Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 65 (2012)

Limits of Random Oracles in Secure Computation

Mohammad Mahmoody* Hemanta K. Maji† Manoj Prabhakaran‡

December 4, 2012

Abstract

The seminal result of Impagliazzo and Rudich (STOC 1989) gave a black-box separation between one-way functions and public-key encryption: informally, a public-key encryption scheme cannot be constructed using one-way functions as the sole source of computational hardness. In addition, this implied a black-box separation between one-way functions and protocols for certain Secure Function Evaluation (SFE) functionalities (in particular, Oblivious Transfer). Surprisingly, however, *since then there has been no further progress in separating one-way functions and SFE functionalities* (though several other black-box separation results were shown). In this work, we present the complete picture for deterministic 2-party SFE functionalities. We show that one-way functions are black-box separated from *all such SFE functionalities*, except the ones which have unconditionally secure protocols (and hence do not rely on any computational hardness), when secure computation against semi-honest adversaries is considered. In the case of security against active adversaries, a black-box one-way function is indeed useful for SFE, but we show that it is useful only as much as access to an ideal commitment functionality is useful.

Technically, our main result establishes the limitations of random oracles for secure computation. We show that a two-party deterministic functionality f has a secure function evaluation protocol in the random oracle model that is (statistically) secure against semi-honest adversaries if and only if f has a protocol *in the plain model* that is (perfectly) secure against semi-honest adversaries. Further, in the setting of active adversaries, a deterministic SFE functionality f has a (UC or standalone) statistically secure protocol in the random oracle model if and only if f has a (UC or standalone) statistically secure protocol in the commitment-hybrid model.

Our proof is based on a "frontier analysis" of two-party protocols, combining it with (extensions of) the "independence learners" of Impagliazzo-Rudich/Barak-Mahmoody. We make essential use of a combinatorial property, originally discovered by Kushilevitz (FOCS'89), of functions that have semi-honest secure protocols in the plain model (and hence our analysis applies only to functions of polynomial-sized domains, for which such a combinatorial characterization is known).

Keywords: Secure Function Evaluation, Random Oracle Model, One-Way Function, Random Permutation Oracle, Ideal Cipher, Symmetric Primitives, Black-Box Separation.

*Cornell, mohammad@cs.cornell.edu. Supported in part by NSF Award CCF-0746990, AFOSR Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US government.

†Univ. of California, Los Angeles. Work partially done when at Univ. of Illinois, Urbana-Champaign. Partially supported by NSF grants CNS 07-47027 and CNS 07-16626. hmaji@cs.ucla.edu.

‡Univ. of Illinois, Urbana-Champaign. Supported by NSF grant CNS 07-47027. mmp@cs.uiuc.edu.

ISSN 1433-8092

Contents

1 Introduction

How useful is a random oracle in two-party secure function evaluation (SFE)? One obvious use of a random oracle is for implementing commitment. We show that, remarkably, for 2-party SFE[1] *a random oracle by itself is only as useful as a commitment functionality*.

This result has important implications in understanding the "complexity" of secure function evaluation functionalities vis a vis computational primitives like one-way functions. An important goal in cryptography is to understand the qualitative complexity of various cryptographic primitives. In the seminal work of Impagliazzo and Rudich [IR89] a formal framework was established to qualitatively separate cryptographic primitives like symmetric-key encryption and public-key encryption from each other. Understanding that such a separation exists has been hugely influential in theoretical and practical cryptographic research in the subsequent decades: to optimize on both security and efficiency dimensions, a cryptographic construction would be based on symmetric-key primitives when possible, and otherwise is shown to "require" public-key primitives.

Beyond encryption, the result in [IR89] already implies the separation of certain SFE functionalities (in particular, Oblivious Transfer) from one-way functions. Surprisingly, however, *since then there has been no further progress on separating SFE functionalities and one-way functions* (though several other black-box separation results have emerged [Sim98, GKM+00, GMR01, BPR+08, KSY11, MM11]). In this work, we present the complete picture for deterministic 2-party SFE functionalities: we show that in the case of security against semi-honest adversaries, *all of them* are black-box separated from one-way functions, except the ones which are trivial (which have unconditionally perfectly secure protocols). In the case of active adversaries, a black-box one-way function is indeed useful for SFE, but we show that it is useful only as much as access to a commitment functionality is useful (and explicitly characterize the functions for evaluating which it is useful).

Our work could be viewed as a confluence of two largely disjoint lines of work — one on black-box one-way functions, and one on the structure of secure function evaluation functionalities. The former line essentially started with [IR89]. The latter can be traced back to concurrent work [CK89, Bea89, Kus89] which combinatorially characterized which finite (2-party) functionalities have (perfectly) semi-honest secure protocols. This property, called decomposability [Kus89] will be important for us. Several later works obtained such combinatorial characterizations of SFE functionalities in different contexts (e.g., [Kil91, BMM99, KKMO00, MPR10, KM11, Kre11]).

An important ingredient of our proof is the "frontier analysis" approach from [MPR09, MOPR11]. As we shall see, frontier analysis provides a powerful means to explicitly work with otherwise-subtle conditional probabilities, especially as arising in 2-party protocols. In essence, it is simply a means to explicitly keep track of the order in which various events occur in a protocol (or more generally, in a sequence of random variables). But as we shall see, having an explicit mental picture lets us define frontiers and reason about their properties that are *a priori* not obvious (see Figure 2 in Section 5.1, for instance). The proof in [CI93] could in fact be viewed as an instance of frontier analysis (and is one of the earliest ones that the authors are aware of). An instance of such an approach in a non-cryptographic setting is present in the recent work of Barak et al. [BBCR10], who consider frontiers in a protocol where significant amounts of "new and relevant" information is revealed, and use this to reduce the total amount of communication.

[1] We restrict our treatment to SFE functionalities with finite (or at most polynomial-sized) domains. This is because, even without random oracles, a tight characterization of realizable functionalities is known only with this restriction.

1.1 Our Results

We summarize our main results below. Our main result is the following.

Theorem 1.1. *A deterministic two-party function f, with a polynomially large domain, has a semi-honest secure protocol against computationally unbounded adversaries in the random oracle model if and only if f has a perfectly semi-honest secure protocol in the plain model.*

We remark that such f can be explicitly characterized as decomposable functions as defined in [Kus89] (if f is symmetric), or more generally, as those for which the symmetric function f' obtained as the "common information" part of f[2] is decomposable and f and f' are "isomorphic."[3]

In this theorem, as is conventional in much of the work on the combinatorial structure of SFE functionalities, we restrict ourselves to functions whose domain size is polynomial in the security parameter. A full combinatorial characterization of semi-honest securely realizable functions (even in the plain model) is known only with this restriction. In particular, there are undecomposable functions, with super polynomial domain size, which are semi-honest securely realizable. Henceforth, unless mentioned otherwise, whenever we consider a function we shall assume that its domain size is polynomial in the security parameter.

The above result — that random oracles are useless for 2-party SFE — does *not* hold in the case of security against active adversaries. In particular, note that the commitment functionality $\mathcal{F}_{\text{COM}}$, can be constructed UC-securely in a black-box manner from random oracles, and so, all the functions which can be UC-securely computed in the $\mathcal{F}_{\text{COM}}$ hybrid can also be UC-securely computed in the random oracle model. But we shall show that this is all that a random oracle is useful for in 2-party SFE. This follows from Theorem 1.1 and a compiler from [MPR09] that turns semi-honest secure protocols to UC-secure protocols in the $\mathcal{F}_{\text{COM}}$-hybrid model (see proof in Section 6).

Theorem 1.2. *A deterministic two-party function f, with a polynomially large domain, has a statistically UC-secure (and equivalently, a statistically standalone-secure) protocol in the random oracle model if and only if f has a statistically UC-secure (and equivalently, a statistically standalone-secure) protocol in the $\mathcal{F}_{\text{COM}}$ hybrid.*

We remark that such f can be characterized as those for which, on removing all "redundant inputs"[4] one at a time, we obtain a function of the kind in Theorem 1.1.

Blackbox Separations. *Black-box constructions* form a general framework of obtaining a (more complex) cryptographic primitive $\mathcal{Q}$ (e.g., pseudorandom generators) from another (perhaps simpler) cryptographic primitive $\mathcal{P}$ (e.g., one-way functions) while $\mathcal{P}$ is used in the implementation of $\mathcal{Q}$ only as a black-box and the security of $\mathcal{Q}$ is proved based on the security of $\mathcal{P}$ also through a black-box argument. Apart from being the most common kind of reductions used in cryptographic constructions (with "provable security"), black-box

[2] For a deterministic two-party function $f : \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}_A \times \mathcal{Z}_B$, the common information function f' is defined as follows (see for e.g., [MOPR11]): consider the bipartite graph consisting of nodes of the form $(x, a) \in \mathcal{X} \times \mathcal{Z}_A$ and $(y, b) \in \mathcal{Y} \times \mathcal{Z}_B$, with an edge between (x, a) and (y, b) iff $f(x, y) = (a, b)$. Then f' maps (x, y) to the connected component containing (x, a) and (y, b) where $f(x, y) = (a, b)$. Intuitively, $f'(x, y)$ reveals only that part of the information about (x, y) that f reveals to "commonly" to both Alice and Bob (and so they know that it is known to the other party as well).

[3] f_0 and f_1 are isomorphic if there is a UC and semi-honest secure protocol for evaluating either function which uses a single instance of the other function with no other communication. In particular, if either function has a semi-honest secure protocol in the random oracle model (respectively, plain model), then the other one has such a protocol too.

[4] Alice's input x to f is said to be redundant (for security against active adversaries) if there is an input $x' \neq x$ that dominates x: i.e., Alice can substitute x' for x without Bob noticing while still being able to calculate her correct output.

reductions provides us with a framework to understand "complexity" of cryptographic primitives. This line of research was initiated in the seminal work of Impagliazzo and Rudich [IR89] who showed that public-key cryptography is strictly more complex than symmetric-key cryptography (say, one-way functions) under this framework.

