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

Kai-Min Chung
Academia Sinica

T-H. Hubert Chan
HKU

Ting Wen
HKU

Elaine Shi*
CMU

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.

*Author ordering is randomized.

Contents

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 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 [Cle86] 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 [GMW87, BGW88, CCD88]. 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 [Blu83]: 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 [Cle86]. 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. [CGL$^+$18] initiated the study of game-theoretic fairness in a multi-party setting. Unfortunately, Chung et al. [CGL$^+$18] proved broad impossibility results (in the corrupt majority setting) for a particular formulation of the multi-party coin toss problem for natural game-theoretic fairness notions. Specifically, suppose that n parties want to toss a *binary* coin, and each player has preference for either the bit 0 or 1. If the outcome agrees with a player's preference, it obtains a utility 1; otherwise, it obtains a utility of 0. Chung et al. [CGL$^+$18] 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[1]. The broad impossiblity results of Chung et al. [CGL+18] for the binary case no longer apply. A folklore approach henceforth called the tournament-tree protocol [MB17, BZ17] 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 [Blu83]; 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 Section B.3).

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 [BBBF18, BBF18]. Also, recall that in the honest majority setting, the standared multi-party computation literature gives us constant-round solutions [BMR90, DI05] 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 [RSZ99,RZ98,Fei99,Dod06]. 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

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

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 1.1 (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 [Cle86]. 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 Section 2.6.

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 [PS17] 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[2]. 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-priori motivation for Alice or Bob to deviate is indeed small. In Section C, 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 1.2. In Appendix A.3, 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. [GKM$^+$13, GKTZ15, GTZ15]. 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 Section 8, 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 1.3 (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. [CGL$^+$18] 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 Section D.

[2]Pass and Shi [PS17] do not consider the threat of a coalition targeting an individual.

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 [MB17, BZ17, BK14, ADMM16]. 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 [MMZ+] 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 [BK14, ADMM16] 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 [MB17, BZ17] 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[3].

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.

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

[3] 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.

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. [CGL$^+$18] 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 $(1-\epsilon)/n$.

It is not hard to see that approximate maximin-fairness and approximate CSP-fairness are not equivalent — we give more explanations in Section A.1.

Remark 2.1 (Coalition-resistant notions of equilibrium)**.** In our definitions, we consider the deviation of a single coalition. This definitional approach is standard in game theory [Fei99, Zuc96, RZ98, Dod06, AGV15, ADH19, AA11, GW12, AAG$^+$10, GHS83, KKM85], 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 2.2 (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 [DPS19, KRDO17].

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 2.3. In the above strawman, one can also replace the committee election with a single iteration of Feige's lightest bin protocol [Fei99]. 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 $\mathcal{F}_{\text{mpc}}^{\eta}$ 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 $\mathcal{F}_{\text{mpc}}^{\eta}$ ideal functionality. Later in Section 5.2, we will instantiate it with an actual, constant-round cryptographic protocol using bounded concurrent MPC techniques [Pas04]. Technically, the real-world cryptographic instantiation does not securely emulate $\mathcal{F}_{\text{mpc}}$ by a standard simulation-based notion; nonetheless, we prove in Section 7 that the fairness properties we care about in the ideal-world protocol (using idealized cryptography) extend to the real-world protocol (using actual cryptography).

$\mathcal{F}_{\text{mpc}}^{\eta}$: 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 $\mathcal{F}_{\text{mpc}}$ 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 $\mathcal{F}_{\text{mpc}}^{\epsilon/2}$ and receive `ok` from $\mathcal{F}_{\text{mpc}}$.
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[4]: 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 6 and 7), 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.

[4]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.

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 [Fei99] and the second phase uses a combinatorial object called a sampler [Vad12] 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 [Vad12][5]. 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 $f : \{0,1\}^v \to \{0,1\}$. 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[6]. 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 [RVW00, Vad12, GUV09].

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

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

[6] 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.

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 5:

- *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 [Fei99] 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.

Roadmap of our formal results. We present a formal description of the final construction to Section 5, including how to instantiate the idealized cryptography. As mentioned, we actually prove our final construction secure under a better solution concept called *sequential approximate fairness*, which we define formally in Section 3. The full proofs will be presented in Sections 6 and 7, respectively.

2.6 Additional Related Work

Leader election in other models. Leader election has also been considered in other *incomparable* computation models and these results do not directly lend to solving the problem phrased in this paper. Leader election was considered in the full information model [RSZ99, RZ98, Fei99, Dod06], culminating in the famous work of Feige who showed how to design an $O(\log^* n)$-round protocol that elects an honest leader with some *small constant* probability, assuming that the *majority* of the players are honest [Fei99]. Their notion is much weaker than the game theoretic notion considered in our work, which may be suitable in some distributed computing applications [BPV06], but not in the type of decentralized applications that supplied the original motivation of our work. Moreover, in the full-information model, leader election is impossible with a majority coalition even under their weak notion of security [Fei99]. Interestingly, our work benefitted from the techniques from this line of work: we used one iteration of Feige's lightest bin protocol [Fei99] to elect a polylogarithmically sized preliminary committee.

Alistarh et al. [AGV15] showed an $O(\log^* n)$-round protocol in an asynchronous message passing model tolerating minority *crash* faults. A sequence of works [AA11, GW12, AAG$^+$10] showed how to construct a sub-logarithmic-round protocol that tolerates *crash* faults in an asynchronous shared memory model. Various works [GHS83, KKM85] have also considered distributed algorithms for leader election where processes communicate over some graph structure but these works did not model faulty behavior and tolerated adversaries that can arbitrarily deviate from the protocol. As mentioned earlier, leader election has also been studied in a full information model [Fei99, Zuc96, RZ98, Dod06] achieving a much weaker notion which we call *resilience*, under honest majority.

Abraham et al. [ADH19] studied an incomparable game theoretic notion for leader election: in their formulation, all users prefer to have a leader than not having a leader, and users may have different preferences regarding who the leader is. Their work considered the *ex post* Nash equilibrium notion which concerns only about a single player deviating but not coalitions that jointly deviate.

Game theory meets cryptography. We review some related work at the intersection of game theory and cryptography. We recap some literature review from the recent work [CGL$^+$18]. Historically, game theory [Nas51, J.A74] and multi-party computation [GMW87, Yao82] have been investigated by separate communities. Some recent efforts have investigated the marriage of game theory and cryptography (see the excellent surveys by Katz [Kat08] and by Dodis and Rabin [DR07]). This line of work has focused on two broad types of questions.

First, a line of works [HT04, KN08, ADGH06, OPRV09, AL11, ACH11] investigated how to define game-theoretic notions of security (as opposed to cryptography-style security notions) for multi-party computation tasks such as secret sharing and secure function evaluation. Existing works have considered *a fundamentally different notion of utility* than us: specifically, these works have made (a subset to all of) the following assumptions about players' utility: players prefer to compute the function correctly; furthermore, they prefer to learn secrets, and prefer that other players do not learn secrets. These works then investigated how to design protocols such that rational players will be incentivized to follow the honest protocol. Inspired by this line of work, Garay et al. propose a new paradigm called Rational Protocol Design (RPD) [GKM$^+$13], and this paradigm was developed further in several subsequent works [GKTZ15, GTZ15]. We show that the *non-sequential* approximate fairness notions are equivalent to some approximate RPD-inspired interpretation in Appendix A.3.

Second, a line of work has asked how cryptography can help traditional game theory. Particularly, many classical works in game theory [Nas51, J.A74] assumed the existence of a trusted mediator — and recent works considered how to realize this trusted mediator using cryptography [DHR00, IML05, GK12, BGKO12].

Gradwohl et al. [GLR10] also proposed a notion of sequential rationality for cryptographic protocols. Their notion is incomparable and cannot be applied to leader election. They defined a notion of sequential equilibrium over extensive games in which computationally bounded players take turns making moves. Their notion captures a notion of games just like in the game theory literature with the additional twist that players now may be computationally bounded. Their definitions are not applicable to protocols where players can make moves in the same round. In fact, their paper hints that defining a sequential notion of equilibrium for a protocol-like setting is an open problem. A few recent works [PS17, CGL$^+$18], also inspired by blockchain applications, considered a similar notion of "cooperative-strategy-proofness" in multi-party protocols such as consensus [PS17] and coin toss [CGL$^+$18]. Our notion of maximin fairness is also inspired by Chung et al. [CGL$^+$18]. These earlier works did not consider sequential rationality in their formulation; moreover, their upper bound results do not directly lend to solving the leader election problem.

Approximate strong fairness. As mentioned in Section 1, the *de facto* notion of fairness in the multi-party computation literature is unbiasability or strong fairness. Approximate variants of unbiasabilty has also been investigated. Cleve proved that in the presence of a coalition consisting of half or more players, it is not possible to achieve $\Omega(\frac{1}{r})$-unbiasable coin toss. For the two-player setting, Moran et al. [MNS16] showed how to obtain an r-protocol that achieves $\Omega(\frac{1}{r})$-unbiasable, matching Cleve's lower bound. For the multi-player setting, our understanding of ϵ-unbiasability

is relatively little. Some very recent works [AO16, BOO10, HT14] made encouraging progress on this front; however, the number of parties needs to very tiny (e.g., constant) for these protocols to have polynomial round complexity. We were not able to rely on multi-party, ϵ-unbiasable coin toss to get our game theoretically fair leader election result, for two reasons: 1) our trade-off curve between round complexity and the fairness slack ϵ is exponentially better than ϵ-unbiasability; and 2) as mentioned, ϵ-unbiasability for the multi-party case is still somewhat poorly understood.

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 3.1); or conditioned on deviating, A does not do noticeably better (condition 2 in Definition 3.1).

Definition 3.1 (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. $\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : W^A \mid \mathsf{Dev}^{A(S)}\right] \leq \frac{1}{1-\epsilon} \cdot \Pr\left[\mathsf{Exec}^{A(\Pi)}(1^\kappa) : W^A \mid \mathsf{Dev}^{A(S)}\right] + \mathsf{negl}(\kappa)$.

In the above, the left-hand-side $\Pr\left[\mathsf{Exec}^{A(S)}(1^\kappa) : W^A \mid \mathsf{Dev}^{A(S)}\right]$ 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 $\Pr\left[\mathsf{Exec}^{A(\Pi)}(1^\kappa) : W^A \mid \mathsf{Dev}^{A(S)}\right]$ 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[7]. Intuitively,

[7]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.*

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 3.2 (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:

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

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

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

Fact 3.3 (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 Section A.4. □

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 3.4. *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 Section A.4. □

3.2 Fairness of the Tournament Tree Protocol

Instantiated with a suitable cryptographic commitment protocol (described in Section 4.3), 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 3.5 (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 Section 4.3, 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 Section B. □

4 Preliminaries

4.1 Averaging Sampler

Our protocol uses the following combinatorial construction.

Definition 4.1 (Averaging sampler)**.** Let $\mathsf{Samp} : \{0,1\}^u \to \{\{0,1\}^v\}^d$ be a *deterministic* algorithm which on input a u-bit string, outputs a sequence of d sample points from $\{0,1\}^v$. We say that Samp is an $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-averaging sampler iff for any function $f : \{0,1\}^v \to \{0,1\}$, the following holds:

$$\Pr_{x \xleftarrow{\$} \{0,1\}^u} \left[\left| \frac{1}{d} \sum_{i=1}^{d} f(\mathsf{Samp}_i(x)) - \mathbf{E}f \right| \geq \epsilon_{\mathrm{s}} \right] \leq \delta_{\mathrm{s}},$$

where $\mathsf{Samp}_i(x) \in \{0,1\}^v$ denotes the i-th output of the sampler Samp, and $\mathbf{E}f$ denotes the expectation of $f(z)$ if z is sampled at random from $\{0,1\}^v$. We say that Samp is *explicit* if given $x \in \{0,1\}^u$ and $i \in [d]$, $\mathsf{Samp}_i(x)$ can be computed in time $\mathsf{poly}(u, \log d)$.

It is known that one could construct averaging samplers from seeded extractors [Vad12]. Therefore, using known seeded extractor constructions [RVW00, GUV09], we have the following theorem:

Theorem 4.2 (Explicit averaging sampler [RVW00, Vad12, GUV09])**.** *Let $c > 1$ and $\widetilde{c} > 1$ be suitable universal constants. For any $v, \delta_{\mathrm{s}}, \epsilon_{\mathrm{s}} > 0$, as long as the input length $u \geq \log \frac{1}{\delta_{\mathrm{s}}} + c \cdot v$, and number of sample points $d \geq \left(\frac{u}{\epsilon_{\mathrm{s}}}\right)^{\widetilde{c}}$, there is an explicit $(\epsilon_{\mathrm{s}}, \delta_{\mathrm{s}})$-averaging sampler $\mathsf{Samp} : \{0,1\}^u \to \{\{0,1\}^v\}^d$.*

4.2 Constant-Round, Bounded Concurrent Secure Computation

Jumping ahead, we will first describe our protocol using idealized cryptography. When replacing the idealized cryptography with real-life instantiations, we will need to use a special type of constant-round zero-knowledge proofs that provide concurrent zero-knowledge and simulation soundness as long as there is an a-priori bound on the number of concurrent invocations [Pas04]. We use a paradigm suggested by Pass [Pas04], that is, first, design the protocol in the **IdealZK**-hybrid world and prove it secure, and then, replace the **IdealZK** with the special zero-knowledge proofs they construct. The resulting real-world protocol asymptotically preserves the round complexity, and moreover can be proven to securely emulate the original **IdealZK**-hybrid protocol. We provide the necessary preliminaries below.

IdealZK functionality. The $\mathbf{IdealZK}^{\mathcal{R}}_{i,j}$ functionality, parametrized with the prover's identity $i \in [n]$, and the verifier's identity $j \in [n]$, and the relations $\mathcal{R}$ corresponding to an NP language, is defined as follows:

Upon receiving (x, w) from the prover i, send $(x, \mathcal{R}(x, w))$ to the verifier j and the adversary.

When the NP relation $\mathcal{R}$ is obvious from the context, we often write $\mathbf{IdealZK}_{i,j}$ for simplicity.

