# An Optimally Fair Coin Toss* 

Tal Moran ${ }^{\dagger}$<br>Efi Arazi School of Computer Science, IDC Herzliya, Herzliya, Israel talm@idc.ac.il<br>Moni Naor<br>$\ddagger$ Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, 76100 Rehovot, Israel<br>moni.naor@weizmann.ac.il<br>Gil Segev<br>§School of Computer Science and Engineering, Hebrew University of Jerusalem, 91904 Jerusalem, Israel segev@cs.huji.ac.il<br>Communicated by Yevgeniy Dodis.<br>Received 17 May 2009<br>Online publication 21 February 2015


#### Abstract

We address one of the foundational problems in cryptography: the bias of coin-flipping protocols. Coin-flipping protocols allow mutually distrustful parties to generate a common unbiased random bit, guaranteeing that even if one of the parties is malicious, it cannot significantly bias the output of the honest party. A classical result by Cleve (Proceedings of the 18th annual ACM symposium on theory of computing, pp 364-369, 1986) showed that for any two-party $r$-round coin-flipping protocol there exists an efficient adversary that can bias the output of the honest party by $\Omega(1 / r)$. However, the best previously known protocol only guarantees $O(1 / \sqrt{r})$ bias, and the question of whether Cleve's bound is tight has remained open for more than 20 years. In this paper, we establish the optimal trade-off between the round complexity and the bias of two-party coin-flipping protocols. Under standard assumptions (the existence of oblivious transfer), we show that Cleve's lower bound is tight: We construct an $r$-round protocol with bias $O(1 / r)$.


[^0]Keywords. Coin flipping, Round complexity, Optimal bias.

## 1. Introduction

A coin-flipping protocol allows mutually distrustful parties to generate a common unbiased random bit. Such a protocol should satisfy two properties. First, when all parties are honest and follow the instructions of the protocol, their common output is a uniformly distributed bit. Second, even if some of the parties collude and deviate from the protocol's instructions, they should not be able to significantly bias the common output of the honest parties.
When a majority of the parties are honest, efficient and completely fair coin-flipping protocols are known as a special case of secure multiparty computation with an honest majority [9] (assuming a broadcast channel). When an honest majority is not available, and in particular when there are only two parties, the situation is more complex. Blum's two-party coin-flipping protocol [12] guarantees that the output of the honest party is unbiased only if the malicious party does not abort prematurely (note that the malicious party can decide to abort after learning the result of the coin flip). This satisfies a rather weak notion of fairness in which once the malicious party is labeled as a "cheater" the honest party is allowed to halt without outputting any value. Blum's protocol relies on any one-way function [26,34], and Impagliazzo and Luby [28] showed that one-way functions are in fact essential even for such a seemingly weak notion. While this notion suffices for some applications, in many cases fairness is required to hold even if one of the parties aborts prematurely (consider, for example, an adversary that controls the communication channel and can prevent communication between the parties). In this paper, we consider a stronger notion: Even when the malicious party is labeled as a cheater, we require that the honest party outputs a bit.

Cleve's Impossibility Result The latter notion of fairness turns out to be impossible to achieve in general. Specifically, Cleve [14] showed that for any two-party $r$-round coin-flipping protocol there exists an efficient adversary that can bias the output of the honest party by $\Omega(1 / r)$. Cleve's lower bound holds even under arbitrary computational assumptions: The adversary only needs to simulate an honest party and decide whether or not to abort early depending on the output of the simulation. However, the best previously known protocol (with respect to bias) only guaranteed $O(1 / \sqrt{r})$ bias [5,14], and the question of whether Cleve's bound was tight has remained open for over 20 years.

Fairness in Secure Computation The bias of coin-flipping protocols can be viewed as a particular case of the more general framework of fairness in secure computation. Typically, the security of protocols is formalized by comparing their execution in the real model to an execution in an ideal model where a trusted party receives the inputs of the parties, performs the computation on their behalf, and then sends all parties their respective outputs. Executions in the ideal model guarantee complete fairness: Either all parties learn the output, or neither party does. Cleve's result, however, shows that without an honest majority complete fairness is generally impossible to achieve, and therefore, the formulation of secure computation (see [19]) weakens the ideal model to
one in which fairness is not guaranteed. Informally, a protocol is "secure-with-abort" if its execution in the real model is indistinguishable from an execution in the ideal model allowing the ideal-model adversary to choose whether the honest parties receive their outputs (this is the notion of security satisfied by Blum's coin-flipping protocol).

Recently, Katz [29] suggested an alternate relaxation: keep the ideal model unchanged (i.e., all parties always receive their outputs), but relax the notion of indistinguishability by asking that the real model and ideal model are distinguishable with probability at most $1 / p(n)+\nu(n)$, for a polynomial $p(n)$ and a negligible function $v(n)$ (we refer the reader to Sect. 2 for a formal definition). Protocols satisfying this requirement are said to be $1 / p$-secure, and intuitively, such protocols guarantee complete fairness in the real model except for probability $1 / p$. In the context of coin-flipping protocols, any $1 / p$-secure protocol has bias at most $1 / p$. However, the definition of $1 / p$-security is more general and applies to a larger class of functionalities.

### 1.1. Our Contributions

In this paper, we establish the optimal trade-off between the round complexity and the bias of two-party coin-flipping protocols. We prove the following theorem:

Theorem 1.1. Assuming the existence of oblivious transfer, for any polynomial $r=$ $r(n)$ there exists an $r$-round two-party coin-flipping protocol that is $1 /(4 r-c)$-secure, for some constant $c>0$.

We prove the security of our protocol under the simulation-based definition of $1 / p$ security, ${ }^{1}$ which for coin-flipping protocols implies, in particular, that the bias is at most $1 / p$. We note that our result not only identifies the optimal trade-off asymptotically, but almost pins down the exact leading constant: Cleve showed that any $r$-round two-party coin-flipping protocol has bias at least $1 /(8 r+2)$, and we manage to achieve bias of at most $1 /(4 r-c)$ for some constant $c>0$.

Our approach holds in fact for a larger class of functionalities. We consider the more general task of sampling from a distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ : Party $P_{1}$ receives a sample from $\mathcal{D}_{1}$, and party $P_{2}$ receives a correlated sample from $\mathcal{D}_{2}$ (in coin flipping, for example, the joint distribution $\mathcal{D}$ produces the values $(0,0)$ and $(1,1)$ each with probability $1 / 2$ ). Before stating our result in this setting, we introduce a standard notation: We denote by $\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)$ the statistical distance between the joint distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ and the direct product of the two marginal distributions $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$. We prove the following theorem which generalizes Theorem 1.1:

Theorem 1.2. Assuming the existence of oblivious transfer, for any polynomially sampleable distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ and polynomial $r=r(n)$ there exists an $r$-round two-party protocol for sampling from $\mathcal{D}$ that is $\frac{\mathrm{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)}{2 r-c}$-secure, for some constant $c>0$.

[^1]Our approach raises several open questions that are fundamental to the understanding of coin-flipping protocols. These questions include identifying the minimal computational assumptions that are essential for reaching the optimal trade-off (i.e., one-way functions vs. oblivious transfer), extending our approach to the multiparty setting, and constructing a more efficient variant of our protocol that can result in a practical implementation. We elaborate on these questions in Sect. 5 and hope that our approach and the questions it raises can make progress toward resolving the complexity of coin-flipping protocols.

### 1.2. Related Work

Coin-Flipping Protocols When security with abort is sufficient, simple variations of Blum's protocol are the most commonly used coin-flipping protocols. For example, an $r$-round protocol with bias $O(1 / \sqrt{r})$ can be constructed by sequentially executing Blum's protocol $O(r)$ times and outputting the majority of the intermediate output values [5,14]. We note that in this protocol an adversary can indeed bias the output by $\Omega(1 / \sqrt{r})$ by aborting prematurely. One of the most significant results on the bias of coin-flipping protocols gave reason to believe that the optimal trade-off between the round complexity and the bias is in fact $\Theta(1 / \sqrt{r})$ (as provided by the latter variant of Blum's protocol): Cleve and Impagliazzo [15] showed that in the fail-stop model, any two-party $r$-round coin-flipping protocol has bias $\Omega(1 / \sqrt{r})$. In the fail-stop model, adversaries are computationally unbounded, but they must follow the instructions of the protocol except for being allowed to abort prematurely.

Coin-flipping protocols were also studied in a variety of other models. Among those are collective coin flipping in the "perfect information model" in which parties are computationally unbounded and all communication is public [2,10,18,35,36], and protocols based on physical assumptions, such as quantum computation $[1,3,4]$ and tamper-evident seals [33].

Fair Computation The techniques underlying our protocols were directly inspired by a recent line of research devoted for achieving various forms of fairness in secure computation. Specifically, the technique of choosing a secret "threshold round," before which no information is learned, and after which aborting the protocol is essentially useless was suggested by Moran and Naor [33] as part of a coin-flipping protocol based on tamperevident seals, by Katz [29] for partially fair protocols using a simultaneous broadcast channel and by Gordon et al. [20] for completely fair protocols for a restricted (but yet rather surprising) class of functionalities.

Various techniques for hiding a meaningful round in game-theoretic settings were suggested by Halpern and Teague [25], Gordon and Katz [21], and Kol and Naor [30]. Katz [29] also introduced the technique of distributing shares to the parties in an initial setup phase (which is only secure-with-abort), and these shares are then exchanged by the parties in each round of the protocol.

### 1.3. Subsequent Work

Multiparty Coin Flipping Beimel et al. [7] generalized our two-party protocol to the multiparty case where fewer than $2 / 3$ of the parties are corrupt. They showed an $m$-party
$r$-round protocol that can tolerate $t$ corrupt parties (where $m / 2 \leq t \leq 2 m / 3$ ) with bias $O\left(2^{2^{k+1}} / r^{\prime}\right)$, where $r^{\prime}=r-O(k)$ and $k=2 t-m$. In particular, when $t$ and $m$ are constant (but $r$ grows with $n$ ), this gives a protocol with $O(1 / r)$ bias.

