{C}|}$, if none of the bad events bad events $\mathsf{param\_error}$ or bad_1 or bad_2 happens; otherwise, Υ_i equals 0.*

Then, $\mathbf{E}[\Upsilon_i] \geq \frac{1}{n}\left(1 - 2^{-0.48r}\right)$, where the expectation is taken over the randomness used in the entire execution.

Proof. For ease of notation, the rest of the proof conditions on the event that during the preliminary committee election, $\mathsf{param\_error}$ does not happen; observe that this bad event happens with probability at most $\exp(-\Omega(2^{7r}))$, by Lemma 1 and Fact 3. Hence, at the end, we just need to multiply any conditional expectation by a factor of $1 - \exp(-\Omega(2^{7r}))$. Recall that we identify an event with its $\{0, 1\}$-indicator random variable.

We next give a lower bound on $\mathbf{E}[M_i | \overline{\mathsf{bad}_1}]$. Since y_H is opened in the last but second step and as long as $|A| < (1 - 2\eta)n$, the reconstruction of z is fully determined before selecting input to the sampler, we may equivalently imagine that y_H is chosen at the end, independently of x, y_A, and z. Since the event bad_1 does not happen, there are at least $d - (1+\epsilon_0) \cdot d \cdot \frac{n}{2^v} = d(1 - (1+\epsilon_0)\frac{n}{2^v}) \geq d(1 - 2^{-0.49r})$ available slots for the honest players' final v-ids, where the inequality follows from $1 + \epsilon_0 \leq 2^{0.01r}$.

For each such slot, player i can get it if it chooses this slot and none of the other honest players choose it; this happens with probability $\frac{1}{2^v} \cdot (1 - \frac{1}{2^v})^{h-1} \geq \frac{1}{2^v}(1 - \frac{n}{2^v}) = \frac{1}{2^v}(1 - 2^{-0.5r})$. Therefore, conditioned on any choice of x, y_A, z, by just using the randomness of y_H, we can conclude that $\mathbf{E}_{y_H}[M_i | \overline{\mathsf{bad}_1}] \geq \frac{d}{2^v} \cdot (1 - 2^{-0.49r})(1 - 2^{-0.5r}) \geq \frac{d}{2^v}(1 - 2^{-0.485r})$, where the last inequality holds for large enough $r = \Omega(1)$.

Since this holds conditioned any any choice of x, y_A, z, we have the desired lower bound on $\mathbf{E}[M_i | \overline{\mathsf{bad}_1}]$.

We next give a lower bound for the following quantity:

$$\mathbf{E}[M_i \cdot \overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}] = \mathbf{E}[M_i | \overline{\mathsf{bad}_1}] \cdot \Pr[\overline{\mathsf{bad}_1}] - \mathbf{E}[M_i \cdot \overline{\mathsf{bad}_1} \cdot \mathsf{bad}_2] \geq \frac{d}{2^v}(1 - 2^{-0.485r}) \cdot \Pr[\overline{\mathsf{bad}_1}] - d \Pr[\mathsf{bad}_2]$$

We use $\mathbf{E}[M_i \cdot \overline{\mathsf{bad}_1} \cdot \mathsf{bad}_2] \leq d \Pr[\mathsf{bad}_2] \leq d \cdot \Pr[\mathsf{bad}_2] \leq d \cdot \exp(-\varOmega(2^{6r})) \leq \frac{d}{2^v} \cdot \exp(-\varOmega(2^{5r}))$ where the last inequality holds because $2^v = n \cdot 2^{0.5r}$ and we assume that $r \geq C_0 \log \log n$ for some suitably large constant C_0. Therefore, we have $\mathbf{E}[M_i \cdot \overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}] \geq \frac{d}{2^v}\left(1 - 2^{-0.485r}\right) \cdot \left(1 - \exp(-\varOmega(2^{6r}))\right) - \frac{d}{2^v} \cdot \exp(-\varOmega(2^{5r})) \geq \frac{d}{2^v}(1 - 2^{-0.483r})$. Finally, we have

