$\log^*$-Round Game-Theoretically-Fair Leader Election*

Ilan Komargodski[1], Shin'ichiro Matsuo[2], Elaine Shi[3], and Ke Wu[3]

[1] School of Computer Science and Engineering, Hebrew University of Jerusalem and NTT Research, 91904 Jerusalem, Israel
ilank@cs.huji.ac.il

[2] NTT Research, 94085 Sunnyvale, USA
and Department of Computer Science, Georgetown University, 20057 D.C., USA
shinichiro.matsuo@ntt-research.com

[3] Computer Science Department, Carnegie Mellon University, 15213 Pittsburgh, USA
kew2@andrew.cmu.edu, runting@gmail.com

Abstract. It is well-known that in the presence of majority coalitions, *strongly fair* coin toss is impossible. A line of recent works have shown that by relaxing the fairness notion to game theoretic, we can overcome this classical lower bound. In particular, Chung et al. (CRYPTO'21) showed how to achieve approximately (game-theoretically) fair leader election in the presence of majority coalitions, with round complexity as small as $O(\log \log n)$ rounds.

In this paper, we revisit the round complexity of game-theoretically fair leader election. We construct $O(\log^* n)$ rounds leader election protocols that achieve $(1-o(1))$-approximate fairness in the presence of $(1-o(1))n$-sized coalitions. Our protocols achieve the same round-fairness trade-offs as Chung et al.'s and have the advantage of being conceptually simpler. Finally, we also obtain game-theoretically fair protocols for committee election which might be of independent interest.

* The full version of this paper is available at [KMSW22].

1 Introduction

Suppose that Murphy, Murky, and Moody co-authored a paper that proved a ground-breaking theorem and the paper got accepted at the prestigious CRYPTO'22 conference. Murphy, Murky, and Moody want to run a coin toss protocol over the Internet to elect a winner who will present the paper at the conference. Since everyone wants to go to the beautiful beaches of Santa Barbara, all of them want to be the winner. They each are worried that the other coauthors might deviate from the honest protocol to gain an unfair advantage. There is both good and bad news. The bad news is that due to a famous lower bound by Cleve [Cle86], there is no *strongly fair* coin toss protocol when half of the parties may be corrupt and misbehaving — roughly speaking, strong fairness requires that the coalition cannot bias the outcome of the coin toss whatsoever. The good news is that a more recent line of work [CCWS21, GGS, CGL+18, WAS22] has shown that a relaxed fairness notion called *game-theoretic* fairness is indeed possible for the leader election problem, even when an arbitrary number of parties may be corrupt. To see why, first observe that the original Blum's coin toss protocol [Blu83] actually gives a game-theoretically fair leader election scheme for $n = 2$ parties. Imagine that each party first commits to a random coin, they then open their coin, and the XOR of the two bits is used to elect a random winner. If one party fails to commit or correctly open, it is eliminated and the remaining party is declared the winner. Blum's coin toss satisfies *game-theoretic* fairness in the following sense. As long as the commitment scheme is not broken, a corrupt layer cannot bias the coin *to its own favor* no matter how it deviates from the protocol. Note that Blum's protocol is not strongly fair since a corrupt party can indeed bias the coin, but only to the other player's advantage.

For the more general case of the n parties, we can use a folklore tournament-tree protocol to accomplish the same purpose. Suppose that n is a power of 2 for simplicity. We first divide the n parties into $n/2$ pairs, and each pair elects a winner using Blum's coin toss. The winner survives to the next round, where we again divide the surviving $n/2$ parties into $n/4$ pairs. The protocol continues after a final winner is elected after $\log_2 n$ rounds. At any point in the protocol, if a party fails to commit or correctly open its commitment, it is eliminated and its opponent survives to the next round.

The recent work of Chung et al. [CCWS21] argued that this simple tournament tree protocol satsfies a strong notion of game-theoretic fairness as explained below. Suppose that the winner obtains a utility of 1 and everyone else obtains a utility of 0. As long as the commitment scheme is not broken, the tournament tree protocol guarantees that 1) no coalition of any size can *increase its own expected utilty* no matter what (polynomially-bounded) strategy it adopts; and 2) no coalition of any size can *harm any individual honest player's expected utility*, no matter what (polynomially-bounded) strategy it adopts. Recent work in this space [CCWS21, GGS, CGL+18, WAS22] calls the former notion cooperative-strategy-proofness (or *CSP-fairness* for short), and calls the latter notion *maximin fairness*. Philosophically, CSP-fairness guarantees that any rational, profit-seeking individual or coalition has no incentive to deviate from the honest protocol; and maximin fairness ensures that any paranoid individual who wants to maximally protect itself in the worst-case scenario has no incentive to deviate either. In summary, the honest protocol is an equilibrium and also the best response for every player and coalition. Therefore, prior works [CGL+18, CCWS21, WAS22, GGS] have argued that game-theoretic notions of fairness are compelling and worth investigating because 1) they are arguably more natural (albeit stricly weaker) than the classical strong fairness notion in practical applications; and 2) the game-theoretic relaxation allows us to circumvent classical impossibility results pertaining to strong fairness in the presence of majority coalitions [Cle86].

Having established the general feasibility of game-theoretically fair leader election in the presence of majority-sized coalitions, Chung et al. [CCWS21] asked the following natural question: *what is the round complexity of game-theoretically fair leader election in the presence of majority coalitions*? Specifically, can we asymptotically outperform the logarithmic round complexity of the folklore tournament tree protocol? They then gave a partial answer to this question, showing that for any desired round complexity parameter $\Theta(\log\log n) \leq R \leq \log n$, there is an $O(R)$-round n-party leader election protocol that achieves $\left(1 - \frac{1}{2^{\Theta(R)}}\right)$-fairness against coalitions of size up to $\left(1 - \frac{1}{2^{\Theta(R)}}\right) n$. In particular, their result statement adopts an approximate notion of game-theoretic fairness. Roughly speaking, a protocol is $(1-\epsilon)$-fair if it satisfies the aforementioned game theoretic fairness (including CSP-fairness and maximin fairness) up to an ϵ slack. More specifically, we want that the coalition's expected utility cannot exceed $1/(1-\epsilon)$ times its normal utility had everyone behaved honestly, and we require that any honest individual's expected utility cannot drop below $(1-\epsilon)$ times its normal utility had everyone behaved honestly. Chung et al.'s result [CCWS21] enables a smooth and mathematically quantifiable tradeoff between the efficiency of the protocol and its resilience to strategic behavior. However, their result requires the protocol to have at least $\Theta(\log\log n)$ rounds to give any meaningful fairness guarantee. Indeed, a more careful examination suggests that their framework has a *sharp* cutoff at $\Theta(\log\log n)$ rounds, i.e., the approach fundamentally fails when we want round complexity to be less than $\log\log n$. Therefore, an obvious gap in our understanding is the following:

> *In the presence of majority-sized coalitions, can we achieve any meaningful fairness guarantee for small-round protocols whose round complexity is less than* $\log\log n$?

1.1 Our Results and Contributions

In this paper, we revisit the round complexity of game-theoretically fair leader election. We make the following contributions. First, we show positive results in the style of Chung et al. [CCWS21], but now for a broader range of parameters as explained in the following Theorem 1.1. In particular, our result shows that under standard cryptographic assumptions, there is a $O(\log^* n)$-round leader election protocol that achieves $(1-o(1))$-game-theoretic-fairness, in the presence of $(1-o(1))\cdot n$-sized coalitions.

Second, we give conceptually simpler constructions than those of Chung et al. [CCWS21], which also result in simpler analyses. More specifically, Chung et al.'s construction relies on combinatorial objects called extractors, which we get rid of in our construction. We believe that our conceptually simpler constructions can lend to better understanding and make it easier for future work to extend our framework. Interestingly, our constructions are inspired and have structural resemblance to Feige's famous lightest bin leader election protocol [Fei99]. We stress, however, that Feige's protocol itself does not satisfy game-theoretic fairness, but rather, achieves only a much weaker notion of resilience, i.e., an honest party is elected leader with constant probability. At a very high level, our approach augments Feige's protocol lightest-bin protocol with a "commit and open" and a "virtual identity" mechanism, and we prove that the resulting protocol satisfies the desired game-theoretic properties.

Third, we also present results for the more generalized problem of fair committee election, where the goal is to elect a committee of size c. The leader election problem can be viewed as a special case of committee election where $c = 1$. Our main results are summarized in the following theorems.

Theorem 1.1 (Game-theoretically fair leader election)**.** *Assume the existence of enhanced trapdoor permutations, and collision-resistant hash functions. Fix n and let $\log^* n \leq R \leq C \log n$ be the round complexity we want to achieve for some constant C. Then there exists an $O(R)$-round leader election that achieves $(1 - \frac{1}{2^{\Theta(R)}})$-game-theoretic fairness against a non-uniform p.p.t. coalition of size at most $(1 - \frac{L}{\Theta(R)})n$, where L is the smallest integer such that $\log^{(L)} n \leq 2^R$.*

For readers who are familiar with the line of work on approximate *strong* fairness [Cle86, MNS09, AO16, BOO10, HT14], an interesting observation is that for game-theoretic fairness, the efficiency-fairness tradeoff is exponentially better than that of strong fairness. Specifically, it is known that any R-round protocol cannot achieve $\Omega(1/R)$ *strong* fairness[4] against an $n/2$-sized coalition, whereas we show that R-round protocols can achieve $(1 - 1/2^{\Theta(R)})$-fairness.

Theorem 1.2 (Game-theoretically fair committee election)**.** *Assume the existence of enhanced trapdoor permutations and collision-resistant hash functions. Fix n and c. Let L^* be the smallest integer such that $\log^{(L^*)} n \leq c$. Then for any $L^* \leq R \leq C_0 \log n$ for some constant C_0, we have that*

- *If $c \geq 2^R$, there exists an $O(R)$-round committee election that achieves $(1 - \frac{1}{c^{\Theta(1)}})$-game-theoretic fairness against a non-uniform p.p.t. coalition of size at most $(1 - \frac{L^*}{\Theta(R)})n$.*
- *If $c < 2^R$, there exists an $O(R)$-round committee election that achieves $(1 - \frac{1}{2^{\Theta(R)}})$-game-theoretic fairness against a non-uniform p.p.t. coalition of size at most $(1 - \frac{L}{\Theta(R)})n$, where L is the smallest integer such that $\log^{(L)} n \leq 2^R$.*

[4] The approximate strong fairness line of work defines what we call $(1 - \epsilon)$-fairness as ϵ-fairness (but for the notion of strong fairness instead). Following the notations of Chung et al. [CCWS21], we flipped this notation to make it more intuitive: with our notation, 1-fair is more fair than 0-fair which agrees with our intuition.

Below are some interesting examples with respect to different committee size c and the round complexity R.

