

COIN FLIPPING WITH CONSTANT BIAS IMPLIES ONE-WAY FUNCTIONS*

IFTACH HAITNER[†] AND ERAN OMRI[‡]

Abstract. It is well known (cf. Impagliazzo and Luby [in *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science*, 1989, pp. 230–235]) that the existence of almost all “interesting” cryptographic applications, i.e., ones that cannot hold information theoretically, implies one-way functions. An important exception where the above implication is not known, however, is the case of coin-flipping protocols. Such protocols allow honest parties to mutually flip an unbiased coin, while guaranteeing that even a cheating (efficient) party cannot bias the output of the protocol by much. Impagliazzo and Luby proved that coin-flipping protocols that are safe against negligible bias do imply one-way functions, and, very recently, Maji, Prabhakaran, and Sahai [in *Proceedings of the 2001 51st Annual IEEE Symposium on Foundations of Computer Science*, 2010, pp. 613–622] proved the same for constant-round protocols (with any nontrivial bias). For the general case, however, no such implication was known. We make progress towards answering the above fundamental question, showing that (strong) coin-flipping protocols safe against a constant bias (concretely, $\frac{\sqrt{2}-1}{2} - o(1)$) imply one-way functions.

Key words. coin-flipping protocols, one-way functions

AMS subject classifications. 94A60, 68Q99

DOI. 10.1137/120887631

1. Introduction. A central focus of modern cryptography has been to investigate the weakest possible assumptions under which various cryptographic primitives exist. This direction of research has been quite fruitful, and minimal assumptions are known for a wide variety of primitives. In particular, it has been shown that one-way functions (i.e., easy-to-compute but hard-to-invert functions) imply pseudorandom generators, pseudorandom functions, symmetric-key encryption/message authentication, commitment schemes, and digital signatures [10, 11, 16, 14, 24, 25, 12, 26], where one-way functions (OWFs) were also shown to be implied by each of these primitives [19].

An important exception for which we have failed to prove the above rule is that of coin-flipping protocols. A coin-flipping protocol [4] allows the honest parties to mutually flip an unbiased coin, where even a cheating (efficient) party cannot bias the outcome of the protocol by much. While one-way functions are known to imply coin-flipping protocols [4, 24, 16], the other direction is less clear: Impagliazzo and Luby [19] showed that *negligible*-bias coin-flipping protocols (i.e., such that an efficient cheating strategy cannot make the common output to be 1, or to be 0, with

*Received by the editors August 24, 2012; accepted for publication (in revised form) November 6, 2013; published electronically March 4, 2014. A preliminary version of this paper appeared in *Proceedings of the 2011 52nd Annual IEEE Symposium on Foundations of Computer Science*, 2011, pp. 110–119 [15].

<http://www.siam.org/journals/sicomp/43-2/88763.html>

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel (iftachh@cs.tau.ac.il). This author’s research was partly supported by the Check Point Institute for Information Security, the Israeli Centers of Research Excellence (I-CORE) program (Center 4/11), Israel Science Foundation (grant 1076/11), and US-Israel BSF (grant 2010196).

[‡]Department of Computer Science and Mathematics, Ariel University, Ariel 40700, Israel (omrier@ariel.ac.il). This author’s research was supported by the European Research Council as part of the ERC project LAST. Most of the work was done while this author was at Bar-Ilan University.

probability greater than $\frac{1}{2} + \text{neg}(n)$ imply one-way functions. Very recently, Maji, Prabhakaran, and Sahai [21] proved the same implication for $(\frac{1}{2} - 1/\text{poly}(n))$ -bias *constant-round* protocols, where n is the security parameter of the protocol. We have no such implications, however, for any other choice of parameters.

1.1. Our result. In this work, we make progress towards answering the above fundamental question, showing that (strong) coin-flipping protocols, safe against a constant bias, imply one-way functions. We note that previous works [19, 21] also applied to weak coin-flipping protocols (in some restricted setting).¹ Our result, however, does not apply to weak coin-flipping protocols, and we consider only strong coin-flipping protocols (see section 6 for further discussion). We prove the following theorem.

THEOREM 1 (informal). *The existence of a $(\frac{\sqrt{2}-1}{2} - o(1))$ -bias coin-flipping protocol (of any round complexity) implies one-way functions.*

1.2. Related results. As mentioned above, Impagliazzo and Luby [19] showed that negligible-bias coin-flipping protocols imply one-way functions, and Maji, Prabhakaran, and Sahai [21] proved the same for $(\frac{1}{2} - 1/\text{poly}(n))$ -bias constant-round protocols. The authors of [21] also proved that $(\frac{1}{4} - o(1))$ -bias coin-flipping protocols imply that $\text{BPP} \neq \text{NP}$. Finally, it is well known that $(\frac{1}{2} - v(n))$ -bias coin-flipping protocols, for any $v(n) > 0$, imply that $\text{BPP} \neq \text{PSPACE}$. All of the above results extend to *weak* coin-flipping protocols: in such protocols, each party has a *different* predetermined value towards which it cannot bias the output coin.² A quick overview of the techniques underlying the above results can be found in section 1.3.3. A summary of the known results is given in Table 1.

Information theoretic coin-flipping protocols (i.e., those whose security holds against all powerful adversaries) were shown to exist in the quantum world; Mochon [22] presents an ε -bias quantum weak coin-flipping protocol for any $\varepsilon > 0$. Chailloux and Kerenidis [5] present a $(\frac{\sqrt{2}-1}{2} - \varepsilon)$ -bias quantum strong coin-flipping protocol for any $\varepsilon > 0$ (which is optimal [20]). A key step in [5] is a reduction from strong to weak coin-flipping protocols, which also holds in the classical world (see section 6 for further discussion). Multiparty coin flipping in the full information model was considered by Ben-Or and Linial [3]. They gave both upper and lower bounds on the bias that a single party can inflict on a coin-flipping protocol.

A related line of work considers *fair* coin-flipping protocols. In this setting the honest party is required to always output a bit, whatever the other party does. In particular, a cheating party might bias the output coin just by aborting. We know that one-way functions imply fair $(1/\sqrt{m})$ -bias coin-flipping protocols [1, 7], where m is the round complexity of the protocols, and this quantity is known to be tight for $O(n/\log n)$ -round protocols with fully black-box reductions [9], where n is the input length of the one-way function. Oblivious transfer, on the other hand, implies fair $1/m$ -bias protocols [23, 2], which is known to be tight [7].

1.3. Our technique. Let (A, B) be a balanced coin-flipping protocol (i.e., the common output of the honest parties is a uniformly chosen bit), and let f be the

¹See section 2.2 for the definition of (strong) coin-flipping protocols and weak coin-flipping protocols.

²While such protocols are strictly weaker than full-fledged coin-flipping protocols, they are still useful in many settings, for instance, when Alice and Bob are trying to decide who is doing the dishes.

TABLE 1
Results summary.

Coin-flipping protocol type	Implication	Paper
$(\frac{1}{4} - o(1))$ -bias	$BPP \neq NP$	Maji et al. [21]
Nontrivial bias, <i>constant round</i>	$BPP \neq NP$	Zachos [27]
Nontrivial bias	$BPP \neq PSPACE$	Folklore
Negligible bias	Existence of OWFS	Impagliazzo and Luby [19]
$(\frac{1}{2} - o(1))$ -bias, <i>constant round</i>	Existence of OWFS	Maji et al. [21]
$(\frac{\sqrt{2}-1}{2} - o(1))$ -bias, <i>strong coin flipping</i>	Existence of OWFS	<i>This work</i> ³

following efficiently computable function:

$$f(r_A, r_B, i) = \text{Trans}(r_A, r_B)_i, \text{Out}(r_A, r_B),$$

where r_A and r_B are the random coins of A and B , respectively, $\text{Trans}(r_A, r_B)_i$ is the first i messages exchanged in the execution $(A(r_A), B(r_B))$, and $\text{Out}(r_A, r_B)$ is the common output of this execution (i.e., the coin). Assuming that one-way functions do not exist, it follows that distributional one-way functions do not exist either [19], and therefore there exists an efficient inverter Inv that, given a random output y of f , samples a random preimage of y . Concretely, for any $p \in \text{poly}$ there exists a PPT Inv such that the following holds:

$$(1) \quad \text{SD}((X, f(X)), (\text{Inv}(f(X')), f(X'))) \leq 1/p(|X|),$$

where X and X' are uniformly distributed over the domain of f , and SD stands for statistical distance. In the following we show how to use the above Inv to bias the output of (A, B) .

Note that, given a random partial transcript \mathbf{t} of (A, B) , the call $\text{Inv}(\mathbf{t}, 1)$ returns a random pair of random coins for the parties that is (1) consistent with \mathbf{t} , and (2) yields a common output 1. In other words, one can use Inv to sample a random continuation of \mathbf{t} which leads to a 1-leaf—a full transcript of (A, B) in which the common output is 1. As we show below, such capability is an extremely useful tool for a dishonest party trying to bias the outcome of this protocol. In particular, we consider the following cheating strategy \mathcal{A} for A (a cheating strategy \mathcal{B} for B is analogously defined): given that the partial transcript is \mathbf{t} , \mathcal{A} uses Inv to sample a pair of random coins (r_A, r_B) that is consistent with \mathbf{t} and leads to a 1-leaf (\mathcal{A} aborts if Inv fails to provide such coins), and then acts as the honest A does on the random coins r_A , given the transcript \mathbf{t} . Namely, at each of its turns \mathcal{A} takes the first step of a random continuation that leads to a 1-leaf.

Assuming that Inv behaves as its ideal variant that returns a uniform random preimage on *any* transcript, it is not that hard to prove (see outline in section 1.3.1) that either \mathcal{A} or \mathcal{B} can significantly bias the outcome of the protocol. Proving that the same holds with respect to the real inverter, however, is not trivial. Algorithm Inv is guaranteed to work well only on *random* transcript/output pairs, as induced by a random output of f (namely, a transcript/output pair defined by a random *honest* execution of (A, B)). A random execution of (\mathcal{A}, B) or of (A, \mathcal{B}) (i.e., with one party

³As mentioned above, while previous results also apply to weak coin flipping, our result applies only to strong coin flipping.

being controlled by the adversary) might, however, generate a query distribution that is very far from that induced by f .

Fortunately, we manage to prove (and this is the crux of our proof; see outline in section 1.3.2) that the following holds: we call a query *nontypical* if its probability mass with respect to the execution of $(\mathcal{A}, \mathcal{B})$ (or of (\mathbf{A}, \mathbf{B})) is much larger than its mass with respect to the output distribution of f . We first show that even if both \mathcal{A} and \mathcal{B} totally fail on such nontypical queries, then either \mathcal{A} or \mathcal{B} can significantly bias the outcome of the protocol assuming access to the ideal sampler. Since on typical queries the real sampler should perform almost as well as its ideal version, we conclude that the cheating probability of either \mathcal{A} or \mathcal{B} is high, also when the cheating strategies are using the real sampler.

1.3.1. When using the ideal sampler. Consider a mental experiment in which the cheating strategies \mathcal{A} and \mathcal{B} (both using the ideal sampler) are interacting with each other. It is not hard to see that the common output of $(\mathcal{A}, \mathcal{B})$ in this case is always 1. Moreover, the transcript distribution induced by such an execution is that of a random execution of the “honest” protocol (\mathbf{A}, \mathbf{B}) conditioned so that the common output is 1 (i.e., a random 1-leaf). In particular, the probability of each 1-leaf in a random execution of $(\mathcal{A}, \mathcal{B})$ is twice its probability in (\mathbf{A}, \mathbf{B}) .

The probability of a 1-leaf \mathbf{t} happening is the product of the probabilities that in each stage of the protocol the relevant party sends the “right” message. Such a product can be partitioned into two parts: the part corresponding to the actions of \mathcal{A} , and the part corresponding to the actions of \mathcal{B} . In particular, either \mathcal{A} or \mathcal{B} contributes a factor of value at least $\sqrt{2}$ to the probability of \mathbf{t} . Namely, the probability of a 1-leaf \mathbf{t} in either $(\mathcal{A}, \mathcal{B})$ or in (\mathbf{A}, \mathbf{B}) is $\sqrt{2}$ times its probability in (\mathbf{A}, \mathbf{B}) . Summing over all 1-leaves, it follows that the common output of either $(\mathcal{A}, \mathcal{B})$ or (\mathbf{A}, \mathbf{B}) is 1 with probability at least $\sqrt{2} \cdot \frac{1}{2} = 1/\sqrt{2}$. That is, either \mathcal{A} or \mathcal{B} can bias the output of (\mathbf{A}, \mathbf{B}) by at least $\frac{1}{\sqrt{2}} - \frac{1}{2} = \frac{\sqrt{2}-1}{2}$.

