The Fairest of Them All: Best-of-Both-Worlds Fair Leader Election

Chen-Da Liu-Zhang[1], Elisaweta Masserova[2], Kelsey Melissaris[3], Lawrence Roy[4], and Sophia Yakoubov[5]

[1] Lucerne University of Applied Sciences and Arts
[2] Carnegie Mellon University
[3] Université Paris Cité, CNRS, IRIF
[4] IBM Research Zürich
[5] Aarhus University

Abstract. Fair leader election is the problem of sampling an unbiased random party out of a set of n parties. The seminal impossibility result of Cleve (STOC 1986) applies to this setting, and shows that the result of the election cannot be unbiased against a dishonest majority. Recent work has achieved the weaker property of game-theoretic fairness, where behaving dishonestly cannot improve one's chances of being elected. However, such protocols allow even a single corrupt party to bias the outcome.

In this work, we introduce best-of-both-worlds (BoBW) fairness, where a small number t_S of corrupt parties cannot bias the result, and even a larger number t_{GT} of corrupt parties (dishonest majority) still cannot improve their collective election chances (CSP-fairness) nor hurt the election chances of any given honest party (MM-fairness). We present a protocol for BoBW fair leader election for $t_S + t_{GT} < n$. We then prove a generalization of Cleve's impossibility, and show that this protocol is optimal. Next, in the setting of statistical security we find that BoBW fairness is uninteresting, as it is only possible for an honest majority, where we can get also unbiased leader election. Finally, in the setting of perfect security, we present a BoBW protocol achieving CSP-fairness and strong fairness for $t_S + 1 = t_{CSP} = \frac{n}{3}$, and prove that this is also optimal.

Table of Contents

The Fairest of Them All: Best-of-Both-Worlds Fair Leader Election . . . 1
Chen-Da Liu-Zhang, Elisaweta Masserova, Kelsey Melissaris, Lawrence Roy, and Sophia Yakoubov
1 Introduction . . . 3
1.1 Our Contributions . . . 4
1.2 Related Work . . . 6
1.3 Model and Notation . . . 7
2 Technical Overview . . . 7
2.1 Computational BoBW Leader Election . . . 8
2.2 Statistical Leader Election . . . 10
2.3 Perfect Leader Election . . . 10
3 Definitions of Fair Sampling . . . 13
3.1 Preference Profiles and Utility Functions . . . 13
3.2 Strategies . . . 14
3.3 Fairness . . . 14
3.4 Leader Election . . . 16
4 Computational Best of Both Worlds Leader Election . . . 16
4.1 Verifiable Secret Sharing . . . 16
4.2 BoBW Two-Party Leader Election . . . 20
4.3 BoBW Multi-Party Leader Election . . . 22
5 BoBW Leader Election Impossible for $t_S + t_{GT} \geqslant n$ . . . 24
5.1 Generalized Cleve . . . 25
5.2 Lower Bound for $n = 3$ . . . 27
5.3 Proof of Theorem 4 . . . 28
5.4 Corollaries . . . 29
6 Statistical Strongly Fair Leader Election for $t < \frac{n}{2}$ . . . 29
7 Statistical Leader Election Impossible for $t \geqslant \frac{n}{2}$ . . . 30
7.1 Two Parties . . . 30
7.2 n Parties . . . 31
7.3 Corollaries . . . 32
8 Perfect CSP-fair Leader Election for $n = 3t$ . . . 33
8.1 CSP-Fair Leader Election for $n = 3$, $t = 1$ . . . 33
8.2 CSP-Fair Leader Election for $n = 3t$ . . . 34
9 Lower Bounds for Perfect Leader Election . . . 36
9.1 Impossibility of CSP Leader Election for $n < 3t$ . . . 36
9.2 Impossibility of MM Fair Leader Election for $n \leqslant 3t$ . . . 38
10 AI Methodology . . . 41

1 Introduction

Random sampling is an important task in many cryptographic computations. When a majority of participants is honest, this can be done in such a way that the output is guaranteed to match the target distribution. However, when only a minority of participants is honest (and a corrupt majority may deviate from the sampling protocol), Cleve [Cle86] shows that the corrupt majority will always be able to bias the output in at least one direction.

One useful sampling task is leader election, where a set of n parties wish to select one among themselves to play a special role (e.g. to present their new result at a conference). Though a seemingly limited functionality, leader election has a wide range of applications (beyond selecting speakers): it can, e.g., be used in blockchains to determine who should mine the next block, and in large-scale MPC to determine who should participate in committee-based computation. The result of Cleve shows that, without an honest majority, the leader cannot be guaranteed to be perfectly random.

Recent work [CGL+18,CCWS21,WAS22,KMSW22,TSW24] uses parties' preferences for certain outcomes to circumvent Cleve's lower bound by designing sampling protocols that guarantee *game-theoretic fairness*, where deviating parties are only able to bias the output in a way that hurts themselves. There are two existing notions of game-theoretic fairness: CSP-fairness, and Maximin-fairness. In a *CSP-fair* protocol, corrupt parties are unable to increase their own utility; in a *Maximin- (MM-) fair* protocol, corrupt parties are unable to decrease the utility of any honest parties. In the context of leader election, CSP-fairness implies that a set of cheating parties cannot increase their own probability of being elected, and MM-fairness additionally implies that they cannot hurt any honest party's chances.

The downside of such game-theoretically fair protocols is that they are designed for a corrupt majority, and even if most of the participants do end up being honest, the protocols only give the weaker fairness guarantee. The number of corrupt participants may be unknown when a leader-election protocol is deployed, so it is desirable to have a single leader election protocol that provides the strongest guarantee (unbiasability, see below) against small coalitions, while still giving meaningful protection (game theoretic fairness, see below) against much larger ones. In this paper, we introduce *Best of Both Worlds* (BoBW) fairness, which captures exactly such a dual guarantee.

Notions of Fairness in Leader Election. We consider several notions of fairness. Strong fairness and several notions of game theoretic fairness are already defined in the literature, all parameterized by a single threshold t.

t-Strong Fairness (Unbiasability): t or fewer cooperating corrupt parties should be unable to change the output distribution.

Game Theoretic Fairness [CGL+18]: Game-theoretic notions of fairness include CSP fairness, maximin (MM) fairness, and their combination — GT fairness.

- **t-CSP Fairness:** t or fewer cooperating corrupt parties should be unable to increase the probability of a corrupt leader being chosen.
- **t-MM Fairness:** t or fewer cooperating corrupt parties should be unable to decrease the probability of any honest party being chosen as leader.
- **t-GT Fairness:** both t-CSP and t-MM fairness should be satisfied.

In this paper, we introduce the notion of best-of-both-worlds (BoBW) fairness. BoBW fairness is parameterized by two thresholds: t_{GT} and t_S. A leader election protocol is (t_S, t_{GT})-BoBW fair if, as long as at most t_S parties are corrupt, the distribution of the chosen leader is uniform over all participants (i.e., strong fairness is satisfied), but if more than t_S but at most t_{GT} parties are corrupt, the protocol satisfies one of the notions of game-theoretic fairness.

(t_S, t_{GT})**-BoBW Fairness (New):** t_S-strong and t_{GT}-CSP/MM/GT fairness should be simultaneously satisfied.

1.1 Our Contributions

We introduce the notion of best-of-both-worlds (BoBW) fairness and determine its feasibility, specifically in the context of leader election. We also determine the exact corruption thresholds for which information theoretic leader election with existing notions of fairness is possible.

Computational. For computational leader election protocols, strong fairness is known to be unachievable if $t \geqslant \frac{n}{2}$ [Cle86] and achievable otherwise [GMW87]. All notions of game theoretic fairness are known to be achievable for all thresholds [CCWS21]. It remains to determine the feasibility of BoBW fair protocols. We present a protocol that achieves (t_S, t_{GT})-BoBW fairness as long as $t_S < \frac{n}{2}$ and $t_S + t_{GT} < n$, where n is the number of participants (Theorem 3). This may, at first glance, seem sub-optimal, but we prove a matching lower bound: if $t_S + t_{GT} \geqslant n$, no such protocol exists (Theorem 4). Figure 1 illustrates our results.

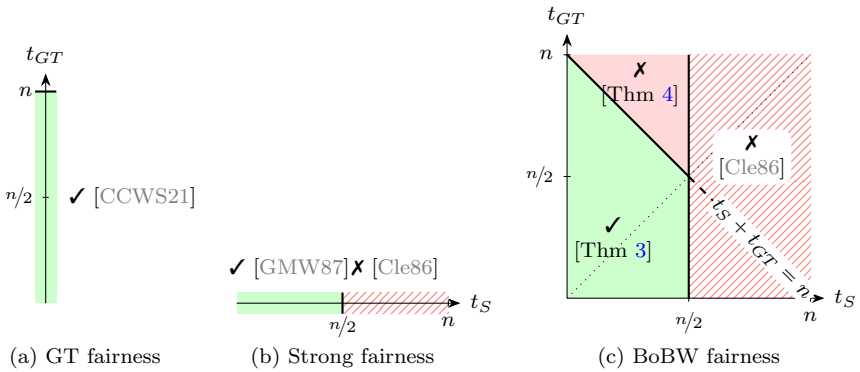

(a) GT fairness (b) Strong fairness (c) BoBW fairness

Fig. 1: Feasibility of computationally fair leader election as a function of the strong-fairness threshold t_S and the GT-fairness threshold t_{GT}, assuming secure point-to-point channels and a broadcast channel. (a) t_{GT}-GT fairness is feasible for every $t_{GT} < n$ [CCWS21]. (b) t_S-strong fairness is feasible iff $t_S < n/2$ [GMW87,Cle86]. (c) (t_S, t_{GT})-BoBW fairness is feasible (green) iff $t_S < n/2$ and $t_S + t_{GT} < n$: Theorem 3 gives the protocol, Theorem 4 rules out $t_S + t_{GT} \geqslant n$ (red), and Cleve's bound rules out $t_S \geqslant n/2$ (hatched).

Statistical. We show that in the statistical setting, game-theoretic fairness buys nothing beyond strong fairness: for $t < \frac{n}{2}$, *strongly* fair leader election is possible (Theorem 6) using general-purpose statistical secure computation, which is feasible for the same corruption thresholds [RB89]. For $t \geqslant \frac{n}{2}$ even the weakest of our notions, CSP-fairness, is impossible (Theorem 7). Consequently, BoBW fairness is uninteresting in the statistical setting: wherever a non-trivial game-theoretic guarantee can be achieved, strong fairness can be achieved as well. Table 1 summarizes the landscape.

Perfect. In the computational and statistical settings, CSP, MM and GT fairness are achievable for exactly the same thresholds (and, in fact, by the same protocols). The perfect setting is the only one in which the three notions separate. Perfect *strongly* fair leader election is possible for $n > 3t$ via perfectly secure general-purpose computation [BGW88,CCD88]; so, as in the statistical setting,

	CSP-fairness (Def. 2)	MM-fairness (Def. 3)	GT-fairness (Def. 4)	Strong fairness (Def. 1)
$n > 2t$	←	←	←	✓ [Thm 6]
$n \leqslant 2t$	✗ [Thm 7]	→	→	✗ [Cle86]

Table 1: Feasibility of statistically fair leader election with n parties and up to t corruptions, assuming secure point-to-point channels and a broadcast channel.

game-theoretic fairness can only buy something at the boundary $n \leqslant 3t$. We show that it does, but only for CSP fairness, and only at a single threshold:

- We present a perfectly CSP-fair leader election protocol for $n = 3t$ (Theorem 9).
- We prove that perfect CSP fairness with bounded number of rounds is impossible for $n < 3t$ (Theorem 10), so our construction is tight.
- We prove that perfect MM fairness — and hence GT and strong fairness — with bounded number of rounds is impossible already for $n \leqslant 3t$ (Theorem 11). That is, unlike CSP fairness, MM fairness gives nothing beyond strong fairness in the perfect setting.

Our protocol for $n = 3t_{GT}$ is, in addition, strongly fair whenever at most $t_{GT} - 1$ parties are corrupt, so it achieves S/CSP $(t_{GT} - 1, t_{GT})$-BoBW fairness (Corollary 5). This is the best one can hope for: $t_S \geqslant \frac{n}{3}$ is ruled out by our impossibility of perfect strong fairness for $n \leqslant 3t_S$, and $t_{GT} > \frac{n}{3}$ is ruled out by Theorem 10. For S/MM BoBW fairness, on the other hand, the perfect setting is as uneventful as the statistical one: wherever MM fairness is achievable, so is strong fairness. Table 2 summarizes the landscape.

	CSP-fairness (Def. 2)	S/CSP BoBW-fairness (Def. 6; $t_{GT} = t$, $t_S = t - 1$)	MM-fairness (Def. 3)	GT-fairness (Def. 4)	Strong fairness (Def. 1)
$n > 3t$	←	←	←	←	✓[BGW88]
$n = 3t$	✓ [Thm 9]	✓ [Cor. 5]	✗ [Thm 11]	→	→
$n < 3t$	✗ [Thm 10]	→	✗ [Cor. 6] →	✗ [Cor. 7] →	✗ [Cor. 8]

Table 2: Feasibility of perfectly fair leader election with n parties and up to t corruptions in the information-theoretic setting.

Notably, the CSP possibility and MM impossibility at $n = 3t$ make leader election the first natural sampling task for which game-theoretic fairness provably helps in the information-theoretic setting, and the first for which CSP and MM fairness separate.

1.2 Related Work

Strong Fairness and Coin Tossing. Fair coin tossing was introduced by Blum [Blu81]. With an honest majority, every efficiently samplable distribution can be sampled with strong fairness: computationally for $t < \frac{n}{2}$ [GMW87], statistically for $t < \frac{n}{2}$ given a broadcast channel [RB89], and perfectly for $t < \frac{n}{3}$ [BGW88,CCD88]. Without an honest majority, Cleve [Cle86] showed that any R-round coin-tossing protocol admits a fail-stop attack with bias $\Omega(1/R)$; a long line of work [MNS09,BOO10,HT14] studies how closely this can be matched. Our generalized Cleve theorem (Theorem 5) extends Cleve's attack from a single shared bit to a pair of correlated bits, which is what leader election naturally induces. A separate line of work in the full-information model [AN90,RZ98,Fei99] studies coin flipping and leader election against large coalitions without any cryptography; there, the goal is to keep the adversary's influence bounded rather than negligible, and honest parties are not guaranteed a $1/n$ chance of winning.

Game-Theoretic Fairness. Chung, Guo, Lin, Pass and Shi [CGL+18] introduced the game-theoretic notions of fairness used in this paper, CSP and maximin fairness, and characterized their feasibility for two-sided coin tossing against coalitions of size $n-1$. Wu, Asharov and Shi [WAS22] completed this characterization for majority-sized coalitions, and Zhang and Wu [ZW25] and Thyagarajan, Soni and Wu [TSW24] extend game-theoretic fairness to arbitrary preference profiles and to sampling from general distributions, respectively. Most closely related to our constructions is the work of Chung, Chan, Wei and Shi [CCWS21] on game-theoretically fair leader election, which gives a tournament-tree protocol that is CSP- and MM-fair against any coalition of size $t < n$ (see also [BZ17] for the same tournament idea in the context of Bitcoin lotteries, and [GGS22]). They further show that $\Omega(\log n)$ rounds are necessary for exact fairness, and trade rounds for an approximation of fairness; Komargodski, Matsuo, Shi and Wu [KMSW22] reduce the round complexity to $O(\log^* n)$ for approximate fairness and consider committee election. All of these protocols are computationally secure. Our protocols (Section 4.3) retain the tournament structure of [CCWS21].

Information-Theoretic Game-Theoretic Fairness. Liu-Zhang, Masserova, Ribeiro and Thyagarajan [LZMRT26] recently showed that, for two-sided coin tossing, game-theoretic fairness does not help in the statistical setting: statistically CSP-fair coin tossing is impossible for $t \geqslant \frac{n}{2}$, i.e., exactly where strong fairness is impossible. Our statistical impossibility (Theorem 7) transfers this result to leader election via a party-emulation argument.

Best of Both Worlds. The idea of a single protocol that offers a strong guarantee against few corruptions and a weaker but still meaningful guarantee against many has a long history in secure computation, where it is also referred to as "best of both worlds" or hybrid security. Chaum [Cha90] combined unconditional security against minorities with computational security against majorities. Ishai et al. [IKLP06] and Katz [Kat07] studied combining guaranteed output delivery for $t < \frac{n}{2}$ with security with abort for larger t, showing that the two thresholds must satisfy $t_{\mathsf{god}} + t_{\mathsf{abort}} < n$. Lucas, Raub and Maurer [LRM10] give a general framework for such hybrid guarantees. Our notion of BoBW fairness is in the same spirit, with strong fairness in the role of the stronger guarantee and game-theoretic fairness in the role of the weaker one.

Rational Cryptography. Our work belongs to the line of research on protocols for rational rather than fully malicious participants [HT04,ADGH06,KN08,GK12]. One approach [IML05] uses external incentives such as rewards or deposits; game-theoretic fairness, which we follow, instead exploits the parties' natural preferences over the protocol's output, which makes it well suited to sampling tasks such as leader election where these preferences are known.

1.3 Model and Notation

We assume that n parties $\{P_1, \ldots, P_n\}$ are connected via secure point-to-point channels and have access to a broadcast channel. All communication occurs in synchronous rounds, and we consider a rushing adversary that sends messages that can depend on honest messages from the same round. We use t_{GT} to denote the threshold for which game-theoretic fairness should hold, and t_S to denote the threshold for which strong fairness should hold. We use $[n]$ to denote the set $\{1, \ldots, n\}$.

Throughout, the view of a party consists of its random tape, inputs, outputs, and all messages sent and received by that party, and the view of the adversary contains of the view of every corrupt party.

2 Technical Overview

As a starting point for our protocols, we take the weighted game-theoretic two-party leader election protocol of [CCWS21], re-stated as Protocol 1. If P_0 has weight k_0 and P_1 has weight k_1, then P_0 should win with probability $\frac{k_0}{k_0+k_1}$ and P_1 should win with probability $\frac{k_1}{k_0+k_1}$.

Protocol 1: $\Pi_{2-GT}(k_0, k_1)$: Game-Theoretically Fair Two-Party Leader Election

Commit: Each party P_i for $i \in [0, 1]$ does the following:

1. $b_i \leftarrow \{0, \ldots, k_0 + k_1 - 1\}$.
2. $(c_i, d_i) \leftarrow \mathsf{Commit}(b_i)$.
3. Broadcast c_i.

If party i fails to send a commitment, then the protocol terminates and party $1 - i$ is chosen as the leader, or, equivalently, party i is eliminated.

Open: Upon receiving c_{1-i}, each party $i \in [0, 1]$ broadcasts d_i.

Output: Everyone does the following:

1. For $i \in [0, 1]$:
   (a) $b_i \leftarrow \mathsf{Open}(c_i, d_i)$.
   (b) If $b_i \notin \{0, \ldots, k_0 + k_1 - 1\}$, then the protocol terminates and party $1 - i$ is chosen as the leader, or, equivalently, party i is eliminated.
2. Let $b = (b_0 + b_1) \bmod (k_0 + k_1)$. P_0 is chosen as the leader if $b < k_0$, and P_1 is chosen as the leader otherwise.