$$\begin{aligned}
\mathbf{E}[\varUpsilon_i | \overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}] = \mathbf{E}\left[\frac{M_i}{|\mathcal{C}|} | \overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}\right] &\geq \frac{\mathbf{E}[M_i | \overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}]}{(1+\epsilon_0) \cdot d \cdot \frac{n}{2^v}} \\
&\geq \frac{1}{n}(1 - 2^{-0.483r})(1 - \epsilon_0) \cdot \Pr[\overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}]^{-1} \\
&\geq \frac{1}{n}(1 - 2^{-0.481r}) \cdot \Pr[\overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}]^{-1}.
\end{aligned}$$

Hence, we have the lower bound $\mathbf{E}[\varUpsilon_i] \geq \mathbf{E}[\varUpsilon_i \cdot \overline{\mathsf{bad}_1} \cdot \overline{\mathsf{bad}_2}] \geq \frac{1}{n}(1 - 2^{-0.481r})$.

Finally, recalling so far we have assume that $\mathsf{param\_error}$ does not happen. Therefore, multiplying the above by $(1 - \Pr[\mathsf{param\_error}]) = 1 - \exp(-\varOmega(2^{7r}))$ gives the desired lower bound for the expectation of $\varUpsilon_i$.

Lemma 5 (Sufficient honest players without collision). *Suppose* $n = g + t < V$. *There are* V *bins, of which* t *bins are* bad *and the rest are* good. *Suppose each of* g *balls is thrown into a bin uniformly at random independently. Let* Z *be the number of good bins containing exactly one ball. For any* $0 < \alpha < 1$, *except with probability* $\exp(-\varTheta(\alpha^2 g(1 - \frac{n}{V})))$, *we have* $Z \geq g(1 - \frac{2n}{V} - 2\alpha)$.

Proof. Consider throwing the g balls one by one independently into the bins. For $1 \leq i \leq g$, let $X_i \in \{0, 1\}$ be the indicator random variable for the event that when the i-th ball is thrown, it goes to an empty good bin. Observe that no matter what happens to the first $i - 1$ balls, the event $X_i = 1$ happens with probability at least $1 - \frac{n}{V}$. Hence, $S := \sum_{i=1}^{g} X_i$ stochastically dominates the binomial distribution $\mathsf{Binom}(g, 1 - \frac{n}{V})$ with g trials and success rate $1 - \frac{n}{V}$. By stochastic dominance and the Chernoff bound,

$$\Pr\left[S \leq (1 - \alpha) \cdot g(1 - \frac{n}{V})\right] \leq \exp\left(-\varTheta(\alpha^2 g(1 - \frac{n}{V}))\right)$$

Hence, except with probability $\exp(-\varTheta(\alpha^2 g(1 - \frac{n}{V})))$, we have that $S \geq (1 - \alpha) \cdot g(1 - \frac{n}{V}) \geq g(1 - \frac{n}{V} - \alpha)$.

Finally, observe what happens to the number Z of good bins having exactly one ball as the g balls are thrown one by one. When $X_i = 1$, Z increases by 1; when $X_i = 0$, Z either remains the same or decreases by 1. Hence, at the end, the number Z of good bins having exactly one ball satisfies $Z \geq S - (g - S) = 2S - g$. The result follows.

Lemma 6 (Sufficient honest players in the final committee). *Suppose that $|A| < (1-2\eta)n$. Let $G \subseteq H$ denote an arbitrary subset of honest players with $g = |G|$, where $\frac{g}{n} \geq 1/2^r$. Except with probability $\exp(-\Omega(2^r))$, the number of players from G that are in the final committee*[13] *is at least $g \cdot \frac{d}{2^v} \cdot (1-2^{-0.48r})$.*

As a direct corollary, no matter how large A is, as long as the coalition A adopts the honest strategy, then, for any subset $G \subseteq [n]$ of at least $n/2^r$ players, except with probability $\exp(-\Omega(2^r))$, the number of players from G that are in the final committee is at least $g \cdot \frac{d}{2^v} \cdot (1-2^{-0.48r})$.