Recently, Haitner and Tsfadia [24] managed to break the $2 / 3$ barrier with the construction of a protocol for three parties with bias $O\left(\log ^{2} r / r\right)$. Their protocol requires a new technique-rather than relying on a "special round" at which the value of the game changes abruptly for the parties (i.e., the expected output of the game is $1 / 2$ before this round, and either 0 or 1 afterward); the new protocol "smoothly" changes the value of the game. Haitner and Tsfadia show that this is necessary when there are more than two parties and arbitrary coalitions may be corrupted. Their results seem somewhat limited, however, to three parties; dealing with more than three parties and arbitrary corruptions remains an open problem (Haitner and Tsfadia conjecture that their technique can be extended to a logarithmic number of parties).

1/p-Secure Computation Our results were recently extended by Gordon and Katz [22] to deal with the more general case of randomized functions, and not only distributions. Gordon and Katz showed that any efficiently computable randomized function $f: X \times$ $Y \rightarrow Z$ where at least one of $X$ and $Y$ is of polynomial size has an $r$-round protocol that is $O\left(\frac{\min \{|X|,|Y|\}}{r}\right)$-secure. In addition, they showed that even if both domains are of superpolynomial size but the range $Z$ is of polynomial size, then $f$ has an $r$-round protocol that is $O\left(\frac{|Z|}{\sqrt{r}}\right)$-secure. Gordon and Katz also showed a specific function $f: X \times Y \rightarrow Z$ where $X, Y$, and $Z$ are of size super-polynomial which cannot be $1 / p$-securely computed for any $p>2$ assuming the existence of exponentially hard one-way functions.

These results were further extended by Beimel et al. [6] to the multiparty setting. For a constant number of parties and any function $f$ with range of size polynomial in the security parameter, Beimel et al. show that there exists a protocol for $1 / p$-secure computation of $f$ tolerating any number of corrupt parties. Moreover, if the number of corrupted parties is less than $2 / 3$, they construct a protocol that can handle a superconstant number of parties (up to $\log n$ ), as long as the domain of $f$ is of constant size. In the other direction, Beimel et al. show that no protocols exist for $1 / p$-secure computation of functions with polynomial-size domain if the number of parties is super-constant.

Coin Flipping Versus One-Way Functions Our work shows how to construct an optimally fair coin-flipping protocol based on oblivious transfer; however, it has left open to find the minimal assumptions necessary for optimally fair coin flipping. The work of Dachman-Soled et al. [16] took a step toward answering this question by showing that any black-box construction of an optimally fair coin-flipping protocol based on a one-way function with $n$-bit input and output needs $\Omega(n / \log n)$ rounds. Subsequently, Dachman-Soled et al. [17] took another step toward understanding the complexity of optimally fair coin flipping by showing that this task (with an arbitrary number of rounds) cannot be based on one-way functions in a black-box way, as long as the protocol is "oblivious" to the implementation of the one-way function (we refer the reader to [17] for more details on their notion of obliviousness).

Surprisingly, an even more fundamental question is still open: Are one-way functions necessary at all for constructing fair coin-flipping protocols? While Impagliazzo and

Luby [28] gave a positive answer for protocols with negligible, only recently have we begin to close the gap with regard to protocols that guarantee only non-negligible bias. Maji et al. [32] showed that one-way functions are required for constant-round protocols with bias $1 / 2-o(1)$. Haitner and Omri [23] showed that strong coin-flipping protocols with bias $\frac{\sqrt{2}-1}{2}$ and any number of rounds imply one-way functions. Berman et al. [11] subsequently improved this result, proving that coin-flipping protocols with any constant bias (and any number of rounds) imply one-way functions (the stronger result also applies to weak coin-flipping protocols). The question of whether protocols with non-constant bias and non-constant rounds can be constructed without one-way functions is still open, to the best of our knowledge.

### 1.4. Paper Organization

The remainder of this paper is organized as follows. In Sect. 2, we review several notions and definitions that are used in the paper (most notably, the definition of $1 / p$-secure computation). In Sect. 3, we describe a simplified variant of our protocol and prove its security. In Sect. 4, we describe a more refined and general variant of our protocol which settles Theorems 1.1 and 1.2. Finally, in Sect. 5, we discuss several open problems.

## 2. Preliminaries

In this section, we review the definitions of coin-flipping protocols, $1 / p$-indistinguishability and $1 / p$-secure computation (taken almost verbatim from [22,29]), security with abort, and one-time message authentication.

### 2.1. Coin-Flipping Protocols

A two-party coin-flipping protocol is defined via two probabilistic polynomial-time Turing machines ( $P_{1}, P_{2}$ ), referred to as parties, that receive as input a security parameter $1^{n}$. The parties exchange messages in a sequence of rounds, where in every round each party both sends and receives a message (i.e., a round consists of two moves). At the end of the protocol, $P_{1}$ and $P_{2}$ produce output bits $c_{1}$ and $c_{2}$, respectively. We denote by $\left(c_{1} \mid c_{2}\right) \leftarrow\left\langle P_{1}\left(1^{n}\right), P_{2}\left(1^{n}\right)\right\rangle$ the experiment in which $P_{1}$ and $P_{2}$ interact (using uniformly chosen random coins), and then $P_{1}$ outputs $c_{1}$ and $P_{2}$ outputs $c_{2}$. It is required that for all sufficiently large $n$, and every possible pair $\left(c_{1}, c_{2}\right)$ that may be output by $\left\langle P_{1}\left(1^{n}\right), P_{2}\left(1^{n}\right)\right\rangle$, it holds that $c_{1}=c_{2}$ (i.e., $P_{1}$ and $P_{2}$ agree on a common value). This requirement can be relaxed by asking that the parties agree on a common value with sufficiently high probability. ${ }^{2}$

The security requirement of a coin-flipping protocol is that even if one of $P_{1}$ and $P_{2}$ is corrupted and arbitrarily deviates from the protocol's instructions, the bias of the honest party's output remains bounded. Specifically, we emphasize that a malicious party is allowed to abort prematurely, and in this case, it is assumed that the honest party is

[^2]notified on the early termination of the protocol. In addition, we emphasize that even when the malicious party is labeled as a cheater, the honest party must output a bit. For simplicity, the following definition considers only the case in which $P_{1}$ is corrupted, and an analogous definition holds for the case that $P_{2}$ is corrupted:

Definition 2.1. A coin-flipping protocol $\left(P_{1}, P_{2}\right)$ has bias at most $\epsilon(n)$ if for every probabilistic polynomial-time Turing machine $P_{1}^{*}$ it holds that

$$
\left|\operatorname{Pr}\left[\left(c_{1} \mid c_{2}\right) \leftarrow\left\langle P_{1}^{*}\left(1^{n}\right), P_{2}\left(1^{n}\right)\right\rangle: c_{2}=1\right]-\frac{1}{2}\right| \leq \epsilon(n)+v(n),
$$

for some negligible function $v(n)$ and for all sufficiently large $n$.

## 2.2. $1 / p$-Indistinguishability and $1 / p$-Secure Computation

$1 / p$-Indistinguishability A distribution ensemble $X=\{X(a, n)\}_{a \in \mathcal{D}_{n}, n \in \mathbb{N}}$ is an infinite sequence of random variables indexed by $a \in \mathcal{D}_{n}$ and $n \in \mathbb{N}$, where $\mathcal{D}_{n}$ is a set that may depend on $n$. For a fixed polynomial $p(n)$, two distribution ensembles $X=\{X(a, n)\}_{a \in \mathcal{D}_{n}, n \in \mathbb{N}}$ and $Y=\{Y(a, n)\}_{a \in \mathcal{D}_{n}, n \in \mathbb{N}}$ are computationally $1 / p$ indistinguishable, denoted $X \stackrel{1 / p}{\approx} Y$, if for every non-uniform polynomial-time algorithm $D$ there exists a negligible function $\nu(n)$ such that for all sufficiently large $n \in \mathbb{N}$ and for all $a \in \mathcal{D}_{n}$ it holds that

$$
|\operatorname{Pr}[D(X(a, n))=1]-\operatorname{Pr}[D(Y(a, n))=1]| \leq \frac{1}{p(n)}+v(n) .
$$

$1 / p$-Secure Computation A two-party protocol for computing a functionality $\mathcal{F}=$ $\left\{\left(f^{1}, f^{2}\right)\right\}$ is a protocol running in polynomial time and satisfying the following functional requirement: If party $P_{1}$ holds input $\left(1^{n}, x\right)$, and party $P_{2}$ holds input $\left(1^{n}, y\right)$, then the joint distribution of the outputs of the parties is statistically close to ( $f^{1}(x, y), f^{2}(x, y)$ ). In what follows we define the notion of $1 / p$-secure computation [22,29]. The definition uses the real/ideal paradigm (following [13, 19,27]) where we consider a completely fair ideal model (as typically considered in the setting of honest majority) and require only $1 / p$-indistinguishability rather than indistinguishability (we note that, in general, the notions of $1 / p$-security and security with abort are incomparable). We consider active adversaries, who may deviate from the protocol in an arbitrary manner, and static corruptions.

Security of Protocols (Informal) The security of a protocol is analyzed by comparing what an adversary can do in a real protocol execution to what it can do in an ideal scenario that is secure by definition. This is formalized by considering an ideal computation involving an incorruptible trusted party to whom the parties send their inputs. The trusted party computes the functionality on the inputs and returns to each party its respective output. Loosely speaking, a protocol is secure if any adversary interacting in the real protocol (where no trusted party exists) can do no more harm than if it was involved in the above-described ideal computation.

Execution in the Ideal Model The parties are $P_{1}$ and $P_{2}$, and there is an adversary $\mathcal{A}$ who has corrupted one of them. An ideal execution for the computation of $\mathcal{F}=\left\{f_{n}\right\}$ proceeds as follows:

Inputs: $P_{1}$ and $P_{2}$ hold the security parameter $1^{n}$ and inputs $x \in X_{n}$ and $y \in Y_{n}$, respectively. The adversary $\mathcal{A}$ receives an auxiliary input aux.
Send inputs to trusted party: The honest party sends its input to the trusted party. The corrupted party may send an arbitrary value (chosen by $\mathcal{A}$ ) to the trusted party. Denote the pair of inputs sent to the trusted party by $\left(x^{\prime}, y^{\prime}\right)$.
Trusted party sends outputs: If $x^{\prime} \notin X_{n}$, the trusted party sets $x^{\prime}$ to some default element $x_{0} \in X_{n}$ (and likewise if $y^{\prime} \notin Y_{n}$ ). Then, the trusted party chooses $r$ uniformly at random and sends $f_{n}^{1}\left(x^{\prime}, y^{\prime} ; r\right)$ to $P_{1}$ and $f_{n}^{2}\left(x^{\prime}, y^{\prime} ; r\right)$ to $P_{2}$.
Outputs: The honest party outputs whatever it was sent by the trusted party, the corrupted party outputs nothing, and $\mathcal{A}$ outputs any arbitrary (probabilistic polynomial-time computable) function of its view.