The protocol has a standard "Commit-and-Open" structure, and therefore uses a commitment scheme as a building block. A commitment scheme is a cryptographic tool with which a party can produce a commitment to a value that hides that value (hiding), and can later open the commitment to only that value (binding). For the two-party protocol above, we assume ideal commitments: a corrupt party must fix its contribution independently of the honest party's hidden contribution. Consequently, their sum is uniform modulo the sum of the weights, giving the prescribed winning probabilities. Note that a corrupted party can still bias the output by aborting, but this bias is permissible by the definition of game theoretic fairness as such misbehavior only decreases the probability that the misbehaving party is elected.

Protocol 1 extends to game-theoretically fair n-party leader election via a tournament [CCWS21]. Every party starts with weight 1; throughout, a party's weight is the number of original parties it "represents," and the winner of a match between parties with weights k_0 and k_1 is chosen with probabilities proportional to those weights. The tournament proceeds in $\lceil \log n \rceil$ iterations. In each iteration, the parties that have not yet been eliminated are paired off (if their number is odd, one

party sits out and advances with its weight unchanged), and each pair runs Protocol 1 with their current weights. The party that is elected in a match advances, and takes on the sum of both parties' weights; the other party is eliminated. After the last iteration, a single party with weight n remains; it is the elected leader. Since in an honest execution a party's probability of surviving each match equals its share of the combined weight, every party is elected with probability exactly $\frac{1}{n}$.

The resulting tournament is MM fair (and thus also CSP and GT fair) because an honest party's expected weight never decreases: in every match it plays, the corrupt opponent can either play honestly, in which case the honest party keeps the combined weight with probability exactly its share of that weight, or abort, in which case the honest party wins outright and keeps the combined weight with probability 1. This remains true if the opponent decides whether to abort after learning the outcome: aborting can only replace a prospective loss for the honest party by a win, or leave a prospective win unchanged.

2.1 Computational BoBW Leader Election

2.1.1 Protocol for $t_S + t_{GT} < n$ To achieve (t_S, t_{GT})-BoBW fairness for $t_S + t_{GT} < n$, we retain the tournament and change how each match is resolved.

Protocol 1 does not achieve strong fairness: a corrupt participant can bias the outcome towards her opponent by aborting, and two corrupt participants can choose the winner by selecting their b's accordingly. To prevent these two biasing attacks when t_S or fewer parties are corrupt, we involve all n parties in both the Commit and Open phases of each individual match in the tournament, thereby increasing the number of misbehaving parties required to mount a biasing attack.

As a building block, we introduce *dual-reconstruction verifiable secret sharing (DR-VSS)*, parameterized by a privacy threshold t. Its sharing phase fixes a value, which remains hidden when the dealer is honest and at most t parties are corrupt. There are two ways to recover it: joint reconstruction succeeds whenever at least $t+1$ parties are honest, while an honest dealer can always reveal the value without the other parties' cooperation. Whenever either reconstruction procedure succeeds, it reveals the same fixed value.

We build DR-VSS from an ideal non-malleable commitment functionality, using bi-variate polynomials. We then use it to obtain a (t_S, t_{GT})-BoBW two-party leader election protocol, which decides a match between P_i and P_j with weights k_i and k_j. All n parties participate, and the protocol proceeds in a commit phase and an open phase.

In the commit phase, every party P_l, $l \in \{0, \ldots, n-1\}$, picks a uniformly random $b_l \leftarrow \{0, \ldots, k_i + k_j - 1\}$ and shares it among all parties using DR-VSS. The privacy threshold of the sharing depends on the party's role. The bystanders, i.e., the parties P_l with $l \notin \{i, j\}$, share with privacy threshold t_S. This leaves enough honest parties for commitment — and guaranteed reconstruction — whether t_S or t_{GT} parties are corrupt, but if more than t_S parties are corrupt, the adversary may learn b_l before reconstruction. The two players P_i and P_j share with privacy threshold t_{GT}. Their values remain hidden under either threshold, but if more than t_S parties are corrupt, the adversary may prevent reconstruction; an honest dealer, however, can always open her value via dual reconstruction.

In the open phase, the parties reconstruct each shared value in turn, obtaining b'_l for $l \in \{0, \ldots, n-1\}$. For bystanders, reconstruction is guaranteed to succeed. If the reconstruction of a player's value fails, that player opens b'_l herself; since dual-reconstruction VSS binds her to the value she dealt, she can only open the correct value. If she refuses, she is eliminated and the other player wins the match. Otherwise, we let $b = \left(\sum_{l=0}^{n-1} b'_l\right) \bmod (k_i + k_j)$; P_i is elected leader if $b < k_i$, and P_j otherwise.

Strong fairness for t_S corruptions. In this regime every contribution is private and uniquely reconstructable, so every value is fixed and hidden before the open phase, and no party can prevent any value from being reconstructed. The sum of the contributions includes at least one honest, uniformly random contribution, so b is uniform and the adversary has no capacity to bias it at all.

Game-theoretic fairness for t_{GT} corruptions. In a match between two corrupt candidates, the outcome does not affect GT fairness. In a match between two honest candidates, all contributions are recoverable: bystanders' because they are shared with a low enough threshold, and candidates' because they will act as honest dealers. In a match between a corrupt and an honest candidate, the corrupt candidate may refuse to reveal its contribution after learning the lottery's outcome, but doing so awards the match to the honest candidate. Thus, the honest candidate's winning probability cannot decrease, and the corrupt candidate's cannot increase.

2.1.2 Lower Bound for $t_S + t_{GT} \geqslant n$ Our construction tolerates any $t_S < \frac{n}{2}$ and $t_{GT} < n - t_S$. At first, the second constraint may look like an artifact of the construction; strong fairness for $t_S < \frac{n}{2}$ and GT fairness for any $t_{GT} < n$ are achievable on their own, so why would asking for both simultaneously force the thresholds to trade off?

In Section 5 we prove that this trade-off is inherent; Theorem 4 states that there does not exist a polynomial-round leader election protocol that achieves S/CSP (t_S, t_{GT})-BoBW fairness for $t_S + t_{GT} \geqslant n$. This limitation holds even in the case of *fail-stop* corruptions, the setting in which corrupt parties follow the protocol honestly apart from choosing their own randomness (known as *semi-malicious* behavior) and aborting when they wish (known as *fail-stop* behavior). Both S/MM BoBW and full BoBW fairness imply S/CSP BoBW fairness, and therefore the same bound applies for all three notions. We now briefly convey the intuition behind the proof of Theorem 4.

Observe that a general n-party S/CSP (t_S, t_{GT})-BoBW fair leader election protocol Π_n for $t_S + t_{GT} \geqslant n$ (and, by Cleve's bound [Cle86], $t_S < \frac{n}{2}$) implies a three-party leader election protocol Π_3 between Alice, Bob and Carol with a restricted form of $(1, 2)$-BoBW fairness, as discussed in Section 5.3. In Π_3 Alice and Bob each play t_S of the n parties in a single execution of Π_n, Carol executes the remaining $n - 2t_S$ parties, and the winner of Π_3 is whichever of the three played the winner of Π_n. Alice alone (and likewise Bob alone) plays t_S parties, so Π_3 is strongly fair against each of them; Alice together with Carol (and likewise Bob with Carol) plays $n - t_S \leqslant t_{GT}$ parties, so Π_3 is CSP fair against each of these coalitions. These four guarantees are all that transfer: Carol alone plays $n - 2t_S$ parties, which may exceed t_S, and Alice and Bob together play $2t_S$ parties, which may exceed t_{GT}, so Π_3 need not be strongly fair against Carol or CSP fair against the Alice–Bob coalition. Fortunately, the four guarantees that do transfer suffice. In Π_3, Alice and Bob each win with probability $p = \frac{t_S}{n}$, and it suffices to prove (Lemma 1) that no poly-round three-party leader election protocol can be strongly fair against Alice alone and against Bob alone while being CSP fair against the Alice–Carol and the Bob–Carol coalitions. The proof by contradiction is as follows.

Protocol Π_3 can be further compiled into a two-party protocol Π_2 between Alice and Bob, in which each locally executes Carol with public randomness (sampled and sent in the first round by, say, Alice). Alice's output bit a indicates whether she wins, and Bob's output bit b whether he wins. Note that the output bits are *negatively correlated*: $\mathsf{cov}(a, b) = -p^2$. To Π_2 one can apply Theorem 5, a generalized variant of Cleve's impossibility proven herein,[6] that states that any R-round two-party correlated bit sampling protocol between Alice and Bob admits an attack that allows a failstop Alice to bias b, or a failstop Bob to bias a, with bias of magnitude at least $\frac{|\mathsf{cov}(a,b)|}{4R+2}$. The proof proceeds as Cleve's: describe $4R + 2$ aborting attacks available to Alice and Bob, show that the sum S of their biases satisfies $|\mathsf{cov}(a, b)| \leqslant S \leqslant (4R + 2)s$ for the largest bias s of any

[6] Recall that Cleve's lower bound is stated for two parties that sample and output a single common bit.

individual attack, and conclude $s \geqslant \frac{|\mathsf{cov}(a,b)|}{4R+2}$. For Π_2 this gives $s \geqslant \frac{p^2}{4R+2}$, which is non-negligible for polynomial R.

The attacks of Theorem 5 come in two kinds. The *naïve* attacks abort at a fixed round and use only the attacker's own view; in Π_3 they are attacks by Alice alone or Bob alone, so any bias contradicts strong fairness against that party. The remaining attacks decide when to abort based on what the attacker would output if the other party aborted, which in Π_2 depends on Carol's (public) randomness; in Π_3 they are therefore attacks by the Alice–Carol or Bob–Carol coalition. Here bias alone is not a contradiction: CSP fairness only forbids bias in the coalition's favor, and the attack might push winning probability *towards* the honest party. Therefore, third, we prove (Lemma 2, the Push-Pull Lemma) that if there exists an attack with which a coalition can *push* probability of success onto some honest party (away from the coalition's favor), then there exists an attack with which the same coalition can *pull* probability of success from that honest party (towards the coalition's favor). The proof mixes the attack with its mirror image, which yields an attack by Alice alone (resp. Bob alone) and hence is bounded by strong fairness. Therefore, there exists an attack that biases the output in the coalition's favor and contradicts CSP-fairness.

2.2 Statistical Leader Election

The feasibility result for $t < \frac{n}{2}$ (Theorem 6) is immediate from statistically secure MPC with guaranteed output delivery, so we only sketch the lower bound (Theorem 7), which we prove in two steps.

Step 1: two parties. For two parties, leader election corresponds to solving coin tossing where parties have opposite preferences. Electing P_1 or P_2 is just a relabeling of the coin 0 or 1, and under this relabeling the leader election preference profile ("I prefer myself") becomes the coin-tossing profile in which P_1 prefers 0 and P_2 prefers 1. Every strategy of one protocol is a strategy of the other, and every expected utility is the same number in both; hence a two-party protocol is CSP-, MM-, or strongly fair leader election if and only if its relabeling is CSP-, MM-, or strongly fair coin tossing (Lemma 3). Liu-Zhang, Masserova, Ribeiro and Thyagarajan [LZMRT26] recently proved that statistically CSP-fair two-party coin tossing with one corruption is impossible for every preference by showing that one of the two parties can force the honest party's output to its own preferred bit with all but negligible probability. Two-party leader election with one corruption is therefore impossible as well (Corollary 3).

Step 2: n parties. An n-party protocol tolerating $\frac{n}{2}$ corruptions yields a two-party protocol tolerating one corruption by standard party emulation. Split the parties into two groups of $\frac{n}{2}$ parties each, and let Alice run all parties of the first group and Bob all parties of the second, forwarding the messages that cross between the groups; Alice is elected if the leader of the n-party protocol lies in her group, and Bob otherwise. In an honest execution each of them is elected with probability $\frac{1}{2}$, so this is a two-party leader election protocol. A corrupt Alice is exactly a coalition of $\frac{n}{2}$ parties in the original protocol, and her probability of being elected is exactly that coalition's utility, so CSP-fairness against $\frac{n}{2}$ corruptions transfers to CSP-fairness of the two-party protocol against one corruption, contradicting Step 1. For odd n, the groups have sizes $\lceil \frac{n}{2} \rceil$ and $\lfloor \frac{n}{2} \rfloor$ and Alice is elected with probability slightly above $\frac{1}{2}$; nevertheless, the argument goes through unchanged because the attack of [LZMRT26] only uses that the two parties agree on their output, and still lets one of them win with all but negligible probability.

2.3 Perfect Leader Election

For $n > 3t$, perfectly secure MPC gives strongly fair leader election directly. We show that perfect CSP-fairness can also be achieved at the boundary $n = 3t$.

We first present a protocol for three parties and one corruption. We then extend the idea to $n = 3t$ parties using the weighted tournament from our computational construction, before explaining why this is the limit for perfect CSP-fairness and why perfect MM-fairness cannot reach even this boundary.

2.3.1 Perfect CSP-Fairness at $n = 3t$

A first example: three parties and one corruption. Consider the following attempt at generating a uniform leader. Party P_1 privately sends a uniform $r \in \{0, 1, 2\}$ to P_2. Next, P_3 broadcasts a uniform $s \in \{0, 1, 2\}$, after which P_1 and P_2 broadcast their respective copies of r. If the copies are valid and agree, the leader is $P_{1+((r+s) \bmod 3)}$. Clearly, when everyone follows the protocol, each party is elected with probability $1/3$.

The difficulty is deciding what to do when the two copies disagree. Such a disagreement tells us that either P_1 or P_2 misbehaved, but does not tell us which one. For our purposes, however, identifying the deviator is unnecessary: we can simply elect P_3! Since at most one party is corrupt, this outcome gives the deviator utility zero. By the same reasoning, if P_3 refuses to broadcast $s \in \{0, 1, 2\}$, we can simply elect P_1.

Notice that this protocol does not protect every honest party: A corrupt P_1 can deliberately cause a disagreement and ensure that P_2 is never elected. This is consistent with CSP-fairness: while the deviation harms an honest party, it *does not benefit the deviator*. It would, however, violate MM-fairness.

Extending the idea to $n = 3t$ parties. The three-party example uses the fact that a disagreement between two parties leaves a third party who must be honest. With several corruptions, electing an arbitrary party outside the disagreement need not punish the coalition. We therefore use the freedom to redistribute probability among honest parties in a different way.

To extend the three-party intuition, we return to the tournament framework used in our computational construction. The question is how to implement the matches between the two candidates with perfect security. Simply delegating the computation of our computational version to a perfectly secure MPC does not resolve the problem: this would require $n > 3t$.

Our solution is to have the parties *outside* a match choose its winner. For a match between P_i and P_j, the other $3t - 2$ original parties run a perfectly secure MPC protocol with guaranteed output delivery against $t - 1$ corruptions. Whenever at least one candidate is corrupt, excluding the candidates leaves at most $t - 1$ corrupt MPC participants. As $3t - 2 > 3(t - 1)$, the MPC is secure in the matches that can advance a corrupt candidate.

If both candidates are honest, the MPC may contain all t corrupt parties and thus exceed its security threshold. In this case, it does not matter for CSP-fairness which of the two is elected winner, but all honest parties must still agree on the winner. To ensure agreement, we require the MPC participants to broadcast their outputs after a fixed timeout, using $\perp$ if they have not obtained a valid output by that time. A candidate supported by at least $2t - 1$ matching reports wins; otherwise, a fixed candidate from the pair wins.

Whenever the MPC's security threshold is respected, at least $2t-1$ honest participants report its output, so this additional step preserves the winner selected by the MPC. Even when the threshold is exceeded, two candidates cannot both receive $2t - 1$ reports among the $3t - 2$ participants. Since all honest parties observe the same broadcasts, the rule therefore ensures agreement on one of the two honest candidates.

Why the tournament remains CSP-fair. We apply the earlier weight argument to the coalition's total weight. Matches between honest candidates leave this weight unchanged, while matches involving a corrupt candidate preserve its conditional expectation. Its initial value is the coalition's size, and its final value is n precisely when the coalition wins. Thus, the coalition's winning probability equals its initial fraction of the total weight, establishing CSP-fairness.

This replacement of the match-selection procedure also yields a **perfect best-of-both-worlds guarantee**. With at most $t-1$ corruptions, every MPC instance is secure, even when both candidates are honest. The weight argument then protects every party's winning probability, giving strong fairness. With t corruptions, it continues to protect the coalition's utility.

2.3.2 Limits of Perfect CSP- and MM-Fairness We now show that perfect CSP-fairness is impossible for $n < 3t$, and that perfect MM-fairness is impossible even at $n = 3t$. As in the computational lower bound, a CSP contradiction requires a deviation that benefits the coalition. Here, an additional difficulty is deciding *when* to attempt the deviation.

Even finding an execution on which a deviation turns a coalition's loss into a win does not settle this issue. The coalition sees only its own view and may be unable to recognize that execution. The same deviation, applied on other executions consistent with its view, might turn wins into losses. We therefore need to ensure that the adversary attempts its deviation only when there is no winning outcome to sacrifice.

Filtering deviations using local knowledge. Our approach is to identify a point at which the coalition's view certifies that it would lose if everyone continued honestly. An adversary that deviates only on such views cannot reduce its own utility: outside these views it behaves honestly, and on them its honest utility is already zero. Any positive probability of winning after the deviation therefore gives a strict improvement.

To find the right point, consider the earliest round r after which, in every honest execution, every size-t coalition knows whether the elected leader belongs to it. Bounded termination ensures that such a round exists, and CSP-fairness rules out this knowledge at the start of the protocol. Immediately before round r, some coalition A consequently has the same view in two honest executions X, Y, but wins in X and loses in Y.

We will combine these executions by replacing the pre-round-r states of some parties in a corrupt coalition with states from another honest execution. To justify this replacement using CSP-fairness, we need an adversary that performs it only when its own view certifies that honest continuation would elect someone outside the coalition.

This introduces a subtlety: the replacement concerns states immediately before round r, but the coalition may learn whether it wins only after that round. To resolve this, we consider a communication schedule with one sender per round. If the sender belongs to the coalition D, sending its own messages gives D no new information, so it already knows whether it wins before the round. Otherwise, D can wait for the round's messages and check whether its resulting view certifies a loss. If so, it installs the replacement pre-round states and processes those same messages on them; otherwise, it continues honestly. Since the sender is outside D, its messages are unaffected by the replacement, so this delayed intervention reproduces the intended hybrid.

This gives our *filtered state replacement* property: a fixed replacement of a coalition's pre-round-r states cannot turn an honest loss into a win. Crucially, the adversary need not recognize a complete execution. Its triggering view certifies a loss in every honest execution consistent with that view, and it behaves honestly whenever the trigger is absent.

Why CSP-fairness requires $n \geqslant 3t$. To see where the threshold arises, focus on $2t < n < 3t$. Let A be the size-t coalition above, and denote the leaders of X, Y by $a \in A$ and $b \notin A$, respectively. Partition the remaining parties into B, C with $|B| = t$ and $|C| = n - 2t < t$, choosing B to contain b.

We combine the states of these executions into the hybrid

$$H = \left(\mathsf{view}_A^X = \mathsf{view}_A^Y,\ \mathsf{view}_B^Y,\ \mathsf{view}_C^X\right),$$

where all views are taken immediately before round r.