Proof. Let $V = 2^v$, and so $\frac{n}{V} = \frac{1}{2^{0.5r}}$. Since $|A| < (1-2\eta)n$, the mask z to be reconstructed later is fully determined before selecting input x to the sampler — in this case, we can imagine that y_G is chosen and revealed at the end, independent of x, $y_{[n]\backslash G}$, and z. Setting $\alpha := \frac{1}{2^r}$ in Lemma 5, we have, except with probability $p \leq \exp\left(-\Omega(\frac{1}{2^{2r}} \cdot g \cdot (1-2^{-0.5r}))\right) \leq \exp\left(-\Omega(\frac{n}{2^{3r}})\right)$, the number of players in G whose final v-id has no collision is at least $Z := g(1-2\cdot 2^{-0.5r}-2\cdot 2^{-r}) \geq \frac{g}{2}$. Recall that $r \leq C_1 \log n$, and, as long as the constant C_1 is sufficiently small, we have that $n > 2^{4r}$, and thus $p \leq \exp(-\Omega(2^r))$.

Setting $\epsilon_0 := \epsilon_{\mathrm{s}} \cdot \frac{2^v}{|Z|} \leq 2 \cdot 2^{-6r} \cdot 2^{1.5r}$, and using Lemma 2, we can show that except with probability $\exp(-\Omega(2^r))$, the number of players from G in the final committee is at least $(1-\epsilon_0) \cdot d \cdot \frac{Z}{2^v} \geq g \cdot \frac{d}{2^v} \cdot (1-2^{-0.48r})$.

5.4 Maximin Fairness

In this section, we will prove the following lemma.

Lemma 7 (Ideal-world protocol: maximin fairness). *The ideal-world protocol (i.e., instantiated with $\mathcal{F}_{\mathrm{comm}}$ and $\mathcal{F}_{\mathrm{mpc}}$) satisfies $(1-2^{-0.4r}) = (1-2^{-\Theta(r)})$-sequential-maximin-fairness against any non-uniform p.p.t. coalition*[14] *of size at most $(1-2\eta)n = (1-2^{-\Theta(r)})n$.*

Proof. Due to a lemma proven in the online full version [15], we can do a round-by-round analysis. Let r^* be the first round in which the coalition deviates. Let $\widetilde{r}$ be the round in which all players reconstruct the mask vector z. Throughout, we may assume that $A < (1-2\eta)n$. Further, for each round r^*, we may assume that $\Pr[\mathsf{Dev}^{r^*}]$ is non-negligible where Dev^{r^*} denotes the event that A deviates first in round r^*. We want to show that conditioned on this non-negligible probability event Dev^{r^*}, A cannot conditionally harm an honest individual noticeably, or conditionally increase its own winning probability noticeably.

Easy case: $r^* > \widetilde{r}$. This means the coalition A will deviate only in the tournament tree protocol, whose sequential maximin fairness holds according to Theorem 3. This means each honest player can only be hurt negligibly more.

[13] Throughout, a player with multiplicity μ in the final committee is counted μ times.

[14] Recall that the tournament-tree protocol is still instantiated with real cryptography.

Easy case: $r^* = \widetilde{r}$. As mentioned earlier, as long as $|A| < (1-2\eta)n$, in this round, no matter what A does, reconstruction of z is guaranteed and the reconstructed value is unique.

Slightly more complicated case: $r^* = \widetilde{r} - 1$. This is the case when the coalition A deviates in the round in which the unmasked v-ids y are opened. Since we are using an ideal $\mathcal{F}_{\text{comm}}$, the only possible deviation in round $r^* = \widetilde{r} - 1$ is if some member of the coalition $i \in A$ fails to open its committed its y_i value.

We consider two cases.