Theorem 1.1 is proven in the computationally unbounded setting, and the honest-but-curious adversaries implicit in our proofs use super-polynomial computational power (even if the honest parties were polynomial time). However, similar to the results in [IR89], this can be translated to a statement about black-box separation of semi-honest SFE protocols (for functions without perfectly secure protocols) from one-way functions, in a probabilistic polynomial time (PPT) setting. Intuitively, this is so because a random oracle is a strong one-way function (but for the drawback that it does not have a small code to implement it); so, if one-way function is the sole computational primitive needed for a construction, and it is used in a black-box manner, then it should be possible to base the construction on a random oracle instead. Hence, ruling out secure protocols in the random oracle model in the computationally unbounded setting would rule out protocols in the PPT setting that base their security on one-way functions in a black-box manner. The technicalities depend on the formal definition of black-box reduction. We follow the definitions in [RTV04], with slight technical modifications, to state our results. A formal statement appears in Theorem 7.2. We summarize this result informally below.

Theorem 1.3. *(Informal.) For a deterministic two-party function f, with a polynomially large domain, there is a fully black-box reduction of semi-honest secure function evaluation of f to one-way functions if and only if f has a perfectly semi-honest secure protocol in the plain model.*

Though we state the result for one-way functions, in fact, any collection of primitives that can be constructed from a random oracle (or ideal cipher) or a random permutation oracle[5] in a black-box manner – one-way functions, one-way permutations, collision resistant hash functions, block-ciphers (including exponentially hard versions of these primitives) – is useless for 2-party SFE, if the primitives are used in a fully black-box manner.

As in the case of Theorem 1.1, the above statement can be extended to the case of security against active adversaries.

Theorem 1.4. *(Informal.) For a deterministic two-party function f, with a polynomially large domain, there is a fully black-box reduction of UC (or stand-alone) secure function evaluation of f to one-way functions if and only if f has a statistically UC (or stand-alone) secure protocol in the $\mathcal{F}_{\text{COM}}$-hybrid model.*

Note that, though commitment is already known to be black-box equivalent to one-way functions, *statistical* (standalone) security in the $\mathcal{F}_{\text{COM}}$-hybrid is, on the face of it, more restrictive than standalone security in the PPT setting using fully black-box commitments. Further, the theorem holds for not only one-way functions, but also the other computational primitives mentioned above.

1.2 Related Work

Impagliazzo and Rudich [IR89] showed that random oracles are not useful against a computationally unbounded adversary for the task of secure key agreement. This analysis was recently simplified and sharpened in [BM09]. These results and techniques are one starting point for our result.

[5]We point out that Theorem 1.1 extends to a random permutation oracle, as argued in [IR89]: otherwise, we can construct an efficient distinguisher between a length preserving random oracle and a length preserving random permutation oracle for "long" inputs, and this can be shown to be impossible (as it is improbable to find collisions in a random oracle).

Following [IR89] many other black-box separation results followed (e.g., [Sim98, GMR01, BPR+08, KSY11, MM11]). In particular, Gertner et al. [GKM+00] insightfully asked the question of comparing oblivious-transfer (OT) and key agreement (KA) and showed that OT is strictly more complex (in the sense of [IR89]). Another trend of results has been to prove lower-bounds on the efficiency of the implementation reduction in black-box constructions (e.g., [KST99, GGKT05, LTW05, HHRS07, BM07, BM09, HHRS07]). A complementary approach has been to find black-box reductions when they exist (e.g., [IL89, Ost91, OW93, Hai08, HNO+09]). Also, results in the black-box separation framework of [IR89, RTV04] have immediate consequences for computational complexity theory. Indeed, separations in this framework can be interpreted as new worlds in Impagliazzo's universe [Imp95].

Frontier analysis is possibly implicit in previous works on proving impossibility or lower bounds for protocols. For instance, the analysis in [CI93] very well fits our notion of what frontier analysis is. The analysis of protocols in [CK89, Bea89, Kus89] also have some elements of a frontier analysis, but of a rudimentary form which was sufficient for analysis of perfect security. In [MPR09] frontier analysis was explicitly introduced and used to prove several protocol impossibility results and characterizations. [KMR09] also presented similar results and used somewhat similar techniques (but relied on analyzing the protocol by rounds, instead of frontiers, and suffered limitations on the round complexity of the protocols for which the impossibility could be shown). We also rely on results from [MOPR11] to extend the result to general SFE functionalities as opposed to symmetric SFE functionalities.

In a concurrent work, Haitner, Omri, and Zarosim [HOZ12, HOZ13] show that random oracles are essentially useless for *randomized inputless* functionalities. Similar to our setting, such lower-bounds imply black-box separations from one-way functions. In contrast to their work, our results pertain to *deterministic* functionalities *with inputs*.

1.3 Technical Overview

We rely on a careful combination of the techniques in the black-box separation literature (in particular [IR89, BM09, DLMM11]) and new frontier analysis techniques. Below we briefly explain the overall approach and point out some of the highlights.

A clear starting point of our investigation is the "independence learner" of [IR89, BM09] which shows, in a protocol between Alice and Bob involving private queries to a random oracle, how to make several (but polynomially many) additional queries to the random oracle and make Alice's and Bob's views (conditioned on their inputs) independent of each other. However,from this independence property it is not immediate to conclude that random oracles are useless in SFE protocols. One conjecture (which we are not able to prove) would be that the effect of the random oracle can be "securely simulated" in the plain model, and then any protocol in the random oracle model can be compiled into a plain-model protocol that is as secure as the original one. This would avoid the need to rely on combinatorial characterizations of SFE functionalities, and indeed show that random oracles are useless for virtually any protocol (up to small, but non-negligible errors inherent in the independence learner). However, in this work we do not obtain such a compiler. In particular, *we do not rule out the possibility that in fact random oracles could have unsimulatable effects, and may aid in secure computation of randomized functionalities, or functionalities with super-polynomial input domains*.[6]

[6] An earlier version of this work (presented in [Maj11]), pursued this approach, and appeared to succeed. However, on closer scrutiny a major gap was found in the case when both Alice and Bob can have private inputs, which we have not been able to repair. Indeed, based on our current understanding, *we do not conjecture* that the random oracle can be compiled away from all protocols involving private inputs to both the parties.

This leads us to the techniques used in showing that a symmetric SFE functionality f is semi-honest securely realizable if and only if it is decomposable. The strongest version of this result was proven using frontier analysis in [MPR09]. However, as we shall see, we need a significantly more sophisticated argument here.

1.3.1 Frontier Analysis Meets Random Oracles

First we describe why naïve attempts at generalizing the argument used to characterize functions with SFE protocols in the plain model [MPR09] fail in the random oracle setting.

The plain model result crucially relies on the following "locality" property. When Alice sends the next message in a plain model protocol, she can reveal (i.e., add to the transcript) new information only about her own input but *not* about Bob's inputs. So, during the execution of the protocol, Alice and Bob would alternately reveal information about their inputs x and y respectively. Suppose we define two frontiers: F_X, where (significant, additional) information about x is first revealed, and F_Y where (significant, additional) information about y is first revealed in the transcript. By the locality property, F_X consists of nodes where Alice has just sent out a message, and F_Y consists of nodes where Bob has just sent out a message. Firstly, for the sake of correctness, information about x and y need to be revealed by the end of the protocol, and hence, F_X and F_Y are almost "full" frontiers (i.e., there is only a small probability that an execution finishes without passing through both frontiers).[7] To draw a contradiction we rely on the property that, *for an undecomposable function, it will be insecure for either party to reveal information about their input first*. In terms of the frontiers, this says that it will be insecure if, a (significantly probable) portion of F_X appears above F_Y, or if a (significantly probable) portion of F_Y appears above F_X. Combined with the fact that both frontiers are almost full, this rules out secure protocols for undecomposable functions.

Handling the Random Oracle. In the presence of a random oracle, we lose the locality property (that Alice's message is independent of Bob's input, conditioned on the transcript). It becomes possible that a correlation is established between Alice's and Bob's views via the common random oracle, even conditioned on the transcript. Indeed, given a random oracle, a secure protocol for even OT is possible unless the curious parties query the oracle on points other than what is prescribed by the protocol. Hence, to be meaningful in the presence of an oracle, we must define the information revealed by a transcript as what a curious eavesdropper making additional (polynomially bounded) queries to the oracle, can learn. This is where the independence learner "Eve" of [IR89, BM09] is relevant. Intuitively, Eve attempts to learn as much as possible (staying within a budget of polynomially many oracle queries), by making all "important" queries to the oracle after each message in the protocol. By including the information obtained by Eve into the transcript itself, we can ensure that the frontiers do correspond to points where certain information is revealed, conditioned on the information obtained by Eve. Being a semi-honest setting, it is not relevant when these queries are performed; but for our frontier analysis, it will be important to consider the curious eavesdropper as running concurrently with the protocol, querying the oracle as many times as it wants, after each message in the protocol.

Main Challenge. Once the transcript is augmented with Eve's view, one could hope that the previous analysis from [MPR09] can be applied. Indeed, in this augmented protocol, the locality property is restored.

[7] As we shall see, *for undecomposable functions*, this must hold even if there are inputs for one party (say Bob) for which the function becomes constant. That is, F_Y needs to be crossed even for executions in which Bob's input is a value y for which the function $f(\cdot, y)$ is constant. This is because, by undecomposability, for certain values of Alice's input x, and another input y' for Bob, $f(x, y) = f(x, y')$ where $f(\cdot, y')$ is not constant, and then by security, the execution with input (x, y) has to be close to the execution with input (x, y'). In the latter, information about y needs to be revealed.

However, now we have introduced new messages in the transcript (namely Eve's interaction with the random oracle), and these messages could be correlated with *both* Alice's and Bob's inputs! This is the core issue that we need to tackle.[8]

Our Solution. Now we give an intuitive (but imprecise) description of our proof. As above, we shall define the frontiers F_X and F_Y where information about x and (respectively) about y is first revealed in the (augmented) transcript. Now, information about x or y could be revealed when Alice sends out a message, Bob sends out a message, or Eve obtains its answers from the oracle. We will be able to rule out information about x being revealed by a message from Bob, or information about y being revealed by a message from Alice (this corresponds to Claim 5.6), but this leaves open the possibility that an answer for an Eve query to the oracle reveals information about x and y simultaneously.