The above intuition can be made formal in proving the success probability of the cheating strategies with respect to the ideal sampler. In the actual treatment, given in section 4, however, we use a slightly different approach and prove this result inductively up the protocol tree. This inductive approach is useful when considering the attacks with respect to the real sampler. A similar approach was used by Ben-Or and Linial [3] to bound from below the success probability of the best attack strategy in any multiparty coin-flipping protocol in the full information model. Below, we mention that the mental experiment in which \mathcal{A} and \mathcal{B} interact with each other does give the intuition for what we call the “compensation lemma” (see Lemma 19), which plays a crucial role in proving the success of the attacks when using the *real* sampler.

1.3.2. Using the real sampler. By our earlier discussion, it suffices to prove that the following holds: either \mathcal{A} or \mathcal{B} can significantly bias the output of the protocol when given access to the ideal sampler, even if both cheating strategies are assumed to fail completely when asking nontypical queries. To this end, we partition the nontypical queries into two: (1) queries $(\mathbf{t}, 1)$ such that the probability to visit \mathbf{t} in $(\mathcal{A}, \mathcal{B})$ or (\mathbf{A}, \mathbf{B}) is much larger than this probability with respect to f (i.e., superpolynomial in n larger than $\Pr[f(X) = (\mathbf{t}, *)]$), and (2) queries $(\mathbf{t}, 1)$ such that the probability of ending in a 1-leaf conditioned on \mathbf{t} is small (i.e., $\Pr[f(X) = (\mathbf{t}, 1) \mid f(X) = (\mathbf{t}, *)]$ is small).

Indeed, the main lemma of this work, Lemma 13, considers the cheating strategies \mathcal{A} and \mathcal{B} with access to the *ideal* sampler. Lemma 13 states that either \mathcal{A} or \mathcal{B} can

bias the output of (A, B) by almost $\frac{\sqrt{2}-1}{2}$, even if they completely fail on nontypical queries. The proof of Theorem 1 (restated as Theorem 14) then follows by using the fact that, on typical queries, replacing the ideal sampler with the real inverter may cause very little effect. In the following, we outline the proof of Lemma 13, focusing on the first type of nontypical queries, which we find to be the more interesting case.

For $q \in \mathbb{N}$, let $\text{UnBal}_{\mathcal{A}}$ contain the transcripts whose weights induced by (\mathcal{A}, B) are at least q times larger than their weights in the honest protocol ($\text{UnBal}_{\mathcal{B}}$ is defined similarly). Using the intuition of the mental experiment described in section 1.3.1, one can show that the probability of every transcript \mathbf{t} induced by a random execution of $(\mathcal{A}, \mathcal{B})$ is at most twice its probability in a random (honest) execution of (A, B) . Hence, the following “compensation effect” happens: if the probability of a transcript \mathbf{t} in (a random execution of) (A, B) is q times *larger* than its probability in (A, B) , then the probability of \mathbf{t} in (\mathcal{A}, B) is q times *smaller* than this value. We conclude that $\text{UnBal}_{\mathcal{A}}$ is visited by $\mathcal{B}^{\text{ideal}}$ with probability at most $1/q$.

To show that both \mathcal{A} and \mathcal{B} can be assumed to fail completely when asking queries in $\text{UnBal}_{\mathcal{A}}$ (the argument for $\text{UnBal}_{\mathcal{B}}$ is analogous), we consider another mental experiment. In this mental experiment, we replace the probabilities of ending up with a 1-leaf, upon reaching a transcript in $\text{UnBal}_{\mathcal{A}}$, by associating a new value to each such transcript. These values are no longer probability measures. Specifically, for all $\mathbf{t} \in \text{UnBal}_{\mathcal{A}}$, we replace the probability that (\mathcal{A}, B) ends up in a 1-leaf conditioned on \mathbf{t} with the value $1/\sqrt{q}$, and replace the probability that (A, \mathcal{B}) ends up in a 1-leaf conditioned on \mathbf{t} with the value \sqrt{q} (this is only a mental experiment, so we can allow these values to be larger than 1). Using a similar approach to that used in section 1.3.1, we can prove that, in the above experiment, it is still true that either \mathcal{A} or \mathcal{B} biases the output of (A, B) by $\frac{\sqrt{2}-1}{2}$.

Finally, we note that we can safely fail both cheating strategies on $\text{UnBal}_{\mathcal{A}}$ almost without changing their overall success probability in the above experiment. Specifically, \mathcal{A} will not suffer much since it visits these nodes with probability at most 1 and gains only $1/\sqrt{q}$ upon visiting them. On the other hand, \mathcal{B} will not suffer much since it visits these nodes with probability at most $1/q$ and gains only \sqrt{q} upon visiting them (hence, these nodes contribute at most $1/\sqrt{q}$ to its overall success). Observe that the probabilities induced by an execution of (\mathcal{A}, B) (or of (A, \mathcal{B})) on typical transcripts in the real scenario, as well as the success probability of the adversary upon visiting these transcripts, are exactly the same as in the above mental experiment. We conclude that either \mathcal{A} or \mathcal{B} biases the output of (A, B) by $\frac{\sqrt{2}-1}{2} - 1/\text{poly}$, even assuming that both cheating strategies totally fail on nontypical queries.

1.3.3. Perspective. The sampling strategy we used above was inspired by the “smooth sampling” approach used by [6, 13, 17] in the setting of parallel repetition of interactive arguments to sample a random winning strategy for the cheating prover. Such an approach can be thought of as a “hedged greedy” strategy, using the recent terminology of Maji, Prabhakaran, and Sahai [21], as it does not necessarily choose the *best* move at each step (the one that maximizes the success probability of the honest strategy), but rather *hedges* its choice according to the relative success probability. The authors of [21] used a different hedged greedy strategy to bias any coin-flipping protocol by $\frac{1}{4} - o(1)$. They then showed how to implement this strategy using an NP-oracle, yielding that $(\frac{1}{4} - o(1))$ -bias coin-flipping protocols imply $\text{BPP} \neq \text{NP}$. Their proof, however, does not follow through using a one-way function inverter and, thus, does not yield that such protocols imply that one-way functions do not exist.

Impagliazzo and Luby [19] used a more conservative method to bias a coin-flipping protocol by $\frac{1}{\sqrt{m}}$ (where m is the protocol's round complexity). Their cheating strategy (which, in turn, was inspired by [8]) follows the prescribed one (i.e., acts honestly), while deviating from it at most once through the execution. In particular, at each step it estimates its potential gain from deviating from the prescribed strategy. If this gain is large enough, it deviates from the prescribed strategy and then continues as the honest party would. Since their strategy needs to estimate only the potential gain *before* deviating from the prescribed strategy, it is rather straightforward to prove that it can be implemented using a one-way function inverter (in particular, the query distribution induced by their strategy is simply the output distribution of the one-way function).

Finally, we mention that the cheating strategy used by the authors of [21] to prove their result for constant-round protocols takes an approach that is very different from the above. Specifically, their cheating strategy uses a one-way function inverter to implement (with close resemblance) the well-known recursive PSPACE-attack on such protocols. Unlike the above greedy strategies, the running time of this recursive approach is doubly exponential in the round complexity of the protocol (which is still efficient for constant-round protocols).

Paper organization. General notation and definitions used throughout the paper are given in section 2. Our adversarial strategy to bias any coin-flipping protocol is presented in section 3. In section 4 we analyze this strategy assuming access to an ideal sampler. Finally, in section 5 we extend this analysis to the real sampler.

2. Preliminaries.

2.1. Notation. We use calligraphic letters to denote sets, uppercase for random variables, and lowercase for values. For an integer $n \in \mathbb{N}$, we let $[n] = \{1, \dots, n\}$.

A function $\mu: \mathbb{N} \rightarrow [0, 1]$ is *negligible* if $\mu(n) = n^{-\omega(1)}$, where neg denotes an arbitrary negligible function. (In particular, $f(n) = \text{neg}(n)$ means that f is negligible, while $f(n) > \text{neg}(n)$ means that f is not negligible.) We let poly denote an arbitrary polynomial, and let PPT denote the set of probabilistic algorithms (i.e., Turing machines) that run in *strict* polynomial time. Given a two-party protocol (A, B) and inputs i_A and i_B , we let $\text{Out}(A(i_A), B(i_B))$ and $(A(i_A), B(i_B))$ denote the (joint) output and transcript, respectively, of the execution of (A, B) with inputs i_A and i_B .

Given a random variable X , we write $x \leftarrow X$ to indicate that x is selected according to X . Similarly, given a finite set \mathcal{S} , we let $s \leftarrow \mathcal{S}$ denote that s is selected according to the uniform distribution on \mathcal{S} . We adopt the convention that when the same random variable occurs several times in an expression, all occurrences refer to a single sample. For example, $\Pr[f(X) = X]$ is defined to be the probability that when $x \leftarrow X$, we have $f(x) = x$. We write U_n to denote the random variable distributed uniformly over $\{0, 1\}^n$. Given a measure M over a set \mathcal{S} , the support of M is defined as $\text{Supp}(M) := \{s \in \mathcal{S}: M(s) > 0\}$. The statistical distance of two distributions P and Q over a finite set \mathcal{U} , denoted by $\text{SD}(P, Q)$, is defined as $\frac{1}{2} \cdot \sum_{u \in \mathcal{U}} |P(u) - Q(u)|$. We use the following notion of measure dominance.

DEFINITION 2 (dominating measure). *A measure M is said to δ -dominate a measure M' if the following hold:*

1. $\text{Supp}(M') \subseteq \text{Supp}(M)$, and
2. $M(y) \geq \delta \cdot M'(y)$ for every $y \in \text{Supp}(M')$.

We stress that it is meaningful to consider δ -dominance even for $0 < \delta \leq 1$. Indeed, for our purposes, it will suffice for a measure M to $1/\text{poly}$ -dominate a measure M' .

2.2. Coin-flipping protocols. In a (strong) coin-flipping protocol the honest execution outputs an unbiased coin, where no (efficient) cheating party can bias the outcome by much. This intuitive description is captured using the following definition.

DEFINITION 3. A polynomial-time protocol (A, B) is a δ -bias (strong) coin-flipping protocol if the following hold:

1. $\Pr[\text{Out}(A, B)(n) = 0] = \Pr[\text{Out}(A, B)(n) = 1] = \frac{1}{2}$, and
2. for any PPTs \mathcal{A} and \mathcal{B} , any $c \in \{0, 1\}$, and all large enough n ,

$$\Pr[\text{Out}(\mathcal{A}, B)(n) = c], \Pr[\text{Out}(A, \mathcal{B})(n) = c] \leq \frac{1}{2} + \delta(n).$$

In the case that $\delta(n) = \text{neg}(n)$, we simply say that (A, B) is a coin-flipping protocol.

It is common to also consider protocols with weaker correctness guarantee than the one we defined above, where with some small probability the output of the protocol is neither 0 nor 1. All the results we present in this paper can be easily generalized to handle such relaxations.

Remark 4 ((partially) fair coin flipping). There are settings in which honest parties are required to always output a bit $c \in \{0, 1\}$, even if the other party arbitrarily deviates from the prescribed protocol (and specifically, upon premature abort by the other party). Constructing coin-flipping protocols in this setting is a much more challenging task. Specifically, it is known that constructing an m -round δ -bias coin-flipping protocol for $\delta \in o(1/m)$ is unconditionally impossible. We mention that since any δ -bias coin-flipping protocol in this setting (i.e., with partial fairness) is also a δ -bias coin-flipping protocol in our setting, our results hold for such protocols as well.