– For committee size $c = 1$, i.e., leader election, and round complexity $R = O(\log^* n)$, our protocol achieves $\Theta(1)$-game-theoretic fairness against a coalition of size $\Theta(n)$ assuming $\log * n$ is a constant;
– For committee size $c = 1$, i.e., leader election, and round complexity $R = \log\log\log n$, out protocol achieves $(1 - \frac{1}{\mathsf{poly}\log\log n})$-fairness against a coalition of size $n - \frac{n}{\Theta(\log\log\log n)}$.
– For committee size $c = \mathsf{poly}\log\log n$ and for constant round complexity $R = \Theta(1)$, our protocol achieves $(1 - \frac{1}{\mathsf{poly}\log\log n})$-fairness against $\Theta(n)$-sized coalition.

In this paper, we consider the standard notions of approximate CSP-fairness and maximin-fairness. The standard notion of approximate CSP-fairness is also sometimes referred to as *approximate coalition-resistant Nash equilibrium* in some earlier works such as Fruitchain [PS17]. It is also known [CCWS21] that the standard notion of approximate CSP-fairness (or maximin-fairness) is equivalent in some sense to approximate notions of fairness formulated by the more classical Rational Protocol Design (RPD) paradigm [GKM+13, GTZ15, GKTZ15].

Although the standard notion of approximate fairness seems the most natural one, Chung et al. [CCWS21] pointed out that when defining approximate fairness, one can in fact adopt a strengthened notion which they call sequential fairness. Their game-theoretically fair leader election result is in fact stated for the sequential notion. In this sense, our result is incomparable to theirs: they consider a stronger solution concept but their approach inherently cannot give any meaningful result for protocols of $o(\log\log n)$ rounds. By contrast, we consider the more standard non-sequential notion and we are able to generalize the smooth tradeoff between efficiency and fairness shown by Chung et al. [CCWS21] to a broader range of parameters.

1.2 Additional Related Work

Game theory meets cryptography. Some recent efforts have instigated the intersection of the game theory [Nas51, Aum74] and multi-party computation [GMW19, Yao82]. See [Kat08, DR+07] for a survey. There have been two classes of questions that have attracted a lot of interests.

Some work [HT04, KN08, ADGH06, OPRV09, AL11, ACH11] explore how to define game-theoretic notions of security, as opposed to cryptography security notions for distributed computing tasks such as secure function evaluation. Existing works in this line considered a different notion of utility than our work. Their utility functions are often defined assuming that players prefer to compute the function correctly, or prefer to learn others' secret data and prefers that other players do not gain knowledge about their own secrets. Garay et al. propose a paradigm called Rational Protocol Design [GKM+13] and develop this paradigm in subsequent works [GTZ15, GKTZ15]. As mentioned in Section 1, the standard notion of approximate CSP-fairness (or maximin fairness) is in some sense equivalent to the approximate notion of fairness formulated in RPD paradigm.

Another line of work explores how cryptography can help traditional game theory. Many works in game theory assumed the existence of a trusted mediator, which can be realized under cryptography [DHR00, IML05, GK12, BGKO11].

Recently, there has been renewed interest in the connection between game theory and cryptography. Besides the work of Chung et al. [CCWS21] that inspires our work, and [GGS] that generalized the lower bound of the round complexity of game-theoretically fair leader election, the recent work [CGL+18, WAS22] have also suggested game-theoretically fair multi-party binary-coin toss. Binary-coin toss considers tossing a binary coin among n players, while in leader election, we consider tossing an n-way coin among n players. These two formulations are different and they exhibit starkly different theoretical landscape.

Leader election in other models. Leader election has been studied extensively. A line of work [BK14, ADMM14] considered how to achieve "financially-fair" n-party lottery over cryptocurrencies. Their game-theoretic notion of fairness is similar to ours, yet they rely on collateral and penalty mechanisms to achieve fairness. As a comparison, our fairness can be achieved without relying on additional assumptions such as collateral and penalty. Moreover, [ADGH06] studied an incomparable game-theoretic notion for leader election. In their notions, all users prefer to have a leader, and users may have different preferences of who the leader is.

Besides, leader election was considered in the full information model [RZ01, RSZ02, Fei99, Dod06]. Their notion of security concentrates on electing an honest leader with some *small constant* probability, assuming honest majority [Fei99]. This notion is much weaker than the game-theoretic notion considered in our work, which are more suitable in some decentralized applications, where honest majority assumption is not applicable. Moreover, in the full-information model, leader election is impossible against a majority coalition even under this weak notion of security. Interestingly, our committee election protocol actually builds on Feige's lightest bin protocol [Fei99].

Approximate strong fairness. As mentioned in Section 1, the *de facto* notion of fairness considered in the multi-party computation literature is strong fairness or unbiasability. The celebrated result of Cleve [Cle86] showed that it is not possible to achieve $\Omega(\frac{1}{R})$-unbiasable coin toss against a coalition consisting of half or more players. Moran et al. [MNS09] showed how to obtain an R-round protocol that achieves $\Omega(\frac{1}{R})$-unbiasability in the two-party setting, that matches Cleve's lower bound. Recent work [AO16, BOO10, HT14] have been making encouraging progress on building fair multi-party coin toss. However, they rely on constant number of players to ensure polynomial round complexity. We cannot directly rely on multi-party unbiasable coin toss to build game-theoretically fair leader election because our trade-off curve between round complexity and the fairness slack ϵ is exponentially better than that of the unbiasability.

2 Technical Roadmap

2.1 Electing Poly-logarithmically Sized Committees: Achieving CSP-Fairness

We start by observing that a single iteration of Feige's lightest-bin protocol [Fei99] can elect a committee of size $c \geq \mathsf{poly}\log n$ while satisfying *CSP-fairness against relatively large coalitions*. Feige's ingenious protocol works as follows (we describe a single iteration of the protocol): each player $i \in [n]$ chooses a random bin b_i among a total of $B = n/c$ bins, and broadcasts its choice b_i. At this moment, we identify the lightest bin, and everyone who has placed itself in the lightest bin is elected as a committee member. A simple analysis shows that this protocol satisfies CSP-fairness against relatively large coalitions. Specifically, the lightest bin cannot exceed a capacity of $c = n/B$. Moreover, applying the standard Chernoff bound and the union bound, we know that with probability at least $1 - n \cdot \exp(-\Omega(\epsilon^4 \cdot c))$, a good event that every bin has at least $(1-\epsilon^2) \cdot (1-\beta) \cdot c$ honest players must happen, where $\beta \cdot n$ is the maximum coalition size for $\beta \in (0,1)$. Now we show that if the coalition has size larger than $\epsilon \cdot n$, then Feige's lightest bin is $(1-\Theta(\epsilon))$-CSP-fair. Given that the good event happens, the expected fraction of corrupted players in the committee is at most $1-(1-\epsilon^2)\cdot(1-\beta) \leq \frac{\beta}{1-2\epsilon}$. For large n, it is easy to see that the good event happens with $1 - \mathsf{negl}(n)$ probability and the expected fraction of coalition in the committee is at most $\frac{\beta}{1-\Theta(\epsilon)}$. For small n, however, the calculation is more involved, as we will describe below. The overall expected fraction of the coalition in the committee is at most $\frac{\beta}{1-2\epsilon} + \delta$, where $\delta = n \cdot \exp(-\Omega(\epsilon^4 \cdot c))$ is the probability that the good event does not happen. To guarantee that the expected fraction of the coalition in the committee is at most $\frac{\beta}{1-\Theta(\epsilon)}$, we need the failure probability $\delta \leq \beta \cdot \Theta(\epsilon)$. The expected fraction of the coalition in the committee is thus $\frac{\beta}{1-2\epsilon} + \delta \leq \beta(\frac{1}{1-2\epsilon} + \Theta(\epsilon)) \leq \frac{\beta}{1-\Theta(\epsilon)}$. For example, if we pick $\epsilon = \frac{1}{\log n}$ and $c = (\log n)^{10}$, then the probability that the good event does not happen is at most $n \exp\{-\Omega((\log n)^6)\} \leq \epsilon^2 \leq \beta \cdot \epsilon$ for any $n \geq 3$. Henceforth the protocol satisfies $(1-\Theta(\epsilon))$ -CSP-fairness as long as the coalition contains at least ϵn players.

Unfortunately, the protocol does not satisfy CSP-fairness for small coalitions. For example, a single individual $i \in [n]$ (i.e., a coalition of size 1) can examine all others' bin choices and then decide to place itself in the lightest bin. In this case, if the lightest bin (not counting player i) is at least 2 lighter than the second lightest bin, player i is elected into the committee. This happens with a probability at least $\frac{6}{5} \cdot \frac{c}{n}$ for large n , which is significantly higher than the normal probability c/n that player i ought to be elected in an all-honest execution.

Commit-and-reveal lightest bin. We introduce commit-and-reveal version of Feige's lightest bin protocol which achieves CSP-fairness not just against large coalitions, but also against small coalitions as well. The idea is quite simple — below we describe the scheme assuming ideal commitments, although in our formal technical sections we will instantiate the commitments using standard non-malleable commitments. Everyone first commits to a random bin number among $B = n/c$ bins. They then open their commitments. Those who land in the lightest bin are declared the committee, and like before, anyone who fails to commit or correctly open is kicked out. Using the same argument as before, we can show that the commit-and-reveal lightest bin protocol also achieves $(1 - \Theta(\epsilon))$-CSP-fairness against coalitions of size at least ϵn .

We now argue why it also satifies CSP-fairness against small coalitions of size $\beta n < \epsilon n$. Intuitively, the coalition's best strategy is to pick a bin with the fewest number of honest players (henceforth called the *honest-lightest* bin), and place as many coalition members in it as possible while still maintaining that it is the lightest. However, the coalition does not know which one is the honest-lightest bin when committing to its own bin choices. In fact, even when conditioned on the coalition's view during the commitment phase, each bin is the honest-lightest bin with equal probability. No matter how the coalition spreads its members across the bins, the expected number of coalition members in a *randomly chosen* bin is at most $\beta \cdot n/B = \beta \cdot c$. Further, with $1 - n \cdot \exp(-\Omega(\epsilon^4 \cdot c))$ probability, the good event that honest-lightest bin should have at least $(1 - \epsilon^2)(1 - \beta)c$ honest players happens. Therefore, the coalition's expected representation on the committee cannot exceed $\frac{\beta}{(1-\epsilon^2)(1-\beta)} \leq \frac{\beta}{1-2\epsilon}$ given that the good event happens. Overall, the expected fraction of the coalition in the committee is at most $\frac{\beta}{1-2\epsilon} + \delta$, where $\delta = n \cdot \exp(-\Omega(\epsilon^4 \cdot c))$ is the probability that the good event does not happen. Still, as long as $\delta \leq \beta\epsilon$, by the same analysis as before, the expected fraction of the coalition in the committee is at most $\frac{\beta}{1-\Theta(\epsilon)}$.

2.2 Electing Poly-logarithmically Sized Committees: Achieving Maximin Fairness

Although simple and cute, the commit-and-reveal lightest bin protocol does not satisfy maximin fairness. For example, a $\Theta(n)$-sized coalition can target a victim player $i \in [n]$ and prevent it from being elected with high probability using the following strategy. During the commitment phase, spread the coalition members evenly across all bins. During opening, first observe which bin (denoted b^*) player i lands in. Then, all coalition members fail to open except those whose choice was b^*.