Bounded concurrent secure computation. In an n-party **IdealZK**-hybrid protocol, the players can invoke $\mathbf{IdealZK}^{\mathcal{R}}_{i,j}$ between any prover $i \in [n]$ and any verifier $j \in [n]$, and for arbitrary NP relations $\mathcal{R}$. Without loss of generality, in every round, there can be at most n^2 concurrent invocations of **IdealZK**. Given an **IdealZK**-hybrid-world, n-player protocol, we can instantiate the **IdealZK** with actual cryptography using the elegant techniques suggested by Pass [Pas04].

Theorem 4.3 (Constant-round, bounded concurrent secure computation [Pas04])**.** *Assume the existence of enhanced trapdoor permutations, and collision-resistant hash functions. Then, given an* **IdealZK***-hybrid n-player protocol* Π_{hyb}*, there exists a real-world protocol* Π_{real} *(instantiated with actual cryptography) such that the following hold:*

- **Simulatability***: For every real-world non-uniform p.p.t. adversary* $\mathcal{A}$ *controlling an arbitrary subset of up to* $n-1$ *players in* Π_{real}*, there exists a non-uniform probabilistic expected polynomial-time adversary* $\widetilde{\mathcal{A}}$ *in the protocol* Π_{hyb}*, such that for any input* $(x_1, \ldots, x_n)$*, every auxiliary string* $z \in \{0,1\}^*$*,*

$$\mathsf{Exec}^{\Pi_{\text{real}},\mathcal{A}}(1^\kappa, x_1, \ldots, x_n, z) \equiv_c \mathsf{Exec}^{\Pi_{\text{hyb}},\widetilde{\mathcal{A}}}(1^\kappa, x_1, \ldots, x_n, z).$$

  In the above, the notation $\mathsf{Exec}^{\Pi_{\text{real}},\mathcal{A}}$ *(or* $\mathsf{Exec}^{\Pi_{\text{hyb}},\widetilde{\mathcal{A}}}$*) outputs each honest players' outputs as well as the corrupt players' (arbitrary) outputs, and* $\equiv_c$ *means computational indistinguishability.*

- **Round efficiency***: The round complexity of* Π_{real} *is at most a constant factor worse than that of* Π_{hyb}*.*

4.3 Publicly Verifiable Concurrent Non-Malleable Commitment

A publicly verifiable commitment scheme $(\mathsf{C}, \mathsf{R}, \mathsf{V})$ consists of a pair of interacting Turing machines called the committer C and the receiver R respectively, and a deterministic, polynomial-time public audit function denoted V. Suppose the commitment protocol completes successfully and produces some transcript Γ (which includes an ordered sequence of all bits transmitted between C and R), then $\mathsf{V}(\Gamma)$ outputs either a bit $b \in \{0,1\}$ to accept or $\perp$ which indicates rejection. If a bit $b \in \{0,1\}$ is output (and not $\perp$), we call it the *accepting bit*. Henceforth we assume that the protocol has two phases, a commitment phase and an opening phase, and that all algorithms receive a security parameter κ as input. Henceforth we often use the notation $\langle \mathsf{C}^*(z), \mathsf{R}^*(z') \rangle$ to indicate a (possibly randomized) execution between C^* that is invoked with the input z and R^* that is invoked with the input z'.

Perfectly correct. We require a strong notion of correctness, that is, for either $b \in \{0,1\}$, and for any $\kappa \in \mathbb{N}$, if C is honest and receives the input bit b, then for any (possibly unbounded) non-aborting R^*, with probability 1, the execution $\langle \mathsf{C}(1^\kappa, b), \mathsf{R}^*(1^\kappa) \rangle$ will successfully complete with the accepting bit b. Note that here we assume that if the malicious receiver R^* sends malformed messages outside the appropriate range, it is treated the same way as aborting. This notion of correctness implies that an honest sender can always complete the protocol correctly opening its bit, and an arbitrarily malicious (non-aborting) receiver cannot cause it to be stuck.

Perfectly binding. We can denote the transcript as $\Gamma := (\Gamma_0, \Gamma_1) \in \{0,1\}^{\ell(\kappa)}$ where the former term denotes the transcript by the end of the commitment phase and the latter term denotes the transcript of the opening phase, and $\ell(\cdot)$ is a fixed polynomial function in κ that denotes the maximum length of the transcript in each phase. We require that for any $\kappa \in \mathbb{N}$, any $\Gamma_0, \Gamma_1, \Gamma_1' \in \{0,1\}^{\ell(\kappa)}$, if $\mathsf{V}(1^\kappa, \Gamma_0, \Gamma_1) = b$ and $\mathsf{V}(1^\kappa, \Gamma_0, \Gamma_1') = b'$ where $b, b' \in \{0,1\}$, then it must be that $b = b'$. In other words, the transcript by the end of the commitment phase determines at most one bit that can be successfully opened (even when both the committer and receiver are corrupt and unbounded).

Computationally hiding. Henceforth let $p(1^\kappa, v)$ denote the probability that R^* outputs 1 at the end of the commitment phase in the execution $\langle \mathsf{C}(1^\kappa, v), \mathsf{R}^* \rangle$ where C runs the honest committer. We say that a commitment scheme is computationally hiding, iff for every non-uniform p.p.t. R^*, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\kappa \in \mathbb{N}$, for every $v_1, v_2 \in \{0,1\}^\kappa$, it must hold that $|p(1^\kappa, v_1) - p(1^\kappa, v_2)| \leq \mathsf{negl}(\kappa)$.

Concurrent non-malleability. We use the definition of concurrent non-malleability by Lin et al. [LPV08] — note that this notion implies computationally hiding. To define concurrent non-malleability, we will consider a man-in-the-middle adversary $\mathcal{A}$ that participates in m left interactions and m right interactions: on the left it interacts with an honest committer who runs the commitment phase of the protocol and commits to values $v_1, \ldots, v_m$ using identities $\mathsf{id}_1, \ldots, \mathsf{id}_m$; on the right $\mathcal{A}$ interacts with an honest receiver attempting to commit to a sequence of values $v'_1, \ldots, v'_m$, again using identities of its choice $\mathsf{id}'_1, \ldots, \mathsf{id}'_m$. If any of the right commitments are invalid its value is set to $\perp$. For any $i \in [m]$, if $\mathsf{id}'_i = \mathsf{id}_j$ for some $j \in [m]$, v'_i is set to $\perp$. Now, let $\mathsf{mitm}^{\mathcal{A}}(1^\kappa, v_1, \ldots, v_m, z)$ denote a random variable that describes the values $v'_1, \ldots, v'_m$ and the view of $\mathcal{A}$ in the above experiment — note that $v'_1, \ldots, v'_m$ are well-defined if the commitment is perfectly binding.

Definition 4.4 (Concurrent non-malleability)**.** A commitment scheme is said to be concurrent non-malleable (w.r.t. commitment) if for every polynomial $p(\cdot)$ and every non-uniform p.p.t. adversary $\mathcal{A}$ that participates in at most $m = p(\kappa)$ concurrent executions, there exists a polynomial-time simulator $\mathcal{S}$ such that the following ensembles are computationally indistinguishable:

$$\left\{\mathsf{mitm}^{\mathcal{A}}(1^\kappa, v_1, \ldots, v_m, z)\right\}_{v_1,\ldots,v_m \in \{0,1\}, \kappa \in \mathbb{N}, z \in \{0,1\}^*} \quad \text{and}$$
$$\left\{\mathcal{S}(1^\kappa, z)\right\}_{v_1,\ldots,v_m \in \{0,1\}, \kappa \in \mathbb{N}, z \in \{0,1\}^*}$$

Theorem 4.5 (Publicly verifiable concurrent non-malleable commitment [LP15])**.** *Assume that one-way permutations exist. Then there exists a constant-round, publicly verifiable commitment scheme that is perfectly correct, perfectly binding, and concurrent non-malleable.*

Proof. Lin and Pass [LP15] construct a concurrent non-malleable commitment scheme starting from any non-interactive or 2-round commitment scheme. If we instantiate this commitment with the perfectly-correct and perfectly-binding non-interactive commitment scheme of Blum [Blu83] based on one-way permutations, then the resulting protocol will inherit these properties. □

5 Formal Description of Our Scheme

5.1 Description of Our Scheme Assuming Idealized Cryptography

Our scheme makes use of an (ϵ_s, δ_s)-averaging sampler which we define in Section 4.1. We will first describe our scheme assuming idealized commitments $\mathcal{F}_{\text{comm}}$ and an ideal MPC functionality $\mathcal{F}_{\text{mpc}}$ described earlier in Section 2.4. Later in Section 5.2, we will instantiate the ideal cryptographic primitives with actual cryptography. In the scheme below, committing to a value is performed by sending it to $\mathcal{F}_{\text{comm}}$, and opening is performed by instructing $\mathcal{F}_{\text{comm}}$ 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 of Theorem 4.2.
- 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 condition required by Theorem 4.2), simply abort with the exception $\mathsf{param\_error}$ and output player 1 as the leader.

   (a) Every player sends `share` to $\mathcal{F}_{\mathrm{mpc}}^{\eta}$, and receives `ok` from $\mathcal{F}_{\mathrm{mpc}}^{\eta}$.

   (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 $\mathcal{F}_{\mathrm{mpc}}^{\eta}$, and receive a mask vector z from $\mathcal{F}_{\mathrm{mpc}}^{\eta}$.

   (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 $\mathcal{F}_{\mathrm{comm}}$ and $\mathcal{F}_{\mathrm{mpc}}^{\eta}$ are instantiated with actual cryptography later in Section 5.2, 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.

5.2 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 Section 4.2. 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 [Pas04] 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 Section 4.3. 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 [Pas04] 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[8]. 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 mes-

[8]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.

sage $\{\overline{\mathsf{coins}}_{i,k}\}_{k\in[n]}$ 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.

1. For every player $j \in S$, if the current player i posted $(\mathtt{ok}, j)$ 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 5.1 (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 6 and 7. □

6 Proofs for the Ideal-World Protocol

6.1 Bounding the Preliminary Committee's Size

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

Lemma 6.1 (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 6.2. *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 6.1, 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. □

6.2 Terminology and Notations

We first present proofs for our protocol in Section 5 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 Section B, since we will use the tournament-tree protocol's fairness properties as a blackbox in our proofs. Later in Section 7, 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.

6.3 Composition of the Final Committee

Lemma 6.3 (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*[9] $\epsilon_0 = \epsilon_s \cdot \frac{2^v}{|\mathcal{N}|}$*. Then, conditioned on no* 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* $[1-\epsilon_0, 1+\epsilon_0] \cdot d \cdot \frac{|\mathcal{N}|}{2^v}$*.*

Alternatively, suppose there is some upper bound $|\mathcal{N}| \leq N$*, and we set* $\epsilon_0 = \epsilon_s \cdot \frac{2^v}{N}$*. 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 Theorem 4.2, 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 $\epsilon_s = \epsilon_0 \cdot \frac{|\mathcal{N}|}{2^v}$, by the property of the (ϵ_s, δ_s)-averaging sampler, except for at most $2^{|\mathcal{U}|} \cdot \delta_s = 2^{0.5\psi|\mathcal{U}|}$ 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 6.3 immediately implies the following bound on the final committee size.

Lemma 6.4 (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* $\epsilon_0 = \epsilon_s \cdot \frac{2^v}{n} = 2^{-5.5r}$*. 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* 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 6.1, 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 6.2, param_error does not happen except with $\exp(-\Omega(2^{7r}))$ probability. Conditioned on these bad events not happening, we now use Lemma 6.3. 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:

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

□

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.

[9] 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.

- Event $\mathsf{param\_error}$. Recall that this happens when the preliminary comittee selected does not have the desirable properties; by Lemma 6.1 and Fact 6.2, this bad event happens with probability at most $\exp(-\Omega(2^{7r}))$.
- Event bad_1: out of the d samples from the (ϵ_s, δ_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 6.4. Assuming the honest fraction $\psi \geq 2\eta$, by Lemma 6.4, $\Pr[\mathsf{bad}_1] \leq \exp(-\Omega(2^{6r}))$. 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 6.4 implies that $\Pr[\mathsf{bad}_2] \leq \exp(-\Omega(2^{6r}))$. Observe that bad_2 depends on x, y, and z.

Lemma 6.5 (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 6.1 and Fact 6.2. 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 $\mathbf{E}[M_i\cdot\overline{\mathsf{bad}_1}\cdot\mathsf{bad}_2] \leq d\Pr[\mathsf{bad}_2] \leq d\cdot\Pr[\mathsf{bad}_2] \leq d\cdot\exp(-\Omega(2^{6r})) \leq \frac{d}{2^v}\cdot\exp(-\Omega(2^{5r}))$ 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 $\mathbf{E}[M_i\cdot\overline{\mathsf{bad}_1}\cdot\overline{\mathsf{bad}_2}] \geq \frac{d}{2^v}\left(1-2^{-0.485r}\right)\cdot\left(1-\exp(-\Omega(2^{6r}))\right) -$

$\frac{d}{2^v} \cdot \exp(-\Omega(2^{5r})) \geq \frac{d}{2^v}(1 - 2^{-0.483r})$. Finally, we have

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

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

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

Lemma 6.6 (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 $\exp(-\Theta(\alpha^2 g(1 - \frac{n}{V})))$, 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,

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

Hence, except with probability $\exp(-\Theta(\alpha^2 g(1 - \frac{n}{V})))$, 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 6.7 (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*[10] *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, $y_{[n]\setminus G}$, and z. Setting $\alpha := \frac{1}{2^r}$ in Lemma 6.6, 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

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

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 $\epsilon_0 := \epsilon_s \cdot \frac{2^v}{|Z|} \leq 2 \cdot 2^{-6r} \cdot 2^{1.5r}$, and using Lemma 6.3, we can show that except with probability $\exp(-\Omega(2^r))$, the number of players from G in the final committee is at least $(1-\epsilon_0) \cdot d \cdot \frac{Z}{2^v} \geq g \cdot \frac{d}{2^v} \cdot (1 - 2^{-0.48r})$. □

6.4 Maximin Fairness

In this section, we will prove the following lemma.

Lemma 6.8 (Ideal-world protocol: maximin fairness)**.** *The ideal-world protocol (i.e., instantiated with $\mathcal{F}_{\text{comm}}$ and $\mathcal{F}_{\text{mpc}}$) satisfies $(1-2^{-0.4r}) = (1-2^{-\Theta(r)})$-sequential-maximin-fairness against any non-uniform p.p.t. coalition[11] of size at most $(1-2\eta)n = (1-2^{-\Theta(r)})n$.*

Proof. Due to Lemma A.1, 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 3.5. This means each honest player can only be hurt negligibly more.