A weaker variant of coin-flipping protocols (that we do not consider in this paper) is that of weak coin-flipping protocols. Such protocols are useful in the case that parties have (a priori known) opposite preferences.

DEFINITION 5. A polynomial time protocol (A, B) is a weak δ -bias coin-flipping protocol if the following hold:

1. $\Pr[\text{Out}(A, B)(n) = 0] = \Pr[\text{Out}(A, B)(n) = 1] = \frac{1}{2}$, and
2. there exist bits $c_A \neq c_B \in \{0, 1\}$ such that the following holds for any PPTs \mathcal{A} and \mathcal{B} , and large enough n :

$$\Pr[\text{Out}(\mathcal{A}, B)(n) = c_A], \Pr[\text{Out}(A, \mathcal{B})(n) = c_B] \leq \frac{1}{2} + \delta(n).$$

In the case that $\delta(n) = \text{neg}(n)$, we simply say that (A, B) is a weak coin-flipping protocol.

2.3. One-way functions and distributional one-way functions. An efficiently computable function is one-way if it is hard to invert it on a random output.

DEFINITION 6 (one-way functions). A polynomially computable function $f : \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ is one-way if the following holds for any PPT A :

$$\Pr_{y \leftarrow f(U_n)} [A(y) \in f^{-1}(y)] = \text{neg}(n).$$

A seemingly weaker requirement is being distributional one-way, meaning that it is hard to sample a random preimage of a random output. The following definition is due to Impagliazzo and Luby [19].

DEFINITION 7 (distributional one-way functions). A polynomially computable function $f : \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ is distributional one-way if there exists $p \in \text{poly}$

such that the following holds for any PPT A :

$$\text{SD}((U_n, f(U_n)), (A(f(U'_n)), f(U'_n))) \geq \frac{1}{p(n)}.$$

Clearly, any one-way function is also a distributional one-way function. While the other implication is not necessarily always true, it was shown in Impagliazzo and Luby [19] that the existence of distributional one-way functions implies that of (standard) one-way functions. In particular, [19] proved that if one-way functions do not exist, then any efficiently computable function has an inverter of the following form. The full proof of this implication can be found in [18].

DEFINITION 8 (γ -inverter). *Let $f : \mathcal{D} \rightarrow \mathcal{R}$ be a deterministic function. An algorithm Inv is called a γ -inverter of f if the following holds:*

$$\text{SD}((U, f(U)), (\text{Inv}(f(U')), f(U'))) \leq \gamma,$$

where U, U' are uniformly distributed in \mathcal{D} .

We call a 0-inverter of f an *ideal inverter* of f . Alternatively, an ideal inverter of f is an algorithm that on $y \in \mathcal{R}$ returns a uniformly chosen element (preimage) in $f^{-1}(y)$.

LEMMA 9 (see [19, Lemma 1]). *Assume that one-way functions do not exist; then for any polynomially computable function $f : \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ and any $p \in \text{poly}$, there exists a PPT Inv that is a $1/p(n)$ -inverter of f for infinitely many n 's.*

Note that nothing is guaranteed when invoking a good inverter for f (i.e., γ -inverter for some small γ) on an arbitrary distribution D . Yet, the following lemma states that if D is dominated by the output distribution of f , then such good inverters are useful.

LEMMA 10. *Let $f : \mathcal{D} \rightarrow \mathcal{R}$ be a deterministic function, and let Ideal be an ideal inverter of f . Let A be an oracle-aided algorithm that makes at most m oracle queries to Ideal , where all of A 's queries are in \mathcal{R} . For $i \in [m]$, let the random variable Q_i describe the i th query of A , where Q_i is set to \perp if the i th query is not asked, and define the measure M_i as follows:*

$$M_i(y) = \begin{cases} \Pr[Q_i = y], & y \in \mathcal{R}, \\ 0 & \text{otherwise.} \end{cases}$$

The probability is taken over the randomness of the algorithm A and the randomness of the ideal inverter Ideal . Let U denote the uniform distribution over \mathcal{D} and suppose that $f(U)$ δ -dominates M_i for all $i \in [m]$ (according to Definition 2); then the following holds for any γ -inverter Inv of f :

$$\text{SD}(A^{\text{Ideal}}, A^{\text{Inv}}) \leq \frac{\gamma \cdot m}{\delta}.$$

Proof. We prove the lemma in two steps. In the first step, we prove that, for the case that $m = 1$, it holds that

$$\text{SD}(A^{\text{Ideal}}, A^{\text{Inv}}) \leq \frac{\gamma}{\delta}.$$

Then, we prove the case where $m > 1$ by reducing it back to the case of $m = 1$ (by considering a slightly different algorithm).

The case of $m = 1$. Since A behaves exactly the same in both cases up to the point at which the query is asked to the inverter, and since upon the same query/answer pair, the algorithm still behaves exactly the same in both cases, we may bound our discussion to the statistical distance between the distributions on query/answer pairs according to each random process. We denote by Q the random variable describing the query (note that this random variable is identically distributed in both scenarios). We denote by X_{Ideal} the random variable describing the answer given in an execution with the ideal inverter, and denote by X_{Inv} the random variable describing the answer given in an execution with the γ -inverter Inv . Since in the case $Q = \perp$ we are guaranteed that $X_{\text{Ideal}} = X_{\text{Inv}} = \perp$, we have that

$$\begin{aligned} \text{SD}(A^{\text{Ideal}}, A^{\text{Inv}}) &\leq \frac{1}{2} \cdot \sum_{q \in \mathcal{R}, a \in \{0,1\}^*} \left| \Pr_{\text{Ideal}}[Q = q \wedge \text{Ideal}(q) = a] - \Pr_{\text{Inv}}[Q = q \wedge \text{Inv}(q) = a] \right| \\ &= \frac{1}{2} \cdot \sum_{q \in \mathcal{R}, a \in \{0,1\}^*} \left| \Pr[Q = q] \cdot \Pr[\text{Ideal}(q) = a] \right. \\ &\quad \left. - \Pr[Q = q] \cdot \Pr[\text{Inv}(q) = a] \right| \\ &= \frac{1}{2} \cdot \sum_{q \in \mathcal{R}, a \in \{0,1\}^*} \left| \Pr[Q = q] \cdot (\Pr[\text{Ideal}(q) = a] - \Pr[\text{Inv}(q) = a]) \right|. \end{aligned}$$

Since M_1 is δ -dominated by $f(U)$, we have that $\Pr[Q = q] = M_i(q) \leq \frac{\Pr[f(U)=q]}{\delta}$ for every $q \in \mathcal{R}$. Hence,

$$\begin{aligned} \text{SD}(A^{\text{Ideal}}, A^{\text{Inv}}) &\leq \frac{1}{2} \cdot \sum_{q \in \mathcal{R}, a \in \{0,1\}^*} \left| \frac{\Pr[f(U) = q]}{\delta} \cdot (\Pr[\text{Ideal}(q) = a] - \Pr[\text{Inv}(q) = a]) \right| \\ &= \frac{1}{\delta} \cdot \frac{1}{2} \cdot \sum_{q \in \mathcal{R}, a \in \{0,1\}^*} \left| \Pr[f(U) = q] \cdot \Pr[\text{Ideal}(q) = a] - \Pr[f(U) = q] \cdot \Pr[\text{Inv}(q) = a] \right| \\ &= \frac{1}{\delta} \cdot \text{SD}((U, f(U)), (\text{Inv}(f(U')), f(U'))) \leq \frac{\gamma}{\delta}. \end{aligned}$$

The case of $m > 1$. Define a sequence $\{H_i\}_{i=0}^m$ of hybrid random variables by letting the i th hybrid describe the output of the algorithm A in an execution where the first i queries of A are answered by Ideal and the remaining $m - i$ queries are answered by Inv . Specifically, we have that $H_m \equiv A^{\text{Ideal}}$ and that $H_0 \equiv A^{\text{Inv}}$. By the triangle inequality, we have that

$$\begin{aligned} \text{SD}(A^{\text{Ideal}}, A^{\text{Inv}}) &= \text{SD}(H_0, H_m) \\ &\leq \sum_{i=1}^m \text{SD}(H_{i-1}, H_i). \end{aligned}$$

Therefore, it suffices to prove that for every i it holds that $\text{SD}(H_{i-1}, H_i) \leq \frac{\gamma}{\delta}$. Observe that except for the i th query, both random variables describe the same random process, which can be viewed as an algorithm that asks only a single oracle query. Furthermore, since the first $i - 1$ queries were asked to the ideal inverter, it holds that the i th query is indeed distributed as Q_i . Thus, the lemma follows by applying the proof for the case where $m = 1$. \square

3. The attack. Let (A, B) be a coin-tossing protocol. In the following we define adversarial strategies for both A and B to bias the output of the protocol towards 1. The strategies for biasing the output towards 0 are defined analogously.

3.1. Notation. We associate the following random variables with an (honest) execution of (A, B) . Throughout, we let n be the security parameter of the protocol and omit it whenever its value is clear from the context. We assume for simplicity that the protocol's messages are single bits, and naturally view a valid execution of the protocol as a path in the binary tree $\mathcal{T} = \mathcal{T}_n$, whose nodes are associated with all possible (valid) transcripts. The root of \mathcal{T} , corresponding to the empty transcript, is denoted by the empty string λ , and the children of a node \mathbf{t} (if they exist) are denoted by $\mathbf{t} \circ 0$ and $\mathbf{t} \circ 1$ (\circ stands for string concatenation), corresponding to the two (possibly partial) executions with these transcripts. A node with no descendants (associated with a full transcript) is called a *leaf*, where we assume for simplicity that a nonleaf node has exactly two descendants. Given a node \mathbf{t} , we let $|\mathbf{t}|$ denote its depth, and for $i \in [|\mathbf{t}|]$ we let \mathbf{t}_i denote the prefix of length i of \mathbf{t} , which describes the i th node on the path from λ to \mathbf{t} (e.g., $\mathbf{t}_0 = \lambda$).

We call a transcript \mathbf{t} an A node (resp., B node), if this is A 's (resp., B 's) turn to send the next message, where without loss of generality the root λ is an A node. We also assume that the parties always exchange $m = m(n)$ messages, and that each party uses $s = s(n)$ random coins, denoted r_A and r_B , respectively. Given a pair of s -long random strings (r_A, r_B) , we let $\text{Trans}(r_A, r_B) = (A(r_A), B(r_B))$ (i.e., the full transcript induced by the execution of $(A(r_A), B(r_B))$).

For $\mathbf{t} \in \mathcal{T}$, let $\text{Uni}(\mathbf{t})$ denote a random sample of (r_A, r_B) , conditioned on $\text{Trans}(r_A, r_B)|_{|\mathbf{t}|} = \mathbf{t}$. Given random coins $r_A \in \{0, 1\}^s$, we let $\text{Next}^A(r_A; \mathbf{t})$ be the next message sent by A with random coins r_A after seeing the transcript \mathbf{t} , and we define the random variable $A(\mathbf{t})$ as $\text{Next}^A(R_A; \mathbf{t})$, where $(R_A, *) \leftarrow \text{Uni}(\mathbf{t})$. ($\text{Next}^B(r_B; \mathbf{t})$ and $B(\mathbf{t})$ are defined analogously.) Finally, we assume without loss of generality that the transcript of an (honest) execution of the protocol always defines an output, 0 or 1 (consistent for both parties). For a leaf \mathbf{t} , we let $V_{\mathbf{t}}$ be the output of the protocol determined by this leaf, where if \mathbf{t} is an internal node, we define $V_{\mathbf{t}}$ as

$$(2) \quad V_{\mathbf{t}} = \mathbb{E}_{(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})} [V_{\text{Trans}(r_A, r_B)}].$$

Namely, $V_{\mathbf{t}}$ is the probability that (A, B) outputs 1, conditioned on \mathbf{t} being the current transcript.

