Can Optimally-Fair Coin Tossing Be Based on One-Way Functions?

Dana Dachman-Soled¹, Mohammad Mahmoody², and Tal Malkin³

¹ University of Maryland
danadach@ece.umd.edu

² University of Virginia
mohammad@cs.virginia.edu

³ Columbia University and Bar-Ilan University
tal@cs.columbia.edu

Abstract. Coin tossing is a basic cryptographic task that allows two distrustful parties to obtain an unbiased random bit in a way that neither party can bias the output by deviating from the protocol or halting the execution. Cleve [STOC'86] showed that in any r round coin tossing protocol one of the parties can bias the output by $\Omega(1/r)$ through a “fail-stop” attack; namely, they simply execute the protocol honestly and halt at some chosen point. In addition, relying on an earlier work of Blum [COMPCON'82], Cleve presented an r -round protocol based on one-way functions that was resilient to bias at most $O(1/\sqrt{r})$. Cleve’s work left open whether “optimally-fair” coin tossing (i.e. achieving bias $O(1/r)$ in r rounds) is possible. Recently Moran, Naor, and Segev [TCC'09] showed how to construct optimally-fair coin tossing based on oblivious transfer, however, it was left open to find the *minimal* assumptions necessary for optimally-fair coin tossing. The work of Dachman-Soled et al. [TCC'11] took a step toward answering this question by showing that any black-box construction of optimally-fair coin tossing based on a one-way functions with n -bit input and output needs $\Omega(n/\log n)$ rounds.

In this work we take another step towards understanding the complexity of optimally-fair coin-tossing 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. Namely, we consider a natural class of black-box constructions based on one-way functions, called *function oblivious*, in which the output of the protocol does not depend on the specific implementation of the one-way function and only depends on the randomness of the parties. Other than being a natural notion on its own, the known coin tossing protocols of Blum and Cleve (both based on one-way functions) are indeed function oblivious. Thus, we believe our lower bound for function-oblivious constructions is a meaningful step towards resolving the fundamental open question of the complexity of optimally-fair coin tossing.

Keywords: Coin-Tossing, One-Way Functions, Black-Box Separations.

1 Introduction

In this work, we address the fundamental problem of secure, two-party coin-tossing, where two mutually distrustful parties wish to generate a common random bit. A secure

coin-tossing scheme has the following complementary properties: (1) Security—even if one of the parties deviates arbitrarily from the protocol, the output bit of the honest party should be almost completely unbiased (namely be equal to 1 with probability that is at most negligibly far from $1/2$) and (2) Correctness—when both parties follow the protocol they are guaranteed to output the same random bit. Unfortunately, a classic result by Cleve [Cle86] shows that even if the party which deviates from the protocol misbehaves only by choosing whether or not to abort early (this is known as a fail-stop adversary), then secure coin-tossing cannot be achieved. In particular, Cleve proved that for any coin tossing protocol running for $\hat{r}$ rounds, there exists an efficient fail-stop adversary that can bias the output bit of the honest party by at least $\Omega(1/\hat{r})$.

It turns out that in a weaker model we can, in fact, construct secure coin-tossing protocols. An early result by Blum [Blu82] uses one-way functions (OWF) to construct a *weak* coin tossing protocol where no party can increase the probability of 0 or 1 by more than negligible by deviating from the protocol or halting. In a weak coin tossing protocol whenever a party aborts, the other party is not required output anything. Weak coin tossing can be useful for scenarios where Alice and Bob each have a preferred outcome in mind (e.g., Alice wants 0 and Bob wants 1) simply because if any party aborts the other one can take their desired outcome as the output. However, note that the weak coin tossing definition doesn't preclude the possibility that a malicious party can cause the output to always be either 0 or abort; indeed this is the case for Blum's protocol (where a malicious party can discover first the emerging output and then choose whether to abort or continue). In contrast, in a *strong* coin tossing protocol (which is the focus of our work), the protocol always requires an output. A strong coin tossing protocol with bias at most δ is one where each honest party always announces an output (even if the other party aborted), and yet no malicious party can bias the honest party's output (in any direction) by more than δ . The weak coin tossing protocol of Blum was used as a building block by Cleve [Cle86] to construct a *strong* coin tossing protocol that (for any polynomial $\hat{r}$) runs for $\hat{r}$ rounds and for which no efficient adversary can bias the output bit by more than $O(1/\sqrt{\hat{r}})$. In our work, whenever not explicitly mentioned, we are referring to *strong* coin tossing protocols.

The question of closing the gap between this best known upper bound ($O(1/\sqrt{\hat{r}})$ based on OWF) and lower bound ($\Omega(1/\hat{r})$ regardless of any assumption) remained unresolved for more than two decades. A few years ago, the gap was closed by Moran et al. [MNS09] who constructed a protocol for coin tossing whose bias matches the lower-bound of [Cle86]. Specifically, for any $\hat{r}$ they constructed an $O(\hat{r})$ -round protocol with the property that no efficient adversary can bias the output by more than $O(1/\hat{r})$. Thus, they demonstrated that the $O(1/\hat{r})$ lower-bound is tight. We call a protocol which achieves bias $O(1/\hat{r})$ *optimally-fair*, because no protocol can achieve asymptotically-lower bias. The protocol of [MNS09], however, uses general secure computation and thus requires the strong assumption that protocols for oblivious transfer exist. In contrast, the coin tossing protocol of Blum [Blu82] and the protocol of [Cle86] achieving bias of $O(1/\sqrt{\hat{r}})$ can be constructed from any one-way function, and in fact, rely only on the existence of a commitment scheme. This leads us to our main question:

Can optimally-fair coin tossing be based on one-way functions?

This question was already asked by Moran et al. [MNS09] as a challenging open problem. Indeed, the question of whether one-way functions suffice for optimally-fair coin-tossing seems to be a difficult problem and remains open, despite much effort. A partial answer to the main question above was presented in the work of Dachman-Soled et al. [DSLMM11]. Informally, they show that if C is a black-box construction of optimally-fair coin tossing based on one-way functions with input and output length n , then the number of rounds of interaction in C is at least $\Omega(n/\log n)$. Thus, they rule out such black-box constructions with a “small” number of rounds. However, their results say nothing about constructions with a higher number of rounds. For example, they do not rule out the possibility of constructing coin-tossing protocols from one-way functions of input size n , which have $\hat{r} = n^3$ number of rounds and for which no efficient adversary can bias the output by more than $1/n^3 = 1/\hat{r}$.

In this work, we make an important step towards answering our main question. In particular, we manage to remove the limitation on the round complexity in the impossibility result of [DSLMM11]. Indeed, we consider protocols with an arbitrary polynomial number of rounds: $\hat{r} = \text{poly}(n)$. However, we introduce another limitation: our impossibility results only rule out protocols which possess the following property.

Definition 1 (Function-Obliviousness:). *A coin-tossing protocol $C^f = \langle A^f, B^f \rangle$ based on one-way functions is called function-oblivious if the outcome of the coin tossing protocol $\langle A^f(r_A), B^f(r_B) \rangle$, when both parties are honest, depends only on the random tapes r_A, r_B of the two parties and not on the choice of one-way function f .*

Function-obliviousness captures the intuition that the one-way function f is being used only to achieve *security* for the coin-tossing protocol but does not affect *correctness*. In this work, we rule out (fully) black-box constructions of optimally-fair coin-tossing protocols which are function-oblivious from one-way functions.

Theorem 1 (Main Theorem, Informal). *There is no (fully) black-box and function-oblivious construction of optimally-fair coin-tossing protocols from one-way functions.*

Our result is incomparable to that of [DSLMM11]: we restrict ourselves to function-oblivious protocols but handle protocols with *arbitrary* polynomial number of rounds.

We believe that function-obliviousness is a natural assumption on coin-tossing protocols. Indeed, the known one-way-function based coin tossing protocols of Blum [Blu82] and Cleve [Cle86], as well as any other protocols based only on commitment schemes, are function-oblivious. The notion of function-obliviousness as defined in Definition 1 can be directly generalized to other pairs of primitives as well, and so understanding the limits of oblivious black-box constructions could be considered as a first step towards understanding the full power of black-box constructions. Thus, introducing the notion of oblivious black-box constructions, as a natural form of black-box constructions, is a conceptual contribution of our work. On a technical level, to deal with function-oblivious protocols, we introduce several new techniques which were not needed/applicable in the case of black-box $O(n/\log n)$ -round protocols. These techniques also may be of independent interest and indicate that we are making progress on a fundamental question by considering this class of protocols. Thus, we believe that our partial negative result is meaningful and improves our understanding of the relative complexity of one-way functions and optimally-fair coin-tossing protocols.

An important remaining open question is whether function-obliviousness is necessary for our result, or we can completely rule out any black-box construction of optimally-fair coin tossing from one-way functions. Our results together with those of [DSLMM11] indicate that if any such construction exists, it must have many ($\omega(n/\log n)$) rounds, and must use the one-way function in a novel way, not only for commitment but to determine the coin toss outcome even when both parties are honest.

1.1 Black-Box Separations

One of the main goals of modern cryptography has been to identify the minimal assumptions necessary to construct secure cryptographic primitives. For example, [Yao82, GM84, Rom90, HILL99, GGM86, LR88, IL89, NY89, Nao91] have shown that private key encryption, pseudorandom generators, pseudorandom functions and permutations, bit commitment, and digital signatures exist if and only if one-way functions exist. On the other hand, some cryptographic primitives such as public key encryption, oblivious transfer, and key agreement are not known to be equivalent to one way functions. Thus, it is natural to ask whether the existence of one-way functions implies these primitives. However, it seems unclear how to formalize such a question; since it is widely believed that both one-way functions and public key encryption exist, this would imply in a trivial logical sense that the existence of one-way functions implies the existence of public key encryption. Thus, we can only hope to rule out restricted types of constructions that are commonly used to prove implications in cryptography. Impagliazzo and Rudich [IR89] were the first to develop a framework and techniques to rule out the existence of an important class of reductions between primitives known as black-box reductions. Intuitively, this is a reduction where the primitive is treated as an oracle or a “black-box”. There are actually several flavors of black-box reductions (fully black-box, semi black-box and weakly black-box [RTV04]). In our work, we only deal with fully black-box reduction, and so we will focus on this notion here.