We let $\operatorname{IDEAL}_{\mathcal{F}, \mathcal{A}(\mathrm{aux})}(x, y, n)$ be the random variable consisting of the view of the adversary and the output of the honest party following an execution in the ideal model as described above.

Execution in the Real Model We now consider the real model in which a two-party protocol $\pi$ is executed by $P_{1}$ and $P_{2}$ (and there is no trusted party). The protocol execution is divided into rounds; in each round, one of the parties sends a message. The honest party computes its messages as specified by $\pi$. The messages sent by the corrupted party are chosen by the adversary, $\mathcal{A}$, and can be an arbitrary (polynomial-time) function of the corrupted party's inputs, random coins, and the messages received from the honest party in previous rounds. If the corrupted party aborts in one of the protocol rounds, the honest party behaves as if it had received a special $\perp$ symbol in that round.

Let $\pi$ be a two-party protocol computing $\mathcal{F}$. Let $\mathcal{A}$ be a non-uniform probabilistic polynomial-time machine with auxiliary input aux. We let $\operatorname{REAL}_{\pi, \mathcal{A} \text { (aux) }}(x, y, n)$ be the random variable consisting of the view of the adversary and the output of the honest party, following an execution of $\pi$ where $P_{1}$ begins by holding input $\left(1^{n}, x\right)$ and $P_{2}$ begins by holding input $\left(1^{n}, y\right)$.

Security as Emulation of an Ideal Execution in the Real Model Having defined the ideal and real models, we can now define security of a protocol. Loosely speaking, the definition asserts that a secure protocol (in the real model) emulates the ideal model (in which a trusted party exists). This is formulated as follows:
Definition 2.2. ( $1 / p$-secure computation) Let $\mathcal{F}$ and $\pi$ be as above and fix a function $p=p(n)$. Protocol $\pi$ is said to $1 / p$-securely compute $\mathcal{F}$ if for every non-uniform probabilistic polynomial-time adversary $\mathcal{A}$ in the real model, there exists a non-uniform probabilistic polynomial-time adversary $\mathcal{S}$ in the ideal model such that

$$
\begin{aligned}
& \left\{\operatorname{IDEAL} \mathcal{F}_{\mathcal{F}, \mathcal{S}(\mathrm{aux})}(x, y, n)\right\}_{(x, y) \in X \times Y, \mathrm{aux} \in\{0,1\}^{*}} \\
& \stackrel{1 / p}{\approx}\left\{\operatorname{REAL}_{\pi, \mathcal{A}(\mathrm{aux})}(x, y, n)\right\}_{(x, y) \in X \times Y, \mathrm{aux} \in\{0,1\}^{*}}
\end{aligned}
$$

and the same party is corrupted in both the real and ideal models.

### 2.3. Security With Abort

In what follows we use the standard notion of computational indistinguishability. That is, two distribution ensembles $X=\{X(a, n)\}_{a \in \mathcal{D}_{n}, n \in \mathbb{N}}$ and $Y=\{Y(a, n)\}_{a \in \mathcal{D}_{n}, n \in \mathbb{N}}$ are computationally indistinguishable, denoted $X \stackrel{c}{=} Y$, if for every non-uniform polynomial-time algorithm $D$ there exists a negligible function $v(n)$ such that for all sufficiently large $n \in \mathbb{N}$ and for all $a \in \mathcal{D}_{n}$ it holds that

$$
|\operatorname{Pr}[D(X(a, n))=1]-\operatorname{Pr}[D(Y(a, n))=1]| \leq v(n) .
$$

Security with abort is the standard notion for secure computation where an honest majority is not available. The definition is similar to the definition of $1 / p$-security presented in Sect. 2.2, with the following two exceptions: (1) The ideal-model adversary is allowed to choose whether the honest parties receive their outputs (i.e., fairness is not guaranteed) and (2) the ideal model and real model are required to be computationally indistinguishable. Specifically, the execution in the real model is as described in Sect. 2.2, and the execution in the ideal model is modified as follows:

Inputs: $P_{1}$ and $P_{2}$ hold the security parameter $1^{n}$ and inputs $x \in X_{n}$ and $y \in Y_{n}$, respectively. The adversary $\mathcal{A}$ receives an auxiliary input aux.
Send inputs to trusted party: The honest party sends its input to the trusted party. The corrupted party controlled by $\mathcal{A}$ may send any value of its choice. Denote the pair of inputs sent to the trusted party by $\left(x^{\prime}, y^{\prime}\right)$.
Trusted party sends output to corrupted party: If $x^{\prime} \notin X_{n}$, the trusted party sets $x^{\prime}$ to some default element $x_{0} \in X_{n}$ (and likewise if $y^{\prime} \notin Y_{n}$ ). Then, the trusted party chooses $r$ uniformly at random, computes $z_{1}=f_{n}^{1}\left(x^{\prime}, y^{\prime} ; r\right)$ and $z_{2}=$ $f_{n}^{2}\left(x^{\prime}, y^{\prime} ; r\right)$ to $P_{2}$, and sends $z_{i}$ to the corrupted party $P_{i}$ (i.e., to the adversary $\mathcal{A})$.
Adversary decides whether to abort: After receiving its output, the adversary sends either "abort" or "continue" to the trusted party. In the former case, the trusted party sends $\perp$ to the honest party
$P_{j}$, and in the latter case, the trusted party sends $z_{j}$ to $P_{j}$.
Outputs: The honest party outputs whatever it was sent by the trusted party, the corrupted party outputs nothing, and $\mathcal{A}$ outputs any arbitrary (probabilistic polynomial-time computable) function of its view.
We let $\operatorname{IDEAL}_{\mathcal{F}, \mathcal{A}(\mathrm{aux})}^{\text {abort }}(x, y, n)$ be the random variable consisting of the view of the adversary and the output of the honest party following an execution in the ideal model as described above.

Definition 2.3. (Security with abort) Let $\mathcal{F}$ and $\pi$ be as above. Protocol $\pi$ is said to securely compute $\mathcal{F}$ with abort if for every non-uniform probabilistic polynomial-time adversary $\mathcal{A}$ in the real model, there exists a non-uniform probabilistic polynomial-time adversary $\mathcal{S}$ in the ideal model such that

$$
\begin{aligned}
& \left\{\operatorname{IDEAL}_{\mathcal{F}, \mathcal{S}(\mathrm{aux})}^{\text {abort }}(x, y, n)\right\}_{(x, y) \in X \times Y, \mathrm{aux} \in\{0,1\}^{*}} \\
& \quad \stackrel{c}{=}\left\{\operatorname{REAL}_{\pi, \mathcal{A}(\mathrm{aux})}(x, y, n)\right\}_{(x, y) \in X \times Y, \mathrm{aux} \in\{0,1\}^{*}}
\end{aligned}
$$

### 2.4. One-Time Message Authentication

Message authentication codes provide assurance to the receiver of a message that it was sent by a specified legitimate sender, even in the presence of an active adversary who controls the communication channel. A message authentication code is defined via triplet (Gen, Mac, Vrfy) of probabilistic polynomial-time Turing machines such that:

1. The key generation algorithm Gen receives as input a security parameter $1^{n}$ and outputs an authentication key $k$.
2. The authentication algorithm Mac receives as input an authentication key $k$ and a message $m$, and outputs a tag $t$.
3. The verification algorithm Vrfy receives as input an authentication key $k$, a message $m$, and a tag $t$, and outputs a bit $b \in\{0,1\}$.
The functionality guarantee of a message authentication code is that for any message $m$ it holds that $\operatorname{Vrfy}(k, m, \operatorname{Mac}(k, m))=1$ with overwhelming probability over the internal coin tosses of Gen, Mac, and Vrfy. In this paper, we rely on message authentication codes that are one-time secure. That is, an authentication key is used to authenticate a single message. We consider an adversary that queries the authentication algorithm on a single message $m$ of her choice and then outputs a pair ( $m^{\prime}, t^{\prime}$ ). We say that the adversary forges an authentication tag if $m^{\prime} \neq m$ and $\operatorname{Vrfy}\left(k, m^{\prime}, t^{\prime}\right)=1$. Message authentication codes that are one-time secure exist in the information-theoretic setting; that is, even an unbounded adversary has only a negligible forgery probability. Constructions of such codes can be based, for example, on pair-wise independent hash functions [37].

## 3. A Simplified Protocol

In order to demonstrate the main ideas underlying our approach, in this section we present a simplified protocol. The simplification is twofold: First, we consider the specific coinflipping functionality (as in Theorem 1.1), and not the more general functionality of sampling from an arbitrary distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ (as in Theorem 1.2). Second, the coin-flipping protocol will only be $1 /(2 r)$-secure and not $1 /(4 r)$-secure.

We describe the protocol in a sequence of refinements. We first informally describe the protocol assuming the existence of a trusted third party. The trusted third party acts as a "dealer" in a preprocessing phase, sending each party an input that it uses in the protocol. In the protocol, we make no assumptions about the computational power of the parties. We then eliminate the need for the trusted third party by having the parties execute a secure-with-abort protocol that implements its functionality (this can be done in a constant number of rounds).

The Protocol The joint input of the parties, $P_{1}$ and $P_{2}$, is the security parameter $1^{n}$ and a polynomial $r=r(n)$ indicating the number of rounds in the protocol. In the preprocessing phase, the trusted third party chooses uniformly at random a value $i^{*} \in$ $\{1, \ldots, r\}$ that corresponds to the round in which the parties learn their outputs. In every round $i \in\{1, \ldots, r\}$, each party learns one bit of information: $P_{1}$ learns a bit $a_{i}$, and $P_{2}$ learns a bit $b_{i}$. In every round $i \in\left\{1, \ldots, i^{*}-1\right\}$ (these are the "dummy" rounds), the values $a_{i}$ and $b_{i}$ are independently and uniformly chosen. In every round $i \in\left\{i^{*}, \ldots, r\right\}$, the parties learn the same uniformly distributed bit $c=a_{i}=b_{i}$ which is their output in the protocol. If the parties complete all $r$ rounds of the protocol, then $P_{1}$ and $P_{2}$ output $a_{r}$ and $b_{r}$, respectively. ${ }^{3}$ Otherwise, if a party aborts prematurely, the other party outputs the value of the previous round and halts. That is, if $P_{1}$ aborts in round $i \in\{1, \ldots, r\}$, then $P_{2}$ outputs the value $b_{i-1}$ and halts. Similarly, if $P_{2}$ aborts in round $i$, then $P_{1}$ outputs the value $a_{i-1}$ and halts.