To achieve maximin fairness, we are inspired by a virtual identity technique originally proposed by Chung et al. [CCWS21], but unfortunately, directly applying this idea to the lightest bin does not work. At a high level, a strawman idea is as follows:

1. Every player $i \in [n]$ selects a random virtual identity v_i from a sufficiently large space, and commits to the pair (i, v_i).
2. Every player $i \in [n]$ selects a random bin b_i among $B = n/c$ bins, and commits to the pair (v_i, b_i) where v_i is its secret virtual identity.
3. Everyone $i \in [n]$ opens their commitment of (v_i, b_i). The virtual identities contained in the lightest bin will be elected committee.

4. Everyone opens their real-virtual identity mapping (i, v_i). This will allow everyone to compute the real identities of those elected to the committee.

Now, as long as the coalition does not know an honest player i's virtual ID, it does not know who to target during the commit-and-reveal lightest bin steps (Steps 2 and 3). Therefore, as long as the good event that each bin contains at least $(1-\epsilon)(1-\beta)c$ honest players happens, an honest player i will be elected into the committee with probability at least $\frac{(1-\epsilon)(1-\beta)c}{(1-\beta)n} = \frac{(1-\epsilon)c}{n}$. By law of total probability, the probability that an honest player i gets elected into the committee with probability at least $\frac{(1-\epsilon)(1-\delta)c}{n}$, where $1-\delta$ is the probability that the good event happens. Henceforth, as long as $\delta \leq \epsilon$, an honest player i gets elected into the committee with probability at least $\frac{(1-\Theta(\epsilon))c}{n}$.

Unfortunately, this idea does not work if the coalition can eavesdrop on the network channel and observe who sent which (bin, virtual ID) pair in the commit-and-reveal lightest bin protocol. This would allow the coalition to immediately learn the correspondance between virtual and real identities.

To salvage this idea, our high-level idea is simple but realizing it turns out to be somewhat subtle as we explain later. First, if we are willing to assume the existence of an idealized anonymous communication network where players can post messages anonymously, then we can overcome the aforementioned problem by running Steps 2 and 3 over an anonymous communication network. Therefore, it suffices to find a suitable anonymous communication protocol to realize anonymous communication. Although anonymous communication has been extensively studied in the literature [Cha81,Cha88,Abe99,CGF10,DMS04,SGR99,ZZZR05], in our setting, it is tricky to adopt existing schemes directly. The main technicality is that in the presence of a majority coalition, we cannot guarantee the liveness of the anonymous communication protocol.

To overcome this problem, one naïve idea is to rely on an anonymous communication protocol with identifiable abort, and if the protocol fails, we kick out an offending player and retry. Unfortunately, the vanilla notion of identifiable abort does not work for us because we cannot afford to kick out offending players one by one since we are aiming for small round complexity. Our idea is to devise an anonymous communication protocol not just with identifiable abort, but with with *plentiful identifiable aborts*. In other words, if the protocol fails, we want to kick out sufficiently many players, such that we can eventually succeed without too many retries.

Therefore, we adapt an anonymous communication protocol inspired by DC-nets [Cha88] to achieve such a plentiful identifiable abort notion. Assuming an upper bound of βn on the coalition size, our protocol kicks out at least $(1-\beta)n$ players in the event of failure. Thus the round complexity is at most $\frac{1}{1-\beta}$. For example, if $\beta = 99\%$, we can still succeed in $O(1)$ rounds.

We present a formal description and proof of our anonymous communication protocol in Section 6.2 in supplementary materials. We give a formal description of our poly-logarithmically-sized committee election protocol and prove its security in Section 4.

2.3 Leader Election

Although the lightest bin protocol via anonymous broadcast (denoted as LBin-V below) achieves CSP-fairness and maximin-fairness simultaneously, it cannot be directly used to select a leader, i.e., $c = 1$. Indeed, the fairness of LBin-V depends on the occurrence of the good event that each bin has at least $(1-\epsilon^2)(1-\beta)c$ number of honest players, where $\beta \cdot n$ is the maximum coalition size for $\beta \in (0, 1)$. If we are to choose a leader directly using LBin-V, then the probability that this good event happens is 0, which makes our protocol unfair.

To construct a leader election protocol, we compose the committee election LBin-V for multiple iterations. In each iteration: we choose a log-sized committee. In the first iteration we choose a $\mathsf{poly}\log$-sized committee $\mathcal{C}_1$, and then in the second iteration we choose a $\mathsf{poly}\log\log$ sized committee $\mathcal{C}_2$ from $\mathcal{C}_1$, and so on. As analyzed earlier, each iteration of LBin-V is $(1-\Theta(\epsilon))$-game-theoretically fair given that the failure probability δ that the good event does not happen in this iteration is small compare to $\beta \cdot \epsilon$.

However, as the committee size becomes smaller in each iteration, the probability that the good event does not happen becomes larger. In the last few rounds, when the committee becomes constant size, the probability that the good event does not happen becomes a constant. Therefore, we need to cut off at some point and instead run the "almost perfect" tournament tree protocol. As shown in Chung et al. [CCWS21], the tournament tree protocol among c players chooses a leader in $O(\log c)$ rounds and is $(1-\mathsf{negl})$-game-theoretically fair. If we want to achieve a round complexity of R, then we can stop running LBin-V when the committee size becomes smaller than $2^{\Theta(R)}$ and run the tournament tree protocol among the committee to elect a leader.

Now suppose that we run L iterations of committee election LBin-V and get a committee of size $2^{\Theta(R)}$. Then we need to guarantee that the round complexity of these L iterations of LBin-V is at most $O(R)$. By the analysis above, if we kick out $(1-\beta)n$ players in each anonymous communication protocol, the round complexity of each LBin-V is at most $\frac{1}{1-\beta}$. This requires that the fraction of coalition $\beta \leq 1 - \frac{L}{\Theta(R)}$.

Now since the probability that the good event does not happen increases in each iteration, the probability that there is an iteration in which the good event does not happen is dominated by $L \cdot \delta_L$, where $\delta_L = \exp\{-\epsilon^4 \cdot 2^{-\Theta(R)}\}$ is the probability that good event does not happen in the last iteration. As long as this probability is smaller than $\beta \cdot \epsilon$, the protocol is $(1-\Theta(\epsilon))$-fair. Picking $\epsilon = \frac{1}{2^R}$ suffices. Therefore, if we run LBin-V multiple iterations to elect a committee $\mathcal{C}$ of size is $2^{\Theta(R)}$, and then run the tournament tree protocol among $\mathcal{C}$ to elect a leader, our leader election protocol achieves $(1-\frac{1}{2^{\Theta(R)}})$-game-theoretic fairness.

In Section 5, we give a generalized protocol that combines multiple iterations of LBin-V and the tournament tree protocol to elect an arbitrary-sized committee, including the special case of committee size 1, i.e., leader election.

3 Preliminaries

Notation. Throughout, we use λ to denote the security parameter. The notation $\log^{(\ell)} n$ means taking logarithm ℓ times over n. For example, $\log^{(3)} n \equiv \log\log\log n$. Moreover, we use $\log^* n$ to denote the smallest integer ℓ such that $\log^{(\ell)} n \leq 1$. For an event E, we denote $\overline{E}$ as the event that E does not happen. For a vector X of length M, we use $X[j]$ for $j \in [M]$ to denote the j-th element of X. By *t-out-of-n SS*, we refer to a Shamir secret sharing protocol in which any $t+1$ players can reconstruct the secret, while any t players know nothing about the secret [Sha79]. We use the acronym p.p.t. for non-uniform probabilistic polynomial time. We use $\{X_\lambda\}_\lambda \equiv_c \{Y_\lambda\}_\lambda$ to denote that two distribution ensembles $\{X_\lambda\}_\lambda$ and $\{Y_\lambda\}_\lambda$ are computationally indistinguishable, i.e., for all non-uniform p.p.t. $\mathcal{A}$, there exists a negligible function $\mathsf{negl}(\cdot)$, such that for any $\lambda \in \mathbb{N}$, $|\Pr[x \xleftarrow{\$} X_n, \mathcal{A}(x) = 1] - \Pr[y \xleftarrow{\$} Y_n, \mathcal{A}(y) = 1]| < \mathsf{negl}(\lambda)$.

3.1 Probability Tools

Lemma 3.1 (Chernoff bound, Corollary A.1.14 [AS16])**.** *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables. Let $\mu = \mathbb{E}\left[\sum_{i=1}^n X_i\right]$. Then, for any $\epsilon \in (0,1)$, it holds that*

$$\Pr\left[\sum_{i=1}^n X_i \leq (1-\epsilon)\mu\right] \leq e^{-\epsilon^2\mu/2}.$$

3.2 Fairness Notions for Committee Election

Since a leader is a special case of a 1-sized committee, we will define correctness and fairness with respect to committee election protocol.

In a (c,n)-committee election protocol, n players interact through pairwise private channels and a public broadcast channel. We assume that each player has identity $1, 2, \ldots, n$, respectively. We assume that all communication channels are authenticated, i.e., messages carry the sender's identity. Moreover, the network is synchronous, and the protocol proceeds in rounds.

The protocol execution is parametrized with the security parameter λ. We assume that the coalition (adversary) A performs a *rushing* attack. In every round r, it waits for all honest players (those not in A) to send messages in round r and decide what messages the players in the coalition send in round r. At the end of the committee election, the protocol outputs a set of at most c players called the *committee*. The output is defined as a deterministic, polynomial-time function over all *public messages posted to the broadcast channel*. Since we assume that all players wish to be selected into the committee, the utility function we consider is as follows: each player elected into the committee gains a utility of 1, while everyone else gains a utility of 0. If all players behave honestly, the committee is chosen uniformly at random.

Correctness. We say that a (c, n)-committee election protocol is correct, if in an all honest execution, every subset $C \subset [n]$ of size c has an equal probability of being elected as the committee, where the probability is taken over the randomness of (an honest execution) the protocol.

For the fairness notion, we recall the definitions proposed by Chung et al. [CCWS21]. The first notion of fairness (CSP-fairness) protects against a malicious coalition from increasing its utility. The second notion (maximin-fairness) protects against a malicious coalition from decreasing the utility of any honest party. Each of these notions is natural and useful on its own, and in some sense, they complement each other. A protocol that satisfies both simultaneously is called *game-theoretically fair.*

Approximate CSP-fairness. The CSP-fairness requires that no coalition can increase its own expected utility by more than a $(1-\epsilon)$ multiplicative factor, no matter how it deviates from the honest protocol.

Definition 3.2 ($(1-\epsilon)$-CSP-fair committee election)**.** *A (c, n)-committee election is $(1-\epsilon)$-CSP-fair against a non-uniform probabilistic polynomial time (p.p.t.) coalition A of size βn, iff no matter what strategy A adopts,*

$$\mathbb{E}[\widetilde{\beta}] \leq \frac{\beta}{1-\epsilon},$$

where $\widetilde{\beta}$ is the fraction of players in the coalition among the committee, where the expectation is taken over the randomness of the protocol.