To address this, we pursue the following intuition: suppose no information about y has been revealed so far, and Alice sends out a message; suppose some information about x is revealed not immediately by this message, but after Bob (and Eve) carry out oracle queries and respond to Alice's message (but before Alice responds again). (Our concern is that this information could depend on x and y simultaneously.) Then we demonstrate a *curious Bob* strategy that can learn the same information about x, irrespective of his actual input y. The intuition behind this strategy is the following: consider the point immediately after Alice sent out her message. Bob samples for himself a view conditioned on an alternate input y' such that an actual execution with input (x, y') reveals information about x that should not be revealed when Bob's input is y. Bob can simulate the execution with input y' for himself, starting from this point until the next message from Alice, without interacting with Alice; however, the oracle Bob has access to is conditioned on the actual pair of inputs (x, y), and not (x, y'). Clearly, it will be pointless to use this oracle directly to simulate the execution with input (x, y'). A crucial observation at this point (this corresponds to Claim 5.7) is that, it is highly unlikely for an oracle query that is not in Eve's view to be present in both Alice's view and Bob's view (or the sampled view for Bob). This lets Bob simulate an oracle conditioned on (x, y') as follows: if an oracle query is already answered in the sampled view for Bob (with input y'), use it (it is likely not to have been asked by Alice); else, if an oracle query is present in the original view for Bob (but not present in the sampled view, and neither in Eve's view), then "undo" the effect of the query in Bob's view by sampling a new answer for it (again, it is unlikely to have been asked by Alice); if not, use the actual oracle (thus ensuring that any queries already present in Alice's view are consistently answered). This allows curious Bob to seamlessly replace the actual oracle with an oracle consistent with inputs (x, y'), even though he does not know x or Alice's view of the oracle. What facilitates this, in addition to the fact that Eve captures all intersection queries, is the special "modular" nature of the random oracle.

This essentially means that when information about x is revealed, information about y must have been revealed already by the time the last message was sent by Alice (even if the information about x is revealed only during subsequent queries to the oracle by Bob or Eve). Further, as mentioned above, since Alice could not have revealed information about y, this information about y must have been revealed strictly before the last message from Alice, and in particular, strictly before the information about x was revealed. This is captured in Claim 4.2 which implies that (in terms of the simplified presentation above) F_X can be reached only strictly after passing through a node in F_Y.

[8]This is the issue that was not correctly handled in a previous attempt by the authors (in [Maj11]), in trying to compile away the random oracle. The current frontier analysis based approach avoids subtle probabilistic reasoning which is invariably fraught with dangers of false intuition.

Some Technical Issues. Formalizing the above intuitive description presents several challenges. The most important aspect is the appropriate definition of the frontier, and the statement regarding the ordering of the frontiers. For the above curious Bob to have an advantage, the information revealed about x should have been after the last message from Alice. For each node u we define $\mathsf{Apred}(u)$ to correspond to the last message from Alice; however for a node u which itself corresponds to a message from Alice (where the argument relies on the locality property and not the above curious Bob strategy) $\mathsf{Apred}(u)$ is defined as its parent node. Another important issue is that, above we argued in terms of "the probability of reaching a segment." However, this probability depends on the inputs. (The set of nodes in the frontier does not change; only the distribution over them changes.) Whether these probabilities are similar or different depends on whether the inputs have already been distinguished or not. Note that we use properties of these distributions to reason about the ordering of the frontiers, and these distributions themselves depend on the ordering of the frontiers! Much of our technical difficulties arise from circumnavigating potential circularities.

1.3.2 Using the Independence Learner

As mentioned above, a crucial tool for analyzing protocols using a random oracle is to show that by making polynomially many queries to the oracle, an eavesdropper Eve can get sufficient information such that conditioned on this, Alice and Bob's views in the protocol are almost always close to being independent (up to an inverse polynomially small error). This is a delicate argument implicitly proved in [BM09] building on ideas from [IR89], and was first explicitly described in [DLMM11]. The view of such an Eve is part of the augmented transcript, with respect to which the frontiers are defined.

A subtle issue to address when extending this Eve to our case is that Alice and Bob receive inputs from an arbitrary environment and Eve does not see the inputs. In particular, Alice and Bob could receive correlated inputs, and we cannot claim that their views, conditioned on Eve's view, are (almost always, close to being) independent. However, we can create an Eve which is oblivious to the actual inputs, but for *every* input pair (x, y) of inputs, when the protocol is executed with these inputs, Alice's and Bob's views conditioned on Eve's view are (almost always, close to being) independent. For this, we take Eve to be as defined in [BM09] (presented in Lemma A.1), but applied to an inputless protocol obtained by considering our original protocol but with inputs (x, y) that are chosen initially at random (say as part of the randomness of the two parties). Initially this Eve considers the actual input to be of significant probability (since the inputs come from a polynomially large domain). In analyzing this Eve, we rely on an argument that with significant probability, at any round of the protocol, this Eve will consider the actual input to be a likely input (Lemma 2.1).

In our analysis sketched above, there are two guarantees from this Eve that we rely on, captured in Claim 5.6 and Claim 5.7, as described below.

1) Alice's Message Independent of Bob's Input. Firstly, recall that the purpose of introducing Eve's view into the transcript was to restore the "locality property" – i.e., Alice's messages, conditioned on Eve's view, are independent of Bob's view. More precisely, we will need the guarantee that at a point where Alice is about to send a message, if two inputs of Bob, y and y' are both somewhat likely, then Alice's message is almost independent of which of these two inputs Bob has. This is stated in Claim 5.6, and follows from Lemma A.2 proven in Appendix A. Note that we need this to hold (and this holds) only at points where both of Bob's inputs y and y' are somewhat likely. (In using this claim, the points considered will be above the frontier F_Y so that all inputs for Bob are significantly probable.)

2) Collisions of Private Queries Unlikely. The second place where we rely on Eve's properties is in arguing that the curious Bob strategy outlined above works: i.e., that when curious Bob samples a view for himself after Alice sends a message, it is unlikely that there will be an oracle query in either his actual view or in the freshly sampled view that occurs in Alice's actual view, but is not present in Eve's view. This is stated in Claim 5.7 and follows from Lemma A.3 proven in Appendix A. We need this to occur only when the "fake" input y' used for the sampled view is somewhat likely. (Again, the claim will be applied only to points above the frontier F_Y, and all inputs are somewhat likely there.) We remark that, just for the actual views, similar statements were already explicitly proven in [IR89, BM09], bounding the probability of an "intersection query" that is not present in Eve's view. The additional twist in our case is that we need to also consider the view sampled for a "fake" input; further, Bob's views we consider are not at the point Eve finishes a round of oracle queries, but after a subsequent message from Alice.

It is important to note that Bob's views considered here consist of the oracle queries he made only up to the point he sent his previous message to Alice (even though the views include the last message from Alice). Lemma A.3 would *not* be true, if instead we consider Bob's views including oracle queries he makes after receiving Alice's last message. The reason is that the last message sent from Alice can simply tell Bob that Alice has asked a random new query q and Bob might make the same query immediately afterwards. This way, the information that was gathered by Eve till the end of the previous round (before Alice sent her message) is incapable of catching this intersection query.

2 Preliminaries

In this section we introduce some basic notation, conventions and definitions. (Further conventions needed shall be introduced in their respective sections).

2.1 Secure Evaluation of 2-Party Functions

2-Party Functions. A (deterministic) *2-party function* $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathcal{Z}_A \times \mathcal{Z}_B$ maps a pairs of inputs (x, y) (associated with Alice and Bob respectively) to a pair of outputs (a, b) (for the two parties, respectively). For most part in our proofs, we shall be dealing with *symmetric* 2-party functions which produce two identical outputs (or equivalently, a single output given to both parties).

For symmetric functions, an *Alice-cut* is a partition $(X, \bar{X})$ of the input space $\mathcal{X}$ such that for any $x \in X$, $\bar{x} \in \bar{X}$ and $y \in \mathcal{Y}$ $f(x, y) \neq f(\bar{x}, y)$. The *functions associated with an Alice-cut* $(X, \bar{X})$ are the two restrictions of f, restricted to domain $X \times \mathcal{Y}$ and to domain $\bar{X} \times \mathcal{Y}$. A Bob-cut and functions associated with it are defined similarly.

Now, we define *decomposable* functions f in the following recursive manner [Kus89, Bea89]:

1. A constant function is decomposable.
2. If f has an Alice-cut or a Bob-cut and the two functions associated with that cut are both decomposable then f is decomposable.

A function is *undecomposable* if it is *not* decomposable. Moreover, it is said to be *undecomposable at the top-most level*, if $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathcal{Z}$ does not have an Alice-cut or Bob-cut (refer Appendix B for some examples).

Secure Function Evaluation. A Secure Function Evaluation (SFE) functionality is associated with a 2-party function f: the ideal SFE functionality accepts x from Alice, y from Bob, computes $f(x, y) = (a, b)$ and gives a to Alice and b to Bob. We shall refer to the SFE functionality and the two-party function associated with it, interchangeably. For most part, we shall consider protocols for SFE functionalities that are secure against semi-honest adversaries. Our final theorems consider the two standard notions of security against active adversaries as well, namely, standalone security and Universally Composable (UC) security. Mostly we work with statistical security, which places no computational limitations on the parties or environment; but we do state consequences for our results for security in the computational setting as well. We omit a detailed description of the standard security definitions. As it turns out, in our results, there would be no distinction between UC security and standalone security. (Readers unfamiliar with the details of the definitions may ignore the few places in our proofs where we discuss the two notions separately, to establish their similarity.)

Security Definitions. Security of protocols is defined under the standard simulation paradigm. We consider semi-honest security in which the adversary and the simulator are semi-honest (a.k.a. passive or honest-but-curious), and also active-security. In the latter case security can be considered in the standalone setting or the universally composable setting. The statistical difference between the views of the environment in the real and ideal executions, maximized for each simulator over all environments, and then minimized over all simulators, will be called the "security error" of a protocol.