Informally, a fully black-box reduction from a primitive $\mathcal{Q}$ to a primitive $\mathcal{P}$ is a pair of oracle PPT machines $(G; S)$ such that the following two properties hold:

Correctness: For every implementation f of primitive $\mathcal{P}$, $g = G^f$ implements $\mathcal{Q}$.

Security: For every implementation f of primitive $\mathcal{P}$, and every adversary $\mathcal{A}$, if $\mathcal{A}$ breaks G^f (as an implementation of $\mathcal{Q}$) then $S^{\mathcal{A};f}$ breaks f .

We remark that an implementation of a primitive is any specific scheme that meets the syntactical requirements of that primitive (e.g., an implementation of a public-key encryption scheme provides samplability of key pairs, encryption with the public-key, and decryption with the private key). Correctness thus states that when G is given oracle access to any valid implementation of $\mathcal{P}$, the result is a valid implementation of $\mathcal{Q}$. Furthermore, security states that any adversary breaking G^f yields an adversary breaking f . The reduction here is *fully* black-box [RTV04] in the sense that the adversary S breaking f uses $\mathcal{A}$ in a black-box manner.

Separation from One-Way Functions. A common technique to separate a cryptographic primitive $\mathcal{P}$ from one-way functions is to show that any implementation of $\mathcal{P}$ in the random oracle can be broken by an attacker that asks “a few” (more specifically $2^{o(n)}$) queries to the random oracle (e.g. see [BM07] or [DSLMM11]). The reason, roughly speaking, is that if a $2^{o(n)}$ attacker Adv exists, the security reduction could turn Adv into a $2^{o(n)}$ -query attack to invert the random oracle, which is not possible [IR89, GT00].¹ We will also take this approach in this work.

1.2 Related Work

Cleve and Impagliazzo [CI93] showed that the bias $O(1/\sqrt{\hat{r}})$ is optimal when the attacker is *computationally unbounded*, and so their result does not resolve our main question.² However, using the result of [CI93], Dachman-Soled et al. [DSLMM11] gave a partial answer to the question of whether optimally-fair coin-tossing can be constructed in a black-box manner from one-way functions. As mentioned previously, they showed that if C is a construction of optimally-fair coin tossing based on one-way functions with input and output length n , then the number of rounds of interaction in C is at least $\Omega(n/\log n)$. More specifically, [DSLMM11] shows how to “compile out” the random oracle from the coin-tossing protocol by asking $(\text{poly}(n))^{\hat{r}}$ oracle queries by the parties where $\hat{r}$ is the number of rounds of the protocol. Note that whenever $\hat{r} = o(n/\log n)$, the number of queries asked by the parties will be $(\text{poly}(n))^{\hat{r}} = 2^{o(n)}$, and using the result of [CI93] for the no-oracle protocols one obtains a $2^{o(n)}$ -query attacker for one of the parties A or B in the with-oracle protocol which, as explained above, leads to a black-box separation. Unfortunately, the techniques of [DSLMM11] do not seem to extend to the case where $\hat{r} = \omega(n/\log n)$ since in this case, the adversarial strategy will require $(\text{poly}(n))^{\omega(n/\log n)} = 2^{\omega(n)}$ number of queries which is in fact enough to successfully invert the random oracle and does not lead to contradiction.

We also mention two other works, which deal with seemingly unrelated problems to ours, but leverage similar techniques. The first work of Haitner et al. [HOZ13] considers the question of constructing protocols for semi-honest no-input two-party computation in the random oracle model. They show that any semi-honest no-input two-party functionality which can be realized in the random oracle model, is trivial in the sense that, essentially, it can also be realized in the information-theoretic semi-honest setting with no random oracle. Note, however, that our coin-tossing in the semi-honest setting is trivial and our setting deals with malicious adversaries, so the result of [HOZ13] does not apply to our case. Mahmoody et al. [MMP13] consider semi-honest, deterministic functionalities with polynomial-sized domains and show that any such functionality which can be realized in the random oracle model is “trivial” in the same sense as above. As in our work, both of the above works utilize the “Eve” algorithm of [BM09] and rely on its specific properties (as described in Lemma 1). Moreover, some of the techniques in the work of [MMP13], where one of the players resamples a fake view and proceeds to compute using this fake view, are similar to our techniques. See the Technical Overview Section (Section 1.3) for additional details.

¹ Note that this technique works even if the attacker is not efficient.

² Although, as we will see, we use the results and approach of [CI93] as a starting point.

1.3 Technical Overview

We consider two-party coin tossing protocols $\mathcal{C} = \langle A, B \rangle$ with $\hat{r}$ rounds (i.e., $2\hat{r}$ messages). Let C denote the outcome of coin tossing protocol $\mathcal{C} = \langle A, B \rangle$ and let T_j denote the transcript of the protocol immediately after message j is sent. Moreover, since we consider the setting where one party may abort early, we denote by C_j the output of the other party, when one party aborts *before* sending the j -th message.

Cleve and Impagliazzo [CI93] showed that for any coin tossing protocol we have that with probability at least $1/5$ over the choice of random tapes of the parties, there is some point in the execution of the protocol that $|\mathbf{E}[C \mid T_j] - \mathbf{E}[C \mid T_{j+1}]| \geq \Omega(1/\sqrt{\hat{r}})$. Moreover, in [CI93], it was observed that if for some x, y , we have that an x -fraction of the executions of the coin-tossing protocol with uniformly chosen random tapes for the two parties reach a point where $|\mathbf{E}[C \mid T_j] - \mathbf{E}[C_j \mid T_j]| \geq y$ then the party who sends the message j has a strategy for biasing the output towards either 0 or 1 by $\Omega(x \cdot y)$ by aborting before sending the j -th message when the above event occurs. Thus, the fact that with probability $1/5$ there is some point in the execution such that $|\mathbf{E}[C \mid T_j] - \mathbf{E}[C \mid T_{j+1}]| \geq \Omega(1/\sqrt{\hat{r}})$ immediately implies a strategy for either A or B for imposing bias $\Omega(1/\sqrt{\hat{r}})$.

In our work, we extend the [CI93] observation in two ways. First, we allow a party to condition not just on the current transcript, but on its entire view and abort before sending message j when $|\mathbf{E}[C \mid V_{A,j}] - \mathbf{E}[C_j \mid V_{A,j}]| \geq y$, where A is the party sending the j -th message and $V_{A,j}$ denotes her partial view right before sending the j -th message. Additionally, we allow a party to abort immediately *after* sending a message. More specifically, we allow a party to abort immediately after sending message j when $|\mathbf{E}[C \mid T_j] - \mathbf{E}[C_{j+2} \mid T_j]| \geq y$. Although this is technically equivalent to waiting to get the $(j+1)$ -st message and aborting immediately after, it will be conceptually helpful to think of the party as aborting immediately after sending the j -th message. In what follows, we refer to the above described strategies as “[CI93]-type strategies”. We do not consider our extensions of the [CI93] strategies as our main technical contribution, but we consider them as useful tools for our proofs.

In our setting, we would like to apply the result of [CI93] in the random oracle model (see Section 1.1). The reason is that, it is well-known that, roughly speaking, (even inefficient) attacks in the random oracle model imply black-box separations from one-way functions. One might correctly say that it is in fact *impossible* to break a coin-tossing protocol in the random oracle model through a fail-stop attack, simply because the parties can trivially use oracle’s answer to a fixed query as their output. However, recall that: (1) our goal is to obtain black-box separations from one-way functions and the mentioned trivial protocol does not work when the random oracle is substituted with an actual one-way function, and (2) we are in fact focusing on function-oblivious protocols that prevent using the random oracle for obtaining the output.

Unfortunately, a straightforward implementation of [CI93] in the random oracle model (where expectations are taken also over the choice of oracle) fails due to the fact that in order for the [CI93] techniques to go through, it must be the case that $\mathbf{E}[C_j \mid T_{j-1}] = \mathbf{E}[C_j \mid T_j]$ (or at least that $|\mathbf{E}[C_j \mid T_{j-1}] - \mathbf{E}[C_j \mid T_j]| = o(1/\sqrt{\hat{r}})$) in all rounds. However, due to dependencies between parties’ views created by the random oracle (in addition to the dependencies created by the transcript) it may in fact be the

case that $|\mathbf{E}[C_j \mid T_{j-1}] - \mathbf{E}[C_j \mid T_j]| = \Omega(1/\sqrt{\hat{r}})$ in the random oracle model. I.e., the distribution over views of B till end of round j , denoted by $V_{B,j}$, conditioned on T_{j-1} may be very far from the the distribution over $V_{B,j}$ conditioned on T_j . This is due to the fact oracle answers received by A during the computation of T_j can affect the distribution over views of B even though A sends the j 'th message.

A natural approach to solving the above problem, would be to leverage the results of [IR89, BM09] on finding so-called “intersection queries” in 2-party protocols. Intersection queries are queries made by both parties A, B during an execution of a protocol in the random oracle model. Intuitively, it is these intersection queries that cause dependencies between the views of A and B. Moreover, in [IR89, BM09], it was shown that an eavesdropping adversary “Eve” can with high probability find all these intersection queries made by A and B, while making only a polynomial number of queries to the random oracle. Intuitively, one could hope that by running the “Eve” protocol of [IR89, BM09] after each pass of the protocol (which we call running the Eve protocol alongside the main protocol) these intersection queries could be found before they are made, thus eliminating dependencies between the views of A and B. It turns out that there is a subtle problem here: In order for [CI93] techniques to go through, we must prevent intersection queries even between Eve queries made alongside the j -th message (sent by A) and private queries that were made previously by B. Unfortunately, this property is not guaranteed by the Eve algorithm of [IR89, BM09]. However, [DSLMM11] showed that the Eve algorithm can be modified (becoming far less efficient) to guarantee that the above does not occur. Using such (inefficient) Eve [DSLMM11] still managed to rule out optimally-fair coin-tossing protocols with $O(n/\log n)$ number of rounds, where n is the input-output length of the one-way function. In this work, we shall find a different approach that allows us to deal with an arbitrary polynomial number of rounds.