This hybrid has two useful interpretations. Starting from X, a corrupt B can replace its states with those from Y. Since the original leader a is outside B, filtered state replacement implies that the hybrid's leader d is outside B as well. Starting instead from Y, replacing only C's states again gives H. Here we use the fact that $|C| < t$: we can corrupt C together with d, padding the coalition to size t while excluding b. Its original leader b is outside the coalition, but its new leader d is inside, contradicting the same replacement property. Since $n > 2t$, some party is honest in both interpretations; its identical continuation ensures that the hybrid has the same leader in each.

The case $n \leqslant 2t$ admits a simpler two-group argument: the complement of A fits inside a corrupt coalition, so replacing its states directly turns X into Y. At $n = 3t$, however, the three-group argument loses exactly the extra corruption it needs: C already has size t, and cannot be enlarged to include a leader outside it. Indeed, CSP-fairness only requires the hybrid's leader to lie outside both B and C, which leaves the possibility of electing someone in A.

Why MM-fairness cannot reach the boundary. MM-fairness protects each honest party from having its winning probability reduced. This gives a stronger consequence of filtering. Suppose a coalition's view certifies that honest continuation would elect a particular honest party a. A deviation attempted only on this view cannot increase a's winning probability: whenever it is attempted, a would already win with certainty. If it ever changes the winner, it strictly reduces a's overall winning probability. Thus, the replacement must preserve the identity of the honest leader, rather than merely keep the winner outside the coalition.

For this argument, we choose the critical round using knowledge of the leader's identity, rather than knowledge of coalition membership. We again obtain a size-t coalition A with identical pre-round views in two executions, now with distinct leaders a, b. For $n = 3t$, partition the other parties into size-t groups B, C so that $a \notin B$ and $b \notin C$, and form the same type of hybrid. Viewed as a replacement by B starting from X, it must elect a. Viewed as a replacement by C starting from Y, it must elect b. The parties in A are honest in both interpretations and have identical continuations, making these requirements incompatible. Unlike the CSP argument, this contradiction needs no extra corruption. Together with the impossibility below $3t$, it rules out perfect MM-fairness, and hence perfect GT-fairness and strong fairness, for leader election with $n \leqslant 3t$.

3 Definitions of Fair Sampling

This paper uses the definitions of cooperative-strategy-proof (CSP) and maximin (MM) fairness introduced in the seminal work on game theoretic fairness [CGL+18]. It also defines a new notion of fairness — *best of both worlds (BoBW)* fairness. All three definitions are stated in Section 3.3.

3.1 Preference Profiles and Utility Functions

At a high level, a distributed sampling protocol is game theoretically fair if there is no incentive for a set of malicious parties, called a *coalition*, to deviate from the prescribed protocol. This formulation

assumes the existence of a public preference profile Pref which associates to each of the n parties $i \in [n]$ a (set of) preferred output(s) $\mathsf{Pref}[i]$. The preference profile defines a binary utility function indicating the satisfaction of the preference of party i, as in Equation (1).

$$\mathsf{util}_i^{\mathsf{Pref}}(x) = \begin{cases} 1 & x \in \mathsf{Pref}[i] \\ 0 & x \notin \mathsf{Pref}[i] \end{cases} \tag{1}$$

For a subset of parties $J \subseteq [n]$ the joint utility is the sum of the individual utility functions $\mathsf{util}_J^{\mathsf{Pref}}(x) := \sum_{j \in J} \mathsf{util}_j^{\mathsf{Pref}}(x)$.

3.2 Strategies

A protocol Π defines the code (or alternatively, the strategy) $\mathring{\pi}_i$ of each honest party $i \in [n]$ in an execution of the protocol. More generally, a strategy is denoted π_j with a subscript indicating the party $j \in [n]$ or set of parties $J \subseteq [n]$ following that strategy. (We use the subscript $-i$ to denote $[n] \backslash \{i\}$, and the subscript $-J$ to denote $[n] \backslash J$.) We let $\Pi\langle \pi_J, \pi_{-J} \rangle$ denote an execution of the protocol between the parties in J and all other parties. We use $x \leftarrow \Pi\langle \cdots \rangle$ to denote that x is the output of the execution.

The expected utility of a set of parties J adopting strategy π_J in an execution of Π is defined in Equation (2). The parameters Pref (the preference profile) and Π (the protocol) are not included when they are clear from context.

$$\mathsf{U}_J^{\mathsf{Pref},\Pi}(\pi_J, \pi_{-J}) = \mathsf{E}\left[\mathsf{util}_J^{\mathsf{Pref}}(x) \;\middle|\; x \leftarrow \Pi\langle \pi_J, \pi_{-J} \rangle \right] \tag{2}$$

3.3 Fairness

3.3.1 Strong Fairness The traditional cryptographic notion of strong fairness [Cle86], i.e., unbiasability, is restated in Definition 1 such that it is compatible with the formalism above. Roughly, the definition requires that the advantage of an adversary in biasing the output away from the intended distribution be negligible.

Definition 1 (t-Strong Fairness). *For a distribution D, a protocol Π between a set of n parties is strongly fair against coalitions of size t iff for all $J \subset [n]$ of size at most t and all strategies π_J, there exists a function $f(\cdot)$ such that for all $x \in \mathsf{supp}(D)$:*

$$\left| \mathsf{Pr}\left[x = X | X \leftarrow D\right] - \mathsf{Pr}\left[x = X | X \leftarrow \Pi\langle \mathring{\pi}_{-J}, \pi_J \rangle\right] \right| \leqslant f(\lambda)$$

The protocol satisfies perfect strong fairness if $f(\lambda) = 0$, statistical strong fairness if $f(\lambda)$ is negligible, and computational strong fairness if $f(\lambda)$ is negligible and the strategy π_J is required to be efficient, i.e. probabilistic polynomial time (PPT).

3.3.2 Existing Notions of Game Theoretic Fairness Cooperative strategy proofness (CSP-fairness) is the requirement that a coalition cannot increase its expected utility. Maxi-min fairness (MM-fairness) is the requirement that a coalition cannot decrease the utility of any individual honest party. A protocol is referred to as *game theoretically fair* for (t, Pref) if the protocol is both CSP- and MM- fair, as per Definitions 2 and 3. Notably, strong fairness implies MM- fairness, as preserving the entire output distribution preserves every honest party's winning probability.

Definition 2 (t-CSP Fairness). *A protocol Π between a set of n parties with preference profile Pref is coalition-strategy-proof (CSP) fair against coalitions of size t iff for all $J \subset [n]$ of size at most t and all strategies π_J, there exists a function $f(\cdot)$ such that:*

$$\mathsf{U}_J(\pi_J, \mathring{\pi}_{-J}) \leqslant \mathsf{U}_J(\mathring{\pi}_{[n]}) + f(\lambda). \tag{3}$$

The protocol satisfies perfect CSP fairness $f(\lambda) = 0$, statistical CSP fairness if $f(\lambda)$ is negligible, and computational CSP fairness if $f(\lambda)$ is negligible and the strategy π_J is required to be PPT.

Definition 3 (t-MM Fairness). *A protocol Π between a set of n parties with preference profile Pref is maxi-min (MM) fair against coalitions of size t iff for all $J \subset [n]$ of size at most t and all strategies π_J, there exists a function $f(\cdot)$ such that for all $i \in [n] \backslash J$:*

$$\mathsf{U}_i(\mathring{\pi}_{-J}, \pi_J) \geqslant \mathsf{U}_i(\mathring{\pi}_{[n]}) - f(\lambda).$$

The protocol satisfies perfect MM fairness $f(\lambda) = 0$, statistical MM fairness if $f(\lambda)$ is negligible, and computational MM fairness if $f(\lambda)$ is negligible and the strategy π_J is limited to be PPT.

Definition 4 (t-GT Fairness). *A protocol Π between a set of n parties with preference profile Pref is (perfect, statistical, computational) game theoretically fair against coalitions of size t iff Π is both (perfect, statistical, computational) t-CSP fair and t-MM fair as in Definitions 2 and 3, respectively.*

3.3.3 Novel Notion: Best of Both Worlds Fairness In this paper, we introduce a new notion of fairness: *best-of-both-worlds (BoBW)* fairness. BoBW fairness is parametrized by two thresholds: t_S and t_{GT}. BoBW simultaneously requires t_S-strong fairness (Definition 1) and t_{GT}-GT fairness (Definition 4).

Definition 5 ((t_S, t_{GT})-BoBW Fairness). *For a distribution D the protocol Π between a set of n parties with preference profile Pref is (perfect, statistical, computational) BoBW fair for thresholds $t_S, t_{GT} < n$ iff Π is both (perfect, statistical, computational) t_S-strongly fair (Definition 1) and t_{GT}-GT fair (Definition 4).*

We also consider BoBW-fair protocols which, for the lower threshold t_{GT}, only achieve one of CSP or MM fairness, but not both. We refer to those protocols as *S/CSP (t_S, t_{GT})-BoBW fair* or *S/MM (t_S, t_{GT})-BoBW fair*, respectively.

Definition 6 (S/CSP (t_S, t_{GT})-BoBW Fairness). *For a distribution D the protocol Π between a set of n parties with preference profile Pref is (perfect, statistical, computational) S/CSP BoBW fair for thresholds $t_S, t_{GT} < n$ iff Π is both (perfect, statistical, computational) t_S-strongly fair (Definition 1) and t_{GT}-CSP fair (Definition 2).*

Definition 7 (S/MM (t_S, t_{GT})-BoBW Fairness). *For a distribution D the protocol Π between a set of n parties with preference profile Pref is (perfect, statistical, computational) S/MM BoBW fair for thresholds $t_S, t_{GT} < n$ iff Π is both (perfect, statistical, computational) t_S-strongly fair (Definition 1) and t_{GT}-MM fair (Definition 3).*

3.4 Leader Election

A weighted leader election protocol associates a weight $w_i \in \mathbb{N}$ to each party $i \in [n]$ and samples (elects) a random party from the group according to the assigned weights, outputting each winner $j \in [n]$ with probability $\frac{w_j}{\sum_{i\in[n]} w_i}$. When a leader election is not specified to be weighted then the weights are equal, and the output is uniform over $[n]$. At the end of the protocol, honest parties are in agreement on who is the leader. In any leader election protocol the preference profile is such that every participant wishes to be chosen; that is,

$$\mathsf{Pref}_{\mathsf{LE}}[i] = \{i\}. \tag{4}$$

Remark 1. For leader election t-MM fairness implies t-CSP fairness, and thereby t-GT fairness. This implication is an immediate consequence of the structure of the preference profile. For all outputs, at most one party's preference is satisfied; that is, for all $x \in [n]$ and all strategies π,

$$\mathsf{U}_{[n]}^{\mathsf{Pref}_{\mathsf{LE}}}\left(\pi_{[n]}\right) = \sum_{i\in[n]} \mathsf{util}_i^{\mathsf{Pref}_{\mathsf{LE}}}(x) = 1. \tag{5}$$

Consequently, a t-MM protocol Π for such a preference profile satisfies the following for all subsets of parties $J \subseteq [n]$ of cardinality at most t,

$$\begin{aligned}
\mathsf{U}_{[n]}\left(\mathring{\pi}_{[n]}\right) &= \mathsf{U}_J\left(\mathring{\pi}_{[n]}\right) + \mathsf{U}_{-J}\left(\mathring{\pi}_{[n]}\right) && (6)\\
&\leqslant \mathsf{U}_J\left(\mathring{\pi}_{[n]}\right) + \sum_{i\in -J} \mathsf{U}_i\left(\mathring{\pi}_{-J}, \pi_J\right) + (n - |J|)\cdot f(\lambda) && (7)\\
&= \mathsf{U}_J\left(\mathring{\pi}_{[n]}\right) + \mathsf{U}_{[n]}\left(\mathring{\pi}_{-J}, \pi_J\right) - \mathsf{U}_J\left(\mathring{\pi}_{-J}, \pi_J\right) + (n - |J|)\cdot f(\lambda) && (8)\\
\implies \mathsf{U}_J\left(\mathring{\pi}_{-J}, \pi_J\right) &\leqslant \mathsf{U}_J\left(\mathring{\pi}_{[n]}\right) + (n - |J|)\cdot f(\lambda). && (9)
\end{aligned}$$

where Equations (6) and (8) are from the definition of U in Equation (2), Equation (7) is from Definition 3 of t-MM, and the implication of Equation (9), which is Definition 2 of t-CSP, is from the property in Equation (9).

4 Computational Best of Both Worlds Leader Election

In Section 4.2, we present a BoBW-fair two-party leader election protocol. Normally, two-party leader election is thought of as being run between those two parties; instead, we run it among all n parties, and require $t_S \leqslant t_{GT}$[7] and $t_S + t_{GT} < n$. To this end, in Section 4.1 we introduce a building block: a verifiable secret sharing (VSS) with two reconstruction modes, which we dub dual-reconstruction (DR) VSS. In Section 4.3, we plug the two-party protocol into a tournament to obtain BoBW-fair n-party leader election. Finally, in Section 5, we prove that the bounds for which these constructions work are tight: for $t_S + t_{GT} \geqslant n$, no such protocols can exist.

4.1 Verifiable Secret Sharing

4.1.1 Tools and Building Blocks. We describe the necessary tools and building blocks that we need for our dual-reconstruction VSS scheme.

Ideal commitments with designated openers. We use an ideal commitment functionality $\mathcal{F}_{\mathsf{com}}$. The functionality takes as parameters a message space MS, a committer P_i, a set $\mathcal{O}$ of *openers*, and a set $\mathcal{P}$ of receivers.

[7] If $t_S \geqslant t_{GT}$, running a strongly fair protocol for t_S immediately satisfies (t_S, t_{GT})-BoBW fairness.

Functionality $\mathcal{F}_{\mathsf{com}}(\mathsf{MS}, P_i, \mathcal{O}, \mathcal{P})$: Ideal Commitment with Designated Openers

The functionality with session identifier sid is parameterized by a message space MS, a committer P_i, a set of openers $\mathcal{O} \subseteq \mathcal{P}$ with $P_i \in \mathcal{O}$, and the set of receivers $\mathcal{P}$. It interacts with the parties in $\mathcal{P}$ and the ideal adversary $\mathcal{A}^*$ via the queries below:

On $(\texttt{set}, \mathcal{O}^*)$ **from** $\mathcal{A}^*$**.** If P_i is corrupted, set $\mathcal{O} = \mathcal{O}^*$.
On $(\texttt{commit}, m)$ **from** P_i**.** Ignore the query if $m \notin \mathsf{MS}$ or if there is a record $(\texttt{committed}, \cdot)$. Otherwise,
- record $(\texttt{committed}, m)$.
- send $(\texttt{committed}, m)$ to every party in $\mathcal{O}$.
- send $(\texttt{committed})$ to every party in $\mathcal{P} \backslash \mathcal{O}$ and to $\mathcal{A}^*$.

On $(\texttt{open})$ **from** $P \in \mathcal{O}$**.** Retrieve the record $(\texttt{committed}, m)$ and ignore the query if the record does not exist. Send $(\texttt{opened}, m)$ to every party in $\mathcal{P}$ and to $\mathcal{A}^*$.

The functionality can be realized using non-malleable concurrent commitment schemes [LP11], and by letting the committer broadcast the commitment and sending the opening information to every opener in the commit phase, and broadcast the opening information in the open phase.

Symmetric bivariate polynomials. Throughout this section $\mathbb{F} := \mathbb{Z}_p$ for a prime $p > n$, so that $0, 1, \ldots, n$ are distinct field elements; we identify the elements of $\mathbb{F}$ with the integers $0, 1, \ldots, p-1$. A *symmetric bivariate polynomial of degree* t is a polynomial $S(x, y) = \sum_{k,\ell=0}^{t} a_{k,\ell} x^k y^\ell$ over $\mathbb{F}$ with $a_{k,\ell} = a_{\ell,k}$ for all k, ℓ. For $i \in [1, n]$ we call the univariate polynomial $S(x, i)$ the *i-th row* of S; rows satisfy $S(j, i) = S(i, j)$.

For a set $T \subseteq [1, n]$ and $i \in T$ we write $\lambda_i^T := \prod_{j \in T \setminus \{i\}} \frac{j}{j-i}$ for the Lagrange coefficients for evaluation at 0, so that $g(0) = \sum_{i \in T} \lambda_i^T \cdot g(i)$ for every univariate polynomial g of degree less than $|T|$. For $s \in \mathbb{F}$ let $\mathcal{S}_s$ denote the set of symmetric bivariate polynomials S of degree t with $S(0, 0) = s$. We use the standard fact that for every $s \in \mathbb{F}$ and every set $C \subseteq [1, n]$ with $|C| \leqslant t$, the distribution of the rows $(S(x, i))_{i \in C}$ of a uniformly random $S \in \mathcal{S}_s$ does not depend on s.

4.1.2 Definition of DR-VSS. We introduce our definition of a dual-reconstruction verifiable secret sharing (DR-VSS) protocol. A DR-VSS protocol among n parties $\mathcal{P} = \{1, \ldots, n\}$ with a designated party $D \in \mathcal{P}$ consists of a triplet of protocols (Share, JRec, DRec).

Share$(m) \to (\langle m \rangle_1, \ldots, \langle m \rangle_n)$ is a protocol that the dealer D with input m runs with the share-holders to create n shares of m.
JRec$(\langle m \rangle_1, \ldots, \langle m \rangle_n) \to m$ is a dealer-independent reconstruction protocol run by the share-holders and returns a message m.
DRec$(\langle m \rangle_1, \ldots, \langle m \rangle_n) \to m$ is a dealer-powered reconstruction protocol run by the dealer and the share-holders and returns a message m.

Our definition of DR-VSS has a parameter t and guarantees that there are two ways to open the secret after the sharing phase: the committed value is guaranteed to be opened when there are at least $t + 1$ honest parties or when the dealer is honest.

Definition 8 (t-DR-VSS). *Protocol Π = (Share, JRec, DRec) between a set of n parties and a designated dealer D with input $s \in \mathbb{F}$ is a secure t-DR-VSS protocol if the following properties are satisfied:*

(Privacy) *If D is honest and at most t parties are corrupted, the adversary's view during protocol* Share *reveals no information about s. More formally, the adversary's views for any two inputs $s, \tilde{s} \in \mathbb{F}$ are (perfectly, statistically, computationally) indistinguishable.*

(Commitment) *At the end of the sharing phase there is a value $s' \in \mathbb{F}$, determined by the transcript of the sharing phase, such that in every execution of* JRec *and in every execution of* DRec *all honest parties output the* same *value, which lies in $\{s', \perp\}$. If D is honest, $s' = s$. Moreover:*

(Joint Reconstruction) *If at least $t+1$ parties are honest, all honest parties output s' at the end of* JRec*.*

(Dealer Reconstruction) *If D is honest, all honest parties output s' at the end of* DRec*.*

4.1.3 Protocol Description. We now describe a t-DR-VSS protocol. Note that Feldman VSS [Fel87] satisfies Definition 8 but relies on the additional property of commitment homomorphism. With Protocol 2 we propose a simple alternative which can leverage any instantiation of $\mathcal{F}_{\mathsf{com}}$.