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 3.5), 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

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

it may increase the final committee size (in comparison with the case when A continues to play honestly). Due to Lemma 6.7, 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 6.4, 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 6.5 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 Υ_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 Υ_i is defined to be 0. (see Lemma 6.5).

  Therefore, deviating in round $\widetilde{r}$ will not reduce any honest individual's conditional winning probability by a $1-2^{-\Theta(r)}$ 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 5.1.

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 6.5, 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:

$$\mathbf{E}\left[tr \leftarrow \mathsf{LEIdeal} : \Upsilon_i | \mathsf{Dev}^{r^*}(tr)\right] \geq \frac{1}{n}\cdot\left(1-2^{-0.48r}\right). \tag{2}$$

   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 $\mathsf{poly}(n)\cdot\exp(-\Omega(2^{6r}))$, which is is still negligible, because we assume that $r = \Omega(\log\log n)$ is sufficiently large. Indeed, for sufficiently large n, $\mathsf{poly}(n)\cdot\exp(-\Omega(2^{6r})) \leq \exp(-\Omega(2^{5.99r}))$, and the proof works as before.

2. We next consider the proof of Lemma 6.7, 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 6.7 states that $\Pr[\mathsf{bad}_3] \leq \exp(-\Omega(2^r))$. Since Dev^{r^*} has non-negligible probability, we have $\Pr[\mathsf{bad}_3|\mathsf{Dev}^{r^*}] \leq \mathsf{poly}(n)\cdot\exp(-\Omega(2^r)) \leq \exp(-\Omega(2^{0.99r})) \leq \mathsf{negl}(\kappa)$, where the last inequalities hold for large enough $n \geq \kappa$ because $r \geq \Omega(\log\log n)$.

This implies that an honest continuation of the execution would lead to a conditional expectation of Υ_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 $(1 - 2\eta)n = (1 - 2^{-\Theta(r)})n$ in size. □

6.5 CSP Fairness for a Large Coalition

In this section, we prove the following lemma.

Lemma 6.9 (Ideal-world protocol: CSP fairness for a large coalition)**.** *The ideal-world protocol (i.e., instantiated with $\mathcal{F}_{\text{comm}}$ and $\mathcal{F}_{\text{mpc}}$) satisfies $(1 - 2^{-\Theta(r)})$-sequential-CSP-fairness against a non-uniform p.p.t. coalition of size at least ηn.*

Proof. We divide into cases by the first round of deviation r^* as before.

Easy case: $r^* > \widetilde{r}$. This means the coalition A will deviate only in the tournament tree protocol, whose sequential CSP fairness holds according to Theorem 3.5. This means that the coalition can gain only negligibly more.

Easy case: $r^* = \widetilde{r}$. We consider two cases. First, if $|A| < (1 - 2\eta)n$, then, in the round $\widetilde{r}$, no matter what A does, reconstruction of z is guaranteed and the reconstructed value is unique.

Second, if $|A| \geq (1 - 2\eta)n$, then the mask z the sharing-phase transcript binds to is revealed to A; moreover, A is allowed to open the mask to an arbitrary value during the round $\widetilde{r}$. Recall that we condition on the event Dev^{r^*}, which happens with non-negligible probability.

Hence, except with negligible (conditional) probability, had A continued to play honestly, the number of corrupt players in the final committee is at least $|A| \cdot \frac{d}{2^v} \cdot (1 - 2^{-0.48r})$ by the argument in Lemma 6.7, and the final committee's size is at most $n \cdot \frac{d}{2^v} \cdot (1 + 2^{-0.48r})$ by Lemma 6.4.

Therefore, except with negligible probability, A can improve its representation on the final committee by at most a $1/(1 - 2^{-\Theta(r)})$ multiplicative factor, because the conditional probability of the adversary winning is at most 1 even when it deviates.

Slightly more complicated case $r^* = \widetilde{r} - 1$. This is the round in which the players' unmasked v-ids are opened. We consider the following cases:

- Suppose that $|A| \geq (1 - 2\eta)n$. The same argument for the case $|A| \geq (1 - 2\eta)n$ and $r^* = \widetilde{r}$ applies here, too.
- Suppose that $|A| \in [\eta n, (1 - 2\eta)n)$. In this case, A already knows the z that will be opened later. A knows all players' unmasked v-ids too at the beginning of this round. A's best strategy is the internal-collision-avoidance strategy described earlier in the proof of Lemma 6.8.

  Due to Lemma 6.7, except with negligible probability, had A continued to play honestly, there are at least $|A| \cdot \frac{d}{2^v}(1 - 2^{-0.48r})$ coalition members in the final committee.

  Next, we condition on the non-negligible event Dev^{r^*} and analyze the gain by A if it actually deviates at round r^*. Observe that this decision could have been made before the input x to

the sampler is chosen, and the adversary can have at most $|A|$ distinct final v-ids after avoiding internal collision. Hence, we can use the argument in Lemma 6.3 with $\epsilon_0 = \epsilon_s \cdot \frac{2^v}{|A|} \leq 2^{-6r} \cdot 2^{0.3r}$ to conclude that, except with negligible (conditional) probability, there are at most $|A| \cdot \frac{d}{2^v}(1+\epsilon_0)$ coalition members in the final committee.

Therefore, except with negligible probability (conditioning on Dev^{r^*}), A can increase its representation in the final committee by a multiplicative factor of at most $\frac{1+\epsilon_0}{1-2^{-0.48r}} \leq \frac{1}{1-2^{-\Theta(r)}}$.

Remaining case $r^* < \widetilde{r} - 1$.

- First, consider the easy case when $|A| \geq (1-2\eta)n$. Recall that we condition on Dev^{r^*}, which has non-negligible probability.

  By the argument of Lemma 6.7, except with $\exp(-\Omega(2^r)) \cdot \Pr[\mathsf{Dev}^{r^*}]^{-1} = \mathsf{negl}(\kappa)$ probability (conditioning on Dev^{r^*}), if the coalition A continues to behave honestly, the number of players from A that are in the final committee is at least $|A| \cdot \frac{d}{2^v} \cdot (1 - 2^{-\Theta(r)})$.

  Moreover, by Lemma 6.4, except with $\exp(-\Omega(2^r)) \cdot \Pr[\mathsf{Dev}^{r^*}]^{-1} = \mathsf{negl}(\kappa)$ probability (conditioning on Dev^{r^*}), the final committee size is at most $n \cdot \frac{d}{2^v} \cdot (1 + 2^{-\Theta(r)})$. This means that conditioned on Dev^{r^*}, except with negligible probability, had A continued to behave honestly, the fraction of corrupt players in the final committee is at least $1 - 2^{-\Theta(r)}$ which is very close to 1 already. Hence, conditioning on Dev^{r^*}, the coalition's winning probability cannot increase by a multiplicative factor of $1/(1 - 2^{-\Theta(r)})$.

- Henceforth, it suffices to consider the case when $|A| < (1-2\eta)n$. We next use the sequential maximin fairness result in Lemma 6.8. From (2), we have for each $i \notin A$, no matter what p.p.t. strategy A adopts, the following holds:

$$\mathbf{E}\left[tr \leftarrow \mathsf{LEIdeal} : \Upsilon_i(tr) | \mathsf{Dev}^{r^*}(tr)\right] \geq \frac{1}{n} \cdot \left(1 - 2^{-Cr}\right), \tag{3}$$

  where $C = 0.48$. Let $h := n - |A|$, and we have,

$$\mathbf{E}\left[tr \leftarrow \mathsf{LEIdeal} : W^A(tr) | \mathsf{Dev}^{r^*}(tr)\right] \leq 1 - \frac{h}{n} \cdot (1 - 2^{-Cr}) = 1 - \frac{h}{n} + \frac{h}{n} \cdot 2^{-Cr} \tag{4}$$

  If $1 - h/n \geq \frac{1}{2^{Cr/2}} \geq \eta$, then the above expression can be upper bounded by $1 - \frac{h}{n} + \frac{h}{n} \cdot 2^{-Cr} \leq 1 - \frac{h}{n} + (1 - \frac{h}{n}) \cdot 2^{-Cr/2} \leq (1 - \frac{h}{n}) \cdot \frac{1}{1-2^{-Cr/2}}$.

  We next condition on the event Dev^{r^*} (which has non-negligible probability) and analyze what happens had A acted honestly. Except with negligible (conditional) probability, 1) there would be at least $|A| \cdot \frac{d}{2^v} \cdot (1 - 2^{-0.48r})$ coalition members in the final committee (by Lemma 6.7); and 2) the final committee's size is at most $n \cdot \frac{d}{2^v}(1 + 2^{-\Theta(r)})$ (by Lemma 6.4). Therefore, conditioning on Dev^{r^*}, had A continued to behave honestly, its probability of winning is at least $\frac{|A|}{n}(1 - 2^{-\Theta(r)})$.

□

6.6 CSP Fairness for a Small Coalition

In this section, we prove the following lemma.

Lemma 6.10 (Ideal-world protocol: CSP fairness for a small coalition)**.** *The ideal-world protocol (i.e., instantiated with $\mathcal{F}_{\text{comm}}$ and $\mathcal{F}_{\text{mpc}}$) satisfies $(1 - 2^{-\Theta(r)})$-sequential-CSP-fairness against a non-uniform p.p.t. coalition of size less than ηn.*

Proof. We divide into cases by the first round of deviation r^* as before.

Easy case: $r^* > \widetilde{r}$**.** Same as Section 6.5.

Easy case: $r^* = \widetilde{r}$**.** For a coalition of size less than ηn, what it does has no effect on the reconstruction of z.

Slightly more complicated case: $r^* \leq \widetilde{r} - 1$**.** A small coalition of size less than ηn has no information about z at this point; further, the z that the sharing-phase transcript binds to is guaranteed to be reconstructed later. In this case, no matter what x is, since Dev^{r^*} is independent of z, suppose that exactly $\alpha \leq |A|$ number of coalition members open their unmasked v-ids, then, conditioning on the event that the final committee is non-empty, by the randomness of z and the symmetry between players, A's expected fraction in the final committee is exactly $\alpha/(\alpha+h)$, where $h = n - |A|$.

If the default winner is not in A when the final committee is empty, obviously A does not have incentive to withhold their unmasked v-ids. However, even if the default winner is in A when the final committee is empty, the adversary can take advantage of this only if no honest player is in the final committee; because of independence of z among the players, the adversary A will not be hurt if all players in A open their y_A.

It follows that conditioning on any trace up to round $\widetilde{r}$, the adversary A should open its y_A.

However, the adversary A may deviate before round $\widetilde{r}$ to influence the probability that the final committee is empty. Nevertheless, Lemma 6.7 implies that no matter what the adversary does, the probability that the final committee is empty is less than $\exp(-\Omega(2^r))$. Hence, conditioning on any non-negligible Dev^{r^*}, the probability that the final committee is empty is still negligible.

□

7 Extending the Fairness Guarantees to the Real-World Protocol

7.1 Analysis of the IdealZK-Hybrid Protocol Π_{mpc}

Theorem 7.1 (Π_{mpc} protocol of Section 5.2)**.** *Suppose that the commitment scheme* comm *is perfectly binding and computationally hiding. Then, the* **IdealZK***-hybrid-world protocol in Section 5.2 satisfies the following security properties:*

1. *If $|A| < (1-2\eta)n$, at the end of the sharing phase, there exists a unique string* coins*, such that during the reconstruction phase, all honest players will output* coins*.*
2. *If $|A| < \eta n$, then there exists a p.p.t. simulator* Sim*, such that*

$$(\mathsf{coins}, \mathsf{view}_A^{\text{sharing}}) \equiv_c (U, \mathsf{Sim}(1^\kappa))$$

where coins *denotes the unique coin that the transcript of the sharing phase binds to,* $\mathsf{view}_A^{\text{sharing}}$ *denotes the adversary's view in the sharing phase, U is a random string of the same length as* coins*, and $\equiv_c$ means computational indistinguishability.*

Proof. For the first claim, if $|S| < \eta n$, then the claim is trivial. Henceforth we focus on the case when $|S| \geq \eta n$. Suppose that $|A| < (1-2\eta)n$, we know that everyone in S must have obtained ok from at least $(1-\eta)n$ players, it must be that at least ηn of these ok messages come from honest players. This means that everyone $j \in S$ must have supplied a valid opening of its commitment message to **IdealZK**, and moreover the opening is a valid $\lceil \eta n \rceil$-out-of-n secret sharing which determines the coin value coins. Furthermore, all the ηn honest players who have sent ok for j's

commitment message must have received a valid opening of its own share, and will send its share's opening during the reconstruction phase. Therefore, reconstruction must succeed and reconstruct to coins.

We now prove the second claim. Let ZKHyb denote an execution of Π_{mpc} where the simulator acts on behalf of all honest players and interacts with the adversary, and moreover the simulator simulates **IdealZK** for the adversary too. We now consider a hybrid experiment Hyb which is almost identical to ZKHyb, except that we replace the way how the simulator computes an honest player i's commitment messages: for each $j \in [n] \backslash A$, let $C_{i,j} := \mathsf{comm}(\mathbf{0}, \rho_{i,j})$; for each $j \in A$, $C_{i,j} := \mathsf{comm}(\mathsf{coins}_{i,j}, \rho_{i,j})$; and the commitment message of i is $(\{C_{i,j}\}_{j\in[n]})$. The **IdealZK** instances interacting with the adversary (simulated by the challenger) vouches for this new commitment message $(\{C_{i,j}\}_{j\in[n]})$, too. Hyb is computationally indistinguishable from ZKHyb due to the computational hiding property of the commitment scheme comm.

Now, consider Hyb. Since $|A| < \eta n$, $|S| \geq \eta n$ must hold since all honest players will send $\mathtt{ok}$ for all honest players. Recall that we are using $\lceil \eta n \rceil$-out-of-n secret sharing; therefore, for every honest i, the shares $\{\mathsf{coins}_{i,j}, \rho_{i,j}\}_{j\in A}$ the adversary receives is independent of coins_i. If $|A| < \eta n$, and we know that $|S| \geq \eta n$, then there must exist an honest player $u \in S$, and coins_u is independent of $\mathsf{view}_A^{\mathrm{sharing}}$. Therefore, in Hyb, $\mathsf{view}_A^{\mathrm{sharing}}$ is independent of the coin $\mathsf{coins} := \oplus_{j\in S}\mathsf{coins}_j$ to be reconstructed later. In other words, Hyb essentially defines a simulator Sim, such that $(\mathsf{coins}, \mathsf{view}_A^{\mathrm{sharing}}) \equiv (U, \mathsf{Sim})$ where $\equiv$ means identically distributed, and U is a uniform random string of appropriate length.