Our Approach. In the following, $\mathcal{D}$ denotes the distribution over views of A, B running the coin-tossing protocol $\mathcal{C}$ with uniformly random coins. Additionally, for joint random variables X, Y , we denote by $X \mid Y$ the distribution over X drawn from $\mathcal{D}$, conditioned on Y . M_i denotes the i -th message of the coin-tossing protocol and T_i denotes the transcript which includes both the messages $M_1, \dots, M_i$ of the protocol $\mathcal{C}$ as well as Eve_i , the information “Eve” has learned by making her queries alongside in the first i messages. Finally, the partial view of a party after the i -th message is sent, denoted by $V_{A,i}$ or $V_{B,i}$ includes its random tape r_A or r_B , its queries to the oracle and the responses, as well as the transcript $M_1, \dots, M_i$ of the $\mathcal{C}$ protocol.

We consider the “middle value”: $MV = \mathbf{E}_{V_{B,j}, Eve_j \mid T_{j+1}}[\mathbf{E}[C \mid V_{B,j}, Eve_j]]$, where A sends the $(j+1)$ 'st message of the protocol. We shall clarify that for brevity, in our notation above Eve_j is consistent with T_{j+1} from which we are sampling $V_{B,j}$ (even though this his not explicitly mentioned). Intuitively, this means that we sample views of B, $V_{B,j}$, conditioned on the transcript at the $j+1$ -st pass (which A knows before B) and look at the expectation of the outcome of the coin-toss conditioned on these views, $V_{B,j}$. Then we take the expectation over these expected values. Here, we give some intuition as to why MV is significant for our analysis. Observe that B, given its real view can of course compute $\mathbf{E}[C \mid V_{B,j}, Eve_j] = \mathbf{E}[C \mid V_{B,j}, T_j]$, which is the expected value of the outcome of the coin-toss, C , from B's point-of-view. Note that A

cannot compute this value since conditioned on its view, it does not know the real $V_{B,j}$. Before computing its message in the $j + 1$ -st pass, A would only be able to compute $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T_j}[\mathbf{E}[C \mid V_{B,j}, \text{Eve}_j]]$, the expectation when views of B, $V_{B,j}$ are sampled conditioned on only T_j , which is equivalent to $\mathbf{E}[C | T_j]$. However, after computing T_{j+1} , A gets more information about B's view and thus can get a better estimate of B's real expected value by sampling views of B, $V_{B,j}$ conditioned on T_{j+1} . It is this advantage which we leverage in the final strategy in order to allow A to impose bias.

In the following we give some more details of our approach. First, using a similar argument to that of [CI93] shows that one of the following cases occurs:

1. With probability at least $1/20$ there is some point in the execution such that $\mathbf{E}[C | T_j] - MV \geq \Omega(1/\sqrt{r})$.
2. With probability at least $1/20$ there is some point in the execution such that $MV - \mathbf{E}[C | T_{j+1}] \geq \Omega(1/\sqrt{r})$.
3. With probability at least $1/20$ there is some point in the execution such that $MV - \mathbf{E}[C | T_j] \geq \Omega(1/\sqrt{r})$.
4. With probability at least $1/20$ there is some point in the execution such that $\mathbf{E}[C | T_{j+1}] - MV \geq \Omega(1/\sqrt{r})$.

For each of the above cases, we need to come up with corresponding strategies that allow A or B to impose bias on the final outcome. It turns out that Cases 1 and 2 give rise to adversarial strategies for biasing towards 0, while Cases 3 and 4 will give rise to adversarial strategies for biasing towards 1. In the following, we give some intuition for the analysis of cases 1 and 2; cases 3 and 4 are entirely analogous.

It is not difficult to see (details can be found in Section 4) that if Case 1 occurs, then one of the following will occur:

- With probability $\Omega(1/\hat{r}^{1/4})$ there is a (first) point where $\mathbf{E}[C \mid T_j] - \mathbf{E}[C_{j+2} \mid T_j] \geq \Omega(1/\sqrt{\hat{r}})$.
- With probability $\Omega(1/\hat{r}^{1/2})$ there is a (first) point where $\mathbf{E}[C \mid T_j] - \mathbf{E}[C_{j+2} \mid T_j] \geq \Omega(1/\hat{r}^{1/4})$.
- With probability $\Omega(1/\hat{r}^{1/4})$ there is a (first) point where $\mathbf{E}[C \mid V_{B,j}, \text{Eve}_j] - \mathbf{E}[C_{j+2} \mid V_{B,j}, \text{Eve}_j] \geq \Omega(1/\sqrt{\hat{r}})$.
- With probability $\Omega(1/\hat{r}^{1/4})$ there is a (first) point where $\mathbf{E}[C \mid V_{B,j}, \text{Eve}_j] - \mathbf{E}[C_{j+2} \mid V_{B,j}, \text{Eve}_j] \geq \Omega(1/\sqrt{\hat{r}})$.

By directly using [CI93]-type strategies we can impose bias of $\Omega(1/\hat{r}^{3/4})$ when any of the above items occurs. On the other hand, if Case 2 occurs then we have either:

- with probability at least $1/40$ there is a (first) point where $\mathbf{E}[C_{j+1} \mid T_{j+1}] - \mathbf{E}[C \mid T_{j+1}] \geq \Omega(1/\sqrt{\hat{r}})$, or:
- with probability at least $1/40$ there is a (first) point where $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T_{j+1}}[\mathbf{E}[C \mid V_{B,j}, \text{Eve}_j]] - \mathbf{E}[C_{j+1} \mid T_{j+1}] \geq \Omega(1/\sqrt{\hat{r}})$.

Again, if the first item above occurs, we can impose bias of $\Omega(1/\sqrt{\hat{r}})$ using [CI93]-type strategies. However, in order to utilize the second item above, which we refer to as Case (2b), to impose $\omega(1/\hat{r})$ bias, we need quite a bit of additional work. More specifically, we show that in order to leverage Case (2b), it is sufficient to present a way to simulate a fake transcripts T_{j+1} , which we denote by T'_{j+1} such that:

- Real transcripts T_{j+1} and fake transcripts T'_{j+1} are distributed nearly identically.
- The expected value of outcomes conditioned on views of B sampled w.r.t. the real transcript $T_{j+1} = t_{j+1}$ is nearly the same as the expected value of outcomes conditioned on views of B sampled w.r.t. $T'_{j+1} = t_{j+1}$. Formally, we have that: $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]] \approx \mathbf{E}_{V_{B,j}, \text{Eve}_j | T'_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]]$.
- T'_{j+1} reveals almost no information about the real $V_{A,j+1}$.

In what follows, we give some intuition as to how the simulated T'_{j+1} is constructed. We note that some of the techniques we use here are similar to those used by [MMP13]. In order to construct T'_{j+1} , we critically use independence of the views of A and B (once the Eve queries have been made). We sample a fake view for A, $V'_{A,j+1}$, conditioned only on T_j and use it to compute a fake next message M'_{j+1} . Then we run the Eve algorithm (pretending that M'_{j+1} is the real $j+1$ -st message) carefully choosing which queries to answer w.r.t. the real oracle and which queries to “lie” about. The main idea (although the actual algorithm is slightly more complicated) is the following: All queries made by Eve that are in $V'_{A,j+1}$ are answered according to $V'_{A,j+1}$, all queries made by Eve that are in the real $V_{A,j+1}$ and not in $V'_{A,j+1}$ are answered uniformly at random. All remaining queries are asked to the oracle and the response from the oracle is returned. Now, intuitively, items (1) and (2) above hold since by independence, it is highly likely that all “modified” Eve queries (i.e. queries that appear in $V_{A,j+1}$ or $V'_{A,j+1}$) *do not* intersect with the real $V_{B,j}$. For item (3), recall that T'_{j+1} is computed by “ignoring” the real $V_{A,j+1}$, sampling a new $V'_{A,j+1}$ and continuing with the computation as though $V'_{A,j+1}$ were the real view. Intuitively, T'_{j+1} is close to independent of $V_{A,j+1}$ (conditioned on T_j) and so knowledge of T'_{j+1} does not give additional information on $V_{A,j+1}$ beyond what is already given by T_j .

Properties (1) and (2) are used to argue that if with high probability there is a first point where $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]] - \mathbf{E}[C_{j+1} | T_{j+1}]$ is large (as occurs in Case (2b)) then with high probability there is a first point where $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T'_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]] - \mathbf{E}[C_{j+1} | T'_{j+1}]$ is large (see Claim 1 for the precise statement).

Property (3) is used to argue that $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T'_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]]$ is close to $\mathbf{E}[C | T'_{j+1}]$ (see Claim 2 for the precise statement). To give some intuition into why this holds, note that $\mathbf{E}[C | T'_{j+1}]$ can be re-written as: $\mathbf{E}_{V_{A,j+1}, V_{B,j} | T'_{j+1}}[\mathbf{E}[C | V_{A,j+1}, V_{B,j}]]$. Now, the quantity $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T'_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]]$ is nearly the same, except views of B, $V_{B,j}$ are sampled conditioned on T'_{j+1} , but views of A, $V_{A,j+1}$, are sampled conditioned only on $(V_{B,j}, \text{Eve}_j)$ (which in particular includes T_j). Intuitively, this reflects the fact that T'_{j+1} does not provide additional information about the real view of A, $V_{A,j+1}$ over what is contained in T_j . However, the fact that T'_{j+1} does not leak additional information on $V_{A,j+1}$, is not sufficient to argue that $\mathbf{E}_{V_{B,j}, \text{Eve}_j | T'_{j+1}}[\mathbf{E}[C | V_{B,j}, \text{Eve}_j]]$ is close to $\mathbf{E}[C | T'_{j+1}]$. This is because T'_{j+1} still contains additional Eve queries which, although they do not provide additional information about $V_{A,j+1}$, do provide additional overall information about the oracle. Thus, in order for item (3) to hold, we need the additional “function-obliviousness” property (see Property 1) which guarantees that the outcome C of the coin-toss does not depend on the oracle, but only on the random tapes of the two parties. We note that this is the only place in the proof where the “function-obliviousness” property

is used. Thus, as long as we can sample partial views of A and B , $V_{A,j+1}, V_{B,j}$ according to the correct distribution, we can compute the expected value of the coin toss $C(V_{A,j+1}, V_{B,j}) = C(r_A, r_B)$.

When the above are combined, we get that with high probability there is a first point where $\mathbf{E}[C \mid T'_{j+1}] - \mathbf{E}[C_{j+1} \mid T'_{j+1}]$ is large, which means that an adversary can impose bias by adopting a [CI93]-type strategy.