We can in fact work with a (weaker) game based definition of semi-honest security which only requires that if $f(x, y) = f(x, y')$ Alice's views in the two executions with inputs (x, y) and (x, y') should be (statistically) indistinguishable from each other; similarly Bob's views for executions with inputs (x, y) and (x, y) should be indistinguishable, if $f(x, y) = f(x', y)$. This definition is identical to the simulation based definition in the computationally unbounded setting; but when considering the PPT setting (for black-box separation results), the weaker security definition makes our results stronger, and more amenable to being framed in terms of the definitions in [RTV04].

2.2 Random Oracles

An oracle $\mathcal{O}$ is specified by a function (from queries to answers) chosen according to a specified distribution. This choice is made before answering any query, however for the sake of analysis of the protocol we can choose the randomness of the oracle *along the way* as the parties interact (this is also known as the *lazy evaluation* of the oracle). In this paper, we shall use $\mathcal{O}$ which are random oracles, i.e. every query is independently mapped to an image chosen uniformly at random.

Security Parameter of $\mathcal{O}$. We shall associate a security parameter κ with the queries to the oracle, and will invariably require that the length of the queries and their answers is polynomial in κ (e.g., $\mathcal{O}$ for the security parameter κ could be a random function from $\{0, 1\}^\kappa$ to $\{0, 1\}^\kappa$). For simplicity, any protocol using the oracle would make all queries with the same security parameter as the protocol's own security parameter.

Query Operator. For any view V of some oracle algorithm interacting with $\mathcal{O}$, we denote the set of oracle queries made by the algorithm according to the view V by $\mathcal{Q}(V)$.

2.3 Frontiers

Consider a (possibly infinite) sequence of correlated random variables $(\mathbf{m}_1, \mathbf{m}_2, \dots)$. We consider a natural representation of such a sequence as a rooted tree, with each level corresponding to a random variable $\mathbf{m}_i$ and each node v at depth t in the tree is uniquely identified with an assignment of values $(m_1, m_2, \dots, m_t)$ to $(\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_t)$, such that $(m_1, m_2, \dots, m_{t-1})$ is equal to the values identified with its parent node. Then we can identify the sequence of values of these random variables with a unique path in this tree, starting at the root.

We can identify a set of nodes S in this tree with the event that the path corresponding to the values taken by the random variables intersects S. A *frontier* on this tree corresponds to a set F of nodes which is "prefix-free" (i.e., no two nodes in F are on the same path starting at the root). We often define a frontier using a predicate, as the set of nodes which satisfy the predicate but do not have an ancestor which satisfies the predicate (i.e., the predicate is satisfied for the "first time"). Note that the frontier event is deterministic given a node in the tree (though the event could be in terms of the probability of other events at that node).

The tree naturally defines an "ancestor" partial order of the nodes in the tree: we say $u \preceq v$ if u occurs somewhere on the path from the root of the tree to v (u could be identical to v). If $u \preceq v$, but $u \neq v$, then we write $u \prec v$.

Invariably, we consider this tree only with sequence of random variables corresponding to the messages exchanged in a protocol (but possibly augmented by additional messages added for analysis). Though not necessary, it will be convenient to consider the underlying process as consisting of picking a uniformly random input and then executing the protocol. However, clearly, the tree and frontiers can be used to represent any sequence of random variables.

As a simple illustration of the routine arguments we carry out over such a tree, we state and prove a simple lemma (which gets used later in the paper). In Lemma 6.4 of [IR89] it was shown how to obtain an upper-bound on the conditional probability of an unlikely event under a sequence of leaking information. The following lemma can be thought of as a "dual" statement showing that if the event is noticeable, when it actually happens, then it remains noticeable conditioned on a sequence of leakages. More formally we prove the following.

Lemma 2.1. *Consider a sequence of correlated random variables* $(\mathbf{m}_1, \mathbf{m}_2, \dots)$ *For any event* X *jointly distributed with these variables, let* S *be the event that there exists* t *such that* $\mathsf{P}[X \mid (m_1, m_2, \dots, m_t)] < \theta$. *Then it holds that* $\mathsf{P}[S \mid X] < \theta/\mathsf{P}[X]$.

Proof. Consider the tree representing the sequence of random variables $(\mathbf{m}_1, \mathbf{m}_2, \dots)$. The event S corresponds to a subset of nodes in this tree: $S = \{v \mid \mathsf{P}[X \mid v] < \theta\}$. Define U to be the frontier of nodes in S that do not have a strict ancestor in S; namely, $U = \{v \mid v \in S \text{ and for all } u \text{ s.t. } u \prec v, u \notin S\}$. Note that $\mathsf{P}[S \mid X] = \mathsf{P}[U \mid X]$. Further,

$$\mathsf{P}[U \mid X] = \sum_{u \in U} \mathsf{P}[u \mid X] = \sum_{u \in U} \mathsf{P}[X \mid u]\mathsf{P}[u]/\mathsf{P}[X] < \theta \mathsf{P}[U]/\mathsf{P}[X] \leq \theta/\mathsf{P}[X]. \qquad \square$$

A corollary to Lemma 2.1 is that in a protocol execution, the actual inputs of Alice and Bob will not become "unlikely" conditioned on the transcript, except with small probability.

3 Transcript Tree and Other Notation

In this section, first we define the tree notation that is used throughout our analysis. We shall also define the frontiers on this tree that are central to our analysis.

Augmented Protocol Execution. We shall consider two-party protocols Π where Alice and Bob interact to evaluate a (symmetric) function $f : \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}$ on their respective local inputs $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. We shall assume that $|\mathcal{X}|$ and $|\mathcal{Y}$ are both polynomial in the security parameter. Alice and Bob have access to a random oracle $\mathcal{O}$. We "augment" the protocol Π with a "public query strategy" Eve, which can see the publicly generated transcript and can also query the random oracle. For simplicity, we consider Eve to be deterministic (as will be the case in our instantiation of Eve). Later, we will instantiate Eve from Lemma A.1 (applied to an inputless protocol obtained by using uniformly randomly chosen inputs for Π).

When Alice is supposed to generate the next message, she queries the random oracle at some points. Based on her local view, she then generates the next message of the protocol using her next message generation algorithm. Similarly, Bob also generates the next message of the protocol during his turns. Eve, on the other hand, simply performs several queries to the random oracle and announces all her queries and their corresponding answers at the end of her turn. For concreteness we shall assume that the protocol starts with Alice sending a message. Alice and Bob take turns alternately, with Eve getting a turn after every Alice or Bob message (i.e., the messages will be sent by Alice, Eve, Bob, Eve, and again Alice, Eve and so on.).

We shall refer to this protocol as the "augmented protocol" (Π, Eve).

Augmented Transcript Tree $\mathbb{T}^+$. Our analysis considers the transcript tree $\mathbb{T}^+$ of an execution of Π augmented with a public query strategy Eve. The $\mathbb{T}^+$ associated with an augmented protocol (Π, Eve), is the tree as defined in Section 2.3 with the sequence of random variables $(\mathbf{m}_1, \mathbf{m}_2, \dots)$ being the messages added to the transcript of the augmented protocol by Alice, Eve and Bob during an execution. In other words, the nodes in the transcript tree are all the possible partial transcripts in the augmented protocol execution, with a directed edge from a node u to a node v, if the partial transcript associated with v is obtained by adding exactly one message (from Alice, Bob or Eve) to the partial transcript associated with u.

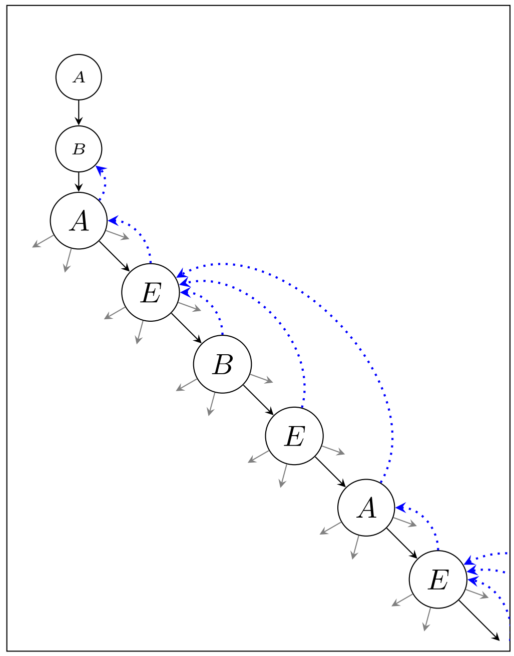

Figure 1: Schematic representation of the nodes in $\mathbb{T}^+$ (including two initial dummy nodes). The nodes are labeled A, B and E, for Alice, Bob and Eve. The dotted lines show the Apred relation.

For convenience we add an initial "dummy" round, in which Alice sends a fixed message followed by Bob sending a fixed message. These correspond to two dummy nodes at the root of $\mathbb{T}^+$. We shall denote by Anodes and Bnodes the sets of Alice and Bob nodes, and by Achildren and Bchildren the sets of (Eve) nodes that are children of, respectively, Alice nodes and Bob nodes. The tree $\mathbb{T}^+$ naturally defines an "ancestor" partial order of the nodes in the tree: we say $u \preceq v$ if u occurs somewhere on the path from the root of the tree to v (u could be identical to v). If $u \preceq v$, but $u \neq v$, then we write $u \prec v$. We define $\mathsf{ancstrs}(v) = \{w | w \preceq v\}$.

An important definition we shall use through out is that of Apred and Bpred nodes.

Definition 3.1 (Apred)**.** For every node v in the transcript tree, except the initial dummy Alice node, we define $\mathsf{Apred}(v)$ as follows:

- If $v \in \mathsf{Achildren}$, then $\mathsf{Apred}(v)$ is the parent of v.
- If $v \notin \mathsf{Achildren}$, we define $\mathsf{Apred}(v)$ to correspond to the last message sent by Alice, before the

transcript reached v: i.e., $\mathsf{Apred}(v) = w$ such that $w \in \mathsf{ancstrs}(v) \cap \mathsf{Achildren}$, and for all $w' \in \mathsf{ancstrs}(v) \cap \mathsf{Achildren}$, $w' \preceq w$.