□

7.2 Hybrid Experiments

Throughout, we will use LE to denote a randomized execution of our leader election protocol instantiated with the **IdealZK**-hybrid Π_{mpc} in place of $\mathcal{F}_{\mathrm{mpc}}$, and with non-malleable commitments in place of the idealized commitments. In our proof, we will make use of a couple hybrid experiments denoted LEHyb_y and LEHyb_z which we define below. Later on, when our proof needs to use the independence of y_H w.r.t. other coins, we will use LEHyb_y as stepping stone. When the coalition is smaller than ηn and we need to use the independence of z w.r.t. other coins, we will use LEHyb_z as stepping stone.

Hybrid LEHyb_y. Hybrid experiment LEHyb_y essentially runs LE but replaces the NMC step with the NMC's simulator.

Hybrid experiment LEHyb_y

1. Elect the preliminary committee $\mathcal{U}$ using lightest bin as before.
2. Like before, let $\mathsf{Samp} : \{0,1\}^{|\mathcal{U}|} \to \{\{0,1\}^v\}^d$, and if the parameter check fails, abort with the exception $\mathsf{param\_error}$ and output player 1 as the leader.
3. Run the sharing phase of Π_{mpc} which is in the **IdealZK**-hybrid world. At the end of this step, the transcript uniquely binds to some z.
4. Each honest player $i \in H$ chooses a random $y_i \in \{0,1\}^v$, and invokes n NMC instances and run the commit phase with n receivers. Run the simulator for the NMC which outputs 1) n values each corrupt player is trying to commit; and 2) the view of the adversary view_A. If the same corrupt player $j \in A$ is committing to different values, we simply let its committed

value be $y_j := 0^v$; else, we let y_j be the value output by the simulator. Use the simulated view view_A to reset the adversary's internal state.

5. 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 (same as in LE).

Hybrid LEHyb_z. Hybrid experiment LEHyb_z essentially runs LE but replaces the Π_{mpc} with its simulator Sim (for the case when $|A| < \eta n$).

Hybrid experiment LEHyb_z

1. Elect the preliminary committee $\mathcal{U}$ using lightest bin as before.
2. Like before, let $\mathsf{Samp} : \{0,1\}^{|\mathcal{U}|} \to \{\{0,1\}^v\}^d$, and if the parameter check fails, abort with the exception $\mathsf{param\_error}$ and output player 1 as the leader.
3. Run the simulator Sim of Π_{mpc} which outputs $\mathsf{view}_A^{\mathrm{sharing}}$ and sample a random mask $z \in \{0,1\}^{vn}$. Use $\mathsf{view}_A^{\mathrm{sharing}}$ to reset the adversary's internal state.
4. Every player commits to a random unmasked v-id $y_i \in \{0,1\}^v$ by creating n instances of the NMC scheme, one for each of the n receivers.
5. 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 (same as in LE).
6. Every player i now opens its committed y_i by posting the openings of n instances to the broadcast channel. If the openings are inconsistent, treat the opened value $y_i := 0^v$; else, let y_i be the opened value (same as in LE).

Remark 7.2. Observe that in LEHyb_y and LEHyb_z, the experiments stop before the opening phase of the NMC, and of the $\mathcal{F}_{\mathrm{mpc}}$, respectively. This is because using the concurrent non-malleability definition of NMC and due to Theorem 7.1, the simulatability of the NMC and the $\mathcal{F}_{\mathrm{mpc}}$ holds only for the commitment/sharing phase. In our proofs later, we will rely on the simulatability of the commitment/sharing phase, and then, we will define polynomial-time computable variables over the view of the experiment before opening, and argue computational indistinguishability of these variables in LEHyb_y (or LEHyb_z) and in our **IdealZK**-hybrid-world leader election protocol denoted LE later.

7.3 Maximin Fairness

In this section, we prove the following theorem.

Theorem 7.3 (Real-world protocol: maximin fairness)**.** *Suppose that the* NMC *scheme satisfies perfectly binding and concurrent non-malleability as defined in Appendix 4.3, and that* Π_{mpc} *satisfies Theorem 7.1. Then, the real-world protocol satisfies* $(1 - 2^{-\Theta(r)})$*-sequential-maximin-fairness against any non-uniform p.p.t. coalition of size at most* $(1 - 2^{-\Theta(r)})n$.

Proof. Let LE be a randomized execution of our leader election protocol instantiated with the **IdealZK**-hybrid Π_{mpc} in place of $\mathcal{F}_{\mathrm{mpc}}$, and with non-malleable commitments in place of the

idealized commitments. Due to Theorem 4.3, it suffices to prove our sequential approximate fairness notions for LE, and the same guarantees would extend to the real-world protocol when the **IdealZK** is instanatiated with the bounded concurrent zero-knowledge proofs suggested by Pass [Pas04].

Like the proof of Lemma 6.8, we will divide into several cases based on the first round of deviation r^*. Let $\widetilde{r}$ be the round in which players open the mask z.

Case $r^* > \widetilde{r}$. The analysis in the proof of Lemma 6.8 still applies.

Case $r^* = \widetilde{r}$. The analysis in the proof of Lemma 6.8 still applies.

Case $r^* = \widetilde{r} - 1$. For the case when $|A| \geq \eta n$, the analysis in the proof of Lemma 6.8 still applies. Below we focus on the case when $|A| < \eta n$. Fix an arbitrary honest individual i, we want to derive an upper bound on $\mathbb{E}[\Upsilon_i]$ conditioned on Dev^{r^*}.

Observation 7.4. Observe that Lemma 6.5 actually holds when Υ_i is defined solely as a polynomial-time computable function over x, y, z and independent of additional random coins of A. In other words, the proof of Lemma 6.5 actually defines Υ_i assuming the worst case for the honest player i when x, y and z are fixed: essentially, we do not penalize the coalition for internal collision, that is, the lemma holds when all members of A open their unmasked v-ids, but if two or more members of A happen to have the same final v-id, we pretend a-posterior that only one of these coalitions members opened.

Due to this observation as well as Theorem 7.1 (the case for $|A| < \eta n$), we may equivalently lower bound on $\mathbb{E}[\Upsilon_i(x, y, z)]$ conditioned on Dev^{r^*} in LEHyb_z, and the quantity $\mathbb{E}[\Upsilon_i(x, y, z)]$ in LE should only be negligibly apart. Now, observe that in LEHyb_z, Lemma 6.5 still holds, even when conditioned on Dev^{r^*} (which is independent of z), if we simply redo the argument using the randomness of z_H in place of y_H.

Now, if A had continued to play honestly, then Υ_i is also a polynomial-time computable function over (x, y, z). In LEHyb_z, it is not hard to see that had A continued to play honestly, $\mathbb{E}[\Upsilon_i]$ is $1/n - \mathsf{negl}(\kappa)$ conditioned on Dev^{r^*}. Therefore, due to Theorem 7.1 (for the case where $|A| < \eta n$), in LE, $\mathbb{E}[\Upsilon_i]$ is at most is $1/n + \mathsf{negl}(\kappa)$ conditioned on Dev^{r^*}. More specifically, Υ_i can only take polynomially many possible values, and if in LEHyb_z and LE, their expectation is more than negligibly apart, then there must exist at least one value whose probability differs by more than a negligible amount in LEHyb_z and LE, respectively. We can then construct an efficient distinguisher which outputs 1 upon encountering this specific value, and otherwise outputs 0. Such an efficient distinguisher can distinguish LEHyb_z and LE with non-negligible probability, thus violating Theorem 7.1.

This means that conditioned on deviating in round $r^* = \widetilde{r} - 1$, A cannot reduce i's conditional winning probability by a multiplicative $1 - 2^{-\Theta(r)}$ factor.

Case $r^* < \widetilde{r} - 1$. For the case when $|A| \geq (1 - 2\eta)n$, the analysis in the proof of Lemma 6.8 still applies. Below we focus on the case $|A| < (1 - 2\eta)n$. Let $i \notin A$ be an arbitrary honest individual. In the proof of Lemma 6.8 for this case, the part that an honest continuation of the execution would lead to the stated bound on the conditional expectation of Υ_i still holds for LE. It remains to show that even in LE, conditioned on Dev^{r^*}, the expectation of Υ_i is at least $\frac{1}{n}(1 - 2^{-\Theta(r)})$. To show this, due to Observation 7.4, it suffices to prove the statement for $\Upsilon_i(x, y, z)$ in LEHyb_y, and the conditional expectation of $\Upsilon_i(x, y, z)$ in LE is only negligibly apart due to the concurrent

non-malleability of NMC. To conclude the proof, observe that the proof of Lemma 6.5 holds in LEHyb_y even when conditioning on Dev^{r^*}, since in LEHyb_y, y_H is independent of x, y_A and z. □

7.4 CSP Fairness

The following theorem states sequential approximate CSP fairness for our real-world protocol for coalitions that are not too small.

Theorem 7.5 (Real-world protocol: CSP fairness for a large coalition)**.** *Suppose that the* NMC *scheme satisfies perfectly binding and concurrent non-malleability as defined in Appendix 4.3, and that* Π_{mpc} *satisfies Theorem 7.1. The real-world protocol satisfies* $(1 - 2^{-\Theta(r)})$*-sequential-CSP-fairness against a non-uniform p.p.t. coalition of size at least* ηn*.*

Proof. Due to Theorem 4.3, it suffices to prove our sequential approximate fairness notions for LE, and the same guarantees would extend to the real-world protocol when the **IdealZK** is instanatiated with the bounded concurrent zero-knowledge proofs suggested by Pass [Pas04]. Like the proof of Lemma 6.9, we will divide into several cases based on the first round of deviation r^*. Let $\widetilde{r}$ be the round in which players open the mask z. The analysis in the proof of Lemma 6.9 for all cases where $r^* \leq \widetilde{r} - 1$ still hold for LE.

Below we focus on the case when $r^* < \widetilde{r} - 1$. For the sub-case when $|A| \geq (1 - 2\eta)n$, the analysis in Lemma 6.9 applies to LE. Therefore, we focus on the sub-case when $|A| > (1 - 2\eta)n$. Here, just like in the proof of Lemma 6.9, we can still conclude that in LE, conditioning on the non-negligible probability event Dev^{r^*}, had A continued to behave honestly, its probability of winning is at least $\frac{|A|}{n}(1 - 2^{-\Theta(r)})$. We can also show that conditioned on Dev^{r^*}, A's winning probability is at most $\frac{|A|}{n}/(1 - 2^{-\Theta(r)})$ using a similar argument as in the proof of Lemma 6.9 since the argument essentially relies on sequential maximin fairness which we have also proven earlier for LE. □

The following theorem states sequential approximate CSP fairness for our real-world protocol for a very small coalition.

Theorem 7.6 (Real-world protocol: CSP fairness for a small coalition)**.** *Suppose that the* NMC *scheme is perfectly binding and that* Π_{mpc} *satisfies Theorem 7.1. The real-world protocol satisfies* $(1 - 2^{-\Theta(r)})$*-sequential-CSP-fairness against a non-uniform p.p.t. coalition of size less than* ηn*.*

Proof. Due to Theorem 4.3, it suffices to prove our sequential approximate fairness notions for LE, and the same guarantees would extend to the real-world protocol when the **IdealZK** is instanatiated with the bounded concurrent zero-knowledge proofs suggested by Pass [Pas04]. Like the proof of Lemma 6.10, we will divide into several cases based on the first round of deviation r^*. Let $\widetilde{r}$ be the round in which players open the mask z. The analysis in the proof of Lemma 6.10 for all cases where $r^* \geq \widetilde{r}$ still hold for LE.

For the case $r^* \leq \widetilde{r} - 1$, we can consider the hybrid LEHyb_z. Essentially, at the point when (a subset of) coalition members open their unmasked v-ids, we can use the values x, y, and z to determine the adversarial fraction on the final committee, and conditioned on Dev^{r^*}, the coalition's winning probability is the same as the expected adversarial fraction. the analysis of Lemma 6.10 holds for the Therefore, the analysis of Lemma 6.10 still holds in LEHyb_z. Now, by Theorem 7.1, the same holds for LE too. □

7.5 Proof of Our Main Theorem

Given the analysis in this section, the proof of Theorem 5.1 follows from Theorems 7.3, 7.5, and 7.6, and observing that all the cryptographic primitives we need can be instantiated from enhanced trapdoor permutations and collision resistant hashing [Pas04, LP15, CGL$^+$18].

8 Lower Bound on Round Complexity

In this section, we show that the folklore tournament tree protocol for leader election (described in Appendix B) has tight round complexity $\Theta(\log n)$ under certain conditions. Our protocol can circumvent this lower bound because it relaxes both requirements.

- *Perfect fairness.* We prove the lower bound for protocols achieving perfect fairness, i.e., $\epsilon = 0$. Observe that for the special case $\epsilon = 0$, maximin fairness is the same as CSP fairness, and moreover, the sequential and non-sequential versions are equivalent too.
- *Open immediately after commit.* The tournament tree protocol follows the paradigm in which values committed by players in one round are immediately open in the next round. In contrast, in our protocol, values committed by players in one round are opened in later rounds, and not immediately the next round.

Theorem 8.1 (Lower bound on round complexity)**.** *Under the above assumptions, a perfectly fair protocol electing a leader uniformly at random from n players must take $\Omega(\log n)$ rounds.*

Proof Intuition. Since the formal proof is technical, we explain some intuitions, which explain why we need the assumptions and also show the difficulties if one wishes to prove a lower bound under looser assumptions.

1. *Hereditary Properties for Sub-Protocols.* To rephrase our lower bound, the goal is to show that if there is a perfectly fair protocol with ℓ rounds, then at most 2^ℓ players have a non-zero chance of winning.