Unfortunately, the actual argument is somewhat more complicated than what is described above, because once the adversarial party A' playing the role of A has computed the simulated transcript T'_{j+1} and the associated information (which we denote by K_{j+1}), A' cannot just throw away the additional information in K_{j+1} and start afresh when computing expectations in the $j + 3$ -rd pass. This is because just *the fact that A' has not aborted* itself gives information that might impact the expected value of the coin toss. Thus, A' cannot decide to abort by conditioning only on T'_{j+3} , but must additionally condition on its extra knowledge K_{j+1} , which it obtained in the previous round, when deciding whether or not to abort. Therefore, all the information in the K_j variables must be used when A' computes subsequent expectations. Moreover, when $V'_{A,j+1}$ is sampled (in order to compute T'_{j+1}), it must be consistent not only with T_j but also with the additional knowledge K_j collected thus far. See Section 4.1 for the precise description and analysis of the final adversarial strategy.

2 Preliminaries

Definition 2 (Black-Box Coin Tossing from One-Way Functions). *For (interactive) oracle algorithms A, B we call $C = \langle A, B \rangle$ a black-box construction of coin tossing with bias at most δ based on one-way functions with input/output length n , if the following properties hold:*

- *The parties A and B get access to private randomness r_A, r_B and common input 1^n and run in time $\text{poly}(n)$ and interact for $\hat{r}(n) = \text{poly}(n)$ number of rounds. The transcript of their interaction determines an output a . Also, if during the protocol, A (resp. B) receives the special message $\perp$ (denoting that the other party has stopped playing in the protocol) then A (resp. B) outputs a bit a (resp b) on their own which will be the output of the protocol.*
- **Completeness:** *For any function f , if A and B are given oracle access to f and execute the protocol honestly, then the output is an unbiased random bit.*
- **Security:** *There is an oracle algorithm S running in polynomial time over its input length with the following property. Given any adversary $\mathcal{A}$ (playing on behalf of A or B) that achieves bias $\delta(n)$ over common input 1^n w.r.t a function f , $S^{f, \mathcal{A}}(1^n, 1^{1/\delta(n)})$ breaks the security of f as a one-way function.*

2.1 The Eavesdropper Algorithm Eve

In this section, we recall the Eve algorithm, first introduced by Impagliazzo and Rudich [IR89] in the context of separating one-way function and key agreement. The Eve algorithm of [IR89] was later improved by Barak and Mahmoody [BM13]. In our

work, we will use the Eve algorithm of [BM13] in a black-box manner. Thus, we do not describe the algorithm itself, and simply state the properties we will need from the Eve algorithm of [BM13] in the following lemma.

Lemma 1 (Implied by Theorem 4.2 in [BM13]). *Let $\mathcal{C} = \langle A, B \rangle$ be an oracle protocol in which the parties A, B ask at most m queries each from the oracle $\mathcal{O}$. Then there is an Eve algorithm who only gets to see the public messages and asks her own oracle queries after each message is sent and on input parameter $\epsilon < 1/100$:*

- **poly(m/ϵ)-Efficiency:** *Eve is deterministic and, over the randomness of the oracle and A and B 's private randomness, the expected number of Eve queries from the oracle $\mathcal{O}$ is at most $(10m/\epsilon)^{10}$.*
- **$(1 - \epsilon)$ -Security:** *Let $T_i = M_1, \dots, M_i || \text{Eve}_i$ be the transcript of messages sent between A and B so far, including the additional information that Eve has learned till the end of the i 'th pass. Let $(V_A, V_B) \mid T_i$ be the joint distribution over the views (V_A, V_B) of A and B only conditioned on T_i . By $V_A \mid T_i$ and $V_B \mid T_i$ we refer to the projections of $\mathcal{D}(T_i)$ over its first or second components. Then, with probability at least $1 - \epsilon$ over the randomness of A, B , and the random oracle $\mathcal{O}$, the following holds at all moments during the protocol when Eve is done with her learning phase in that round:*
 1. *The statistical distance between $V_A \mid T_i \times V_B \mid T_i$ and $\mathcal{D}(T_i)$ is at most ϵ . Namely: $\Delta(V_A \mid T_i \times V_B \mid T_i, (V_A, V_B) \mid T_i) \leq \epsilon$.*
 2. *For every oracle query $q \notin \text{Eve}_i$ it holds that $\Pr_{(V_A, V_B) \mid T_i}[q \in Q_{V_A} \cup Q_{V_B}] \leq \epsilon$.*

In the following, we will run the Eve algorithm with input parameter $\epsilon = \frac{1}{3m\hat{r}^4}$.

For simplicity of the notation and when it is clear from the context, in the following, for probabilities and expected values taken over $(V_A, V_B) \sim \mathcal{D}$, instead of writing $\mathbf{E}_{(V_A, V_B) \sim \mathcal{D}}$ or $\Pr_{(V_A, V_B) \sim \mathcal{D}}$, we simply write $\mathbf{E}$ and $\Pr$.

We consider coin-tossing protocols $\mathcal{C}$, where the Eve algorithm is run alongside the protocol and Eve queries are made immediately after every message M_j is sent. We denote by Eve_j the set of queries made by the Eve algorithm up to and including the queries made immediately after the j -th message is sent. We denote by T_j , the transcript of the protocol with the Eve queries made alongside. Thus, $T_j = M_1, \dots, M_j || \text{Eve}_j$.

3 Types of Coin Tossing Protocols We Consider

Consider a coin-tossing protocol $\mathcal{C} = \langle A, B \rangle$ with $\hat{r} = \hat{r}(k) = \text{poly}(k)$ rounds and $2\hat{r}$ passes. For $1 \leq w \leq \hat{r}$, let C_{2w-1} denote the output of party B in the case that A aborts before sending the $2w - 1$ -st message. Similarly, For $1 \leq w \leq \hat{r}$, let C_{2w} denote the output of party A in the case that B aborts before sending the $2w$ -th message.

Let $V_{A,j}$ (resp. $V_{B,j}$) denote the partial view of A (resp. B) up to and including pass j . In particular, $V_{A,j}$ consists of the transcript M_j thus far as well as the random tape r_A of A and the queries and responses, $Q_{V_{A,j}}$, that have been made by A thus far. $V_{B,j}$ and $Q_{V_{B,j}}$ are defined analogously.

We consider the distribution $\mathcal{D}$ to be the distribution over pairs of complete views $(V_{A,2\hat{r}}, V_{B,2\hat{r}})$ (also denoted simply by V_A, V_B) generated by a run of $\mathcal{C}$ with a random oracle. More specifically, a draw from $\mathcal{D}$ is obtained as follows:

- Draw $O \sim \mathcal{Y}$, $r_A, r_B \leftarrow \{0, 1\}^{p(n)}$, for some polynomial $p(\cdot)$ and execute $\mathcal{C} = \langle A, B \rangle$ with O, r_A, r_B .
- Output the views (V_A, V_B) resulting from the execution of $\mathcal{C} = \langle A, B \rangle$ above.

We prove our result for so-called *instant* constructions as defined in [DSLMM11]. Instant constructions are coin-tossing protocols where for $1 \leq w \leq \hat{r}$, A (resp. B) computes the value C_{2w} (resp. C_{2w+1}) before sending message M_{2w-1} (resp. M_{2w}). Thus, in case a party (say A) aborts before sending message M_{2w+1} , then B can simply output its precomputed value C_{2w+1} which depends only on B 's view at the point right after B computed message M_{2w} without making any additional oracle queries. It is not hard to see that the restriction of instant constructions can be removed as was shown by [DSLMM11]. This is a subtle argument relying on the fact that our ultimate goal is to rule out separations from one-way functions and not random oracles (since in the random oracle model coin tossing is trivial). In the following we sketch the argument of [DSLMM11] on why assuming the protocol to be instant is w.l.o.g.

Instant vs. General Protocols. Dealing with non-instant protocols can be done exactly as it was done in [DSLMM11], so in this work we focus on instant protocols and leave the full discussions on dealing with non-instant protocols for the full version of the paper. However, here we give a sketch of how this can be done. Firstly, note that any general coin tossing protocol using an oracle can be made “almost instant” without losing the security as follows. Whenever a party A (or B) wants to send a message M_i , they also go ahead and ask any oracle query that they would need to ask in case the other party halts the execution of the protocol and not sent M_{i+1} . This way, the protocol becomes almost instant because the only time that the instant property might be violated is when the first message is aborted by Alice in which case, Bob might still need to query the oracle to decide the output. However, as shown in [DSLMM11], it is always possible to “fix” a “small” set S of queries of the random oracle in a way that (1) Bob does not ask any query to decide the output if he gets aborted in the first message, and (2) the protocol remains as secure. Roughly speaking, the set S is determined (and its answers are fixed) as follows. The set S contains any query q that has a “non-negligible” chance of being asked by Bob in case of not receiving the first message. It is easy to show that $|S| \leq \text{poly}(n)$, and by sampling (and fixing) the answer of the queries in S , Bob will not need to ask any oracle queries in case of getting aborted in the first round. Finally, observe that a partially-fixed random oracle is still one-way and so one can apply the argument of our work for the instant protocols to the final instant protocol.

We consider coin-tossing protocols that are so-called “function oblivious.” As defined in Definition 1, these are coin-tossing protocols such that the outcome of protocol $\mathcal{C}^f = \langle A^f, B^f \rangle$ when both parties are honest depends only on the random tapes r_A, r_B of the two parties and not on the choice of one-way function f . We denote by $C(r_A, r_B)$ the output of protocol $\mathcal{C}$ when run with random tapes r_A, r_B . When the settings of r_A, r_B are clear from context, we denote the output of the protocol by C .

We are now ready to state our main theorem.

Theorem 2 (Main Theorem, Formal). *There is no (fully) black-box construction of an $\hat{r} = \hat{r}(n)$ -round, function-oblivious coin-tossing protocol $\mathcal{C}^f = \langle A^f, B^f \rangle$ with bias $o(1/\hat{r}^{3/4})$ from one-way functions.*

4 Proof of the Main Theorem

Towards proving Theorem 2, we begin with the following fact, which follows straightforwardly from [CI93].