In our proof, we will also make use of another fairness notion:

Definition 3.3 ($(1-\epsilon, \delta)$-CSP-fair committee election)**.** *A (c, n)-committee election is $(1-\epsilon, \delta)$-CSP-fair against a non-uniform probabilistic polynomial time (p.p.t.) coalition A of size βn, if there exists an event* GOOD*, where* $\Pr[\mathsf{GOOD}] \geq 1-\delta$*, such that no matter what strategy A adopts,*

$$\mathbb{E}[\widetilde{\beta} \mid \mathsf{GOOD}] \leq \frac{\beta}{1-\epsilon},$$

where $\widetilde{\beta}$ is the fraction of the coalition's representation in the committee, and the expectation is taken over the randomness of the protocol.

Analogously, we define $(1-\epsilon)$-maximin-fair and $(1-\epsilon, \delta)$-maximin-fair committee election, which requires that the probability that an honest individual gets into the committee is large enough given that the good event happens.

Approximate maximin-fairness. Maximin-fairness requires that no coalition can harm any honest individual by more than a $(1-\epsilon)$ multiplicative factor, no matter how it deviates from the honest protocol.

Definition 3.4 ($(1-\epsilon)$-maximin-fair committee election)**.** *A (c, n)-committee election is $(1-\epsilon)$-maximin-fair against a non-uniform probabilistic polynomial time (p.p.t.) coalition A of size βn, iff for any honest individual i, the probability that i gets into the committee is*

$$\Pr[i \textit{ is in the committee}] \geq \frac{(1-\epsilon)c}{n},$$

no matter what strategy A adopts. The probability is taken over the randomness of the protocol.

Definition 3.5 ($(1-\epsilon, \delta)$-maximin-fairness)**.** *A (c, n)-committee election is $(1-\epsilon, \delta)$-maximin-fair against a non-uniform probabilistic polynomial time (p.p.t.) coalition A of size βn, if there exists an event GOOD, where $\Pr[\mathsf{GOOD}] \geq 1-\delta$, such that no matter what strategy A adopts,*

$$\Pr[i \textit{ is in the committee} \mid \mathsf{GOOD}] \geq \frac{(1-\epsilon)c}{n},$$

for any honest individual i. The probability is taken over the randomness of the protocol.

Although committee election is a constant-sum game, these two notions of fairness are non-equivalent. See Section **??** for more explanation.

Finally, we define *game-theoretical fairness*. This notion of fairness requires CSP and maximin-fairness simultaneously.

Definition 3.6 ($(1-\epsilon)$-game-theoretical fairness)**.** *A (c, n)-committee election is $(1-\epsilon)$ game-theoretically fair committee election against a non-uniform probabilistic polynomial time (p.p.t.) coalition A, iff it is $(1-\epsilon)$-CSP-fair and $(1-\epsilon)$-maximin-fair against A.*

Definition 3.7 ($(1-\epsilon, \delta)$-game-theoretical fairness)**.** *A (c, n)-committee election is $(1-\epsilon)$ game-theoretically fair committee election against a non-uniform probabilistic polynomial time (p.p.t.) coalition A, iff it is $(1-\epsilon, \delta)$-CSP-fair and $(1-\epsilon, \delta)$-maximin-fair against A.*

By definition, a $(1-\epsilon)$-game-theoretically fair committee election is also $(1-\epsilon, 0)$-game-theoretically fair. Next we give the translation from $(1-\epsilon, \delta)$-CSP/maximin-fair to $(1-\epsilon)$-CSP/maixin-fair.

Lemma 3.8. *Let n be the number of parties and fix a parameter c. Let CElect be an R-round $(1-\epsilon, \delta)$-CSP-fair (c, n)-committee election protocol against a coalition of size βn. Then the above leader election protocol is $(1-\epsilon_1)$-CSP-fair against a coalition of size βn, with a round complexity $R + O(\log c)$, where*

$$\epsilon_1 = \frac{\beta\epsilon + \delta(1-\epsilon)}{\beta + \delta(1-\epsilon)} + \mathsf{negl}(\lambda).$$

Lemma 3.9. *Let n be the number of parties and fix a parameter c. Let* CElect *be an R-round $(1-\epsilon, \delta)$-maximin-fair (c, n)-committee election protocol against a coalition of size βn. Then the above leader election protocol is $(1-\epsilon_2)$-maximin-fair, with a round complexity $R + O(\log c)$, where*

$$\epsilon_2 = \epsilon + \delta + \mathsf{negl}(\lambda).$$

The proofs of these two lemmas are available in the full version.

Hybrid vs. real worlds. For ease of presentation and modulatiry purposes, we shall sometimes consider protocols in a hybrid setting where we assume some "generic" functionality is given for free. This is called a "hybrid world". That is, we say that a protocol is in the $\mathcal{F}$-hybrid world if players interacting in this protocol have access to an ideal functionality $\mathcal{F}$. A protocol in the (plain) real world is a protocol without any ideal functionalities or setup assumptions. Specifically for us, we say that a (c, n)-committee election protocol achieves $(1-\epsilon)$-game-theoretic fairness against a coalition A in the $\mathcal{F}$-hybrid world, if the protocol achieves $(1-\epsilon)$-game-theoretic fairness against this coalition A, assuming the ideal functionality $\mathcal{F}$.

3.3 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, the committer C, the receiver R, and a deterministic, polynomial-time public verifier V. We assume that the protocol has two phases, a commitment phase and an opening phase. The public verifier, upon receiving a transcript $\varGamma$ of the commitment protocol, outputs either a bit $b \in \{0, 1\}$ to accept or $\perp$ to reject. We use $\langle \mathsf{C}^*(z), \mathsf{R}^*(z') \rangle$ to denote an execution between C^* on input $z, 1^\lambda$, and R^* on input $z', 1^\lambda$, where λ is the security parameter.

Correctness. Correctness guarantees that an honest committer always completes the protocol and correctly opens its input bit; and will not be stuck by a malicious, non-aborting receiver. Formally, for $b \in \{0, 1\}$, for any $\lambda \in \mathbb{N}$, if C is honest and receives input bit b, then $\langle \mathsf{C}(z), \mathsf{R}^*(z') \rangle$ will complete with the accepting bit b with probability 1, for any non-aborting R^*. If the messages sent by R^* are outside the valid range, it is treated as aborting.

Perfect Binding. Perfect binding guarantees that the commitment phase will determine only one bit that can be successfully opened. Formally, let $(\varGamma_c, \varGamma_o) \in \{0, 1\}^{\ell(\lambda)}$ be the transcripts of the commitment phase and the opening phase, respectively, where $\ell(\lambda)$ is a fixed polynomial function denoting the maximum length of the transcripts. Then for any $\lambda \in \mathbb{N}$, any transcripts $\varGamma_c, \varGamma_o, \varGamma'_o$, if $\mathsf{V}(1^\lambda, \varGamma_c, \varGamma_o) = b$ and $\mathsf{V}(1^\lambda, \varGamma_c, \varGamma'_o) = b'$, where $b, b' \in \{0, 1\}$, it must be that $b = b'$.

Computationally Hiding. Computationally hiding guarantees that at the end of the commitment phase, the receiver learns only a negligible amount of information about the input that the committer commits to. Formally, let $p_\lambda(v)$ denote the probability that R^* outputs 1 at the end of the commitment phase in an execution $\langle \mathsf{C}^*(1^\lambda, v), \mathsf{R}^*(1^\lambda)\rangle$, then for any non-uniform p.p.t. R^*, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$ and every $v_1, v_2 \in \{0,1\}^\lambda$, it holds that $|p_\lambda(v_1) - p_\lambda(v_2)| \le \mathsf{negl}(\lambda)$.

Concurrent Non-malleability. We follow the definition of Lin et al. [LPV08]. Consider a man-in-the-middle adversary A that participate on the left m interactions with an honest committer who runs commitment phase committing to values $v_1, \ldots, v_m$ with identity $\mathsf{id}_1, \ldots, \mathsf{id}_m$, and on the right m interactions with an honest receiver trying to commit to values $v'_1, \ldots, v'_m$ with identity $\mathsf{id}'_1, \ldots, \mathsf{id}'_m$. If any of the right commitments are invalid its value is set to $\perp$. For every $i \in [m]$, if $\mathsf{id}'_j = \mathsf{id}_i$ for some $j \in [m]$, then v'_j is set to be $\perp$. Let $\mathsf{mitm}^{\mathcal{A}}(1^\lambda, v_1, v_2, \ldots, v_m, z)$ denote the view of $\mathcal{A}$ and the values $v'_1, \ldots, v'_m$.

Definition 3.10. *A commitment scheme is concurrent non-malleable if for every polynomial $p(\cdot)$, for every non-uniform p.p.t. adversary $\mathcal{A}$ that participates in at most $m = p(\lambda)$ concurrent executions, there exists a polynomial time simulator $\mathcal{S}$ such that*

$$\{\mathsf{mitm}^{\mathcal{A}}(1^\lambda, v_1, v_2, \ldots, v_m, z)\}_{v_1,\ldots,v_m \in \{0,1\}, z \in \{0,1\}^*, \lambda \in \mathbb{N}} \equiv_c$$
$$\{\mathcal{S}(1^\lambda, z)\}_{v_1,\ldots,v_m \in \{0,1\}, z \in \{0,1\}^*, \lambda \in \mathbb{N}}.$$

Theorem 3.11 ([LPV08]). *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.*

In this paper, we will only consider bounded concurrency. Without loss of generality, the number of concurrent calls to public verifiable concurrent non-malleable commitment in our protocol is upper bounded by n^2, where n is the number of players.

4 Game-Theoretically Fair Committee Election

In this section, we present our game-theoretically fair committee election that extends Feige's lightest bin protocol. Later, in Section 5, we will use it as a building block to get our committee election protocol that achieves game-theoretic fairness for arbitrary committee size.

4.1 Electing Poly-logarithmically Sized Committees: Achieving CSP-Fairness

In this section, we give a CSP-fair committee election protocol. This is the first step towards our game-theoretically fair committee election (that needs to be CSP-fair and maximin fair, simultaneously).

Our CSP-fair protocol is a commit-and-reveal variant of Feige's well-known lightest bin protocol [Fei99]. Specifically, we require all parties to (cryptographically) commit to their bin choices and only afterward to reveal their choices. The parties whose choices correspond to the lightest bin are the committee. The commitments that we use are *interactive*. To commit to a string, a player invokes n instances of NMC, one for each of the n receivers. To open the commitments, the committer posts the openings for all n instances in the broadcast channel, and the opening is correct iff all of the n instances are correctly opened to the same string. Without loss of generality, we assume that the committer only needs to send one message in the opening phase. Moreover, we assume that messages are posted to the broadcast channel, and it can be checked publicly if a commitment is correctly opened. This is why we also require public verifiability of the commitment scheme. We say that a player fails to commit if the player fails to commit in an instance, where the receiver is non-aborting.

$\mathsf{LBin\text{-}C}$: **Commit-and-Reveal Lightest Bin**

Parameters: Let c be an upper bound of the size of the required committee and n is the number of players. Fix $B = \lceil \frac{n}{c} \rceil$ as the number of bins. For simplicity, we assume c divides n.

Building blocks: A publicly verifiable concurrent non-malleable commitment as in Section 3.3, NMC.