   The commit-open framework implies that after any initial rounds in which no player aborts, any sub-protocol reached remains perfectly fair (Lemma 8.4). Observe that the tournament tree protocol can fall under a more restrictive model: it is possible that all players commit to all messages in the first round, and the commitments are open round by round subsequently. However, by doing so, any sub-protocol reached will have a dependence on pre-committed messages, which is not the case for the overall protocol. Hence, more sophisticated techniques would be needed to make the induction proof work for other frameworks.

2. *Replacement Lemma.* The most technical part of the inductive step is the replacement lemma (Lemma 8.9). Suppose that U is a subset of players, each of which has a chance to win in the (perfectly fair) protocol φ, but somehow all get eliminated under some message configuration in the first round. Then, it is possible for players in U to change their messages (while players not in U keep their messages) in the first round of φ such that all players in U will still have a chance to win after the first round.

   Then, one can see how this lemma can help in an induction proof. Suppose, for contradiction's sake, that there is an ℓ-round perfectly fair protocol in which $2^\ell + 1$ players have non-zero chance of winning. Consider any message configurations sent in the first round. Observe that the resulting sub-protocol has $\ell - 1$ rounds, which, by the induction hypothesis, means that

at most $2^{\ell-1}$ players still have a non-zero chance of winning. Hence, there is a subset U of $2^{\ell-1}+1$ players eliminated in the first round.

The replacement lemma says that those players in U could have acted differently in the first round such that all of them would still have a non-zero chance of winning after the first round. However, the induction hypothesis says that after the first round, there could be at most $2^{\ell-1}$ players that can win with non-zero probability, thereby reaching the desired contradiction.

Observe that the commit-open framework assumption is crucial in the above argument. Otherwise, we could not have used the induction hypothesis.

8.1 Notation

In addition to the definitions in Section 3, we introduce further concepts to facilitate our lower bound proof.

Winning probability vector. Recall that the goal is that there is a set $[n]$ of players who wish to collaboratively select a leader (or winner). Our lower bound works for the more general case when each player $i \in [n]$ has some *winning probability* $p_i \in [0,1]$ such that $\sum_{i\in[n]} p_i = 1$. Indeed, it is easy to modify the tournament tree protocol in Appendix B to support rational probabilities.

Given $p \in [0,1]^n$, its support is $\sigma(p) := \{i \in [n] : p_i > 0\}$ and we denote $\|p\|_0 := |\sigma(p)|$.

Commitment scheme and adversary. Our lower bound assumes an ideal concurrent non-malleable commit scheme as described in Appendix 4.3. Hence, no assumption on the computational power of the adversary is needed. However, we do need that the adversary is non-uniform, i.e., it can depend on a given protocol.

Communication model under the commit-and-immediately-open framework. One round of the protocol is captured by a function $\varphi : \times_{i\in[n]}\Omega_i \cup \{\perp\} \to \mathcal{O}$, where Ω_i is the input set for player i and $\mathcal{O}$ is the set of outcomes. We consider the *commit-open* framework in which one round consists of two phases.

Commit Phase. Each honest player $i \in [n]$ samples an element x_i from its input set Ω_i uniformly at random (which is without loss of generality), produces a commitment c_i of x_i, and broadcasts the commitment c_i to every other player. A coalition A of dishonest players may collude and sample their inputs according to any arbitrary joint distribution on $\Omega_A := \times_{i\in A}\Omega_i$. A dishonest player may also broadcast anything arbitrarily, but we assume that whatever a player broadcasts can be observed by everyone.

Open Phase. Each honest player $i \in [n]$ opens the element x_i that it has committed in the previous phase (which can be verified under the commitment scheme). We assume that the adversary is *rushing*, i.e., after observing all honest players' revealed inputs, a set A of dishonest players can decide to reveal the inputs of which subset of A.

Outcome. For each player i, if its c_i in the commitment phase and its x_i in the open phase can be verified, then its input to the function φ is x_i, otherwise its input to φ is $\perp$. The outcome of this round is then produced by φ.

Multi-round protocols for committee selection. We use $\mathcal{P}^{(1)}$ to denote the set of one-round protocols to elect a leader[12]. Each one round protocol is captured by some function φ :

[12]In the case that all players deviate from the protocol, it might be the case that the protocol might not return any leader. To insist on exactly one leader elected, the protocol can have some default leader, say player 1.

$\times_{i\in[n]}\Omega_i \cup \{\perp\} \to [n]$ in the commit-open framework. We also use the convention $\mathcal{P}^{(0)} := [n]$, which we identify with $\{x \in \{0,1\}^n : \sum_{i\in[n]} x_i = 1\}$.

Multi-round protocols to select exactly one leader is defined recursively. Suppose for some $\ell \geq 1$, the set $\mathcal{P}^{(\ell)}$ of ℓ-round protocols is already defined. Then, each $(\ell+1)$-round protocol in $\mathcal{P}^{(\ell+1)}$ is captured by some function $\varphi : \times_{i\in[n]}\Omega_i \cup \{\perp\} \to \mathcal{P}^{(\ell)}$ in the commit-open framework, where each outcome is an ℓ-round protocol.

Additional Assumptions. For simplicity, we also assume that if a player's input to the protocol is $\perp$ in some round, then all its inputs in subsequent rounds are also $\perp$, and that player has no chance to be a winner.

Definition 8.2 (Perfect Fairness and Tight Protocol)**.** Given a winning probability vector $\vec{p} := (p_i : i \in [n])$, a (multi-round) protocol among players in $[n]$ achieves $\vec{p}$ with perfect fairness, if for each $i \in [n]$, even when all players in $[n] \setminus \{i\}$ collude, provided that player i remains honest (in every round), the probability that player i wins is at least p_i.

We say that a protocol φ attains $\vec{p}$ optimistically, if for each player i, the winning probability is p_i, provided that all players are honest; in this case, we write $\mathcal{P}(\varphi) := \vec{p}$ and $\mathcal{P}_i(\varphi) = p_i$.

We say that a protocol is tight if there is some $\vec{p}$ such that it achieves $\vec{p}$ both optimistically and with perfect fairness; in this case, we simply say that the tight protocol achieves $\vec{p}$.

Fact 8.3. *For any $\vec{p} \in [0,1]^n$ such that $\sum_{i\in[n]} p_i = 1$, any (multi-round) protocol φ electing exactly one leader that achieves $\vec{p}$ with perfect fairness must be tight.*

Lemma 8.4 (Tightness is Hereditary with No Aborts)**.** *Suppose φ is a tight multi-round protocol, and φ' is some intermediate protocol encountered after several rounds during which no player aborts. Then, φ' is also tight.*

Proof. On the contrary, assume that the protocol φ' is not fair with respect to $\mathcal{P}(\varphi')$. This means that there is some honest player i such that it wins with probability strictly less than $\mathcal{P}_i(\varphi')$ if all other players collude against it (under the non-uniform adversarial model).

Then, the original protocol φ is not tight, because all other players can collude against i as follows. All players play honestly unless their sequence of inputs during initial rounds leads to the intermediate protocol φ', at which point they collude against i. □

8.2 Technical Proof

To simplify the notation, for each player $i \in [n]$, we use Ω_i to denote its set of inputs in *every* round. To avoid running into non-measurable and other technical issues, we restrict to the case that Ω_i is finite, and an honest player i should sample from Ω_i uniformly at random (which is without loss of generality) in each round. One consequence is that every configuration of inputs over all the rounds happens with non-zero probability. We prove the following statement by induction on ℓ. Observe that Theorem 8.1 immediately follows from the special case when all p_i's are equal.

Theorem 8.5 (Lower Bound)**.** *Suppose some tight ℓ-round protocol selecting exactly one winner achieves $\vec{p} \in [0,1]^n$ and $\sum_{i\in[n]} p_i = 1$. Then, $\|\vec{p}\|_0 \leq 2^\ell$.*

When players abort. The following definition captures the expected reward of honest users when other users abort.

Definition 8.6 (Closure Vector)**.** Suppose φ is an ℓ-round protocol. The closure operator $\mathcal{P}^\varphi$ is defined on the inputs of a subset of players as follows. Suppose $T \subseteq [n]$ and denote $\overline{T} := [n] \setminus T$.

Then, for some inputs $x \in \Omega_T$ from players in T, the closure operator $\mathcal{P}^\varphi(x) \in [0,1]^T$ is defined such that for each $i \in T$, its i-th coordinate is $\mathcal{P}_i^\varphi(x) := \max_{y \in \Omega_{\overline{T}}} \mathcal{P}_i(\varphi(x,y))$.

Recall that $\mathcal{P}(\varphi(x,y))$ is the winning vector that the $(\ell-1)$-round protocol $\varphi(x,y)$ achieves optimistically.

Lemma 8.7 (Tight Protocol and Closure Vector)**.** *Suppose φ is an ℓ-round protocol that is tight, and suppose $S \subseteq [n]$. Then, for any $x \in \Omega_T$, the $(\ell-1)$-round protocol $\varphi(x,\perp)$ is perfectly fair with respect to the closure vector $\mathcal{P}^\varphi(x) \in [0,1]^S$.*

Proof. Suppose there is some $x_0 \in \Omega_T$ such that $\varphi(x_0,\perp)$ is not perfectly fair with respect to $\mathcal{P}^\varphi(x_0)$. This means that there exists some $i \in T$ and $y_0 = \arg\max_{y \in \Omega_{\overline{T}}} \mathcal{P}_i(\varphi(x_0,y))$ such that in the $(\ell-1)$-round protocol $\varphi(x_0,\perp)$, the probability that player i wins is less than $\mathcal{P}_i(\varphi(x_0,y_0))$, if all other players collude against i.

It follows that in the supposedly tight protocol φ, player i can be hurt by all other players (who form a non-uniform adversary) as follows. Suppose in the first round, all players sample their inputs according to the specified distribution. If $(x_0,y_0) \in \Omega_T \times \Omega_{\overline{T}}$ is sampled, then the (rushing) adversary instructs players in $\overline{T}$ to abort; subsequently, in the protocol $\varphi(x_0,\perp)$, all players other than i will collude to make sure that i wins with probability less than $\mathcal{P}_i(\varphi(x_0,y_0))$. This violates the tightness of φ. □

We prove Theorem 8.5 by induction on ℓ.

8.2.1 Base Case

The base case $\ell = 0$ is trivially true, because with no communication there can only be (at most) one default winner. However, as a warmup, we will also prove the case $\ell = 1$ (which the reader might skip). As we shall see, the general inductive step is more complicated.

Lemma 8.8. *For any $\vec{p} \in [0,1]^n$ such that $\sum_{i\in[n]} p_i = 1$ and $\|p\|_0 \geq 3$, there is no tight 1-round protocol that achieves $\vec{p}$.*

Proof. Assume the contrary and suppose there is some tight 1-round protocol that achieves some $\vec{p} \in [0,1]^{[n]}$ with $\|p\|_0 \geq 3$. Without loss of generality, we can assume $n = 3$, because we can merge extra players into one (which makes the adversary slightly weaker as the merged players can only abort together). Hence, we have for all $i \in [3]$, $p_i > 0$. Moreover, $\sum_{i\in[3]} p_i = 1$. Since the winner is determined after one round, we can simplify the notation and assume $\varphi(\cdot) \in [3]$.

Fix any $x \in \Omega_1$ and $y \in \Omega_2$. Because of the fairness of φ, $\Pr_{Z \leftarrow \Omega_3}[\varphi(x,y,Z) = 3] \geq p_3$; in fact, equality holds because φ is also tight. Hence, we have $\{3\} \subsetneq \varphi(x,y,\Omega_3)$, where the proper subset follows from $p_3 < 1$.

Next, we argue that $\varphi(x,y,\Omega_3)$ cannot contain both 1 and 2. Otherwise, there exists z and z' such that $\varphi(x,y,z) = 1$ and $\varphi(x,y,z') = 2$. By Lemma 8.7, the outcome $\varphi(x,y,\perp)$ is fair with respect to $\mathcal{P}^\varphi(x,y) = (1,1)$, which means both 1 and 2 are winners. Hence, we can conclude that $\varphi(x,y,\Omega_3)$ is either $\{1,3\}$ or $\{2,3\}$.

For the rest of the proof, we fix some $x_0 \in \Omega_1$ and define for $i \in \{1,2\}$, $G^{(i)} := \{y \in \Omega_2 : \varphi(x_0,y,\Omega_3) = \{i,3\}\}$. Note that $\Omega_2 = G^{(1)} \cup G^{(2)}$; moreover, $p_1, p_2 > 0$ implies that both $G^{(1)}$ and $G^{(2)}$ are non-empty. For any $y \in \Omega_2$, define $Q_y := \{z \in \Omega_3 : \varphi(x_0,y,z) = 3\}$, which is non-empty, because of fairness (to player 3, who is supposed to win with probability $p_3 > 0$). Observe that for any $y, y' \in G^{(1)}$, we must have $Q_y = Q_{y'}$. Otherwise, there exists z such that $\varphi(x_0,y,z) = 1$ and $\varphi(x_0,y',z) = 3$; by Lemma 8.7, the outcome $\varphi(x_0,\perp,z)$ is fair with respect to $\mathcal{P}^\varphi(x_0,z) = (1,1)$, i.e., we have the impossible situation that both 1 and 3 are winners. Hence, we

write $Q^{(1)} := Q_y$, for all $y \in G^{(1)}$. Recall that because φ is tight, we must have for all $y \in G^{(1)}$, $\Pr_{Z \leftarrow \Omega_3}[\varphi(x_0, y, Z) = 3] = \Pr[Z \in G^{(1)}] = p_3$.

For any $y \in G^{(2)}$ and $z \in \Omega_3 \setminus Q^{(1)}$, we observe that $\varphi(x_0, y, z) = 2$. Otherwise, we have $\varphi(x_0, y, z) = 3$ and picking any $y' \in G^{(1)}$, the choice of $z \notin Q^{(1)}$ implies that $\varphi(x_0, y', z) = 1$. Again, by Lemma 8.7, the outcome $\varphi(x_0, \perp, z)$ has the impossible situation of containing both 1 and 3 as winners. Hence, it follows that for all $y \in G^{(2)}$, $Q_y \subseteq Q^{(1)}$. However, since φ is tight, we also have $p_3 = \Pr_{Z \leftarrow \Omega_3}[Z \in Q_y] \leq \Pr[Z \in Q^{(1)}] = p_3$, where the last equality is from above. Hence, we must have $Q_y = Q^{(1)}$ for all $y \in G^{(2)}$.