Fact 1. *Let C be a coin-tossing protocol and let $\{Y_1, \dots, Y_{2\hat{r}}\}$ be a set of random variables, where Y_j is associated with some state of protocol C immediately after the j -th message (M_j, Eve_j) has been computed.*

- For $1 \leq w \leq \hat{r}$, set $j = 2w - 2$ and define the indicator variable $I_{\text{Val}_{j+1}^A}$ in the following way: $I_{\text{Val}_{j+1}^A} = 1$ if $|\mathbf{E}[C|Y_{j+1}] - \mathbf{E}[C_{j+1}|Y_{j+1}]| \geq \beta$ and for $1 \leq \ell \leq w$, $I_{\text{Val}_{2\ell-1}^A} = 0$. Otherwise $I_{\text{Val}_{j+1}^A} = 0$.
- For $1 \leq w \leq \hat{r}$, set $j = 2w$ and define the indicator variable $I_{\text{Val}_{j+1}^A}$ in the following way: $I_{\text{Val}_j^B} = 1$ if $|\mathbf{E}[C|Y_j] - \mathbf{E}[C_{j+2}|Y_j]| \geq \beta$ and for $1 \leq \ell \leq w$, $I_{\text{Val}_{2\ell}^B} = 0$. Otherwise $I_{\text{Val}_j^B} = 0$.

If for some (α, β)

$$\sum_{w=1}^{\hat{r}} \Pr[I_{\text{Val}_{2w-1}^A} = 1] \geq \alpha$$

then player A has a fail-stop strategy for imposing bias $\pm 1/2 \cdot \alpha \cdot \beta$ on C by aborting before sending message M_{2j-1} either when $\mathbf{E}[C|Y_{2w-1}] - \mathbf{E}[C_{2w-1}|Y_{2w-1}] \geq \beta$ or when $\mathbf{E}[C|Y_{2w-1}] - \mathbf{E}[C_{2w-1}|Y_{2w-1}] \leq -\beta$. An analogous claim holds for player B.

If for some (α, β) ,

$$\sum_{w=1}^{\hat{r}} \Pr[I_{\text{Val}_{2w}^B} = 1] \geq \alpha$$

then player B has a fail-stop strategy for imposing bias $\pm 1/2 \cdot \alpha \cdot \beta$ on C by aborting after sending message M_{2j} either when $\mathbf{E}[C|Y_{2w}] - \mathbf{E}[C_{2w+2}|Y_{2w}] \geq \beta$ or when $\mathbf{E}[C|Y_{2w}] - \mathbf{E}[C_{2w+2}|Y_{2w}] \leq -\beta$. An analogous claim holds for player A.

The following fact is implicit in [CI93]:

Fact 2. *With prob. at least $1/5$ over choice of random tapes and oracle there is some point in the execution such that $|\mathbf{E}[C|T_j] - \mathbf{E}[C|T_{j+1}]| \geq \Omega(1/\sqrt{r})$.*

Let us choose the quantity

$$MV = \mathbf{E}_{V_{B,j}, \text{Eve}_j | T_{j+1}}[\mathbf{E}[C|V_{B,j}, \text{Eve}_j]]$$

as the "middle value."

Thus, it must be the case that either with prob. at least $1/10$ there is some point in the execution such that $|\mathbf{E}[C|T_j] - MV| \geq \Omega(1/\sqrt{r})$. OR with prob. at least $1/10$ there is some point in the execution such that $|MV - \mathbf{E}[C|T_{j+1}]| \geq \Omega(1/\sqrt{r})$.

In particular, there are four possible cases.

1. With probability $1/20$ there is some point s.t. $\mathbf{E}[C|T_j] - MV \geq \Omega(1/\sqrt{r})$.
2. With probability $1/20$ there is some point s.t. $MV - \mathbf{E}[C|T_{j+1}] \geq \Omega(1/\sqrt{r})$.
3. With probability $1/20$ there is some point s.t. $MV - \mathbf{E}[C|T_j] \geq \Omega(1/\sqrt{r})$.
4. With probability $1/20$ there is some point s.t. $\mathbf{E}[C|T_{j+1}] - MV \geq \Omega(1/\sqrt{r})$.

Note that Cases 1 and 2 will give rise to adversarial strategies for biasing towards 0, while Cases 3 and 4 will give rise to adversarial strategies for biasing towards 1. In the following, we analyze only cases 1 and 2; cases 3 and 4 are entirely analogous.

Lemma 2. *Assume Case 1 occurs with probability at least $1/20$, then there is a strategy that biases the output by $\Omega(1/\sqrt{r})$.*

Proof. Assume that Case (1) occurs. Then this means that with prob. at least $1/20$ there is some point in the execution such that

$$\mathbf{E}[C | T_j] - MV = \mathbf{E}[C | T_j] - \mathbf{E}_{V_{B,j}, Eve_j | T_{j+1}}[\mathbf{E}[C | V_{B,j}, Eve_j]] \geq \Omega(1/\sqrt{r}).$$

Fix T_{j+1} such that Case (1) occurs. Note that T_{j+1} completely defines T_j and so the quantity above can be calculated for every valid T_{j+1} . Now, for each such T_{j+1} we must have that one of the following two subcases occurs:

- (1a) $\Pr_{V_{B,j}, Eve_j | T_{j+1}}[\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C | T_j] \geq \Omega(1/\hat{r}^{1/4})] \geq \Omega(1/\sqrt{\hat{r}})$ OR
- (1b) $\Pr_{V_{B,j}, Eve_j | T_{j+1}}[\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C | T_j] \geq \Omega(1/\sqrt{r})] \geq \Omega(1/\hat{r}^{1/4})$.

To see this, assume towards contradiction that neither item (1a) nor item (1b) occur. Then this means that when $V_{B,j}$ is sampled conditioned on T_{j+1} , we have that the contribution from $V_{B,j}$ such that $\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C | T_j] \geq \Omega(1/\sqrt{r})$ and $\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C | T_j] \leq \Omega(1/\hat{r}^{1/4})$ is at most $o(1/\hat{r}^{1/4}) \cdot 1/\hat{r}^{1/4} = o(1/\sqrt{\hat{r}})$. Additionally, the contribution from $V_{B,j}$ such that $\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C | T_j] \geq \Omega(1/\hat{r}^{1/4})$ is at most $o(1/\sqrt{\hat{r}}) \cdot 1 = o(1/\sqrt{\hat{r}})$. This is a contradiction to Case 1 occurring.

Now, if item (1a) occurs then this means that either with probability $\Omega(1/\hat{r}^{1/4})$ we have that $\mathbf{E}[C | T_j] - \mathbf{E}[C_{j+2} | T_j] \geq \Omega(1/\sqrt{\hat{r}})$ occurs OR that with probability $\Omega(1/\hat{r}^{1/4})$ we have that $\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C | T_j] \geq \Omega(1/\sqrt{\hat{r}})$ AND $\mathbf{E}[C | T_j] - \mathbf{E}[C_{j+2} | T_j] \leq o(1/\sqrt{\hat{r}})$.

In the first case, B can employ the following strategy:

Abort immediately **after** sending message M_j if:

$$\mathbf{E}[C|T_j] - \mathbf{E}[C_{j+2}|T_j] \geq \Omega(1/\sqrt{\hat{r}})$$

By Fact 1 the strategy above imposes bias of at least $\Omega(1/\hat{r}^{3/4})$ towards 0.

In the second case, note that since C_{j+2} is a function of only $V_{A,j+1}$, we have by the properties of the Eve algorithm given in Lemma 1 we have that with probability $1 - O(1/\hat{r}^2)$, we have that $|\mathbf{E}[C_{j+2} | T_j] - \mathbf{E}[C_{j+2} | V_{B,j}, Eve_j]| \leq O(1/\hat{r}^2)$. Thus, in this case, we have that with probability $\Omega(1/\hat{r}^{1/4})$, $\mathbf{E}[C | V_{B,j}, Eve_j] - \mathbf{E}[C_{j+2} | V_{B,j}, Eve_j] \geq \Omega(1/\sqrt{\hat{r}})$ and in this case, B can employ the following strategy:

Abort immediately **after** sending message M_j if:

$$\mathbf{E}[C \mid V_{B,j}, \text{Eve}_j] - \mathbf{E}[C_{j+2} \mid V_{B,j}, \text{Eve}_j] \geq \Omega(1/\sqrt{\hat{r}})$$

Thus by Fact 1 the above strategy imposes bias of at least $\Omega(1/\hat{r}^{3/4})$ towards 0.
The analysis for item (1b) is entirely analogous.

Lemma 3. *Assume Case 2 occurs with probability at least $1/20$, then there is a strategy that biases the output by $\Omega(1/\hat{r}^{3/4})$.*

Proof. Case (2) implies that one of the following occurs with prob. at least $1/40$:

(2a)

$$\mathbf{E}[C_{j+1} \mid T_{j+1}] - \mathbf{E}[C \mid T_{j+1}] \geq \Omega(1/\sqrt{\hat{r}})$$

(2b)

$$MV - \mathbf{E}[C_{j+1} \mid T_{j+1}] = \mathbf{E}_{V_{B,j}, \text{Eve}_j | T_{j+1}}[\mathbf{E}[C \mid V_{B,j}, \text{Eve}_j]] - \mathbf{E}[C_{j+1} \mid T_{j+1}] \geq \Omega(1/\sqrt{\hat{r}})$$

Note that in Case (2a), A can employ the following strategy:

Abort **before** sending message M_{j+1} if:

$$\mathbf{E}[C \mid T_{j+1}] - \mathbf{E}[C_{j+1} \mid T_{j+1}] \geq \Omega(1/\sqrt{\hat{r}})$$

and thus, by Fact 1 imposes bias $\Omega(1/\sqrt{\hat{r}})$ towards 0. Thus, to complete the lemma, we need to show that if Case (2b) occurs with probability at least $1/40$ then there is a strategy for imposing bias of $\Omega(1/\hat{r}^{3/4})$ on the outcome.

Since this case becomes more complicated, we devote the following section to show how to deal with Case (2b).

4.1 Analysis for Case (2b)

The Protocol C' . The modified protocol C' will execute the regular C protocol with Eve queries made alongside. B behaves as in the original protocol. A' behaves as A in the original protocol and additionally computes extra state information K_i and related values in each round i .

For each pass $0 \leq j \leq 2\hat{r} - 1$, we consider the distribution $\mathcal{D}_{\text{extend}, j+1}$, which is a distribution over a tuple consisting of partial views $V_{A, j+1}, V_{B, j+1}$, transcripts (with Eve queries alongside) T_{j+1} , and additional knowledge K_{j+1} generated by a random execution of C' with random oracle $\mathcal{O}$.