Note that $\mathsf{Apred}(v) \in \mathsf{Achildren} \cup \mathsf{Anodes}$ and $\mathsf{Apred}(v) \prec v$. Further, for any node v, the sequence $v, \mathsf{Apred}(v), \mathsf{Apred}(\mathsf{Apred}(v)), \ldots$ ends at the initial dummy Alice node.[9] Figure 1 pictorially summarizes the Apred relation.

Similarly, for every node v (except the initial dummy Alice and Bob nodes), we define $\mathsf{Bpred}(v)$ as either the maximal element of $\mathsf{ancstrs}(v) \cap \mathsf{Bchildren}$ (if $v \notin \mathsf{Bchildren}$) or the parent of v (if $v \in \mathsf{Bchildren}$). Note that $\mathsf{Bpred}(v) \in \mathsf{Bchildren} \cup \mathsf{Bnodes}$ and $\mathsf{Bpred}(v) \prec v$.

For any partial transcript w, we define the views of Alice, Bob and Eve consistent with the partial transcript w. The Eve view consistent with w is represented by $V_E(w)$. We represent the distribution of Alice views and Bob views conditioned on w, when their local inputs are x and y, respectively, by $\mathbf{V}_{A,x}(w)$ and $\mathbf{V}_{B,y}(w)$ (the bold face emphasizing that these are distributions). The probability is over the choice of random tapes for Alice and Bob and the random oracle. We emphasize that the local views of parties contain only those query-answer pairs which were generated during next message generation of messages already present in w. So, if Alice sends the next message in a round and the resulting transcript was w, then Bob's views consistent with w will contain only query-answer pairs which were generated in previous rounds. Bob's view gets updated with new query-answer pairs when he sends the next message in the protocol.

Strictly Above a Set: $u \prec F$ and $F_1 \prec F_2$. We shall abuse the $\prec$ notation slightly, and use it in the following senses too: if u is a node and F is a set of nodes, we write $u \prec F$ (read as u is strictly above F) if u can be reached from the root without passing through any node in F (i.e., there is no $v \in F$ such that $v \preceq u$); note that for u to be strictly above F, it is not necessary to have any $v \in F$ such that $u \prec v$. For two sets of nodes F_1, F_2, we define the *event* $F_1 \prec F_2$ to occur if the transcript path of an execution passes through a node $v \in F_1$ strictly before passing through any node in F_2 (it may or may not pass through a node in F_2 afterwards).

4 Overview of Our Analysis

Here we sketch the technical details of our frontier analysis (see Section 1.3 for a motivating discussion, and Section 5 for the remaining details).

Suppose Π is a 2-party protocol using a random oracle $\mathcal{O}$ that ν_0-securely realizes a symmetric SFE functionality f that is not row or column decomposable at the top level (i.e., not even the first step of decomposition is possible; as we shall see, it is enough to rule out protocols for such functionalities). Let Eve be the public query strategy described in Lemma A.1, with an adjustable parameter ε as described there. ($\varepsilon = 1/\operatorname{poly}(\kappa)$ will be tuned later in the proof.) Note that in Lemma A.1, the protocol considered has no inputs; in order to define Eve from this, we use an inputless protocol obtained by running Π with private inputs chosen uniformly at random (as part of Alice's and Bob's local randomness). We shall modify the protocol so that at the end of the protocol, Alice adds the output of the protocol to the transcript. (The simulation error ν_0 at most doubles by this modification.) We consider the transcript tree $\mathbb{T}^+$ as described above, for this protocol Π augmented with Eve.

[9] We added dummy Alice and Bob nodes at the root level to ensure that Apred and Bpred is well-defined for all the original nodes. Note that no information is exchanged until after the protocol passes these dummy nodes, and so these nodes will not be part of any of our frontiers defined later.

Intuitively, we will be arguing that if some information about x has been revealed by the time the transcript reaches a node v, some information about x or y must have already been revealed when it reached $\mathsf{Apred}(v)$. Similarly, for information about y to be revealed at v, some information about x or y should already have been revealed at $\mathsf{Bpred}(v)$. Together these requirements yield a contradiction. To formalize this, we shall define a frontier F_X (and symmetrically F_Y) that consists of nodes v such that the "extra information" revealed about x at v since reaching $\mathsf{Apred}(v)$ is significant.

More precisely we define the following two frontiers on this tree, in terms of two parameters δ and θ (for concreteness, consider $\delta = \frac{1}{N}$, where N is the depth of the tree $\mathbb{T}^+$, and $\theta = \frac{1}{32|\mathcal{X}||\mathcal{Y}|}$).

- $F_X^\theta = \{v | v$ is the first node on the path from root to v s.t. $\exists y \in \mathcal{Y}, x, x' \in \mathcal{X}, \mathsf{P}[y|v] \geq \theta$ and $\mathsf{P}[v|\mathsf{Apred}(v); x, y] > (1+\delta)\mathsf{P}[v|\mathsf{Apred}(v); x', y]\}$
- $F_Y^\theta = \{v | v$ is the first node on the path from root to v s.t. $\exists x \in \mathcal{X}, y, y' \in \mathcal{Y}, \mathsf{P}[x|v] \geq \theta$ and $\mathsf{P}[v|\mathsf{Bpred}(v); x, y] > (1+\delta)\mathsf{P}[v|\mathsf{Bpred}(v); x, y']\}$,

Here, $\mathsf{P}[v|w; x, y]$ denotes the probability (over the random tapes of the parties and the oracle $\mathcal{O}$) of reaching a node v in $\mathbb{T}^+$, conditioned on having reached the node w, when the parties run the protocol honestly with inputs x and y respectively. We shall also write $\mathsf{P}[x|v]$ and $\mathsf{P}[y|v]$ to denote the probabilities of x and y being the inputs for Alice and Bob, respectively, conditioned on a protocol execution with a *uniformly random input pair* reaching the node v.[10] Intuitively, the quantity $\max_{x,x',y} |\log \mathsf{P}[v|w; x, y] - \log \mathsf{P}[v|w; x', y]|$ measures the amount of information about Alice's input that is revealed at v, since passing through $w = \mathsf{Apred}(v)$. This quantity is "significant" if it is beyond a threshold $\log(1+\delta)$ (where, for concreteness, $\delta = 1/N$, N being the depth of $\mathbb{T}^+$) and if it is realized by a y which is somewhat likely (i.e., $\mathsf{P}[y|v] \geq \theta$). In our proofs, it will be useful to consider frontiers F_X^0 and F_Y^0 which are defined identically as F_X^θ and F_Y^θ, but with $\theta = 0$, i.e. these frontiers are considered without the restriction of $\mathsf{P}[y|v] \geq \theta$ and $\mathsf{P}[x|v] \geq \theta$ respectively.

Based on the correctness and the security of the protocol, *and using the fact that f is undecomposable at the top level*, we shall first prove that these frontiers are almost "full frontiers" (when ν_0, the security error for Π, is negligible and θ is set sufficiently small):

Claim 4.1. *On an execution over $\mathbb{T}^+$ with a random input pair (x, y), for any value of θ, the probability that the transcript does not pass through F_X^θ (or symmetrically, F_Y^θ) is at most* $\text{poly}(|\mathcal{X}||\mathcal{Y}|) \cdot \theta + O(\nu_0)$.

This is proven as Claim 5.3. Given that these frontiers exist, next we prove a restriction on how they can occur relative to each other, leading to our final contradiction. Intuitively, the claim states the following: suppose a transcript passes through a node $u \in F_X^\theta$; in a secure protocol not only should u occur only at or below the frontier F_Y^θ, but even $\mathsf{Apred}(u)$ should occur only at or below F_Y^θ; that is a node in F_Y^θ should occur *strictly above* u. (Similarly, for $v \in F_Y^\theta$ and the frontier F_X^θ.)

Claim 4.2. *Consider running the execution on $\mathbb{T}^+$ with a random input (x, y) where ε is the parameter of the Independence Learner Eve. The probability that the transcript passes through a node $u \in F_X^\theta$ such that* $\mathsf{Apred}(u) \prec F_Y^\theta$ *is at most*

$$\text{poly}\left(\frac{N|\mathcal{X}||\mathcal{Y}|}{\theta}\right) \cdot (\varepsilon^{\Omega(1)} + \nu_0) + \text{poly}(|\mathcal{X}||\mathcal{Y}|) \cdot \theta.$$

[10] In all our equations, we use the convention that the probability of an event conditioned on a zero-probability event is zero. Alternately, we can avoid this by assuming, adding a negligible security error, that for any pair of inputs, any node in $\mathbb{T}^+$ is reached with positive probability.

Similarly, the probability that the transcript passes through a node $v \in F_Y^\theta$ *such that* $\mathsf{Bpred}(v) \prec F_X^\theta$ *is bounded by the same quantity.*

Once we prove this claim (as Claim 5.1), the required contradiction follows easily: by setting θ small enough (but $\Omega(1/\operatorname{poly}(|\mathcal{X}||\mathcal{Y}|))$), and choosing ε for the independence learner appropriately (note that this does not affect N), the bounds in above claims can all be driven below, say, any constant (for sufficiently large values of the security parameter). Thus with positive probability the transcript must pass through $u \in F_X^\theta$ and $v \in F_Y^\theta$, with $v \prec u$ and $u \prec v$, giving us the desired contradiction.

To prove Claim 4.2, technically, it is more convenient to bound the probability of encountering $\widetilde{F_X} = \{u | u \in F_X^\theta \text{ and } \mathsf{Apred}(u) \prec F_Y^0\}$ (instead of $u \in F_X^\theta$ such that $\mathsf{Apred}(u) \prec F_Y^\theta$). The difference between these two events can be bounded relatively easily (see the proof in Section 5.1 for details). In particular, for this we use the above Claim 4.1 (with $\theta = 0$) and a bound on the probability of F_X^0 appearing strictly above F_X^θ and F_Y^θ (proven as Claim 5.4):

$$\mathsf{P}[F_Y^0 \prec (F_X^\theta \cup F_Y^\theta)] \leq \theta \operatorname{poly}(|\mathcal{X}||\mathcal{Y}|) \tag{1}$$

Intuitively, the bound above says that if F_Y^0 is encountered strictly above F_X^θ, then it is very likely to occur together with F_Y^θ; hence when a part of F_X^θ occurs at or above F_Y^θ (so that its Apred is strictly above F_Y^θ) it is very likely to be at or above F_Y^0 too. To upper bound the probability of the former, it is enough to upper bound the probability of the latter.