Protocol:

1. Round 1: Every player i randomly chooses a bin $b_i \in [B]$, invokes n NMC instances and run the commit phase with n receivers to commit to b_i. The messages are sent in a broadcast channel. Exclude those players who fail to commit.
2. Round 2: Every player i runs the opening phase with n receivers to open its bin choice b_i. Exclude those players who fail to open all n instances correctly.
3. Let $\widehat{b}$ be the lightest bin after exclusion (break ties with lexicographically the smallest bin). The players who choose bin $\widehat{b}$ constitute the committee.

Theorem 4.1. *Assume that* NMC *is publicly verifiable concurrent non-malleable commitment as in Section 3.3. For* $n, c \in \mathbb{N}$, $\epsilon \in (0, 1/2)$, *and* $\beta \in (0, 1)$, *the protocol* $\mathsf{LBin\text{-}C}$ *is a constant round* $(1-2\epsilon, \delta)$*-CSP-fair* (c, n)*-committee election protocol against a coalition* $\mathcal{K}$ *of size* βn, *where*

$$\delta = \frac{n}{c} \exp\left\{ -\frac{\epsilon^4}{2}(1-\beta)c \right\}. \tag{1}$$

Proof. Fix n, c, ϵ, and β as in the statement. Define GOOD to be the event that each bin has at least $(1-\epsilon^2)(1-\beta)c$ honest players. Let $\widetilde{\beta}$ denote the fraction of players in $\mathcal{K}$ among the committee. Then, we have the following lemma.

Lemma 4.2. $\mathbb{E}\left[\widetilde{\beta} \mid \mathsf{GOOD}\right] \le \frac{\beta}{1-2\epsilon}$.

For now assume that Lemma 4.2 holds and we explain why Theorem 4.1 follows from it. The proof of Lemma 4.2 appears right afterwards. By Chernoff bound (Lemma 3.1) and the union bound,

$$\Pr\left[\mathsf{GOOD}\right] \ge 1 - \frac{n}{c}\exp\left\{-\frac{\epsilon^4}{2}(1-\beta)c\right\}. \tag{2}$$

Combing Lemma 4.2 and (2), LBin-C is a $(1-2\epsilon,\delta)$-CSP-fair committee election protocol by Definition 3.5. □

Proof sketch of Lemma 4.2 We split into two cases. First, assume that $\beta \ge \epsilon$. In this case, the claim follows directly from the assumption that GOOD holds: The fraction of players in $\mathcal{K}$ among the committee must satisfy $\widetilde{\beta} \le 1-(1-\epsilon^2)(1-\beta) = \beta\left(1+\frac{\epsilon^2}{\beta}-\epsilon^2\right) \le \frac{\beta}{1-2\epsilon}$ as required.

Now, we focus on the case where $\beta < \epsilon$. By the perfect binding property, at the end of commit phase, player i's bin choice are fixed. Let $\{b_i\}_{i=1}^n$ denote the bin choices of n players at the end of the commit phase. To compute $\mathbb{E}[\widetilde{\beta} \mid \mathsf{GOOD}]$, we define a random variable γ, which depends only on $\{b_i\}_{i=1}^n$, that upper bounds $\widetilde{\beta}$ in an execution of LBin-C. Let $\widetilde{b} \in [B]$ be the index of the bin that contains least number of honest players; and $b^* \in [B]$ be index of the lightest bin at the end of the commit phase. Note that by the way the protocol works, $\widetilde{b}$ and b^* depends only on $\{b_i\}_{i=1}^n$. Below, for $l \in [B]$, we use h_l to denote the number of honest players in bin l, and f_l to denote the number of players in $\mathcal{K}$ in bin l.

Given the bin choices $\{b_i\}_{i=1}^n$ at the end of the commit phase, the fraction of players in $\mathcal{K}$ among the committee is at most $\gamma := \frac{f_{\widetilde{b}}}{h_{b^*}+f_{b^*}}$. This is because by the perfect binding property and public verifiability of the commitment scheme, the only way the coalition can deviate is essentially to refuse to open some of their bin choices in the opening phase and get excluded at the end of Round 2, in order to change the lightest bin. To maximize the fraction of the coalition in the committee, the best strategy for the coalition is to choose bin $l = \widetilde{b}$, which contains the least number of honest players. Since the number of honest players in bin l is h_l, the fraction of players in $\mathcal{K}$ in bin $l = \widetilde{b}$, after excluding the misbehaved players, is at most $1 - \frac{h_{\widetilde{b}}}{f_{b^*}+h_{b^*}} \le \frac{f_{\widetilde{b}}}{h_{b^*}+f_{b^*}}$.

Therefore, to upper bound $\mathbb{E}[\widetilde{\beta} \mid \mathsf{GOOD}]$, it suffices to bound $\mathbb{E}[\gamma \mid \mathsf{GOOD}]$. Since when GOOD happens, the number of honest players in every bin is at least $(1-\epsilon^2)(1-\beta)c$, we have that

$$\mathbb{E}[\gamma \mid \mathsf{GOOD}] \le \frac{1}{(1-\epsilon^2)(1-\beta)n}\sum_{l=1}^{B}\mathbb{E}\left[f_l \mid \widetilde{b} = l, \mathsf{GOOD}\right].$$

By the non-malleability of the commitment scheme, $\mathbb{E}[\gamma \mid \mathsf{GOOD}]$ in the protocol should be negligibly close to the conditional expectation of γ in an idealized world where the bin choices of the players in $\mathcal{K}$ are independent from the honest players' bin choices, i.e, $\widetilde{b}$ is independent from f_ℓ. Therefore,

$$\mathbb{E}[\gamma \mid \mathsf{GOOD}] = \frac{1}{(1-\epsilon^2)(1-\beta)n} \sum_{l=1}^{B} \mathbb{E}\left[f_l\right] + \mathsf{negl}(\lambda) \le \frac{\beta}{(1-\epsilon^2)(1-\epsilon)} + \mathsf{negl}(\lambda),$$

where the last inequality comes from the assumption that $\beta < \epsilon$. Putting together, the expectation $\mathbb{E}\left[\widetilde{\beta} \mid \mathsf{GOOD}\right]$ in the committee election $\mathsf{LBin\text{-}C}$ is at most $\frac{\beta}{(1-\epsilon^2)(1-\epsilon)} + \mathsf{negl}(\lambda) \le \frac{\beta}{1-2\epsilon}$.

4.2 Electing Poly-logarithmically Sized Committees: Achieving Maximin-Fairness

In Section 4.1 we gave a commit-and-reveal variant of Feige's lightest bin protocol for committee election and showed that it is CSP-fair. The protocol is, however, not maximin-fair. While the adversary cannot gain too much utility by deviating from the protocol, it can still harm the utility of an honest individual. Specifically, consider the following adversarial strategy. The coalition generates commitments so that the coalition's representations in each bin are equal. Then, when it wants to target at a specific player i to not participate in the committee, it waits to see which bin l was chosen by that honest party and then it refuses to reveal commitments from some other bin l' which will then be lighter than the bin l chosen by honest player i. This attack prevents an honest individual i from being elected into the committee.

By the properties of the commitment scheme and how our protocol works, this is the only useful attack for the adversary. Thus, we modify our protocol to withstand this attack by masking the identity of parties. Namely, we hide which bin choice belongs to which party. We achieve this by requiring players to choose a random virtual ID and use it throughout the execution. Players will only reveal their virtual IDs at the end of the protocol, after the lightest bin has been fixed. A-priori, it seems hard to implement such a system because once a party sends its message, everybody knows who sent it (recall that we are in the broadcast model). We overcome this by implementing an "anonymous" broadcast channel on top of our existing broadcast channel.

Thus, we first describe our anonymous broadcast functionality $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$. Then, we show that in a $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$-hybrid model, we can build a committee election protocol that ensures CSP-fairness and maximin-fairness simultaneously.

Anonymous Broadcast Functionality Let $\mathcal{O}$ be the set of all players involving in the protocol. Our anonymous broadcast functionality $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ works as follows.

$\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$: **Anonymous broadcast with t-identifiable abort**

Parameters: $\mathcal{O}$ is the set of players involving in the protocol and t is a bound on the number of misbehaved players to exclude.

Functionality:

1. **Input**: Every player i sends a single message m_i or $\perp$ to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$.
2. **Output**: $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ computes a multiset $\mathsf{Out} = \{m_i : i \in \mathcal{O} \text{ and } m_i \neq \perp\}$. If the number of corrupted players is smaller than t, send $(\mathsf{ok}, \mathsf{Out})$ to everyone in $\mathcal{O}$. Otherwise, send Out to the adversary $\mathcal{A}$.
   - If receives ok from $\mathcal{A}$, $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ sends $(\mathsf{ok}, \mathsf{Out})$ to every honest player in $\mathcal{O}$.
   - Otherwise, it receives a set $\mathcal{D}$ of corrupted IDs of size at least t from the adversary $\mathcal{A}$, and then send $(\mathsf{fail}, \mathcal{D})$ to every honest player in $\mathcal{O}$.

We say that an adversary A is *admissible* if 1) it sends only one message for each corrupt player, and 2) it either sends ok, or a set of corrupted players of size at least t in Step 2.

The functionality exhibits several appealing properties that are important for us. Specifically, in the ideal functionality $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, it holds that:

1. Each player can only send one message.
2. The coalition has to choose their messages independently from honest players' messages.
3. The coalition cannot tell which honest player sends which message.
4. The output is either $(\mathsf{ok}, \mathsf{Out})$, or $(\mathsf{fail}, \mathcal{D})$ with a set $\mathcal{D}$ of size at least t.

Formal Description of the Protocol Here we present the formal description of our lightest bin via anonymous broadcast protocol in the $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$-hybrid model.

$\mathsf{LBin\text{-}V}(c, n, \beta)$: **Lightest Bin via Anonymous Broadcast**

Parameters: Let c be an upper bound of the required committee and n is the number of players. Fix $B = \lceil \frac{n}{c} \rceil$ as the number of bins. For simplicity, we assume c divides n. Let $\mathcal{O}$ be initialized as $[n]$ that denotes the set of active players. $\beta \cdot n$ is the maximum size of the coalition for $\beta \in (0, 1)$.

Building blocks: A publicly verifiable concurrent non-malleable commitment as in Section 3.3, NMC.

Protocol:

1. Every player i randomly chooses a string $v_i \leftarrow \{0,1\}^\lambda$ as its virtual ID, invokes n instances of NMC, and runs the commit phase with n receivers to commit to (i, v_i). Exclude those players who fail to commit.
2. Each player randomly chooses a bin $b_i \leftarrow [B]$ with fresh randomness, and sets $m_i = (b_i, v_i)$. Broadcast m_i using $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ with $t = \lfloor (1-\beta) n \rfloor$.

– If the output is $(\mathsf{fail}, \mathcal{D})$, exclude the players in $\mathcal{D}$ from $\mathcal{O}$ (namely, set $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$). Then, the remaining players (i.e., those in the updated $\mathcal{O}$) re-run step 2.
– If the output is $(\mathsf{ok}, \mathsf{Out})$, go to the next step.

