

# Game-Theoretically Fair Distributed Coin Tossing With Private Preferences

Pedro Branco\*, Pratik Soni†, Sri AravindaKrishnan Thyagarajan‡, Ke Wu§

## Abstract

Secure coin-tossing is typically modeled as an *input-less* functionality, where parties with no private inputs jointly generate a fair coin. In the dishonest majority setting, however, a strongly fair coin-tossing protocol is impossible. To circumvent this barrier, recent work has adopted the weaker notion of *game-theoretic fairness*, where adversaries are rational parties with preferences for specific outcomes, seeking to bias the coin in their favor. Yet these preferences may encode secret information, making prior protocols that assume preferences are public, fundamentally incompatible with privacy.

We *initiate* a comprehensive study of *privacy-preserving* game-theoretically fair coin-tossing, where the preferences of honest parties remain private. We propose a simulation-based security framework and a new ideal functionality that reconciles *both* preference-privacy and game-theoretic fairness. A key ingredient is a certifying authority that authenticates each party’s preference and publishes only aggregate statistics, preventing misreporting while hiding parties’ preferences. The functionality guarantees that every honest party receives an output: either a uniform coin; or, if an adversary deviates, a coin that strictly decreases the adversarial coalition’s expected utility.

Within this framework, we construct a protocol realizing our ideal functionality under standard cryptographic assumptions that works for both binary and general  $m$ -sided coin-tossing. Our schemes tolerate the same optimal (or nearly optimal) corruption thresholds as the best known protocols with public preferences (Wu-Asharov-Shi, EUROCRYPT ’22; Thyagarajan-Wu-Soni, CRYPTO ’24). Technically, our protocols combine authenticated preferences with an anonymous communication layer that decouples identities from preference-dependent actions, together with a deviation-penalty mechanism that enforces game-theoretic fairness.

Our work is the first to reconcile game-theoretic fairness with preference privacy, offering new definitional tools and efficient protocols for rational multi-party computation in dishonest majority settings.

---

\*Bocconi University, [pedrodemelobranco@gmail.com](mailto:pedrodemelobranco@gmail.com)

†Kahlert School of Computing, University of Utah, [psoni@cs.utah.edu](mailto:psoni@cs.utah.edu)

‡School of Computer Science, University of Sydney, [t.srikrishnan@gmail.com](mailto:t.srikrishnan@gmail.com)

§Electrical Engineering and Computer Science, University of Michigan, [kewucse@umich.edu](mailto:kewucse@umich.edu)

# 1 Introduction

Designing protocols for generating useful randomness in a distributed setting is fundamental to both Cryptography and Distributed Computing. Consider three friends, Alice, Bob, and Charlie, who want to toss a coin over a group call to decide on a cafe to meet. Each prefers one of the only two cafes, but wants to appear neutral while keeping their preference private and trying to bias the outcome towards their preference. This scenario involves two key requirements: (a) *Fairness* stipulates that the chosen coin is uniform (i.e., is 0 with probability 0.5), and (b) *preference-privacy* that requires that no coalition of malicious parties learns any additional information about the honest participants’ preferences, beyond the chosen coin.

Coin-tossing is typically formalized as an *input-less* functionality (where parties have no input); hence, all prior works focus exclusively on fairness. When the majority of the parties act honestly, the above notion of fairness that the chosen coin must be uniform (generally referred to as *strong fairness*) can indeed be achieved [Yao82, GMW87, BGW88, CCD88, RB89]. In fact, existing protocols that achieve strong fairness work independently of the parties’ preferences, hence vacuously achieving preference-privacy.

However, the situation changes dramatically in the presence of a dishonest majority. In the two-party setting, the celebrated result by Cleve [Cle86] rules out the existence of a strongly fair coin-tossing protocol against a single corruption. Moreover, Cleve’s negative result can be strengthened to rule out strong fairness in the presence of a dishonest majority for  $n \geq 2$  parties.

**Fairness in the Dishonest Majority Setting.** In light of this, the recent work of Chung, Guo, Lin, Pass, and Shi [CGL<sup>+</sup>18a] (henceforth, CGLPS18) developed *game-theoretic* notion of fairness that models adversaries as rational parties who are looking to increase their utilities. In particular, they showed that for several settings where coin-tossing is used as a subroutine, each party has an implicit game-theoretic notion of utility. For example, if  $P_0$  and  $P_1$  were running the coin-tossing protocol to elect a leader amongst them, each party  $P_b$  would prefer that the sampled coin be  $b$  so that they get elected. More formally, a party  $P_b$  gains utility 1 if the outcome matches its preference  $b$  or utility 0 otherwise. Then, CGLPS18’s game-theoretic fairness requires that any party cannot *increase* their expected utility by deviating from the protocol, compared to its utility in an honest run.<sup>1</sup>

CGLPS18 observed that a variant of Blum’s two-party coin-tossing protocol [Blu81] already satisfies this notion of game-theoretic fairness against single corruption, hence circumventing Cleve’s impossibility. This led to subsequent works on the feasibility of *multi*-party sampling protocols that achieve game-theoretic fairness against dishonest majority coalitions: for binary coin-toss [WAS21] gave a complete characterization of tolerable corruption thresholds, for leader election (or  $n$ -sided coin-toss) [CCWS21, KMSW22] explore round optimal protocols, and [TSW24] give a complete characterization of tolerable corruption thresholds for the general case of sampling from any efficiently sampleable distribution.

**Public Preferences - A Privacy Concern.** While coin-tossing is typically modeled as an input-less functionality, in the context of game-theoretic fairness, each party actually holds individual preferences as inputs. These preferences could encode information that, depending on the application setting, parties may want to keep secret. Therefore, a truly secure coin-tossing protocol

---

<sup>1</sup>CGLPS18 refer to this notion as Cooperative-Strategy-Proofness (CSP)-fairness. Other notions of fairness such as maximin fairness considered in the literature become infeasible in the strict dishonest majority setting (our focus), hence considered beyond the scope of this work.

must not only ensure fair coin sampling but also preserve the privacy of honest parties’ preferences, akin to secure multi-party computation protocols.

Despite this, all prior works on designing game-theoretically fair protocols completely ignore such privacy concerns. Moreover, the designed protocols inherently require the preferences of each party to be known *publicly*. As an example, consider the binary coin-tossing protocol of [WAS21], where the  $n$  parties are grouped according to which of the two outcomes they prefer, and these groups act in different phases of the protocol. Consequently, while being part of a group is necessary for the protocol specification, it inadvertently publicly reveals all parties’ preferences.

At a definitional level, game-theoretic fairness does require the parties to have pre-determined preferences to make parties’ utility functions well-defined. Yet there is no inherent reason that these preferences must be publicly known. This observation motivates the question:

*Can we build  $n$ -party protocols for tossing an  $m$ -sided coin that satisfies both game-theoretic fairness and preference-privacy against a dishonest majority?*

## 1.1 Our Results

In this work, we answer the above question affirmatively. Surprisingly, even defining a meaningful yet feasible security notion capturing both game-theoretic fairness and preference-privacy is quite non-trivial. We propose a new definitional framework for authenticated-preference coin-tossing and provide our new ideal coin-tossing functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  that combines game-theoretic fairness and preference-privacy. Below, we will describe our choice of definition and describe the protocols that realize the ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ . Finally we conclude with some philosophical discussions on the new definitional framework in Section 1.2.

### 1.1.1 The Need for Authenticated Preferences

Recall that in the public preference setting, coin-tossing protocols and fairness can be naturally defined with respect to a preference profile. However, this is not the case for private preference. CGLPS18 briefly explored the case of private preference. They considered universal  $n$ -party private-preference coin-tossing protocols that is defined to work for every possible preference profile. Unfortunately, CGLPS18 show that a universal protocol that satisfies game-theoretic fairness against malicious adversaries becomes impossible. The obstacle lies in *dynamic misreporting*: in any protocol where parties are free to declare preferences during execution, a malicious coalition can report preferences that differ from their true ones, while still measuring utility against their actual (hidden) preferences. This creates a definitional inconsistency: To evaluate fairness, one must compare the adversary’s expected utility to its “honest baseline” utility. But if the coalition can misreport its preferences, the honest baseline no longer reflects what the adversary truly values. As a result, even if a protocol is fair with respect to declared preferences, a coalition can still bias the outcome toward its true preference by strategic misreporting. CGLPS18 formalized this gap and proved that fairness with private preferences is impossible in the malicious model. Feasibility can only be recovered under weaker semi-malicious assumptions, where adversaries are forced to commit to their true preferences.

To overcome this barrier, parties must be held accountable to their declared preferences. Without such accountability, the impossibility of CGLPS18 persists.

We address this by introducing *authenticated preferences*, inspired by prior works on authenticated MPC [DGPS22]. Before joining the protocol, each party interacts with a *certifying authority* that signs the party’s stated preference. A preference is considered authentic or certified only if it

carries a signature verifiable under the authority’s public verification key. Once certified, a party is bound to this preference throughout the protocol execution, preventing adaptive “switching”.

In contrast to prior work on authenticated MPC, we require the authority to be *stateful*: after processing all certification requests, it publishes an aggregate statistics of the reported preferences. For example, in the binary coin-tossing case, the authority releases  $(n_0, n_1)$ , where  $n_b$  is the number of parties that report preference  $b$  to the authority. While this assumption may appear strong at first glance, we will explain why publishing aggregates is crucial to feasibility of both preference-privacy and game-theoretic fairness (see Section 1.2).

Naturally, a malicious party may still misreport to the authority at the time of certification. We do not attempt to prevent this entirely, but two factors mitigate its impact. First, in many applications the authority (e.g., a regulator, membership registry, or government agency) can independently validate preferences, making misreporting risky. Second, certification must be obtained before the adversary learns honest parties’ preferences, therefore preventing adversaries from misreporting strategically based on any information they learned about honest parties.

Thus, while authentication does not eliminate every form of dishonesty, it decisively rules out the *dynamic misreporting attack* that underlies the impossibility result of CGLPS18. This enables a meaningful definitional framework for game-theoretic fairness with private preferences. We emphasize that to the best of our knowledge, some form of setup or authority is unavoidable in this setting (as dynamic misreporting makes purely distributed solutions impossible). Our model uses a certifying authority in the weakest possible way – it is only involved in an initial setup phase and never in the coin-tossing itself. We direct the reader to Section 1.2 for more discussions.

### 1.1.2 Defining Fairness and Preference-Privacy

A natural approach for capturing both security properties of fairness and preference-privacy would be to define them independently. Particularly, for preference-privacy, we can adopt the simulation-based definition commonly used in secure multi-party computation and zero-knowledge proofs [Yao82, GMW87, Can01]. On the other hand, game-theoretic fairness is typically captured via a property-based definition that guarantees the joint utility of the adversarial coalition must be negligibly close to the coalition’s expected honest utility [CGL<sup>+</sup>18a].

In this work, we instead take a *unified* approach and propose a new simulation-based definition that reconciles both game-theoretic fairness and privacy. Our definition follows the standard real-and-ideal-world paradigm that is considered the gold standard notion for defining security for any multi-party computation [Yao82, GMW87, Can01]. Specifically, we introduce an ideal coin-tossing functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  that appropriately models both game-theoretic fairness and preference-privacy.

**Ideal Coin-tossing Functionality with Authenticated Preferences.** In Section 3, we formally describe our ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, (n_0, \dots, n_{m-1})]$  for  $m$ -sided coin-tossing. Our ideal functionality is additionally parameterized by the verification key  $\text{mvk}$  of the certifying authority and the aggregate information  $(n_0, \dots, n_{m-1})$  that is published by the certifying authority, where  $n_i$  denotes the number of parties reporting preference  $i$  to the authority.<sup>2</sup>

For simplicity, let us consider tossing a binary coin and consider an adversary  $\mathcal{A}$  that controls a set  $\mathcal{K} \subseteq [n]$  of at most  $t$  parties. Let  $\mathcal{H} = [n] \setminus \mathcal{K}$  be the set of honest parties. We assume that all parties identities (e.g., public-keys) are publicly known. For such an  $\mathcal{A}$ , the ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}[2, \text{mvk}, (n_0, n_1)]$  works as follows:

<sup>2</sup>In fact, our ideal functionality only requires  $n_0$  and  $n - n_0$  where 0 is the least preferred outcome. We state all aggregates here for simplicity.

1. **Reporting by Honest Parties.** Parties  $i \in \mathcal{H}$  sends their  $(\text{pref}_i, \sigma_i)$  to the ideal functionality which verifies the validity of  $\sigma_i$  as a signature on  $(i, \text{pref}_i)$  under key  $\text{mvk}$ . If  $\sigma_i$  is invalid or  $\text{pref}_i \notin \{0, 1\}$ , then set  $\text{pref}_i = \perp$ . For  $b \in \{0, 1\}$ , functionality computes  $h_b$  as the number of parties  $i \in \mathcal{H}$  with  $\text{pref}_i = b$ , and sets  $k_b = n_b - h_b$ . It then sends  $(h_0, h_1)$  to  $\mathcal{A}$  as leakage.<sup>3</sup>
2. **Reporting by Corrupt Parties.** Parties  $i \in \mathcal{K}$  send their  $(\text{pref}_i, \sigma_i)$  to the ideal functionality which verifies the validity of  $\sigma_i$  as a signature on  $(i, \text{pref}_i)$  under key  $\text{mvk}$ . If  $\sigma$  is invalid or  $\text{pref}_i \notin \{0, 1\}$ , then set  $\text{pref}_i = \perp$ .
3. The functionality proceeds to compute a coin  $c' \in \{0, 1\}$  that  $\mathcal{A}$  does *not* prefer as follows: if  $k_0 > k_1$ , then set  $c' = 1$ . Otherwise, set  $c' = 0$ .
4. Then, the ideal functionality samples a uniform random coin  $c \xleftarrow{\$} \{0, 1\}$  and sends it to  $\mathcal{A}$ . If  $\mathcal{A}$  accepts the coin, the ideal functionality returns  $c$  to all honest parties. Otherwise, it sends the coin  $c'$  that  $\mathcal{A}$  does *not* prefer to all honest parties.

Our functionality is reminiscent of the ideal coin-tossing functionality that captures simulation-security with abort [Lin16], albeit for authenticated inputs. Specifically, the security with abort notion allows  $\mathcal{A}$  to prevent honest parties from learning the coin by “aborting” the ideal functionality. In comparison, our ideal functionality *always* delivers an output to all honest parties, irrespective of whether  $\mathcal{A}$  choosing to abort. If  $\mathcal{A}$  does not abort, the delivered coin is the uniformly chosen coin  $c$ ; whereas if  $\mathcal{A}$  aborts, the ideal functionality delivers the coin  $c'$  that necessarily reduces  $\mathcal{A}$ ’s utility. In addition, our ideal functionality enables embedding preference authentication into the coin-tossing protocol, thus forcing adversaries to use authenticate preferences.

Our goal is to design a protocol that realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[2, \text{mvk}, (n_0, n_1)]$  in the presence of the certifying authority. Philosophically, realizing  $\mathcal{F}_{\text{toss}}^{\text{auth}}[2, \text{mvk}, (n_0, n_1)]$  ensures that no coalition controlled by  $\mathcal{A}$  can either increase its expected joint utility or learn any additional information about honest parties’ preferences beyond what is revealed by  $(n_0, n_1)$  and the computed coin.