The dealer commits to a symmetric bivariate polynomial S of degree t with $S(0,0) = s$ by committing separately to each value $S(i,j)$ with $1 \leqslant i \leqslant j \leqslant n$ in an instance of $\mathcal{F}_{\mathsf{com}}$ whose openers are D, P_i and P_j. Through the functionality, P_i thus learns the committed values of its row $S(x,i)$, and can check that the row forms a polynomial of degree at most t; if not, P_i complains, and the dealer responds by opening the entire row, which everyone can then check. Any dealer that fails to adequately respond to a complaint by opening is disqualified.

At the end of a sharing phase in which the dealer is not disqualified, every honest party P_i holds a row of degree at most t, and the rows of honest parties are consistent with each other.

Protocol 2: t-DR-VSS in the $\mathcal{F}_{\mathsf{com}}$-hybrid model

Parties $P_1, \ldots, P_n$ with dealer D. For $1 \leqslant i \leqslant j \leqslant n$, let $\mathcal{F}_{\mathsf{com}}^{\{i,j\}} := \mathcal{F}_{\mathsf{com}}(\mathbb{F}, D, \{D, P_i, P_j\}, \mathcal{P})$ with session identifier $(\mathsf{sid}, \{i,j\})$, and let $\mathcal{F}_{\mathsf{com}}^{\{j,i\}} := \mathcal{F}_{\mathsf{com}}^{\{i,j\}}$. The dealer's input is $s \in \mathbb{F}$. We say that $\{i,j\}$ *is opened* once $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ has delivered a message $(\mathtt{opened}, v_{i,j})$; set $v_{j,i} := v_{i,j}$. Row i is *complete* if $\{i,j\}$ is opened for every $j \in [1,n]$, and *valid* if it is complete and the n points $(j, v_{i,j})_{j \in [1,n]}$ lie on a polynomial of degree at most t.

Sharing Phase (Share):

1. **(Round 1, D: commit.)**
   (a) Sample a uniformly random symmetric bivariate polynomial of degree t with $S(0,0) = s$. For $i \in [1,n]$ let $f_i(x) := S(x,i)$.
   (b) For every $1 \leqslant i \leqslant j \leqslant n$ query $(\mathtt{commit}, S(i,j))$ to $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$.

   Every P_k receives $(\mathtt{committed})$ from every $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ with $k \notin \{i,j\}$, and $(\mathtt{committed}, v_{k,j})$ from $\mathcal{F}_{\mathsf{com}}^{\{k,j\}}$ for every $j \in [1,n]$. *P_k learns its row*
2. **(Round 2, each P_i: verify and complain.)** If, for some $j \in [1,n]$, P_i did not receive $(\mathtt{committed}, \cdot)$ from $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$, or if the points $(j, v_{i,j})_{j \in [1,n]}$ do not lie on a polynomial of degree at most t, broadcast $(\mathsf{complain}, i)$.
3. **(Round 3, D: resolve.)** For every i such that $(\mathsf{complain}, i)$ was broadcast, query $(\mathtt{open})$ to $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ for every $j \in [1,n]$.
4. **(Local, each P_k: decide.)**

(a) Set $\mathsf{disq} := 1$ if some $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ did not deliver $(\texttt{committed})$ or $(\texttt{committed}, \cdot)$ in Round 1, or if for some i with $(\mathsf{complain}, i)$ broadcast row i is not valid after Round 3; otherwise set $\mathsf{disq} := 0$.

(b) Output $\langle s \rangle_k := (\mathsf{disq}, (v_{k,j})_{j \in [1,n]})$.

Joint Reconstruction (JRec): each P_k does the following.

5. If $\mathsf{disq} = 1$, output 0. Otherwise query $(\texttt{open})$ to $\mathcal{F}_{\mathsf{com}}^{\{k,j\}}$ for every $j \in [1,n]$. *P_k opens its row*
6. Let $V \subseteq [1,n]$ be the set of indices i such that row i is valid, and for $i \in V$ let f_i be the polynomial of degree at most t through the points $(j, v_{i,j})_{j \in [1,n]}$. If $|V| \leqslant t$, output $\perp$. Otherwise let T be the set of the $t+1$ smallest elements of V and output

$$s' := \sum_{i \in T} \lambda_i^T \cdot f_i(0).$$

Dealer Reconstruction (DRec):

7. If $\mathsf{disq} = 1$, every party outputs 0. Otherwise D queries $(\texttt{open})$ to $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ for every $1 \leqslant i \leqslant j \leqslant n$. *$D$ opens everything*
8. Each P_k: if every row $i \in [1,n]$ is valid, let f_i be the corresponding polynomials, let $T := [1, t+1]$ and output $\sum_{i \in T} \lambda_i^T \cdot f_i(0)$. Otherwise output $\perp$.

Theorem 1. *Protocol 2 is a secure t-DR-VSS protocol (Definition 8) in the $\mathcal{F}_{\mathsf{com}}$-hybrid model for any t.*

Proof (of Theorem 1). We first fix some notation and make two observations, then verify the properties of Definition 8.

The committed value. Each instance of $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ holds at most a single record $(\texttt{committed}, v_{i,j})$, and every $(\texttt{committed}, \cdot)$ and $(\texttt{opened}, \cdot)$ message it ever delivers carries this value $v_{i,j}$; by construction $v_{i,j} = v_{j,i}$. If $\mathsf{disq} = 0$, then every $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ delivered $(\texttt{committed})$ in Round 1, so all records exist and are fixed from the end of Round 1 on; set $v_{i,j} := \perp$ for instances without a record. Call an index $i \in [1,n]$ *good* if $v_{i,j} \neq \perp$ for all j and the points $(j, v_{i,j})_{j \in [1,n]}$ lie on a polynomial f_i^* of degree at most t, and let G be the set of good indices. For $i, j \in G$ we have $f_i^*(j) = v_{i,j} = v_{j,i} = f_j^*(i)$, so if $|G| \geqslant t+1$ then there is a unique symmetric bivariate polynomial S^* of degree t with $S^*(x, i) = f_i^*(x)$ for all $i \in G$. We define

$$s' := \begin{cases} S^*(0,0) & \text{if } \mathsf{disq} = 0 \text{ and } |G| \geqslant t+1, \\ 0 & \text{otherwise.} \end{cases}$$

Since disq is determined by the broadcast messages of the sharing phase and the records are fixed at the end of Round 1, s' is determined at the end of the sharing phase.
We make two observations.

(O1) If $\mathsf{disq} = 0$, then every honest P_i holds values $(v_{i,j})_{j \in [1,n]}$ of its row at the end of the sharing phase, and $i \in G$. Indeed, P_i received $(\texttt{committed}, v_{i,j})$ from $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ for every j (otherwise $\mathsf{disq} = 1$). If these values did not lie on a polynomial of degree at most t, P_i would have complained in Round 2, and since row i can only ever be opened to the very same values, it could not be valid after Round 3, so disq would be 1.

(O2) If D is honest, then $v_{i,j} = S(i,j)$ for all i,j, so $G = [1,n]$ and $S^* = S$; every $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ delivers in Round 1; no honest party complains, since its row is a row of S; every complaint is resolved in Round 3 by opening a row of S, which is valid; hence $\mathsf{disq} = 0$ and $s' = S(0,0) = s$.

Commitment – agreement and outputs in $\{s', \perp\}$. Everything that determines an honest party's output in JRec or DRec – the broadcasts of the sharing phase, which fix disq, and the $(\texttt{opened}, \cdot)$ messages, which $\mathcal{F}_{\mathsf{com}}$ delivers to all parties – is seen identically by all honest parties, so they compute the same V, the same polynomials f_i, and hence the same output. It remains to show that any non-$\perp$ output equals s'. If $\mathsf{disq} = 1$, every honest party outputs $0 = s'$ (Steps 5 and 7). So let $\mathsf{disq} = 0$. Since an instance $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ only ever opens to its recorded value $v_{i,j}$, a valid row i satisfies $i \in G$ and $f_i = f_i^*$; that is, $V \subseteq G$. In JRec, a non-$\perp$ output requires $|V| \geqslant t+1$, hence $|G| \geqslant t+1$ and $T \subseteq G$, and the output is $\sum_{i \in T} \lambda_i^T f_i^*(0) = \sum_{i \in T} \lambda_i^T S^*(0,i) = S^*(0,0) = s'$, since $S^*(0,y)$ is a polynomial of degree at most t in y. In DRec, a non-$\perp$ output requires every row to be valid, hence $G = [1,n]$, and the output is $\sum_{i \in [1,t+1]} \lambda_i^{[1,t+1]} S^*(0,i) = S^*(0,0) = s'$ by the same argument.

Commitment – Joint Reconstruction with $t+1$ honest parties. Let $H \subseteq [1,n]$ be the set of honest parties, and assume that $|H| \geqslant t+1$. If $\mathsf{disq} = 1$ then every honest party outputs $0 = s'$ in Step 5. Otherwise, by (O1), $H \subseteq G$. In Step 5 every honest P_i opens $\{i,j\}$ for every j, so row i is complete, and it is valid because $i \in G$; hence $H \subseteq V$ and $|V| \geqslant t+1$, so no honest party outputs $\perp$, and by the previous paragraph every honest party outputs s'.

Commitment – Dealer Reconstruction with honest dealer. Suppose D is honest. By (O2), $\mathsf{disq} = 0$ and $s' = s$, and in Step 7 the dealer opens every $\{i,j\}$, after which every row i is valid with $f_i = S(x,i)$. Then, every honest party outputs $\sum_{i \in [1,t+1]} \lambda_i^{[1,t+1]} S(0,i) = S(0,0) = s$ in Step 8.

Privacy. Suppose D is honest and let $C \subseteq [1,n]$ be the set of corrupt parties, $|C| \leqslant t$. We determine the adversary's view of the sharing phase. In Round 1 the adversary receives $(\texttt{committed})$ from every $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$, which carries no information, and $(\texttt{committed}, S(i,j))$ from every $\mathcal{F}_{\mathsf{com}}^{\{i,j\}}$ with $i \in C$ or $j \in C$, i.e., exactly the rows $(S(x,i))_{i \in C}$. In Round 2 the honest parties do not complain, by (O2). In Round 3 the dealer opens the rows of the complaining parties, all of which lie in C; the adversary thus receives $(\texttt{opened}, S(i,j))$ only for pairs with $i \in C$, which it already knows. Hence the adversary's view is a function of the rows $(S(x,i))_{i \in C}$ of a uniformly random $S \in \mathcal{S}_s$ and of the adversary's own randomness, and by the fact recalled above the distribution of these rows does not depend on s. The views for any two inputs $s, \tilde{s}$ are therefore identically distributed.

4.2 BoBW Two-Party Leader Election

For $t_S \leqslant t_{GT}$ and $t_S + t_{GT} < n$, we implement a BoBW leader election that elects one leader out of two designated parties P_i and P_j with weights $k_i, k_j \in \mathbb{N}$, i.e., with probabilities $\frac{k_i}{k_i+k_j}$ and $\frac{k_j}{k_i+k_j}$, respectively. The protocol is run among all n parties and makes use of the verifiable secret sharing from Section 4.1.

Protocol 3: (t_S, t_{GT})-BoBW Two-Party Leader Election between parties P_i and P_j with weights k_i, k_j

Parties $P_1, \ldots, P_n$; candidates P_i and P_j with weights $k_i, k_j \in \mathbb{N}$, $k_i + k_j \leqslant n$; let $k := k_i + k_j$. For $l \in [n]$, let DRVSS_l denote an instance of Protocol 2 with dealer P_l, share-holders $P_1, \ldots, P_n$, session identifier (sid, l) and threshold t_l, where $t_l := t_{GT}$ for $l \in \{i, j\}$ and $t_l := t_S$ otherwise.

Commit:

1. Each party P_l picks $x_l \leftarrow \{0, \ldots, k-1\}$ uniformly at random.
2. Each party P_l runs Share of DRVSS_l as the dealer with input x_l; every party participates as a share-holder in every instance. *the candidates share with threshold t_{GT}, all other parties with threshold t_S*

Open:

3. For every $l \in [n]$, all parties run JRec and then DRec of DRVSS_l; let y_l^{jrec} and y_l^{drec} be the respective outputs.

Output: Each party does the following:

4. If $(y_i^{\mathsf{jrec}}, y_i^{\mathsf{drec}}) = (\perp, \perp)$, output P_j. Otherwise, if $(y_j^{\mathsf{jrec}}, y_j^{\mathsf{drec}}) = (\perp, \perp)$, output P_i. *a candidate who refuses to open loses; if both refuse, P_j wins*
5. Otherwise, for every $l \in [n]$ let $y_l := y_l^{\mathsf{drec}}$ if $y_l^{\mathsf{drec}} \neq \perp$, let $y_l := y_l^{\mathsf{jrec}}$ if $y_l^{\mathsf{drec}} = \perp \neq y_l^{\mathsf{jrec}}$, and let $y_l := 0$ if both are $\perp$. Let $y := \sum_{l \in [n]} y_l \bmod k$. Output P_i if $y < k_i$, and P_j otherwise.

Theorem 2. *Let $t_S \leqslant t_{GT}$ and $t_S + t_{GT} < n$, let $i \neq j \in [n]$, and let $k_i, k_j \in \mathbb{N}$ with $k_i + k_j \leqslant n$. In the $\mathcal{F}_{\mathsf{com}}$-hybrid model, Protocol 3 is a perfect (t_S, t_{GT})-BoBW two-party leader election where parties P_i and P_j have weights k_i, k_j.*

Proof (of Theorem 2). We have to show that all honest parties output the same element of $\{i, j\}$, that Protocol 3 is t_S-strongly fair (Definition 1), and that it is t_{GT}-GT fair (Definition 4). By Remark 1, for the latter it suffices to show t_{GT}-MM fairness (Definition 3).

Fix a coalition $J \subseteq [n]$ with $|J| \leqslant t_{GT}$ and a strategy π_J; let $\mathcal{H} := [n] \backslash J$ be the honest parties. We write out for the leader chosen by the honest parties and $k := k_i + k_j$. The n instances DRVSS_l are run in the $\mathcal{F}_{\mathsf{com}}$-hybrid model with distinct session identifiers, and the guarantees of Theorem 1 are perfect, so they hold simultaneously for all instances; we use them in the following form.

(V1) For every $l \in [n]$, at the end of the commit phase there is a value $s'_l \in \mathbb{F}$, determined by the transcript of the commit phase, such that every honest party's outputs satisfy $y_l^{\mathsf{jrec}}, y_l^{\mathsf{drec}} \in \{s'_l, \perp\}$, and all honest parties obtain the *same* pair $(y_l^{\mathsf{jrec}}, y_l^{\mathsf{drec}})$. If $l \in \mathcal{H}$, then $s'_l = x_l$. This is the commitment property of DRVSS_l.
(V2) If $l \in \mathcal{H}$, then $y_l^{\mathsf{drec}} = x_l$ (dealer reconstruction with honest dealer); in particular $(y_l^{\mathsf{jrec}}, y_l^{\mathsf{drec}}) \neq (\perp, \perp)$. If $l \notin \{i, j\}$, then $y_l^{\mathsf{jrec}} = s'_l$, since $t_l = t_S$ and there are at least $n - t_{GT} \geqslant t_S + 1$ honest parties (joint reconstruction). If moreover $|J| \leqslant t_S$, then $y_l^{\mathsf{jrec}} = s'_l$ for every $l \in [n]$, since $t_l \leqslant t_{GT}$ and there are at least $n - t_S \geqslant t_{GT} + 1$ honest parties.
(V3) If $l \in \mathcal{H}$ and $|J| \leqslant t_l$, the adversary's view of the entire commit phase is independent of x_l. This holds by the (perfect) privacy of DRVSS_l, since the x_l are sampled independently and the honest parties' messages in the other instances do not depend on x_l. In particular, this applies to every honest l if $|J| \leqslant t_S$, and to every honest candidate $l \in \{i, j\} \cap \mathcal{H}$ if $|J| \leqslant t_{GT}$.

Agreement. By (V1), all honest parties hold the same pairs $(y_l^{\mathsf{jrec}}, y_l^{\mathsf{drec}})_{l\in[n]}$, so they compute the same values y_l and y and output the same $\mathsf{out} \in \{i, j\}$. Moreover, by (V1), $y_l \in \{s'_l, 0\}$ for every l, and $y_l = s'_l$ whenever $(y_l^{\mathsf{jrec}}, y_l^{\mathsf{drec}}) \neq (\perp, \perp)$.

The sum Y. Let

$$Y := \sum_{l\in[n]} s'_l \bmod k = \Big(\sum_{l\in\mathcal{H}} x_l + \sum_{l\in J} s'_l\Big) \bmod k,$$

where we used (V1) for the honest dealers.

Strong fairness for t_S. Suppose $|J| \leqslant t_S$. By (V2), no pair $(y_l^{\mathsf{jrec}}, y_l^{\mathsf{drec}})$ equals $(\perp, \perp)$, so Step 4 does not apply, and $y_l = s'_l$ for every l by agreement; hence $y = Y$, and $\mathsf{out} = i$ if and only if $Y < k_i$. Since Y is uniform over $\mathbb{Z}_k$ (all VSS values are reconstructed and the values by J are independent of the honest values), we get $\Pr[\mathsf{out} = i] = \frac{k_i}{k}$ and $\Pr[\mathsf{out} = j] = \frac{k_j}{k}$, exactly the distribution D_{k_i,k_j}; this is t_S-strong fairness with $f(\lambda) = 0$. Applied to $J = \varnothing$, it also shows $\mathsf{U}_i(\mathring{\pi}_{[n]}) = \frac{k_i}{k}$ and $\mathsf{U}_j(\mathring{\pi}_{[n]}) = \frac{k_j}{k}$, while $\mathsf{U}_l(\mathring{\pi}_{[n]}) = 0$ for $l \notin \{i, j\}$.

MM fairness for t_{GT}. Suppose $|J| \leqslant t_{GT}$ and let $l \in \mathcal{H}$. If $l \notin \{i, j\}$, then $\mathsf{util}_l^{\mathsf{Pref}_{\mathsf{LE}}}(\mathsf{out}) = 0$ in every execution, so $\mathsf{U}_l(\mathring{\pi}_{-J}, \pi_J) = 0 = \mathsf{U}_l(\mathring{\pi}_{[n]})$ and there is nothing to show. So let, without loss of generality, $l = i$ be an honest candidate. By (V2), $(y_i^{\mathsf{jrec}}, y_i^{\mathsf{drec}}) \neq (\perp, \perp)$, and we distinguish two cases.

> **Case 1: $(y_j^{\mathsf{jrec}}, y_j^{\mathsf{drec}}) = (\perp, \perp)$.** Then $\mathsf{out} = i$ by Step 4. By (V2), this case requires $j \in J$.
> **Case 2: no pair of a candidate equals $(\perp, \perp)$.** By (V2), no pair of a non-candidate equals $(\perp, \perp)$ either, so $y_l = s'_l$ for all l, $y = Y$, and $\mathsf{out} = j$ if and only if $Y \geqslant k_i$.

Since P_i is honest, Y is uniform over $\mathbb{Z}_k$ by the claim above, and therefore

$$\Pr[\mathsf{out} = j] = \Pr[\text{Case 2} \wedge Y \geqslant k_i] \leqslant \Pr[Y \geqslant k_i] = \frac{k_j}{k}.$$

Intuitively, the coalition may decide, possibly after learning Y during the open phase, whether Case 1 occurs, but doing so elects the honest candidate.