Similarly, we associate the following random variables with an execution of (\mathcal{A}, B) , where \mathcal{A} is a cheating strategy for A : we denote the random coins used by \mathcal{A} by $r_{\mathcal{A}}$, and for $\mathbf{t} \in \mathcal{T}$ we let $\text{Uni}^{\mathcal{A}}(\mathbf{t})$ denote a random sample of $(r_{\mathcal{A}}, r_B)$, conditioned on $(\mathcal{A}(r_{\mathcal{A}}), B(r_B))|_{|\mathbf{t}|} = \mathbf{t}$. Given random coins $r_{\mathcal{A}} \in \{0, 1\}^s$, we let $\text{Next}^{\mathcal{A}}(r_{\mathcal{A}}; \mathbf{t})$ be the next message sent by \mathcal{A} with random coins $r_{\mathcal{A}}$ after seeing the transcript \mathbf{t} , and we define the random variable $\mathcal{A}(\mathbf{t})$ as $\text{Next}^{\mathcal{A}}(R_{\mathcal{A}}; \mathbf{t})$, where $(R_{\mathcal{A}}, *) \leftarrow \text{Uni}^{\mathcal{A}}(\mathbf{t})$. ($\text{Next}^B(r_B; \mathbf{t})$ and $B(\mathbf{t})$ are defined analogously.) Finally, we define $V_{\mathbf{t}}^{\mathcal{A}}$ as

$$(3) \quad V_{\mathbf{t}}^{\mathcal{A}} = \mathbb{E}_{(r_{\mathcal{A}}, r_B) \leftarrow \text{Uni}^{\mathcal{A}}(\mathbf{t})} [V_{(\mathcal{A}(r_{\mathcal{A}}), B(r_B))}],$$

where we set $V_{(\mathcal{A}(r_{\mathcal{A}}), B(r_B))} = 0$ if $(\mathcal{A}(r_{\mathcal{A}}), B(r_B))$ aborts. Namely, $V_{\mathbf{t}}^{\mathcal{A}}$ is a lower bound on the probability that (\mathcal{A}, B) outputs 1, with the condition that \mathbf{t} is the current transcript. ($V_{\mathbf{t}}^B$ is defined analogously.)

3.2. The adversary \mathcal{A} . We now present an adversarial strategy \mathcal{A} for A, designed to bias the outcome of the protocol towards 1 (the adversarial strategy \mathcal{B} for B is defined analogously). In each round \mathcal{A} uses a “sampling oracle” **Samp** to sample a value for the coins of A, and then acts as the (honest) A would, given these coins and the current transcript. Roughly speaking, the objective of **Samp** is to return a random pair of coins (r_A, r_B) consistent with \mathbf{t} (i.e., $\text{Trans}(r_A, r_B)|_{\mathbf{t}} = \mathbf{t}$), which leads to a 1-node (i.e., $V_{\text{Trans}(r_A, r_B)} = 1$). In the following we analyze the success probability of \mathcal{A} when using different implementations for **Samp**. Specifically, in section 4 we consider an “ideal sampler” (which is not necessarily efficient). Then, in section 5, we consider a more realistic implementation of the sampler (specifically, using the inverter that will stem from the assumption that one-way functions do not exist). Before describing and analyzing each of these samplers, we first give the formal description of \mathcal{A} .

ALGORITHM 11 (adversary \mathcal{A}).

Input: Security parameter n .

Oracle: **Samp**.

Operation: Let \mathbf{t} be the current transcript.

1. Halt if \mathbf{t} is a leaf node.
2. Let $(r_A, *) \leftarrow \text{Samp}(\mathbf{t})$. Abort if $r_A = \perp$.
3. Send $\text{Next}^A(r_A; \mathbf{t})$ to B.

Given an instantiation of **Samp**, we view $\mathcal{A}^{\text{Samp}}$ as a random algorithm whose random coins are those used by **Samp** (independent coins for each call).

4. Using the ideal sampler. Our “ideal sampler” **Ideal** is defined as follows: on input $\mathbf{t} \in \mathcal{T}$, **Ideal** returns a random sample $(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})$, conditioned on $V_{\text{Trans}(r_A, r_B)} = 1$. In the case that $V_{\mathbf{t}} = 0$, **Ideal** returns \perp . The following lemma asserts that at least one of the parties has a good cheating strategy given oracle access to this sampler.

LEMMA 12. For any $n \in \mathbb{N}$ and any transcript $\mathbf{t} \in \mathcal{T}_n$, it holds that

$$V_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} \cdot V_{\mathbf{t}}^{\mathcal{B}^{\text{Ideal}}} \geq V_{\mathbf{t}}.$$

Assuming that $V_{\lambda_n} = 1/2$, it holds that either $V_{\lambda}^{\mathcal{A}^{\text{Ideal}}} \geq 1/\sqrt{2}$ or $V_{\lambda}^{\mathcal{B}^{\text{Ideal}}} \geq 1/\sqrt{2}$. Namely, either $\mathcal{A}^{\text{Ideal}}$ or $\mathcal{B}^{\text{Ideal}}$ can bias the output of the protocol by $\frac{1}{\sqrt{2}} - \frac{1}{2}$.

Proof of Lemma 12. We prove the lemma using induction up the protocol tree. The proof is immediate for a leaf node and for an internal node \mathbf{t} with $V_{\mathbf{t}} = 0$ (i.e., when **Ideal**(\mathbf{t}) = \perp). In the following \mathbf{t} is a fixed internal node (with $V_{\mathbf{t}} > 0$). Thus, for the sake of simplicity of notation, we let $\mathcal{A} = \mathcal{A}^{\text{Ideal}}$, $V = V_{\mathbf{t}}$, and $V^{\mathcal{A}} = V_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}}$. Similarly, for $j \in \{0, 1\}$ we let $V_j = V_{\mathbf{t}j}$ and $V_j^{\mathcal{A}} = V_{\mathbf{t}j}^{\mathcal{A}^{\text{Ideal}}}$. ($V^{\mathcal{B}}$ and $V_j^{\mathcal{B}}$ are defined analogously.) We need to prove that $V^{\mathcal{A}} \cdot V^{\mathcal{B}} \geq V$, assuming that $V_j^{\mathcal{A}} \cdot V_j^{\mathcal{B}} \geq V_j$ for both $j \in \{0, 1\}$. We assume that \mathbf{t} is an A node (the other case is analogous). Let $\beta = \Pr[\mathbf{A}(\mathbf{t}) = 1]$ (i.e., the probability that the next message of the honest A is 1). Note that $V = \beta \cdot V_1 + (1 - \beta) \cdot V_0$, and that

$$\begin{aligned} \Pr[\mathcal{A}(\mathbf{t}) = 1] &= \frac{\Pr_{(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})}[\text{Trans}(r_A, r_B)|_{\mathbf{t}+1} = \mathbf{t} \circ 1 \mid V_{\text{Trans}(r_A, r_B)} = 1]}{\Pr_{(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})}[\text{Trans}(r_A, r_B)|_{\mathbf{t}+1} = \mathbf{t} \circ 1 \wedge V_{\text{Trans}(r_A, r_B)} = 1]} \\ &= \frac{\beta \cdot V_1}{V}, \end{aligned}$$

where the last equation is by a simple chain rule, i.e., since

$$\begin{aligned}\beta &= \Pr_{(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})} [\text{Trans}(r_A, r_B)_{|\mathbf{t}|+1} = \mathbf{t} \circ 1], \text{ and} \\ V_1 &= \Pr_{(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})} [V_{\text{Trans}(r_A, r_B)} = 1 \mid \text{Trans}(r_A, r_B)_{|\mathbf{t}|+1} = \mathbf{t} \circ 1].\end{aligned}$$

Thus,

- $V^{\mathcal{A}} = \Pr[\mathcal{A}(\mathbf{t}) = 1] \cdot V_1^{\mathcal{A}} + \Pr[\mathcal{A}(\mathbf{t}) = 0] \cdot V_0^{\mathcal{A}} = \frac{\beta \cdot V_1}{V} \cdot V_1^{\mathcal{A}} + \frac{(1-\beta) \cdot V_0}{V} \cdot V_0^{\mathcal{A}}$, and
- $V^{\mathcal{B}} = \Pr[\mathcal{B}(\mathbf{t}) = 1] \cdot V_1^{\mathcal{B}} + \Pr[\mathcal{B}(\mathbf{t}) = 0] \cdot V_0^{\mathcal{B}} = \beta \cdot V_1^{\mathcal{B}} + (1 - \beta) \cdot V_0^{\mathcal{B}}$.

Using the induction hypothesis, we get that

$$\begin{aligned}V^{\mathcal{A}} \cdot V^{\mathcal{B}} &= \left(\frac{\beta \cdot V_1}{V} \cdot V_1^{\mathcal{A}} + \frac{(1-\beta) \cdot V_0}{V} \cdot V_0^{\mathcal{A}} \right) \cdot (\beta \cdot V_1^{\mathcal{B}} + (1-\beta) \cdot V_0^{\mathcal{B}}) \\ &= \frac{\beta^2 \cdot V_1}{V} \cdot V_1^{\mathcal{A}} \cdot V_1^{\mathcal{B}} + \frac{(1-\beta)^2 \cdot V_0}{V} \cdot V_0^{\mathcal{A}} \cdot V_0^{\mathcal{B}} \\ &\quad + \frac{\beta(1-\beta)}{V} \cdot (V_1 \cdot V_1^{\mathcal{A}} \cdot V_0^{\mathcal{B}} + V_0 \cdot V_0^{\mathcal{A}} \cdot V_1^{\mathcal{B}}) \\ &\geq \frac{\beta^2 \cdot V_1^2}{V} + \frac{(1-\beta)^2 \cdot V_0^2}{V} + \frac{\beta \cdot (1-\beta) \cdot V_1 \cdot V_0}{V} \cdot \left(\frac{V_1^{\mathcal{A}}}{V_0^{\mathcal{A}}} + \frac{V_0^{\mathcal{A}}}{V_1^{\mathcal{A}}} \right).\end{aligned}$$

Since $V^2 = (\beta \cdot V_1 + (1-\beta) \cdot V_0)^2 = \beta^2 \cdot V_1^2 + (1-\beta)^2 \cdot V_0^2 + 2 \cdot \beta \cdot (1-\beta) \cdot V_1 \cdot V_0$, it follows that

$$\begin{aligned}V^{\mathcal{A}} \cdot V^{\mathcal{B}} &\geq \frac{\beta^2 \cdot V_1^2}{V} + \frac{(1-\beta)^2 \cdot V_0^2}{V} + \frac{\beta \cdot (1-\beta) \cdot V_1 \cdot V_0}{V} \cdot \left(\frac{V_1^{\mathcal{A}}}{V_0^{\mathcal{A}}} + \frac{V_0^{\mathcal{A}}}{V_1^{\mathcal{A}}} \right) \\ &= V + \frac{\beta \cdot (1-\beta) \cdot V_1 \cdot V_0}{V} \cdot \left(\frac{V_1^{\mathcal{A}}}{V_0^{\mathcal{A}}} + \frac{V_0^{\mathcal{A}}}{V_1^{\mathcal{A}}} - 2 \right).\end{aligned}$$

Thus, to prove that $V^{\mathcal{A}} \cdot V^{\mathcal{B}} \geq V$ it suffices to prove that $\frac{V_1^{\mathcal{A}}}{V_0^{\mathcal{A}}} + \frac{V_0^{\mathcal{A}}}{V_1^{\mathcal{A}}} \geq 2$. We will show that $(\frac{V_1^{\mathcal{A}}}{V_0^{\mathcal{A}}} + \frac{V_0^{\mathcal{A}}}{V_1^{\mathcal{A}}} - 2) \cdot V_1^{\mathcal{A}} \cdot V_0^{\mathcal{A}} \geq 0$, which also suffices, since $V_1^{\mathcal{A}}$ and $V_0^{\mathcal{A}}$ are both positive values. Indeed,

$$\begin{aligned}\left(\frac{V_1^{\mathcal{A}}}{V_0^{\mathcal{A}}} + \frac{V_0^{\mathcal{A}}}{V_1^{\mathcal{A}}} - 2 \right) \cdot V_1^{\mathcal{A}} \cdot V_0^{\mathcal{A}} &= (V_1^{\mathcal{A}})^2 + (V_0^{\mathcal{A}})^2 - 2V_1^{\mathcal{A}} \cdot V_0^{\mathcal{A}} \\ &= (V_1^{\mathcal{A}} - V_0^{\mathcal{A}})^2 \\ &\geq 0. \quad \square\end{aligned}$$

5. Moving to an efficient sampler. Our goal in this section is to use the above analysis of the success probability of our adversaries when given access to the *ideal* sampler, for analyzing their success probability when given access to an *efficient* sampler. The accuracy of such an inverter will be parametrized by a function $1/p$ for some $p \in \text{poly}$. In the following we fix such p .