3. Let b^* be the lightest bin. Every player opens its virtual ID (i, v_i). Let U_{b^*} be the set of virtual IDs that are *unique* and choose the lightest bin b^*. Those who open the (i, v_i) successfully with $v_i \in U_{b^*}$ are chosen to be the committee.

Note that in $\mathsf{LBin\text{-}V}$, players do not need to commit to their bin choices and then open, since the functionality $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ guarantees that the malicious coalition has to choose their messages, i.e., bin choices, independently from honest players' messages. In the following theorem we show that the protocol $\mathsf{LBin\text{-}V}$ described above is both maximin-fair and CSP-fair in the $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$-hybrid model.

Theorem 4.3. *Assume that* NMC *is a publicly verifiable concurrent non-malleable commitment as in Section 3.3. For any* $n, c \in \mathbb{N}$ *and* $\epsilon \in (0, 1/2), \beta \in (0, 1)$, *the committee election protocol* $\mathsf{LBin\text{-}V}(c, n, \beta)$ *is a* $(1-\epsilon, \delta)$*-maximin-fair and a* $(1-2\epsilon, \delta)$*-CSP-fair* (c, n)*-committee election* [5] *in the* $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$*-hybrid model, against a coalition* $\mathcal{K}$ *of size* βn, *where*

$$\delta = \frac{2n}{(1-\beta)c} \exp\left\{-\frac{\epsilon^4}{2}(1-\beta)c\right\} + \mathsf{negl}(\lambda).$$

Moreover, the round complexity of $\mathsf{LBin\text{-}V}$ *is at most* $\frac{2}{1-\beta} + 2$.

Proof. Fix n, c, ϵ, and β as in the statement. Let Unique be the event that honest players choose unique virtual IDs, and their virtual IDs do not collide with any players in the coalition. Let GOOD be the event that in every execution of $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ in Step 2, each bin has at least $(1-\epsilon^2)(1-\beta)c$ honest players.

We use the following lemma to prove maximin-fairness and CSP-fairness. The proof to the lemma appears afterward.

Lemma 4.4. $\Pr[\mathsf{Unique}, \mathsf{GOOD}] \geq 1 - \delta$.

<u>Maximin-fairness</u> Let H_i denote the event that an honest player i is chosen into the committee. The claimed maximin-fairness follows from the following lemma. The proof of the lemma appears below.

Lemma 4.5. $\Pr[\mathsf{H}_i \mid \mathsf{Unique}, \mathsf{GOOD}] \geq (1-\epsilon)c/n$.

Combining Lemmas 4.4 and 4.5, we have that $\mathsf{LBin\text{-}V}$ is a $(1-\epsilon, \delta)$-maximin-fair committee election protocol against a coalition of size βn by Definition 3.7.

<u>CSP-fairness</u> Let $\widetilde{\beta}$ denote the fraction of the coalition in the committee. Now, the claimed CSP-fairness follows from the following lemma. The proof of the lemma appears below.

[5] Theorem 4.3 implies that the protocol $\mathsf{LBin\text{-}V}$ is a $(1-2\epsilon, \delta)$-game-theoretic fairness by Definition 3.7.

Lemma 4.6. $\mathbb{E}\left[\widetilde{\beta} \mid \mathsf{GOOD}, \mathsf{Unique}\right] \leq \frac{\beta}{1-2\epsilon}$.

Combining Lemmas 4.4 and 4.6, we have that $\mathsf{LBin\text{-}V}$ is a $(1-2\epsilon, \delta)$-CSP-fair committee election protocol against a coalition of size βn by Definition 3.5. □

Proof sketch of Lemma 4.4 By the non-malleability property of the commitment scheme, $\Pr[\mathsf{Unique}]$ in an real execution of the protocol should be negligibly close to this probability in an idealized world where the virtual IDs of players in $\mathcal{K}$ are chosen independently from honest players' virtual IDs. Therefore, $\Pr[\mathsf{Unique}] = 1 - \mathsf{negl}(\lambda)$. By Chernoff's bound (Lemma 3.1) and the union bound over B bins, in a single execution of $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, each bin contains at least $(1-\epsilon^2)(1-\beta)c$ honest players with probability $p = 1 - \frac{n}{c}\exp\left\{-\frac{\epsilon^4}{2}(1-\beta)c\right\}$.

Each time $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ is invoked, it either outputs ok or wipes out a set of players in the coalition of size at least t. Since $t = \lfloor(1-\beta)n\rfloor$, we will run at most $\frac{\beta n}{\lfloor(1-\beta)n\rfloor} < \frac{2}{1-\beta}$ rounds of $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$. Hence, $\Pr[\mathsf{GOOD}, \mathsf{Unique}] \geq p^{\frac{2}{1-\beta}}(1-\mathsf{negl}(\lambda))$. Lemma 4.4 thus follows.

Proof sketch of Lemma 4.5 In $\mathsf{LBin\text{-}V}$, the players choose their bins in Step 2 with their virtual IDs and broadcast the bin choices using $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$. By the property of the functionality, in each execution of $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, the coalition has to choose their bins independently from honest players' bin choices. If the coalition chooses to fail a call to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, honest players will choose bins with *fresh* randomness in the next call to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$. Therefore, the coalition's strategy S_l of whether to fail the l-th call to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ in Step 2 depends only on the output of the first l calls $\mathsf{Out}_1, \ldots, \mathsf{Out}_l$ to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, and the view $\mathsf{view}_K^{\mathsf{comm}}$ of the coalition $\mathcal{K}$ in Step 1. Still, we use $\mathcal{H}$ to denote the set of honest players, where $|\mathcal{H}| = n - \beta n$.

Let L denote the total number of $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ calls. Now consider the l-th call to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$. Let $\mathsf{H}_{i,j}$ denote the event that honest player i chooses bin j in that $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ call. Since honest players choose their bins independently in different calls to $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, it follows that

$$\Pr\left[\mathsf{H}_{i,j} \mid \mathsf{Out}_1, \ldots, \mathsf{Out}_L, \mathsf{view}_K^{\mathsf{comm}}, S_1, \ldots, S_L\right] = \Pr\left[\mathsf{H}_{i,j} \mid \mathsf{Out}_l, \mathsf{view}_K^{\mathsf{comm}}\right].$$

By the non-malleability and the anonymity of $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$, the map between the honest virtual ID and the honest players' identity remains hidden from the coalition $\mathcal{K}$. For $j \in [B]$, we use V_j to denote the set $V_j = \{v_i : (v_i, j) \in \mathsf{Out}_l\}$, i.e., the set of virtual IDs choosing bin j. Then we have

$$\Pr[\mathsf{H}_{i,j} \mid \mathsf{Out}_l = \{(v_i, b_i)\}_{i\in[n]}, \mathsf{view}_K^{\mathsf{comm}} = v] \geq \frac{h_j}{|\mathcal{H}|} - \mathsf{negl}(\lambda),$$

where h_j is the number of honest players in bin j. Given the assumption that GOOD happens, $h_j \geq (1-\epsilon^2)(1-\beta)c$ for every $j \in [B]$. Let view_K denote the view of the adversary at the end of Step 2, which includes $\mathsf{view}_K^{\mathsf{comm}}$, all the outputs $\mathsf{Out}_1, \ldots, \mathsf{Out}_L$, as well as $\mathcal{A}$'s strategy $S_1, \ldots, S_L$. Then, the lightest bin b^* is deterministic given view_K. For any $i \in \mathcal{H}$ we have

$$\Pr[\mathsf{H}_i \mid \mathsf{GOOD}] = \sum_{j \in [B]} \Pr\left[\mathsf{H}_{i,j}, b^* = j \mid \mathsf{GOOD}\right] \geq \frac{(1-\epsilon^2)c}{n} - \mathsf{negl}(\lambda).$$

Therefore, at the end of Step 2, the probability that an honest player i's virtual ID is in the lightest bin b^* is at least $(1-\epsilon^2)c/n - \mathsf{negl}(\lambda)$. This implies that the honest player i will be elected into the committee with a probability at least $(1-\epsilon^2)c/n - \mathsf{negl}(\lambda) \geq (1-\epsilon)c/n$, given that GOOD and Unique happens.

Proof sketch of Lemma 4.6 The proof to this Lemma is similar to the proof of Lemma 4.2, except that honest players' bin choices are now hidden from the coalition by the anonymous broadcast functionality $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$.

5 Fairness Amplification Though Iteration

This section gives our final game-theoretically fair committee election and leader election protocols to select arbitrary committee size with good fairness parameters. The committee election protocol $\mathsf{LBin\text{-}V}$ introduced in Section 4.2 does not achieve fairness with good parameter for arbitrary committee size. For example, if we want to choose a $\log\log n$-sized committee from n players using $\mathsf{LBin\text{-}V}$, the probability that the GOOD event does not happen is upper bounded by $\frac{n}{\log\log n}\exp\{-\frac{\epsilon^4}{2}\log\log n\}$, which is even larger than 1. This makes $\mathsf{LBin\text{-}V}$ not fair enough for electing a small sized-committee.

Therefore, to build a fair committee election protocol that works for arbitrary committee size, we compose $\mathsf{LBin\text{-}V}$ for multiple iterations, and combine it with the tournament tree protocol if necessary.

We first give the formal description of the tournament tree protocol and its "almost perfect" fairness. Then we give our final committee election protocol that achieves game-theoretic fairness for arbitrary committee size.

5.1 Preliminary: Fairness of Tournament Tree Protocol

This section gives a formal description of the tournament tree protocol.

Tournament tree protocol $\mathsf{Tourn}(\mathcal{O})$

Let n be the size of $\mathcal{O}$.

- If $n = 1$, return the single player in $\mathcal{O}$.
- Otherwise, let $n_1 = \lfloor \frac{n}{2} \rfloor$ and $n_2 = \lceil \frac{n}{2} \rceil$. Let $\mathcal{O}_1$ be the first n_1 players in $\mathcal{O}$ and $\mathcal{O}_2$ be the remaining players.
- In parallel, run $\mathsf{Tourn}(\mathcal{O}_1)$ and $\mathsf{Tourn}(\mathcal{O}_2)$, and denote the output as O_1 and O_2, respectively.

- The final winner is determined by the duel protocol between O_1 and O_2 such that O_i wins with probability n_i/n. This is described below.

Duel Protocol between O_1 and O_2

Let $\frac{k_1}{k_1+k_2}$ and $\frac{k_2}{k_1+k_2}$ be the probability that player O_1 and O_2 wins, respectively.

- Let $k = k_1 + k_2$, and $\ell = \lceil \log k \rceil$. Each player O_i commits to an ℓ-bit random string that represents some $s_i \in \mathbb{Z}_{k-1}$ for $i = 1, 2$.
- Each player O_i opens its commitment and reveals s_i. If $s_1+s_2 \mod k \in \{0, \ldots, k_1 - 1\}$, player O_1 wins. Otherwise, O_2 wins.
- If a player aborts or fails to open the commitment correctly, it is treated as forfeiting and the other player wins.