### 1.1.3 Binary Coin-Toss: Realizing $\mathcal{F}_{\text{toss}}^{\text{auth}}[2, \text{mvk}, (n_0, n_1)]$

For a given set of  $n$  parties, we use  $\mathcal{P}$  to denote the *preference profile*, a tuple mapping parties with their reported preferences. Further, we assume the existence of authenticated pairwise private channels and a broadcast channel. We assume that the parties have already requested for certifications on their preferences, and the aggregate information  $(n_0, n_1)$  was already published. Without loss of generality, we assume that 0 is the least preferred outcome (i.e.,  $n_0 < n_1$ ).

With the above notation, we design a protocol  $\Pi_{n,2}$  that securely realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[2, \text{mvk}, (n_0, n_1)]$ , which allows us to toss a unbounded polynomial number of coins in sequence.

**Theorem 1.1** (Binary coin-toss, informal). *Let  $\mathcal{P}$  be an arbitrary preference profile for  $n$  parties where  $n_0 < n_1$ . Assuming bounded concurrent zero-knowledge arguments, semantically secure public-key encryption and secure digital signature scheme, there exists a protocol  $\Pi_{n,2}$  that securely realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[2, \text{mvk}, (n_0, n_1)]$  against  $t$  corruptions, where*

$$t = \begin{cases} n - \lfloor \frac{3}{2}n_0 \rfloor, & \text{if } n \geq \frac{7}{2}n_0; \\ \lfloor \frac{2}{3}n - \frac{5}{6}n_0 \rfloor + \lceil \frac{1}{2}n_0 \rceil, & \text{otherwise.} \end{cases}$$

<sup>3</sup>At first glance, one might think that any adversary  $\mathcal{A}$  controlling a coalition already knows  $h_b$ , since it can combine its own preferences with the publicly available aggregation  $n_b$ . However, in the simulation-based setting the real-world adversary need not reveal these preferences to the simulator. Although our simulator extracts the coalition’s preferences, an adversary that aborts prematurely prevents the simulator from extracting all corrupted parties preferences, leaving the simulator unable to determine the true value of  $h_b$  needed to complete the simulation. In our protocols, knowledge of  $h_b$  is essential, as the simulator relies on it to set up the honest parties’ roles correctly.

Semantically-secure PKE and digital signatures can be instantiated from a variety of concrete algebraic assumptions. Moreover, bounded concurrent zero-knowledge proofs can be constructed from enhanced trapdoor permutations and collision resistant hash functions in [Pas04].

#### 1.1.4 General $m$ -sided Coin-Toss: Realizing $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{Leak}^{\text{agg}}(\cdot)]$

Following the technique of Thyagarajan, Wu, and Soni [TSW24], our protocol can be generalized to multi-sided coin-tossing. Given an arbitrary preference profile  $\mathcal{P}$ , without loss of generality, assume that 0 is the least preferred outcome. Let  $n_0$  denote the number of players preferring outcome 0.

**Theorem 1.2** ( $m$ -sided coin-toss, informal). *Let  $\mathcal{P}$  be an arbitrary preference profile for  $n$  parties where 0 is the least preferred outcome. Assuming bounded concurrent zero-knowledge arguments, semantically secure public-key encryption and secure digital signature scheme, there exists a protocol  $\Pi_{n,m}$  that securely realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, n_0, n - n_0]$  against  $t$  corruptions, where*

$$t = \begin{cases} \lceil \frac{1}{2}n_0 \rceil + n - m \cdot n_0, & \text{if } n \geq \frac{4m-1}{2}n_0 \\ \lceil \frac{1}{2}n_0 \rceil + \lfloor \frac{m}{2m-1}n - \frac{3m-1}{2(2m-1)}n_0 \rfloor, & \text{otherwise.} \end{cases}$$

As a comparison, for binary coin-tossing, our protocol matches the corruption threshold of [WAS22]; for  $m$ -sided coin-tossing, it matches that of [TSW24]. Unlike these prior works, which require public knowledge of all parties' preferences and proves game-theoretic fairness under a property-based definition, our protocol achieves the same thresholds while simultaneously ensuring preference-privacy and realizing a simulation-based notion of fairness, albeit under the assumption of authenticated preferences.

**Some Concrete Parameters.** The above relation between  $t, n, n_0$  is arguably complex. Below, we give some concrete parameters for better understanding the thresholds.

- When the number of players preferring each outcome is the same, i.e.,  $n = m \cdot n_0$ , our protocol  $\Pi_{n,m}$  tolerates any  $\lceil \frac{n}{2} \rceil$ -sized coalition.

For binary coin-toss, our protocol  $\Pi_{2,m}$  tolerates any  $\lceil \frac{n}{2} \rceil$ -sized coalition when  $n = 2n_0$ .

- When  $n$  is much larger than  $m \cdot n_0$ , say  $n = 10 \cdot m \cdot n_0$ , our protocol tolerates any  $0.9n$ -sized coalitions.

For binary coin-toss, our protocol  $\Pi_{2,m}$  tolerates any  $0.85n$ -sized coalition when  $n \geq 10n_0$ .

- When  $n_0 = 1$ , our protocol  $\Pi_{n,m}$  tolerates  $(n - m)$ -sized coalition.

For binary coin-toss, our protocol  $\Pi_{2,m}$  tolerates any  $(n - 1)$ -sized coalition.

## 1.2 Discussion on our Definition

Our definition introduces a certifying authority that both authenticates preferences and publishes aggregate statistics. While these assumptions may appear strong, we view them as the minimal and natural ingredients needed to overcome the impossibility of fairness with private preferences (as shown by CGLPS18). Below we explain why each aspect of our model is necessary, how it aligns with practice, and why it enables the strongest definitional guarantees currently achievable.

**Authenticated preferences.** Authentication is essential to rule out the dynamic misreporting attack that underlies the malicious-case impossibility. Without being bound to a certified preference, a coalition could declare different preferences opportunistically and bias the protocol toward its true (hidden) desires, a strategy implicit in Cleve’s impossibility. For instance, consider five parties where a coalition of size three has true preferences  $(0, 0, 1)$  (jointly preferring 0). The coalition misreports  $(0, 1, 1)$  (jointly preferring 1) instead. While Cleve’s strategy indicates that the coalition must be able to bias the outcome (towards some bit), game-theoretic fairness rule out all strategies that bias towards 1 based on their reported preferences  $(0, 1, 1)$ . Therefore, the coalition is able to bias the outcome towards 0, their true joint preferences.

By contrast, authenticated preferences rules out this adaptation and ensures that once a party has obtained a signature on its stated choice, it is committed to that choice throughout the protocol. Of course, parties may still misreport to the authority in the certification stage, but in many applications (e.g., when the authority is a regulator or membership registry) the truthfulness of preferences can be verified. More importantly, adversaries must commit before seeing others’ preferences, preventing them from tailoring lies strategically.

**On Stateful Authority and Aggregate Information.** The role of the authority extends beyond authentication: it is also stateful and publish aggregate statistics (e.g., the number of parties preferring 0 versus 1). If this information is not published, two challenges arise. First, many fair-coin protocols, including ours, internally require aggregate preference information. Absent the authority, the parties would need to compute this information themselves, which in the dishonest-majority setting requires guaranteed output delivery MPC, known to be impossible. Second, one might design protocols that ignore parties that fail to provide valid authenticated preferences (e.g., using commit-and-prove with identifiable abort), but such designs remain vulnerable: the coalition could provide unauthenticated preferences for enough parties (in the coalition) such that the “alive” portion of an adversarial coalition prefers a different value than the full coalition’s preference. This “alive” coalition can then perform the Cleve’s strategy to bias the outcome towards the whole coalition’s preference as described above. The publication of aggregate information prevents this adaptation.

One might ask whether fairness could be achieved with only authentication from the authority, but without requiring it to provide aggregate statistics. While this remains an open theoretical possibility, we are unaware of any prior constructions, and we conjecture that some aggregate information is inherent to any meaningful notion of fairness in this setting. Indeed, even defining a coherent ideal functionality without  $(n_0, n_1)$  appears challenging, since the fairness condition must be evaluated against the distribution of preferences.

**Coin-tossing in Presence of an Authority.** Our model does not delegate coin-tossing itself to the authority, despite its trusted role in setup. The reason is both scalability and practicality: we envision the authority’s role as a one-time initialization step, after which parties can independently toss an unbounded sequence of coins. This avoids requiring the authority’s continuous availability while still preserving fairness and privacy. In this respect, our use of an authority parallels authenticated MPC, where the authority ensures correct inputs but does not participate in the actual computation.

**On Composability.** Our framework provides the first simulation-based formulation of game-theoretic fairness with private preferences that enjoys composability. Previous property-based notions lacked such guarantees, limiting their utility as subroutines in larger protocols. By aligning



with the real/ideal paradigm familiar from secure computation, our definition ensures that fair sampling protocols can be modularly embedded into higher-level tasks.

Taken together, these design choices reflect a deliberate balance: while our model imposes additional structure through a certifying authority, this structure appears indispensable for reconciling fairness, privacy, and malicious security. In this sense, we view our framework as the strongest currently feasible approach, and as a foundation for future refinements that may further reduce assumptions.

### 1.3 Additional Related Work

Exploring the connection between game theory and multi-party computation is not new. A line of work [HT04, KN08, ADGH06, OPRV09, AL11, ACH11] explored rational multi-party computation. However, the utility they consider differs from the line of work on game-theoretic fairness. Specifically, they define the utility based on the assumption that players prefer to compute the function correctly and learn others' secrets while not leaking their secrets.

Recently, with the success of blockchains and digital payments, many works explored a notion called financial fairness [BK14, KMS<sup>+</sup>16, ADMM16, KB14, KVV16, DEF18, ATM<sup>+</sup>22, CMST22] based on the use of collaterals. Parties are required to put in collaterals beforehand and get punished if they misbehave. Interestingly, our work and the previous work on game-theoretic fairness [CGL<sup>+</sup>18b, WAS22, CCWS21, KMSW22, TSW24] guarantees incentive compatibility even without the use of any collateral.

A complementary line of works focus on designing multi-party coin-tossing that achieve security with abort in the dishonest majority setting. In such constructions, the emphasis is on building a round optimal protocol in the plain model that tolerates  $n - 1$  corruptions. Yao's [Yao82] proposed a constant round protocol for securely computing any two-party functionality (including coin-tossing) against semi-honest adversaries. The following work of [GMW87] then presented a general-purpose compiler for lifting a semi-honest multi-party protocol to a malicious security, but incurs polynomial rounds. Since this a long line of works [BMR90, Pas04, ACJ17, BHP17, IPS08, Wee10, Goy11, IKSS21] (not exhaustively cited) has optimized the round complexity of coin-tossing (and more generally, secure computation), culminating in round optimal protocols based on standard cryptographic assumptions [BGJ<sup>+</sup>18, HHPV21, RCCG<sup>+</sup>20].

A line of work has studied how to enforce the integrity of parties' inputs in secure computation via *authenticated inputs*. Dutta et al. [DGPS22] give a general compiler that transforms any honest-majority, linear-secret-sharing based MPC into one with authenticated inputs, achieving strong security guarantees such as identifiable abort without costly, protocol-specific preprocessing. While our focus differs, aiming to reconcile game-theoretic fairness with preference-privacy in dishonest majority coin-tossing, our use of authenticated preferences follows the same insight: binding parties to certified inputs is often essential to overcome fundamental impossibility results.

**Organization.** In Section 2, we illustrate the overview of the techniques used in our protocol. The formal model is given in Section 3, and the coin-tossing protocol is presented in Section 4.

## 2 Technical Overview

In this section, we present a brief overview of our techniques. For a better understanding, we will first explain how to construct a binary coin-tossing protocol that securely realizes the ideal functionality against semi-malicious coalitions, who can program its randomness used in the protocol



or abort early, but otherwise follows the protocol honestly. Then we explain how to upgrade to malicious security and multi-sided coin-toss.

Recall that parties get their preference certified by an authority prior to the coin-tossing protocol. We call those who prefer outcome  $b \in \{0, 1\}$  as  $b$ -supporters. If the outcome of the coin-tossing protocol matches the party’s authenticated preference, the party gets utility 1 and 0 otherwise. The expected joint utility of a coalition is the sum of all coalition member’s expected utility.

Before presenting our protocol, we first outline the public preference protocol from previous works and explain where the preference of each party is being leaked. Then we show how to hide the parties’ preferences using anonymous broadcast. Jumping ahead, the anonymous broadcast required in our protocol can be instantiated using only Shamir secret sharing [Sha79].

**Why Previous Protocols Leak Preferences.** Previous game-theoretic fair protocols [WAS22, TSW24] assumes that each party’s preference is publicly known. They work in the following framework: Let  $G_b$  denote the set of  $b$ -supporters, for  $b \in \{0, 1\}$ .

1. Each 0-supporters first interacts with other 0-supporters in group  $G_0$  to agree on a coin  $c_0$  using Shamir’s secret sharing, during which  $G_1$  audits. Similarly, each 1-supporter interacts with other 1-supporters in  $G_1$  to decide a coin  $c_1$  using secret sharing.
2. Groups  $G_0$  reconstructs coin  $c_0$ , and then group  $G_1$  reconstructs coin  $c_1$ .

The outcome of the protocol is then defined based on  $c_0$  and  $c_1$ .

In the above protocol, 0-supporters only interact with 0-supporters, while 1-supporters only interact with 1-supporters via secret sharing. Clearly, protocols with such structure cannot protect parties’ preferences: For each party to send messages, they need to know the recipients’ identities and preferences. Otherwise, how would they know where to send the message?

**Our Protocol: Intuition.** In a nutshell, our protocol adds an anonymous network mechanism on top of the previous framework to protect the identity of both the sender and the receiver of any message being sent in the protocol, so that it is infeasible to link a party with its preference. Only the aggregate information is revealed. Specifically, we need the following properties:

1. *Anonymity*: Neither the sender nor the receiver can identify each other’s real identity.
2. *Privacy*: For any message, its content, sender identity, and receiver identity remain hidden from all other parties.

Property 1 above ensures that honest parties’ preferences are unlinkable via anonymous communication. Property 2 guarantees that all the private messages remain hidden from the adversary to ensure game-theoretic fairness.

*Why Natural Approaches Fail.* A natural approach is to use an anonymous broadcast channel such as DC-nets [Cha88] that hides the sender’s identity for any broadcast messages. A straw-man attempt may instruct a party to send all messages via this anonymous broadcast channel. However, broadcasting all messages directly compromises Property 2. Therefore, we need the senders to first encrypt a message under the receiver’s public key for some public-key encryption scheme,<sup>4</sup> and then broadcast it through the anonymous broadcast channel. In this way, only the correct receiver can decrypt the message, while no other parties learn the message content, thus achieving Property 2.

There is still one main challenge: how does a sender know who to send messages to? In the previous public-preference protocols [WAS21, TSW24],  $b$ -supporters only interact with  $b$ -supporters. Even with anonymous broadcast channel, the sender needs to know which public keys corresponds to parties with the same preferences to encrypt the message. Thus, the sender needs to know

---

<sup>4</sup>This public key is different from the public key that represents parties’ identities. This public key is used to encrypt the corresponding message and is independent from the public key representing a party’s identity.

the public keys for those who have the same preferences, while honest parties’ preferences need to remain hidden.