- First, suppose that $|A| \geq \eta n$. This means that the adversarial coalition already knows the committed mask z at the end of the sharing phase. In this case, the z mask to be reconstructed is uniquely determined at the end of the sharing phase. In the round $r^* = \widetilde{r} - 1$, to harm any specific honest individual, A's best strategy is the following: for every final v-id in the space $\{0,1\}^v$, if one or more player(s) in A happen(s) to have that final v-id, make exactly one of them open its y_i value, such that there is no internal collision among the coalition A. Due to the sequential fairness of the tournament-tree protocol (i.e., Theorem 3), conditioned on the history of the protocol till the end of round $\widetilde{r}$, every honest final committee member's winning probability is at least $\frac{1}{|\mathcal{C}|} - \mathsf{negl}(\kappa)$, no matter how A behaves in any round greater than $\widetilde{r}$. Therefore, avoiding internal collision but otherwise opening every final v-id is A's best strategy for harming any specific honest player.
  Note that opening the coalition members' unmasked v-ids in an internal-collision-avoiding manner like above does not change whether any honest individual is included in the final committee, but it may increase the final committee size (in comparison with the case when A continues to play honestly). Due to Lemma 6, and since A has acted honestly so far, except with negligible probability, the final committee size is at least $\frac{nd}{2^v}(1 - 2^{-0.48r})$.
  Now, suppose A excludes its members from the final committee due to internal collision. Observe that actually this decision could have been made before the input x to the Samp is chosen. Since there are at most n finalized v-ids with no collision, by Lemma 3, except with $\exp(-2^{\Omega(r)})$ probability (which is negligible if $r \geq C_0 \log\log n$ for a sufficiently large C_0), the final committee has size at most $\frac{nd}{2^v}(1 + 2^{-5.5r})$.
  Therefore, except with negligible probability, for any honest i, the coalition A can only reduce Υ_i by a $1 - 2^{-\Theta(r)}$ factor.
- Second, suppose that $|A| < \eta n$. In this case, A has no information about the mask z, and Dev^{r^*} is independent of z. Further, z is guaranteed to be reconstructed later. In this case, we can reprove Lemma 4 almost identically except that instead of using the randomness y_H, we now use the randomness z_H; further, notice that bad_1 is independent of z_H, and even when conditioning on the non-negligible probability event Dev^{r^*}, the probabilities of bad_1 and bad_2 are still negligible. Therefore, we get that even when conditioning on Dev^{r^*}, for any honest i, the expectation of Υ_i is at least $\frac{1}{n} \cdot (1 - 2^{-0.48r})$ no matter how A behaves during round $\widetilde{r}$ and after. Had A continued to play honestly, using the randomness of z, we know that even when conditioning on

Dev^{r^*}, the expectation of $\varUpsilon_i$ is at least $1/n - \mathsf{negl}(\kappa)$ where the $\mathsf{negl}(\kappa)$ term is due to the negligibly small probability of bad_1 and bad_2 in which case $\varUpsilon_i$ is defined to be 0. (see Lemma 4).
Therefore, deviating in round $\widetilde{r}$ will not reduce any honest individual's conditional winning probability by a $1 - 2^{-\varTheta(r)}$ multiplicative factor.

Remaining case: $r^* < \widetilde{r} - 1$. The rest of the proof focuses on this remaining case. Recall that we assume $\Pr[\mathsf{Dev}^{r^*}] \geq \frac{1}{\mathsf{poly}(n)}$. Let $\mathsf{LEIdeal}$ denote a randomized execution of our ideal-world leader-election protocol described in Section 4.1.