Since $\mathsf{util}_i^{\mathsf{Pref}_{\mathsf{LE}}}(\mathsf{out}) = 1$ if and only if $\mathsf{out} = i$, and $\mathsf{out} \in \{i, j\}$, we have:

$$\mathsf{U}_i(\mathring{\pi}_{-J}, \pi_J) = 1 - \Pr[\mathsf{out} = j] \geqslant 1 - \frac{k_j}{k} = \frac{k_i}{k} = \mathsf{U}_i(\mathring{\pi}_{[n]}),$$

which is t_{GT}-MM fairness with $f(\lambda) = 0$. With Remark 1 this gives t_{GT}-GT fairness, which completes the proof.

4.3 BoBW Multi-Party Leader Election

To achieve BoBW-fair n-party leader election, we invoke the tournament protocol from [CCWS21,BZ17], where in each step of the tournament we run a BoBW two-party leader election (Protocol 3) with the candidates' current weights.

Protocol 4: Computational BoBW Leader Election for $t_S + t_{GT} < n$

Building Blocks: For $i \neq j \in [n]$ and $k_i, k_j \in \mathbb{N}$, let $\Pi_{\mathsf{2LE}}(i, j, k_i, k_j)$ be a (t_S, t_{GT}, i, j)-BoBW two-party leader election with weights k_i, k_j.

Protocol Execution:

1. Initialize the set of candidates as $\mathcal{C} \leftarrow \{P_1, \ldots, P_n\}$ and set $w_i \leftarrow 1$ for every $i \in [n]$.
2. While $|\mathcal{C}| > 1$, perform the following steps:
   i. Partition $\mathcal{C}$ into ordered pairs according to a fixed public rule. If $|\mathcal{C}|$ is odd, the unpaired candidate advances to the next iteration with its weight unchanged.
   ii. For each pair (P_i, P_j), all n parties execute $\Pi_{\mathsf{2LE}}(i, j, w_i, w_j)$, which elects P_i with probability $\frac{w_i}{w_i+w_j}$ and P_j with probability $\frac{w_j}{w_i+w_j}$.
   iii. Let P be the winner of the pair and assign its new weight the sum $w_i + w_j$.
   iv. Replace $\mathcal{C}$ with the set consisting of all pair winners and the unpaired candidate, if any.

Output Computation: Every party outputs the unique candidate remaining in $\mathcal{C}$.

Theorem 3. *Let $t_S \leqslant t_{GT}$ and $t_S + t_{GT} < n$, and let every $\Pi_{\mathsf{2LE}}(i, j, k_i, k_j)$ be a (t_S, t_{GT}, i, j)-BoBW two-party leader election with weights k_i, k_j. Then Protocol 4 is a (t_S, t_{GT})-BoBW fair leader election protocol (Definition 5) for the uniform distribution over $[n]$. In particular, with Protocol 3 as Π_{2LE}, it is perfectly (t_S, t_{GT})-BoBW fair in the $\mathcal{F}_{\mathsf{com}}$-hybrid model.*

Proof (of Theorem 3). In every execution of $\Pi_{\mathsf{2LE}}(i, j, w_i, w_j)$ all honest parties output a common winner in $\{P_i, P_j\}$, and for every coalition J and strategy π_J:

(2LE-S) if $|J| \leqslant t_S$, then P_i wins with probability $\frac{w_i}{w_i+w_j} \pm \mathsf{negl}(\lambda)$ (strong fairness for D_{w_i,w_j});
(2LE-MM) if $|J| \leqslant t_{GT}$ and P_i is honest, then P_i wins with probability at least $\frac{w_i}{w_i+w_j} - \mathsf{negl}(\lambda)$ (MM fairness, together with (2LE-S) for $J = \varnothing$, which gives $\mathsf{U}_i(\mathring{\pi}_{[n]}) = \frac{w_i}{w_i+w_j} \pm \mathsf{negl}(\lambda)$).

Here $\mathsf{negl}(\lambda)$ is 0 in the perfect case. Note that all n parties, including those already eliminated, participate in every match, so the bound on $|J|$ applies to every match.

We show t_S-strong fairness for the uniform distribution over $[n]$ and t_{GT}-GT fair for $\mathsf{Pref}_{\mathsf{LE}}$. Fix a coalition J with $|J| \leqslant t_{GT}$, a strategy π_J, and let $\mathcal{H} := [n] \backslash J$. Since all honest parties agree on the winner of every match, they agree on the candidate set, the weights and the pairing throughout, and they output the same $\mathsf{leader} \in [n]$.

Weights. The protocol runs for exactly $R := \lceil \log_2 n \rceil$ iterations. Let $W_i(r)$ denote the weight of P_i after iteration r, i.e., its current weight w_i if it is still a candidate and 0 if it has been eliminated, and let H_r denote the history of the execution through iteration r (all messages and all parties' internal states). Initially $W_i(0) = 1$. The total weight of the candidates is always n, since a match transfers the combined weight of the pair to the winner, and after iteration R a single candidate remains, so

$$W_i(R) = n \cdot \mathbb{1}[\mathsf{leader} = P_i]. \tag{10}$$

One iteration. Fix $r \in [1, R]$ and a history h through iteration $r-1$. Conditioned on $H_{r-1} = h$, the behaviour of the coalition in the matches of iteration r is a strategy for Π_{2LE} (namely π_J with h hard-wired), so (2LE-S) and (2LE-MM) apply to every match of iteration r conditioned on $H_{r-1} = h$. Now consider P_i in iteration r. If P_i was eliminated earlier or is unpaired, then $W_i(r) = W_i(r-1)$. If P_i is paired with P_j, then $W_i(r) = w_i + w_j$ if P_i wins and 0 otherwise, so

$\mathsf{E}[W_i(r) \mid H_{r-1}] = (w_i + w_j) \cdot \Pr[P_i \text{ wins} \mid H_{r-1}]$. Since $w_i + w_j \leqslant n$, (2LE-MM) and (2LE-S) give

$$\mathsf{E}[W_i(r) \mid H_{r-1}] \geqslant W_i(r-1) - n \cdot \mathsf{negl}(\lambda) \quad \text{if } i \in \mathcal{H}, \text{ and} \tag{11}$$
$$\mathsf{E}[W_i(r) \mid H_{r-1}] = W_i(r-1) \pm n \cdot \mathsf{negl}(\lambda) \quad \text{if } |J| \leqslant t_S, \text{ for every } i \in [n]. \tag{12}$$

Election probabilities. Taking expectations in Equations (11) and (12) and iterating over the R iterations gives $\mathsf{E}[W_i(R)] \geqslant 1 - Rn \cdot \mathsf{negl}(\lambda)$ for honest i, and $\mathsf{E}[W_i(R)] = 1 \pm Rn \cdot \mathsf{negl}(\lambda)$ for every i if $|J| \leqslant t_S$. By Equation (10), $\mathsf{E}[W_i(R)] = n \cdot \Pr[\mathsf{leader} = P_i]$, and $Rn \cdot \mathsf{negl}(\lambda)$ is negligible, hence

$$\Pr[\mathsf{leader} = P_i] \geqslant \tfrac{1}{n} - \mathsf{negl}(\lambda) \quad \text{for every } i \in \mathcal{H}, \text{ and} \tag{13}$$
$$\Pr[\mathsf{leader} = P_i] = \tfrac{1}{n} \pm \mathsf{negl}(\lambda) \quad \text{for every } i \in [n] \text{ if } |J| \leqslant t_S. \tag{14}$$

Conclusion. Equation (14) with $J = \varnothing$ shows that an honest execution elects every party with probability $\frac{1}{n}$ (up to a negligible error, and exactly if Π_{2LE} is perfectly strongly fair), so $\mathsf{U}_i(\mathring{\pi}_{[n]}) = \frac{1}{n} \pm \mathsf{negl}(\lambda)$. For $|J| \leqslant t_S$, Equation (14) is exactly t_S-strong fairness (Definition 1) for the uniform distribution over $[n]$. For $|J| \leqslant t_{GT}$ and honest i, Equation (13) gives $\mathsf{U}_i(\mathring{\pi}_{-J}, \pi_J) = \Pr[\mathsf{leader} = P_i] \geqslant \mathsf{U}_i(\mathring{\pi}_{[n]}) - \mathsf{negl}(\lambda)$, i.e., t_{GT}-MM fairness (Definition 3); since the protocol never aborts, Remark 1 yields t_{GT}-CSP fairness and hence t_{GT}-GT fairness.

5 BoBW Leader Election Impossible for $t_S + t_{GT} \geqslant n$

In this section, we prove that BoBW fairness is unachievable if $t_S + t_{GT} \geqslant n$. We show that this holds even if corrupt participants follow the protocol honestly apart from being (a) *semi-malicious*, meaning that the adversary chooses their randomness, and (b) *fail-stop*, meaning that the adversary may instruct them to abort at any time.

Theorem 4 (Lower Bound for $t_S + t_{GT} \geqslant n$). *Let $t_S + t_{GT} \geqslant n$ for $1 \leqslant t_S, t_{GT} < n$. Then no PPT polynomial-round n-party leader election protocol can be S/CSP (t_S, t_{GT})-BoBW fair (Definition 5), even if corrupt participants are fail-stop and semi-malicious but otherwise follow protocol instructions.*

Roadmap. In order to prove Theorem 4, we proceed as follows:

1. In Section 5.1, we state and prove Theorem 5, a generalized version of Cleve's lower bound for two parties that sample correlated bits rather than a single shared bit.
2. In Section 5.2, we state and prove Lemma 1, showing that a polynomial-round three-party leader election protocol cannot be $(1, 2)$-BoBW fair, even when fairness is only required against four of the six possible coalitions.
3. In Section 5.3, we prove Theorem 4, extending Lemma 1 to arbitrary n and $t_S + t_{GT} \geqslant n$.

Finally, in Section 5.4, we observe that the impossibility of S/CSP BoBW fairness (Theorem 4) implies the impossibility of both S/MM BoBW fairness and BoBW fairness.

Remark 2 (Protocol Efficiency). Throughout this section, the attacks we construct only require the attacker to run the honest code of some parties locally, so they are efficient whenever the protocol is. Hence, if the protocol is efficient, the proofs rule out computational fairness guarantees; without any efficiency assumption on the protocol or the adversary, they rule out statistical and perfect fairness guarantees. In both cases we require the protocol to have polynomially many rounds, since this is what makes the bias of Theorem 5 non-negligible.

5.1 Generalized Cleve

Theorem 5 (Generalized Cleve). *Assume that we have an R-round two party protocol between Alice and Bob, where Alice outputs bit a and Bob outputs bit b. Alice will try to bias b, while Bob will try to bias a. Let $u = \mathsf{cov}(a, b) = \mathsf{E}[a \cdot b] - \mathsf{E}[a] \cdot \mathsf{E}[b]$ be the covariance of a and b in the case where Alice and Bob both behave honestly. Then there is an attack allowing either Alice to bias b or Bob to bias a, with bias of magnitude at least $\frac{|u|}{4R+2}$.*

Proof (of Theorem 5). Assume, without loss of generality, that Alice speaks first in each round[8].

Let

- a_r be Alice's output if Bob aborts in round r (computable by Alice at the start of round r), and
- b_r be Bob's output if Alice aborts in round r (so Bob knows b_{r+1} once he has Alice's round-r message).

Note that a_1 and b_1 are independent (since no messages have been exchanged yet), and a_{R+1} and b_{R+1} are distributed exactly as Alice and Bob's output in an honest execution of the protocol (since parties corrupted by fail-stop adversaries either aborted or behaved honestly).

We define the following attacks and their associated biases.

Alice's naïve attack NA_r: Alice always aborts after round r. We have

$$\Pr[b = 1 \mid NA_r] = \mathsf{E}[b_{r+1}],$$

so the bias of this attack has magnitude

$$\mathsf{Bias}_{NA_r} = |\mathsf{E}[b_{r+1}] - \mathsf{E}[b]|.$$

We use NA_0 to denote the attack of aborting at the start, without sending any messages.

Bob's naïve attack NB_r: Bob always aborts after round r.

$$\Pr[a = 1 \mid NB_r] = \mathsf{E}[a_{r+1}],$$

so the bias of this attack has magnitude

$$\mathsf{Bias}_{NB_r} = |\mathsf{E}[a_{r+1}] - \mathsf{E}[a]|.$$

We use NB_0 to denote the attack of aborting at the start, without sending any messages.

Alice's attack A_r: at the start of round r, Alice aborts now if $a_r = 1$, and otherwise aborts in round $r + 1$. Then $b = b_r$ when $a_r = 1$ and $b = b_{r+1}$ when $a_r = 0$, so

$$\begin{aligned}\Pr[b = 1 \mid A_r] &= \mathsf{E}[a_r b_r + (1 - a_r) b_{r+1}] \\ &= \mathsf{E}[b_{r+1} + a_r (b_r - b_{r+1})] \\ &= \mathsf{E}[b_{r+1}] + \underbrace{\mathsf{E}\big[a_r \, (b_r - b_{r+1})\big]}_{=:\, \beta_r}.\end{aligned}$$

That is, β_r is exactly how much A_r differs from NA_r. The bias of this attack has magnitude

$$\mathsf{Bias}_{A_r} = |\mathsf{E}[b_{r+1}] + \beta_r - \mathsf{E}[b]|.$$

[8] In the rest of this paper we assumed that parties may speak simultaneously in a given round, whereas here we assume that Alice speaks first in each round. However, any protocol between Alice and Bob with synchronous rounds can be compiled to one in which Alice speaks first, with the only effect being that Alice can no longer be "rushing" meaning that Alice no longer has the opportunity to choose her message in any given round r based on Bob's message in the same round. If the original protocol was fair, the compiled protocol is also fair; so, by proving that the compiled protocol cannot be fair, we prove that the original protocol was not fair.

Bob's attack B_r: after Alice's round-r message, Bob aborts if $b_{r+1} = 1$, and otherwise aborts in round $r + 1$. Symmetrically,

$$\begin{aligned}\Pr[a = 1 \mid B_r] &= \mathsf{E}[b_{r+1}a_r + (1 - b_{r+1})a_{r+1}] \\ &= \mathsf{E}[a_{r+1}] + \underbrace{E\big[b_{r+1}\,(a_r - a_{r+1})\big]}_{=:\,\alpha_r}.\end{aligned}$$

The bias of this attack has magnitude

$$\mathsf{Bias}_{B_r} = |\mathsf{E}[a_{r+1}] + \alpha_r - \mathsf{E}[a]|.$$

Since there are $2R$ attacks $A_1, \ldots, A_R$ and $B_1, \ldots, B_R$, and $2(R + 1)$ attacks $NA_0, \ldots, NA_R$ and $NB_0, \ldots, NB_R$, this makes for $4R + 2$ attacks total.

Let S denote the sum of the bias magnitudes of all of the attacks described above. Summing up the biases, we get

$$\begin{aligned}S =& \mathsf{Bias}_{NA_0} + \mathsf{Bias}_{NB_0} + \sum_{r=1}^{R} \big(\mathsf{Bias}_{NA_r} + \mathsf{Bias}_{A_r} + \mathsf{Bias}_{NB_r} + \mathsf{Bias}_{B_r}\big) \\ =& \mathsf{Bias}_{NA_0} + \mathsf{Bias}_{NB_0} \\ &+ \sum_{r=1}^{R} |\mathsf{E}[b_{r+1}] - \mathsf{E}[b]| + \sum_{r=1}^{R} |\mathsf{E}[b_{r+1}] + \beta_r - \mathsf{E}[b]| \\ &+ \sum_{r=1}^{R} |\mathsf{E}[a_{r+1}] - \mathsf{E}[a]| + \sum_{r=1}^{R} |\mathsf{E}[a_{r+1}] + \alpha_r - \mathsf{E}[a]| \\ \geqslant& \mathsf{Bias}_{NA_0} + \mathsf{Bias}_{NB_0} + \Big|\sum_{r=1}^{R} \big(\beta_r + \alpha_r\big)\Big|,\end{aligned}$$

using the triangle inequality in the last step. Expanding the remaining sum, we get

$$\begin{aligned}\sum_{r=1}^{R}(\beta_r + \alpha_r) &= \sum_{r=1}^{R} \big(\mathsf{E}[a_r b_r - a_r b_{r+1}] + \mathsf{E}[b_{r+1}a_r - b_{r+1}a_{r+1}]\big) \\ &= \sum_{r=1}^{R} \big(\mathsf{E}[a_r b_r] - \mathsf{E}[a_{r+1}b_{r+1}]\big) \\ &= \mathsf{E}[a_1 b_1] - \mathsf{E}[a_{R+1}b_{R+1}] \\ &= \mathsf{E}[a_1]\mathsf{E}[b_1] - \mathsf{E}[ab]\end{aligned}$$

using independence of a_1, b_1 in the last step. Now, observe that

$$\begin{aligned}\Big|\mathsf{E}[a_1]\mathsf{E}[b_1] - \mathsf{E}[a]\mathsf{E}[b]\Big| &= \Big|(\mathsf{E}[a_1] - \mathsf{E}[a])\mathsf{E}[b_1] + \mathsf{E}[a](\mathsf{E}[b_1] - \mathsf{E}[b])\Big| \\ &\leqslant |\mathsf{E}[a_1] - \mathsf{E}[a]|\mathsf{E}[b_1] + \mathsf{E}[a]|\mathsf{E}[b_1] - \mathsf{E}[b]| \\ &\leqslant |\mathsf{E}[a_1] - \mathsf{E}[a]| + |\mathsf{E}[b_1] - \mathsf{E}[b]| \\ &= \mathsf{Bias}_{NB_0} + \mathsf{Bias}_{NA_0},\end{aligned}$$

using the fact that all expectations lie between 0 and 1 in the second-to-last step. We can conclude that

$$\begin{aligned}|\mathsf{E}[a]\mathsf{E}[b] - \mathsf{E}[ab]| &\leqslant |\mathsf{E}[a_1]E[b_1] - \mathsf{E}[ab]| + \mathsf{Bias}_{NB_0} + \mathsf{Bias}_{NA_0} \\ &= |\sum_{r=1}^{R}(\beta_r + \alpha_r)| + \mathsf{Bias}_{NB_0} + \mathsf{Bias}_{NA_0} \\ &\leqslant S.\end{aligned}$$

Since there are $4R + 2$ attacks, it follows that $S \leqslant (4R + 2)s$, where s is the largest bias. So,

$$\begin{aligned}(4R + 2)s &\geqslant \left|\mathsf{E}[a]\mathsf{E}[b] - \mathsf{E}[ab]\right| \\ \Rightarrow s &\geqslant \frac{\left|\mathsf{E}[a]\mathsf{E}[b] - \mathsf{E}[ab]\right|}{4R + 2} = \frac{|\mathsf{cov}(a, b)|}{4R + 2};\end{aligned}$$

in other words, the bias of at least one attack must have magnitude at least $\frac{|\mathsf{cov}(a,b)|}{4R+2}$, proving Theorem 5.

5.2 Lower Bound for $n = 3$

Continuing towards proving Theorem 4, we first prove the statement for three parties.