Lemma 5.1 (Theorem 3.5 of Chung et al. [CCWS21])**.** *Let n be the number of players and λ be the security parameter. Then, the tournament-tree protocol, when instantiated with a suitable publicly verifiable, non-malleable commitment scheme as defined in Section 3.3, satisfies $(1 - \mathsf{negl}(\lambda))$-CSP-fairness and $(1 - \mathsf{negl}(\lambda))$-maximin-fairness against coalition of arbitrarily sizes. Moreover, the round complexity is $O(\log n)$.*

5.2 Our Final Game-Theoretically Fair Committee Election

In this section, we give our fair committee election protocol that works for arbitrary committee size. Our final protocol runs multiple iterations of LBin-V and combines it with the tournament tree protocol if necessary. The $\mathcal{F}^{t,\mathcal{O}}_{\text{anon}}$ ideal functionality in LBin-V can be instantiated in real-world cryptography, with only a constant round blowup. The instantiation will be given in Section 6 in supplementary materials.

Let c be the upper bound of the committee size we want to achieve. The final committee election is given below.

Committee election protocol $\mathsf{CElect}(n, c)$

Parameter: Let c be the upper bound of the committee size and R be the round complexity we want to achieve. The initial committee is $\mathcal{C}_0 = [n]$, $c_0 = n$. The fraction of the coalition is $\beta_0 = \beta$. If $c \geq 2^R$, let $L \leq R$ be the smallest integer such that $\log^{(L)} n \leq c^{0.1}$ and $\epsilon = \frac{1}{c^{0.1}}$; otherwise, set $L \leq R$ be the smallest integer such that $\log^{(L)} n \leq 2^R$ and $\epsilon = \frac{1}{2^R}$.

Protocol

1. For $\ell = 1, \ldots, L - 1$:
   - Let $c_\ell = (\log^{(\ell)} n)^{10}$, $\mathcal{O} = \mathcal{C}_{\ell-1}$, $\beta_\ell = \beta_{\ell-1}(1 - \epsilon^2) + \epsilon^2$.
   - Run $\mathsf{LBin\text{-}V}(c_\ell, \mathcal{C}_{\ell-1}, \beta_{\ell-1})$. That is, we choose a committee $\mathcal{C}_\ell$ of size $c_\ell = (\log^{(\ell)} n)^{10}$ from $\mathcal{C}_{\ell-1}$.

– $\ell = \ell + 1$.

2. If $c \geq 2^R$, set $c_L = c$; otherwise, set $c_L = 2^{11R}$. Run the committee election protocol $\mathsf{LBin\text{-}V}(c_L, \mathcal{C}_{L-1}, \beta_{L-1})$ to elect a committee $\mathcal{C}_L$ of size at most c_L.
3. If $c_L \geq c$, run c number of parallel instances of $\mathsf{Tourn}^{sid}(\mathcal{C}_L)$ for $sid \in [c]$. Let the final committee be the set of elected leaders in these c instances of tournament tree protocol.

Note that in the protocol, β_ℓ is just a parameter that passes to $\mathsf{LBin\text{-}V}$, together with c and $\mathcal{O}$. It is *not* the real fraction of the coalition in committee $\mathcal{C}_\ell$. Instead, it is the upper bound of the real fraction of the coalition in $\mathcal{C}_\ell$ if *good* event happens in each round up to ℓ. The parameter β_ℓ is only used to set the parameter t of $\mathcal{F}^{t,\mathcal{O}}_{\text{anon}}$ in the ℓ-th $\mathsf{LBin\text{-}V}$ call.

Theorem 5.2. *Assume the existence of enhanced trapdoor permutations and collision-resistant hash functions. Fix n and c. Let L^* be the smallest integer such that $\log^{(L^*)} n \leq c$. Then for any $L^* \leq R \leq C_0 \log n$ for some constant C_0, we have that*

– *If $c \geq 2^R$, there exists an $O(R)$-round committee election that achieves $(1 - \frac{1}{c^{\Theta(1)}})$-game-theoretic fairness against a non-uniform p.p.t. coalition of size at most $(1 - \frac{L^*}{\Theta(R)})n$.*

– *If $c < 2^R$, there exists an $O(R)$-round committee election that achieves $(1 - \frac{1}{2^{\Theta(R)}})$-game-theoretic fairness against a non-uniform p.p.t. coalition of size at most $(1 - \frac{L}{\Theta(R)})n$, where L is the smallest integer such that $\log^{(L)} n \leq 2^R$.*

Our final leader election protocol can be gained directly by picking $c = 1$ in Theorem 5.2.

Theorem 5.3. *Assume the existence of enhanced trapdoor permutations, and collision-resistant hash functions. Fix n and let $\log^* n \leq R \leq C \log n$ be the round complexity we want to achieve for some constant C. Then there exists an $O(R)$-round leader election that achieves $(1 - \frac{1}{2^{\Theta(R)}})$-game-theoretic fairness against a non-uniform p.p.t. coalition of size at most $(1 - \frac{L}{\Theta(R)})n$, where L is the smallest integer such that $\log^{(L)} n \leq 2^R$.*

The full proof of Theorem 5.2 and Theorem 5.3 are available in the full version.

6 Instantiation of the Ideal Functionalities

In this section, we show how to instantiate the ideal functionalities $\mathcal{F}^{t,\mathcal{O}}_{\text{anon}}$ used in committee election $\mathsf{LBin\text{-}V}$. Recall that the ideal functionality $\mathcal{F}^{t,\mathcal{O}}_{\text{anon}}$ receives one message from each player and either sends the set of all messages it receives to everyone or a set of corrupt players of size at least t to everyone. We will first give our protocol in a $\mathsf{IdealZK}^*$-hybrid model in which players have access to an ideal zero-knowledge proof functionality. Then we use the elegant techniques of

Pass [Pas04] to instantiate the protocol with real-world cryptography. Next. we will first describe the $\mathsf{IdealZK}^*$ functionality in Section 6.1, and then we will give our protocol in the $\mathsf{IdealZK}^*$-hybrid world in Section 6.2.

6.1 Ideal Zero-Knowledge Functionality $\mathsf{IdealZK}^*$

Basically, $\mathsf{IdealZK}^*$ either sends $\mathsf{success}$ to everyone indicating that the proof is correct, or the identity of the prover/verifier who leads to the failure of the proof. Formally,

Ideal Zero-knowledge Functionality $\mathsf{IdealZK}^*[x, L, i, j]$

The functionality involves n parties $1, \ldots, n$, and is parametrized with a statement x, the language L, the prover's identity i and the verifier's identity j.

1. If both the prover i and the verifier j are corrupted, receive a bit b from the prover i. If $b = 1$, send $(\mathsf{success}, i, j)$ to everyone.
2. Receive ok or $\perp$ from the verifier j.
3. If received $\perp$ from the verifier, send (fail, j) to everyone.
4. Receive w or $\perp$ from the prover.
5. If $\mathcal{R}(x, w) = 1$, send $(\mathsf{success}, i, j)$ to everyone. Otherwise send (fail, i) to everyone.

In an n-party $\mathsf{IdealZK}^*$-hybrid protocol, the players can invoke the ideal zero-knowledge functionality $\mathsf{IdealZK}^*[x, L, i, j]$ between any prover $i \in [n]$ and any verifier $j \in [n]$, and for arbitrary NP language L. Without loss of generality, in every round, there can be at most n^2 concurrent invocations of $\mathsf{IdealZK}^*$. Given an n-party $\mathsf{IdealZK}^*$-hybrid protocol, we can instantiate $\mathsf{IdealZK}^*$ with actual cryptography using the elegant techniques suggested by Pass [Pas04].

Theorem 6.1. *(Constant-round, bounded concurrent secure computation [Pas04]). Assume the existence of enhanced trapdoor permutations and collision-resistant hash functions. Then, given an n-party $\mathsf{IdealZK}^*$-hybrid protocol Π^*, in which the number of concurrent calls of $\mathsf{IdealZK}^*$ is upper bounded by a priori known bound $m = \mathsf{poly}(\lambda)$, there exists a real-world protocol Π 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, there exists a non-uniform p.p.t. adversary $\mathcal{A}^*$ in the protocol Π^*, such that for any input $(x_1, \ldots, x_n)$, every auxiliary string $z \in \{0,1\}^*$,*

$$\mathsf{Exec}^{\Pi,\mathcal{A}}(1^\lambda, x_1, \ldots, x_n, z) \equiv_c \mathsf{Exec}^{\Pi^*,\mathcal{A}^*}(1^\lambda, x_1, \ldots, x_n, z).$$

In the above, the notation $\mathsf{Exec}^{\Pi,\mathcal{A}}$ (or $\mathsf{Exec}^{\Pi^*,\mathcal{A}^*}$) outputs each honest players' outputs as well as the corrupt players' (arbitrary) outputs.

– **Round efficiency**: *The round complexity of Π is at most a constant factor worse than that of Π^*.*

This real-world protocol is fulfilled by replacing the $\mathsf{IdealZK}^*$ instance with the bounded concurrent zero-knowledge proofs. All the zero-knowledge proof messages are posted to the broadcast channel. The full proof of Theorem 6.1 is available in the full version.

Now it suffices to show how to replace $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ with a protocol $\mathsf{Anon}^{t,\mathcal{O}}$ in the $\mathsf{IdealZK}^*$-hybrid world. In the protocol, we will omit the language L when it is clear from the context.

6.2 Implementing Anonymous Broadcast Functionality

In this section, we describe how to implement $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ in the $\mathsf{IdealZK}^*$-hybrid model. The protocol makes use of a perfect binding, statistically hiding commitment scheme comm. Also, every player keeps track of two sets, $\mathcal{D}_s$ and $\mathcal{D}_r$, the set of players who fail to share and the set of players who fail to reconstruct, respectively, to guarantee the identifiable abort property. Still, we use $\mathcal{K}$ to represent the set of corrupted players, $\mathcal{H}$ to represent the set of honest players. The number of parallel sessions is set to be λ. The protocol $\mathsf{Anon}^{t,\mathcal{O}}$ is given below.

$\mathsf{Anon}^{t,\mathcal{O}}$: instantiating $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ in the $\mathsf{IdealZK}^*$ -hybrid world

Parameters: Let $M = 2n$ be the number of slots. Let $\mathcal{D}_s$, $\mathcal{D}_r$ and Out be initially empty sets. Without loss of generality we assume that $\mathcal{O} = [n]$.

Building blocks: A perfectly binding, computationally hiding commitment scheme comm.

Input: Each player has an input $m_i \in \mathbb{F}$ for a finite field $\mathbb{F}$ with size larger than 2^λ. The sum of tuples is computed entry-wise, i.e., $(a_1, b_1, c_1) + (a_2, b_2, c_2) = (a_1 + a_2, b_1 + b_2, c_1 + c_2)$.