More specifically, in the preprocessing phase the trusted third party chooses $i^{*} \in$ $\{1, \ldots, r\}$ uniformly at random and defines $a_{1}, \ldots, a_{r}$ and $b_{1}, \ldots, b_{r}$ as follows: First, it chooses $a_{1}, \ldots, a_{i^{*}-1} \in\{0,1\}$ and $b_{1}, \ldots, b_{i^{*}-1} \in\{0,1\}$ independently and uniformly at random. Then, it chooses $c \in\{0,1\}$ uniformly at random and lets $a_{i^{*}}=\cdots=$ $a_{r}=b_{i^{*}}=\cdots=b_{r}=c$. The trusted third party creates secret shares of the values $a_{1}, \ldots, a_{r}$ and $b_{1}, \ldots, b_{r}$ using an information-theoretically secure 2-out-of-2 secretsharing scheme, and these shares are given to the parties. For concreteness, we use the specific secret-sharing scheme that splits a bit $x$ into $\left(x^{(1)}, x^{(2)}\right)$ by choosing $x^{(1)} \in\{0,1\}$ uniformly at random and letting $x^{(2)}=x \oplus x^{(1)}$. In every round $i \in\{1, \ldots, r\}$, the parties exchange their shares for the current round, which enables $P_{1}$ to reconstruct $a_{i}$ and $P_{2}$ to reconstruct $b_{i}$. Clearly, when both parties are honest, the parties produce the same output bit which is uniformly distributed.

[^3]```
Functionality ShareGen \(_{r}\)
Input: Security parameter \(1^{n}\).
Computation:
    1. Choose \(i^{*} \in\{1, \ldots, r\}\) uniformly at random.
    2. Define values \(a_{1}, \ldots, a_{r}\) and \(b_{1}, \ldots, b_{r}\) as follows:
        - For \(1 \leq i \leq i^{*}-1\) choose \(a_{i}, b_{i} \in\{0,1\}\) independently and uniformly at random.
    - Choose \(c \in\{0,1\}\) uniformly at random, and for \(i^{*} \leq i \leq r\) let \(a_{i}=b_{i}=c\).
    3. For \(1 \leq i \leq r\), choose \(\left(a_{i}^{(1)}, a_{i}^{(2)}\right)\) and \(\left(b_{i}^{(1)}, b_{i}^{(2)}\right)\) as random secret shares of \(a_{i}\) and \(b_{i}\), respectively.
    4. Compute \(k_{1}^{a}, \ldots, k_{r}^{a}, k_{1}^{b}, \ldots, k_{r}^{b} \leftarrow \operatorname{Gen}\left(1^{n}\right)\). For \(1 \leq i \leq r\), let \(t_{i}^{a}=\operatorname{Mac}_{k_{i}^{a}}\left(i \| a_{i}^{(2)}\right)\) and \(t_{i}^{b}=\)
        \(\operatorname{Mac}_{k_{i}^{b}}\left(i \| b_{i}^{(1)}\right)\).
Output:
    1. Party \(P_{1}\) receives the values \(a_{1}^{(1)}, \ldots, a_{r}^{(1)},\left(b_{1}^{(1)}, t_{1}^{b}\right), \ldots,\left(b_{r}^{(1)}, t_{r}^{b}\right)\), and \(k^{a}=\left(k_{1}^{a}, \ldots, k_{r}^{a}\right)\).
    2. Party \(P_{2}\) receives the values \(\left(a_{1}^{(2)}, t_{1}^{a}\right), \ldots,\left(a_{r}^{(2)}, t_{r}^{a}\right), b_{1}^{(2)}, \ldots, b_{r}^{(2)}\), and \(k^{b}=\left(k_{1}^{b}, \ldots, k_{r}^{b}\right)\).
```

Fig. 1. The ideal functionality ShareGen $_{r}$.
Eliminating the Trusted Third Party We eliminate the need for the trusted third party by relying on a possibly unfair sub-protocol that securely computes with abort the functionality ShareGen ${ }_{r}$, formally described in Fig. 1. Such a protocol with a constant number of rounds can be constructed assuming the existence of oblivious transfer (see, for example, [31]). In addition, our protocol also relies on a one-time message authentication code (Gen, Mac, Vrfy) that is information-theoretically secure. The functionality ShareGen ${ }_{r}$ provides the parties with authentication keys and authentication tags, so each party can verify that the shares received from the other party were the ones generated by ShareGen ${ }_{r}$ in the preprocessing phase. A formal description of the protocol is provided in Fig. 2.
Proof of Security The following theorem states that the protocol is $1 /(2 r)$-secure. We then conclude the section by showing our analysis is in fact tight: There exists an efficient adversary that can bias the output of the honest party by essentially $1 /(2 r)$.

Theorem 3.1. For any polynomial $r=r(n)$, if protocol $\pi$ securely computes ShareGen $_{r}$ with abort, then protocol CoinFlip ${ }_{r}$ is $1 /(2 r)$-secure.

Proof. We prove the (1/2r)-security of protocol CoinFlip ${ }_{r}$ in a hybrid model where a trusted party for computing ShareGen ${ }_{r}$ with abort is available. Using standard techniques (see [13]), it then follows that when replacing the trusted party computing ShareGen $r$ with a sub-protocol that security computes ShareGen $r$ with abort, the resulting protocol is $1 /(2 r)$-secure.

Specifically, for every polynomial-time hybrid-model adversary $\mathcal{A}$ corrupting $P_{1}$ and running CoinFlip $r_{r}$ in the hybrid model, we show that there exists a polynomial-time ideal-model adversary $\mathcal{S}$ corrupting $P_{1}$ in the ideal model with access to a trusted party computing the coin-flipping functionality such that the statistical distance between these two executions is at most $1 /(2 r)+v(n)$, for some negligible function $v(n)$. For simplicity,