**Our Approach: Masked Profile.** To achieve this, we ask each party  $i$  to randomly generate a virtual ID  $\text{vID}_i$  and submits  $(\text{vID}_i, \text{pref}_i)$  to the certificate authority to get authentication. This virtual ID will be used throughout the protocol as party  $i$ ’s ID, and we need to hide the mapping between parties’ real identities and their virtual identities.

At the beginning of the protocol, all  $n$  parties enter a *commit phase* by anonymously broadcasting a tuple of their virtual IDs, an encryption key used to encrypt messages to party  $i$ , and its preference  $(\text{vID}_i, \text{pk}_i, \text{pref}_i)$ . At the end of commit phase, parties agree on a “masked preference profile”  $\mathcal{P} = \{(\text{vID}_1, \text{pk}_1, \text{pref}_1), \dots, (\text{vID}_n, \text{pk}_n, \text{pref}_n)\}$ , which itself is a multiset that hides the mapping between party  $i$  and its tuple  $(\text{vID}_i, \text{pk}_i, \text{pref}_i)$ . Some corrupted parties may fail to broadcast a valid tuple and their preferences are set to  $\perp$ . Therefore, each  $\text{pref}_i \in \{0, 1, \perp\}$  in  $\mathcal{P}$ .

With this masked preference profile, one can imagine a public-preference scenario where the preference of each virtual ID is publicly known, except that now some parties may have  $\perp$ -preferences. Parties can now interact with others using a similar framework as for public-preferences using virtual IDs and anonymous broadcast channel, but adapt to also account for  $\perp$ -preferences. For example, if a 0-supporter with virtual ID  $\text{vID}_i$  wants to send a message to another 0-supporter with virtual ID  $\text{vID}_j$ , party  $\text{vID}_i$  encrypts the message using  $\text{pk}_j$  and broadcast it anonymously. This achieves both desired properties.

*Instantiating Anonymous Broadcast.* When instantiating the protocol in the real world, we cannot assume the availability of an ideal anonymous broadcast channel that always succeeds. In fact, when defending against majority-sized coalitions, the anonymous broadcast may fail. To handle this, we adopt the identifiable abort approach from [KMSW22]. Consider the following anonymous broadcast with identifiable abort functionality: either the anonymous broadcast succeeds or all honest parties receive the identities of some misbehaving parties. The parties kick out those misbehaved parties and retry the anonymous broadcast with the same message. This functionality guarantees that anonymity is preserved, and the broadcast will eventually succeed, i.e., honest parties’ messages will eventually be successfully broadcast. Such ideal functionality can be instantiated using only secret sharing against semi-malicious coalitions.

**Our Protocol Description.** We now instantiate  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  using the masked profile and anonymous broadcast with identifiable abort. The protocol proceeds in three phases and is parameterized with two parameters  $t_0$  and  $t_1$ . We will explain how to choose the parameters afterwards.

Each party  $i$  samples a virtual identifier  $\text{vID}_i \xleftarrow{\$} \{0, 1\}^\lambda$  and a key pair  $(\text{pk}_i, \text{sk}_i)$  that will be used solely for encryption. It obtains  $\text{auth}_i$  on  $(\text{vID}_i, \text{pref}_i)$  from  $\text{agg-auth}$ . Here, the authority’s signature binds the virtual identity to the certified preference.

1. *Commit phase:* Each party anonymously broadcasts a tuple  $(\text{vID}_i, \text{pref}_i, \text{pk}_i)$ . If fail, remove the misbehaved parties and retry with the same message. At the end of this phase, all honest parties agree on a *masked profile*

$$\mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}, \quad \text{pref}_i \in \{0, 1, \perp\},$$

This profile reveals only the aggregate information  $(n_0, n_1)$ , and hides the link between parties’ real identity and  $(\text{vID}_i, \text{pk}_i, \text{pref}_i)$  by anonymity of the anonymous broadcast.

2. *Sharing phase:* Let  $G_b = \{\text{vID}_i : (\text{vID}_i, \text{pk}_i, \text{pref}_i) \in \mathcal{P}, \text{pref}_i = b\}$  for  $b \in \{0, 1\}$ . Those who have  $\perp$ -preferences are treated as if they aborted at the beginning of the sharing phase. The protocol now runs *two* group-toss sub-protocol, one for  $G_0$  and one for  $G_1$ .

Each party  $\text{vID}_i \in G_b$  randomly picks a secret  $s_i \xleftarrow{\$} \{0, 1\}$  and generates a  $(t_b + 1)$ -out-of- $n_b$  Shamir sharing  $\{s_{i,j}\}$ . Note that although  $G_b$  is the group of valid  $b$ -supporters, parties generate  $n_b$  number of shares using the true aggregation from the certified authority, each party will only send shares to those in  $G_b$ . For every recipient  $\text{vID}_j \in G_b$ , encrypt the shares using the corresponding receivers' public keys and use anonymous broadcast with identifiable abort to broadcast all encrypted shares. Retry until success. Denote  $G'_b$  as those who successfully shared.

3. *Reconstruction phase:* Each  $\text{vID}_i \in G_b$  decrypts the incoming shares for itself, sums the decrypted shares from parties in  $G'_b$  and *anonymously broadcasts* its reconstruction shares to reconstruct  $c_b = \bigoplus_{\text{vID}_i \in G'_b} s_i$ . Kick out misbehaving parties and retry the anonymous broadcast until succeed. If the reconstruction of  $c_0$  fails, set  $c_0 = 0$ ; if the reconstruction of  $c_1$  fails, the protocol directly outputs 0. Otherwise if both reconstruction succeed, output  $c_0 \oplus c_1$ .

In a nutshell, the security follows from the anonymity that hides the mapping between all messages and the senders, thus making the virtual IDs unlinkable from real identities. Only the aggregation information is revealed. The protocol achieves game-theoretic fairness against any coalition satisfying the following three constraints

1. The coalition cannot control both coins.
2. If the coalition can control coin  $c_1$ , then the coalition cannot fail the reconstruction of coin  $c_0$ .
3. If the coalition can fail the reconstruction of  $c_1$ , then it must not prefer 0.

At a high level, Constraint 3 guarantees that a coalition is never incentivized to fail the reconstruction of  $c_1$ . Given that  $c_1$  is always successfully reconstructed, Constraints 1 and 2 guarantees that the output is a uniformly random coin. The parameters  $t_0$  and  $t_1$  are chosen to maximize the coalition tolerance.

**Upgrade to Malicious Security.** To achieve malicious security, we ask each party to prove the validity of their computation in each phase using bounded concurrent zero-knowledge proofs [Pas04]. Specifically,

- In the commit phase, each party  $i$  proves that they have a valid certification from the authority for their broadcast preference  $(\text{vID}_i, \text{pref}_i)$ .
- In the sharing and reconstruction phase, each party proves that the shares are correctly computed.

**Extension to  $m$ -Sided Coin-Toss.** We adopt the standard two-group reduction [TSW24]: treat 0 as the least-preferred outcome and aggregate all other outcomes into  $\text{oth}$ . The protocol is identical to the above, except that coins are produced by  $G_0$  and  $G_{\text{oth}}$  and combined as  $(c_0 + c_{\text{oth}}) \bmod m$ . Our ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  leaks only  $(h_0, h_{\text{oth}})$ ; the same anonymity, privacy, and fairness arguments apply, with thresholds chosen to satisfy the feasibility constraints in the main theorem.

### 3 Model and Preliminary

**Notations.** We assume each party is a probabilistic polynomial time (PPT) Turing machine. Throughout the paper, we may use  $\mathcal{A}$  to denote the corrupted coalition and adversary controlling the coalition interchangeably. We use  $\lambda$  to denote the security parameter. We call a function  $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}$  negligible in  $n$  if for every  $k \in \mathbb{N}$ , there exists  $n_0 \in \mathbb{N}$  s.t. for every  $n \geq n_0$  it holds that  $|\text{negl}(n)| \leq n^{-k}$ . In this work,  $[n]$  denotes the set of integers  $\{1, \dots, n\}$ .

We use  $U_m$  to denote the uniform distribution over  $\{0, \dots, m-1\}$ . Given two distribution ensembles  $\{X_n\}_n$  and  $\{Y_n\}_n$ , we say that the two ensembles are computationally indistinguishable, denoted as  $\{X_n\}_n \approx \{Y_n\}_n$ , iff there exists a negligible function  $\text{negl}(n)$  such that for any non-uniform PPT adversary  $\mathcal{A}$ , for any  $n$ , we have  $|\Pr[x \leftarrow X_n : \mathcal{A}(x) = 1] - \Pr[y \leftarrow Y_n : \mathcal{A}(y) = 1]| \leq \text{negl}(n)$ . We use  $\equiv$  to denote that two distributions are identical.

#### 3.1 Coin-Tossing Protocol With Certified Authority

In an  $n$ -party coin-tossing protocol with a certified authority **agg-auth**, each of the  $n$  parties first gets its preference certified by the certificate authority. Each party  $i$  has a preference  $\text{pref}_i$  over the outcomes in  $\{0, \dots, m-1\}$ . We use the set  $\mathcal{P} = \{(\text{ID}_1, \text{pref}_1), \dots, (\text{ID}_n, \text{pref}_n)\}$  to denote the *preference profile* of the parties, where  $\text{pref}_i$  denotes party  $i$ 's preference. Without loss of generality, we assume that 0 is the least preferred outcome, i.e.,  $n_0 \leq n_\ell$  for any  $\ell \in \{0, \dots, m-1\}$ , where  $n_\ell$  denotes the number of  $\ell$ -supporters in the preference profile. We also refer to those non-0-supporters as **oth**-supporters and use  $n_{\text{oth}} = n - n_0$  to denote the number of **oth**-supporters.

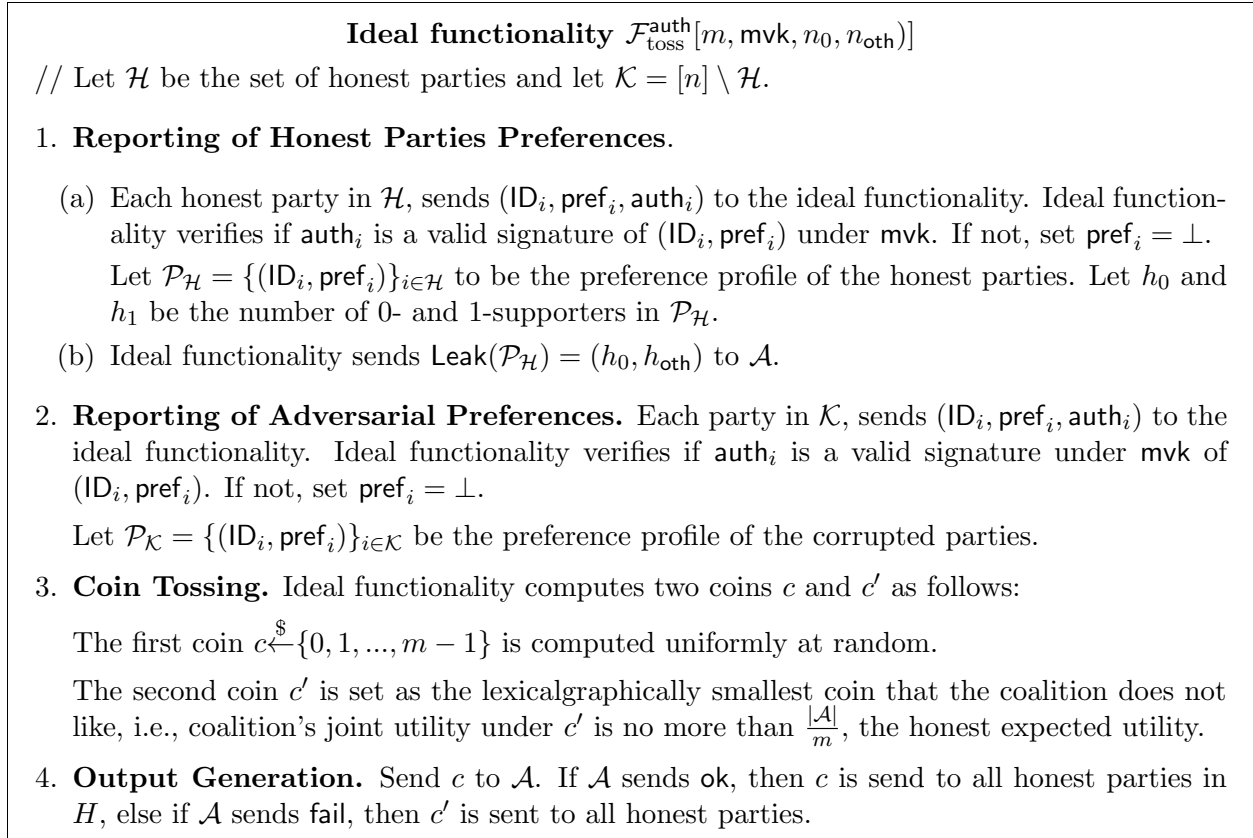
In this work, we consider a certified authority parametrized with a digital signature scheme  $\Pi_{\text{Sig}} = (\text{Gen}, \text{Sign}, \text{Ver})$ , with which the authority signs parties' inputs. The authority **agg-auth** samples a pair of  $(\text{msk}, \text{mvk}) \leftarrow \Pi_{\text{Sig}}.\text{Gen}(1^\lambda)$ , receives  $(\text{ID}_i, \text{pref}_i)$  from party  $i$  and sends  $\text{auth}_i = \Pi_{\text{Sig}}.\text{Sign}(\text{msk}, (\text{ID}_i, \text{pref}_i))$  to party  $i$ . It adds  $(\text{ID}_i, \text{pref}_i)$  into a set  $\mathcal{P}$  and ignores any messages with duplicate  $\text{ID}_i$  already in  $\mathcal{P}$ . At the end, it publishes  $(\text{mvk}, n_0, n_{\text{oth}})$  on some public bulletin board, where  $n_0$  and  $n_{\text{oth}}$  denote the number of 0-preferences and non-0 preferences in  $\mathcal{P}$ .

With certified preferences, the  $n$  parties communicate with each other through a pairwise private channel and a public broadcast channel. We assume that the coin-tossing protocol proceeds in rounds. All the communication channels are synchronous and authenticated, i.e., each message carries the identity of the true sender. A coalition  $\mathcal{A}$ , also referred to as the adversary for cryptographic analysis, can perform *rushing attacks*: It waits for honest parties' messages in any round  $r$  and then decides what messages to send to the honest parties in round  $r$ . If a party sends ill-formed messages, it is treated as abort.

We require *public verifiability* of the outcome: at the end of the protocol, the coin outcome is a deterministic, polynomial-time computable function based on the public aggregate information posted by the certified authority **agg-auth** and all messages posted in the broadcast channel.

**Utility and Strategies.** A party gets utility 1 if the outcome agrees with its certified preference and 0 otherwise. A coalition's utility is the sum of the utilities of all coalition members.

A *malicious* coalition can adopt a mix of the following strategies: 1) program its randomness based on its current view, or 2) abort the protocol in some round  $r$ , or 3.) send arbitrary messages in each round  $r$ .