Preparation Phase Run the following for λ independent, parallel sessions:

1. Player i uniformly randomly choose a nonce $\mathsf{mid}_i \in \mathbb{F}$. It then uniformly randomly chooses a slot $l_i \leftarrow [M]$ and computes a vector $\mathbf{S}_i \in (\mathbb{F}^3)^M$ such that $\mathbf{S}_i[l] = (0, 0, 0)$ if $l \neq l_i$, and $\mathbf{S}_i[l] = (m_i, \mathsf{mid}_i, 1)$ if $l = l_i$.
2. Player i then splits $\mathbf{S}_i$ into $(n-t)$-out-of-n Shamir secret shares. Let $\mathbf{X}_{i,j}$ be the j-th share of $\mathbf{S}_i$. Let $\widehat{\mathbf{X}}_{i,j} = \mathsf{comm}(\mathbf{X}_{i,j}, r_{i,j})$ where $r_{i,j}$ are fresh randomness. Broadcast the commitments $\{\widehat{\mathbf{X}}_{i,j}\}_{j \in [n]}$.
3. If a player i fails to broadcast the commitments, add i to the set $\mathcal{D}_s$.

Validation Phase For $sid \in [\lambda]$, let $*^{sid}$ denote the variable $*$ in session sid. Player i invoke $\mathsf{IdealZK}^*[\mathsf{stmt}_i, i, j]$ for each $j \in [n]$, with the statement $\mathsf{stmt}_i = \{\widehat{\mathbf{X}}_{i,j}^{sid}\}_{j \in [n], sid \in [\lambda]}$, and send the witness $w = (m_i, \mathsf{mid}_i, \{\mathbf{S}_i^{sid}\}_{sid \in [\lambda]}, \{\mathbf{X}_{i,j}^{sid}, r_{i,j}^{sid}\}_{j \in [n], sid \in [\lambda]})$ to prove that

– For each $sid \in [\lambda]$, for each $j \in [n]$, $(\mathbf{X}_{i,j}^{sid}, r_{i,j}^{sid})$ is the correct opening of $\widehat{\mathbf{X}}_{i,j}^{sid}$;
– For each $sid \in [\lambda]$, $\{\mathbf{X}_{i,j}^{sid}\}_{j\in[n]}$ forms a valid $(n-t)$-out-of-n secret sharing of $\mathbf{S}_i^{sid}$;
– For each $sid \in [\lambda]$, the vector $\mathbf{S}_i^{sid}$ contains only one non-zero slot $(m_i, \mathsf{mid}_i, 1)$.

For each $i \in [n]$, if there exists a j that $\mathsf{IdealZK}^*[\mathsf{stmt}_i, i, j]$ outputs (fail, i), i.e., the prover fails to prove the statement to receiver j, add i to the set $\mathcal{D}_s^{sid}$ for all $sid \in [\lambda]$.

Sharing phase Continue the following for λ independent, parallel sessions:

1. For $j \in [n]$, player i sends $(\mathbf{X}_{i,j}, r_{i,j})$ to player j.
2. Player i does the following: for every $j \in [n] \setminus \mathcal{D}_s$, if it receives a message $(\mathbf{X}_{j,i}, r_{j,i})$ that is a correct opening with respect to $\widehat{\mathbf{X}}_{j,i}$, record $(\mathbf{X}_{j,i}, r_{j,i})$ and broadcast (ok, i, j). Otherwise, broadcast $(\mathsf{complain}, i, j)$ to complain about j.
3. Player i does the following: for all j such that there is a complain $(\mathsf{complain}, j, i)$ in Step 2, player i broadcasts the corresponding opening $(i, j, \mathbf{X}_{i,j}, r_{i,j})$.
4. Unless player i broadcasts all correct openings for those players who has sent $(\mathsf{complain}, j, i)$ to complain about i, add i to the set $\mathcal{D}_s$.
5. Player i does the following: for $j \in [n]\setminus\mathcal{D}_s$, if player i sent $(\mathsf{complain}, i, j)$ in Step 2, and j broadcast a correct opening $(\mathbf{X}_{j,i}, r_{j,i})$ in Step 3. then record the correct opening $(\mathbf{X}_{j,i}, r_{j,i})$.

Reconstruction Phase Run the following for λ independent, parallel sessions:

1. Player i computes $\mathbf{Y}_i = \sum_{j\in[n]\setminus\mathcal{D}_s} \mathbf{X}_{j,i}$ and broadcast $\mathbf{Y}_i$. If a player j fails to broadcast, add j to the set $\mathcal{D}_r$.
2. Player i does the following for each $j \in [n]$: invoke $\mathsf{IdealZK}^*[\mathsf{stmt}'_i, i, j]$ with the statement $\mathsf{stmt}'_i = (\mathcal{D}_s, \mathbf{Y}_i, \{\widehat{\mathbf{X}}_{j,i}\}_{j\in[n]\setminus\mathcal{D}_s})$. It sends the witness $w' = (\{\mathbf{X}_{j,i}, r_{j,i}\}_{j\in[n]\setminus\mathcal{D}_s})$ to the ideal functionality $\mathsf{IdealZK}^*$ to prove that
   – For any $j \in [n] \setminus \mathcal{D}_s$, $(\mathbf{X}_{j,i}, r_{j,i})$ is a correct opening of $\widehat{\mathbf{X}}_{j,i}$;
   – $\mathbf{Y}_i = \sum_{j\in[n]\setminus\mathcal{D}_s} \mathbf{X}_{j,i}$.
3. If there exists a j such that $\mathsf{IdealZK}^*[\mathsf{stmt}'_i, i, j]$ outputs (fail, i), i.e., the prover fails to prove the statement to receiver j, add i to the set $\mathcal{D}_r$.
4. If $|\mathcal{D}_r| \geq t$, everyone stores $(\mathsf{fail}, \mathcal{D}_r \cup \mathcal{D}_s)$ for the reconstruction phase of this session.
5. Otherwise, every player uses all broadcast shares $\{\mathbf{Y}_i\}_{i\in[n]\setminus\mathcal{D}_r}$ to reconstruct the sum $\mathbf{S} = \sum_{i\notin\mathcal{D}_s} \mathbf{Y}_i$. Store $(\mathsf{ok}, \mathbf{S})$ for the reconstruction phase of this session.

Output Phase For each $sid \in [\lambda]$, we use $(\mathsf{fail}, \mathcal{D}^{sid})$ or $(\mathsf{ok}, \mathbf{S}^{sid})$ to denote the value each player stores in the reconstruction phase of session sid. Each player i does the following:

1. If there is a $sid \in [\lambda]$ such that player i stores $(\mathsf{fail}, \mathcal{D}^{sid})$ for that session, outputs $(\mathsf{fail}, \cup_{sid \in [\lambda]} \mathcal{D}^{sid})$, where $\mathcal{D}^{sid} = \emptyset$ for those successfully reconstructed sessions.
2. Otherwise, each player does the following: We say that (m, mid) *appears* in session sid if there exists a slot $l \in [M]$ such that $\mathbf{S}^{sid}[l] = (m, \mathsf{mid}, 1)$. For each pair (m, mid) that appears in a majority number of sessions, add a copy of m to Out.
3. Output $(\mathsf{ok}, \mathsf{Out})$.

Theorem 6.2. *If the commitment scheme* comm *is perfectly binding and computationally hiding, then* $\mathsf{Anon}^{t,\mathcal{O}}$ *securely realizes* $\mathcal{F}_{\text{anon}}^{t,\mathcal{O}}$ *in the* $\mathsf{IdealZK}^*$*-hybrid model as long as* $|\mathcal{O}| - t \geq |\mathcal{K}|$*. Moreover,* $\mathsf{Anon}^{t,\mathcal{O}}$ *runs in constant number of rounds.*

The full proof of the above theorem is available in the full version.

Acknowledgement

This work is in part supported by NSF awards under the grant numbers 2044679 and 1704788, a Packard Fellowship, and a generous gift from Nikolai Mushegian.

References

Abe99. Masayuki Abe. Mix-networks on permutation networks. In *ASIACRYPT*, 1999.

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 Joe Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *PODC*, 2006.

ADMM14. Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *S&P*, 2014.

AL11. Gilad Asharov and Yehuda Lindell. Utility dependence in correct and fair rational secret sharing. *Journal of Cryptology*, 24(1):157–202, 2011.

AO16. Bar Alon and Eran Omri. Almost-optimally fair multiparty coin-tossing with nearly three-quarters malicious. In *TCC*, 2016.

AS16. Noga Alon and Joel H Spencer. *The probabilistic method*. John Wiley & Sons, 2016.

Aum74. Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.

BGKO11. Amos Beimel, Adam Groce, Jonathan Katz, and Ilan Orlov. Fair computation with rational players. *Cryptology ePrint Archive*, 2011.

BK14. Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *CRYPTO*, 2014.

Blu83. Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *SIGACT News*, 1983.

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

CCWS21. Kai-Min Chung, T-H Hubert Chan, Ting Wen, and Elaine Shi. Game-theoretic fairness meets multi-party protocols: The case of leader election. In *CRYPTO*, 2021.

CGF10. Henry Corrigan-Gibbs and Bryan Ford. Dissent: accountable anonymous group messaging. In *CCS*, 2010.

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.

Cha81. David L Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.

Cha88. David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of cryptology*, 1(1):65–75, 1988.

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.

DMS04. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, 2004.

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

DR+07. Yevgeniy Dodis, Tal Rabin, et al. Cryptography and game theory. *Algorithmic game theory*, pages 181–207, 2007.

Fei99. Uriel Feige. Noncryptographic selection protocols. In *FOCS*, 1999.

GGS. Rati Gelashvili, Guy Goren, and Alexander Spiegelman. Short paper: On game-theoretically-fair leader election.

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

GKM+13. Juan 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.

GMW19. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game, or a completeness theorem for protocols with honest majority. In *STOC*. 2019.

GTZ15. Juan Garay, Björn Tackmann, and Vassilis Zikas. Fair distributed computation of reactive functions. In *DISC*, 2015.

HT04. Joseph Halpern and Vanessa Teague. Rational secret sharing and multi-party computation. In *STOC*, 2004.

HT14. Iftach Haitner and Eliad Tsfadia. An almost-optimally fair three-party coin-flipping protocol. *STOC*, 2014.

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

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

KMSW22. Ilan Komargodski, Shin'ichiro Matsuo, Elaine Shi, and Ke Wu. log*-round game-theoretically-fair leader election. Cryptology ePrint Archive, Paper 2022/791, 2022. https://eprint.iacr.org/2022/791.

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

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

MNS09. Tal Moran, Moni Naor, and Gil Segev. An optimally fair coin toss. In *TCC*, 2009.

Nas51. John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

OPRV09. Shien Jin Ong, David C Parkes, Alon Rosen, and Salil 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.

RSZ02. Alexander Russell, Michael Saks, and David Zuckerman. Lower bounds for leader election and collective coin-flipping in the perfect information model. *SIAM Journal on Computing*, 31(6):1645–1662, 2002.

RZ01. Alexander Russell and David Zuckerman. Perfect information leader election in log* n+ o (1) rounds. *Journal of Computer and System Sciences*, 63(4):612–626, 2001.

SGR99. Paul Syverson, D Goldschlag, and M Reed. Onion routing for anonymous and private internet connections. *Communications of the ACM*, 42(2):5, 1999.

Sha79. Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.

WAS22. Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. In *Eurocrypt*, 2022.

Yao82. Andrew C Yao. Protocols for secure computations. In *FOCS*, 1982.

ZZZR05. Li Zhuang, Feng Zhou, Ben Y Zhao, and Antony Rowstron. Cashmere: Resilient anonymous routing. In *NSDI*, 2005.