```
Protocol CoinFlip \({ }_{r}\)
Joint input: Security parameter \(1^{n}\).
Preliminary phase:
    1. Parties \(P_{1}\) and \(P_{2}\) run protocol \(\pi\) for computing \(\operatorname{ShareGen}_{r}\left(1^{n}\right)\) (see Figure 1).
    2. If \(P_{1}\) receives \(\perp\) from the above computation, it outputs a uniformly chosen bit and halts. Likewise, if
    \(P_{2}\) receives \(\perp\) it outputs a uniformly chosen bit and halts. Otherwise, the parties proceed.
    3. Denote the output of \(P_{1}\) from \(\pi\) by \(a_{1}^{(1)}, \ldots, a_{r}^{(1)},\left(b_{1}^{(1)}, t_{1}^{b}\right), \ldots,\left(b_{r}^{(1)}, t_{r}^{b}\right)\), and \(k^{a}=\left(k_{1}^{a}, \ldots, k_{r}^{a}\right)\).
    4. Denote the output of \(P_{2}\) from \(\pi\) by \(\left(a_{1}^{(2)}, t_{1}^{a}\right), \ldots,\left(a_{r}^{(2)}, t_{r}^{a}\right), b_{1}^{(2)}, \ldots, b_{r}^{(2)}\), and \(k^{b}=\left(k_{1}^{b}, \ldots, k_{r}^{b}\right)\).
In each round \(i=1, \ldots, r\) do:
    1. \(P_{2}\) sends the next share to \(P_{1}\) :
        (a) \(P_{2}\) sends \(\left(a_{i}^{(2)}, t_{i}^{a}\right)\) to \(P_{1}\).
        (b) \(P_{1}\) receives \(\left(\hat{a}_{i}^{(2)}, \hat{t}_{i}^{a}\right)\) from \(P_{2}\). If \(\mathrm{Vrfy}_{k_{i}^{a}}\left(i \| \hat{a}_{i}^{(2)}, \hat{t}_{i}^{a}\right)=0\) (or if \(P_{1}\) received an invalid message or
        no message), then \(P_{1}\) outputs \(a_{i-1}\) and halts (if \(i=1\) it outputs a uniformly chosen bit).
        (c) If \(\mathrm{Vrfy}_{k_{i}^{a}}\left(i \| \hat{a}_{i}^{(2)}, \hat{t}_{i}^{a}\right)=1\) then \(P_{1}\) reconstructs \(a_{i}\) using the shares \(a_{i}^{(1)}\) and \(\hat{a}_{i}^{(2)}\).
    2. \(P_{1}\) sends the next share to \(P_{2}\) :
    (a) \(P_{1}\) sends \(\left(b_{i}^{(1)}, t_{i}^{b}\right)\) to \(P_{2}\).
    (b) \(P_{2}\) receives \(\left(\hat{b}_{i}^{(1)}, \hat{t}_{i}^{b}\right)\) from \(P_{1}\). If \(\operatorname{Vrfy}_{k_{i}^{b}}\left(i \| \hat{b}_{i}^{(1)}, \hat{t}_{i}^{b}\right)=0\) (or if \(P_{2}\) received an invalid message or
        no message), then \(P_{2}\) outputs \(b_{i-1}\) and halts (if \(i=1\) it outputs a uniformly chosen bit).
    (c) If \(\operatorname{Vrfy}_{k_{i}^{b}}\left(i \| \hat{b}_{i}^{(1)}, \hat{t}_{i}^{b}\right)=1\) then \(P_{2}\) reconstructs \(b_{i}\) using the shares \(b_{i}^{(1)}\) and \(b_{i}^{(2)}\)
```

Output: $P_{1}$ and $P_{2}$ output $a_{r}$ and $b_{r}$, respectively.

Fig. 2. The coin-flipping protocol CoinFlip ${ }_{r}$.
in the remainder of the proof we ignore the aspect of message authentication in the protocol and assume that the only malicious behavior of the adversary $\mathcal{A}$ is early abort. This does not result in any loss of generality, since there is only a negligible probability of forging an authentication tag.

On input ( $1^{n}$, aux), the ideal-model adversary $\mathcal{S}$ invokes the hybrid-model adversary $\mathcal{A}$ on ( $1^{n}$, aux) and queries the trusted party computing the coin-flipping functionality to obtain a bit $c$. The ideal-model adversary $\mathcal{S}$ proceeds as follows:

1. $\mathcal{S}$ simulates the trusted party computing the $\mathrm{ShareGen}_{r}$ functionality by sending $\mathcal{A}$ independently and uniformly chosen shares $a_{1}^{(1)}, \ldots, a_{r}^{(1)}, b_{1}^{(1)}, \ldots, b_{r}^{(1)}$. If $\mathcal{A}$ aborts (i.e., if $\mathcal{A}$ sends abort to the simulated ShareGen $_{r}$ after receiving the shares), then $\mathcal{S}$ outputs $\mathcal{A}$ 's output and halts.
2. $\mathcal{S}$ chooses $i^{*} \in\{1, \ldots, r\}$ uniformly at random.
3. In every round $i \in\left\{1, \ldots, i^{*}-1\right\}, \mathcal{S}$ chooses a random bit $a_{i}$ and sends $\mathcal{A}$ the share $a_{i}^{(2)}=a_{i}^{(1)} \oplus a_{i}$. If $\mathcal{A}$ aborts, then $\mathcal{S}$ outputs $\mathcal{A}$ 's output and halts.
4. In every round $i \in\left\{i^{*}, \ldots, r\right\}, \mathcal{S}$ sends $\mathcal{A}$ the share $a_{i^{*}+1}^{(2)}=a_{i^{*}+1}^{(1)} \oplus c$ (recall that $c$ is the value received from the trusted party computing the coin-flipping functionality). If $\mathcal{A}$ aborts, then $\mathcal{S}$ outputs $\mathcal{A}$ 's output and halts.
5. At the end of the protocol, $\mathcal{S}$ outputs $\mathcal{A}$ 's output and halts.

We now consider the joint distribution of $\mathcal{A}$ 's view and the output of the honest party $P_{2}$ in the ideal model and in the hybrid model. There are three cases to consider:

1. $\mathcal{A}$ aborts before round $i^{*}$. In this case, the distributions are identical: In both models, the view of the adversary is the sequence of shares and the sequence of messages up to the round in which $\mathcal{A}$ aborted, and the output of $P_{2}$ is a uniformly distributed bit which is independent of $\mathcal{A}$ 's view.
2. $\mathcal{A}$ aborts in round $i^{*}$. In this case, $\mathcal{A}$ 's view is identical in both models, but the distributions of $P_{2}$ 's output given $\mathcal{A}$ 's view are not identical. In the ideal model, $P_{2}$ outputs the random bit $c$ that was revealed to $\mathcal{A}$ by $\mathcal{S}$ in round $i^{*}$ (recall that $c$ is the bit received from the trusted party computing the coin-flipping functionality). In the hybrid model, however, the output of $P_{2}$ is the value $b_{i^{*}-1}$ which is a random bit that is independent of $\mathcal{A}$ 's view. Thus, in this case the statistical distance between the two distributions is $1 / 2$. However, this case occurs with probability at most $1 / r$ since in both models $i^{*}$ is independent of $\mathcal{A}$ 's view until this round (that is, the probability that $\mathcal{A}$ aborts in round $i^{*}$ is at most $1 / r$ ).
3. $\mathcal{A}$ aborts after round $i^{*}$ or does not abort. In this case, the distributions are identical: The output of $P_{2}$ is the same random bit that was revealed to $\mathcal{A}$ in round $i^{*}$.
Note that $\mathcal{A}$ 's view in the hybrid and ideal models is always identically distributed (no matter what strategy $\mathcal{A}$ uses to decide when to abort) and that the only difference is in the distribution of the honest party's output conditioned on $\mathcal{A}$ 's view. Specifically, for any particular round $i$ the event in which the adversary aborts in round $i$ occurs with exactly the same probability in both models. ${ }^{4}$ Thus, the above three cases imply that the statistical distance between the two distributions is at most $1 /(2 r)$.

Claim 3.2. In protocol CoinFlip ${ }_{r}$ there exists an efficient adversarial party $P_{1}^{*}$ that can bias the output of $P_{2}$ by $\frac{1-2^{-r}}{2 r}$.

Proof. Consider the adversarial party $P_{1}^{*}$ that completes the preprocessing phase and then halts in the first round $i \in\{1, \ldots, r\}$ for which $a_{i}=0$. We denote by Abort the random variable corresponding to the round in which $P_{1}^{*}$ aborts, where Abort $=\perp$ if $P_{1}^{*}$ does not abort. In addition, we denote by $c_{2}$ the random variable corresponding to the output bit of $P_{2}$. Notice that if $P_{1}^{*}$ aborts in round $j \leq i^{*}$, then $P_{2}$ outputs a random bit, and if $P_{1}^{*}$ does not abort, then $P_{2}$ always outputs 1. Therefore, for every $i \in\{1, \ldots, r\}$ it holds that

$$
\begin{aligned}
\operatorname{Pr}\left[c_{2}=1 \mid i^{*}=i\right]= & \sum_{j=1}^{i} \operatorname{Pr}\left[\text { Abort }=j \mid i^{*}=i\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=j \wedge i^{*}=i\right] \\
& +\operatorname{Pr}\left[\text { Abort }=\perp \mid i^{*}=i\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=\perp \wedge i^{*}=i\right] \\
= & \sum_{j=1}^{i} \operatorname{Pr}\left[a_{1}=\cdots=a_{j-1}=1, a_{j}=0\right] \\
& \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=j \wedge i^{*}=i\right]
\end{aligned}
$$

[^4]\[

$$
\begin{aligned}
& +\operatorname{Pr}\left[a_{1}=\cdots=a_{i}=1\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=\perp \wedge i^{*}=i\right] \\
= & \sum_{j=1}^{i} \frac{1}{2^{j}} \cdot \frac{1}{2}+\frac{1}{2^{i}} \cdot 1 \\
= & \frac{1}{2}+\frac{1}{2^{i+1}} .
\end{aligned}
$$
\]

This implies that

$$
\begin{aligned}
\operatorname{Pr}\left[c_{2}=1\right] & =\sum_{i=1}^{r} \operatorname{Pr}\left[i^{*}=i\right] \operatorname{Pr}\left[c_{2}=1 \mid i^{*}=i\right] \\
& =\sum_{i=1}^{r} \frac{1}{r}\left(\frac{1}{2}+\frac{1}{2^{i+1}}\right) \\
& =\frac{1}{2}+\frac{1-2^{-r}}{2 r}
\end{aligned}
$$

## 4. The Generalized Protocol

In this section, we present a more refined and generalized protocol that settles Theorems 1.1 and 1.2. The improvements over the protocol presented in Sect. 3 are as follows:
Improved security guarantee: In the simplified protocol, party $P_{1}$ can bias the output of party $P_{2}$ (by aborting in round $i^{*}$ ), but party $P_{2}$ cannot bias the output of party $P_{1}$. This is due to the fact that party $P_{1}$ always learns the output before party $P_{2}$ does. In the generalized protocol, the party that learns the output before the other party is chosen uniformly at random (i.e., party $P_{1}$ learns the output before party $P_{2}$ with probability $1 / 2$ ). This is achieved by having the parties exchange a sequence of $2 r$ values $\left(a_{1}, b_{1}\right), \ldots,\left(a_{2 r}, b_{2 r}\right)$ (using the same secret-sharing exchange technique as in the simplified protocol) with the following property: For odd values of $i$, party $P_{1}$ learns $a_{i}$ before party $P_{2}$ learns $b_{i}$, and for even values of $i$, party $P_{2}$ learns $b_{i}$ before party $P_{1}$ learns $a_{i}$. Thus, party $P_{1}$ can bias the result only when $i^{*}$ is odd, and party $P_{2}$ can bias the result only when $i^{*}$ is even. The key point is that the parties can exchange the sequence of $2 r$ shares in only $r+1$ rounds by combining some of their messages. ${ }^{5}$ Figure 3 gives a graphic overview of the protocol.