Assuming that one-way functions do not exist, our efficient sampler is defined as follows: let $f : \{0, 1\}^{s(n)} \times \{0, 1\}^{s(n)} \times \{0, \dots, m(n)\}$ be defined as

$$f(r_A, r_B, i) = \text{Trans}(r_A, r_B)_i, V_{\text{Trans}(r_A, r_B)}.$$

Namely, $f(r_A, r_B, i)$ outputs the i th node in the execution of $(A(r_A), B(r_B))$ and the outcome coin induced by the leaf (i.e., full transcript) of this execution. Given a node

$\mathbf{t} \in \mathcal{T}_n$, the sampler Real_p returns $\text{Inv}_f(\mathbf{t}, 1)$, where Inv_f is the distributional inverter for f guaranteed by Lemma 9 with respect to accuracy parameter $1/p$.⁴

Notice that while Lemma 9 tells us that Inv_f samples well over a random output of f , the distribution induced by the calls of $\mathcal{A}^{\text{Real}_p}$ might be very different from this distribution. While we cannot bound the difference between these two distributions, we prove that there exists a high-probability event conditioned on these distributions being close enough. Loosely speaking, we first show that Lemma 12 still (almost) holds even if both $\mathcal{A}^{\text{Ideal}}$ and $\mathcal{B}^{\text{Ideal}}$ fail on their “nontypical” queries to Ideal —the calls that happen with probability that is very different from the one induced by f . Since Real_p performs similarly to Ideal on the typical queries, it follows that $V_\lambda^{\mathcal{A}^{\text{Real}_p}} \cdot V_\lambda^{\mathcal{B}^{\text{Real}_p}}$ is almost as large as V_λ , and, therefore, either $\mathcal{A}^{\text{Real}_p}$ or $\mathcal{B}^{\text{Real}_p}$ can significantly bias the outcome of the protocol.

In Lemma 13, we formally capture the above intuition regarding $\mathcal{A}^{\text{Ideal}}$ and $\mathcal{B}^{\text{Ideal}}$ (with access to the ideal sampler). We denote by $w(\mathbf{t})$ the probability that the node \mathbf{t} is visited in a random execution of (A, B) , and by $w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t})$ the probability of this visit in a random execution of $(\mathcal{A}^{\text{Ideal}}, B)$. ($w^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t})$ is defined analogously.) Recall that we omit n from the notation whenever its value is clear from the context. Specifically, we let λ denote the root of \mathcal{T}_n , and $V_\lambda = \mathbb{E}[\text{Out}(A, B)(1^n)]$.

LEMMA 13. *Let (A, B) be a coin-tossing protocol as above. For any $q \in \text{poly}$ and for any $n \in \mathbb{N}$, there exists a set $\mathcal{E} \subseteq \{\mathbf{t} \in \mathcal{T}_n : V_{\mathbf{t}} > 0\}$ such that the following holds:*

1. *For any $\mathbf{t} \in \mathcal{E}$, it holds that $\max\{w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}), w^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t})\} \in O(q(n)^5 \cdot w(\mathbf{t}) \cdot V_{\mathbf{t}})$, and*
2. *$V_\lambda^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \cdot V_\lambda^{\mathcal{B}^{\mathcal{E}^{\text{Ideal}}}} \geq V_\lambda - \frac{1}{q(n)}$, where $\mathcal{A}_{\mathcal{E}}$ acts as \mathcal{A} does, but aborts if a node outside of \mathcal{E} is reached. ($\mathcal{B}_{\mathcal{E}}$ is defined analogously.)*

Proving Lemma 13 is the main contribution of this section, but first let us use it for proving Theorem 14.

THEOREM 14 (restating Theorem 1). *Let (A, B) be a coin-tossing protocol with $V_\lambda = \mathbb{E}[\text{Out}(A, B)(1^n)]$. Assuming that one-way functions do not exist, then for any $g \in \text{poly}$ there exists a pair of efficient (cheating) strategies \mathcal{A} and \mathcal{B} such that the following holds: for infinitely many n 's, for each $j \in \{0, 1\}$ either $\Pr[(\mathcal{A}(j), B)(1^n) = j]$ or $\Pr[(\mathcal{B}(j), B)(1^n) = j]$ is greater than $\sqrt{V_n^j} - \frac{1}{g(n)}$, where $V_n^1 = V_\lambda$ and $V_n^0 = 1 - V_\lambda$.*

In particular, for the case of $V_\lambda = \frac{1}{2}$, one party can “bias the outcome” of (A, B) by almost $\frac{1}{\sqrt{2}} - \frac{1}{2}$.

Proof. We focus on $j = 1$, where the proof for $j = 0$ follows analogously. We prove the theorem by considering the success probabilities of the adversaries $\mathcal{A}^{\text{Real}_p}$ and $\mathcal{B}^{\text{Real}_p}$ (with access to an efficient inverter Inv_f) on the set $\mathcal{E} \subseteq \mathcal{T}_n$ guaranteed by Lemma 13, namely, the success probabilities of $\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}$ and $\mathcal{B}_{\mathcal{E}}^{\text{Real}_p}$. We show that if Inv_f is “good enough,” then they will do almost as well as $\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}$ and $\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}$ would. Towards this end, we show that the distribution induced by f on a random input $(1/\text{poly})$ -dominates (according to Definition 2) both query distributions induced by $\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}$ and $\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}$. Thus, we can apply Lemma 10 to show that each adversary behaves almost identically when given access to Ideal as when given access to Inv_f . Finally, we remark that, while $\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}$ and $\mathcal{B}_{\mathcal{E}}^{\text{Real}_p}$ may not be efficient (since they need to abort on $\mathbf{t} \notin \mathcal{E}$), they serve as a mental experiment and provide lower bounds on the success probabilities of $\mathcal{A}^{\text{Real}_p}$ and $\mathcal{B}^{\text{Real}_p}$, respectively. We next give the formal argument.

⁴We assume for simplicity that the security parameter of the protocol is determined by its (even partial) transcript, and, therefore, the domain of f in the calls to Inv_f is well defined.

Let $g'(n) := \frac{g(n)}{\sqrt{V^1}}$, where we assume without loss of generality that $g(n) \geq \frac{1}{\sqrt{V^1}}$ (otherwise, the statement is trivial). Let $D_f(y)$ be the probability that a random output of f equals y . Note that the following holds for any $\mathbf{t} \in \mathcal{T}_n$:

$$\begin{aligned} D_f(\mathbf{t}, 1) &:= \Pr[f(U_{2s(n)}, I_n) = (\mathbf{t}, 1)] \\ &= \Pr[I_n = |\mathbf{t}|] \cdot \Pr[\text{Trans}(U_{2s(n)})_{|\mathbf{t}|} = \mathbf{t} \wedge V_{\text{Trans}(U_{2s(n)})} = 1] \\ &= \frac{1}{m(n) + 1} \cdot w(\mathbf{t}) \cdot V_{\mathbf{t}}, \end{aligned}$$

where I_n is uniformly distributed over $\{0, \dots, m(n)\}$. Let $\mathcal{E} \subseteq \mathcal{T}_n$ be the set guaranteed by Lemma 13 with respect to $q(n) = 2 \cdot g'(n)$. It follows that

$$(4) \quad \max\{w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}), w^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t})\} \in O(q(n)^5 \cdot m(n) \cdot D_f(\mathbf{t}, 1))$$

for any $\mathbf{t} \in \mathcal{E}$. In other words, the distributions induced by the queries of $\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}$ and $\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}$ on the range of f are δ -dominated by the distribution of a random output of f for $\delta = 1/O(q(n)^5 \cdot m(n))$.

Fix $n \in \mathbb{N}$ such that the inverter Inv_f (guaranteed by Lemma 9) is a $1/p(n)$ -inverter for f , and let Real_p be the sampler described above (i.e., $\text{Real}_p(\mathbf{t})$ returns $\text{Inv}_f(\mathbf{t}, 1)$). For $\text{Samp} \in \{\text{Ideal}, \text{Real}_p\}$, let $E^{\mathcal{A}_{\mathcal{E}}^{\text{Samp}}}$ be the algorithm that emulates a random execution of $(\mathcal{A}_{\mathcal{E}}^{\text{Samp}}, \mathcal{B})$ and outputs the outcome of this execution, where $\mathcal{A}_{\mathcal{E}}$ is as in Lemma 13 ($E^{\mathcal{B}_{\mathcal{E}}^{\text{Samp}}}$ is defined analogously). For $i \in \{0, \dots, m(n)\}$, let Q_i be the value of the i th Ideal -query made in the execution of $E^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}}$ (set to \perp if no such call was made). Relation (4) yields that $\Pr[Q_i = (\mathbf{t}, 1)] \in O(q(n)^5 \cdot m(n) \cdot D_f(\mathbf{t}, 1))$ for any $i \in [m(n)]$ and for any $\mathbf{t} \in \mathcal{E}$. Thus, Lemma 10 yields that

$$\text{SD}(E^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}}, E^{\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}}) \in \frac{O(q(n)^5 \cdot m(n)^2)}{p(n)} < 1/8g'(n)$$

for the proper choice of p . Therefore,

$$(5) \quad V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}} = \Pr[E^{\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}} = 1] \geq \Pr[E^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} = 1] - 1/4g'(n) = V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} - 1/4g'(n).$$

Doing the analogous calculation for $V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Real}_p}}$ and using Lemma 13, it follows that

$$\begin{aligned} (6) \quad V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}} \cdot V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Real}_p}} &\geq (V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} - 1/4g'(n)) \cdot (V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}} - 1/4g'(n)) \\ &= V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} \cdot V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}} - \frac{V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} + V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}}}{4g'(n)} + \frac{1}{2g'(n)^2} \\ &\geq V^1 - \frac{1}{q(n)} - \frac{1}{2g'(n)} = V^1 - \frac{1}{g'(n)}. \end{aligned}$$

Since $V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}} \geq V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}}$ and $V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Real}_p}} \geq V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Real}_p}}$ (on the nodes in \mathcal{E} the strategies $\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}$ and $\mathcal{A}^{\text{Real}_p}$ act identically, and $\mathcal{A}_{\mathcal{E}}^{\text{Real}_p}$ fails on the other nodes), it follows that

$$(7) \quad V_{\lambda}^{\mathcal{A}^{\text{Real}_p}} \cdot V_{\lambda}^{\mathcal{B}^{\text{Real}_p}} \geq V^1 - \frac{1}{g'(n)}.$$

It follows that

$$V^1 - \frac{1}{g'(n)} = V^1 - \frac{\sqrt{V^1}}{g(n)} \geq V^1 - \frac{2 \cdot \sqrt{V^1}}{g(n)} + \frac{1}{g(n)^2} = \left(\sqrt{V^1} - \frac{1}{g(n)} \right)^2,$$

where the inequality holds since $g(n) \geq \frac{1}{\sqrt{V^1}}$. In particular, either $V_{\lambda}^{\mathcal{A}^{\text{Real}_p}}$ or $V_{\lambda}^{\mathcal{B}^{\text{Real}_p}}$ is larger than $\sqrt{V^1} - \frac{1}{g(n)} \geq \sqrt{V^1} - \frac{1}{g(n)}$, which completes the proof of the theorem. \square

5.1. Proving Lemma 13. Towards proving Lemma 13 we identify the nodes (queries) in $\mathcal{T} = \mathcal{T}_n$ that are potentially “nontypical” (i.e., either $V_{\mathbf{t}}$ is small or $\max\{w^{\mathcal{A}^{\text{ideal}}}(\mathbf{t}), w^{\mathcal{B}^{\text{ideal}}}(\mathbf{t})\}$ is large) and prove that by modifying $\mathcal{A}^{\text{ideal}}$ or $\mathcal{B}^{\text{ideal}}$ to totally fail on such nodes, we hardly change their overall success probability. Alternatively, if $\mathcal{A}^{\text{ideal}}$ and $\mathcal{B}^{\text{ideal}}$ abort whenever they reach a nontypical node (as do $\mathcal{A}_{\mathcal{E}}^{\text{ideal}}$ and $\mathcal{B}_{\mathcal{E}}^{\text{ideal}}$), then they will only give away a $1/\text{poly}$ fraction of their success probability. The proof then follows by taking \mathcal{E} to be the set of “typical” nodes in \mathcal{T} .