This implies that for any $z \in Q^{(1)}$, for any $y \in \Omega_2 = G^{(1)} \cup \Omega_2 = G^{(2)}$, $\varphi(x_0, y, z) = 3$.

This means that players 1 and 3 can collude such that player 1 first picks any $x_0 \in \Omega_1$ and player 3 picks any $z \in Q^{(1)}$ (depending on x_0) to hurt player 2 (who is supposed to win with non-zero probability $p_2 > 0$). □

8.2.2 Inductive Step

Lemma 8.4 says that if a multi-round protocol φ is tight and no player aborts during the initial rounds, then the resulting protocol is also tight. However, in the case that some players abort, then the remaining protocol might not be tight (but will still be fair with respect to some closure vector as stated in Lemma 8.7). This is the reason why the inductive step is tricky, because if the protocol is not tight, then the induction hypothesis cannot be readily applied. As we shall see, the inductive step uses the technical result in Lemma 8.9.

Notation. Given any (multi-round) protocol φ, we use $\mathsf{S}(\varphi)$ to denote the support (set of non-zero coordinates) of the vector $\mathcal{P}(\varphi)$ achieved optimistically by φ.

Lemma 8.9 (Replacement Lemma)**.** *Suppose φ is a tight (multi-round) protocol and $\omega \in \Omega_{[n]}$ is some input such that $T = \mathsf{S}(\varphi(\omega))$. Then, for any $U \subseteq ([n] \setminus T) \cap \mathsf{S}(\varphi)$, there exists some $\omega' \in \Omega_{[n]}$ such that every player in $[n] \setminus U$ has the same input in ω and ω', and $U \subseteq \mathsf{S}(\varphi(\omega'))$.*

In other words, suppose U is a subset of players, each of which has a chance to win in the protocol φ, but somehow all get eliminated under some input configuration $\omega \in \Omega_{[n]}$ in the first round. Then, it is possible for players in U to change their inputs (while players not in U keep their inputs as in ω) in the first round of φ such that all players in U will still have a chance to win after the first round.

Proof. For any subset $U \subseteq [n]$, for any input vectors $x \in \Omega_{[n]}$ and $y \in \Omega_U$, we use $(x : y) \in \Omega_{[n]}$ to denote the vector such that $(x : y)_i = y_i$ if $i \in U$ and x_i otherwise. Moreover, we use $(x : y : z)$ to mean $((x : y) : z)$.

We prove the statement by induction on $|U|$.

Base case. Pick any player $j \in ([n] \setminus T) \cap \mathsf{S}(\varphi)$. Since player j is supposed to win with some positive probability $p_j > 0$, there must exist some $x_j \in \Omega_j$ such that $j \in \mathsf{S}(\varphi(\omega : x_j))$. Otherwise, all other players (who form a non-uniform adversary) can collude against j by choosing their inputs according to ω, in which case j wins with zero probability.

Inductive step. Suppose that for some $U \subseteq (N \setminus T) \cap \mathsf{S}(\varphi)$ (where $|U| \geq 1$), there exists some $x_U \in \Omega_U$ such that $U \subseteq \mathsf{S}(\varphi(\omega : x_U))$.

Pick any player $j \in (N \setminus (T \cup U)) \cap \mathsf{S}(\varphi)$; if there is no such player, the proof is finished.

Observe that for all $x \in \Omega_U$, $j \notin \mathsf{S}(\varphi(\omega : x))$. Otherwise, consider the closure vector $\mathcal{P}^{\varphi}(\omega_{N \setminus U})$ (corresponding to players in U aborting), whose support must include both T and j. However, the

closure vector must dominate $\mathcal{P}_{N\setminus U}(\varphi(\omega))$, whose coordinates in T already sum to 1, and so the closure vector cannot have a non-zero coordinate at j, otherwise Lemma 8.7 will be violated.

Next, fix any $y \in \Omega_j$. Observe that for any $x \in \Omega_U$, for any $i \in U$, we must have:

$\mathcal{P}_i(\varphi(\omega : x)) \geq \mathcal{P}_i(\varphi(\omega : x : y))$. Otherwise, consider the closure vector $\mathcal{P}^{\varphi}((\omega : x)_{N\setminus\{j\}})$ (corresponding to j aborting), which must dominate $\mathcal{P}(\varphi(\omega : x))$ whose j-th coordinate is 0 and other coordinates already sum to 1 (because Lemma 8.4 says $\varphi(\omega : x)$ is also tight); therefore, the i-th coordinate in $\mathcal{P}(\varphi(\omega : x))$ must also dominate $\mathcal{P}_i(\varphi(\omega : x : y)$, otherwise Lemma 8.7 will be violated.

Because φ is tight, it follows that if all players in U sample $X \leftarrow \Omega_U$ honestly, we have for each $i \in U$:

$p_i = E[\mathcal{P}_i(\varphi(\omega : X))] \geq E[\mathcal{P}_i(\varphi(\omega : X : y)] = p_i$. This implies that for all $y \in \Omega_j$, all $x \in \Omega_U$, and all $i \in U$, $\mathcal{P}_i(\varphi(\omega : x)) = \mathcal{P}_i(\varphi(\omega : x : y)$.

Finally, consider the case when all players except j play according to $(\omega : x_U)$ in the first round of φ. We have just proved that no matter what input $y \in \Omega_j$ the player j chooses, for all $i \in U$, we have $\mathcal{P}_i(\varphi(\omega : x_U : y)) = \mathcal{P}_i(\varphi(\omega : x_U)) > 0$, which implies that $U \subseteq \mathsf{S}(\varphi(\omega : x_U : y))$. Since j is supposed to win with probability $p_j > 0$ if it plays honestly, there must exist some $y_j \in \Omega_j$ such that $j \in \mathsf{S}(\varphi(\omega : x_U : y_j))$.

Therefore, we have shown that $U \cup \{j\} \subseteq \mathsf{S}(\varphi(\omega : x_U : y_j))$, which completes the inductive step and also the proof. □

We state the induction hypothesis formally as follows.

Induction hypothesis. Suppose for some $\ell \geq 2$, any vector $\vec{p} \in [0,1]^N$ satisfying $\sum_{i\in N} p_i = 1$ that can be achieved by a tight $(\ell-1)$-round protocol selecting 1 winner must satisfy $\|\vec{p}\|_0 \leq 2^{\ell-1}$.

Lemma 8.10 (Inductive Step)**.** *Assuming the above induction hypothesis, any vector $\vec{p} \in [0,1]^N$ satisfying $\sum_{i\in N} p_i = 1$ that can be achieved by a tight ℓ-round protocol φ selecting 1 winner must satisfy $\|\vec{p}\|_0 \leq 2^{\ell}$.*

Proof. Assume on the contrary that there is some tight ℓ-round protocol achieving $\vec{p} \in [0,1]^N$ such that $\sum_{i\in N} p_i = 1$ and $\|\vec{p}\|_0 \geq 2^{\ell} + 1$. By merging extra players as in the proof of Lemma 8.8, we can assume that $|N| = 2^{\ell} + 1$ such that for all $i \in N$, $p_i > 0$.

Pick any $\omega \in \Omega_{[n]}$. By the induction hypothesis, $T := \mathsf{S}(\varphi(\omega))$ has size at most $2^{\ell-1}$, because $\varphi(\omega)$ is a tight $(\ell-1)$-round protocol (by Lemma 8.4). We next apply Lemma 8.9 with $U = N \setminus T$, which says there is some $\omega' \in \Omega_{[n]}$ such that $U \subseteq \mathsf{S}(\varphi(\omega'))$. However, observe that $\varphi(\omega')$ is also a tight $(\ell-1)$-round protocol and $|U| \geq 2^{\ell-1} + 1$. This contradicts the induction hypothesis. □

Acknowledgment

Elaine Shi would like to thank Rafael Pass and abhi shelat for explaining Feige's lightest bin protocol several years ago. She would like to thank Yevgeniy Dodis and Daniel Wichs for explaining seeded extractors and samplers, and for insightful technical discussions that partly inspired this work; she also would like to thank Iddo Bentov and Andrew Miller for discussions on how to run fair lotteries on blockchains. She also would like to thank Benedikt Bunz for helpful discussions. Elaine Shi is partially supported by NSF under the award numbers CNS-1601879 and CNS-1561209, a Packard Fellowship, and an ONR YIP award. T-H. Hubert Chan is partially supported by the Hong Kong RGC under the grants 17200418 and 17201220.

References

[AA11] Dan Alistarh and James Aspnes. Sub-logarithmic test-and-set against a weak adversary. In David Peleg, editor, *DISC*, 2011.

[AAG+10] Dan Alistarh, Hagit Attiya, Seth Gilbert, Andrei Giurgiu, and Rachid Guerraoui. Fast randomized test-and-set and renaming. In *DISC*, 2010.

[ACH11] Gilad Asharov, Ran Canetti, and Carmit Hazay. Towards a game theoretic view of secure computation. In *Eurocrypt*, 2011.

[ADGH06] Ittai Abraham, Danny Dolev, Rica Gonen, and Joseph Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *PODC*, 2006.

[ADH19] Ittai Abraham, Danny Dolev, and Joseph Y. Halpern. Distributed protocols for leader election: A game-theoretic perspective. *ACM Trans. Econ. Comput.*, 7(1), February 2019.

[ADMM16] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and undefinedukasz Mazurek. Secure multiparty computations on bitcoin. *Commun. ACM*, 59(4):76–84, March 2016.

[AGV15] Dan Alistarh, Rati Gelashvili, and Adrian Vladu. How to elect a leader faster than a tournament. In *PODC*, 2015.

[AL11] Gilad Asharov and Yehuda Lindell. Utility dependence in correct and fair rational secret sharing. *Journal of Cryptology*, 24(1), 2011.

[AO16] Bar Alon and Eran Omri. Almost-optimally fair multiparty coin-tossing with nearly three-quarters malicious. In *TCC (B1)*, pages 307–335. Springer, 2016.

[BBBF18] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *CRYPTO*, 1 2018.

[BBF18] Dan Boneh, Benedikt Bünz, and Ben Fisch. A survey of two verifiable delay functions. Cryptology ePrint Archive, Report 2018/712, 2018.

[BGKO12] Amos Beimel, Adam Groce, Jonathan Katz, and Ilan Orlov. Fair computation with rational players. In *Eurocrypt*, 2012.

[BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *STOC*, 1988.

[BK14] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *CRYPTO*, pages 421–439, 2014.

[Blu83] Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *SIGACT News*, 15(1):23–27, January 1983.

[BMR90] D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols. In *STOC*, 1990.

[BOO10] Amos Beimel, Eran Omri, and Ilan Orlov. Protocols for multiparty coin toss with dishonest majority. In Tal Rabin, editor, *CRYPTO*, 2010.

[BPV06] Michael Ben-Or, Elan Pavlov, and Vinod Vaikuntanathan. Byzantine agreement in the full-information model in o(log n) rounds. In *STOC*, pages 179–186, 2006.

[BZ17] Massimo Bartoletti and Roberto Zunino. Constant-deposit multiparty lotteries on bitcoin. In *Financial Cryptography and Data Security*, 2017.

[CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *STOC*, 1988.

[CGL+18] Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, 2018.

[Cle86] Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *STOC*, 1986.

[DHR00] Yevgeniy Dodis, Shai Halevi, and Tal Rabin. A cryptographic solution to a game theoretic problem. In *CRYPTO*, 2000.

[DI05] Ivan Damgård and Yuval Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In *CRYPTO*, 2005.

[Dod06] Yevgeniy Dodis. Fault-tolerant leader election and collective coin-flipping in the full information model. Manuscript, 2006.

[DPS19] Phil Daian, Rafael Pass, and Elaine Shi. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In *FC*, 2019.

[DR07] Yevgeniy Dodis and Tal Rabin. Cryptography and game theory. In *AGT*, 2007.

[Fei99] U. Feige. Non-cryptographic selection protocols. In *FOCS*, 1999.

[GHS83] R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 1983.

[GK12] Adam Groce and Jonathan Katz. Fair computation with rational players. In *Eurocrypt*, 2012.

[GKM+13] Juan A. Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *FOCS*, 2013.

[GKTZ15] Juan Garay, Jonathan Katz, Björn Tackmann, and Vassilis Zikas. How fair is your protocol? a utility-based approach to protocol optimality. In *PODC*, 2015.

[GLR10] Ronen Gradwohl, Noam Livne, and Alon Rosen. Sequential rationality in cryptographic protocols. In *FOCS*, 2010.

[GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *ACM symposium on Theory of computing (STOC)*, 1987.

[GTZ15] Juan A. Garay, Björn Tackmann, and Vassilis Zikas. Fair distributed computation of reactive functions. In *DISC*, volume 9363, pages 497–512, 2015.

[GUV09] Venkatesan Guruswami, Christopher Umans, and Salil P. Vadhan. Unbalanced expanders and randomness extractors from parvaresh-vardy codes. *J. ACM*, 56(4):20:1–20:34, 2009.

[GW12] George Giakkoupis and Philipp Woelfel. On the time and space complexity of randomized test-and-set. In *PODC*, 2012.

[HT04] Joseph Halpern and Vanessa Teague. Rational secret sharing and multiparty computation. In *STOC*, 2004.

[HT14] Iftach Haitner and Eliad Tsfadia. An almost-optimally fair three-party coin-flipping protocol. In *STOC*, page 408–416, 2014.

[IML05] Sergei Izmalkov, Silvio Micali, and Matt Lepinski. Rational secure computation and ideal mechanism design. In *FOCS*, 2005.

[J.A74] Robert J.Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1), 1974.

[Kat08] Jonathan Katz. Bridging game theory and cryptography: Recent results and future directions. In *TCC*, 2008.

[KKM85] 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.

[KN08] Gillat Kol and Moni Naor. Cryptography and game theory: Designing protocols for exchanging information. In *TCC*, 2008.

[KRDO17] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Crypto*, 2017.

[LP15] Huijia Lin and Rafael Pass. Constant-round nonmalleable commitments from any one-way function. *J. ACM*, 62(1):5:1–5:30, 2015.

[LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *TCC*, 2008.

[MB17] Andrew Miller and Iddo Bentov. Zero-collateral lotteries in bitcoin and ethereum. In *EuroS&P Workshops*, 2017.