[^5]

Fig. 3. Overview of the generalized protocol.
We note that modifying the original protocol by just having ShareGen randomly choose which party starts would also halve the bias (since with probability $1 / 2$ the adversary chooses a party that cannot bias the outcome at all). However, this is vulnerable to a trivial dynamic attack: The adversary decides which party to corrupt after seeing which party was chosen to start.
A larger class of functionalities: We consider the more general task of sampling from a distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ : Party $P_{1}$ receives a sample from $\mathcal{D}_{1}$, and party $P_{2}$ receives a correlated sample from $\mathcal{D}_{2}$ (in coin flipping, for example, the joint distribution $\mathcal{D}$ produces the values $(0,0)$ and $(1,1)$ each with probability $1 / 2$ ). Our generalized protocol can handle any polynomially sampleable distribution $\mathcal{D}$.
In the following, we describe the generalized protocol Sampling ${ }_{r}$ (Sect. 4.1) and then prove its security (Sect. 4.2).

### 4.1. Description of the Protocol

Joint Input Security parameter $1^{n}$ and a polynomially sampleable distribution $\mathcal{D}=$ $\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$.

## Functionality ShareGen ${ }_{r}$

Input: Security parameter $1^{n}$, and a polynomially-sampleable distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$.
Computation:

1. Choose $i^{*} \in\{1, \ldots, 2 r\}$ uniformly at random.
2. Define values $a_{1}, \ldots, a_{2 r}$ and $b_{1}, \ldots, b_{2 r}$ as follows:

- For $1 \leq i \leq i^{*}-1$ sample $a_{i} \leftarrow \mathcal{D}_{1}$ and $b_{i} \leftarrow \mathcal{D}_{2}$ independently.
- Sample $\left(c^{(1)}, c^{(2)}\right) \leftarrow \mathcal{D}$, and for $i^{*} \leq i \leq r$ let $\left(a_{i}, b_{i}\right)=\left(c^{(1)}, c^{(2)}\right)$.

3. For $1 \leq i \leq 2 r$, choose $\left(a_{i}^{(1)}, a_{i}^{(2)}\right)$ and $\left(b_{i}^{(1)}, b_{i}^{(2)}\right)$ as random secret shares of $a_{i}$ and $b_{i}$, respectively.
4. Compute $k_{1}^{a}, \ldots, k_{2 r}^{a}, k_{1}^{b}, \ldots, k_{2 r}^{b} \leftarrow \operatorname{Gen}\left(1^{n}\right)$. For $1 \leq i \leq 2 r$, let $t_{i}^{a}=\operatorname{Mac}_{k_{i}^{a}}\left(i \| a_{i}^{(2)}\right)$ and $t_{i}^{b}=$ $\operatorname{Mac}_{k_{i}^{b}}\left(i \| b_{i}^{(1)}\right)$.
Output:
5. Party $P_{1}$ receives the values $a_{1}^{(1)}, \ldots, a_{2 r}^{(1)},\left(b_{1}^{(1)}, t_{1}^{b}\right), \ldots,\left(b_{2 r}^{(1)}, t_{2 r}^{b}\right)$, and $k^{a}=\left(k_{1}^{a}, \ldots, k_{2 r}^{a}\right)$.
6. Party $P_{2}$ receives the values $\left(a_{1}^{(2)}, t_{1}^{a}\right), \ldots,\left(a_{2 r}^{(2)}, t_{2 r}^{a}\right), b_{1}^{(2)}, \ldots, b_{2 r}^{(2)}$, and $k^{b}=\left(k_{1}^{b}, \ldots, k_{2 r}^{b}\right)$.

Fig. 4. The ideal functionality ShareGen $_{r}$ of the generalized two-party protocol.

## Preliminary phase:

1. Parties $P_{1}$ and $P_{2}$ run protocol $\pi$ for computing $\operatorname{ShareGen}_{r}\left(1^{n}, \mathcal{D}\right)$ (see Fig. 4).
2. If $P_{1}$ receives $\perp$ from the above computation, it outputs a random sample from $\mathcal{D}_{1}$ and halts. Likewise, if $P_{2}$ receives $\perp$, it outputs a random sample from $\mathcal{D}_{2}$ and halts. Otherwise, the parties proceed.
3. Denote the output of $P_{1}$ from $\pi$ by $a_{1}^{(1)}, \ldots, a_{2 r}^{(1)},\left(b_{1}^{(1)}, t_{1}^{b}\right), \ldots,\left(b_{2 r}^{(1)}, t_{2 r}^{b}\right)$ and $k_{1}^{a}, \ldots, k_{2 r}^{a}$.
4. Denote the output of $P_{2}$ from $\pi$ by $\left(a_{1}^{(2)}, t_{1}^{a}\right), \ldots,\left(a_{2 r}^{(2)}, t_{2 r}^{a}\right), b_{1}^{(2)}, \ldots, b_{2 r}^{(2)}$ and $k_{1}^{b}, \ldots, k_{2 r}^{b}$.

## In round 1 do:

1. $P_{2}$ sends a share to $P_{1}$ :
(a) $P_{2}$ sends $\left(a_{1}^{(2)}, t_{1}^{a}\right)$ to $P_{1}$.
(b) $P_{1}$ receives $\left(\hat{a}_{1}^{(2)}, \hat{t}_{1}^{a}\right)$ from $P_{2}$. If $\operatorname{Vrfy}_{k_{1}^{a}}\left(1 \| \hat{a}_{1}^{(2)}, \hat{t}_{1}^{a}\right)=0$, then $P_{1}$ outputs a random sample from $\mathcal{D}_{1}$ and halts. Otherwise, $P_{1}$ reconstructs $a_{1}$ using the shares $a_{1}^{(1)}$ and $\hat{a}_{1}^{(2)}$.
2. $P_{1}$ sends a pair of shares to $P_{2}$ :
(a) $P_{1}$ sends $\left(b_{1}^{(1)}, t_{1}^{b}\right)$ and $\left(b_{2}^{(1)}, t_{2}^{b}\right)$ to $P_{2}$.
(b) $P_{2}$ receives $\left(\hat{b}_{1}^{(1)}, \hat{t}_{1}^{b}\right)$ and $\left(\hat{b}_{2}^{(1)}, \hat{t}_{2}^{b}\right)$ from $P_{1}$. If $\operatorname{Vrfy}_{k_{1}^{b}}\left(1 \mid \hat{b}_{1}^{(1)}, \hat{t}_{1}^{b}\right)=0$, then $P_{2}$ outputs a random sample from $\mathcal{D}_{2}$ and halts. Otherwise, $P_{2}$ reconstructs $b_{1}$ using the shares $b_{1}^{(1)}$ and $b_{1}^{(2)}$.
(c) If $\mathrm{Vrfy}_{k_{2}^{b}}\left(2 \| \hat{b}_{2}^{(1)}, \hat{t}_{2}^{b}\right)=0$, then $P_{2}$ outputs $b_{1}$ and halts. Otherwise, $P_{2}$ reconstructs $b_{2}$ using the shares $b_{2}^{(1)}$ and $b_{2}^{(2)}$.

## In each round $\boldsymbol{j}=2, \ldots, r$ do:

1. $P_{2}$ sends a pair of shares to $P_{1}$ :
(a) $P_{2}$ sends $\left(a_{2 j-2}^{(2)}, t_{2 j-2}^{a}\right)$ and $\left(a_{2 j-1}^{(2)}, t_{2 j-1}^{a}\right)$ to $P_{1}$.
(b) $P_{1}$ receives $\left(\hat{a}_{2 j-2}^{(2)}, \hat{t}_{2 j-2}^{a}\right)$ and $\left(\hat{a}_{2 j-1}^{(2)}, \hat{t}_{2 j-1}^{a}\right)$ from $P_{2}$. If $V^{\prime} f y_{k_{2 j-2}^{a}}$ $\left(2 j-2 \| \hat{a}_{2 j-2}^{(2)}, \hat{t}_{2 j-2}^{a}\right)=0$, then $P_{1}$ outputs $a_{2 j-3}$ and halts. Otherwise, $P_{1}$ reconstructs $a_{2 j-2}$ using the shares $a_{2 j-2}^{(1)}$ and $\hat{a}_{2 j-2}^{(2)}$.
(c) If $\operatorname{Vrfy}_{k_{2 j-1}^{a}}\left(2 j-1| | \hat{a}_{2 j-1}^{(2)}, \hat{t}_{2 j-1}^{a}\right)=0$, then $P_{1}$ outputs $a_{2 j-2}$ and halts. Otherwise, $P_{1}$ reconstructs $a_{2 j-1}$ using the shares $a_{2 j-1}^{(1)}$ and $\hat{a}_{2 j-1}^{(2)}$.
2. $P_{1}$ sends a pair of shares to $P_{2}$ :
(a) $P_{1}$ sends $\left(b_{2 j-1}^{(1)}, t_{2 j-1}^{b}\right)$ and $\left(b_{2 j}^{(1)}, t_{2 j}^{b}\right)$ to $P_{2}$.
(b) $P_{2}$ receives $\left(\hat{b}_{2 j-1}^{(1)}, \hat{t}_{2 j-1}^{b}\right)$ and $\left(\hat{b}_{2 j}^{(1)}, \hat{t}_{2 j}^{b}\right)$ from $P_{1}$. If $\operatorname{Vrfy}_{k_{2 j-1}^{b}}$ $\left(2 j-1| | \hat{b}_{2 j-1}^{(1)}, \hat{t}_{2 j-1}^{b}\right)=0$, then $P_{2}$ outputs $b_{2 j-2}$ and halts. Otherwise, $P_{2}$ reconstructs $b_{2 j-1}$ using the shares $b_{2 j-1}^{(1)}$ and $b_{2 j-1}^{(2)}$.
(c) If $\mathrm{Vrfy}_{k_{2 j}^{b}}\left(2 j \| \hat{b}_{2 j}^{(1)}, \hat{t}_{2 j}^{b}\right)=0$, then $P_{2}$ outputs $b_{2 j-1}$ and halts. Otherwise, $P_{2}$ reconstructs $b_{2 j}$ using the shares $b_{2 j}^{(1)}$ and $b_{2 j}^{(2)}$.

## In round $\boldsymbol{r}+\mathbf{1}$ do:

1. $P_{2}$ sends a share to $P_{1}$ :
(a) $P_{2}$ sends $\left(a_{2 r}^{(2)}, t_{2 r}^{a}\right)$ to $P_{1}$.
(b) $P_{1}$ receives $\left(\hat{a}_{2 r}^{(2)}, \hat{t}_{2 r}^{a}\right)$ from $P_{2}$. If $\mathrm{Vrfy}_{k_{2 r}^{a}}\left(2 r \| \hat{a}_{2 r}^{(2)}, \hat{t}_{2 r}^{a}\right)=0$, then $P_{1}$ outputs $a_{2 r-1}$ and halts. Otherwise, $P_{1}$ reconstructs $a_{2 r}$ using the shares $a_{2 r}^{(1)}$ and $\hat{a}_{2 r}^{(2)}$.
2. $P_{1}$ and $P_{2}$ output the values $a_{2 r}$ and $b_{2 r}$, respectively, and halt.

### 4.2. Proof of Security

We remind the reader that $\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)$ denotes the statistical distance between the joint distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ and the direct product of the two marginal distributions $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$. We prove the following theorem which implies Theorems 1.1 and 1.2:

Theorem 4.1. For any polynomially sampleable distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ and polynomial $r=r(n)$, if protocol $\pi$ securely computes ShareGen $_{r}$ with abort, then Sampling ${ }_{r}$ is $\frac{\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)}{2 r}$-secure.

Proof. As in the proof of Theorem 3.1, we prove the security of the protocol in a hybrid model where a trusted party for computing ShareGen $_{r}$ with abort is available. For every polynomial-time hybrid-model adversary $\mathcal{A}$ corrupting $P_{1}$ and running Sampling ${ }_{r}$ in the hybrid model, we show that there exists a polynomial-time ideal-model adversary $\mathcal{S}$ corrupting $P_{1}$ in the ideal model with access to a trusted party that samples from $\mathcal{D}$
such that the statistical distance between these two executions is at most $\frac{\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)}{2 r}+$ $\nu(n)$, for some negligible function $v(n)$. The proof for the case that $P_{2}$ is corrupted is essentially identical and therefore is omitted. For simplicity, we ignore the aspect of message authentication and assume that the only malicious behavior of $\mathcal{A}$ is early abort. This does not result in any loss of generality, since there is only a negligible probability of forging an authentication tag.