We next give a slightly more detailed overview of the proof. For simplicity, in the discussion below, we (implicitly) assume that V_{λ} is constant (in the formal proof, we deal with any value of V_{λ}). We need to show that the set \mathcal{E} satisfies both of the requirements in Lemma 13. Proving that the first requirement is satisfied will be immediate, simply by the way we define nontypical nodes. To show that \mathcal{E} satisfies the second requirement (i.e., that $\mathcal{A}^{\text{ideal}}$ and $\mathcal{B}^{\text{ideal}}$ can indeed abort on nodes outside \mathcal{E} without losing much), we partition the nontypical nodes into two sets. The first set, denoted **Small**, contains those nodes for which $V_{\mathbf{t}} \in O(\frac{1}{q^2})$. The second set, denoted **UnBal**, contains the nodes whose weights induced by $\mathcal{A}^{\text{ideal}}$ or $\mathcal{B}^{\text{ideal}}$ are $\Omega(q^2)$ times larger than their weights in an honest execution of the protocol. On a very intuitive level, handling the set **Small** is fairly easy: consider a mental experiment in which we (artificially) set a new “success probability” for such nodes by setting $V_{\mathbf{t}}^{\mathcal{A}^{\text{ideal}}} = V_{\mathbf{t}}^{\mathcal{B}^{\text{ideal}}} = \sqrt{V_{\mathbf{t}}}$ for every $\mathbf{t} \in \text{Small}$. Since $V_{\mathbf{t}}^{\mathcal{A}^{\text{ideal}}} \cdot V_{\mathbf{t}}^{\mathcal{B}^{\text{ideal}}} \geq V_{\mathbf{t}}$, the proof of Lemma 12 will still go through with respect to the above experiment. Namely, it will still hold that $V_{\lambda}^{\mathcal{A}^{\text{ideal}}} \cdot V_{\lambda}^{\mathcal{B}^{\text{ideal}}} \geq V_{\lambda}$. To then allow aborting on nodes in **Small**, we observe that neither $\mathcal{A}^{\text{ideal}}$ nor $\mathcal{B}^{\text{ideal}}$ gains much on any node $\mathbf{t} \in \text{Small}$ (at most $\sqrt{V_{\mathbf{t}}} \in O(1/q)$). Hence, even if **Small** is reached with high probability, it contributes an overall success probability of $O(1/q)$.

Handling the unbalanced nodes inside **UnBal**, on the other hand, seems much more challenging. These nodes might have arbitrary expected values (i.e., $V_{\mathbf{t}}$) and are reached by one of the adversaries with high probability. As such, they may contribute significantly to the success probability of the cheating parties. Fortunately, by making a critical use of the query distribution induced by the ideal sampler, we are able to prove the following “compensation lemma”: a node \mathbf{t} whose weight with respect to $\mathcal{A}^{\text{ideal}}$ is k times *larger* than its real weight (i.e., $w^{\mathcal{A}^{\text{ideal}}}(\mathbf{t}) = k \cdot w(\mathbf{t})$) has weight with respect to $\mathcal{B}^{\text{ideal}}$ that is k time *smaller* than its real weight. Hence, the set **UnBal** can be separated into two disjoint subsets **UnBal_A** and **UnBal_B**, where **UnBal_B** is almost never visited by $\mathcal{A}^{\text{ideal}}$, and **UnBal_A** is almost never visited by $\mathcal{B}^{\text{ideal}}$. Now, we handle each of these sets in a manner similar to the way we handled the nodes in **Small** (for simplicity we consider here only the set **UnBal_A**): consider the mental experiment in which for every $\mathbf{t} \in \text{UnBal}_A$ we modify the values of $V_{\mathbf{t}}^{\mathcal{A}^{\text{ideal}}}$ and $V_{\mathbf{t}}^{\mathcal{B}^{\text{ideal}}}$ such that $V_{\mathbf{t}}^{\mathcal{A}^{\text{ideal}}} = 1/q$ and $V_{\mathbf{t}}^{\mathcal{B}^{\text{ideal}}} = q$ (this is only a mental experiment, so we do not care that these values might be larger than 1). Since $V_{\mathbf{t}}^{\mathcal{A}^{\text{ideal}}} \cdot V_{\mathbf{t}}^{\mathcal{B}^{\text{ideal}}} = 1 \geq V_{\mathbf{t}}$, the proof of Lemma 12 still goes through with respect to this experiment as well. Furthermore, we can safely fail both cheating strategies on **UnBal_A** without changing their overall success probability too much. Specifically, $\mathcal{A}^{\text{ideal}}$ will not suffer much because its success probability on these nodes is bounded by $\frac{1}{q}$ (i.e., it has gained at most $O(1 \cdot \frac{1}{q} = \frac{1}{q})$ from these nodes), and $\mathcal{B}^{\text{ideal}}$ will not suffer much since it almost never visits these nodes (i.e., it has gained $O(q \cdot \frac{1}{q^2} = \frac{1}{q})$ from these nodes).

We now work towards formalizing the above discussion. We assume that $V_{\lambda} \geq 1/q$, since otherwise the lemma follows trivially, and start with formally defining the

different subsets of \mathcal{T} that we considered above. We define the relative weights of $\mathbf{t} \in \mathcal{T}$ as $Wrel^{\mathcal{A}^{Ideal}}(\mathbf{t}) = \frac{w^{\mathcal{A}^{Ideal}}(\mathbf{t})}{w(\mathbf{t})}$ and $Wrel^{\mathcal{B}^{Ideal}}(\mathbf{t}) = \frac{w^{\mathcal{B}^{Ideal}}(\mathbf{t})}{w(\mathbf{t})}$. Then let

$$(8) \quad \text{UnBal}_{\mathcal{A}} := \{\mathbf{t} \in \mathcal{T} : Wrel^{\mathcal{A}^{Ideal}}(\mathbf{t}) > 16 \cdot q^3\},$$

$$(9) \quad \text{UnBal}_{\mathcal{B}} := \{\mathbf{t} \in \mathcal{T} : Wrel^{\mathcal{B}^{Ideal}}(\mathbf{t}) > 16 \cdot q^3\},$$

and let $\text{UnBal} = \text{UnBal}_{\mathcal{A}} \cup \text{UnBal}_{\mathcal{B}}$. Finally, we let

$$(10) \quad \text{Small} := \left\{ \mathbf{t} \in \mathcal{T} \setminus \text{UnBal} : V_{\mathbf{t}} < \frac{1}{16 \cdot q^2} \right\}$$

and let $\mathcal{E} = \mathcal{T} \setminus (\text{Small} \cup \text{UnBal})$. The following fact is immediate.

CLAIM 15. For any $\mathbf{t} \in \mathcal{E}$ it holds that $\max\{w^{\mathcal{A}^{Ideal}}(\mathbf{t}), w^{\mathcal{B}^{Ideal}}(\mathbf{t})\} \in O(q^5 \cdot w(\mathbf{t}) \cdot V_{\mathbf{t}})$.

Proof. For any $\mathbf{t} \in \mathcal{E}$, it holds that $\mathbf{t} \notin \text{UnBal}_{\mathcal{A}}$. Hence,

$$w^{\mathcal{A}^{Ideal}}(\mathbf{t}) = Wrel^{\mathcal{A}^{Ideal}}(\mathbf{t}) \cdot w(\mathbf{t}) \leq 16 \cdot q^3 \cdot w(\mathbf{t}) \leq 2^8 \cdot q^5 \cdot w(\mathbf{t}) \cdot V_{\mathbf{t}},$$

where the last inequality follows since $\mathbf{t} \notin \text{Small}$ (and therefore $16 \cdot q^2 \cdot V_{\mathbf{t}} \geq 1$). \square

To prove that \mathcal{E} satisfies the second property of Lemma 13, we present a pair of random variables $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}}$ and $Y_{\mathbf{t}}^{\mathcal{B}^{Ideal}}$ such that the following holds for λ (the root of \mathcal{T}):

1. $Y_{\lambda}^{\mathcal{A}^{Ideal}} \cdot Y_{\lambda}^{\mathcal{B}^{Ideal}} \geq V_{\lambda}$, and
2. $V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{Ideal}}} \geq Y_{\lambda}^{\mathcal{A}^{Ideal}} - 1/2q$ and $V_{\lambda}^{\mathcal{B}^{\mathcal{E}^{Ideal}}} \geq Y_{\lambda}^{\mathcal{B}^{Ideal}} - 1/2q$.

The variables $Y_{\lambda}^{\mathcal{A}^{Ideal}}$ and $Y_{\lambda}^{\mathcal{B}^{Ideal}}$ are defined below, but intuitively they measure the success probability of \mathcal{A}^{Ideal} and \mathcal{B}^{Ideal} , respectively, in the mental experiment where their success probability on internal nodes outside \mathcal{E} is changed according to the informal description above. The above immediately yields that $V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{Ideal}}} \cdot V_{\lambda}^{\mathcal{B}^{\mathcal{E}^{Ideal}}} \geq V_{\lambda} - \frac{1}{q}$, completing the proof of Lemma 13.

Since our goal is to bound (from below) the success probabilities of $\mathcal{A}_{\mathcal{E}}^{Ideal}$ and $\mathcal{B}_{\mathcal{E}}^{Ideal}$, it suffices to restrict the discussion to the nodes in \mathcal{T} that have nonzero probability of being reached in executions with $\mathcal{A}_{\mathcal{E}}^{Ideal}$ and $\mathcal{B}_{\mathcal{E}}^{Ideal}$. This set of nodes defines a tree (which is defined below and denoted \mathcal{T}') that can alternatively be defined as the set of all nodes in \mathcal{T} that have no proper ancestor in $\text{Small} \cup \text{UnBal}$. We use the following random variables.

DEFINITION 16. For $\mathbf{t} \in \mathcal{T}' := \text{Supp}((\mathcal{A}, \mathcal{B})(1^n)) \cap \text{Supp}((\mathcal{A}, \mathcal{B})(1^n)) \subseteq \mathcal{T}$,⁵ we define $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}}$ as follows ($Y_{\mathbf{t}}^{\mathcal{B}^{Ideal}}$ is defined analogously):

- If $\mathbf{t} \in \mathcal{E}$, then
 1. if \mathbf{t} is a leaf, $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}} = V_{\mathbf{t}}$;
 2. otherwise, $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}} = \Pr[\mathcal{A}^{Ideal}(\mathbf{t}) = 1] \cdot Y_{\mathbf{t}_{o1}}^{\mathcal{A}^{Ideal}} + \Pr[\mathcal{A}^{Ideal}(\mathbf{t}) = 0] \cdot Y_{\mathbf{t}_{o0}}^{\mathcal{A}^{Ideal}}$.
- If $\mathbf{t} \in \text{UnBal}$, then
 1. if $\mathbf{t} \in \text{UnBal}_{\mathcal{A}}$, $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}} = \frac{1}{4q}$;
 2. otherwise ($\mathbf{t} \in \text{UnBal}_{\mathcal{B}}$), $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}} = 4q$.
- Otherwise ($\mathbf{t} \in \text{Small}$), $Y_{\mathbf{t}}^{\mathcal{A}^{Ideal}} = \frac{1}{4q}$.

⁵We assume without loss of generality that an honest party aborts if the other party does. Hence, \mathcal{T}' is indeed contained in \mathcal{T} .

We emphasize that the adversaries $\mathcal{A}^{\text{Ideal}}$ and $\mathcal{B}^{\text{Ideal}}$ remain exactly as before, and the random variables $Y_t^{\mathcal{A}^{\text{Ideal}}}$ and $Y_t^{\mathcal{B}^{\text{Ideal}}}$ only enable us to present a refined analysis of their success probabilities. The following fact easily follows from arguments similar to those used in the proof of Lemma 12.