[MMZ+] Deepak Maram, Harjasleen Malvai, Fan Zhang, Nerla Jean-Louis, Alexander Frolov, Tyler Kell, Tyrone Lobban, Christine Moy, Ari Juels, and Andrew Miller. Candid: Can-do decentralized identity with legacy compatibility, sybil-resistance, and accountability. `https://eprint.iacr.org/2020/934`.

[MNS16] Tal Moran, Moni Naor, and Gil Segev. An optimally fair coin toss. *J. Cryptol.*, 29(3):491–513, July 2016.

[Nas51] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2), 1951.

[OPRV09] Shien Jin Ong, David C. Parkes, Alon Rosen, and Salil P. Vadhan. Fairness with an honest minority and a rational majority. In *TCC*, 2009.

[Pas04] Rafael Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *STOC*, 2004.

[PS17] Rafael Pass and Elaine Shi. Fruitchains: A fair blockchain. In *PODC*, 2017.

[RSZ99] Alexander Russell, Michael Saks, and David Zuckerman. Lower bounds for leader election and collective coin-flipping in the perfect information model. In *STOC*, 1999.

[RVW00] Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *FOCS*, 2000.

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

[Vad12] Salil P. Vadhan. Pseudorandomness (foundations and trends in theoretical computer science), 2012.

[Yao82] Andrew Chi-Chih Yao. Protocols for secure computations. In *FOCS*, 1982.

[Zuc96] David Zuckerman. Randomness-optimal sampling, extractors, and constructive leader election. In *STOC*, 1996.

A Additional Details Regarding Sequential Approximate Fairness

A.1 Non-Equivalence of Approximate Maximin Fairness and Approximate CSP Fairness

Even though leader-election is a constant-sum game, approximate maximin fairness and approximate CSP fairness are not the same (both for the non-sequential and sequential formulations). For example, a protocol where a 99% coalition A can surely exclude a specific individual Bob from winning (while not affecting any other honest individual's expected utility) can potentially be $(1 - o(1))$-CSP-fair, because the $1/n$ utility transferred from Bob to the coalition A is relatively small w.r.t. A's default chance of winning (had A played honestly), that is, $|A|/n = 99\%$. On the flip side, a protocol can be $(1 - \epsilon)$-maximin-fair against a small coalition A containing $O(1)$ players for some tiny constant $\epsilon \in (0, 1)$, but since A may be able to steal ϵ/n utility from everyone else, it may be able to significantly increase its own gains by a $\Theta(n)$ factor.

A.2 Equivalent Formulation of Sequential Approximate Fairness

For our sequential approximate fairness proofs, we will check round by round and argue that the adversary has no ϵ incentive to deviate in any round. We show that Definitions 3.1 and 3.2 are equivalent to an alternative version where we essentially rule out the incentive for deviation round by round.

Henceforth given a coalition $A \subset [n]$ playing the strategy S, we use the notation $\mathsf{view}^r_{A(S)}(tr)$ (often abbreviated as $\mathsf{view}^r_A(tr)$) to denote the coalition's view in the protocol up to the beginning of round r, after having observed honest players's round-r messages.

Fix a coalition A and its strategy S. Let $\mathsf{Dev}^{r,A(S)}(tr)$ denote the event that $A(S)$ first deviates in round r given a sample path tr which consists of all players' random coins. We sometimes write Dev^r when $A(S)$ is clear from the context.

Lemma A.1 (An equivalent formulation of sequential approximate fairness)**.** *Let $\epsilon \in (0,1)$. A leader election protocol Π achieves $(1-\epsilon)$-sequential-CSP-fairness against a (non-uniform p.p.t.) coalition $A \subset [n]$ iff there exists a negligible function $\mathsf{negl}(\cdot)$, s.t. at least one of the following holds for every r:*

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

Note that a similar equivalence lemma holds for maximin fairness too, and the proof is almost identical. Below, we only prove the equivalence for CSP-fairness.

Proof. The $\implies$ direction is obvious. We now prove the $\impliedby$ direction. Observe that for any r, for any κ, Definition 3.1 says that at least one of the following holds for some negligible function $\mathsf{negl}(\cdot)$:

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

$$\mathbf{E}_{A(S)}\left[W^A \cdot \mathsf{Dev}^r\right] \leq \frac{1}{1-\epsilon} \cdot \mathbf{E}_{A(\Pi)}\left[W^A \cdot \mathsf{Dev}^r\right] + \mathsf{negl}(\kappa) \cdot \Pr_{A(S)}[\mathsf{Dev}^r] \tag{5}$$

Note that Dev can be partitioned into the disjoint events $\mathsf{Dev}^1, \mathsf{Dev}^2, \ldots$. Then, summing (5) over all r, we have

$$\mathbf{E}_{A(S)}\left[W^A \cdot \mathsf{Dev}\right] \leq \frac{1}{1-\epsilon} \cdot \mathbf{E}_{A(\Pi)}\left[W^A \cdot \mathsf{Dev}\right] + \mathsf{negl}'(\kappa)$$

Therefore, either $\Pr[\mathsf{Dev}]$ is negligible; or if not, it must be that $\Pr_{A(S)}\left[W^A \;\middle|\; \mathsf{Dev}\right] \leq \frac{1}{1-\epsilon} \cdot \Pr_{A(\Pi)}\left[W^A \;\middle|\; \mathsf{Dev}\right] + \mathsf{negl}''(\kappa)$, as required. In the above, negl' and negl'' denote different negligible functions. □

A.3 Comparison with Rational Protocol Design

Garay et al. [GKM+13, GKTZ15, GTZ15] suggest a the Rational Protocol Design (RPD) paradigm for studying game-theoretic cryptographic protocols. In their work, they model a *meta-game*, i.e., a Stackelberg game between the protocol designer and an attacker: the designer first picks a protocol Π, then the attacker can design which coalition to corrupt and its attack strategy after examining this protocol Π. They suggest a solution concept that achieves a subgame perfect equilibrium in this Stackelberg meta-game. The existing works on RPD [GKM+13, GKTZ15, GTZ15] consider MPC-style protocols where the adversary's utility comes from breaking either privacy or correctness.

In comparison with the RPD line of work, we consider a different type of protocol where the most natural utility function is very different. Nonetheless, it is interesting to consider whether we can think of our notion as a meta-game between the protocol designer and the attacker too. We therefore slightly modify the RPD definitions by 1) changing the utility definition; and 2) allowing a multiplicative version of approximation. We show that our non-sequential game-theoretic fairness notions can be interpretted in this modified RPD framework. However, as mentioned in our work, our new, sequential approximate notions provide a better solution concept when approximation is necessary.

For completeness, we describe the slightly modified RPD paradigm in our leader election context. Consider the following Stackelberg meta-game. The protocol designer first chooses some protocol

$\Pi \in \mathsf{ITM}^n$ for the leader election problem where ITM^n denote the space of all possible n-party p.p.t. protocols that *satisfy the correctness definition* specified in Section 3, i.e., under an all-honest execution, every player is equally likely to get elected. The designer then gives the description of the protocol Π to the adversary who controls a coalition $A \subset [n]$, and is additionally targeting some individual $i \notin [n]$ in the case of maximin-fair utility. The adversary now picks a (non-uniform p.p.t.) strategy $S(\Pi)$. We care about these two types of utilities:

- *Maximin-fair utility.* The designer's utility $u_D(\Pi, S)$ is defined to be the player i's probability of being elected in the protocol Π when playing against the coalition A adopting the strategy S; and the adversary's utility $u_A(\Pi, S) = -u_D(\Pi, S)$, i.e., the game is zero-sum.
- *CSP-fair utility.* The adversary's utility $u_A(\Pi, S)$ is the probability that some member of A is elected leader when the coalition A adopts the strategy S, and the designer's utility $u_D(\Pi, S) := -u_A(\Pi, S)$.

Definition A.2 (Subgame perfect equilibrium)**.** Fix some $A \subset [n]$ and possibly some victim $i \notin A$ in the case of maximin-fair utility. Let S be a mapping from efficient protocols to non-uniform p.p.t. strategies. A pair of actions by the designer/adversary $(\Pi, S(\Pi))$ achieves an ϵ-subgame-perfect equilibrium in the above Stackelberg meta-game, iff the following hold:

1. for any (correct) protocol $\Pi' \in \mathsf{ITM}^n$, $u_D(\Pi', S(\Pi')) \leq u_D(\Pi, S(\Pi)) + \epsilon \cdot |u_D(\Pi, S(\Pi))|$; and
2. for any $\widetilde{S}$, $u_A(\Pi, \widetilde{S}(\Pi)) \leq u_A(\Pi, S(\Pi)) + \epsilon \cdot |u_A(\Pi, S(\Pi))|$.

Henceforth, we use $S^\Pi(\Pi)$ (or S^Π for short) to denote the trivial strategy where the coalition A simply acts honestly according to the prescribed protocol Π.

The following simple fact proves an equivalence of our non-sequential approximate fairness notions and the above RPD-inspired formulation. As mentioned earlier in Sections 1 and 3, we believe that our *sequential approximate* notion provides a better solution concept for defining *approximate* game-theoretic fairness in the context of multi-party protocols.

Fact A.3 (Relationship to RPD)**.** *A (correct) leader election protocol Π satisfies non-sequential $(1-\epsilon)$-maximin-fairness against the coalition $A \subseteq [n]$, iff for any $i \notin A$, (Π, S^Π) is an ϵ-subgame-perfect equilibrium the above Stackelberg meta-game (fixing A and i), for the maximin-fair utility.*

Similarly, if a (correct) leader election protocol Π satisfies non-sequential $(1-\epsilon)$-CSP-fairness against the coalition $A \subseteq [n]$, then (Π, S^Π) is a 2ϵ-subgame-perfect equilibrium the above Stackelberg meta-game (fixing A and i), for the CSP-fair utility. Further, given a (correct) leader election protocol Π, if (Π, S^Π) is an ϵ-subgame-perfect equilibrium the above Stackelberg meta-game (fixing A and i), for the CSP-fair utility, then Π satisfies non-sequential $(1-\epsilon)$-CSP-fairness.

Proof. We prove it for maximin fairness and CSP-fairness, respectively.

Maximin fairness. First, consider a leader election protocol Π that is $(1-\epsilon)$-maximin-fair. We want to prove that (Π, S^Π) is an ϵ-subgame-perfect equilibrium of the Stackelberg meta-game w.r.t. the maximin-fair utility. The first condition, i.e., $\forall \Pi' \in \mathsf{ITM}^n$, $u_D(\Pi', S^{\Pi'}) \leq u_D(\Pi, S^\Pi) + \epsilon \cdot |u_D(\Pi, S^\Pi)|$, follows directly from the correctness definition of a leader election protocol. We now prove the second condition, i.e., for any (non-uniform p.p.t.) strategy $\widetilde{S}$, $u_A(\Pi, \widetilde{S}) \leq u_A(\Pi, S^\Pi) + \epsilon \cdot |u_A(\Pi, S^\Pi)|$. Recall that the adversary's utility in this case is 0 or negative. Therefore, flipping the sign on both sides of the inequality, and observing that the Stackelberg meta-game is zero-sum, it suffices to prove that for any (non-uniform p.p.t.) strategy $\widetilde{S}$, $u_D(\Pi, \widetilde{S}) \geq u_D(\Pi, S^\Pi) - \epsilon \cdot u_D(\Pi, S^\Pi) = (1-\epsilon) \cdot u_D(\Pi, S^\Pi)$. This follows directly from non-sequential $(1-\epsilon)$-maximin fairness.

Next, suppose that no matter who the targeted victim $i \notin A$ is, (Π, S^Π) is a ϵ-subgame-perfect equilibrium of the Stackelberg meta-game w.r.t. the maximin-fair utility. This means that for any (non-uniform p.p.t.) strategy $\widetilde{S}$, $u_D(\Pi, \widetilde{S}) \geq u_D(\Pi, S^\Pi) - \epsilon \cdot u_D(\Pi, S^\Pi) = (1-\epsilon) \cdot u_D(\Pi, S^\Pi)$, which is exactly the requirement of non-sequential $(1-\epsilon)$-maximin fairness.

CSP fairness. First, consider a leader election protocol Π that is $(1-\epsilon)$-CSP-fair. We want to prove that no matter who the targeted victim $i \notin A$ is, (Π, S^Π) is a 2ϵ-subgame-perfect equilibrium of the Stackelberg meta-game w.r.t. the CSP-fair utility. As before, the first condition follows directly from the correctness definition of a leader election protocol. We therefore focus on proving the second condition, i.e., for any (non-uniform p.p.t.) strategy $\widetilde{S}$, $u_A(\Pi, \widetilde{S}) \leq u_A(\Pi, S^\Pi) + 2\epsilon \cdot |u_A(\Pi, S^\Pi)|$. In this case, the adversary's utility non-negative. Therefore, it suffices to prove that for any (non-uniform p.p.t.) strategy $\widetilde{S}$, $u_A(\Pi, \widetilde{S}) \leq u_A(\Pi, S^\Pi) + 2\epsilon \cdot u_A(\Pi, S^\Pi) = (1+2\epsilon) \cdot u_A(\Pi, S^\Pi)$. This follows directly from non-sequential $(1-\epsilon)$-maximin fairness which guarantees that $u_A(\Pi, \widetilde{S}) \leq \frac{1}{1-\epsilon} \cdot u_A(\Pi, S^\Pi) < (1+2\epsilon) \cdot u_A(\Pi, S^\Pi)$. Note that the factor of 2 is because we define our multiplicative factor as $\frac{1}{1-\epsilon}$[13].

Next, suppose that (Π, S^Π) is an ϵ-subgame-perfect equilibrium of the Stackelberg meta-game w.r.t. the CSP-fair utility. This means that for any (non-uniform p.p.t.) strategy $\widetilde{S}$, $u_A(\Pi, \widetilde{S}) \leq u_A(\Pi, S^\Pi) + \epsilon \cdot u_A(\Pi, S^\Pi) \leq \frac{1}{1-\epsilon} \cdot u_A(\Pi, S^\Pi)$ which is exactly the requirement of non-sequential $(1-\epsilon)$-CSP fairness. □

A.4 Deferred Proofs from Section 3

Fact A.4 (Restatement of Fact 3.3)**.** *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. We prove it for CSP-fairness and the proof for maximin fairness is essentially the same. Suppose S is some strategy by the coalition and let Dev be the corresponding deviation event. We consider two cases. We use shorthand $\Pr_S[\cdot] := \Pr[\mathsf{Exec}^{A(S)}(1^\kappa) : \cdot]$ and $\Pr_\Pi[\cdot] := \Pr[\mathsf{Exec}^{A(\Pi)}(1^\kappa) : \cdot]$.