On input ( $1^{n}$, aux), the ideal-model adversary $\mathcal{S}$ invokes the hybrid-model adversary $\mathcal{A}$ on ( $1^{n}$, aux) and queries the trusted party who sends to the parties a sample $\left(c_{1}, c_{2}\right)$ drawn from the joint distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$. At this point, $\mathcal{S}$ receives the value $c_{1}$ which was sent to the corrupted $P_{1} . \mathcal{S}$ simulates the trusted party computing the ShareGen ${ }_{r}$ functionality by sending $\mathcal{A}$ independently and uniformly chosen shares $a_{1}^{(1)}, \ldots, a_{2 r}^{(1)}, b_{1}^{(1)}, \ldots, b_{2 r}^{(1)}$. If $\mathcal{A}$ aborts at this point, then $\mathcal{S}$ outputs $\mathcal{A}$ 's output and halts. Otherwise, $\mathcal{S}$ chooses $i^{*} \in\{1, \ldots, 2 r\}$ uniformly at random and proceeds by sending $\mathcal{A}$ shares $a_{1}^{(2)}, \ldots, a_{i^{*}+1}^{(2)}$ (in the order defined by the rounds of the protocol), where the shares are defined as follows:

1. For every $i \in\left\{1, \ldots, i^{*}-1\right\}, \mathcal{S}$ samples a random $a_{i} \leftarrow \mathcal{D}_{1}$ and sets $a_{i}^{(2)}=$ $a_{i}^{(1)} \oplus a_{i}$.
2. For every $i \leq\left\{i^{*}, \ldots, r\right\}, \mathcal{S}$ sets $a_{i^{*}}^{(2)}=a_{i^{*}}^{(1)} \oplus c_{1}$ where $c_{1}$ is the value received from the trusted party.

If at some point during the simulation $\mathcal{A}$ aborts, then $\mathcal{S}$ outputs $\mathcal{A}$ 's output and halts.
We now consider the joint distribution of the adversaries view and the output of the honest party $P_{2}$ in the ideal model and in the hybrid model, and show that the statistical distance between the two distributions is at most $\frac{\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)}{2 r}$. As in the proof of Theorem 3.1, note that the adversary's view is always identically distributed in both cases, and therefore, we only need to consider the distribution of $P_{2}$ 's output given the adversary's view. There are two cases to consider, depending on whether $i^{*}$ is even or odd.
Case 1: $\boldsymbol{i}^{*}=\mathbf{2} \boldsymbol{j}^{*}$ for some $\boldsymbol{j}^{*} \in\{\mathbf{1}, \ldots, r\}$. In this case, $P_{2}$ learns its output in round $j^{*}$ and $P_{1}$ learns its output in round $j^{*}+1$, and we show that the two distributions are identical. There are two cases to consider:

1. $\mathcal{A}$ aborts before round $j^{*}+1$. In both models, if $\mathcal{A}$ aborts before round $j^{*}+1$, then he does not receive the share $a_{i^{*}}^{(2)}=a_{2 j^{*}}^{(2)}$ since this share is sent by $P_{2}$ only in round $j^{*}+1$. Therefore, $\mathcal{A}$ 's view is independent of $P_{2}$ 's output.
2. $\mathcal{A}$ aborts in round $j^{*}+1$ or does not abort. In this case, in both models $\mathcal{A}$ learns $c_{1}$ and $P_{2}$ outputs $c_{2}$, where $\left(c_{1}, c_{2}\right)$ are sampled from the joint distribution $\mathcal{D}=$ ( $\mathcal{D}_{1}, \mathcal{D}_{2}$ ).

Case 2: $\boldsymbol{i}^{*}=2 \boldsymbol{j}^{*}-\mathbf{1}$ for some $\boldsymbol{j}^{*} \in\{\mathbf{1}, \ldots, r\}$. In this case, both parties learn their outputs in round $j^{*}$, but $P_{1}$ learns its output first. Informally, $P_{1}$ can bias $P_{2}$ 's output only by aborting in round $j^{*}$ after receiving $P_{2}$ 's message for this round. More formally, there are three cases to consider:

1. $\mathcal{A}$ aborts before round $j^{*}$. In this case, the distributions are identical: In both models, the view of the adversary is the sequence of shares and the sequence of
messages up to the round in which $\mathcal{A}$ aborted, and the output of $P_{2}$ is a random sample from $\mathcal{D}_{2}$ that is independent of $\mathcal{A}$ 's view.
2. $\mathcal{A}$ aborts in round $j^{*}$. In this case, $\mathcal{A}$ 's view is identical in both models, but the distributions of $P_{2}$ 's output given $\mathcal{A}$ 's view are not identical. In the ideal model, $P_{2}$ outputs the value $c_{2}$ that is correlated with the value $c_{1}$ that was revealed to $\mathcal{A}$ by $\mathcal{S}$ in round $i^{*}$ (i.e., $\left(c_{1}, c_{2}\right)$ is sampled from the joint distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ ). In the hybrid model, however, the output of $P_{2}$ is the value $b_{i^{*}-1}$ which is a random sample from $\mathcal{D}_{2}$ that is independent of $\mathcal{A}$ 's view. Thus, in this case, the statistical distance between the two distributions is $\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)$. However, this case occurs with probability at most $1 / 2 r$ since in both cases $i^{*}$ is odd with probability exactly $1 / 2$ and is independent of $\mathcal{A}$ 's view until this round (that is, the probability that $\mathcal{A}$ aborts in round $j^{*}$ is at most $1 / r$ ).
3. $\mathcal{A}$ aborts in round $j^{*}+1$ or does not abort. In this case, the distribution are identical: In both models, $\mathcal{A}$ learns $c_{1}$ and $P_{2}$ outputs $c_{2}$, where $\left(c_{1}, c_{2}\right)$ are sampled from the joint distribution $\mathcal{D}=\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$.
This implies that the statistical distance between the two distributions is at most $\frac{\operatorname{SD}\left(\mathcal{D}, \mathcal{D}_{1} \otimes \mathcal{D}_{2}\right)}{2 r}$ and concludes the proof of the theorem.

We conclude this section by showing that Theorem 4.1 is tight for the coin-flipping functionality: There exists an efficient adversary that can bias the output of the honest party by essentially $1 /(4 r)$. This adversary is a natural generalization of the adversary presented at the end of Sect. 3.

Claim 4.2. In protocol Sampling $_{r}$ instantiated with the distribution $\mathcal{D}$ that outputs the values $(0,0)$ and $(1,1)$ each with probability $1 / 2$, there exists an efficient adversarial party $P_{1}^{*}$ that can bias the output of $P_{2}$ by $\frac{1-2^{-r}}{4 r}$.

Proof. Consider the adversarial party $P_{1}^{*}$ that completes the preprocessing phase and then halts in the first round $j \in\{1, \ldots, r\}$ for which $a_{2 j-1}=0$. We denote by Abort the random variable corresponding to the round in which $P_{1}^{*}$ aborts, where Abort $=\perp$ if $P_{1}^{*}$ does not abort. In addition, we denote by $c_{2}$ the random variable corresponding to the output bit of $P_{2}$. Notice that if $i^{*}$ is even, then $P_{2}$ outputs 1 with probability $1 / 2$. Now, suppose that $i^{*}=2 j^{*}-1$ for some $j^{*} \in\{1, \ldots, r\}$, then there are two cases to consider:

- If $P_{1}^{*}$ aborts in round $j \leq j^{*}$, then $P_{2}$ outputs a random bit.
- If $P_{1}^{*}$ does not abort, then $P_{2}$ always outputs 1 .

Therefore, for every $j^{*} \in\{1, \ldots, r\}$, it holds that

$$
\begin{aligned}
\operatorname{Pr} & {\left[c_{2}=1 \mid i^{*}=2 j^{*}-1\right] } \\
= & \sum_{j=1}^{j^{*}} \operatorname{Pr}\left[\text { Abort }=j \mid i^{*}=2 j^{*}-1\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=j \wedge i^{*}=2 j^{*}-1\right] \\
& +\operatorname{Pr}\left[\text { Abort }=\perp \mid i^{*}=2 j^{*}-1\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=\perp \wedge i^{*}=2 j^{*}-1\right]
\end{aligned}
$$

$$
\begin{aligned}
= & \sum_{j=1}^{j^{*}} \operatorname{Pr}\left[a_{1}=\cdots=a_{2 j-3}=1, a_{2 j-1}=0\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=j \wedge i^{*}=2 j^{*}-1\right] \\
& +\operatorname{Pr}\left[a_{1}=\cdots=a_{2 j^{*}-3}=a_{2 j^{*}-1}=1\right] \operatorname{Pr}\left[c_{2}=1 \mid \text { Abort }=\perp \wedge i^{*}=2 j^{*}-1\right] \\
= & \sum_{j=1}^{j^{*}} \frac{1}{2^{j}} \cdot \frac{1}{2}+\frac{1}{2^{j^{*}}} \cdot 1 \\
= & \frac{1}{2}+\frac{1}{2^{j^{*}+1}}
\end{aligned}
$$

This implies that

$$
\begin{aligned}
\operatorname{Pr}\left[c_{2}=1\right]= & \operatorname{Pr}\left[i^{*} \text { is even }\right] \operatorname{Pr}\left[c_{2}=1 \mid i^{*} \text { is even }\right] \\
& +\sum_{j^{*}=1}^{r} \operatorname{Pr}\left[i^{*}=2 j^{*}-1\right] \operatorname{Pr}\left[c_{2}=1 \mid i^{*}=2 j^{*}-1\right] \\
= & \frac{1}{2} \cdot \frac{1}{2}+\sum_{j^{*}=1}^{r} \frac{1}{2 r}\left(\frac{1}{2}+\frac{1}{2^{j^{*}+1}}\right) \\
= & \frac{1}{2}+\frac{1-2^{-r}}{4 r}
\end{aligned}
$$

## 5. Open Problems

Identifying the Minimal Computational Assumptions Blum's coin-flipping protocol, as well as its generalization that guarantees bias of $O(1 / \sqrt{r})$, can rely on the existence of any one-way function. We showed that the optimal trade-off between the round complexity and the bias can be achieved assuming the existence of oblivious transfer, a complete primitive for secure computation. A challenging problem is to either achieve the optimal bias based on seemingly weaker assumptions (e.g., one-way functions) or demonstrate that oblivious transfer is in fact essential. Although meaningful progress was recently made in this direction for restricted black-box constructions by Dachman-Soled et al. [16,17] (see Sect. 1.3), the general problem is still left open.