CLAIM 17. For any $\mathbf{t} \in \mathcal{T}'$, it holds that

$$Y_t^{\mathcal{A}^{\text{Ideal}}} \cdot Y_t^{\mathcal{B}^{\text{Ideal}}} \geq V_t.$$

Proof. The proof is by induction up the protocol tree. For a node $\mathbf{t} \notin \mathcal{E}$, the lemma is trivially true since $Y_t^{\mathcal{A}^{\text{Ideal}}} \cdot Y_t^{\mathcal{B}^{\text{Ideal}}} \geq V_t$. For any other node \mathbf{t} (without loss of generality \mathbf{t} is an A node), the proof follows from exactly the same argument as in Lemma 12. This is true since for the base cases nothing has changed, and for an internal node \mathbf{t} it holds that

- $Y_t^{\mathcal{A}^{\text{Ideal}}} = \Pr[\mathcal{A}^{\text{Ideal}}(\mathbf{t}) = 1] \cdot Y_{t_{01}}^{\mathcal{A}^{\text{Ideal}}} + \Pr[\mathcal{A}^{\text{Ideal}}(\mathbf{t}) = 0] \cdot Y_{t_{00}}^{\mathcal{A}^{\text{Ideal}}}$, and
- $Y_t^{\mathcal{B}^{\text{Ideal}}} = \Pr[\mathcal{A}(\mathbf{t}) = 1] \cdot Y_{t_{01}}^{\mathcal{B}^{\text{Ideal}}} + \Pr[\mathcal{A}(\mathbf{t}) = 0] \cdot Y_{t_{00}}^{\mathcal{B}^{\text{Ideal}}}$.

Hence, the proof of the induction step follows exactly as in the proof of Lemma 12, which uses no property of the children of \mathbf{t} other than that they satisfy the induction hypothesis. \square

To complete the proof of Lemma 13, we need to prove that the success probability of both $\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}$ and $\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}$ is not far from the above mental experiment. We prove the following lemma.

LEMMA 18. It holds that $V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} \geq Y_{\lambda}^{\mathcal{A}^{\text{Ideal}}} - 1/2q$ and $V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}} \geq Y_{\lambda}^{\mathcal{B}^{\text{Ideal}}} - 1/2q$.

Proof. The main tool we are using for proving Lemma 18 is the following ‘‘compensation lemma.’’

LEMMA 19 (compensation lemma). Let the relative weights of $\mathbf{t} \in \mathcal{T}$ be as above (i.e., $Wrel^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) = \frac{w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t})}{w(\mathbf{t})}$ and $Wrel^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t}) = \frac{w^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t})}{w(\mathbf{t})}$). The following holds for every $\mathbf{t} \in \mathcal{T}$:

$$Wrel^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot Wrel^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t}) = \frac{V_t}{V_{\lambda}}.$$

Namely, the lemma states that a node \mathbf{t} whose weight with respect to $\mathcal{A}^{\text{Ideal}}$ is k times larger than its typical weight (i.e., $w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) > k \cdot w(\mathbf{t})$) has weight with respect to $\mathcal{B}^{\text{Ideal}}$ that is (close to) k times smaller than its typical weight. The proof of Lemma 19 is given later. We first use the lemma for completing the proof of Lemma 18.

In the following we focus on analyzing the value of $V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}}$ (the proof of $V_{\lambda}^{\mathcal{B}_{\mathcal{E}}^{\text{Ideal}}}$ is done analogously). Let \mathcal{F} be the set of leaves in \mathcal{T}' . That is, \mathcal{F} contains nodes of two types: (i) a leaf \mathbf{t} of the original tree \mathcal{T} (such that there is no ancestor \mathbf{t}' of \mathbf{t} in $\text{Small} \cup \text{UnBal}$), and (ii) a node $\mathbf{t} \in \text{Small} \cup \text{UnBal}$ (such that there is no ancestor $\mathbf{t}' \neq \mathbf{t}$ of \mathbf{t} in $\text{Small} \cup \text{UnBal}$). Furthermore, any execution $(\mathcal{A}^{\text{Ideal}}, \mathcal{B})$ passes through a node in \mathcal{F} . It follows that

$$(11) \quad Y_{\lambda}^{\mathcal{A}^{\text{Ideal}}} = \sum_{\mathbf{t} \in \mathcal{F}} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot Y_t^{\mathcal{A}^{\text{Ideal}}},$$

$$(12) \quad V_{\lambda}^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}} = \sum_{\mathbf{t} \in \mathcal{F}} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot V_t^{\mathcal{A}_{\mathcal{E}}^{\text{Ideal}}}.$$

Let

- $\mathcal{F}_1 = \mathcal{F} \cap \text{UnBal}_{\mathcal{B}}$,
- $\mathcal{F}_2 = \mathcal{F} \cap (\text{UnBal}_{\mathcal{A}} \cup \text{Small})$, and

- $\mathcal{F}_3 = \mathcal{F} \setminus (\mathcal{F}_1 \cup \mathcal{F}_2) = \mathcal{F} \cap \mathcal{E}$.

Lemma 19 yields that $\text{UnBal}_{\mathcal{A}}$ and $\text{UnBal}_{\mathcal{B}}$ are disjoint. It follows that $\mathcal{F}_1, \mathcal{F}_2$, and \mathcal{F}_3 form a partition of \mathcal{F} , and (11) yields that

$$\begin{aligned}
 Y_{\lambda}^{\mathcal{A}^{\text{Ideal}}} &= \sum_{\mathbf{t} \in \mathcal{F}_1} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot Y_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} + \sum_{\mathbf{t} \in \mathcal{F}_2} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot Y_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} + \sum_{\mathbf{t} \in \mathcal{F}_3} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot Y_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} \\
 &\leq \sum_{\mathbf{t} \in \mathcal{F}_1} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot 4q + \sum_{\mathbf{t} \in \mathcal{F}_2} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot \frac{1}{4q} + \sum_{\mathbf{t} \in \mathcal{F}_3} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) \cdot V_{\mathbf{t}}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \\
 &\leq 4q \cdot \sum_{\mathbf{t} \in \mathcal{F}_1} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) + \frac{1}{4q} \cdot \sum_{\mathbf{t} \in \mathcal{F}_2} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) + V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \\
 (13) \quad &\leq 4q \cdot \sum_{\mathbf{t} \in \mathcal{F}_1} w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) + \frac{1}{4q} + V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}},
 \end{aligned}$$

where the first inequality follows from Definition 16, which yields that $Y_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} = 4q$ for any $\mathbf{t} \in \mathcal{F}_1$, that $Y_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} = 1/4q$ for any $\mathbf{t} \in \mathcal{F}_2$, and that $Y_{\mathbf{t}}^{\mathcal{A}^{\text{Ideal}}} = V_{\mathbf{t}}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}}$ for any $\mathbf{t} \in \mathcal{F}_3$. The second inequality follows from (12).

We next consider the probability of visiting \mathcal{F}_1 in a random execution of $(\mathcal{A}^{\text{Ideal}}, \mathcal{B})$. The definition of $\text{UnBal}_{\mathcal{B}}$ yields that $Wrel^{\mathcal{B}^{\text{Ideal}}}(\mathbf{t}) > 16 \cdot q^3$ for any $\mathbf{t} \in \mathcal{F}_1$. Applying Lemma 19 yields that

$$(14) \quad \frac{w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t})}{w(\mathbf{t})} = Wrel^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) < \frac{1}{16 \cdot q^2} \cdot \frac{V_{\mathbf{t}}}{V_{\lambda}}$$

for any $\mathbf{t} \in \mathcal{F}_1$. Since $V_{\mathbf{t}} \leq 1$ and $\frac{1}{V_{\lambda}} \leq q$, we have that $w^{\mathcal{A}^{\text{Ideal}}}(\mathbf{t}) < \frac{w(\mathbf{t})}{16 \cdot q^2}$. Plugging this into (13) yields that

$$\begin{aligned}
 Y_{\lambda}^{\mathcal{A}^{\text{Ideal}}} &< 4q \cdot \sum_{\mathbf{t} \in \mathcal{F}_1} \frac{w(\mathbf{t})}{16 \cdot q^2} + \frac{1}{4q} + V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \\
 &= \frac{4q}{16 \cdot q^2} \cdot \sum_{\mathbf{t} \in \mathcal{F}_1} w(\mathbf{t}) + \frac{1}{4q} + V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \\
 &\leq \frac{1}{4 \cdot q} + \frac{1}{4 \cdot q} + V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \\
 &= \frac{1}{2 \cdot q} + V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}}.
 \end{aligned}$$

Hence, $V_{\lambda}^{\mathcal{A}^{\mathcal{E}^{\text{Ideal}}}} \geq Y_{\lambda}^{\mathcal{A}^{\text{Ideal}}} - 1/2q$, as desired. \square

5.1.1. Putting it all together. We next summarize the arguments that conclude the proof of Lemma 13.

Proof of Lemma 13. Let \mathcal{E} be defined as in the foregoing discussion. Claim 15 asserts that \mathcal{E} satisfies the first requirement of Lemma 13. For the second requirement,

Lemma 18 yields that $V_\lambda^{\mathcal{A}^\varepsilon \text{Ideal}} \geq Y_\lambda^{\mathcal{A}^\text{Ideal}} - \frac{1}{2q}$ and $V_\lambda^{\mathcal{B}^\varepsilon \text{Ideal}} \geq Y_\lambda^{\mathcal{B}^\text{Ideal}} - \frac{1}{2q}$. Hence, we have

$$\begin{aligned} V_\lambda^{\mathcal{A}^\varepsilon \text{Ideal}} \cdot V_\lambda^{\mathcal{B}^\varepsilon \text{Ideal}} &\geq \left(Y_\lambda^{\mathcal{A}^\text{Ideal}} - \frac{1}{2q} \right) \cdot \left(Y_\lambda^{\mathcal{B}^\text{Ideal}} - \frac{1}{2q} \right) \\ &= Y_\lambda^{\mathcal{A}^\text{Ideal}} \cdot Y_\lambda^{\mathcal{B}^\text{Ideal}} - \frac{Y_\lambda^{\mathcal{A}^\text{Ideal}} + Y_\lambda^{\mathcal{B}^\text{Ideal}}}{2q} + \left(\frac{1}{2q} \right)^2 \\ &\geq Y_\lambda^{\mathcal{A}^\text{Ideal}} \cdot Y_\lambda^{\mathcal{B}^\text{Ideal}} - \frac{2}{2q}. \end{aligned}$$

Claim 17 asserts that $Y_\lambda^{\mathcal{A}^\text{Ideal}} \cdot Y_\lambda^{\mathcal{B}^\text{Ideal}} \geq V_\lambda$, and hence, the second requirement is also satisfied; i.e.,

$$V_\lambda^{\mathcal{A}^\varepsilon \text{Ideal}} \cdot V_\lambda^{\mathcal{B}^\varepsilon \text{Ideal}} \geq V_\lambda - \frac{1}{q}. \quad \square$$

5.1.2. The proof of the compensation lemma.

Proof of Lemma 19. For $\mathbf{t} \in \mathcal{T}$ and $c \in \{0, 1\}$, let $\beta_{\mathbf{t}}(c)$ be the probability that the next message is c given that the transcript until now has been \mathbf{t} . That is,

$$(15) \quad \beta_{\mathbf{t}}(c) = \Pr_{(r_A, r_B) \leftarrow \text{Uni}(\mathbf{t})} [\text{Trans}(r_A, r_B)_{|\mathbf{t}+1} = \mathbf{t} \circ c].$$

Recall that $w(\mathbf{t})$ is the probability that \mathbf{t} is a prefix of the full communication transcript in an honest execution of the protocol. Assume that $\mathbf{t} = c_1 c_2 \dots c_\ell$; then $w(\mathbf{t}) = \beta_{\mathbf{t}_0}(c_1) \cdot \beta_{\mathbf{t}_1}(c_2) \cdot \dots \cdot \beta_{\mathbf{t}_{\ell-1}}(c_\ell)$.