Bounding $\mathsf{P}[\widetilde{F_X}]$ (the probability of reaching $\widetilde{F_X}$ with uniformly random inputs) involves several parts:

- **Part 1:** Firstly, we show that we can concentrate on a 2×2 minor of the function f: that is, $\hat{x}_0, \hat{x}_1 \in \mathcal{X}$ and $\hat{y}_0, \hat{y}_1 \in \mathcal{Y}$ such that $f(\hat{x}_0, \hat{y}_0) = f(\hat{x}_1, \hat{y}_0)$ (but $f(\hat{x}_0, \hat{y}_1) \neq f(\hat{x}_1, \hat{y}_1)$ if $\widetilde{F_X}$ has significant probability). We show that there exists a segment $\widehat{F_X} \subseteq \widetilde{F_X}$ such that the inputs $(\hat{x}_0, \hat{y}_1)$ and $(\hat{x}_1, \hat{y}_1)$ are distinguished at $\widehat{F_X}$, and $\mathsf{P}[\widetilde{F_X}] \leq \operatorname{poly}(|\mathcal{X}||\mathcal{Y}|/\theta)\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1]$.[11]

In the rest of the proof we need to bound $\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1]$. The segment $\widehat{F_X}$ splits into two parts: nodes u with $\mathsf{Apred}(u)$ being an Alice node, denoted by $\widehat{S_X}$, and the ones with $\mathsf{Apred}(u)$ being a child of an Alice node, denoted by $\widehat{R_X}$.

- **Part 2:** Using Lemma A.1 we show that Alice's message cannot reveal any (significant) information about Bob's input, given the information already present in the transcript of the augmented execution (in Lemma A.2). This is used to bound $\mathsf{P}[\widehat{S_X}|\hat{x}_0, \hat{y}_1]$. Note that this part is analogous to the argument when no oracle is present, though more involved (without oracles, this property is a trivial consequence of the nature of a protocol).

- **Part 3:** The most involved part is to bound $\mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1]$. Here we want to bound the probability that a distinction between $\hat{x}_0$ and $\hat{x}_1$ is revealed (when Bob's input is $\hat{y}_1$) at a node u that is not a child of an Alice node, but at $w = \mathsf{Apred}(u)$ the distinction between $\hat{y}_0$ and $\hat{y}_1$ has not been made. Since at w, $\hat{y}_0$ and $\hat{y}_1$ is not distinguished by the transcript, in an execution with his actual input being $\hat{y}_0$, on hitting the node w, Bob can *mentally switch his input to* $\hat{y}_1$ — i.e., sample a view (including answers from the oracle) consistent with the transcript and input $\hat{y}_0$. We would like to argue that then Bob can continue the execution of the protocol (till before Alice should send the next message) and check if it hits u or not, to distinguish between $\hat{x}_0$ and $\hat{x}_1$. However, the execution depends on the random oracle which in turn is correlated with

[11] Note that the need for working with F_X^θ and F_Y^θ rather than just F_X^0 and F_Y^0 is that in this part we rely on the "distinguishing input" being somewhat likely. (If $\theta = 0$ the above bound is useless.)

both parties' inputs. So Bob cannot sample a correctly distributed random oracle (since he does not know Alice's input) nor directly use the actual random oracle he has access to (since it is conditioned on his actual input $\hat{y}_0$ and not $\hat{y}_1$).

The main idea here is that the independence guarantee from Lemma A.1 can be used to let Bob "edit" the actual random oracle (conditioned on $(x, \hat{y}_0)$) to simulate a random oracle conditioned on $(x, \hat{y}_1)$ (without knowing x). The editing involves inserting answers consistent with a sampled view (with input $\hat{y}_1$), "deleting" answers not present in this sampled view, but is present in the actual view (with input $\hat{y}_0$) and using the original oracle for queries not answered in the sampled view or the actual view. (See Figure 3 for an illustration.) The "safety condition" in Claim 5.7 assures that the queries from the sampled view that are not in Eve's view (for which the answers from the sampled view are used) and the queries from the original view that are not in Eve's view (for which random answers are used) are both unlikely to be in in Alice's view; this lets us show that the oracle resulting from the editing is correctly conditioned on the input pair $(x, \hat{y}_1)$.

The final (passive) attack involves carrying out the above attack at every node w and checking if the curious exploration hits the segment $\widehat{R_X}$ in any such exploration. We show that if $\mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1]$ has significant probability then it will be more likely for the exploration to hit $\widehat{R_X}$ in the exploration with input $(\hat{x}_0, \hat{y}_0)$ than in the exploration with input $(\hat{x}_1, \hat{y}_0)$, thereby violating the security condition.

Throughout the argument, translating intuitive statements about information and probability is complicated by the fact that the probability of reaching different nodes depends on the inputs themselves. While intuitively, some of these distributions must be close to each other until the frontiers F_X and F_Y are crossed, we cannot often leverage this intuition without being trapped in circular arguments. Nevertheless, going through several carefully chosen intermediate steps, we can relate the advantage obtained by Bob in distinguishing $\hat{x}_0$ and $\hat{x}_1$ when using input $\hat{y}_1$, with that he obtains when using input $\hat{y}_0$ with the above attack.

5 Detailed Proof of Theorem 1.1

In this section we present the remaining details of the proof of Theorem 1.1, that were sketched in Section 4.

Recall the setting introduced in Section 4: f is a deterministic symmetric two-party function which is undecomposable at the top-most level (i.e., not even the first step of decomposition is possible). Suppose Π is a semi-honest secure SFE protocol for f using a random oracle $\mathcal{O}$ with simulation error ν_0. We defined an augmented transcript tree $\mathbb{T}^+$, and frontiers F_X^θ and F_Y^θ in $\mathbb{T}^+$. First, we shall state our main technical claim about these frontiers in Section 5.1, and show how it follows from several sub-claims that are proven in subsequent sections. Based on Claim 5.3 and Claim 5.1, we present the proof of Theorem 1.1 in Section 5.2. The sub-claims used in the proof of Claim 5.1 are proven in Section 5.3, Section 5.4 and Section 5.5.

The technical heart of the proof appears in Section 5.5, which is part of the proof of Claim 5.1.

5.1 Frontier Ordering

In this section we shall prove the claim regarding the frontier ordering, Claim 4.2. The claim bounds the probability (with uniformly random inputs) of the transcript encountering the following part of the frontier F_X^θ:

$$\breve{F}_X = \{u | u \in F_X^\theta \text{ and } \mathsf{Apred}(u) \prec F_Y^\theta\}.$$

Figure 2 shows this part schematically.

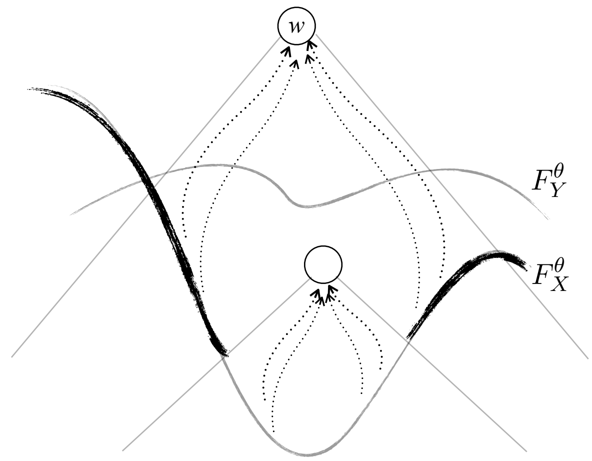

Figure 2: A schematic representation of the segment $\breve{F}_X$ (indicated by thicker line). The dotted lines connect nodes to their Apred nodes (see Figure 1). $\breve{F}_X$ contains those nodes $u \in F_X^\theta$ such that $\mathsf{Apred}(u)$ occurs strictly above F_Y^θ. We seek to upperbound the probability $\mathsf{P}[\breve{F}_X]$ (when the inputs are uniformly chosen).

Claim 5.1. *Let* $\breve{F}_X := \{u | u \in F_X^\theta$ *and* $\mathsf{Apred}(u) \prec F_Y^\theta\}$ *and* $\breve{F}_Y := \{u | u \in F_Y^\theta$ *and* $\mathsf{Bpred}(u) \prec F_X^\theta\}$. *Then there exist polynomials* ξ, ξ' *and* $\widehat{\varepsilon}_0, \widehat{\varepsilon}_1 = \varepsilon^{\Omega(1)} \operatorname{poly}(|\mathcal{X}||\mathcal{Y}|\kappa)$*, such that for any value of* θ,

$$\mathsf{P}[\breve{F}_X] \leq \xi\left(N|\mathcal{X}||\mathcal{Y}|/\theta\right) \cdot \left(\widehat{\varepsilon}_0 + \widehat{\varepsilon}_1 + \nu_0\right) + \xi'(|\mathcal{X}||\mathcal{Y}|) \cdot \theta \tag{2}$$

$$\mathsf{P}[\breve{F}_Y] \leq \xi\left(N|\mathcal{X}||\mathcal{Y}|/\theta\right) \cdot \left(\widehat{\varepsilon}_0 + \widehat{\varepsilon}_1 + \nu_0\right) + \xi'(|\mathcal{X}||\mathcal{Y}|) \cdot \theta \tag{3}$$

Proof. We shall prove Eq. 2 (the second part being symmetrical). That is, we are interested in bounding the probability that, on running the execution on $\mathbb{T}^+$ with uniformly random inputs (x, y), the transcript reaches a node in $\breve{F}_X = \{u | u \in F_X^\theta$ and $\mathsf{Apred}(u) \prec F_Y^\theta\}$. We say that the event $\breve{F}_X$ occurs, if the path from root to the generated transcript passes through a node in $\breve{F}_X$.