Identifying the Exact Trade-Off The bias of our protocol almost exactly matches Cleve's lower bound: Cleve showed that any $r$-round protocol has bias at least $1 /(8 r+2)$, and we manage to achieve bias of at most $1 /(4 r-c)$ for some constant $c>0$. It will be interesting to eliminate the multiplicative gap of $1 / 2$ by either improving Cleve's lower bound or improving our upper bound. We note, however, that this cannot be resolved by improving the security analysis of our protocol since there exists an efficient adversary that can bias our protocol by essentially $1 /(4 r)$ (see Sect. 4), and therefore, our analysis is tight.

Efficient Implementation Our protocol uses a general secure computation step in the preprocessing phase. Although asymptotically optimal, the techniques used in general secure computation often have a large overhead. Hence, it would be helpful to find an efficient sub-protocol to compute the ShareGen functionality that can lead to a $^{\text {a }}$ practical implementation.

The Multiparty Setting Blum's coin-flipping protocol can be extended to an $m$-party $r$-round protocol that has bias $O(m / \sqrt{r})$, and several exciting results for restricted cases of the multiparty setting were recently achieved by Beimel et al. [7] and by Haitner and Tsfadia [24] (see Sect. 1.3). However, the general problem of identifying the optimal trade-off between the number of parties, the round complexity, and the bias is still left open.

## References

[1] D. Aharonov, A. Ta-Shma, U.V. Vazirani, A.C. Yao, Quantum bit escrow, in Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (2000), pp. 705-714
[2] N. Alon, M. Naor, Coin-flipping games immune against linear-sized coalitions. SIAM J. Comput. 22(2), 403-417 (1993)
[3] A. Ambainis, A new protocol and lower bounds for quantum coin flipping. J. Comput. Syst. Sci. 68(2), 398-416 (2004)
[4] A. Ambainis, H. Buhrman, Y. Dodis, H. Rohrig, Multiparty quantum coin flipping, in Proceedings of the 19th Annual IEEE Conference on Computational Complexity (2004), pp. 250-259
[5] B. Averbuch, M. Blum, B. Chor, S. Goldwasser, S. Micali, How to implement Bracha's $O(\log n)$ byzantine agreement algorithm. Manuscript (1985)
[6] A. Beimel, Y. Lindell, E. Omri, I. Orlov. 1/p-secure multiparty computation without honest majority and the best of both worlds, in Advances in Cryptology-CRYPTO '11 (2011), pp. 277-296
[7] A. Beimel, E. Omri, I. Orlov, Protocols for multiparty coin toss with dishonest majority, in Advances in Cryptology-CRYPTO '10 (2010), pp. 538-557
[8] M. Bellare, P. Rogaway, Code-based game-playing proofs and the security of triple encryption. Cryptology ePrint Archive, Report 2004/331 (2004). http://eprint.iacr.org/2004/331.pdf
[9] M. Ben-Or, S. Goldwasser, A. Wigderson, Completeness theorems for non-cryptographic fault-tolerant distributed computation, in Proceedings of the 20th Annual ACM Symposium on Theory of Computing (1988), pp. 1-10
[10] M. Ben-Or, N. Linial, Collective coin flipping. Adv. Comput. Res. Randomness Comput. 5, 91-115 (1989)
[11] I. Berman, I. Haitner, A. Tentes, Coin flipping of any constant bias implies one-way functions, in Proceedings of the 46th Annual ACM Symposium on Theory of Computing (2014), pp. 398-407
[12] M. Blum, Coin flipping by telephone-A protocol for solving impossible problems, in Proceedings of the 25th IEEE Computer Society International Conference (1982), pp. 133-137
[13] R. Canetti, Security and composition of multiparty cryptographic protocols. J. Cryptol. 13(1), 143-202 (2000)
[14] R. Cleve, Limits on the security of coin flips when half the processors are faulty, in Proceedings of the 18th Annual ACM Symposium on Theory of Computing (1986). pp. 364-369
[15] R. Cleve, R. Impagliazzo, Martingales, collective coin flipping and discrete control processes. http:// www.cpsc.ucalgary.ca/~cleve/pubs/martingales.ps (1993)
[16] D. Dachman-Soled, Y. Lindell, M. Mahmoody, T. Malkin, On the black-box complexity of optimally-fair coin tossing, in Proceedings of the 8th Theory of Cryptography Conference (2011), pp. 450-467
[17] D. Dachman-Soled, M. Mahmoody, T. Malkin, Can optimally-fair coin tossing be based on one-way functions? in Proceedings of the 11th Theory of Cryptography Conference (2014), pp. 217-239
[18] U. Feige, Noncryptographic selection protocols, in Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (1999), pp. 142-153
[19] O. Goldreich, Foundations of Cryptography—Volume 2: Basic Applications. (Cambridge University Press, Cambridge, 2004)
[20] S.D. Gordon, C. Hazay, J. Katz, Y. Lindell, Complete fairness in secure two-party computation, in Proceedings of the 40th Annual ACM Symposium on Theory of Computing (2008), pp. 413-422
[21] S.D. Gordon, J. Katz, Rational secret sharing, revisited, in Proceedings on the 5th International Conference on Security and Cryptography for Networks (2006), pp. 229-241
[22] S.D. Gordon, J. Katz, Partial fairness in secure two-party computation. J. Cryptol. 25(1), 14-40 (2012)
[23] I. Haitner, E. Omri, Coin flipping with constant bias implies one-way functions. SIAM J. Comput. 43(2), 389-409 (2014)
[24] I. Haitner, E. Tsfadia, An almost-optimally fair three-party coin-flipping protocol, in Proceedings of the 46th Annual ACM Symposium on Theory of Computing (2014), pp. 408-416
[25] J.Y. Halpern, V. Teague, Rational secret sharing and multiparty computation, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing (2004), pp. 623-632
[26] J. Håstad, R. Impagliazzo, L.A. Levin, M. Luby, A pseudorandom generator from any one-way function. SIAM J. Comput. 28(4), 1364-1396 (1999)
[27] C. Hazay, Y. Lindell, Efficient Secure Two-Party Protocols-Techniques and Constructions (Springer, Berlin, 2010)
[28] R. Impagliazzo, M. Luby, One-way functions are essential for complexity based cryptography, in Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science (1989), pp. 230-235
[29] J. Katz, On achieving the "best of both worlds" in secure multiparty computation, in Proceedings of the 39th Annual ACM Symposium on Theory of Computing (2007), pp. 11-20
[30] G. Kol, M. Naor, Cryptography and game theory: Designing protocols for exchanging information, in Proceedings of the 5th Theory of Cryptography Conference (2008), pp. 320-339.
[31] Y. Lindell, Parallel coin-tossing and constant-round secure two-party computation. J. Cryptol. 16(3), 143-184 (2003)
[32] H.K. Maji, M. Prabhakaran, A. Sahai, On the computational complexity of coin flipping, in Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (2010), pp. 613-622
[33] T. Moran, M. Naor, Basing cryptographic protocols on tamper-evident seals, in Proceedings of the 32nd International Colloquium on Automata, Languages and Programming, (2005), pp. 285-297
[34] M. Naor, Bit commitment using pseudorandomness. J. Cryptol. 4(2), 151-158 (1991)
[35] A. Russell, D. Zuckerman, Perfect information leader election in $\log ^{*} n+O$ (1) rounds. J. Comput. Syst. Sci. 63(4), 612-626 (2001)
[36] M. Saks, A robust noncryptographic protocol for collective coin flipping. SIAM J. Discrete Math. 2(2), 240-244 (1989)
[37] M.N. Wegman, L. Carter, New hash functions and their use in authentication and set equality. J. Comput. Syst. Sci. 22(3), 265-279 (1981)


[^0]:    *This submission was solicited from TCC 2009.
    $\dagger$ Supported by the European Union's Seventh Framework Programme (FP7) via a Marie Curie Career Integration Grant and by the Israel Science Foundation (Grant No. 1790/13). Most of this work was done while the author was a PhD student at the Weizmann Institute of Science.
    $\ddagger$ Research supported in part by a grant from the I-CORE Program of the Planning and Budgeting Committee, the Israel Science Foundation, BSF, and the Israeli Ministry of Science and Technology.
    § Supported by the European Union's Seventh Framework Programme (FP7) via a Marie Curie Career Integration Grant, by the Israel Science Foundation (Grant No. 483/13), and by the Israeli Centers of Research Excellence (I-CORE) Program (Center No. 4/11). Most of this work was done while the author was a PhD student at the Weizmann Institute of Science.

[^1]:    ${ }^{1}$ In a preliminary version of this work, we proved our results with respect to the definition of bias (see Sect. 2), and motivated by [22,29], we switch to the more general framework of $1 / p$-secure computation.

[^2]:    ${ }^{2}$ Cleve's lower bound [14] holds under this relaxation as well. Specifically, if the parties agree on a common value with probability $1 / 2+\epsilon$, then Cleve's proof shows that the protocol has bias at least $\epsilon /(4 r+1)$.

[^3]:    ${ }^{3}$ An alternative approach that reduces the expected number of rounds from $r$ to $r / 2+1$ is as follows. In round $i^{*}$, the parties learn their output $c=a_{i^{*}}=b_{i^{*}}$, and in round $i^{*}+1$, they learn a special value $a_{i^{*}+1}=b_{i^{*}+1}=$ NULL indicating that they should output the value from the previous round and halt. For simplicity (both in the presentation of the protocol and in the proof of security), we choose to present the protocol as always having $r$ rounds, but this is not essential for our results.

[^4]:    ${ }^{4}$ We state this explicitly because in some settings (that are inherently different from our setting), analyzing the statistical distance between executions in two models by conditioning on a specific event might be problematic when that event does not occur with equal probabilities in both models (see for example, [8, Sect. 2]).

[^5]:    ${ }^{5}$ Recall that each round consists of two moves: a message from $P_{2}$ to $P_{1}$ followed by a message from $P_{1}$ to $P_{2}$.