Consider now an execution of $(\mathcal{A}^\text{Ideal}, \mathcal{B})$. For $c \in \{0, 1\}$, let $\beta_{\mathbf{t}}^{\mathcal{A}^\text{Ideal}}(c)$ be the probability that the next message is c given that the transcript until now has been \mathbf{t} . That is,

$$(16) \quad \beta_{\mathbf{t}}^{\mathcal{A}^\text{Ideal}}(c) = \Pr[\mathcal{A}^\text{Ideal}(\mathbf{t}) = c].$$

Recall that $w^{\mathcal{A}^\text{Ideal}}(\mathbf{t})$ is the probability that the node \mathbf{t} is reached in an execution of $(\mathcal{A}^\text{Ideal}, \mathcal{B})$. It follows that

$$w^{\mathcal{A}^\text{Ideal}}(\mathbf{t}) = \beta_{\mathbf{t}_0}^{\mathcal{A}^\text{Ideal}}(c_1) \cdot \beta_{\mathbf{t}_1}^{\mathcal{A}^\text{Ideal}}(c_2) \cdot \dots \cdot \beta_{\mathbf{t}_{\ell-1}}^{\mathcal{A}^\text{Ideal}}(c_\ell).$$

Note that, if \mathbf{t} is an \mathbf{A} node, then $\beta_{\mathbf{t}}^{\mathcal{A}^\text{Ideal}}(c) = \frac{\beta_{\mathbf{t}}(c) \cdot V_{\mathbf{t} \circ c}}{V_{\mathbf{t}}}$, and otherwise $\beta_{\mathbf{t}}^{\mathcal{A}^\text{Ideal}}(c) = \beta_{\mathbf{t}}(c)$. It follows that

$$\begin{aligned} Wrel^{\mathcal{A}^\text{Ideal}}(\mathbf{t}) &= \frac{w^{\mathcal{A}^\text{Ideal}}(\mathbf{t})}{w(\mathbf{t})} \\ &= \frac{1}{w(\mathbf{t})} \cdot \beta_{\mathbf{t}_0}^{\mathcal{A}^\text{Ideal}}(c_1) \beta_{\mathbf{t}_1}^{\mathcal{A}^\text{Ideal}}(c_2) \cdot \dots \cdot \beta_{\mathbf{t}_{\ell-2}}^{\mathcal{A}^\text{Ideal}}(c_{\ell-1}) \cdot \beta_{\mathbf{t}_{\ell-1}}^{\mathcal{A}^\text{Ideal}}(c_\ell) \\ &= \frac{1}{w(\mathbf{t})} \cdot \frac{\beta_{\mathbf{t}_0}(c_1) \cdot V_{\mathbf{t}_1}}{V_{\mathbf{t}_0}} \beta_{\mathbf{t}_1}(c_2) \cdot \dots \cdot \frac{\beta_{\mathbf{t}_{\ell-2}}(c_{\ell-1}) \cdot V_{\mathbf{t}_{\ell-1}}}{V_{\mathbf{t}_{\ell-2}}} \beta_{\mathbf{t}_{\ell-1}}(c_\ell) \\ &= \frac{\beta_{\mathbf{t}_0}(c_1) \beta_{\mathbf{t}_1}(c_2) \cdot \dots \cdot \beta_{\mathbf{t}_{\ell-2}}(c_{\ell-1}) \beta_{\mathbf{t}_{\ell-1}}(c_\ell)}{w(\mathbf{t})} \cdot \frac{V_{\mathbf{t}_1}}{V_{\mathbf{t}_0}} \cdot 1 \cdot \dots \cdot \frac{V_{\mathbf{t}_{\ell-1}}}{V_{\mathbf{t}_{\ell-2}}} \cdot 1. \end{aligned}$$

Since $w(\mathbf{t}) = \beta_{\mathbf{t}_0}(c_1) \beta_{\mathbf{t}_1}(c_2) \cdot \dots \cdot \beta_{\mathbf{t}_{\ell-2}}(c_{\ell-1}) \beta_{\mathbf{t}_{\ell-1}}(c_\ell)$, we obtain that

$$Wrel^{\mathcal{A}^\text{Ideal}}(\mathbf{t}) = \prod_{i=1}^{\ell/2} \frac{V_{\mathbf{t}_{2i-1}}}{V_{\mathbf{t}_{2i-2}}}.$$

A similar argument shows that

$$W_{\text{rel}}^{\mathcal{B}^{\text{ideal}}}(\mathbf{t}) = \prod_{i=1}^{\ell/2} \frac{V_{\mathbf{t}_{2i}}}{V_{\mathbf{t}_{2i-1}}}.$$

Hence, we conclude that

$$W_{\text{rel}}^{\mathcal{A}^{\text{ideal}}}(\mathbf{t}) \cdot W_{\text{rel}}^{\mathcal{B}^{\text{ideal}}}(\mathbf{t}) = \prod_{i=1}^{\ell/2} \frac{V_{\mathbf{t}_{2i-1}}}{V_{\mathbf{t}_{2i-2}}} \cdot \prod_{i=1}^{\ell/2} \frac{V_{\mathbf{t}_{2i}}}{V_{\mathbf{t}_{2i-1}}} = \prod_{i=1}^{\ell} \frac{V_{\mathbf{t}_i}}{V_{\mathbf{t}_{i-1}}} = \frac{V_{\mathbf{t}}}{V_{\lambda}}. \quad \square$$

6. Discussion and open questions. The main open question is understanding the limits of efficient attacks in breaking coin-flipping protocols. Specifically (assuming one-way functions do not exist), does there exist, for any (correct) coin-flipping protocol, an efficient adversary that biases its output towards 0 *or* towards 1 by $\frac{1}{2} - 1/\text{poly}$, or even by $\frac{\sqrt{2}-1}{2} + \Omega(1)$? In light of the reduction of Chailloux and Kerenidis [5] from $(\frac{\sqrt{2}-1}{2} + \sqrt{2} \cdot \varepsilon + o(\varepsilon))$ -bias *strong* coin-flipping to ε -bias *weak* coin-flipping, a positive answer, even to the weaker form of the above question, would mean that constant-bias weak coin-flipping protocols imply the existence of one-way functions. More specifically, if an adversary can always bias the output of a coin-flipping protocol by $\frac{\sqrt{2}-1}{2} + c$ for some $c \in \Omega(1)$ (assuming one-way functions do not exist), then, for any ε such that $\sqrt{2} \cdot \varepsilon + o(\varepsilon) < c$, the existence of ε -bias weak coin-flipping protocols implies the existence of one-way functions.

While our analysis proves only the existence of an adversary achieving $\frac{\sqrt{2}-1}{2} - o(1)$ bias (and thus has no direct implication for weak coin flipping), it shows that (assuming one-way functions do not exist) for any coin-flipping protocol there exists an efficient adversary that can bias its output *both* towards 0 and towards 1, by $\frac{\sqrt{2}-1}{2} - o(1)$. Hence, our attack accomplishes a harder task than the required one. Interestingly, $\frac{\sqrt{2}-1}{2}$ is the right bound for this more challenging task. That is, there exists a (correct) coin-flipping protocol for which no adversary (not even an unbounded one) can bias the output towards 1 by more than $\frac{\sqrt{2}-1}{2}$.⁶

Acknowledgment. We thank Hemanta Maji for useful discussions.

REFERENCES

- [1] B. AVERBUCH, M. BLUM, B. CHOR, S. GOLDWASSER, AND S. MICALI, *How to Implement Bracha's $O(\log n)$ Byzantine Agreement Algorithm*, unpublished manuscript, 1985.
- [2] A. BEIMEL, E. OMRI, AND I. ORLOV, *Protocols for multiparty coin toss with dishonest majority*, in *Advances in Cryptology—CRYPTO 2010*, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 538–557.
- [3] M. BEN-OR AND N. LINIAL, *Collective coin flipping, robust voting schemes and minima of Banzhaf values*, in *Proceedings of the 26th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 1985, pp. 408–416.
- [4] M. BLUM, *Coin flipping by telephone*, in *Advances in Cryptology—CRYPTO '81*, Springer-Verlag, Berlin, Heidelberg, 1981, pp. 11–15.

⁶For instance, consider the protocol where A (playing first) sets the outcome of the protocol to 0 with probability $1 - \frac{1}{\sqrt{2}}$ and defers the decision to B otherwise. Then, party B, if plays, sets the outcome to 1 with probability $\frac{1}{\sqrt{2}}$ and to 0 otherwise. It is clear that the protocol is correct (i.e., expected outcome for an honest execution is $\frac{1}{2}$) and that there exists no cheating strategy for either A or B that can make the expected outcome of the protocol to be larger than $\frac{1}{\sqrt{2}}$.

- [5] A. CHAILLOUX AND I. KERENIDIS, *Optimal quantum strong coin flipping*, in Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2009, pp. 527–533.
- [6] K.-M. CHUNG AND F.-H. LIU, *Parallel repetition theorems for interactive arguments*, in Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 19–36.
- [7] 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 (STOC), 1986, pp. 364–369.
- [8] R. CLEVE AND R. IMPAGLIAZZO, *Martingales, Collective Coin Flipping and Discrete Control Processes*, manuscript, 1993.
- [9] D. DACHMAN-SOLED, Y. LINDELL, M. MAHMOODY, AND T. MALKIN, *On the black-box complexity of optimally-fair coin tossing*, in Theory of Cryptography, 8th Theory of Cryptography Conference, TCC 2011, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 450–467.
- [10] O. GOLDREICH, S. GOLDWASSER, AND S. MICALI, *On the cryptographic applications of random functions*, in Advances in Cryptology—CRYPTO '84, Springer-Verlag, Berlin, Heidelberg, 1984, pp. 276–288.
- [11] O. GOLDREICH, S. GOLDWASSER, AND S. MICALI, *How to construct random functions*, J. ACM, 33 (1986), pp. 792–807.
- [12] O. GOLDREICH AND L. A. LEVIN, *A hard-core predicate for all one-way functions*, in Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC), 1989, pp. 25–32.
- [13] I. HAITNER, *A parallel repetition theorem for any interactive argument*, in Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2009, pp. 241–250.
- [14] I. HAITNER, M.-H. NGUYEN, S. J. ONG, O. REINGOLD, AND S. VADHAN, *Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function*, SIAM J. Comput., 39 (2009), pp. 1153–1218.
- [15] I. HAITNER AND E. OMRI, *Coin flipping with constant bias implies one-way functions*, in Proceedings of the 2011 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2011, pp. 110–119.
- [16] J. HÅSTAD, R. IMPAGLIAZZO, L. A. LEVIN, AND M. LUBY, *A pseudorandom generator from any one-way function*, SIAM J. Comput., 28 (1999), pp. 1364–1396.
- [17] J. HÅSTAD, R. PASS, D. WIKSTRÖM, AND K. PIETRZAK, *An efficient parallel repetition theorem*, in Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 1–18.
- [18] R. IMPAGLIAZZO, *Pseudo-Random Generators for Cryptography and for Randomized Algorithms*, Ph.D. thesis, Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, 1992.
- [19] R. IMPAGLIAZZO AND M. LUBY, *One-way functions are essential for complexity based cryptography*, in Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 1989, pp. 230–235.
- [20] A. Y. KITAEV, *Quantum coin-flipping*, presentation at the Workshop on Quantum Information Processing, Mathematical Sciences Research Institute, Berkeley, CA, 2002.
- [21] H. K. MAJI, M. PRABHAKARAN, AND A. SAHAI, *On the computational complexity of coin flipping*, in Proceedings of the 2010 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2010, pp. 613–622.
- [22] C. MOCHON, *Quantum Weak Coin Flipping with Arbitrarily Small Bias*, preprint, arXiv:0711.4114v1 [quant-ph], 2007.
- [23] T. MORAN, M. NAOR, AND G. SEGEV, *An optimally fair coin toss*, in Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 1–18.
- [24] M. NAOR, *Bit commitment using pseudorandomness*, J. Cryptology, 4 (1991), pp. 151–158.
- [25] M. NAOR AND M. YUNG, *Universal one-way hash functions and their cryptographic applications*, in Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC), 1989, pp. 33–43.
- [26] J. ROMPEL, *One-way functions are necessary and sufficient for secure signatures*, in Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC), 1990, pp. 387–394.
- [27] S. ZACHOS, *Probabilistic quantifiers, adversaries, and complexity classes: An overview*, in Proceedings of the First Annual IEEE Conference on Computational Complexity, 1986, pp. 383–400.