Lemma 1 (Lower Bound for $n = 3$, $t_{GT} = 2$, $t_S = 1$). *No polynomial-round three-party leader election protocol carried out by Alice, Bob and Carol where Alice and Bob each win with prescribed (non-negligible) probability p can be strongly fair against a corrupt Alice or a corrupt Bob and simultaneously CSP fair against a corrupt Alice-Carol coalition or a corrupt Bob-Carol coalition, even if corrupt participants are fail-stop and semi-malicious (i.e. may choose their randomness) but otherwise follow protocol instructions.*

Proof (of Lemma 1). Towards contradiction, let Π be an R-round three-party leader-election protocol run between Alice, Bob and Carol with strong fairness against a corrupt Alice or a corrupt Bob, and CSP fairness against a corrupt Alice-Carol coalition or a corrupt Bob-Carol coalition. Say Alice gets a bit a (where $a = 1$ iff Alice wins), Bob gets a bit b (where $b = 1$ iff Bob wins), and Carol gets a bit c (where $c = 1$ iff Carol wins).

Now, compile this protocol into a two-party protocol which excludes Carol, by having Alice and Bob each run Carol's code locally with public randomness. Alice provides this randomness alongside her first message; if Alice aborts before sending her first message, Bob chooses Carol's randomness instead[9]. Then $\mathsf{E}[a] = \mathsf{E}[b] = p$ and $\mathsf{E}[ab] = 0$, so $u = \mathrm{cov}(a, b) = -p^2$.

By Theorem 5, some attack NA_r, A_r, NB_r or B_r (as described in the proof of Theorem 5) must induce a bias of magnitude at least $s = \frac{p^2}{4R+2}$ in the other party's bit, which is non-negligible.

However, this alone does not get us all the way to a contradiction, since CSP-fairness is not violated if the bias is positive (i.e., it increases the other party's probability of winning). To complete the argument, we introduce Lemma 2 which states that the existence of an attack with bias s implies the existence of an attack with bias $-s + \mathsf{negl}$.

[9] Bob's resampling of Carol's randomness corresponds, in the three-party protocol, to Carol choosing a fresh random tape when Alice aborts in round 1. This is the only place where a corrupt party deviates by more than aborting, and it is why Lemma 1 is stated for semi-malicious rather than strictly fail-stop corruptions.

Lemma 2 (Push-Pull). *If for some A_r it holds that $\Pr[b = 1 \mid A_r] \geqslant \mathsf{E}[b] + s$ for $0 < s$, then if Π is 1-strong fair, there exists an attack $\overline{A}_r$ by Alice such that $\Pr[b = 1 \mid \overline{A}_r] \leqslant \mathsf{E}[b] - s + \epsilon$ for a negligible ϵ.*

(The statement for attacks by Bob is symmetric.)

Proof (of Lemma 2). Define $\overline{A}_r$ to abort at the start of round r if $a_r = 0$, and to abort in round $r + 1$ otherwise. (That is, $\overline{A}_r$ makes the opposite choice of A_r.)

Randomly choosing between running A_r and $\overline{A}_r$ (with probability $\frac{1}{2}$) is the same as fail-stopping Alice with probability $\frac{1}{2}$ in round r, and fail-stopping Alice with probability $\frac{1}{2}$ in round $r+1$. This is an attack that does not use Carol's view, so can be run by Alice alone; by the 1-strong fairness of Π, such an attack should have at most negligible influence δ over the outcome.

We can conclude that

$$\begin{aligned}\mathsf{E}[b] + \delta &\geqslant \frac{\Pr[b = 1|A_r]}{2} + \frac{\Pr[b = 1|\overline{A}_r]}{2}\\ \Rightarrow 2\mathsf{E}[b] + 2\delta &\geqslant \Pr[b = 1|A_r] + \Pr[b = 1|\overline{A}_r]\end{aligned}$$

Plugging in $\Pr[b = 1|A_r] = \mathsf{E}[b] + s$, we get

$$\begin{aligned}2\mathsf{E}[b] + 2\delta &\geqslant \mathsf{E}[b] + s + \Pr[b = 1|\overline{A}_r]\\ \Rightarrow \mathsf{E}[b] + 2\delta - s &\geqslant \Pr[b = 1|\overline{A}_r].\end{aligned}$$

Setting $\epsilon = 2\delta$, we get exactly the statement in Lemma 2.

The symmetric argument holds for an attack by Bob.

Now, coming back to Lemma 1, we can consider several cases.

1. If either NA_r or NB_r (for some r) has bias of magnitude at least s, this violates 1-strong fairness, since NA_r and NB_r can be carried out by either Alice or Bob alone.
2. If either A_r or B_r has bias of magnitude at least s, this leads to two possibilities:
   (a) The bias is in the attacker's favor, violating 2-CSP fairness.
   (b) The bias works against the attacker; however, by Lemma 2, there then exists a symmetric attack in the attacker's favor, violating 2-CSP fairness.

This concludes the proof of Lemma 1.

5.3 Proof of Theorem 4

To conclude the proof of Theorem 4, we observe that if there existed an n-party protocol Π_n with (t_S, t_{GT})-BoBW fairness for $t_S + t_{GT} \geqslant n$, we could compile it into a 3-party protocol Π_3 with the fairness guarantees ruled out in Lemma 1 (strong fairness against a corrupt Alice or a corrupt Bob and simultaneously CSP fair against a corrupt Alice-Carol coalition or a corrupt Bob-Carol coalition), where Alice and Bob each win with probability $p = \frac{t_S}{n}$.

Assume, without loss of generality, that $t_S < \frac{n}{2}$ (otherwise, the standard lower bound due to Cleve [Cle86] rules out t_S-fair protocols). Alice and Bob in Π_3 would each play the roles of t_S parties each in Π_n; Carol in Π_3 would play the roles of the remaining $n - 2t_S$ parties in Π_n. A party wins in Π_3 if any of the parties whose roles it plays wins in Π_n.

5.4 Corollaries

Corollary 1 (Lower Bound for S/MM BoBW-Fairness, $t_S + t_{GT} \geqslant n$). *S/MM BoBW-fair (Definition 7) leader election in polynomial many rounds is impossible for $t_S + t_{GT} \geqslant n$.*

Proof. Any leader election protocol that satisfies S/MM (t_S, t_{GT})-BoBW fairness (Definition 7) also satisfies S/CSP (t_S, t_{GT})-BoBW fairness (Definition 6) as per Remark 1. So, if a S/MM (t_S, t_{GT})-BoBW-fair protocol existed for $t_S + t_{GT} \geqslant n$, it would also be S/CSP (t_S, t_{GT})-BoBW fair, violating Theorem 4.

Corollary 2 (Lower Bound for BoBW-Fairness, $t_S + t_{GT} \geqslant n$). *BoBW-fair (Definition 5) leader election in polynomial many rounds is impossible for $t_S + t_{GT} \geqslant n$.*

Proof. As with S/MM BoBW-fairness, any protocol that satisfies BoBW-fairness also satisfies S/CSP BoBW-fairness, which would violate Theorem 4.

6 Statistical Strongly Fair Leader Election for $t < \frac{n}{2}$

In this section, we consider leader election in the statistical setting, i.e., against computationally unbounded coalitions. For an honest majority, statistically strongly fair leader election follows directly from general-purpose statistically secure multi-party computation with guaranteed output delivery.[10]

Theorem 6 (Statistical Strong Fairness for $t < \frac{n}{2}$). *Let $t < \frac{n}{2}$. There exists a bounded-round n-party leader election protocol that is statistically t-strongly fair (Definition 1) for the uniform distribution over $[n]$. In particular, the protocol is statistically t-MM fair, t-CSP fair and t-GT fair (Definitions 2 to 4) for the leader election preference profile $\mathsf{Pref}_{\mathsf{LE}}$ of Equation (4).*

Proof (of Theorem 6). Let Π_{mpc} be a statistically secure n-party MPC protocol with guaranteed output delivery tolerating $t < \frac{n}{2}$ malicious corruptions, assuming secure point-to-point channels and broadcast, e.g., the protocol of Rabin and Ben-Or [RB89]. The parties run Π_{mpc} for the (input-less, randomized) functionality $\mathcal{F}_{\mathsf{LE}}$ that samples $x \leftarrow [n]$ uniformly at random and delivers x to all parties. Since $\mathcal{F}_{\mathsf{LE}}$ takes no inputs, Π_{mpc} evaluates a fixed circuit of size and depth independent of λ, so the protocol is bounded-round.

Fix a coalition J of size at most t and an arbitrary (unbounded) strategy π_J. By statistical security with guaranteed output delivery, the joint distribution of the honest parties' outputs in the real execution $\Pi_{\mathsf{mpc}}\langle \pi_J, \mathring{\pi}_{-J} \rangle$ is statistically close to the joint distribution of their outputs in the ideal execution with $\mathcal{F}_{\mathsf{LE}}$ and some ideal-world adversary. In the ideal execution, the adversary has no influence on the output: all honest parties receive the same x, and x is uniform over $[n]$. Hence, in the real execution, all honest parties agree on their output except with negligible probability, and the output x satisfies $\left| \Pr[x = X \mid X \leftarrow \Pi_{\mathsf{mpc}}\langle \pi_J, \mathring{\pi}_{-J} \rangle] - \frac{1}{n} \right| \leqslant \mathsf{negl}(\lambda)$ for every $x \in [n]$, which is precisely statistical t-strong fairness.

For the "in particular" part, statistical strong fairness implies statistical MM fairness: for every honest party $i \notin J$, $\mathsf{U}_i(\mathring{\pi}_{-J}, \pi_J) = \Pr[x = i] \geqslant \frac{1}{n} - \mathsf{negl}(\lambda) = \mathsf{U}_i(\mathring{\pi}_{[n]}) - \mathsf{negl}(\lambda)$. Statistical CSP fairness and GT fairness then follow from Remark 1.

[10] Recall that *guaranteed output delivery* is the property that all honest parties compute the output in every execution of a protocol.

7 Statistical Leader Election Impossible for $t \geqslant \frac{n}{2}$

We now show that Theorem 6 is tight, even with respect to the weakest notion of fairness considered in this paper.

Theorem 7 (Lower Bound for Statistical CSP-Fairness, $t \geqslant \frac{n}{2}$). *Let $n \geqslant 2$ and $t \geqslant \frac{n}{2}$. No bounded-round n-party leader election protocol is statistically t-CSP fair (Definition 2) for the leader election preference profile $\mathsf{Pref}_{\mathsf{LE}}$ of Equation (4).*

Roadmap. We prove Theorem 7 in two steps.

1. In Section 7.1, we observe that two-party leader election and two-party coin tossing are the same problem, since the two tasks induce identical utility functions (Lemma 3). The impossibility of statistically CSP-fair two-party coin tossing due to Liu-Zhang, Masserova, Ribeiro and Thyagarajan [LZMRT26] then rules out statistically 1-CSP fair two-party leader election as an immediate corollary (Corollary 3).
2. In Section 7.2, we show by a party-emulation argument that any statistically t-CSP fair n-party leader election protocol with $t \geqslant \frac{n}{2}$ yields a statistically 1-CSP fair two-party protocol of the kind ruled out by Corollary 3.

Finally, in Section 7.3, we derive the corresponding impossibilities for MM-, GT- and BoBW-fairness.

7.1 Two Parties

A two-party *coin tossing* protocol is a protocol in which both parties output a bit on which they agree (except with negligible probability) and which, in an all-honest execution, is uniform over $\{0,1\}$. We consider the preference profile $(0,1)$ of [LZMRT26], in which the parties prefer different outcomes:

$$\mathsf{Pref}_{\mathsf{CT}}[1] = \{0\} \quad \text{and} \quad \mathsf{Pref}_{\mathsf{CT}}[2] = \{1\}. \tag{15}$$

Two-party leader election with preference profile $\mathsf{Pref}_{\mathsf{LE}}$ and two-party coin tossing with preference profile $\mathsf{Pref}_{\mathsf{CT}}$ are the same task up to renaming of the output.

Lemma 3. *Let $\phi : [2] \to \{0,1\}$ be the bijection $\phi(x) = x - 1$. For a two-party protocol Π with outputs in $[2]$, let Π^{ϕ} denote the protocol in which the parties run Π and output $\phi(x)$, where x is their output in Π. Then for every coalition $J \subseteq [2]$ and every strategy π_J,*

$$\mathsf{U}_J^{\mathsf{Pref}_{\mathsf{CT}},\Pi^{\phi}}\big(\pi_J, \mathring{\pi}_{-J}\big) = \mathsf{U}_J^{\mathsf{Pref}_{\mathsf{LE}},\Pi}\big(\pi_J, \mathring{\pi}_{-J}\big). \tag{16}$$

Consequently, Π is a (perfectly, statistically, computationally) t-CSP fair, t-MM fair, or t-strongly fair leader election protocol if and only if Π^{ϕ} is a (perfectly, statistically, computationally) t-CSP fair, t-MM fair, or t-strongly fair coin tossing protocol, respectively; and symmetrically for ϕ^{-1} applied to a coin tossing protocol.

Proof (of Lemma 3). By Equations (4) and (15), $x \in \mathsf{Pref}_{\mathsf{LE}}[i]$ if and only if $\phi(x) \in \mathsf{Pref}_{\mathsf{CT}}[i]$, i.e., $\mathsf{util}_i^{\mathsf{Pref}_{\mathsf{CT}}}(\phi(x)) = \mathsf{util}_i^{\mathsf{Pref}_{\mathsf{LE}}}(x)$ for all $i, x \in [2]$, and hence the same holds for the joint utility of any coalition. Since Π^{ϕ} differs from Π only by a deterministic local post-processing of the output, every strategy for Π is a strategy for Π^{ϕ} and vice versa, and the output of $\Pi^{\phi}\langle \pi_J, \mathring{\pi}_{-J}\rangle$ is $\phi(x)$ for x the output of $\Pi\langle \pi_J, \mathring{\pi}_{-J}\rangle$. Equation (16) follows by taking expectations, and the CSP and MM conditions (Definitions 2 and 3) are inequalities between such expected utilities. For strong fairness, ϕ maps the uniform distribution over $[2]$ to the uniform distribution over $\{0,1\}$ and preserves point probabilities.

The two-party model of [LZMRT26] is the statistical variant of ours: a computationally unbounded, rushing adversary, each party's output a function of its own view, and CSP-fairness as in Definition 2. Theorem 1 of [LZMRT26] states that no two-party coin tossing protocol with a fixed maximum number of rounds is statistically 1-CSP fair for the profile $(0, 1)$. Its proof ([LZMRT26, Lemma 2]) in fact shows more: for *any* bounded-round two-party protocol in which the honest outputs agree except with negligible probability, whatever their distribution, either P_1 can force an honest P_2 to output 0, or P_2 can force an honest P_1 to output 1, with probability $1 - \mathsf{negl}(\lambda)$. We use this stronger form since, for odd n, the two-party protocol emulated in Section 7.2 does not elect each party with probability exactly $\frac{1}{2}$ (cf. Lemma 1).

Corollary 3 (Lower Bound for Statistical CSP-Fairness, $n = 2$, $t = 1$). *Let Π be a bounded-round two-party protocol in which both parties output a common element of $[2]$ (except with negligible probability), and let p_i be the probability that party i is elected in an all-honest execution. If $1 - p_1$ and $1 - p_2$ are both non-negligible, then Π is not statistically 1-CSP fair (Definition 2). In particular, no bounded-round two-party leader election protocol is statistically 1-CSP fair.*

Proof (of Corollary 3). Π^{ϕ} is a bounded-round two-party protocol whose honest outputs agree except with negligible probability, so by [LZMRT26, Lemma 2] some $i \in [2]$ has a strategy π_i under which an honest P_{3-i} outputs $\phi(i)$, i.e., elects i in Π, with probability at least $1 - \mathsf{negl}(\lambda)$. Thus $\mathsf{U}^{\mathsf{Pref}_{\mathsf{LE}},\Pi}_{\{i\}}(\pi_i, \mathring{\pi}_{3-i}) \geqslant 1 - \mathsf{negl}(\lambda)$, whereas $\mathsf{U}^{\mathsf{Pref}_{\mathsf{LE}},\Pi}_{\{i\}}(\mathring{\pi}_{[2]}) = p_i$; since $1 - p_i$ is non-negligible, this violates Definition 2. For the "in particular" part ($p_1 = p_2 = \frac{1}{2}$), it also suffices to combine Lemma 3 with [LZMRT26, Theorem 1] directly.

7.2 n Parties

We now prove Theorem 7 by a standard party-emulation argument: two parties, each emulating at most t of the parties of a hypothetical n-party protocol, would obtain a two-party protocol contradicting Corollary 3.

Proof (of Theorem 7). Towards contradiction, let Π_n be a bounded-round n-party leader election protocol that is statistically t-CSP fair for some $t \geqslant \frac{n}{2}$. Partition $[n]$ into two sets A and B with $|A| = \lceil \frac{n}{2} \rceil$ and $|B| = \lfloor \frac{n}{2} \rfloor$. Since $t \geqslant \frac{n}{2}$ and t is an integer, $|B| \leqslant |A| = \lceil \frac{n}{2} \rceil \leqslant t$; hence both A and B are coalitions to which the t-CSP fairness of Π_n applies.

The two-party protocol Π_2. We define a two-party protocol Π_2 between parties P_1 and P_2, in which P_1 emulates the parties in A and P_2 emulates the parties in B:

- *Emulation.* P_1 runs the honest code $\mathring{\pi}_i$ of Π_n for every $i \in A$, using independent fresh randomness for each; P_2 does the same for every $i \in B$. Round r of Π_2 corresponds to round r of Π_n.
- *Message delivery.* A point-to-point message of Π_n from i to j with $i, j \in A$ is delivered internally by P_1; likewise for B and P_2. A point-to-point message from $i \in A$ to $j \in B$ is sent by P_1 to P_2 over the point-to-point channel of Π_2, tagged with (i, j); P_2 delivers it to its emulated party j as a message from i. A broadcast by $i \in A$ is delivered internally to all emulated parties in A and sent to P_2 tagged with (i, bc), who delivers it to all emulated parties in B as a broadcast from i. Messages from B to A are handled symmetrically. In each round, each of P_1, P_2 sends the (ordered) list of all tagged messages of that round in a single message; a party that receives a malformed list, or a list containing a tag with an incorrect sender group, ignores the malformed entries (i.e., treats the corresponding emulated messages as not sent).
- *Output.* Once the emulated parties in A have produced their outputs, P_1 lets $x \in [n]$ be the output of the emulated party with the smallest index in A, and outputs 1 if $x \in A$ and 2 if $x \in B$. P_2 proceeds symmetrically with its emulated parties in B.

Honest executions. An all-honest execution of Π_2 is, by construction, an all-honest execution of Π_n in which every message is delivered faithfully and all parties use independent uniform randomness. Hence both parties of Π_2 output a common value except with negligible probability, and the leader x of the emulated execution is uniform over $[n]$, so that

$$p_1 = \Pr[x \in A] = \frac{\lceil n/2 \rceil}{n} \quad \text{and} \quad p_2 = \Pr[x \in B] = \frac{\lfloor n/2 \rfloor}{n}.$$

Since $n \geqslant 2$, $1 - p_1 = \frac{\lfloor n/2 \rfloor}{n} \geqslant \frac{1}{n}$ and $1 - p_2 = \frac{\lceil n/2 \rceil}{n} \geqslant \frac{1}{n}$ are both non-negligible (recall that n is polynomial in λ). Thus Π_2 satisfies the hypotheses of Corollary 3, and it remains to show that Π_2 is statistically 1-CSP fair, which yields the desired contradiction.