**Figure 1:** Ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  for  $m$ -sided coin toss.

### 3.2 Coin-Tossing Ideal Functionality

We adopt a simulation-based definition to capture both game-theoretic fairness and preference privacy. Consider the following ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  parametrized with a parameter  $m$ , the number of sides of the coin to be tossed, the verification key for checking authentication from `agg-auth` and the published aggregate information  $n_0$  and  $n_{\text{oth}}$ .

Let  $\mathcal{P}_{\mathcal{H}}$  denote the preference profile of honest parties. In the case of the public preference profile considered in previous works [CGL<sup>+</sup>18c, WAS22, TSW24], where every party's preference is public, one may think of an “all” leakage where  $\text{Leak}^{\text{all}}(\mathcal{P}_{\mathcal{H}}) = \mathcal{P}_{\mathcal{H}}$ . For private preference, we consider the following leakage in this work:

$$\text{Leak}(\mathcal{P}_{\mathcal{H}}) = (h_0, h_{\text{oth}}),$$

where  $h_0$  and  $h_{\text{oth}}$  denotes the number of 0- and `oth`-supporters in  $\mathcal{P}_{\mathcal{H}}$ .

In Appendix B, we present the previous property-based definition of game-theoretic fairness defined in [CGL<sup>+</sup>18a] and show that the above ideal functionality implies property-based game-theoretic fairness.

**Definition 3.1** (Instantiation of  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ ). Given any  $m \geq 2$ . A protocol securely realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, n_0, n_{\text{oth}}]$  against any  $t$ -sized coalition, iff there exists an expected p.p.t. simulator `Sim` interacting with the above ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{Leak}]$ , such that for any non-uniform p.p.t. adversary  $\mathcal{A}$  control-

ling  $t$  or less number of parties, such that for any preference profile,

$$\{\text{Exec}_{\mathcal{A}}^{\Pi}(1^\lambda, \mathcal{P})\} \approx \{\text{Exec}_{\text{Sim}}^{\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{Leak}]}(1^\lambda, \mathcal{P})\},$$

where  $\text{Exec}_{\mathcal{A}}^{\Pi}(1^\lambda, \mathcal{P})$  outputs the honest parties' output as well as the  $\mathcal{A}$ 's view when interacting in coin toss protocol  $\Pi$ , and  $\text{Exec}_{\text{Sim}}^{\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{Leak}]}(1^\lambda, \mathcal{P})$  outputs the honest parties' output as well as the  $\text{Sim}$ 's output.

### 3.3 Building Blocks

For a better understanding of our construction, we present the syntax of the building block and informally state what kinds of security guarantees they provide. The security of these schemes will be formally defined in Appendix A.

**Shamir Secret Sharing.** A  $k$ -out-of- $n$  secret sharing scheme consists of two algorithms  $\text{Share}$  and  $\text{Reconstruct}$ :

- $(s_1, \dots, s_n) \leftarrow \text{Share}(s)$ : On input a secret  $s$ , outputs  $n$  shares  $s_1, \dots, s_n$  of the secret.
- $s \leftarrow \text{Reconstruct}(I, (s_i)_{i \in I})$ : The reconstruction algorithm takes  $(s_i)_{i \in I}$  and computes the secret only if  $|I| \geq k$ . Else, it outputs fail.

Roughly speaking, the security of Shamir's secret sharing guarantees that any  $k$  or more parties can reconstruct the secret, while any less than  $k$  parties get no information about the secret. Our construction will make use of the *linearity* of secret sharing: if every party gets a share of  $s$  and a share of  $s'$ , then they can reconstruct  $s + s'$  using the sum of their shares.

**Digital signature.** A digital signature scheme  $\Pi_{\text{Sig}}$  consist of three algorithms ( $\text{Gen}, \text{Sign}, \text{Ver}$ ):

- $\text{Gen}(1^\lambda)$ : takes in the security parameter  $\lambda$  and outputs a key pair  $(\text{vk}, \text{sk})$  of verification key  $\text{vk}$  and signing key  $\text{sk}$ .
- $\sigma \leftarrow \text{Sign}(m, \text{sk})$ : takes in a message  $m$  and the signing key  $\text{sk}$ , produces a signature  $\sigma$ .
- $b \leftarrow \text{Ver}(\text{vk}, \sigma, m)$ : takes in a message  $m$ , the signature  $\sigma$ , and the verification key  $\text{vk}$ , output 0 or 1 indicating reject or accept.

A secure digital signature is unforgeable: the adversary, without the knowledge of the signing key, cannot forge a valid signature on some new messages.

**Semantically-Secure Public-Key Encryption.** Public-key encryption consists of three algorithms:

- $\text{Gen}(1^\lambda)$  takes in the security parameter  $\lambda$  and outputs a key pair  $(\text{pk}, \text{sk})$ .
- $\text{ct} \leftarrow \text{Enc}(\text{pk}, \text{msg})$ : take in the public key and a message  $\text{msg}$ , outputs a ciphertext  $\text{ct}$ .
- $\text{msg} \leftarrow \text{Dec}(\text{sk}, \text{ct})$ : take in a ciphertext  $\text{ct}$  and the secret key  $\text{sk}$ , output  $\text{msg}$ .

A semantically secure public-key encryption hides the encrypted message from the adversary.

## 4 Instantiating $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, n_0, n_{\text{oth}}]$

In this section, we present an  $m$ -sided coin-toss protocol that realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, n_0, n_{\text{oth}}]$  (Figure 1).

At a high level, our protocol follows a structure similar to the public-preference coin-toss protocol against semi-malicious adversaries described in [TSW24]: Parties are partitioned into two groups: the set  $G_0$  of 0-supporters and the set  $G_{\text{oth}}$  of all other supporters, denoted as  $\text{oth}$ -supporter. Each  $b$ -supporter will run the Shamir secret sharing among its own group  $G_b$  to jointly decide a coin  $c_b$  for  $b \in \{0, \text{oth}\}$ . The final output is then defined based on  $c_0$  and  $c_{\text{oth}}$ .

To achieve privacy and simulation security, we make the following modifications:

- *Masked preference:* To hide parties' preferences, each party  $i$  will sample a “virtual ID”  $\text{vID}_i$  for themselves and use this virtual ID to communicate *anonymously*. The privacy of the preference is achieved by hiding the mapping between the real identities  $i$  of a party and its virtual ID  $\text{vID}_i$  throughout the protocol. Players first compute a “masked” preference profile  $\mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}$  and interact with other virtual IDs based on the masked preference profile. The  $\text{pk}_i$  here is only used for encrypting messages and is independent from a party's real identity.
- *Support undefined preference:* We make the protocol to work even when some corrupted parties' preferences are  $\perp$ , whereas previous protocol only works when every party's preference is publicly known in clear. This  $\perp$ -preference accounts for corrupted players' misbehavior when reporting their preferences.
- *Anonymous communication:* We make all the messages public by encrypting the private messages using the receiver's public key for encryption. This allows parties to send all messages using anonymous broadcast and therefore protects the mapping between the real identities  $i$  and the corresponding virtual ID  $\text{vID}_i$ .

## 4.1 Building Blocks

**Bounded Concurrent Zero-Knowledge Proofs.** To achieve malicious security, we require the parties to prove validity of their behaviors in each step. We will describe the protocol in a `IdealZK`-hybrid model, where parties have access to an ideal zero-knowledge proof functionality `IdealZK` given below. The functionality `IdealZK` either sends `success` to *everyone* indicating that the proof is correct, or the identity of the prover/verifier who leads to the failure of the proof. Formally,

### **Ideal Zero-knowledge Functionality `IdealZK[x, $\mathcal{L}$ , $i$ , $j$ ]`**

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

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

In an  $n$ -party `IdealZK`-hybrid protocol, parties invoke `IdealZK[x,  $\mathcal{L}$ ,  $i$ ,  $j$ ]` between every pair of prover and verifier for some NP language  $\mathcal{L}$ . The result of the proof, either succeed, or who leads to the failure of the proof, is broadcast to every party. Given an  $n$ -party `IdealZK`-hybrid protocol, `IdealZK` can be instantiated using the elegant techniques suggested by Pass [Pas04].

**Theorem 4.1.** (*Constant-round, bounded concurrent secure computation [Pas04]*). *Assume the existence of enhanced trapdoor permutations and collision-resistant hash functions. Then, given an  $n$ -party `IdealZK`-hybrid protocol  $\Pi^*$ , in which the number of concurrent calls of `IdealZK` is upper bounded by a-priori known bound  $\text{poly}(\lambda)$ , there exists a real-world protocol  $\Pi$  such that the following hold:*




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

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

In the above, the notation  $\text{Exec}_{\mathcal{A}}^{\Pi}$  (or  $\text{Exec}_{\mathcal{A}^*}^{\Pi^*}$ ) outputs each honest players' outputs as well as the corrupt players' (arbitrary) outputs. Moreover, The round complexity of  $\Pi$  is at most a constant factor worse than that of  $\Pi^*$ .

This real-world protocol is fulfilled by replacing the `IdealZK` instance with the bounded concurrent zero-knowledge proofs. All the zero-knowledge proof messages are posted to the broadcast channel.

**Anonymous Broadcast Channel with Identifiable Abort.** Throughout the protocol, we need to hide the mapping between virtual IDs and the real identity of parties via anonymous broadcast. However, ideal anonymous broadcast is impossible against corrupted majority. Therefore, parties interact with the following  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$  ideal functionality to anonymously broadcast a message. The ideal functionality  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$  receives messages from the senders and broadcasts a multiset containing all messages it receives, revealing no information about the sender's identity. If the broadcast fails, parties kick out the misbehaving parties and retry.



**Figure 2:** Anonymous broadcast ideal functionality  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$

**Lemma 4.2** (Theorem 6.2 in [KMSW22]). *Assuming perfectly binding and computationally hiding commitments. there is a protocol that securely realizes  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$  in the `IdealZK`-hybrid model against  $|\mathcal{O}| - t$  corruptions.*

In our protocol, parties will need to invoke `IdealZK` $[*, *, \text{vID}_i, \text{vID}_j]$  using provers' and verifiers' virtual identities  $\text{vID}_i$  and  $\text{vID}_j$ . Therefore, when instantiating the protocol, all the zero-knowledge proof messages will be routed using  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ , which itself can be instantiated using `IdealZK` $[*, *, i, j]$  over real identities for provers  $i$  and verifiers  $j$ .

## 4.2 Protocol Description

At the beginning of the protocol, each party  $i$  holds its preference  $\text{pref}_i$ , its authentication  $\text{auth}_i$  and the aggregate information  $(\text{mvk}, n_0, n_{\text{oth}})$  from the certificate authority. Recall that we assume 0 is the least preferred outcome. Our protocol contains three phases:

1. *Commit phase:* parties randomly samples a virtual ID  $\text{vID}$  and commit to their authenticated preference. At the end of the commit phase, parties agree on a masked preference profile  $\mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}$  where each  $\text{pref}_i \in \{0, \text{oth}, \perp\}$  and all the virtual IDs are unique.

2. *Sharing phase:* Given the masked preference profile, parties with non- $\perp$  preferences are partitioned into the group of 0-supporters  $G_0$  and the group of oth-supporters  $G_{\text{oth}}$ . Each group  $G_b$  for  $b \in \{0, \text{oth}\}$  separately run a sub-protocol to jointly determine a coin  $c_b$ .
3. *Reconstruction phase:* Each group  $G_b$  reconstruct the coin  $c_b$  and the reconstruction outcome determines the final output.

#### 4.2.1 Group Toss Sub-Protocol.

We first describe the sub-protocol for each group  $G_b$  for  $b \in \{0, \text{oth}\}$ , i.e., how each party acts during the sharing and reconstruction phase, given a masked preference profile  $\mathcal{P}$ . Recall that at the end of the commit phase, parties agree on a masked preference profile  $\mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}$ , where  $\text{pref}_i \in \{0, \text{oth}, \perp\}$ . Here, we do not distinguish preferences  $1, \dots, m-1$ , they are all denoted by oth.

Let  $G_0$  be the set of virtual IDs whose preference is 0 and  $G_{\text{oth}}$  be the set of virtual IDs whose preference is oth. Formally,

$$G_0 = \{\text{vID}_i : (\text{vID}_i, *, 0) \in \mathcal{P}\}, \quad G_{\text{oth}} = \{\text{vID}_i : (\text{vID}_i, *, \text{oth}) \in \mathcal{P}\}, \quad (1)$$

where  $*$  stands for wildcards.

**NP Languages.** We specify the following NP languages that will be used in our protocol. Specifically, parties invoke  $\text{IdealZK}$  for  $\mathcal{L}_S$  in the sharing phase, and  $\text{IdealZK}$  for  $\mathcal{L}_R$  in the reconstruction phase to prove correctness of computation.

<p>Language <math>\mathcal{L}_S</math></p> <p>Given a statement <math>x_i = (\text{Out}, \{(\text{vID}_j, \tilde{s}_{i,j})\}_{\text{vID}_j \in G_b}, \mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}, G_b, t_b)</math>, we say <math>x_i \in \mathcal{L}_S</math> if there exists a witness <math>w_i = (s_i, \{s_{i,j}\}_{j \in [n_b]}, \text{sk}_i)</math> such that</p> <ul style="list-style-type: none"> <li>• <math>(\text{vID}_i, \{(\text{vID}_j, \tilde{s}_{i,j})\}_{\text{vID}_j \in G_b})</math> is an entry in <math>\text{Out}</math>.</li> <li>• <math>(\text{pk}_i, \text{sk}_i)</math> is a valid key pair, <math>(\text{vID}_i, \text{pk}_i, \text{pref}_i) \in \mathcal{P}</math>, and <math>\text{vID}_i \in G_b</math>.</li> <li>• For each <math>\text{vID}_j \in G_b</math>, the ciphertext <math>\tilde{s}_{i,j}</math> is correctly computed using <math>\text{pk}_j</math>.</li> <li>• <math>\{s_{i,j}\}_{j \in [n_b]}</math> forms a valid <math>(t_b + 1)</math>-out-of-<math>n_b</math> secret sharing of <math>s_i</math>.</li> </ul>
<p>Language <math>\mathcal{L}_R</math></p> <p>Given a statement <math>x'_i = (\text{Out}', \mathcal{P}, G_b, \mathcal{I}_b^S, v_i, \text{Out}_s)</math>, where <math>\mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}</math> and <math>\text{Out}_s = \{(\text{vID}_i, \tilde{s}_{j,i})\}_{\text{vID}_i \in G_b, \text{vID}_j \in G_b \setminus \mathcal{I}_b^S}</math>, we say <math>x'_i \in \mathcal{L}_R</math> if there exists <math>w'_i = (\{s_{j,i}\}_{\text{vID}_j \in G_b \setminus \mathcal{I}_b^S}, \text{sk}_i)</math> such that</p> <ul style="list-style-type: none"> <li>• <math>(\text{vID}_i, v_i)</math> is an entry in <math>\text{Out}'</math>.</li> <li>• <math>(\text{pk}_i, \text{sk}_i)</math> is a valid key pair, <math>(\text{vID}_i, \text{pk}_i, \text{pref}_i) \in \mathcal{P}</math>, and <math>\text{vID}_i \in G_b</math>.</li> <li>• For each <math>\text{vID}_j \in G_b \setminus \mathcal{I}_b^S</math>, the <math>s_{j,i} = \text{Dec}(\text{sk}_i, \tilde{s}_{j,i})</math> is the correct decryption of <math>\tilde{s}_{j,i}</math>, and <math>(\text{vID}_i, \tilde{s}_{j,i})</math> is in <math>x_j</math>.</li> <li>• <math>v_i = \sum_{j \in G_b \setminus \mathcal{I}_b^S} s_{j,i}</math>.</li> </ul>

Throughout the protocol, we use two sets  $\mathcal{I}_b^S$  to keep track of the parties who misbehaved in the committing and sharing phases and  $\mathcal{I}_b^R$  for those misbehaved in the reconstruction phase. The sub-protocols are parametrized with parameters  $t_0$  and  $t_{\text{oth}}$  to be specified later.