Conditioning on the event Dev^{r^*}, we prove maximin fairness assuming that the coalition A contains no more than a $1 - 2\eta$ fraction of the players. Fix any $i \notin A$. Now, observe the following:

1. Recall that we may assume Dev^{r^*} happens with non-negligible probability. Following the proof of Lemma 4, and observing that before round $\tilde{r}$, the randomness y_H remains hidden and is independent of whatever that has happened so far, we have:
$$\mathbf{E}\left[tr \leftarrow \mathsf{LEIdeal} : \varUpsilon_i | \mathsf{Dev}^{r^*}(tr)\right] \geq \frac{1}{n} \cdot \left(1 - 2^{-0.48r}\right). \tag{2}$$
The only difference in the argument is that both the probabilities $\Pr[\mathsf{bad}_1 | \mathsf{Dev}^{r^*}]$ and $\Pr[\mathsf{bad}_2 | \mathsf{Dev}^{r^*}]$ are at most $\mathsf{poly}(n) \cdot \exp(-\varOmega(2^{6r}))$, which is is still negligible, because we assume that $r = \varOmega(\log \log n)$ is sufficiently large. Indeed, for sufficiently large n, $\mathsf{poly}(n) \cdot \exp(-\varOmega(2^{6r})) \leq \exp(-\varOmega(2^{5.99r}))$, and the proof works as before.
2. We next consider the proof of Lemma 6, but now we conditioned on Dev^{r^*} (which has non-negligible probability). Suppose all players in A actually play honestly. Define bad_3 to be the event that the final committee has size less than $\frac{nd}{2^v} \cdot (1 - 2^{-0.48r})$. Lemma 6 states that $\Pr[\mathsf{bad}_3] \leq \exp(-\varOmega(2^r))$. Since Dev^{r^*} has non-negligible probability, we have $\Pr[\mathsf{bad}_3 | \mathsf{Dev}^{r^*}] \leq \mathsf{poly}(n) \cdot \exp(-\varOmega(2^r)) \leq \exp(-\varOmega(2^{0.99r})) \leq \mathsf{negl}(\kappa)$, where the last inequalities hold for large enough $n \geq \kappa$ because $r \geq \varOmega(\log \log n)$.
This implies that an honest continuation of the execution would lead to a conditional expectation of $\varUpsilon_i$ of at most
$$\frac{d/2^v}{n \cdot \frac{d}{2^v} \cdot (1 - 2^{-0.48r})} + \mathsf{negl}(\kappa) \leq \frac{1}{n} \cdot (1 + 2^{-0.47r}) + \mathsf{negl}(\kappa) \leq \frac{1}{n} \cdot (1 + 2^{-0.46r})$$

Summarizing the above, the ideal protocol is $(1 - 2^{-0.4r})$-sequential-maximin-fair for any coalition that is at most $(1 - 2\eta)n = (1 - 2^{-\varTheta(r)})n$ in size.

Deferred materials. We defer to the online full version [15] 1) proofs of CSP fairness for the ideal-world protocol, 2) proofs for the real-world protocol, and 3) our full lower bound proof. The online full version [15] also contain additional preliminaries, additional proofs for our sequential approximate fairness notion, relationship to the RPD notion [22–24], as well as proofs for the folklore tournament-tree protocol.

Acknowledgment. This work is partially supported by NSF under the award numbers CNS-1601879 and CNS-1561209, a Packard Fellowship, an ONR YIP award, and the Hong Kong RGC under the grants 17200418 and 17201220.