Translating strategies. Let π_1 be an arbitrary (unbounded) strategy for P_1 in Π_2; we show that P_1 cannot increase its expected utility by more than a negligible amount. Define a strategy π_A for the coalition A in Π_n as follows. The coalition runs π_1 internally. Whenever an honest party $j \in B$ sends a point-to-point message to some $i \in A$, or broadcasts a message, the coalition feeds the message to π_1 tagged with (j, i) or (j, bc), respectively, exactly as the honest P_2 would have sent it in Π_2. Whenever π_1 outputs a message to P_2 for a given round, the coalition parses it as P_2 would: for each well-formed entry tagged (i, j) with $i \in A$ and $j \in B$, party i sends the entry's payload to j over the point-to-point channel; for each well-formed entry tagged (i, bc) with $i \in A$, party i broadcasts the payload; malformed entries are dropped, meaning that the corresponding message is not sent. Messages between coalition members are irrelevant to the honest parties and are not sent.

By construction, the joint view of the honest parties in B in $\Pi_n \langle \pi_A, \mathring{\pi}_B \rangle$ is identically distributed to the view of the honest party P_2 in $\Pi_2 \langle \pi_1, \mathring{\pi}_2 \rangle$: in both executions the honest parties in B are run with independent uniform randomness, receive exactly the same messages from A in each round (a broadcast in Π_n delivers the same message to all parties in B, exactly as the single tagged message does in Π_2), and the messages they send are passed to π_1 in the same way. In particular, the output $x \in [n]$ of the honest parties in B (which is common to all of them except with negligible probability) has the same distribution in both executions, and P_2 outputs 1 in Π_2 if and only if $x \in A$ (up to the negligible disagreement event). Therefore,

$$\begin{aligned}
\mathsf{U}_{\{1\}}^{\mathsf{Pref}_{\mathsf{LE}}, \Pi_2}(\pi_1, \mathring{\pi}_2) &= \Pr\left[x \in A \;\middle|\; x \leftarrow \Pi_n \langle \pi_A, \mathring{\pi}_B \rangle \right] \pm \mathsf{negl}(\lambda) \\
&= \mathsf{U}_{A}^{\mathsf{Pref}_{\mathsf{LE}}, \Pi_n}(\pi_A, \mathring{\pi}_B) \pm \mathsf{negl}(\lambda) && (17) \\
&\leqslant \mathsf{U}_{A}^{\mathsf{Pref}_{\mathsf{LE}}, \Pi_n}(\mathring{\pi}_{[n]}) + \mathsf{negl}(\lambda) && (18) \\
&= \frac{|A|}{n} + \mathsf{negl}(\lambda) = \mathsf{U}_{\{1\}}^{\mathsf{Pref}_{\mathsf{LE}}, \Pi_2}(\mathring{\pi}_{[2]}) + \mathsf{negl}(\lambda), && (19)
\end{aligned}$$

where Equation (17) uses that $\mathsf{util}_A^{\mathsf{Pref}_{\mathsf{LE}}}(x) = 1$ if and only if $x \in A$, Equation (18) is the statistical t-CSP fairness of Π_n applied to the coalition A of size $|A| \leqslant t$ and the (unbounded) strategy π_A, and Equation (19) uses that the leader is uniform in an all-honest execution. The symmetric argument, with the coalition B of size $|B| \leqslant t$, shows that $\mathsf{U}_{\{2\}}^{\mathsf{Pref}_{\mathsf{LE}}, \Pi_2}(\mathring{\pi}_1, \pi_2) \leqslant \mathsf{U}_{\{2\}}^{\mathsf{Pref}_{\mathsf{LE}}, \Pi_2}(\mathring{\pi}_{[2]}) + \mathsf{negl}(\lambda)$ for every strategy π_2 of P_2. (For the empty coalition, Definition 2 holds trivially.) Hence Π_2 is statistically 1-CSP fair, contradicting Corollary 3. This concludes the proof of Theorem 7.

7.3 Corollaries

Corollary 4 (Lower Bound for Statistical MM- and GT-Fairness, $t \geqslant \frac{n}{2}$). *Let $n \geqslant 2$ and $t \geqslant \frac{n}{2}$. No bounded-round n-party leader election protocol is statistically t-MM fair (Definition 3) or statistically t-GT fair (Definition 4).*

Proof. By Remark 1, for leader election statistical t-MM fairness implies statistical t-CSP fairness, and t-GT fairness requires t-CSP fairness by definition. Both are therefore ruled out by Theorem 7.

8 Perfect CSP-fair Leader Election for $n = 3t$

For $n > 3t$, perfect *strongly* fair leader election is possible, since perfect general-purpose secure computation is possible [Mas26]. We now characterize perfect CSP-fair leader election for $n \leqslant 3t$ parties and up to t corruptions. First, we obtain a perfect CSP-fair protocol for the special case of $t = 1$, $n = 3$. Then, we show how to generalize this result to $n = 3t$ for any $t \geqslant 1$.

8.1 CSP-Fair Leader Election for $n = 3$, $t = 1$

We start by describing a perfect CSP-fair leader election protocol for three parties that is secure against one corruption.

Protocol 5: Perfect CSP-Fair Leader Election for $n = 3$ and $t = 1$

Protocol Execution:

1. Party P_1 samples $r \leftarrow \{0, 1, 2\}$ uniformly at random and sends it to P_2.
2. P_2 checks whether it received a valid message from P_1. If so, denote it r'. If not, set $r' = \perp$.
3. Party P_3 samples $s \leftarrow \{0, 1, 2\}$ uniformly at random and broadcasts s.
4. Party P_1 broadcasts r. Party P_2 broadcasts r'.

Output Computation: Each party does the following:

- If P_3 failed to broadcast a valid value in the third round, output 1.
- Otherwise, if either P_1 or P_2 failed to broadcast a non-$\perp$ value in the fourth round, or if the values r and r' they broadcast differ, output 3.
- Otherwise, output $((r + s) \mod 3) + 1$.

Theorem 8. *Protocol 5 is a perfect CSP-fair (Definition 2) protocol for leader election with three parties and up to one corruption.*

Proof (of Theorem 8). If everyone behaves honestly, since r and s are sampled from $\{1, 2, 3\}$ uniformly at random, $(r+s) \mod 3$ is uniformly random over $\{0, 1, 2\}$. Thus, every party has $\frac{1}{3}$ chance of being chosen as a leader.

A malicious P_1 cannot increase its expected utility by misbehaving. Say r' is the value recorded by the honest P_2. If $r' = \perp$ (or is of the wrong format), an honest P_2 will broadcast $\perp$ in the fourth round, and thus P_3 will be elected leader. Hence, not sending an $r \in \{0, 1, 2\}$ to P_2 can never increase P_1's utility. Otherwise, r' is fixed prior to an honest P_3 sampling its independent value s uniformly at random. Failing to broadcast (the correct) value r in the fourth round can never increase P_1's utility: in this case, as $r' \neq r$, P_3 will again be elected leader.

A malicious P_2 similarly cannot increase its expected utility by misbehaving: Failing to broadcast the correct value r' received from P_1 does not increase its utility, as in this case P_3 is elected leader.

Finally, a malicious P_3 has no knowledge of r. Hence, its choice of s is independent of r. Misbehaving by failing to broadcast a value $s \in \{0, 1, 2\}$ can never increase its utility, as in this case P_1 is elected leader.

8.2 CSP-Fair Leader Election for $n = 3t$

We now describe a perfect CSP-fair leader election protocol for $n = 3t$ parties that is secure against up to $t > 1$ corruptions.

At a high level, our construction follows the tournament-tree idea used in our computational construction. The main difference from the computational version is in how each match is resolved. To achieve perfect security, we delegate the weighted selection to all $3t-2$ parties *outside* the current pair, including parties eliminated in earlier matches. These parties execute a perfectly secure MPC protocol with guaranteed output delivery against up to $t-1$ corruptions.

The key idea is that whenever at least one competing candidate is corrupt, at most $t-1$ of the remaining $3t-2$ MPC participants are corrupt. Since

$$3t - 2 > 3(t-1),$$

perfectly secure MPC is available, and the match respects the prescribed winning probabilities. If both candidates are honest, the number of corrupt MPC participants may exceed this bound, but either possible winner is honest. We augment the MPC protocol by a public timeout and a voting rule to ensure that in this case the match nevertheless terminates with a common winner drawn from the two candidates. Thus, every match involving a corrupt candidate preserves its prescribed advancement probability, while deviations in other matches cannot increase the probability of a corrupt party becoming the elected leader.

Protocol 6: Perfect CSP-Fair Leader Election for $n = 3t$ Parties, where $t > 1$

Building Blocks: Let Π_{mpc} be a perfectly secure MPC protocol with guaranteed output delivery for $3t-2$ parties and at most $t-1$ corruptions. Let T be an upper bound on its running time under this corruption bound.

Protocol Execution:

1. Initialize the set of candidates as $\mathcal{C} \leftarrow \{P_1, \ldots, P_n\}$ and set $w_i \leftarrow 1$ for every $i \in [n]$.
2. While $|\mathcal{C}| > 1$, perform the following steps:
   i. Partition $\mathcal{C}$ into ordered pairs according to a fixed public rule. If $|\mathcal{C}|$ is odd, the unpaired candidate advances to the next iteration with its weight unchanged.
   ii. For each pair (P_i, P_j), all $3t-2$ parties other than P_i and P_j, including previously eliminated parties, execute Π_{mpc} for a functionality that selects P_i with probability
   $$\frac{w_i}{w_i + w_j}$$
   and P_j with probability
   $$\frac{w_j}{w_i + w_j}.$$
   iii. At time $T+1$ relative to the start of each MPC instance, every participant in that instance broadcasts its output. A party that has not obtained an output in $\{P_i, P_j\}$ broadcasts $\perp$.
   If at least
   $$(3t-2) - (t-1) = 2t-1$$
   participants broadcast the same candidate $P \in \{P_i, P_j\}$, declare P the winner of the pair. Otherwise, declare P_i the winner. Missing or malformed broadcasts are treated as $\perp$.

iv. Assign each pair's winner weight $w_i + w_j$, including when the winner is selected by the fallback rule.

v. Replace $\mathcal{C}$ with the set consisting of all pair winners and the unpaired candidate, if any.

Output Computation: Every party outputs the unique candidate remaining in $\mathcal{C}$.

Theorem 9. *Protocol 6 is a perfect CSP-fair (Definition 2) protocol for leader election with $n = 3t$ parties and up to t corruptions.*

Proof (of Theorem 9). Consider an adversary $\mathcal{A}$ corrupting up to t parties. For each match between some parties P_i and P_j, we consider the following two cases.

Case 1: both P_i and P_j are honest: In this case, the t adversarial parties participating in Π_{mpc} can disturb the execution of the protocol. However, at time $T + 1$ (starting from the start of the corresponding Π_{mpc} instance), all honest parties will broadcast their output (or $\perp$ if they have not received an output in $\{P_i, P_j\}$). Thus, everyone is in agreement on whether at least $2t - 1$ parties selected the same candidate. Hence, there is agreement on whether P_i or P_j won (if neither reached the $2t - 1$ threshold, the fallback rule declares P_i a winner). Either way, the winner of this match is an honest party.
Case 2: at most one party is honest among $\{P_i, P_j\}$: In this case, there are at most $t - 1$ corrupt parties among the $3t - 2$ participants of Π_{mpc}. Hence, the security guarantees of Π_{mpc} apply, and one of $\{P_i, P_j\}$ is selected according to their current weight.

Next, we argue that the expected weight of a malicious party does not change throughout the leader election protocol. Let $W_i(r)$ denote the weight of a malicious party P_i after protocol iteration r (i.e., it is 0 if P_i was eliminated in iteration r or before, and it is its weight w_i otherwise). Let H_r denote the execution history through iteration r.

If P_i has been eliminated before iteration r:

$$\mathsf{E}[W_i(r+1) \mid H_r] = 0 = W_i(r).$$

If P_i is unpaired in iteration r: P_i advances to the next iteration with its weight unchanged. Since $W_i(r)$ is determined by H_r, we obtain

$$\mathsf{E}[W_i(r+1) \mid H_r] = \mathsf{E}[W_i(r) \mid H_r] = W_i(r).$$

If P_i is paired with some P_j in iteration r: by Case 2 above the expected weight of P_i conditional on the execution history H_r after this iteration is

$$\mathsf{E}[W_i(r+1) \mid H_r] = (w_i + w_j) * \frac{w_i}{w_i + w_j} = w_i.$$

Hence, the expected weight of P_i remains the same throughout iterations.

Further, by the protocol description, the total weight of the candidates in $\mathcal{C}$ is always n, hence for the final round r^* it holds that $W_i(r^*) = n \cdot \mathbb{1}[P_i = \mathsf{leader}]$[11], where leader is the elected leader.

[11] $\mathbb{1}[P_i = \mathsf{leader}]$ is an indicator equal to 1 if $P_i = \mathsf{leader}$ and 0 otherwise.

Thus, $\mathsf{E}[n \cdot \mathbb{1}[P_i = \mathsf{leader}]] = \mathsf{E}[W_i(r^*)]$. Given that we only have a bounded number of iterations, it holds that:

$$n \cdot \Pr[P_i = \mathsf{leader}] = \mathsf{E}[W_i(r^*)] = \mathsf{E}[W_i(0)] = 1$$

Thus, $\Pr[P_i = \mathsf{leader}] = \frac{1}{n}$. Hence, the adversary can never increase the probability of corrupt parties being elected winner.

Further, if everyone behaves honestly, we get that the same argument above applies to every party P_i (not just corrupt ones), hence every party has probability $\frac{1}{n}$ of being elected winner.

Corollary 5. *Protocol 6 is a perfect S-CSP (t_S, t_{GT})-BoBW-fair (Definition 5) protocol for leader election with $n = 3t_{GT}$ parties and $t_S = t_{GT} - 1$.*

9 Lower Bounds for Perfect Leader Election

9.1 Impossibility of CSP Leader Election for $n < 3t$

We now show that perfect CSP fair leader election with a bounded number of rounds is impossible for $n < 3t$, $t \geqslant 1$. For this, we utilize the techniques from the proof of the impossibility of perfect coin toss for $n \leqslant 3t$ [Mas26].

Below, we split parties into up to three groups A, B, C. We use leader_X to denote the elected leader of the protocol Π during an honest execution X (an honest execution is determined solely by the randomness of the participating parties). We further use $\mathsf{view}^X_{P,r}$ to denote the joint view of the parties in the group $P \in \{A, B, C\}$ at the end of round r in the honest execution X. The group's view up to round r consists of the randomness of the parties in the group as well as all messages sent and received by the parties in P in rounds up to and including round r.

We now formally prove the following theorem:

Theorem 10 (Lower Bound for CSP-Fairness, $n < 3t$). *Perfect CSP-fair (Definition 2) leader election is impossible for $n < 3t$.*

Proof. Say that there exists a perfect CSP-fair leader election protocol Π for $n < 3t$. W.l.o.g., in the following we assume that in Π in each round at most one party sends messages.

Let r be the earliest round such that, in every honest execution of Π, every size-t coalition knows whether it wins, i.e., whether the elected leader is one of the coalition members, by the end of round r. Such a round exists because the protocol terminates within a fixed number of rounds, and all parties learn the leader.

We know that $r > 0$: Otherwise, every size-t coalition would know whether the elected winner is among the parties in the coalition from its initial randomness alone. Since its honest winning probability is $\frac{t}{n} > 0$, some choice of its randomness would guarantee that it wins. Fixing that randomness would increase its winning probability to 1, contradicting game-theoretic security.

Next, by the minimality of r, there exist a size-t coalition A and two honest executions X, Y such that

$$\mathsf{view}^X_{A,r-1} = \mathsf{view}^Y_{A,r-1},$$

but their respective leaders satisfy

$$\mathsf{leader}_X \in A, \qquad \mathsf{leader}_Y \notin A.$$

In the remainder of the proof, we write view^Z_D for the joint view of D immediately before round r in execution Z.

We first establish a helpful observation that will justify the hybrid arguments below.

Lemma 4 (Filtered state replacement). *Fix a t-size coalition D and an honest execution Z whose leader leader_Z is outside D. Replacing the pre-round-r states of any subset of D with fixed alternative states, and subsequently following the protocol, cannot cause a party in D to be elected with positive probability.*

Proof. For the sake of contradiction, suppose that there exists t-size coalition D, an honest execution Z with $\mathsf{leader}_Z \notin D$, and a state replacement whose continuation elects a party in D with positive probability. We now construct an adversary corrupting D that reproduces this replacement with positive probability, but deviates only when the joint view of the parties in D certifies that honest continuation would elect a leader outside D. Consequently, the adversary strictly increases the probability that the elected leader belongs to D, contradicting CSP security. Let P_s be the sender in round r (where r is chosen as described above).

We consider two cases, depending on whether P_s is in D.

If $P_s \in D$, then D's view after round r is determined by its view before that round: sending a message generated from its own state gives the coalition no new information. Since D knows whether it wins after round r (by our choice of r), it already knows this *before* that round. The adversary therefore follows the honest protocol until its pre-round-r joint view matches the corresponding view in Z, and performs the replacement only upon such a match. Otherwise, it continues honestly.

If $P_s \notin D$, the adversary instead follows the honest protocol through round r and performs the replacement only if its post-round-r joint view matches the corresponding view in Z. It then initializes the modified parties with their replacement pre-round-r states and processes the round-r message on those states. This reproduces the intended hybrid continuation: the sender P_s is outside D, so its message is unaffected by the replacement. Parties in D whose states are not replaced continue normally. If the triggering view is not observed, the adversary continues honestly.

In either case, the triggering view is consistent exclusively with honest executions whose leader lies outside D. Thus, on every execution where the adversary intervenes, its honest winning indicator is 0; on all other executions, it behaves honestly. The deviation therefore cannot decrease its winning probability. If the selected hybrid elected a member of D, execution Z, which occurs with positive probability, would yield a strict increase. This would contradict game-theoretic fairness and proves the claim.

We now proceed to use the filtered state replacement lemma to show the impossibility of CSP-fair leader election for $n < 3t$. We distinguish between two cases, depending on the number of parties n.

Case 1: $2t < n < 3t$**.** Split the parties into three groups A, B, C such that $|A| = |B| = t$, $|C| = n - 2t$. Consider the hybrid

$$H = (\mathsf{view}_A^X = \mathsf{view}_A^Y, \mathsf{view}_B^Y, \mathsf{view}_C^X)$$

where in honest execution X the leader is some $\mathsf{leader}_X \in A$, and in honest execution Y the leader is some $\mathsf{leader}_Y \in B$. H can be produced in two ways:

- If everyone in B is corrupt, H can be achieved by starting from execution X, and letting parties in B change their view to view_B^Y via the filtered state replacement described above. Let ℓ^* be the elected leader (determined by the output of the honest parties) in this continuation. Since $\mathsf{leader}_X \in A$, by the filtered state replacement lemma we get

$$\ell^* \notin B.$$

– If everyone in C is corrupt, H can be achieved by starting from execution Y, and letting parties in C change their view to view_C^X via the filtered replacement. As $|C| \leqslant t-1$, we can choose a size-t coalition D satisfying

$$C \cup \{\ell^*\} \subseteq D, \qquad \mathsf{leader}_Y \notin D.$$

This second attack produces the same hybrid continuation H. Furthermore, since $n > 2t$, there is a party outside $B \cup D$. This party is honest in both attacks and follows the same continuation from H, so it outputs ℓ^* in both. As we require that honest parties produce the same output, all honest parties in the second attack also output ℓ^*.