<p><b>Sub-protocol <math>\text{GroupToss}^b[m, t_b, \mathcal{P}]</math></b></p> <p><i>Ingredients:</i> A public-key encryption scheme <math>\Pi_{\text{Enc}} = (\text{Gen}, \text{Enc}, \text{Dec})</math>.</p>
---

*Inputs:* Each party  $i$  holds its preference  $\text{pref}_i$ , authentication  $\text{auth}_i$ , and the aggregate information  $(\text{mvk}, n_0, n_{\text{oth}})$  from the certificate authority.

*Parameters:* The masked profile  $\mathcal{P} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{i \in [n]}$  where  $\text{pref}_i \in \{0, \text{oth}, \perp\}$  and  $t_b$  is a threshold parameter. Every party initializes  $\mathcal{I}_b^R$  and  $\mathcal{I}_b^S$  as empty sets at the beginning of the protocol.

**Protocol:** If  $|G_b| \leq t_b$ , every party outputs fail. Otherwise, proceeds as follows.

Sharing phase:

- Every  $\text{vID}_i \in G_b$  randomly chooses a secret  $s_i \xleftarrow{\$} \{0, \dots, m-1\}$  and computes the shares  $\{s_{i,j}\}_{j \in [n_b]}$  using  $(t_b + 1)$ -out-of- $n_b$  secret sharing. Without loss of generality, let the first  $|G_b|$  shares be  $\{s_{i,j}\}_{\text{vID}_j \in G_b}$ . Let  $\tilde{s}_{i,j} = \text{Enc}(\text{pk}_j, s_{i,j})$  for  $\text{vID}_j \in G_b$ .
- Initialize  $\mathcal{O} = [n]$ . Repeat the following until succeed:  
 Each  $\text{vID}_i \in G_b$  sets  $M_i = (\text{vID}_i, \{(\text{vID}_j, \tilde{s}_{i,j})\}_{\text{vID}_j \in G_b})$  and all other parties  $\text{vID}_j \notin G_b$  sets  $M_j = \mathbf{0}$  as dummy message. Every party  $\text{vID}_i$  sends  $M_i$  to  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ . If fails, receive a set  $\mathcal{D}$  from  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ , set  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.  
 Let  $\text{Out}$  be the set of non-dummy messages successfully broadcast.
- Each party  $\text{vID}_i$  invokes  $\text{IdealZK}[x_i, \mathcal{L}_S, \text{vID}_i, \text{vID}_j]$  for each  $\text{vID}_j$  in  $\mathcal{P}$ , with statement  $x_i = (\text{Out}, \{(\text{vID}_j, \tilde{s}_{i,j})\}_{\text{vID}_j \in G_b}, \mathcal{P}, G_b, t_b)$  and the witness  $w_i = (s_i, \{s_{i,j}\}_{j \in [n_b]}, \text{sk}_i)$  to prove that  $x_i \in \mathcal{L}_S$ .  
 For any  $\text{vID}_i \in G_b$ , if there exists a  $\text{vID}_j$  such that  $\text{IdealZK}[x_i, \mathcal{L}_S, \text{vID}_i, \text{vID}_j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}_b^S$ .  
 Let  $\text{Out}_s$  be the set of entries  $(\text{vID}_i, \{(\text{vID}_j, \tilde{s}_{i,j})\}_{\text{vID}_j \in G_b})$  for  $\text{vID}_i \in G_b \setminus \mathcal{I}_b^S$  successfully broadcast with valid proofs in  $\text{Out}$ .

Reconstruction phase:

- Parse each entry in  $\text{Out}_s$  as  $(\text{vID}_j, \{(\tilde{s}_{i,j})\}_{\text{vID}_j \in G_b})$ . Each  $\text{vID}_i \in G_b$  computes  $s_{j,i} = \text{Dec}(\text{sk}_i, \tilde{s}_{j,i})$  and the sum of decrypted shares:  $v_i = \sum_{j \in G_b \setminus \mathcal{I}_b^S} s_{j,i}$ .
- Reset  $\mathcal{O} = [n]$ .<sup>a</sup> Parties repeat the following until succeed:  
 Every  $\text{vID}_i \in G_b$  sets  $M'_i = (\text{vID}_i, v_i)$  and all other parties  $\text{vID}_j \notin G_b$  sets  $M'_j = \mathbf{0}$  as dummy message. Every party  $\text{vID}_i$  sends  $M'_i$  to  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ . If fails, receive a set  $\mathcal{D}$  from  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ , set  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.  
 Let  $\text{Out}'$  be the set of non-dummy messages successfully broadcast.
- Each party  $\text{vID}_i$  invokes  $\text{IdealZK}[x_i, \mathcal{L}_R, \text{vID}_i, \text{vID}_j]$  for each  $\text{vID}_j$  in  $\mathcal{P}$ , with statement  $x'_i = (\text{Out}', \mathcal{P}, G_b, \mathcal{I}_b^S, v_i, \text{Out}_s)$  and witness  $w'_i = (\{s_{j,i}\}_{\text{vID}_j \in G_b \setminus \mathcal{I}_b^S}, \text{sk}_i)$  to prove that  $x'_i \in \mathcal{L}_R$ .  
 For any  $\text{vID}_i \in G_b$ , if there exists a  $\text{vID}_j$  such that  $\text{IdealZK}[x_i, \mathcal{L}_R, \text{vID}_i, \text{vID}_j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}_b^R$ .
- If  $|\mathcal{I}_b^R| > n_b - t_b$ , everyone outputs fail. Otherwise, parse each entry in  $\text{Out}'$  as  $(\text{vID}_i, v_i)$  and use the shares  $\{v_i\}_{i \in G_b \setminus \mathcal{I}_b^R}$  to reconstruct the secret  $c_b := \sum_{i \in G_b \setminus \mathcal{I}_b^R} s_i$ .

<sup>a</sup>Parties can also continue with the  $\mathcal{O}$  stored from the sharing phase. We reset here for simplicity.

While each  $b$ -supporter only shares its secret with other  $b$ -supporters, it proves the correctness of

its behavior to all parties, including non- $b$ -supporters, by broadcasting the proof. Note that in the above protocol, all messages are broadcast through  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$  (Figure 2). Based on the description of the ideal functionality, the sets  $\mathcal{I}_b^S$  and  $\mathcal{I}_b^R$  can be uniquely determined from broadcast messages. These two sets are also publicly verifiable, and all honest parties hold the same  $\mathcal{I}_b^S$  and  $\mathcal{I}_b^R$ .

#### 4.2.2 Full Protocol Description

**NP Language:** We specify the following NP language  $\mathcal{L}_{\text{com}}$  used in the commit phase: Given a statement  $\tilde{x}_i = (\widetilde{\text{Out}}, \text{vID}_i, \text{pk}_i, \text{pref}_i, \text{mvk})$  where  $\text{pref}_i \in \{0, \text{oth}\}$ , we say  $\tilde{x}_i \in \mathcal{L}_{\text{com}}$  if  $(\text{vID}_i, \text{pk}_i, \text{pref}_i) \in \widetilde{\text{Out}}$  and

- For  $\text{pref}_i = 0$ , there exists a witness  $(\text{auth}_i, \text{sk}_i)$  such that  $\text{auth}_i$  is a valid signature of message  $(\text{vID}_i, 0)$  under verification key  $\text{mvk}$ , and  $\text{sk}_i$  is the valid secret key of  $\text{pk}_i$ ;
- For  $\text{pref}_i = \text{oth}$ , there exists a witness  $(\text{auth}_i, \text{sk}_i)$  such that  $\text{auth}_i$  is a valid signature of message  $(\text{vID}_i, \ell)$  for  $\ell \in \{1, \dots, m-1\}$  under verification key  $\text{mvk}$ , and  $\text{sk}_i$  is the valid secret key of  $\text{pk}_i$ ;

#### Protocol $\Pi_{n,m}[t_0, t_{\text{oth}}]$

**Authentication:** Each party  $i$  samples  $\text{vID}_i \xleftarrow{\$} \{0, 1\}^\lambda$  and sends  $(\text{vID}_i, \widetilde{\text{pref}}_i)$  to  $\text{agg-auth}$  and gets  $\text{auth}_i$ , where party  $i$ 's preference  $\widetilde{\text{pref}}_i \in \{0, \dots, m-1\}$ .

**Inputs:** If  $\widetilde{\text{pref}}_i = 0$ , set  $\text{pref}_i = 0$ ; otherwise, set  $\text{pref}_i = \text{oth}$ . Each party holds  $(\text{vID}_i, \text{pref}_i, \text{auth}_i)$  and aggregate information  $(\text{mvk}, n_0, n_{\text{oth}})$  from  $\text{agg-auth}$ .

#### Commit phase:

1. Each party  $i$  samples a pair of  $(\text{pk}_i, \text{sk}_i) \leftarrow \Pi_{\text{Enc.Gen}}(1^\lambda)$ .
2. Initialize  $\mathcal{O} = [n]$ . Repeat the following until the broadcast succeed:
 

Every party sends  $(\text{vID}_i, \text{pk}_i, \text{pref}_i)$  to  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ . If fails, receive a set  $\mathcal{D}$  from  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$ , set  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.

Let  $\widetilde{\text{Out}}$  be the set of successfully broadcast messages.
3. Each party  $i$  invokes  $\text{IdealZK}[\tilde{x}_i, \mathcal{L}_{\text{com}}, \text{vID}_i, j]$  for each  $j$  in  $[n]$ , with  $\tilde{x}_i = (\widetilde{\text{Out}}, \text{vID}_i, \text{pref}_i, \text{pk}_i, \text{mvk})$  and witness  $\tilde{w}_i = (\text{auth}_i, \text{sk}_i)$ .
 

For any  $\text{vID}_i$  appeared in  $\widetilde{\text{Out}}$ , if there exists a  $j$  such that  $\text{IdealZK}[\tilde{x}_i, \mathcal{L}_{\text{com}}, \text{vID}_i, j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}^C$ .

4. Let  $\mathcal{P}_O$  be the preferences with a valid proof in  $\widetilde{\text{Out}}$ :

$$\mathcal{P}_O = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i) \in \widetilde{\text{Out}} : \text{vID}_i \notin \mathcal{I}^C\}.$$

5. Let  $n' = |\mathcal{P}_O|$ . Let  $\mathcal{P}_A = \{\text{vID}'_k, \perp, \perp\}_{k \in [n-n']}$ , where  $\text{vID}'_k$  is the  $(\lambda+1)$ -bits presentation of  $k$ . Let  $\mathcal{P} = \mathcal{P}_O \cup \mathcal{P}_A$ , and  $G_0$  and  $G_{\text{oth}}$  as defined in Equation (1).

#### Sharing phase:

1. *Share-0:* Group  $G_0$  runs the sharing phase of  $\text{GroupToss}^0[m, t_0, \mathcal{P}]$ .
2. *Share-oth:* Group  $G_{\text{oth}}$  runs the sharing phase of  $\text{GroupToss}^{\text{oth}}[m, t_{\text{oth}}, \mathcal{P}]$ .

**Reconstruction phase:**

1. *Recon-0*: Group  $G_0$  runs the reconstruction phase of  $\text{GroupToss}^0[m, t_0, \mathcal{P}]$ . If the sub-protocol outputs fail, set  $c_0 = 0$ .
2. *Share-oth*: Group  $G_{\text{oth}}$  runs the reconstruction phase of  $\text{GroupToss}^{\text{oth}}[m, t_{\text{oth}}, \mathcal{P}]$ . If the sub-protocol outputs fail, the protocol outputs 0. Otherwise, output  $(c_0 + c_{\text{oth}}) \bmod m$ .

Below we use  $k_0$  and  $k_{\text{oth}}$  to denote the number of corrupted 0- and oth-supporters, respectively.

**Constraints 4.3.** *The threshold  $k_0$  and  $k_{\text{oth}}$ , and the number of corrupted parties  $k = k_0 + k_{\text{oth}}$  the protocol can tolerate is chosen such that the following three constraints are satisfied.*

- (C1) *The corrupted coalition cannot control both coins:  $k_0 + k_{\text{oth}} \leq t_0 + t_{\text{oth}} + 1$ .*
- (C2) *If the coalition can control coin  $c_{\text{oth}}$ , then the coalition cannot fail the reconstruction of coin  $c_0$ . Equivalently, if  $k_{\text{oth}} \geq t_{\text{oth}} + 1$ , then  $k_0 < n_0 - t_0$ .*
- (C3) *If the coalition can fail the reconstruction of  $c_{\text{oth}}$ , then it must not prefer 0. That is, if  $k_{\text{oth}} \geq n_{\text{oth}} - t_{\text{oth}}$ , then  $k_0 \leq \frac{k_0 + k_{\text{oth}}}{m}$ .*

**Theorem 4.4.** *If the encryption scheme  $\Pi_{\text{Enc}}$  is semantically secure, then the above  $\Pi_{n,m}[t_0, t_{\text{oth}}]$  securely realizes  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, n_0, n_{\text{oth}}]$  in the  $(\mathcal{F}_{\text{anon}}^{\mathcal{O},t}, \text{IdealZK})$ -hybrid world as long as the thresholds  $t_0, t_{\text{oth}}$ , and the number of corrupted parties  $k = k_0 + k_{\text{oth}}$  satisfy Constraints 4.3.*

*Proof.* We will show that for any non-uniform PPT  $\mathcal{A}$  that controls  $k_0$  number of 0-supporter and  $k_{\text{oth}}$  number of oth-supporter interacting with  $\Pi_{\text{toss}}$ , where  $k_0$  and  $k_{\text{oth}}$  satisfy Constraints 4.3, there exists an adversary  $\text{Sim}$  interacting with  $\mathcal{F}_{\text{toss}}^{\text{auth}}[m, \text{mvk}, n_0, n_{\text{oth}}]$  such that  $\mathcal{A}$ 's view in an execution with  $\Pi_{\text{toss}}$  is computationally indistinguishable from its view simulated by  $\text{Sim}$ .

**Commit phase:** Simulator receives  $\text{Leak}(\mathcal{P}_{\mathcal{H}}) = (h_0, h_{\text{oth}})$  from  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ . Emulate honest parties  $i \in \mathcal{H}$  as follows:

- Sample  $\text{vID}_i \xleftarrow{\$} \{0, 1\}^\lambda$  and key pair  $(\text{pk}_i, \text{sk}_i)$  for each honest party. If there exist duplicate virtual IDs for honest parties, the simulator aborts.

Let  $\mathcal{H}$  be the set of honest virtual IDs. Set  $\text{pref}_i = 0$  for the first  $h_0$  number of honest parties and  $\text{pref}_i = \text{oth}$  for the remaining honest parties. Let  $\mathcal{P}_{\mathcal{H}} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}_{\text{vID}_i \in \mathcal{H}}$ .