More specifically, a draw from $\mathcal{D}_{\text{extend}, j+1}$ is obtained as follows:

- Draw $O \sim \mathcal{I}$, $r_{A'}, r_B \leftarrow \{0, 1\}^{p'(n)}$, for some polynomial $p'(\cdot)$ and execute $C' = \langle A', B \rangle$ with $O, r_{A'}, r_B$.
- Output a tuple consisting of the views $V_{A, j+1}, V_{B, j+1}$, transcript T_{j+1} , and additional state information K_{j+1} resulting from the execution of $C' = \langle A', B \rangle$ above.

We are now ready to describe how K_{j+1} is computed: For $j = 0$, the variable K_0 is set to empty. For each round $1 \leq w \leq \hat{r}$, set $j = 2(w - 1)$. A' computes the state information K_{j+1} in the following way:

- Sample a random partial view $V'_{A,j+1}$ from $\mathcal{D}_{extend,j+1}(T_j, K_j)$. Recall that this denotes the distribution $\mathcal{D}_{extend}$, conditioned on the current transcript with Eve queries, T_j , and the additional state information K_j . Note that $V'_{A,j+1}$ includes the next message, M'_{j+1} .
- We run a modified version of the Eve algorithm, called the Eve' algorithm. For each pass ℓ , let $Q_{Eve'_\ell}$ denote the set of queries and responses made by Eve' in the j -th pass. In the $j + 1$ -st pass do the following: Run the Eve algorithm at pass $j + 1$ conditioned on $T_j || M'_{j+1}$ (i.e. as if M'_{j+1} is the real next message). Answer oracle queries made by Eve' in the following way³:
 - If the query q appears in $V'_{A,j+1}$, answer according to $V'_{A,j+1}$ (without querying the oracle).
 - Otherwise, if for some $i \leq j$, the query q appears in $Q_{Eve'_i} \setminus Q_{V'_{A,i}}$, respond according to the value listed in $Q_{Eve'_i}$ (without querying the oracle).
 - Otherwise, if a query q appears in $V_{A,j}$, sample and return a uniformly random string (without querying the oracle).
 - Otherwise, query the oracle and return the oracle's response.
- We denote by T'_{j+1} the *fake transcript* generated. More specifically, $T'_{j+1} = T_j || M'_{j+1} || Q_{Eve'_{j+1}}$.
- Set K_{j+1} and K_{j+2} to be K_j with the variables $V'_{A,j+1}, Q_{Eve'_{j+1}}$ appended.

Intuitively, the point of the protocol C' is that it allows a malicious A to sample fake transcripts T'_{j+1} , which, conditioned on T_j, K_j , are distributed (almost) identically to real transcripts T_{j+1} , but reveal (almost) no additional information about the real $V_{A,j+1}$, beyond what was revealed by T_j, K_j . In particular, a “fake” view $V'_{A,j+1}$, independent of the real $V_{A,j+1}$, is sampled and a fake next message M'_{j+1} is computed. Now, when we run the Eve' algorithm, ideally we would like to answer all oracle queries q appearing in $Q_{V'_{A,j+1}}$ dishonestly according to $V'_{A,j+1}$ and all other queries honestly according to the real oracle. However, there is a subtle issue here: Queries in the real $Q_{V_{A,j+1}}$ may be “incorrectly” distributed if they are answered according to the real oracle. In particular, queries which appear in $Q_{Eve'_i} \setminus Q_{V'_{A,i}}$ and do not appear in $V'_{A,j+1}$, must be answered according to the value listed there (regardless of whether they are in $V_{A,j+1}$). Queries which do not appear in $Q_{V'_{A,j+1}}$ and do not appear in $Q_{Eve'_i} \setminus Q_{V'_{A,i}}$, but do appear in $Q_{V_{A,j+1}}$ are answered uniformly at random.

We are now ready to describe the final adversarial strategy:

- Set $f(\hat{r}) = 1/\sqrt{\hat{r}}$ or $f(\hat{r}) = 1/\hat{r}^{1/4}$.
- Play the role of A' in an execution C' , while interacting with an honest B .
- Abort immediately **before** sending message M_{j+1} if:

$$\mathbf{E}[C|T'_{j+1}, K_j] - \mathbf{E}[C_{j+1}|T'_{j+1}, K_j] = \Omega(f(\hat{r}))$$

³ We assume that Eve' never re-queries a query that is already contained in Eve_j

Fact 1 implies that all we need to show is that the event occurs “frequently.” More specifically, we prove the following lemma, which is sufficient for completing the proof of Case (2b).

Lemma 4. *If Case (2b) occurs with probability $1/40$ then either:*

- *With probability $\Omega(1/\hat{r}^{1/2})$ over executions of C' and choice of oracle $\mathcal{O}$ there is a first message j where*

$$\mathbf{E}[C|T'_{j+1}, K_j] - \mathbf{E}[C_{j+1}|T'_{j+1}, K_j] = \Omega(1/\hat{r}^{1/4})$$

- *With probability $\Omega(1/\hat{r}^{1/4})$ over executions of C' and choice of oracle $\mathcal{O}$ there is a first message j where*

$$\mathbf{E}[C|T'_{j+1}, K_j] - \mathbf{E}[C_{j+1}|T'_{j+1}, K_j] = \Omega(1/\hat{r}^{1/2})$$

Before proving Lemma 4, we introduce the following notation. For $f(\hat{r}) = 1/\sqrt{\hat{r}}$ or $f(\hat{r}) = 1/\hat{r}^{1/4}$ and for every $j = 0, 1 \leq j \leq 2\hat{r}$, we define the indicator random variables $I_{\text{EV}_j}^f$ and $I_{\text{EV}'_j}^f$ which are set before the j -th message is sent during an execution of C' . For $j = 0$, $I_{\text{EV}'_0}^f = 0$ and $I_{\text{EV}_0}^f = 0$. For $j \geq 1$, $I_{\text{EV}'_{j+1}}^f, I_{\text{EV}'_{j+2}}^f$ are set to 1 if:

- $I_{\text{EV}'_j}^f = 1$ OR
- $\mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}'_{j+1}=t'_{j+1}, K_j=k_j}[\mathbf{E}[C|\mathbf{V}_{B,j}, K_j = k_j]] - \mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}'_{j+1}=t'_{j+1}, K_j=k_j}[C_{j+1}(\mathbf{V}_{B,j})] = \Omega(f(\hat{r}))$.

For $j \geq 1$, $I_{\text{EV}_{j+1}}^f, I_{\text{EV}_{j+2}}^f$ are set to 1 if:

- $I_{\text{EV}_j}^f = 1$ OR
- $\mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}'_{j+1}=t_{j+1}, K_j=k_j}[\mathbf{E}[C|\mathbf{V}_{B,j}, K_j = k_j]] - \mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}'_{j+1}=t_{j+1}, K_j=k_j}[C_{j+1}(\mathbf{V}_{B,j})] = \Omega(f(\hat{r}))$.

Note that in the last expression, we condition on $\mathbf{T}'_{j+1} = t_{j+1}$. This means that $\mathbf{T}_{j+1} = t_{j+1}$ is sampled via a run of the protocol C' . Then, the expectation above is computed using this same value of t_{j+1} , but conditioning on the variable $\mathbf{T}'_{j+1}$ being equal to this value.

Let the event B_j be the event that upon a draw from $\mathcal{D}_{\text{extend},j+1}$ there is a query q such that $q \in Q_{\mathbf{V}_{B,j}} \cap (Q_{\mathbf{V}'_{A,j+1}} \cup Q_{\mathbf{V}_{A,j+1}})$ and $q \notin \text{Eve}_j$. Note that by Lemma 1, for each j , the probability that B_j occurs is at most $\frac{3m}{3m\hat{r}^4} = O(1/\hat{r}^4)$. Let $\mathcal{D}_{\text{extend},j+1}^{\text{Good}_j}$ denote the distribution $\mathcal{D}_{\text{extend},j+1}$, conditioned on $\bar{B}_j$.

By $\mathcal{D}_{\text{extend}}$ we denote the distribution $\mathcal{D}_{\text{extend},2\hat{r}}$. Additionally, for joint random variables X, Y , we denote by $X | Y$ the distribution over X drawn from $\mathcal{D}_{\text{extend}}$, conditioned on Y .

We first consider three important properties of the C' protocol which will help us prove the lemma:

Property 1 ($\mathcal{T}_{j+1}$ and $\mathcal{T}'_{j+1}$ are close). The two distributions

$$\mathcal{D}_{\mathcal{T}'_{j+1}}^{\text{Good}_j}(\mathcal{T}_j = t_j, \mathcal{K}_j = k_j) \quad \mathcal{D}_{\mathcal{T}_{j+1}}^{\text{Good}_j}(\mathcal{T}_j = t_j, \mathcal{K}_j = k_j)$$

are identical.

Since $\mathcal{B}_j$ occurs with probability at most $O(1/\hat{r}^4)$, Property 1 immediately implies the following: With probability $1 - O(1/\hat{r}^2)$ over $\mathcal{T}_j = t_j, \mathcal{K}_j = k_j$ drawn from $\mathcal{D}_{\text{extend}}$, we have that

$$\mathcal{D}_{\mathcal{T}'_{j+1}}(\mathcal{T}_j = t_j, \mathcal{K}_j = k_j) \quad \mathcal{D}_{\mathcal{T}_{j+1}}(\mathcal{T}_j = t_j, \mathcal{K}_j = k_j)$$

are $O(1/\hat{r}^2)$ -close.

Property 2 ($\mathcal{V}_{\mathcal{B}}$ conditioned on $\mathcal{T}_{j+1}$ or $\mathcal{T}'_{j+1}$ are close). For every $1 \leq j \leq 2\hat{r}$, the two distributions

$$\mathcal{D}_{\mathcal{V}_{\mathcal{B},j}}^{\text{Good}_j}(\mathcal{T}'_{j+1} = t_{j+1}, \mathcal{K}_j = k_j) \quad \mathcal{D}_{\mathcal{V}_{\mathcal{B},j}}^{\text{Good}_j}(\mathcal{T}_{j+1} = t_{j+1}, \mathcal{K}_j = k_j)$$

are identical.