Observe that we have $\Pr_S[\mathsf{Dev}] = \Pr_\Pi[\mathsf{Dev}]$ and $\Pr_S[W^A \cdot \overline{\mathsf{Dev}}] = \Pr_\Pi[W^A \cdot \overline{\mathsf{Dev}}]$.

- **Case $\Pr[\mathsf{Dev}] \leq \mathsf{negl}(\kappa)$.** Then, we have $\Pr_S[W^A] = \Pr_S[W^A \cdot \mathsf{Dev}] + \Pr_S[W^A \cdot \overline{\mathsf{Dev}}] \leq \mathsf{negl}(\kappa) + \Pr_\Pi[W^A]$.
- **Case $\Pr[\mathsf{Dev}]$ is non-negligible in κ.** In this case, our sequential CSP-fairness definition implies that $\Pr_S[W^A | \mathsf{Dev}] \leq \frac{1}{1-\epsilon} \cdot \Pr_\Pi[W^A | \mathsf{Dev}] + \mathsf{negl}(\kappa)$.

Hence, we have

$$\begin{aligned}
\Pr_S[W^A] &= \Pr_S[W^A \cdot \mathsf{Dev}] + \Pr_S[W^A \cdot \overline{\mathsf{Dev}}] \\
&\leq \frac{1}{1-\epsilon} \cdot \Pr_\Pi[W^A \cdot \mathsf{Dev}] + \mathsf{negl}(\kappa) \cdot \Pr[\mathsf{Dev}] + \Pr_\Pi[W^A \cdot \overline{\mathsf{Dev}}] \\
&\leq \frac{1}{1-\epsilon} \cdot \Pr_\Pi[W^A] + \mathsf{negl}(\kappa)
\end{aligned}$$

[13] We choose to use this style because it is more natural if 1-fair is perfectly fair and 0-fair means not fair at all.

In either case, by choosing $\epsilon' := \epsilon + \mathsf{negl}'(\kappa)$ for some suitable negligible function $\mathsf{negl}'(\cdot)$, we can conclude that $\Pr_S[W^A] \leq \frac{1}{1-\epsilon'} \cdot \Pr_\Pi[W^A]$, where we have used the fact that $\Pr_\Pi[W^A] \geq \frac{1}{n}$. □

Fact A.5 (Restatement of Fact 3.4)**.** *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. As in Fact 3.3, we prove the part concerning CSP-fairness, and the maximin-fairness part is almost identical. Suppose the protocol Π does not satisfy the sequential definition. This means that the coalition has a strategy S such that the deviation probability $p := \Pr_S[\mathsf{Dev}]$ is non-negligible, and moreover for any ρ_1 and ρ_2 that are negligible, we have $\Pr_S[W^A|\mathsf{Dev}] \geq \frac{1}{1-\rho_1} \Pr_\Pi[W^A|\mathsf{Dev}] + \rho_2$. Therefore, we have

$$\begin{aligned}
\Pr_S[W^A] &= \Pr_S[W^A|\mathsf{Dev}] \cdot \Pr[\mathsf{Dev}] + \Pr_S[W^A|\overline{\mathsf{Dev}}] \cdot \Pr[\overline{\mathsf{Dev}}] \\
&\geq \left(\frac{1}{1-\rho_1} \cdot \Pr_\Pi[W^A|\mathsf{Dev}] + \rho_2\right) \cdot \Pr[\mathsf{Dev}] + \Pr_\Pi[W^A|\overline{\mathsf{Dev}}] \cdot \Pr[\overline{\mathsf{Dev}}] \\
&\geq \Pr_\Pi[W^A|\mathsf{Dev}] \cdot \Pr_\Pi[\mathsf{Dev}] + \Pr_\Pi[W^A|\overline{\mathsf{Dev}}] \cdot \Pr_\Pi[\mathsf{Dev}] + \rho_2 \cdot \Pr[\mathsf{Dev}] \\
&\geq \Pr_\Pi[W^A] + \rho_2 \cdot \Pr[\mathsf{Dev}]
\end{aligned}$$

Since the above holds for any negligible function ρ_2, it means that there exists some non-negligible ξ, such that $\Pr_S[W^A] \geq \Pr_\Pi[W^A] + \xi \cdot \Pr[\mathsf{Dev}]$. This means the protocol Π is not $(1-\mathsf{negl}(\cdot))$-CSP-fair for any negligible $\mathsf{negl}(\cdot)$. □

B Tournament Tree Protocol

In this section, we provide details on the well-known tournament tree protocol. Whenever appropriate, we state the security of our protocols assuming that the commitment scheme adopted is the ideal functionality $\mathcal{F}_{\text{comm}}$. For all lemmas stated, the same guarantees are still respected (except for a negligibly small additive loss in security) if we instead instantiate the commitment with a publicly verifiable, non-malleable commitment scheme as defined in Section 4.3. We will illustrate how to achieve this using a hybrid argument (e.g., in the proof of of Theorem B.4).

B.1 Duel Protocol

First, we extend the two-party Blum coin toss protocol [Blu83] (henceforth also called the duel protocol) to one that supports rational distributions. In particular, for $i \in [2]$, we have $k_i \in \mathbb{Z}_{>0}$ such that player i wins with probability $\frac{k_i}{k_1+k_2}$.

To achieve this, we can modify Blum's coin toss as follows:

- Denote $k = k_1 + k_2$ and $\ell := \lceil \log_2 k \rceil$. First, each player $i \in [2]$ commits to an ℓ-bit random string that denotes some $s_i \in \mathbb{Z}_k$.
- Next, each player i opens their commitment and reveals s_i. If $s_1 + s_2 \mod k \in \{0, 1, \ldots, k_1 - 1\}$, then the player 1 wins; else, player 2 wins.

  If a player aborts or fails to correctly open the commitment that reveals an element in $\mathbb{Z}_k$, then it is treated as having forfeited and the other player wins.

It is easy to see that the following fact holds.

Fact B.1. *Suppose that in the duel protocol, the commitment scheme is instantiated with the perfect non-malleable commitment $\mathcal{F}_{\text{comm}}$ functionality, then for each $i \in [2]$, if player i plays honestly, then it wins with probability at least $\frac{k_i}{k_1+k_2}$.*

B.2 Tournament Tree Protocol for Leader Election

Recall that the tournament tree protocol can elect a leader from n players. We want that in an honest execution, every player gets elected with probability $\frac{1}{n}$.

We use $\mathsf{Elect}(S)$ to denote the protocol that elects a leader from the set S of players, where $|S| = n$:

$\mathsf{Elect}(S)$:

Assume: S is a set of size $n \geq 1$.

- If $n = 1$, simply return single player in S; else continue with the following.
- Suppose that the players in S are arranged in a lexicographical order by their respective identities. Denote $n_1 := \lfloor \frac{n}{2} \rfloor$ and $n_2 := \lceil \frac{n}{2} \rceil$. Let S_1 be the first n_1 players in S and $S_2 := S \setminus S_1$.
- In parallel, for $i \in [2]$, run $\mathsf{Elect}(S_i)$ to elect $P_i \in S_i$.
- The final winner is determined by the duel protocol between P_1 and P_2 such that for $i \in [2]$, P_i wins with probability $\frac{n_i}{n}$.

Fact B.2. *The above protocol finishes in $O(\log n)$ number of rounds if the commitment protocol adopted is constant round; and moreover, in an honest execution, every player has an equal chance of being elected leader.*

Remark B.3. Obviously, Theorem B.4 holds if the commitment is instantiated with an idealized commitment functionality $\mathcal{F}_{\text{comm}}$: for a player to commit a value, it simply sends it to $\mathcal{F}_{\text{comm}}$; for the committer to open the value, it sends "`open`" to $\mathcal{F}_{\text{comm}}$, and $\mathcal{F}_{\text{comm}}$ sends the committed value to everyone.

The hybrid argument deals with the case when $\mathcal{F}_{\text{comm}}$ is replaced with a real commitment scheme in Appendix 4.3.

Theorem B.4 (Game theoretic characterization of the tournament-tree protocol)**.** *The above generalized tournament-tree protocol, when instantiated with a suitable publicly verifiable, non-malleable commitment scheme (see Section 4.3), satisfies $(1-\mathsf{negl}(\kappa))$-sequential-cooperative-strategy-proofness and $(1-\mathsf{negl}(\kappa))$-sequential-maximin-fairness against non-uniform p.p.t. coalitions of arbitrarily sizes.*

Proof. Because of Lemma A.1, we fix some arbitrary round r and consider a trace for which A's strategy S wants to deviate at the beginning of round r.

We show that conditioned on any such trace till the beginning of round r, A can only negligibly increase its probability of winning or harm any individual. Observe that the remaining protocol corresponds to a tree, where each leaf corresponds to a player that still has a chance to become the leader and each internal node corresponds to an instance of some duel protocol (whose winning probabilities are already specified).

Suppose there are m internal nodes and we label them using any breadth-first order from the root, i.e., the root (which is the final duel that determines the final winner) receives label 1, and some internal node furthest away from the root receives label m. Recall that each node i corresponds to some duel protocol between players with some winning probabilities p_i and q_i such that $p_i + q_i = 1$.

Fix such a trace till the beginning of round r, and consider a continuation of the randomized execution based on the current trace tr. We now define a sequence of hybrid experiments:

- Let Hyb_0 denote an execution of the real-world protocol continuing from this round.
- For $i \in \{1, 2, \ldots, m\}$, Hyb_i is almost identical to Hyb_{i-1}, except that in node i of the tournament tree, we instead perform the following:
  - Let view_A be the internal state of the adversarial algorithm representing the corrupt coalition right before players start the duel protocol corresponding to node i of the tree. We now call the simulator of the non-malleable commitment scheme (see Section 4.3) feeding it view_A. Now the simulator outputs an updated state view'_A (including the committed values).
  - For the duel protocol at node i, we do the following to replace the real protocol:
    1. if a corrupt player plays with an honest player, then we use the honest player's opening and the committed value output by the simulator to determine the winner (note that this implicitly assumes that the corrupt player does not abort);
    2. if both players are honest, the winner is determined by the honest protocol between the two players;
    3. if both players are corrupt, an arbitrary one can be the winner.

Notice that in Hyb_m, every time a corrupt player in A duels with an honest player, a simulator is called to produce the commitment; and therefore the committed value must be independent of the honest player's committed value. Thus, in Hyb_m, every time a corrupt player in A duels with an honest player, the honest player survives with at least the probability indicated by the duel protocol.

This means that in Hyb_m regardless of A's strategy, any honest player has at least the correct probability of winning, and the coalition A has no unfair advantage of winning.

Let p_i^j denote the probability that player j wins in Hyb_i, and let p_i^A denote the probability that someone in A wins in Hyb_i. Now, due to the security of the non-malleable commitment scheme, in any pair of adjacent hybrid experiments Hyb_{i-1} and Hyb_i, it follows that for any $j \notin A$, $p_{i-1}^j \geq p_i^j - \mathsf{negl}(\kappa)$, and $p_{i-1}^A \leq p_i^A + \mathsf{negl}(\kappa)$. The theorem now follows through a hybrid argument. □

B.3 Failed Idea: Collapsing the Tournament-Tree Protocol to Two Rounds

To achieve better round complexity, 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. For example, consider the following scenario: there are 4 players, Alice, Bob, Charles, and Daisy. Alice and Bob are in the first bracket, and Charles and Daisy are in the second bracket. Suppose that Alice and Bob form a corrupt coalition. They can each commit $(0, 0)$ and $(1, 1)$ respectively. They then wait for Charles and Daisy to open their random coins, and suppose without loss of generality, Daisy survives to the second round, and her random coin in the second round is b. At this moment, the coalition knows which of Alice and Bob should survive to the second round of the tournament,

to defeat Daisy and become the final winner — without loss of generality, suppose that Alice should survive to the next round. The coalition can simply make Bob abort in the first round, to ensure that Alice survives.

C Another Counter-Example that Violates Sequential Fairness

Earlier in Section 1, we gave an example undesirable protocol that violates sequential fairness, but the example may seem somewhat contrived. In this section, we give another counter-example. Consider the following protocol:

1. First, all players use an *idealized* committee election protocol to randomly elect a preliminary committee of size $\log^8 n$. We may assume that this step is done by a trusted authority — even if this step is completely ideal, we can show that the entire protocol is not sequentially fair.

2. Second, all players in the preliminary committee elect a final committee of size $\log^6 n$ as follows: each player in the preliminary committee broadcasts a random bit, and the concatenation of all players' random bits are fed into a random oracle, whose outcome will determine the final committee. If any player fails to post a bit, we will treat its bit as 0.

3. Finally, the final committee runs the tournament-tree protocol to elect the final leader.

This protocol does not satisfy sequential CSP fairness. Suppose that there is coalition consisting of a single corrupt player. There is a non-negligible probability that the corrupt player is elected into the preliminary committee. If this small probability event happens, the corrupt player will be incentivized to deviate in the final committee election: if it looks at the honest coins and then decides its own coin, it can increase its conditional winning probability by a factor of roughly 2x. A good approximate fairness notion should rule out undesirable protocols like this.

D Open Questions

In comparision with the well-studied unbiasability notion [Cle86], our understanding of game-theoretic fairness is still very little. Our work leaves open various exciting future directions. We give some examples below:

- Recall that in comparison with the tournament tree protocol, we relax the fairness notion to approximate, and allow more general usage of cryptography. Can we precisely characterize the minimal assumptions necessary to obtain a sub-logarithmic round, game-theoretically fair leader election protocol?

- For technical reasons, our protocol does not give a meaningful result for $o(\log \log n)$ rounds. Can we have any meaningful fairness guarantee for the $o(\log \log n)$-round regime?

- In contrast with leader election, recalle that Chung et al. [CGL$^+$18] showed that for the binary-coin case, broad impossibility results apply. This stark difference seems intriguing, and begs for further exploration for formulations in between these two cases, and giving a more complete theoretical characterization of the landscape.

- Finally, what is the natural extension of our game-theoretic notions to general multi-party computation, and what kind of interesting results can we obtain?