- Initialize  $\mathcal{O}$  as  $[n]$  and emulate  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$  until it succeeds as follows:

Receive  $\mathcal{P}_{\mathcal{K}} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i)\}$  from  $\mathcal{A}$ . Send a multiset  $\mathcal{P}_{\mathcal{H}} \cup \mathcal{P}_{\mathcal{K}}$  to  $\mathcal{A}$ . If received a set  $\mathcal{D}$  of size at least  $t$ , update  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.

Let  $\widetilde{\text{Out}}$  be the set of messages that are successfully broadcast.

- Emulate the  $\text{IdealZK}[* , \mathcal{L}_{\text{com}}, *, *]$  as follows:
  - For honest prover  $\text{vID}_i$  and honest verifier  $i'$ , send  $(\text{success}, \text{vID}_i, i')$  to  $\mathcal{A}$ ;
  - For honest prover  $\text{vID}_i$  and corrupted verifier  $j$ : If received  $\perp$  from a corrupted verifier, send  $(\text{fail}, j)$  to  $\mathcal{A}$ ; otherwise, send  $(\text{success}, \text{vID}_i, j)$  to  $\mathcal{A}$ ;
  - For corrupt prover  $\text{vID}_j$  and honest verifier  $i$ , send ok to  $\text{IdealZK}$  for the honest verifier, and forward  $\perp$  or witness  $\tilde{w}_j$  received from  $\mathcal{A}$  to  $\text{IdealZK}$ . Send the output of  $\text{IdealZK}$  to  $\mathcal{A}$ .
  - For corrupted prover  $\text{vID}_j$  and corrupted verifier  $j'$ , receive a bit from  $\mathcal{A}$  and send the output of  $\text{IdealZK}$ .

For any  $\text{vID}_i$  appeared in  $\widetilde{\text{Out}}$ , if there exists a  $j$  such that  $\text{IdealZK}[\widetilde{x}_i, \mathcal{L}_{\text{com}}, \text{vID}_i, j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}^C$ .

Let  $\mathcal{K}$  be the set of corrupted virtual IDs in  $\mathcal{P}_{\mathcal{K}}$  but not in  $\mathcal{I}^C$ . At this moment, for any corrupted  $\text{vID}_j \in \mathcal{K}$ , the simulator receives a valid witness  $\widetilde{w}_j = (\text{auth}_j, \text{sk}_j)$ . If  $\mathcal{K} \cap \mathcal{H} \neq \emptyset$ , i.e., a corrupted virtual ID collides with an honest virtual ID, the simulator aborts.

- Let  $\mathcal{P}'_{\mathcal{K}}$  be the corrupted preferences with a valid proof in  $\widetilde{\text{Out}}$ :  $\mathcal{P}'_{\mathcal{K}} = \{(\text{vID}_i, \text{pk}_i, \text{pref}_i) \in \mathcal{P}_{\mathcal{K}} : \text{vID}_i \in \mathcal{K}\}$ . Let  $n' = |\mathcal{P}_{\mathcal{H}} \cup \mathcal{P}'_{\mathcal{K}}|$ . Let  $\mathcal{P}_A = \{\text{vID}'_k, \perp, \perp\}_{k \in [n-n']}$ , where  $\text{vID}'_k$  is the  $(\lambda + 1)$ -bits presentation of  $k$ . Let  $\mathcal{P} = \mathcal{P}_{\mathcal{H}} \cup \mathcal{P}'_{\mathcal{K}} \cup \mathcal{P}_A$ , and  $G_0$  and  $G_{\text{oth}}$  as defined in Equation (1).
- For any  $\text{vID}_j \in \mathcal{K}$ , send  $(\text{vID}_j, \text{pk}_j, \text{pref}_j, \text{auth}_j)$  to  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ ; for other corrupted parties, send  $(\text{vID}'_j, \perp, \perp, \perp)$  to  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ .

The simulator gets a coin  $c$  from the ideal functionality. At this point, the simulator has a preference profile  $\mathcal{P}$  of size  $n$  where each  $\text{pref}_i \in \{0, \text{oth}, \perp\}$ . For all corrupted parties with non-bot preferences, the simulator gets the corresponding secret key  $\text{sk}_j$  from the extraction algorithm.

Let  $\mathcal{H}_0$  be the set of first  $h_0$  number of honest parties whose preferences are 0 and  $\mathcal{H}_{\text{oth}}$  be the set of remaining honest parties whose preferences are  $\text{oth}$ . Let  $\mathcal{K}_b$  to denote the set of honest and corrupted  $b$ -supporters for  $b \in \{0, \text{oth}\}$  whose preference is  $b$  in  $\mathcal{P}$ , respectively. At a high-level, those who have  $\perp$ -preferences are treated as if they aborted at the beginning of the sharing phase.

**Sharing phase:** The simulator computes  $k_b = n_b - h_b$  for  $b \in \{0, \text{oth}\}$ , where  $h_b$  is the number of honest  $b$ -supporters in  $\mathcal{P}_{\mathcal{H}}$ . Recall that  $\mathcal{I}_b^S$  and  $\mathcal{I}_b^R$  records the misbehaved  $b$ -supporters in the sharing phase and the reconstruction phase, respectively, for  $b \in \{0, \text{oth}\}$ .

**Share- $b$ :** Emulate honest parties to simulate execution of the sharing phase in sub-protocol  $\text{GroupToss}^b[m, t_b, \mathcal{P}]$ . If  $|G_b| \leq t_b$ , outputs fail for  $\text{GroupToss}^b[m, t_b, \mathcal{P}]$ .

If  $\text{GroupToss}^0[m, t_0, \mathcal{P}]$  outputs fail, set  $c_0 = 0$  and proceed to simulate Share- $\text{oth}$  step; if  $\text{GroupToss}^{\text{oth}}[m, t_{\text{oth}}, \mathcal{P}]$  outputs fail, send no to  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ . Outputs whatever the adversary  $\mathcal{A}$  outputs.

Otherwise, proceeds as follows:

- *Sharing:* If  $k_b < t_b + 1$ , for any  $j \in \mathcal{K}_b$ , uniformly randomly choose  $s_{i,j}$ ; for any  $\text{vID}_{i'} \in \mathcal{H}_b$ , set  $s_{i,i'} = 0$ . Otherwise, chooses a random coin  $s_i$  and compute the shares  $\{s_{i,j}\}_{j \in [n_b]}$  using  $(t_b + 1)$ -out-of- $n_b$  Shamir secret sharing.

For  $\text{vID}_j \in G_b$ , compute  $\widetilde{s}_{i,j} = \text{Enc}(\text{pk}_j, s_{i,j})$ .

- *Broadcast:* Emulate  $\mathcal{F}_{\text{anon}}^{\mathcal{O}, t}$  until it succeeds as follows:

For honest party  $\text{vID}_i$ , set  $M_i = (\text{vID}_i, \{(\text{vID}_j, \widetilde{s}_{i,j})\}_{\text{vID}_j \in G_b})$  for  $\text{vID}_i \in \mathcal{H}_b$  and  $M_i = \mathbf{0}$  for  $\text{vID}_i \in \mathcal{H} \setminus \mathcal{H}_b$ . Let multiset  $\mathcal{M}_{\mathcal{H}} = \{M_i\}_{\text{vID}_i \in \mathcal{H}}$ . Receive  $\mathcal{M}_{\mathcal{K}} = \{M_j\}$  from  $\mathcal{A}$  for corrupted parties. Send the multiset  $\mathcal{M}_{\mathcal{H}} \cup \mathcal{M}_{\mathcal{K}}$  to  $\mathcal{A}$ . If received a set  $\mathcal{D}$  from  $\mathcal{A}$ , update  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.

- *Validation:* Emulate the  $\text{IdealZK}[* , \mathcal{L}_S, *, *]$  as follows:

- For honest prover  $\text{vID}_i$  and honest verifier  $\text{vID}_{i'}$ , send  $(\text{success}, \text{vID}_i, \text{vID}_{i'})$  to  $\mathcal{A}$ ;
- For honest prover  $\text{vID}_i$  and corrupted verifier  $\text{vID}_j$ : If received  $\perp$  from a corrupted verifier, send  $(\text{fail}, \text{vID}_j)$  to  $\mathcal{A}$ ; otherwise, send  $(\text{success}, \text{vID}_i, \text{vID}_j)$  to  $\mathcal{A}$ ;
- For corrupt prover  $\text{vID}_j$  and honest verifier  $\text{vID}_i$ , send ok to  $\text{IdealZK}$  for the honest verifier, and forward  $\perp$  or witness  $\widetilde{w}_j$  received from  $\mathcal{A}$  to  $\text{IdealZK}$ . Send the output of  $\text{IdealZK}$  to  $\mathcal{A}$ .
- For corrupted prover  $\text{vID}_j$  and corrupted verifier  $\text{vID}_{j'}$ , receive a bit from  $\mathcal{A}$  and send the output of  $\text{IdealZK}$ .

For any  $\text{vID}_i \in G_b$ , if there exists a  $\text{vID}_j$  such that  $\text{IdealZK}[x_i, \mathcal{L}_S, \text{vID}_i, \text{vID}_j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}_b^S$ .

**Reconstruction phase:** At this point, for any  $\text{vID}_j \in \mathcal{K}_b \setminus \mathcal{I}_b^S$  for  $b \in \{0, \text{oth}\}$ , the simulator receives a witness from emulating  $\text{IdealZK}$ , including the secret  $s_j$  and shares for honest parties  $\{s_{j,i}\}_{i \in \mathcal{H}_b}$ .

**Recon-0:** Emulate honest parties for the reconstruction phase in sub-protocol  $\text{GroupToss}^0[m, t_0, \mathcal{P}]$ :

1. Summation: The simulator computes the reconstruction share for the honest parties  $\text{vID}_i \in \mathcal{H}_0$  as follows:

- If  $k_0 \geq t_0 + 1$ : simulator has the shares  $s_{j,i}$  for all  $j \in G_0 \setminus \mathcal{I}_0^S$  and  $i \in \mathcal{H}_0$  from the sharing phase. It computes  $v_i = \sum_{j \in G_0 \setminus \mathcal{I}_0^S} s_{j,i}$ .
- Otherwise if  $k_0 < t_0 + 1$  and  $k_{\text{oth}} \geq t_{\text{oth}} + 1$ : simulator has the secret  $s_j$  for all  $j \in (G_0 \cup G_{\text{oth}}) \setminus (\mathcal{I}^S \cup \mathcal{H}_0)$ . The simulator chooses  $\{s_i\}_{i \in \mathcal{H}_0}$  such that  $\sum_{i \in (G_0 \cup G_{\text{oth}}) \setminus \mathcal{I}^S} s_i = c \pmod m$ . For every  $i \in \mathcal{H}_0$ , compute a  $(t_0 + 1)$ -out-of- $n_0$  secret sharing  $\{y_{i,j}\}$  of  $s_i$  such that for any  $j \in \mathcal{K}_0$ , the share  $y_{i,j} = s_{i,j}$ . Set  $v_i = \sum_{j \in \mathcal{K}_0 \setminus \mathcal{I}_0^S} s_{j,i} + \sum_{j \in \mathcal{H}_0} y_{j,i}$ .
- Otherwise if  $k_0 < t_0 + 1$  and  $k_{\text{oth}} < t_{\text{oth}} + 1$ : Choose a uniformly random  $s_i$  for each  $i \in \mathcal{H}_0$  and compute a  $(t_0 + 1)$ -out-of- $n_0$  secret sharing  $\{y_{i,j}\}$  of each  $s_i$  such that for any  $j \in \mathcal{K}_0$ , the share  $y_{i,j} = s_{i,j}$ . Set  $v_i = \sum_{j \in \mathcal{K}_0 \setminus \mathcal{I}_0^S} s_{j,i} + \sum_{j \in \mathcal{H}_0} y_{j,i}$ .

2. Broadcast: Emulate  $\mathcal{F}_{\text{anon}}^{\mathcal{O}, t}$  until it succeeds as follows:

Set  $M'_i = (\text{vID}_i, v_i)$  for  $\text{vID}_i \in \mathcal{H}_0$  and  $M'_i = \mathbf{0}$  for  $\text{vID}_i \in \mathcal{H}_{\text{oth}}$ . Let  $\mathcal{M}'_{\mathcal{H}} = \{M'_i\}_{\text{vID}_i \in \mathcal{H}}$ . Receive  $\mathcal{M}'_{\mathcal{K}} = \{M'_j | 0\}$  from  $\mathcal{A}$  for corrupted parties. Send the multiset  $\mathcal{M}'_{\mathcal{K}} \cup \mathcal{M}'_{\mathcal{H}}$  to  $\mathcal{A}$ . If received a set  $\mathcal{D}$  from  $\mathcal{A}$ , update  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.

3. Validation: Emulate the  $\text{IdealZK}[* , \mathcal{L}_R, *, *]$  as follows:

- For honest prover  $\text{vID}_i$  and honest verifier  $\text{vID}_{i'}$ , send  $(\text{success}, \text{vID}_i, \text{vID}_{i'})$  to  $\mathcal{A}$ ;
- For honest prover  $\text{vID}_i$  and corrupted verifier  $\text{vID}_j$ : If received  $\perp$  from a corrupted verifier, send  $(\text{fail}, \text{vID}_j)$  to  $\mathcal{A}$ ; otherwise, send  $(\text{success}, \text{vID}_i, \text{vID}_j)$  to  $\mathcal{A}$ ;
- For corrupt prover  $\text{vID}_j$  and honest verifier  $\text{vID}_i$ , send ok to  $\text{IdealZK}$  for the honest verifier, and forward  $\perp$  or witness  $\tilde{w}_j$  received from  $\mathcal{A}$  to  $\text{IdealZK}$ . Send the output of  $\text{IdealZK}$  to  $\mathcal{A}$ .
- For corrupted prover  $\text{vID}_j$  and corrupted verifier  $\text{vID}_{j'}$ , receive a bit from  $\mathcal{A}$  and send the output of  $\text{IdealZK}$ .

For any  $\text{vID}_i \in G_0$ , if there exists a  $\text{vID}_j$  such that  $\text{IdealZK}[x_i, \mathcal{L}_R, \text{vID}_i, \text{vID}_j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}_0^R$ .

4. Reconstruction: Reconstruct  $c_0$  using  $\{v_i\}_{i \in G_0 \setminus \mathcal{I}_0^R}$ . If fails, set  $c_0 = 0$ ; otherwise, record the value of  $c_0$ .

**Recon-oth:** Emulate honest parties to simulate execution of the reconstruction phase in sub-protocol  $\text{GroupToss}^{\text{oth}}[m, t_{\text{oth}}, \mathcal{P}]$ :

1. Summation: The simulator computes the reconstruction share for the honest parties:

- If  $k_{\text{oth}} \geq t_{\text{oth}} + 1$ : then it must be that  $k_0 < t_0 + 1$  according to Constraints 4.3. The simulator Sim would have  $s_{j,i}$  for all  $j \in G_{\text{oth}} \setminus \mathcal{I}_{\text{oth}}^S$ , and it is guaranteed that  $\sum_{i \in (G_0 \cup G_{\text{oth}}) \setminus \mathcal{I}^S} s_i = c \pmod m$ . It computes  $v_i = \sum_{j \in G_{\text{oth}} \setminus \mathcal{I}_{\text{oth}}^S} s_{j,i}$ .