To obtain an upper bound on $\mathsf{P}[\breve{F}_X]$, we first observe that the event $\breve{F}_X$ implies the occurrence of one of the following three events:

1. Event $\overline{F_Y^0}$: the transcript path does not pass through any node in F_Y^0.
2. Event $F_Y^0 \prec (F_X^\theta \cup F_Y^\theta)$: the transcript path passes through a node $z \in F_Y^0$ and $z \prec (F_X^\theta \cup F_Y^\theta)$ (i.e., there is no node $v \in F_X^\theta \cup F_Y^\theta$ such that $v \preceq z$).
3. Event $\widetilde{F_X}$: the path passes through $\widetilde{F_X}$ which is defined similarly to $\breve{F}_X$, but replacing F_Y^θ by F_Y^0. i.e.,

$$\widetilde{F_X} = \{u | u \in F_X^\theta \text{ and } \mathsf{Apred}(u) \prec F_Y^0\}.$$

To see this, suppose $\breve{F}_X$ is encountered, but neither of the first two events occur; then transcript path passes through $u \in \breve{F}_X$, and a node $z \in F_Y^0$, and a node $v \preceq z$ such that $v \in F_X^\theta \cup F_Y^\theta$. We argue that in this case $\mathsf{Apred}(u) \prec z$; then, since F_Y^0 is part of a frontier, $\mathsf{Apred}(u) \prec F_Y^0$ and hence $u \in \widetilde{F_X}$. This is because:

- If $v \in F_X^\theta$, then $u = v$ (since $u \in \breve{F}_X$ and $v \in F_X^\theta$ are on the same path), and $\mathsf{Apred}(u) \prec u = v \preceq z$.
- If $v \in F_Y^\theta$, then $v = z$ (since, $v \in F_Y^\theta$, $z \in F_Y^0$ and $v \preceq z \Rightarrow z = v$) and further, since $u \in \breve{F}_X$ and $v \in F_Y^\theta$ are on the same path, by definition of $\breve{F}_X$, $\mathsf{Apred}(u) \prec v = z$.

Thus, it suffices to upper bound the probabilities of each of these three events. We will be able to easily bound both $\mathsf{P}[\overline{F_Y^0}]$ and $\mathsf{P}[F_Y^0 \prec (F_X^\theta \cup F_Y^\theta)]$, (Claim 4.1 — proven as Claim 5.3 – and Eq. 1 – proven as Claim 5.4 – respectively). The main technical difficulty is in bounding $\mathsf{P}[\widetilde{F_X}]$, which is carried out in Claim 5.5. From these three claims we get

$$\mathsf{P}[\overline{F_Y^0}] = O(\nu_0) \qquad \text{(By Claim 4.1, with } \theta = 0)$$
$$\mathsf{P}[F_Y^0 \prec (F_X^\theta \cup F_Y^\theta)] \le \mathrm{poly}(|\mathcal{X}||\mathcal{Y}|)\theta \qquad \text{(By Eq. 1)}$$
$$\mathsf{P}[\widetilde{F_X}] \le \mathrm{poly}\left(\frac{N|\mathcal{X}||\mathcal{Y}|}{\theta}\right)(\widehat{\varepsilon}_0 + \widehat{\varepsilon}_1 + \nu_0) \qquad \text{(By Claim 5.5)}$$

Adding the three, we get the required bound. □

5.2 Proof of Theorem 1.1

The main part of the proof proves the impossibility of a semi-honest secure SFE protocol, even using random oracles, for a symmetric function f that is undecomposable at the top-level. We shall shortly see that this is enough.

So, suppose f is a 2-party symmetric function that is undecomposable at the top-most level, and Π is a semi-honest secure protocol using a random oracle $\mathcal{O}$, for the SFE funcionality evaluating f, with simulation error ν_0. This is the setting under which the frontiers in $\mathbb{T}^+$ are defined, and Claim 5.3 and Claim 5.1 hold. The proof follows by a deriving a contradiction from these two claims (instantiated with suitable parameters).

We shall set $\theta = \min\{\frac{1}{8\zeta_0(|\mathcal{X}||\mathcal{Y}|)}, \frac{1}{8\xi'(|\mathcal{X}||\mathcal{Y}|)}\}$ where the ζ_0 and ξ' are as in Claim 5.3 and Claim 5.1 (in fact, $\theta = \Theta(\frac{1}{|\mathcal{X}||\mathcal{Y}|})$, by following the proofs of the various claims), and then choose a small enough (but $1/\mathrm{poly}(\kappa)$) value of ε so that $(\widehat{\varepsilon}_0 + \widehat{\varepsilon}_1) \le \frac{1}{8\xi(N|\mathcal{X}||\mathcal{Y}|/\theta)}$ (which is possible since $1/\theta$ is $\mathrm{poly}(\kappa)$ and $\widehat{\varepsilon}_0$ and $\widehat{\varepsilon}_1$ are $\varepsilon^{\Omega(1)}\,\mathrm{poly}(|\mathcal{X}||\mathcal{Y}|\kappa)$), so that (for large enough κ)

$$\mathsf{P}[\overline{F_X^\theta}] + \mathsf{P}[\overline{F_Y^\theta}] \le 2\Big(c_0\nu_0 + \zeta_0(|\mathcal{X}||\mathcal{Y}|)\theta\Big) < \frac{1}{3} \qquad \text{By Claim 5.3}$$
$$\mathsf{P}[\breve{F}_X] < \frac{1}{3} \quad \text{and} \quad \mathsf{P}[\breve{F}_Y] < \frac{1}{3} \qquad \text{By Claim 5.1}$$

So, with non-zero probability, for a random input pair (x, y), the honestly generated transcript passes through both F_X^θ and F_Y^θ, but avoids both events $\breve{F}_X$ and $\breve{F}_Y$. Consider one such transcript τ. Let u and v be the intersection of this path with the frontiers F_X^θ and F_Y^θ. For this transcript τ: $v \preceq \mathsf{Apred}(u)$ (since of $u \notin \breve{F}_X$) and $\mathsf{Apred}(u) \prec u$ (by definition of Apred), i.e. $v \prec u$. Symmetrically, we also get: $u \preceq \mathsf{Bpred}(v)$ and $\mathsf{Bpred}(v) \prec v$, and hence $u \prec v$. This gives us a contradiction as desired.

Extending to all 2-party functions. Above we showed that any symmetric 2-party function that is undecomposable at the top-level does not have an SFE protocol secure against semi-honest adversaries, in the random oracle model. Now we extend this to show that the only 2-party functions for which semi-honest secure protocols exist in the random oracle model are those for which (perfectly) semi-honest secure protocols exist in the plain model. We do this in two steps, first for symmetric 2-party functions and then for general 2-party functions. But first we state a claim that we will need (in the second step).

Claim 5.2. *If a (not necessarily symmetric) 2-party function f_0 has a semi-honest secure protocol in the random oracle model (resp. plain model), it must be "isomorphic" to a symmetric 2-party function f_1 that has a semi-honest secure protocol in the random oracle model (resp. plain model).*

This is because, by a result in [MOPR11], if a 2-party function f_0 is not isomorphic to a certain symmetric 2-party function f_1 (namely, the "common information function of f_0 mentioned in Footnote 2), then f_0 is complete against semi-honest adversaries. But a complete functionality cannot have a semi-honest secure protocol in the plain or random oracle model (as otherwise all functionalities will have semi-honest secure protocols in the random oracle model, contradicting the above results.)

Below are the two steps to complete the proof of Theorem 1.1.

1. Firstly, we argue that if a *symmetric* 2-party function f_1 has a semi-honest secure protocol in the random oracle model, it must be decomposable (and hence has a perfectly semi-honest secure protocol). This is because, if f_1 is undecomposable, then it has a minor f which is undecomposable at the top-level. Further, if f_1 is semi-honest securely realizable using a random oracle, so is every minor of f_1, including f, which contradicts our above result.

2. Next, if a general 2-party function f_0 has a semi-honest secure protocol in the random oracle model, then by Claim 5.2, there is a symmetric 2-party function f_1 that is isomorphic to f_0 and has a semi-honest secure protocol in the random oracle. By the previous point, f_1 has a perfectly semi-honest secure protocol in the plain model, and as f_0 is isomorphic to f_1, so does f_0.

5.3 Bounding probability of events $\overline{F_X^\theta}$ and $\overline{F_Y^\theta}$

In this section we prove Claim 4.1 (restated below).

Claim 5.3. *There exists a constant c_0 and a polynomial ζ_0 such that, on executing the augmented protocol with a random input pair (x, y), $\mathsf{P}[\overline{F_X^\theta}]$ and $\mathsf{P}[\overline{F_Y^\theta}]$ are both at most $c_0\nu_0 + \zeta_0(|\mathcal{X}||\mathcal{Y}|) \cdot \theta$.*

Proof. We shall just show that $\mathsf{P}[\overline{F_X^\theta}] \le p^* = (5 + (1 + \delta)^N)\nu_0 + |\mathcal{X}||\mathcal{Y}|\theta$ (so that $c_0 = (5 + (1 + \delta)^N)$ and $\zeta_0(\alpha) = \alpha$). The bound on $\mathsf{P}[\overline{F_Y^\theta}]$ follows similarly. We shall, in fact, show the stronger result that $\mathsf{P}[\overline{F_X^\theta}|x, y] \le p^*$, for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$.

Let S be the set of all complete transcripts such that none of their ancestors lie in F_X^θ. First, consider any input pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $f(\cdot, y)$ is not a constant function; we shall upper-bound the probability $\mathsf{P}[S|x, y]$ by $p^* - 4\nu_0$.