However, coalition D loses in the original execution Y, whose leader is $\mathsf{leader}_Y \notin D$, and wins in the hybrid, whose leader is $\ell^* \in D$. This contradicts the filtered state replacement lemma.

Case 2: $n \leqslant 2t$**.** Let $B = [n] \backslash A$. Then $|B| = n - t \leqslant t$. Choose a size-t coalition D such that

$$B \subseteq D, \qquad \mathsf{leader}_X \notin D.$$

Such a coalition exists because $t < n$.

Starting from execution X, coalition D replaces the states of parties in B with view_B^Y (other coalition members remain passive). As $\mathsf{view}_A^X = \mathsf{view}_A^Y$, the resulting global snapshot is exactly the snapshot of execution Y immediately before round r. Its continuation therefore elects $\mathsf{leader}_Y \in B \subseteq D$.

Thus, the replacement changes an honest outcome $\mathsf{leader}_X \notin D$ into an outcome $\mathsf{leader}_Y \in D$, again contradicting the filtered-replacement claim.

Thus, both cases yield a contradiction, and hence, there exists no perfect CSP-fair protocol Π for $n < 3t$.

Corollary 6 (Lower Bound for Perfect MM-Fairness, $n < 3t$**).** *Perfect MM-fair (Definition 3) leader election with bounded number of rounds is impossible for* $n < 3t$*.*

Proof. Any leader election protocol that satisfies t-MM fairness also satisfies t-CSP fairness, as per Remark 1. So, if a perfect MM-fair protocol existed for $n < 3t$, it would also be CSP-fair, violating Theorem 10.

Corollary 7 (Lower Bound for Perfect GT-Fairness, $n < 3t$**).** *Perfect GT-fair (Definition 4) leader election with bounded number of rounds is impossible for* $n < 3t$*.*

Proof. As with MM-fairness, any protocol that satisfies t-GT fairness also satisfies t-CSP fairness, which would violate Theorem 10.

Corollary 8 (Lower Bound for Perfect Strong Fairness, $n < 3t$**).** *Perfect strongly fair (Definition 1) leader election with bounded number of rounds is impossible for* $n < 3t$*.*

Proof. As before, any protocol that satisfies t-strong fairness also satisfies t-CSP fairness, which would violate Theorem 10.

9.2 Impossibility of MM Fair Leader Election for $n \leqslant 3t$

Theorem 11 (Lower Bound for Perfect MM-Fairness). *Perfect MM-fair (Definition 3) leader election with bounded number of rounds is impossible for* $n \leqslant 3t$*.*

Proof. By Corollary 6, we know that perfectly secure MM-fair leader election is impossible for $n < 3t$. In the following, we further rule out $n = 3t$.

Suppose, that there exists a perfect MM-fair leader-election protocol Π for $n = 3t$ parties. W.l.o.g., in the following we assume that in Π in each round at most one party sends messages.

Unlike in the CSP proof, where a coalition only needed to know whether the leader belonged to it, here we consider knowledge of the leader's identity. Let r be the earliest round such that, in every honest execution, every size-t coalition knows the identity of the elected leader by the end of round r. As before, such a round exists because the protocol terminates within a fixed number of rounds and all parties learn the leader; and further, $r > 0$.

As before, by the minimality of r, there exist a size-t coalition A and two honest executions X and Y whose joint views for A immediately before round r coincide, but whose leaders differ:

$$\mathsf{view}_A^X = \mathsf{view}_A^Y, \qquad \mathsf{leader}_X \neq \mathsf{leader}_Y.$$

Throughout the remainder of the proof, view_D^Z denotes the joint view of D immediately before round r in execution Z. We additionally write $\mathsf{view}_{D,+}^Z$ for the joint view of D immediately after round r in execution Z. These views include the parties' complete private random tapes. By the choice of r, for every size-t coalition D and every pair of honest executions W, Z, we have

$$\mathsf{view}_{D,+}^W = \mathsf{view}_{D,+}^Z \quad \Longrightarrow \quad \mathsf{leader}_W = \mathsf{leader}_Z.$$

In particular, observing $\mathsf{view}_{D,+}^Z$ certifies that honest continuation from the original states would elect leader_Z.

We now establish a strengthening of CSP's filtered state replacement lemma:

Lemma 5 (Strong filtered state replacement). *Fix a t-size coalition D and an honest execution Z whose leader leader_Z is outside D. Replacing the pre-round-r states of any subset of D with fixed alternative states, and subsequently following the protocol, cannot cause a party $P \neq \mathsf{leader}_Z$ to be elected with positive probability.*

Proof. Suppose otherwise. Fix such a coalition D, execution Z, and a replacement. We will now construct an adversary corrupting D that performs this replacement only when its local view certifies that honest continuation would elect leader_Z. Let P_s be the sender in round r. We consider two cases:

- *If $P_s \in D$:* The coalition's post-round-r view is determined by its pre-round-r view: sending messages computed from its own states gives it no new information. Since D knows the leader after round r, it already knows the leader before that round. More explicitly, for every honest execution W,

  $$\mathsf{view}_D^W = \mathsf{view}_D^Z \quad \Longrightarrow \quad \mathsf{view}_{D,+}^W = \mathsf{view}_{D,+}^Z \quad \Longrightarrow \quad \mathsf{leader}_W = \mathsf{leader}_Z.$$

  Thus, matching the pre-round-r view view_D^Z certifies that honest continuation would elect leader_Z.

  The adversary follows the honest protocol until immediately before round r. If the joint view of parties in D matches view_D^Z, it performs the replacement and subsequently follows the protocol. Otherwise, it continues honestly.
- *If $P_s \notin D$:* The adversary follows the honest protocol through round r. It performs the replacement only if its post-round-r joint view matches $\mathsf{view}_{D,+}^Z$.

Upon such a match, it initializes the modified parties with their replacement pre-round-r states and processes the round-r messages on those states. This reproduces the intended replacement: the sender is outside D, so its round-r messages are unaffected by the replacement. Coalition members whose states are not replaced continue normally. If the triggering view is not observed, the adversary continues honestly.

Let E denote the event that the triggering view is observed. Sample a vector of initial random tapes for all parties. For each such vector, compare two executions: one in which everyone follows the protocol honestly, and one in which the adversary corrupting D follows the strategy described above. Denote the elected leaders in these executions by ℓ_{hon} and ℓ_{adv}, respectively, where the latter is the common output of the honest parties.

Both executions start with the same random tapes. Consequently, they are identical until the adversary performs the state replacement; if the triggering event never occurs, they remain identical throughout. The filtering conditions established above imply

$$E \implies \ell_{\text{hon}} = \mathsf{leader}_Z.$$

Outside E, the adversary continues honestly, and hence $\ell_{\text{adv}} = \ell_{\text{hon}}$. Therefore,

$$\begin{aligned}\Pr[\ell_{\text{adv}} = \mathsf{leader}_Z] &= \Pr[\ell_{\text{hon}} = \mathsf{leader}_Z] - \Pr[E \wedge \ell_{\text{adv}} \neq \mathsf{leader}_Z] \\ &= \frac{1}{n} - \Pr[E \wedge \ell_{\text{adv}} \neq \mathsf{leader}_Z].\end{aligned}$$

The adversary reproduces the stipulated replacement on execution Z, which occurs with positive probability. If that replacement changed the honest parties' output, the final term would be strictly positive. Consequently,

$$\Pr[\ell_{\text{adv}} = \mathsf{leader}_Z] < \frac{1}{n},$$

contradicting MM-fairness because $\ell_Z \notin D$. This proves the claim.

The hybrid argument. Partition $\mathcal{P}\backslash A$ into sets B and C such that

$$|A| = |B| = |C| = t, \qquad \mathsf{leader}_X \notin B, \qquad \mathsf{leader}_Y \notin C.$$

Such a partition exists because $\mathsf{leader}_X \neq \mathsf{leader}_Y$: if $\mathsf{leader}_X \notin A$, place it in C; if $\mathsf{leader}_Y \notin A$, place it in B.

Consider the hybrid pre-round-r state

$$H = \left(\mathsf{view}_A^X = \mathsf{view}_A^Y,\, \mathsf{view}_B^Y,\, \mathsf{view}_C^X\right),$$

after which all parties follow the protocol. We now obtain two incompatible requirements on its continuation:

- Starting from X, corrupt B and replace its states with view_B^Y. This produces H. Since $\mathsf{leader}_X \notin B$, strong filtered state replacement requires every honest party to output leader_X.
- Starting from Y, corrupt C and replace its states with view_C^X. This also produces H. Since $\mathsf{leader}_Y \notin C$, strong filtered state replacement requires every honest party to output leader_Y.

Every party in A is honest in both attacks. The two attacks produce the same hybrid state and the same subsequent execution, so these parties must produce the same output in both. However, the first interpretation requires their output to be leader_X, whereas the second requires it to be $\mathsf{leader}_Y \neq \mathsf{leader}_X$. This is a contradiction.

Hence, perfect MM-fair leader election is impossible for $n = 3t$, and therefore for every $1 \leqslant t < n \leqslant 3t$.

Corollary 9. *Under the same assumptions, perfect GT-fair and perfect strongly fair leader election with bounded number of rounds are impossible for* $1 \leqslant t < n \leqslant 3t$.

Proof. Any leader-election protocol that satisfies t-GT fairness also satisfies t-MM fairness by definition. Likewise, t-strong fairness implies t-MM fairness: preserving every party's winning probability at $1/n$ in particular prevents the adversary from decreasing the winning probability of any honest party. Therefore, the impossibility of perfectly MM-fair leader election also rules out perfectly GT-fair and perfectly strongly fair leader election.

10 AI Methodology

The authors developed the vast majority of the protocols and lower bounds, as well as the proofs for all results without assistance from large language models (LLMs). The exception is Sections 8.2 and 9, where ChatGPT-6 Astra aided the design of the protocol for the general n case and for the lower bound (based on the authors' draft). The authors subsequently used the ChatGPT models 6 Astra as well as Claude models Fable 5.1.and Opus 5.5. LLMs assisted with writing and proof-checking throughout the paper and were used to create the figures. The authors take full responsibility for the correctness, originality, and integrity of all content, including material prepared with LLM assistance and the references.

References

ADGH06. Ittai Abraham, Danny Dolev, Rica Gonen, and Joseph Y. Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In Eric Ruppert and Dahlia Malkhi, editors, *25th ACM Symposium Annual on Principles of Distributed Computing*, pages 53–62, Denver, CO, USA, July 23–26, 2006. Association for Computing Machinery.

AN90. Noga Alon and Moni Naor. Coin-flipping games immune against linear-sized coalitions (extended abstract). In *31st Annual Symposium on Foundations of Computer Science*, pages 46–54, St. Louis, MO, USA, October 22–24, 1990. IEEE Computer Society Press.

BGW88. Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th Annual ACM Symposium on Theory of Computing*, pages 1–10, Chicago, IL, USA, May 2–4, 1988. ACM Press.

Blu81. Manuel Blum. Coin flipping by telephone. In Allen Gersho, editor, *Advances in Cryptology – CRYPTO'81*, volume ECE Report 82-04, pages 11–15, Santa Barbara, CA, USA, 1981. U.C. Santa Barbara, Dept. of Elec. and Computer Eng.

BOO10. Amos Beimel, Eran Omri, and Ilan Orlov. Protocols for multiparty coin toss with dishonest majority. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 538–557, Santa Barbara, CA, USA, August 15–19, 2010. Springer Berlin Heidelberg, Germany.

BZ17. Massimo Bartoletti and Roberto Zunino. Constant-deposit multiparty lotteries on Bitcoin. In Michael Brenner, Kurt Rohloff, Joseph Bonneau, Andrew Miller, Peter Y. A. Ryan, Vanessa Teague, Andrea Bracciali, Massimiliano Sala, Federico Pintore, and Markus Jakobsson, editors, *FC 2017 Workshops*, volume 10323 of *Lecture Notes in Computer Science*, pages 231–247, Sliema, Malta, April 7, 2017. Springer, Cham, Switzerland.

CCD88. David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *20th Annual ACM Symposium on Theory of Computing*, pages 11–19, Chicago, IL, USA, May 2–4, 1988. ACM Press.

CCWS21. Kai-Min Chung, T.-H. Hubert Chan, Ting Wen, and Elaine Shi. Game-theoretic fairness meets multiparty protocols: The case of leader election. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part II*, volume 12826 of *Lecture Notes in Computer Science*, pages 3–32, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland.

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 Amos Beimel and Stefan Dziembowski, editors, *TCC 2018: 16th Theory of Cryptography Conference, Part I*, volume 11239 of *Lecture Notes in Computer Science*, pages 563–596, Panaji, India, November 11–14, 2018. Springer, Cham, Switzerland.

Cha90. David Chaum. The spymasters double-agent problem: Multiparty computations secure unconditionally from minorities and cryptographically from majorities. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO'89*, volume 435 of *Lecture Notes in Computer Science*, pages 591–602, Santa Barbara, CA, USA, August 20–24, 1990. Springer, New York, USA.

Cle86. Richard Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In *18th Annual ACM Symposium on Theory of Computing*, pages 364–369, Berkeley, CA, USA, May 28–30, 1986. ACM Press.

Fei99. Uriel Feige. Noncryptographic selection protocols. In *40th Annual Symposium on Foundations of Computer Science*, pages 142–153, New York, NY, USA, October 17–19, 1999. IEEE Computer Society Press.

Fel87. Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science*, pages 427–437, Los Angeles, CA, USA, October 12–14, 1987. IEEE Computer Society Press.

GGS22. Rati Gelashvili, Guy Goren, and Alexander Spiegelman. Short paper: On game-theoretically-fair leader election. In Ittay Eyal and Juan A. Garay, editors, *FC 2022: 26th International Conference on Financial Cryptography and Data Security*, volume 13411 of *Lecture Notes in Computer Science*, pages 531–538, Grenada, May 2–6, 2022. Springer, Cham, Switzerland.

GK12. Adam Groce and Jonathan Katz. Fair computation with rational players. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 81–98, Cambridge, UK, April 15–19, 2012. Springer Berlin Heidelberg, Germany.

GMW87. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th Annual ACM Symposium on Theory of Computing*, pages 218–229, New York City, NY, USA, May 25–27, 1987. ACM Press.

HT04. Joseph Y. Halpern and Vanessa Teague. Rational secret sharing and multiparty computation: Extended abstract. In László Babai, editor, *36th Annual ACM Symposium on Theory of Computing*, pages 623–632, Chicago, IL, USA, June 13–16, 2004. ACM Press.

HT14. Iftach Haitner and Eliad Tsfadia. An almost-optimally fair three-party coin-flipping protocol. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 408–416, New York, NY, USA, May 31 – June 3, 2014. ACM Press.

IKLP06. Yuval Ishai, Eyal Kushilevitz, Yehuda Lindell, and Erez Petrank. On combining privacy with guaranteed output delivery in secure multiparty computation. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 483–500, Santa Barbara, CA, USA, August 20–24, 2006. Springer Berlin Heidelberg, Germany.

IML05. Sergei Izmalkov, Silvio Micali, and Matt Lepinski. Rational secure computation and ideal mechanism design. In *46th Annual Symposium on Foundations of Computer Science*, pages 585–595, Pittsburgh, PA, USA, October 23–25, 2005. IEEE Computer Society Press.

Kat07. Jonathan Katz. On achieving the "best of both worlds" in secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th Annual ACM Symposium on Theory of Computing*, pages 11–20, San Diego, CA, USA, June 11–13, 2007. ACM Press.

KMSW22. Ilan Komargodski, Shin'ichiro Matsuo, Elaine Shi, and Ke Wu. log*-round game-theoretically-fair leader election. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 409–438, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

KN08. Gillat Kol and Moni Naor. Cryptography and game theory: Designing protocols for exchanging information. In Ran Canetti, editor, *TCC 2008: 5th Theory of Cryptography Conference*, volume 4948 of *Lecture Notes in Computer Science*, pages 320–339, San Francisco, CA, USA, March 19–21, 2008. Springer Berlin Heidelberg, Germany.

LP11. Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. In Lance Fortnow and Salil P. Vadhan, editors, *43rd Annual ACM Symposium on Theory of Computing*, pages 705–714, San Jose, CA, USA, June 6–8, 2011. ACM Press.

LRM10. Christoph Lucas, Dominik Raub, and Ueli M. Maurer. Hybrid-secure MPC: trading information-theoretic robustness for computational privacy. In Andréa W. Richa and Rachid Guerraoui, editors, *29th ACM Symposium Annual on Principles of Distributed Computing*, pages 219–228, Zurich, Switzerland, July 25–28, 2010. Association for Computing Machinery.

LZMRT26. Chen-Da Liu-Zhang, Elisaweta Masserova, João Miguel Lourenço Ribeiro, and Sri Aravinda Krishnan Thyagarajan. Game theory does not always help: The case of statistical multi-party coin tossing. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026, Part III*, volume 16543 of *Lecture Notes in Computer Science*, pages 488–514, Rome, Italy, May 10–14, 2026. Springer, Cham, Switzerland.

Mas26. Elisaweta Masserova. On the impossibility of perfectly secure coin-toss for $n < 3t + 1$. Unpublished manuscript, 2026.

MNS09. Tal Moran, Moni Naor, and Gil Segev. An optimally fair coin toss. In Omer Reingold, editor, *TCC 2009: 6th Theory of Cryptography Conference*, volume 5444 of *Lecture Notes in Computer Science*, pages 1–18. Springer Berlin Heidelberg, Germany, March 15–17, 2009.

RB89. Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *21st Annual ACM Symposium on Theory of Computing*, pages 73–85, Seattle, WA, USA, May 15–17, 1989. ACM Press.

RZ98. Alexander Russell and David Zuckerman. Perfect information leader election in log*n + O(1) rounds. In *39th Annual Symposium on Foundations of Computer Science*, pages 576–583, Palo Alto, CA, USA, November 8–11, 1998. IEEE Computer Society Press.

TSW24. Sri Aravinda Krishnan Thyagarajan, Pratik Soni, and Ke Wu. Game-theoretically fair distributed sampling. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part VIII*, volume 14927 of *Lecture Notes in Computer Science*, pages 207–239, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

WAS22. Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multiparty coin toss. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part I*, volume 13275 of *Lecture Notes in Computer Science*, pages 120–149, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland.

ZW25. Forest Zhang and Ke Wu. Game-theoretically fair coin toss with arbitrary preferences. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology – ASIACRYPT 2025, Part VIII*, volume 16252 of *Lecture Notes in Computer Science*, pages 549–581, Melbourne, VIC, Australia, December 8–12, 2025. Springer, Singapore, Singapore.