- Otherwise, choose  $s_i$  for each  $i \in \mathcal{H}_{\text{oth}}$  such that  $c_0 + (\sum_{i \in G_{\text{oth}} \setminus \mathcal{I}_{\text{oth}}^S} s_i) = c \pmod m$ . Compute a  $(t_{\text{oth}} + 1)$ -out-of- $n_{\text{oth}}$  secret sharing  $\{y_{i,j}\}$  of  $s_i$  such that for any  $j \in \mathcal{K}_{\text{oth}}$ , the share  $y_{i,j} = s_{i,j}$ . Set  $v_i = \sum_{j \in \mathcal{K}_{\text{oth}} \setminus \mathcal{I}_{\text{oth}}^S} s_{j,i} + \sum_{j \in \mathcal{H}_{\text{oth}}} y_{j,i}$ .

2. Broadcast: Emulate  $\mathcal{F}_{\text{anon}}^{\mathcal{O},t}$  until it succeeds as follows:

Set  $M'_i = (\text{vID}_i, v_i)$  for  $\text{vID}_i \in \mathcal{H}_0$  and  $M'_i = \mathbf{0}$  for  $\text{vID}_i \in \mathcal{H}_{\text{oth}}$ . Let  $\mathcal{M}'_{\mathcal{H}} = \{M'_i\}_{\text{vID}_i \in \mathcal{H}}$ . Receive  $\mathcal{M}'_{\mathcal{K}} = \{M'_j\}_0$  from  $\mathcal{A}$  for corrupted parties. Send the multiset  $\mathcal{M}'_{\mathcal{K}} \cup \mathcal{M}'_{\mathcal{H}}$  to  $\mathcal{A}$ . If received a set  $\mathcal{D}$  from  $\mathcal{A}$ , update  $\mathcal{O} = \mathcal{O} \setminus \mathcal{D}$  and retry.

3. Validation: Emulate the  $\text{IdealZK}[* , \mathcal{L}_R, *, *]$  as follows:

- For honest prover  $\text{vID}_i$  and honest verifier  $\text{vID}_{i'}$ , send  $(\text{success}, \text{vID}_i, \text{vID}_{i'})$  to  $\mathcal{A}$ ;
- For honest prover  $\text{vID}_i$  and corrupted verifier  $\text{vID}_j$ : If received  $\perp$  from a corrupted verifier, send  $(\text{fail}, \text{vID}_j)$  to  $\mathcal{A}$ ; otherwise, send  $(\text{success}, \text{vID}_i, \text{vID}_j)$  to  $\mathcal{A}$ ;
- For corrupt prover  $\text{vID}_j$  and honest verifier  $\text{vID}_i$ , send ok to  $\text{IdealZK}$  for the honest verifier, and forward  $\perp$  or witness  $\tilde{w}_j$  received from  $\mathcal{A}$  to  $\text{IdealZK}$ . Send the output of  $\text{IdealZK}$  to  $\mathcal{A}$ .
- For corrupted prover  $\text{vID}_j$  and corrupted verifier  $\text{vID}_{j'}$ , receive a bit from  $\mathcal{A}$  and send the output of  $\text{IdealZK}$ .

For any  $\text{vID}_i \in G_{\text{oth}}$ , if there exists a  $\text{vID}_j$  such that  $\text{IdealZK}[x_i, \mathcal{L}_R, \text{vID}_i, \text{vID}_j]$  outputs  $(\text{fail}, \text{vID}_i)$ , i.e., the prover fails to prove the statement, add  $\text{vID}_i$  to  $\mathcal{I}_{\text{oth}}^R$ .

4. Reconstruction: For those corrupted parties  $\text{vID}_j$  that failed to broadcast a valid proof, add  $\text{vID}_j$  to  $\mathcal{I}_{\text{oth}}^R$ . Reconstruct  $c_{\text{oth}}$  using  $\{v_i\}_{i \in G_{\text{oth}} \setminus \mathcal{I}_{\text{oth}}^R}$ . If reconstruction succeeds, send  $\text{yes}$  to  $\mathcal{F}_{\text{toSS}}^{\text{auth}}$ .

Otherwise, send  $\text{no}$  to  $\mathcal{F}_{\text{toSS}}^{\text{auth}}$ .

Output whatever  $\mathcal{A}$  outputs.

We first show that the probability that the simulator aborts is negligible.

**Claim 4.5.** *There exists a negligible function  $\text{negl}(\cdot)$ , such that the probability that the simulator  $S$  aborts is at most  $\text{negl}(\lambda)$ .*

*Proof.* The simulator aborts if either 1.) It samples duplicate virtual IDs for honest parties; 2.) If any corrupted virtual ID  $\text{vID}_j$  conflicts with an honest virtual ID, i.e., if  $\mathcal{K} \cap \mathcal{H} \neq \emptyset$ . We now compute the probability of these three two separately.

The probability that the simulator samples duplicate virtual ID is  $1 - \binom{2^\lambda}{h} \cdot h! / 2^{\lambda \cdot h}$ , which is negligible.

In the rest of the proof, we focus on proving that the second event happens with a negligible probability. Let  $\mathcal{K}_{\text{id}}$  be the set of corrupted virtual IDs signed by the certificate authority. Observe that  $\Pr[\mathcal{H} \cap \mathcal{K}_{\text{id}} \neq \emptyset]$  is negligible. We only need to show that conditioned on  $\mathcal{H} \cap \mathcal{K}_{\text{id}} = \emptyset$ , for any non-uniform PPT  $\mathcal{A}$ , the probability that  $\mathcal{K} \cap \mathcal{H} \neq \emptyset$  is negligible.

Suppose for the sake of contradiction that there exists a non-uniform  $\mathcal{A}$  with a non-negligible probability resulting in some  $\text{vID}^* \in \mathcal{K} \cap \mathcal{H}$ , conditioned on  $\mathcal{H} \cap \mathcal{K}_{\text{id}} = \emptyset$ . This implies that the adversary is able to produce a valid  $\text{auth}_i$  for  $(\text{vID}^*, *)$  for  $\text{vID}^* \notin \mathcal{K}_{\text{id}}$  to send to  $\text{IdealZK}$ , which breaks the existential unforgeability of signature scheme. Putting together,

$$\begin{aligned} & \Pr[\mathcal{H} \cap \mathcal{K} \neq \emptyset] \\ &= \Pr[\mathcal{H} \cap \mathcal{K} \neq \emptyset \mid \mathcal{H} \cap \mathcal{K}_{\text{id}} = \emptyset] \cdot \Pr[\mathcal{H} \cap \mathcal{K}_{\text{id}} = \emptyset] + \Pr[\mathcal{H} \cap \mathcal{K} \neq \emptyset \mid \mathcal{H} \cap \mathcal{K}_{\text{id}} \neq \emptyset] \cdot \Pr[\mathcal{H} \cap \mathcal{K}_{\text{id}} \neq \emptyset] \\ &= \text{negl}(\lambda). \end{aligned}$$

Therefore, the probability that the simulator aborts is negligible.  $\square$

We now show that the joint distribution of honest parties' output and the adversary's view is indistinguishable in the ideal execution when interacting with  $\text{Sim}$  (denoted as  $\text{Expt}_{\mathcal{A}}^{\text{Ideal}}$ ) is identical to the joint distribution of the output of honest parties and the view of the adversary in a real execution (denoted as  $\text{Expt}_{\mathcal{A}}^{\text{Real}}$ ), conditioned on the simulator not abort. Consider the following hybrids:

**Hybrid  $\text{Hyb}_1$ :** This is an execution of  $\Pi_{n,m}[t_0, t_{\text{oth}}]$ , where the simulator acts on behalf of all honest parties and interact with the adversary. Moreover, the simulator emulates all the  $\text{IdealZK}$  instances for the adversary as follows:

- For honest prover  $i$  and honest verifier  $i'$ , send  $(\text{success}, i, i')$  to  $\mathcal{A}$ ;
- For honest prover  $i$  and corrupted verifier  $j$ : If received  $\perp$  from a corrupted verifier, send  $(\text{fail}, j)$  to  $\mathcal{A}$ ; otherwise, send  $(\text{success}, i, j)$  to  $\mathcal{A}$ ;
- For corrupt prover  $j$  and honest verifier  $i$ , send  $\text{ok}$  to  $\text{IdealZK}$  for the honest verifier, and forward  $\perp$  or witness  $\tilde{w}_j$  received from  $\mathcal{A}$  to  $\text{IdealZK}$ . Send the output of  $\text{IdealZK}$  to  $\mathcal{A}$ .
- For corrupted prover  $j$  and corrupted verifier  $j'$ , receive a bit from  $\mathcal{A}$  and send the output of  $\text{IdealZK}$ .

We have  $\text{Expt}_{\mathcal{A}}^{\text{Real}} \equiv \text{Hyb}_1$  by definition.

**Hybrid  $\text{Hyb}_2$ :** The simulator behaves almost the same as in  $\text{Hyb}_1$ , except that now the simulator receives leakage  $\text{Leak} = (h_0, h_{\text{oth}})$  from the ideal functionality and randomly sample virtual IDs for each honest parties. We have  $\text{Hyb}_2 \equiv \text{Hyb}_1$ .

**Hybrid  $\text{Hyb}_3$ :** The simulator behaves almost the same as in  $\text{Hyb}_4$  except the following: In  $\text{Share-}b$  step, the simulator computes the shares for honest parties  $\text{vID}_i \in \mathcal{H}_b$  as in  $\text{Hyb}_4$ , but the encryptions  $\{\tilde{s}_{i,j}\}_{\text{vID}_j \in G_b}$  are computed differently. For  $b \in \{0, \text{oth}\}$ ,

- For honest  $\text{vID}_i \in \mathcal{H}_b$  and corrupted  $\text{vID}_j \in \mathcal{K}_b$ , let  $\tilde{s}_{i,j} = \text{Enc}(\text{pk}_j, s_{i,j})$ .
- For honest  $\text{vID}_i \in \mathcal{H}_b$  and  $\text{vID}_{i'} \in \mathcal{H}_b$ , let  $\tilde{s}_{i,j} = \text{Enc}(\text{pk}_j, 0)$ .

We have  $\text{Hyb}_3 \approx \text{Hyb}_2$  by semantic security of public-key encryption scheme.

Next, we show that  $\text{Hyb}_3 \equiv \text{Expt}_{\mathcal{A}}^{\text{Ideal}}$ , which concludes the proof. By security of Shamir's secret sharing, the adversary's view in both experiments are identical. We only need to show that for any fixed adversary's view, honest parties' outputs are also identical in both experiments.

By constraint (C3), if the reconstruction of  $c_{\text{oth}}$  fails, and  $\text{Sim}$  sends  $\text{no}$  to  $\mathcal{F}_{\text{toSS}}^{\text{auth}}$ . Then  $\mathcal{F}_{\text{toSS}}^{\text{auth}}$  uses  $k_0 = n_0 - h_0$  and  $k_{\text{oth}} = n_{\text{oth}} - h_{\text{oth}}$  to determine the lexicographically smallest outcome that the coalition does not like, which is 0. Therefore, if the reconstruction of  $c_{\text{oth}}$  fails, honest parties' output will be 0 in both experiments.

In the following, we will only focus on proving that if the reconstruction of  $c_{\text{oth}}$  succeeds, the honest parties' output in both worlds will be  $c$ , the coin generated by  $\mathcal{F}_{\text{toSS}}^{\text{auth}}$ . There are two cases:

- If  $t_{\text{oth}} \geq k_{\text{oth}} + 1$ , by constraints (C1) and (C2), it must be the case that  $t_0 < k_0 + 1$  and  $c_0$  must be successfully reconstructed. Therefore, by the description of  $\text{Recon-0}$  step,  $\{s_i\}_{i \in \mathcal{H}_0}$  are chosen such that  $(\sum_{i \in (G_{\text{oth}} \cup G_0) \setminus \mathcal{I}^S} s_i) \bmod m = c$ . Therefore, if both  $c_0$  and  $c_{\text{oth}}$  are successfully reconstructed, honest parties' output should be  $c$ .
- If  $t_{\text{oth}} < k_{\text{oth}} + 1$ , by the description of the  $\text{Recon-oth}$  step,  $\{s_i\}_{i \in \mathcal{H}_{\text{oth}}}$  are chosen such that  $(c_0 + \sum_{j \in G_{\text{oth}} \setminus \mathcal{I}_{\text{oth}}^S} s_j) \bmod m = c$ . If the reconstruction of  $c_{\text{oth}}$  succeeds, the honest parties will output  $c$  in both experiments.

Therefore, by hybrid argument,  $\text{Expt}_{\mathcal{A}}^{\text{Ideal}} \approx \text{Expt}_{\mathcal{A}}^{\text{Real}}$ .  $\square$

Previous work [TSW24] presented the optimal solutions of  $t_0, t_{\text{oth}}$  and  $k$  that satisfies Constraints 4.3. The optimal parameters are given in the table below. For  $m = 2$ , the coalition tolerance matches the optimal tolerance in [WAS22].

Case	$k_0$	$t_{\text{oth}}$	$t$
If $n_{\text{oth}} \geq \frac{4m-3}{2}n_0$	$\lfloor \frac{n_0}{2} \rfloor$	$n_{\text{oth}} - (m-1)n_0$	$\lceil \frac{n_0}{2} \rceil + n_{\text{oth}} - (m-1)n_0$
Otherwise	$\lfloor \frac{n_0}{2} \rfloor$	$\lfloor \frac{m}{2m-1}n_{\text{oth}} - \frac{m-1}{2(2m-1)}n_0 \rfloor$	$\lceil \frac{n_0}{2} \rceil + \lfloor \frac{m}{2m-1}n_{\text{oth}} - \frac{m-1}{2(2m-1)}n_0 \rfloor$

**Table 1:** Optimal parameter choice for  $m$ -sided coin-toss  $\Pi$  instantiating  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ . Here,  $n_b$  denotes the number of  $b$ -supporters for  $b \in \{0, \text{oth}\}$  receiving certification from the certified authority.