Since $\mathcal{B}_j$ occurs with probability at most $O(1/\hat{r}^4)$, Property 2 immediately implies the following: For every $1 \leq j \leq 2\hat{r}$, we have that with probability $1 - O(1/\hat{r}^2)$ over draws of $\mathcal{T}_{j+1} = t_{j+1}$ and $\mathcal{K}_j = k_j$ from $\mathcal{D}_{\text{extend}}$, the statistical distance between the following:

$$\mathcal{D}_{\mathcal{V}_{\mathcal{B},j}}(\mathcal{T}'_{j+1} = t_{j+1}, \mathcal{K}_j = k_j) \quad \mathcal{D}_{\mathcal{V}_{\mathcal{B},j}}(\mathcal{T}_{j+1} = t_{j+1}, \mathcal{K}_j = k_j)$$

is at most $O(1/\hat{r}^2)$.

Property 3 ($\mathcal{T}'_{j+1}$ does not reveal much information about $\mathcal{V}_{\mathcal{A},j}$). With probability $1 - O(1/\hat{r}^2)$ over $\mathcal{T}_j = t_j, \mathcal{T}'_{j+1} = t_j | m'_{j+1}, \text{eve}'_{j+1}, \mathcal{K}_j = k_j$ drawn from $\mathcal{D}_{\text{extend}}$, we have that

$$\mathcal{D}_{\mathcal{V}_{\mathcal{A},j+1}}(\mathcal{T}'_{j+1} = t'_{j+1}, \mathcal{K}_j = k_j) \quad \mathcal{D}_{\mathcal{V}_{\mathcal{A},j+1}}(\mathcal{T}_j = t_j, \mathcal{K}_j = k_j)$$

are $O(1/\hat{r}^2)$ -close.

We defer the proofs of Properties 1, 2, 3 to the full version and now complete the proof of Lemma 4 via the following claims and facts:

Claim 1. *If Case (2b) occurs with probability $1/40$ then either:*

- With probability $\Omega(1/\hat{r}^{1/2})$ there is a first point where

$$\mathbf{E}_{\mathcal{V}_{\mathcal{B},j}, \text{Eve}_j | \mathcal{T}'_{j+1}, \mathcal{K}_j} [E[\mathcal{C} | \mathcal{V}_{\mathcal{B},j}, \text{Eve}_j, \mathcal{K}_j]] - \mathbf{E}_{\mathcal{V}_{\mathcal{B},j} | \mathcal{T}'_{j+1}, \mathcal{K}_j} [\mathcal{C}_{j+1}(\mathcal{V}_{\mathcal{B},j})] = \Omega(1/\hat{r}^{1/4}).$$

- With probability $\Omega(1/\hat{r}^{1/4})$ there is a first point where

$$\mathbf{E}_{\mathcal{V}_{\mathcal{B},j}, \text{Eve}_j | \mathcal{T}'_{j+1}, \mathcal{K}_j} [E[\mathcal{C} | \mathcal{V}_{\mathcal{B},j}, \text{Eve}_j, \mathcal{K}_j]] - \mathbf{E}_{\mathcal{V}_{\mathcal{B},j} | \mathcal{T}'_{j+1}, \mathcal{K}_j} [\mathcal{C}_{j+1}(\mathcal{V}_{\mathcal{B},j})] = \Omega(1/\hat{r}^{1/2}).$$

Claim 2. *With probability $1 - O(1/\hat{r}^2)$, we have that*

$$|\mathbf{E}[\mathbf{C}|\mathbf{T}'_{j+1}, \mathbf{K}_j] - \mathbf{E}_{\mathbf{V}_{B,j}, \mathbf{Eve}_j|\mathbf{T}'_{j+1}, \mathbf{K}_j}[\mathbf{E}[\mathbf{C}|\mathbf{V}_{B,j}, \mathbf{Eve}_j, \mathbf{K}_j]]| = O(1/\hat{r}^2).$$

Fact 3. *We have the following equivalence:*

$$\mathbf{E}[\mathbf{C}_{j+1}|\mathbf{T}'_{j+1}, \mathbf{K}_j] = \mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}'_{j+1}, \mathbf{K}_j}[\mathbf{C}_{j+1}(\mathbf{V}_{B,j})]$$

The above immediately imply Lemma 4. We now proceed to prove Claims 1 and 2

Proof. (Claim 1) The hypothesis of Claim 1 and Markov's inequality imply that one of the following must occur:

- With probability $\Omega(1/\hat{r}^{1/2})$ there is a first point where

$$\mathbf{E}_{\mathbf{V}_{B,j}, \mathbf{T}_j|\mathbf{T}_{j+1}, \mathbf{K}_j}[\mathbf{E}[\mathbf{C}|\mathbf{V}_{B,j}, \mathbf{T}_j, \mathbf{K}_j]] - \mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}_{j+1}, \mathbf{K}_j}[\mathbf{C}_{j+1}(\mathbf{V}_{B,j})] \geq \Omega(1/\hat{r}^{1/4})$$

- With probability $\Omega(1/\hat{r}^{1/4})$ there is a first point where

$$\mathbf{E}_{\mathbf{V}_{B,j}, \mathbf{T}_j|\mathbf{T}_{j+1}, \mathbf{K}_j}[\mathbf{E}[\mathbf{C}|\mathbf{V}_{B,j}, \mathbf{T}_j, \mathbf{K}_j]] - \mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}_{j+1}, \mathbf{K}_j}[\mathbf{C}_{j+1}(\mathbf{V}_{B,j})] \geq \Omega(1/\hat{r}^{1/2})$$

Let us assume that the first case above occurs. The analysis for the remaining case is entirely analogous. Now, by Claim 2 we have that with probability $1 - O(1/\hat{r}^2)$, over $\mathbf{T}_{j+1} = t_{j+1}, \mathbf{K}_j = k_j$ drawn from $\mathcal{D}_{\text{extend}}$ the distributions $\mathbf{V}_{B,j} \mid \mathbf{T}_{j+1} = t_{j+1}, \mathbf{K}_j = k_j$ and $\mathbf{V}_{B,j} \mid \mathbf{T}'_{j+1} = t_{j+1}, \mathbf{K}_j = k_j$ are $O(1/\hat{r}^2)$ -close.

Thus, we have that with probability $\Omega(1/\hat{r}^{1/2})$ over $\mathbf{T}_{j+1} = t_{j+1}, \mathbf{K}_j = k_j$ drawn from $\mathcal{D}_{\text{extend}}$

$$\mathbf{E}_{\mathbf{V}_{B,j}, \mathbf{T}_j|\mathbf{T}'_{j+1}=t_{j+1}, \mathbf{K}_j}[\mathbf{E}[\mathbf{C}|\mathbf{V}_{B,j}, \mathbf{T}_j, \mathbf{K}_j]] - \mathbf{E}_{\mathbf{V}_{B,j}|\mathbf{T}'_{j+1}=t_{j+1}, \mathbf{K}_j}[\mathbf{C}_{j+1}(\mathbf{V}_{B,j})] = \Omega(1/\hat{r}^{1/4}). \quad (4.1)$$

By definition of $I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}}$ we have by (4.1) that $\Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \text{ for some } 1 \leq j \leq 2\hat{r}] = \Omega(1/\hat{r}^{1/2})$. In the following, we will use this fact to show that $\Pr[I_{\text{EV}'_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \text{ for some } 1 \leq j \leq 2\hat{r}] = \Omega(1/\hat{r}^{1/2})$ as well. This will immediately imply the Claim.

First, for $1 \leq w \leq \hat{r}$, where $j = 2w - 2$, we define

$$\begin{aligned} v_{2w-1} &= \Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \wedge I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \\ y_{2w-1} &= \Pr[I_{\text{EV}'_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \wedge I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0]. \end{aligned}$$

Now, we have by Claim 1 that for every $1 \leq w \leq \hat{r}, j = 2w - 2$, one of the following occurs:

- $\Pr[I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] = O(1/\hat{r}^2)$
- The distributions $\mathbf{T}_{j+1}, \mathbf{K}_j \mid I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0, I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0$ and $\mathbf{T}'_{j+1}, \mathbf{K}_j \mid I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0, I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0$ are at most $O(1/\hat{r}^2)$ -far.

We show that in both cases, we must have that

$$v_{2w-1} = y_{2w-1} \pm O(1/\hat{r}^2). \quad (4.2)$$

In the first case, we clearly must have that $0 \leq v_{2w-1}, y_{2w-1} \leq O(1/\hat{r}^2)$. In the second case, we bound the difference between v_{2w-1}, y_{2w-1} in the following way:

$$\begin{aligned} v_{2w-1} &= \Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \wedge I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \\ &= \Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \mid I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0, I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \cdot \Pr[I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \\ &= \left(\Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \mid I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0, I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \pm O(1/\hat{r}^2) \right) \\ &\quad \cdot \Pr[I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \\ &= \Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \wedge I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0, I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0] \pm O(1/\hat{r}^2) \\ &= y_{2w-1} \pm O(1/\hat{r}^2), \end{aligned}$$

where the third equality follows since $I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}}, I_{\text{EV}'_{j+1}}^{f=1/\hat{r}^{1/4}}$ are completely determined by $\mathsf{T}_{j+1}, \mathsf{K}_j$ and $\mathsf{T}'_{j+1}, \mathsf{K}_j$, respectively.

Now, using the definition of $I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}}$ and (4.2) above, we have that

$$\begin{aligned} \Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \text{ for some } j = 2w - 2, 1 \leq w \leq \hat{r}] &\leq \sum_{w=1}^{\hat{r}} v_{2w-1} + y_{2w-1} \\ &\leq O(1/\hat{r}) + 2 \sum_{w=1}^{\hat{r}} y_{2w-1}. \end{aligned}$$

Moreover, (4.1) implies that

$$\begin{aligned} \Omega(1/\hat{r}^{1/2}) &= \Pr[I_{\text{EV}_{j+1}}^{f=1/\hat{r}^{1/4}} = 1 \text{ for some } j = 2w - 2, 1 \leq w \leq \hat{r}] \\ &\leq O(1/\hat{r}) + 2 \sum_{w=1}^{\hat{r}} y_{2w-1}. \end{aligned}$$

Thus, it must be the case that

$$\begin{aligned} \sum_{w=1}^{\hat{r}} y_{2w-1} &= \sum_{w=1}^{\hat{r}} \Pr[I_{\text{EV}_{2w-1}}^{f=1/\hat{r}^{1/4}} = 1 \wedge I_{\text{EV}_{2w-2}}^{f=1/\hat{r}^{1/4}} = 0 \wedge I_{\text{EV}'_{2w-2}}^{f=1/\hat{r}^{1/4}} = 0] \\ &= \Omega(1/\hat{r}^{1/2}). \end{aligned}$$

Finally, by definition, this implies that with probability $\Omega(1/\hat{r}^{1/2})$ over $\mathcal{D}_{\text{extend}}$ we have some j such that $I_{\text{EV}_j}^{f=1/\hat{r}^{1/4}} = 0, I_{\text{EV}'_j}^{f=1/\hat{r}^{1/4}} = 0$ and

$$\mathbf{E}_{\mathsf{V}_{B,j}, \mathsf{Eve}_j \mid \mathsf{T}'_{j+1}, \mathsf{K}_j} [E[\mathsf{C}[\mathsf{V}_{B,j}, \mathsf{Eve}_j, \mathsf{K}_j]]] - \mathbf{E}_{\mathsf{V}_{B,j} \mid \mathsf{T}'_{j+1}, \mathsf{K}_j} [\mathsf{C}_{j+1}(\mathsf{V}_{B,j})] = \Omega(1/\hat{r}^{1/4}).$$

and so the claim is proved.