References

1. I. Abraham, D. Dolev, and J. Y. Halpern. Distributed protocols for leader election: A game-theoretic perspective. *ACM Trans. Econ. Comput.*, 7(1), Feb. 2019.
2. D. Alistarh and J. Aspnes. Sub-logarithmic test-and-set against a weak adversary. In D. Peleg, editor, *DISC*, 2011.
3. D. Alistarh, H. Attiya, S. Gilbert, A. Giurgiu, and R. Guerraoui. Fast randomized test-and-set and renaming. In *DISC*, 2010.
4. D. Alistarh, R. Gelashvili, and A. Vladu. How to elect a leader faster than a tournament. In *PODC*, 2015.
5. M. Andrychowicz, S. Dziembowski, D. Malinowski, and u. Mazurek. Secure multiparty computations on bitcoin. *Commun. ACM*, 59(4):76–84, Mar. 2016.
6. M. Bartoletti and R. Zunino. Constant-deposit multiparty lotteries on bitcoin. In *Financial Cryptography and Data Security*, 2017.
7. D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols. In *STOC*, 1990.
8. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *STOC*, 1988.
9. I. Bentov and R. Kumaresan. How to use bitcoin to design fair protocols. In *CRYPTO*, pages 421–439, 2014.
10. M. Blum. Coin flipping by telephone a protocol for solving impossible problems. *SIGACT News*, 15(1):23–27, Jan. 1983.
11. D. Boneh, J. Bonneau, B. Bünz, and B. Fisch. Verifiable delay functions. In *CRYPTO*, 1 2018.
12. D. Boneh, B. Bünz, and B. Fisch. A survey of two verifiable delay functions. Cryptology ePrint Archive, Report 2018/712, 2018.
13. D. Chaum, C. Crépeau, and I. Damgård. Multiparty unconditionally secure protocols (extended abstract). In *STOC*, 1988.
14. K. Chung, Y. Guo, W. Lin, R. Pass, and E. Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, 2018.
15. K.-M. Chung, T.-H. H. Chan, T. Wen, and E. S. (random author ordering). Game-theoretic fairness meets multi-party protocols: The case of leader election. Online full version of this paper, `https://eprint.iacr.org/2020/1591`.
16. R. Cleve. Limits on the security of coin flips when half the processors are faulty. In *STOC*, 1986.
17. P. Daian, R. Pass, and E. Shi. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In *FC*, 2019.
18. I. Damgård and Y. Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In *CRYPTO*, 2005.
19. Y. Dodis. Fault-tolerant leader election and collective coin-flipping in the full information model. Manuscript, 2006.
20. U. Feige. Non-cryptographic selection protocols. In *FOCS*, 1999.
21. R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 1983.

22. J. Garay, J. Katz, B. Tackmann, and V. Zikas. How fair is your protocol? a utility-based approach to protocol optimality. In *PODC*, 2015.
23. J. A. Garay, J. Katz, U. Maurer, B. Tackmann, and V. Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *FOCS*, 2013.
24. J. A. Garay, B. Tackmann, and V. Zikas. Fair distributed computation of reactive functions. In *DISC*, volume 9363, pages 497–512, 2015.
25. G. Giakkoupis and P. Woelfel. On the time and space complexity of randomized test-and-set. In *PODC*, 2012.
26. O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *ACM symposium on Theory of computing (STOC)*, 1987.
27. V. Guruswami, C. Umans, and S. P. Vadhan. Unbalanced expanders and randomness extractors from parvaresh-vardy codes. *J. ACM*, 56(4):20:1–20:34, 2009.
28. A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Crypto*, 2017.
29. E. Korach, S. Kutten, and S. Moran. A modular technique for the design of efficient distributed leader finding algorithms. In *PODC*, page 163–174, 1985.
30. D. Maram, H. Malvai, F. Zhang, N. Jean-Louis, A. Frolov, T. Kell, T. Lobban, C. Moy, A. Juels, and A. Miller. Candid: Can-do decentralized identity with legacy compatibility, sybil-resistance, and accountability. `https://eprint.iacr.org/2020/934`.
31. A. Miller and I. Bentov. Zero-collateral lotteries in bitcoin and ethereum. In *EuroS&P Workshops*, 2017.
32. R. Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *STOC*, 2004.
33. R. Pass and E. Shi. Fruitchains: A fair blockchain. In *PODC*, 2017.
34. O. Reingold, S. P. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *FOCS*, 2000.
35. A. Russell, M. Saks, and D. Zuckerman. Lower bounds for leader election and collective coin-flipping in the perfect information model. In *STOC*, 1999.
36. A. Russell and D. Zuckerman. Perfect information leader election in $\log^* n + o(1)$ rounds. In *FOCS*, 1998.
37. S. P. Vadhan. Pseudorandomness (foundations and trends in theoretical computer science), 2012.
38. D. Zuckerman. Randomness-optimal sampling, extractors, and constructive leader election. In *STOC*, 1996.