## References

- [ACH11] Gilad Asharov, Ran Canetti, and Carmit Hazay. Towards a game theoretic view of secure computation. In *Eurocrypt*, 2011. 1.3
- [ACJ17] Prabhanjan Ananth, Arka Rai Choudhuri, and Abhishek Jain. A new approach to round-optimal secure multiparty computation. Cryptology ePrint Archive, Paper 2017/402, 2017. 1.3
- [ADGH06] Ittai Abraham, Danny Dolev, Rica Gonen, and Joseph Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *PODC*, 2006. 1.3
- [ADMM16] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and undefinedukasz Mazurek. Secure multiparty computations on bitcoin. *Commun. ACM*, 59(4):76–84, March 2016. 1.3
- [AL11] Gilad Asharov and Yehuda Lindell. Utility dependence in correct and fair rational secret sharing. *Journal of Cryptology*, 24(1), 2011. 1.3
- [ATM<sup>+</sup>22] Lukas Aumayr, Sri AravindaKrishnan Thyagarajan, Giulio Malavolta, Pedro Moreno-Sanchez, and Matteo Maffei. Sleepy channels: Bi-directional payment channels without watchtowers. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS '22*, page 179–192, New York, NY, USA, 2022. Association for Computing Machinery. 1.3
- [BGJ<sup>+</sup>18] Saikrishna Badrinarayanan, Vipul Goyal, Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Promise zero knowledge and its applications to round optimal mpc. In *Annual International Cryptology Conference*, pages 459–487. Springer, 2018. 1.3
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 1–10, 1988. 1
- [BHP17] Zvika Brakerski, Shai Halevi, and Antigoni Polychroniadou. Four round secure computation without setup. In *Theory of Cryptography: 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I*, page 645–677, Berlin, Heidelberg, 2017. Springer-Verlag. 1.3
- [BK14] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *CRYPTO*, pages 421–439, 2014. 1.3
- [Blu81] Manuel Blum. Coin flipping by telephone. In *CRYPTO*, 1981. 1
- [BMR90] D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing, STOC '90*, page 503–513, New York, NY, USA, 1990. Association for Computing Machinery. 1.3

- [Can01] R. Canetti. Universally composable security: a new paradigm for cryptographic protocols. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145, 2001. [1.1.2](#)
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 11–19. ACM, 1988. [1](#)
- [CCWS21] Kai-Min Chung, T.-H. Hubert Chan, Ting Wen, and Elaine Shi. Game-theoretic fairness meets multi-party protocols: The case of leader election. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021*, pages 3–32, Cham, 2021. Springer International Publishing. [1](#), [1.3](#)
- [CGL<sup>+</sup>18a] Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *Theory of Cryptography: 16th International Conference, TCC 2018, Panaji, India, November 11–14, 2018, Proceedings, Part I*, page 563–596, Berlin, Heidelberg, 2018. Springer-Verlag. [1](#), [1.1.2](#), [3.2](#), [B](#)
- [CGL<sup>+</sup>18b] Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, 2018. [1.3](#)
- [CGL<sup>+</sup>18c] Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, 2018. [3.2](#), [B.1](#)
- [Cha88] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of cryptology*, 1(1):65–75, 1988. [2](#)
- [Cle86] Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 364–369, 1986. [1](#)
- [CMST22] Hao Chung, Elisaweta Masserova, Elaine Shi, and Sri AravindaKrishnan Thyagarajan. Rapidash: Foundations of side-contract-resilient fair exchange. Cryptology ePrint Archive, Paper 2022/1063, 2022. <https://eprint.iacr.org/2022/1063>. [1.3](#)
- [DEF18] Stefan Dziembowski, Lisa Ekekey, and Sebastian Faust. Fairswap: How to fairly exchange digital goods. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, page 967–984, New York, NY, USA, 2018. Association for Computing Machinery. [1.3](#)
- [DGPS22] Moumita Dutta, Chaya Ganesh, Sikhar Patranabis, and Nitin Singh. Compute, but verify: Efficient multiparty computation over authenticated inputs. Cryptology ePrint Archive, Paper 2022/1648, 2022. [1.1.1](#), [1.3](#)
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *ACM symposium on Theory of computing (STOC)*, 1987. [1](#), [1.1.2](#), [1.3](#)
- [Goy11] Vipul Goyal. Constant round non-malleable protocols using one way functions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 695–704, 2011. [1.3](#)

- [HHPV21] Shai Halevi, Carmit Hazay, Antigoni Polychroniadou, and Muthuramakrishnan Venkatasubramanian. Round-optimal secure multi-party computation. *Journal of Cryptology*, 34(3):19, 2021. [1.3](#)
- [HT04] Joseph Halpern and Vanessa Teague. Rational secret sharing and multiparty computation. In *STOC*, 2004. [1.3](#)
- [IKSS21] Yuval Ishai, Dakshita Khurana, Amit Sahai, and Akshayaram Srinivasan. On the round complexity of black-box secure mpc. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part II 41*, pages 214–243. Springer, 2021. [1.3](#)
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer – efficiently. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, pages 572–591, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. [1.3](#)
- [KB14] Ranjit Kumaresan and Iddo Bentov. How to use bitcoin to incentivize correct computations. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 30–41. ACM, 2014. [1.3](#)
- [KMS<sup>+</sup>16] Ahmed E. Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 839–858. IEEE Computer Society, 2016. [1.3](#)
- [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*, pages 409–438, Cham, 2022. Springer Nature Switzerland. [1](#), [1.3](#), [2](#), [4.2](#)
- [KN08] Gillat Kol and Moni Naor. Cryptography and game theory: Designing protocols for exchanging information. In *TCC*, 2008. [1.3](#)
- [KVV16] Ranjit Kumaresan, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Improvements to secure computation with penalties. In *ACM CCS*, 2016. [1.3](#)
- [Lin16] Yehuda Lindell. How to simulate it - a tutorial on the simulation proof technique. Cryptology ePrint Archive, Paper 2016/046, 2016. [1.1.2](#)
- [OPRV09] Shien Jin Ong, David C. Parkes, Alon Rosen, and Salil P. Vadhan. Fairness with an honest minority and a rational majority. In *TCC*, 2009. [1.3](#)
- [Pas04] Rafael Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *STOC*, 2004. [1.1.3](#), [1.3](#), [2](#), [4.1](#), [4.1](#)
- [RB89] Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 73–85. ACM, 1989. [1](#)

- [RCCG<sup>+</sup>20] Arka Rai Choudhuri, Michele Ciampi, Vipul Goyal, Abhishek Jain, and Rafail Ostrovsky. Round optimal secure multiparty computation from minimal assumptions. In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part II 18*, pages 291–319. Springer, 2020. [1.3](#)
- [Sha79] Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979. [2](#)
- [TSW24] Sri AravindaKrishnan Thyagarajan, Pratik Soni, and Ke Wu. Game-theoretically fair distributed sampling. In *Annual International Cryptology Conference*, pages 207–239. Springer, 2024. [1](#), [1.1.4](#), [1.1.4](#), [1.3](#), [2](#), [2](#), [2](#), [3.2](#), [4](#), [4.2.2](#)
- [WAS21] Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. *IACR Cryptol. ePrint Arch.*, page 748, 2021. [1](#), [1](#), [2](#)
- [WAS22] Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 120–149, Cham, 2022. Springer International Publishing. [1.1.4](#), [1.3](#), [2](#), [3.2](#), [4.2.2](#)
- [Wee10] Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 531–540, 2010. [1.3](#)
- [Yao82] Andrew C Yao. Protocols for secure computations. In *FOCS*, 1982. [1](#), [1.1.2](#), [1.3](#)



## Supplementary Materials

### A Preliminary

#### A.1 Shamir Secret Sharing.

A  $k$ -out-of- $n$  secret sharing scheme consists of two algorithms `Share` and `Reconstruct`

- $(s_1, \dots, s_n) \leftarrow \text{Share}(s)$ : On input a secret  $s$ , outputs  $n$  shares  $s_1, \dots, s_n$  of the secret.
- $s \leftarrow \text{Reconstruct}(I, (s_i)_{i \in I})$ : The reconstruction algorithm takes  $(s_i)_{i \in I}$  and reconstructs the secret only if  $|I| \geq k$ . Otherwise, the reconstruction fails and output `fail`.

A  $k$ -out-of- $n$  Shamir secret sharing scheme satisfies the following properties:

- **Correctness**: For any secret  $s$  and any sets  $I \subset \{1, \dots, n\}$  such that  $|I| \geq k$ ,

$$\Pr[(s_1, \dots, s_n) \leftarrow \text{Share}(s) : \text{Reconstruct}(I, (s_i)_{i \in I}) = s] = 1.$$

- **Security**: For any secrets  $s$  and  $s'$ , for any sets  $I$  such that  $I \subset \{1, \dots, n\}$ , and  $|I| \leq k - 1$ ,

$$\{(s_1, \dots, s_n) \leftarrow \text{Share}(s) : (s_i)_{i \in I}\} \approx \{(s'_1, \dots, s'_n) \leftarrow \text{Share}(s') : (s'_i)_{i \in I}\}.$$

Informally, this means that if a coalition has only  $k - 1$  shares, then they learn no information about the secret.

- **Linearity**: For any secrets  $s$  and  $s'$ , for any sets  $I \subset \{1, \dots, n\}$  such that  $|I| \geq k$ , the following holds:

$$\Pr \left[ \begin{array}{l} (s_1, \dots, s_n) \leftarrow \text{Share}(s), \\ (s'_1, \dots, s'_n) \leftarrow \text{Share}(s') \end{array} : \text{Reconstruct}(I, (s_i + s'_i)_{i \in I}) = s + s' \right] = 1.$$

Roughly speaking, this means that one can reconstruct  $s + s'$  using the sum of the shares of  $s$  and  $s'$ .

In this work, we consider Shamir's secret sharing on a finite field  $\mathbb{F}_q$  for some prime  $q$ . Specifically, we will consider a finite field of size  $q > n \cdot m$  in our  $m$ -sided coin-toss protocol among  $n$  parties. The sharing algorithm randomly chooses a degree  $k - 1$  polynomial  $p(\cdot)$  such that  $p(0) = s$ . The reconstruction algorithm takes in  $k$  or more number of shares and interpolates the polynomial to recover  $s$ .

#### A.2 Secure Public-Key Encryption.

A public-key encryption consists of three algorithms:

- $\text{Gen}(1^\lambda)$  takes in the security parameter  $\lambda$  and outputs a key pair  $(\text{pk}, \text{sk})$ .
- $\text{ct} \leftarrow \text{Enc}(\text{pk}, \text{msg})$ : take in the public key and a message  $\text{msg}$ , outputs a ciphertext  $\text{ct}$ .
- $\text{msg} \leftarrow \text{Dec}(\text{sk}, \text{ct})$ : take in a ciphertext  $\text{ct}$  and the secret key  $\text{sk}$ , output  $\text{msg}$ .

A semantically secure public-key encryption satisfies the following properties:

- **Correctness**:  $\Pr[(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) : \text{Dec}(\text{sk}, \text{Enc}(\text{pk}, \text{msg})) = \text{msg}] = 1$ .

- **Security:** We say that the encryption scheme is semantically secure iff

$$\left| \Pr \left[ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\text{msg}_0, \text{msg}_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk}) : \mathcal{A}(1^\lambda, \text{ct}) = 1 \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, \text{msg}_0) \end{array} \right] - \Pr \left[ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (\text{msg}_0, \text{msg}_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk}) : \mathcal{A}(1^\lambda, \text{ct}) = 1 \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, \text{msg}_1) \end{array} \right] \right| \leq \text{negl}(n)$$

### A.3 Digital Signature Scheme

A digital signature scheme  $\Pi_{\text{Sig}}$  consists of the following algorithms:

- $(\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ : takes in the security parameter  $\lambda$  and outputs a key pair  $(\text{vk}, \text{sk})$  of verification key  $\text{vk}$  and signing key  $\text{sk}$ .
- $\sigma \leftarrow \text{Sign}(m, \text{sk})$ : takes in a message  $m$  and the signing key  $\text{sk}$ , produces a signature  $\sigma$ .
- $b \leftarrow \text{Ver}(\text{vk}, \sigma, m)$ : takes in a message  $m$ , the signature  $\sigma$ , and the verification key  $\text{vk}$ , output 0 or 1 indicating reject or accept.

A secure digital scheme satisfies the following properties

- **Correctness:** Except with negligible probability over  $(\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ , it must be that for any message  $m$ ,  $\text{Ver}(\text{vk}, m, \text{Sign}(\text{sk}, m)) = 1$ .
- **Security:** A digital scheme is existentially unforgeable under chosen-message attack iff

$$\Pr \left[ \begin{array}{l} (\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (m^*, \sigma) \leftarrow \mathcal{A}^{\text{Sign}_{\text{sk}}(\cdot)}(1^\lambda, \text{vk}) : \text{Ver}(\text{vk}, m^*, \sigma) = 1 \end{array} \right] \leq \text{negl}(n)$$

## B $\mathcal{F}_{\text{toss}}^{\text{auth}}$ Implies CSP Fairness and Preference Privacy

**CSP-fairness.** CSP-fairness was first introduced in [CGL<sup>+</sup>18a]. It states that no coalition can increase its own expected utility except for a negligible term, no matter how it deviates from the protocol. For a coin-toss protocol  $\Pi$  with trusted setup  $\text{Setup}$  and a coalition  $\mathcal{A}$ , let  $\Pi^{(S_{\mathcal{A}}, H_{-\mathcal{A}})}(\text{pp}, \text{sk}_1, \dots, \text{sk}_n)$  denote an execution of protocol  $\Pi$  when every party holds public parameter  $\text{pp}$ , each party holds  $\text{sk}_i$  as private parameter, the coalition  $\mathcal{A}$  adopts strategy  $S_{\mathcal{A}}$  and other parties behave honestly. Each trace  $tr$  sampled from the random execution  $\Pi^{(S_{\mathcal{A}}, H_{-\mathcal{A}})}(\text{pp}, \text{sk}_1, \dots, \text{sk}_n)$  assigned a joint utility  $\text{util}_{\mathcal{A}}(tr)$  of coalition  $\mathcal{A}$ . We use  $\text{util}_{\mathcal{A}}(S_{\mathcal{A}}, H_{-\mathcal{A}})$  to denote the expected utility of the coalition:

$$\begin{aligned} & \text{util}_{\mathcal{A}}(S_{\mathcal{A}}, H_{-\mathcal{A}}) \\ & := \mathbb{E} \left[ (\text{pp}, \text{sk}_1, \dots, \text{sk}_n) \leftarrow \text{Setup}(1^\lambda, \mathcal{P}), tr \leftarrow \Pi^{(S_{\mathcal{A}}, H_{-\mathcal{A}})}(\text{pp}, \text{sk}_1, \dots, \text{sk}_n) : \text{util}_{\mathcal{A}}(tr) \right]. \end{aligned}$$

**Definition B.1** (Property-based definition, CSP fairness [CGL<sup>+</sup>18c]). An  $m$ -sided coin-toss protocol  $\Pi$  with trusted setup  $\text{Setup}$  for a preference profile  $\mathcal{P}$  is *cooperative-strategy-proofness* (CSP-fairness) against any  $t$ -sized coalition iff for any coalition  $\mathcal{A}$  of size no more than  $t$ , for any non-uniform PPT strategy  $S_{\mathcal{A}}$  that  $\mathcal{A}$  adopts, there exists a negligible function  $\text{negl}(\cdot)$ , such that  $\text{util}_{\mathcal{A}}(S_{\mathcal{A}}, H_{-\mathcal{A}}) \leq \text{util}_{\mathcal{A}}(H_{\mathcal{A}}, H_{-\mathcal{A}}) + \text{negl}(\lambda)$ , where  $H_{\mathcal{A}}$  denotes the honest strategy of  $\mathcal{A}$ .

**Remark B.2** ( $\mathcal{F}_{\text{toss}}^{\text{auth}}$  (Figure 1) implies Definition B.1). The ideal functionality  $\mathcal{F}_{\text{toss}}^{\text{auth}}$  guarantees Definition B.1. The outcome is either a uniformly random coin chosen by the ideal functionality, or a coin is such that the corrupted coalition's joint utility is at most  $\frac{|A|}{m}$ . Clearly, in the ideal world where parties have access to  $\mathcal{F}_{\text{toss}}^{\text{auth}}$ , the corrupted coalition has no strategy to increase its joint utility. Therefore, if a real-world protocol securely instantiates this ideal functionality, it also satisfies Definition B.1.