Proof. (Claim 2) Towards proving the claim, note that by Lemma 1 we have that with probability $1 - O(1/\hat{r}^4)$ for every j

$$V_{A,j}, V_{B,j} \mid T_j \quad V_{A,j} \mid T_j \times V_{B,j} \mid T_j$$

are $O(1/\hat{r}^4)$ -close.

Thus, by applying Markov's inequality, we have that with probability $1 - O(1/\hat{r}^2)$, for every $1 \leq j \leq 2\hat{r}$,

$$V_{A,j}, V_{B,j} \mid T'_{j+1}, K_j \quad V_{A,j} \mid T'_{j+1}, K_j \times V_{B,j} \mid T'_{j+1}, K_j$$

are $O(1/\hat{r}^2)$ -close.

Now, by applying Property 3 we have that for every $1 \leq j \leq 2\hat{r}$ with probability $1 - O(1/\hat{r}^2)$, over draws of $T_j = t_j, T'_{j+1} = t'_j \mid m'_{j+1}, eve'_{j+1}, K_j = k_j$,

$$V_{A,j} \mid T'_{j+1}, K_j \times V_{B,j} \mid T'_{j+1}, K_j \quad V_{A,j} \mid T_j, K_j \times V_{B,j} \mid T'_{j+1}, K_j$$

are $O(1/\hat{r}^2)$ -close.

By combining the above, we have that for every j , with probability $1 - O(1/\hat{r}^2)$, over draws of $T_j = t_j, T'_{j+1} = t'_j \mid m'_{j+1}, eve'_{j+1}, K_j = k_j$,

$$V_{A,j}, V_{B,j} \mid T'_{j+1}, K_j, \quad V_{A,j} \mid T_j, K_j \times V_{B,j} \mid T'_{j+1}, K_j \quad (4.3)$$

are $O(1/\hat{r}^2)$ -close.

Now, let us consider the expression $\mathbf{E}_{V_{B,j}, T_j \mid T'_{j+1}, K_j} [\mathbf{E}[C \mid V_{B,j}, T_j, K_j]]$ and the expression $\mathbf{E}[C \mid T'_{j+1}, K_j]$. If we expand notation, we have that:

$$\mathbf{E}_{V_{B,j}, T_j \mid T'_{j+1}, K_j} [\mathbf{E}[C \mid V_{B,j}, T_j, K_j]] = \mathbf{E}_{V_{B,j}, T_j \mid T'_{j+1}, K_j} [\mathbf{E}_{V_{A,2\hat{r}}, V_{B,2\hat{r}} \mid T_j, K_j, V_{B,j}} [C(V_{A,2\hat{r}}, V_{B,2\hat{r}})]].$$

and that

$$\mathbf{E}[C \mid T'_{j+1}, K_j] = \mathbf{E}_{V_{A,2\hat{r}}, V_{B,2\hat{r}} \mid T'_{j+1}, K_j} [C(\text{View}_{A,2\hat{r}}, V_{B,2\hat{r}})].$$

Due to the function-obliviousness property (see Property 1), we have that $C(V_A, V_B)$ depends only on the random tapes r_A, r_B of A, B , which are contained in the partial views $V_{A,j}, V_{B,j}$, and so

$$\mathbf{E}_{V_{B,j}, T_j \mid T'_{j+1}, K_j} [\mathbf{E}_{V_{A,j}, V_{B,j} \mid T_j, K_j, V_{B,j}} [C(r_A, r_B)]] = \mathbf{E}_{V_{B,j}, T_j \mid T'_{j+1}, K_j} [\mathbf{E}_{V_{A,j} \mid T_j, K_j, V_{B,j}} [C(r_A, r_B)]]$$

and that

$$\mathbf{E}[C \mid T'_{j+1}, K_j] = \mathbf{E}_{V_{A,j}, V_{B,j} \mid T'_{j+1}, K_j} [C(r_A, r_B)].$$

Next, we have by Lemma 1 and Markov's inequality, that with probability $1 - O(1/\hat{r}^2)$,

$$V_{A,j} \mid V_{B,j}, T_j, K_j \quad V_{A,j} \mid T_j, K_j$$

are $O(1/\hat{r}^2)$ -close.

Thus, we have that with probability $1 - O(1/\hat{r}^2)$:

$$\left| \mathbf{E}_{V_{B,j}, T_j | T'_{j+1}, K_j} [\mathbf{E}[C | V_{B,j}, T_j, K_j]] - \mathbf{E}_{V_{B,j}, T_j | T'_{j+1}, K_j} [\mathbf{E}_{V_{A,j} | T_j, K_j} [C(V_{A,j}, V_{B,j})]] \right| = O(1/\hat{r}^2).$$

Equivalently, we have that with probability $1 - O(1/\hat{r}^2)$ over draws of $T_j = t_j, T'_{j+1} = t'_j | m'_{j+1}, eve'_j, K_j = k_j$:

$$\left| \mathbf{E}_{V_{B,j}, T_j | T'_{j+1}, K_j} [\mathbf{E}[C | V_{B,j}, T_j, K_j]] - \mathbf{E}_{V_{A,j} | T_j, K_j \times V_{B,j} | T'_{j+1}, K_j} [C] \right| = O(1/\hat{r}^2).$$

Finally, by applying (4.3), and since $V_{B,j}, T_j$ and $V_{B,j}, Eve_j$ contain the same information, we have that with all but $O(1/\hat{r}^2)$ probability,

$$\left| \mathbf{E}_{V_{B,j}, Eve_j | T'_{j+1}, K_j} [\mathbf{E}[C | V_{B,j}, Eve_j, K_j]] - \mathbf{E}_{V_{A,j}, V_{B,j} | T'_{j+1}, K_j} [C(V_{A,j}, V_{B,j})] \right| = O(1/\hat{r}^2).$$

Equivalently, we have that with all but $O(1/\hat{r}^2)$ probability,

$$\left| \mathbf{E}[C | T'_{j+1}, K_j] - \mathbf{E}_{V_{B,j}, Eve_j | T'_{j+1}, K_j} [\mathbf{E}[C | V_{B,j}, Eve_j, K_j]] \right| = O(1/\hat{r}^2),$$

and so the claim is proved.

References

- [Blu82] Blum, M.: Coin flipping by telephone - a protocol for solving impossible problems. In: COMPCON, pp. 133–137 (1982)
- [BM07] Barak, B., Mahmoody, M.: Lower bounds on signatures from symmetric primitives. In: FOCS: IEEE Symposium on Foundations of Computer Science, FOCS (2007)
- [BM09] Barak, B., Mahmoody-Ghidary, M.: Merkle puzzles are optimal—an $o(n^2)$ -query attack on key exchange from a random oracle. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 374–390. Springer, Heidelberg (2009)
- [BM13] Barak, B., Mahmoody, M.: Merkle’s key agreement protocol is optimal - an $O(n^2)$ -query attack on any key exchange from random oracles (2013), <http://www.cs.cornell.edu/~mohammad/files/papers/MerkleFull.pdf>
- [CI93] Cleve, R., Impagliazzo, R.: Martingales, collective coin flipping and discrete control processes (1993) (unpublished)
- [Cle86] Cleve, R.: Limits on the security of coin flips when half the processors are faulty (extended abstract). In: STOC, pp. 364–369 (1986)
- [DSLMM11] Dachman-Soled, D., Lindell, Y., Mahmoody, M., Malkin, T.: On the black-box complexity of optimally-fair coin tossing. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 450–467. Springer, Heidelberg (2011)
- [GGM86] Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. J. ACM 33(4), 792–807 (1986)
- [GM84] Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. 28(2), 270–299 (1984)

- [GT00] Gennaro, R., Trevisan, L.: Lower bounds on the efficiency of generic cryptographic constructions. In: FOCS, pp. 305–313 (2000)
- [HILL99] Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM J. Comput.* 28(4), 1364–1396 (1999)
- [HOZ13] Haitner, I., Omri, E., Zarusim, H.: Limits on the usefulness of random oracles. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 437–456. Springer, Heidelberg (2013)
- [IL89] Impagliazzo, R., Luby, M.: One-way functions are essential for complexity based cryptography (extended abstract). In: FOCS, pp. 230–235 (1989)
- [IR89] Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: STOC, pp. 44–61 (1989)
- [LR88] Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.* 17(2), 373–386 (1988)
- [MMP13] Mahmoody, M., Maji, H.K., Prabhakaran, M.: Limits of random oracles in secure computation. To Appear in: Innovations in Theoretical Computer Science, ITCS (2013)
- [MNS09] Moran, T., Naor, M., Segev, G.: An optimally fair coin toss. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 1–18. Springer, Heidelberg (2009)
- [Nao91] Naor, M.: Bit commitment using pseudorandomness. *J. Cryptology* 4(2), 151–158 (1991)
- [NY89] Naor, M., Yung, M.: Universal one-way hash functions and their cryptographic applications. In: STOC, pp. 33–43 (1989)
- [Rom90] Rompel, J.: One-way functions are necessary and sufficient for secure signatures. In: STOC, pp. 387–394 (1990)
- [RTV04] Reingold, O., Trevisan, L., Vadhan, S.: Notions of reducibility between cryptographic primitives. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 1–20. Springer, Heidelberg (2004)
- [Yao82] Yao, A.C.-C.: Theory and applications of trapdoor functions. In: FOCS, pp. 80–91 (1982)