Let the frontier $U(y)$ be the set of nodes u where, for the first time on a path from the root, $\mathsf{P}[y|u] < \theta$. Let $L(y) = \{u \in S | u \prec U(y)\}$ be the part of S which is strictly above $U(y)$. Then $\mathsf{P}[S|x, y] \le \mathsf{P}[U(y)|x, y] + \mathsf{P}[L(y)|x, y]$. Firstly,

$$\mathsf{P}[U(y)|x, y] = \sum_{u \in U(y)} \mathsf{P}[u|x, y] = \sum_{u \in U(y)} \mathsf{P}[x, y|u]p[u]/\mathsf{P}[x, y] = |\mathcal{X}||\mathcal{Y}| \sum_{u \in U(y)} \mathsf{P}[x, y|u]p[u]$$
$$\le |\mathcal{X}||\mathcal{Y}| \sum_{u \in U(y)} \mathsf{P}[y|u]p[u] < |\mathcal{X}||\mathcal{Y}|\theta \sum_{u \in U(y)} \mathsf{P}[u] \le |\mathcal{X}||\mathcal{Y}|\theta.$$

For nodes $v \in L(y)$, we have $\mathsf{P}[y|u] \geq \theta$ for all $u \preceq v$. Recall that v does not have an ancestor in F_X^θ. So, it must be the case that, for all $x, x' \in \mathcal{X}$ we have $\mathsf{P}[v|x,y] \leq (1+\delta)^N \mathsf{P}[v|x',y]$. Since $f(\cdot, y)$ is not a constant function, there exists $x' \in \mathcal{X}$ such that $f(x,y) \neq f(x',y)$. We can partition the set $L(y)$ into two sets:

1. $C(y)$: Those transcripts $v \in L(y)$ whose associated output is $f(x,y)$, i.e. those transcripts which provide correct output when the input is (x,y), and

2. $W(y)$: Those transcripts $v \in L(y)$ whose associated output is $\neq f(x,y)$, i.e. those transcript which provide wrong output when the input is (x,y).

Since, the simulation error is at most ν_0, we can conclude that $\mathsf{P}[W(y)|x,y] \leq \nu_0$. Further, observe that the output associated with the transcripts in $C(y)$ are incorrect for input (x',y). Therefore, $\mathsf{P}[C(y)|x',y] \leq \nu_0$. But, $\mathsf{P}[C(y)|x,y] \leq (1+\delta)^N \mathsf{P}[C(y)|x',y] \leq (1+\delta)^N \nu_0$. Now, we can claim that the $\mathsf{P}[L(y)|x,y] \leq (1+(1+\delta)^N)\nu_0$.

Adding these two results, we can conclude that

$$\mathsf{P}[S|x,y] \leq p^* - 4\nu_0$$

Now, we consider any $(x,y) \in \mathcal{X} \times \mathcal{Y}$ such that $f(\cdot, y)$ is a constant function. Since f is undecomposable at the top-most level, there exists $x' \in \mathcal{X}$ and $y' \in \mathcal{Y}$ such that $f(x',y) = f(x',y')$ and $f(\cdot, y')$ is not a constant function. Thus, by security condition, we can conclude that the final transcript distributions induced by (x,y) and (x',y') have at most $4\nu_0$ statistical distance. Thus, to complete the proof of the theorem[12] :

$$\mathsf{P}[S|x,y] \leq \mathsf{P}[S|x',y'] + 4\nu_0 \leq p^* \qquad \square$$

5.4 Bounding probability of event $F_Y^0 \prec (F_X^\theta \cup F_Y^\theta)$

Claim 5.4. *On executing the augmented protocol with a random input pair (x,y), $\mathsf{P}[F_Y^0 \prec (F_X^\theta \cup F_Y^\theta)]$ is at most $(1+(1+\delta)^N)|\mathcal{X}||\mathcal{Y}|\theta$. The same bound holds for $\mathsf{P}[F_X^0 \prec (F_X^\theta \cup F_Y^\theta)]$.*

Proof. Let S be the set of nodes $v \in F_Y^0$ such that for all $u \preceq v$, $u \notin F_X^\theta \cup F_Y^\theta$, i.e. $v \prec (F_X^\theta \cup F_Y^\theta)$. We shall bound $\mathsf{P}[S|x,y]$, for each input pair $(x,y) \in \mathcal{X} \times \mathcal{Y}$. Fix an input pair (x,y). Let $U(x,y)$ be the frontier of nodes v where for the first time $\mathsf{P}[x|v] < \theta$ or $\mathsf{P}[y|v] < \theta$. Let $L(x,y) = \{u | u \in S, \text{ and } u \prec U(x,y)\}$ be the part of S which is strictly above $U(x,y)$. We shall bound $\mathsf{P}[S|x,y] \leq \mathsf{P}[U(x,y)|x,y] + \mathsf{P}[L(x,y)|x,y]$, by bounding the two terms separately.

$$\begin{aligned}\mathsf{P}[U(x,y)|x,y] &= \sum_{v \in U(x,y)} \mathsf{P}[v|x,y] = \sum_{v \in U(x,y)} \mathsf{P}[x,y|v] \cdot \mathsf{P}[v] / \mathsf{P}[x,y] \\ &\leq |\mathcal{X}||\mathcal{Y}| \sum_{v \in U(x,y)} \min\{\mathsf{P}[x|v],\ \mathsf{P}[y|v]\} \cdot \mathsf{P}[v] < \theta |\mathcal{X}||\mathcal{Y}| \sum_{v \in U(x,y)} \mathsf{P}[v] \leq \theta |\mathcal{X}||\mathcal{Y}|\end{aligned}$$

[12] We note that this bound is not restricted only to the uniform distribution over input pairs. In fact, for *any* input pair distribution such that $\mathsf{P}[x,y]$ is a function of the output $f(x,y)$, $\mathsf{P}[S] \leq p^*$.

To bound $\mathsf{P}[L(x,y)|x,y]$, we partition $L(x,y)$ into $L_{\tilde{x}}(x,y) \subseteq L(x,y)$, one for each $\tilde{x} \in \mathcal{X} \setminus \{x\}$, such that for $v \in L_{\tilde{x}}(x,y)$, v is included in F_Y^0 because $\exists y', y''$ such that $\mathsf{P}[v|\mathsf{Bpred}(v); \tilde{x}, y'] > (1+\delta)\mathsf{P}[v|\mathsf{Bpred}(v); \tilde{x}, y'']$. Note that $\tilde{x} \neq x$, otherwise $v \in F_Y^\theta$. By definition of $L(x,y)$, we have $v \prec (\ F_X^\theta \cup F_Y^\theta \cup U(x,y)\ )$, i.e. $v \prec F_Y^\theta$. This implies that:

$$\mathsf{P}[\tilde{x}|v] < \theta$$

Observe that for all $u \preceq v \in L(x,y)$, we have $\mathsf{P}[y|u] \geq \theta$. But $v \prec F_X^\theta$. Which implies:

$$\mathsf{P}[v|x,y] \leq (1+\delta)^N \mathsf{P}[v|\tilde{x},y]$$

Now, $\mathsf{P}[v|\tilde{x},y] \leq \mathsf{P}[v|\tilde{x}]/\mathsf{P}[y|\tilde{x}] = |\mathcal{Y}|\mathsf{P}[v|\tilde{x}]$. So for $v \in L_{\tilde{x}}(x,y)$,

$$\begin{aligned}\mathsf{P}[v|x,y] &\leq (1+\delta)^N |\mathcal{Y}| \mathsf{P}[v|\tilde{x}] = (1+\delta)^N |\mathcal{X}||\mathcal{Y}| \mathsf{P}[\tilde{x}|v]\mathsf{P}[v] \\ &\leq (1+\delta)^N |\mathcal{X}||\mathcal{Y}|\theta \mathsf{P}[v].\end{aligned}$$

Hence, $\mathsf{P}[L(x,y)|x,y] \leq (1+\delta)^N |\mathcal{X}||\mathcal{Y}|\theta \sum_{v \in L(x,y)} \mathsf{P}[v] \leq (1+\delta)^N |\mathcal{X}||\mathcal{Y}|\theta$. Putting this together with the above bound on $\mathsf{P}[U(x,y)|x,y]$ we get, for all (x,y), $\mathsf{P}[S|x,y] \leq (1+(1+\delta)^N)\theta|\mathcal{X}||\mathcal{Y}|$. Hence, $\mathsf{P}[S] \leq (1+(1+\delta)^N)\theta|\mathcal{X}||\mathcal{Y}|$. □

5.5 Bounding the probability of event $\widetilde{F_X}$

This section carries out the technical heart of the proof. For convenience we define $\mu = (1+\delta)^N$, $\delta' = (1+\delta)^{1/(|\mathcal{X}|-1)} - 1$ and $\delta'' = (1+\delta)^{1/(|\mathcal{Y}|-1)} - 1$. Note that with $\delta = \frac{1}{N}$, $\mu = O(1)$ and $\delta', \delta'' = \Omega(\frac{1}{N(|\mathcal{X}|+|\mathcal{Y}|)})$ (where $|\mathcal{X}|, |\mathcal{Y}| > 1$).

Claim 5.5. *There exist $\widehat{\varepsilon}_0, \widehat{\varepsilon}_1 = \varepsilon^{\Omega(1)} \operatorname{poly}(\kappa|\mathcal{X}||\mathcal{Y}|)$, such that the probability of the augmented protocol with uniformly random inputs reaching $\widetilde{F_X}$ is*

$$\mathsf{P}[\widetilde{F_X}] \leq \frac{2\mu^2|\mathcal{X}|^2|\mathcal{Y}|(1+\delta')N}{\theta\delta'} \left(4\nu_0 + \widehat{\varepsilon}_0 + 2\widehat{\varepsilon}_1\right).$$

The same bound, with δ'' instead of δ', holds for $\mathsf{P}[\widetilde{F_Y}]$.

We focus on proving the first part of this claim (the second part being symmetrical). That is, we are interested in bounding the probability that, on executing Π with uniformly random inputs (x,y), the transcript reaches a node in $\widetilde{F_X} = \{u | u \in F_X^\theta \text{ and } \nexists z \in F_Y^0 \text{ s.t. } z \preceq \mathsf{Apred}(u)\}$.

We break the full proof of the claim into three parts:

1. **Part 1.** We shall show that there exist $\widehat{F_X} \subseteq \widetilde{F_X}$ such that $\mathsf{P}[\widehat{F_X}] \geq \mathsf{P}[\widetilde{F_X}]/(|\mathcal{X}||\mathcal{Y}|)^2$, and there are $\hat{x}_0, \hat{x}_1 \in \mathcal{X}$ and $\hat{y}_0, \hat{y}_1 \in \mathcal{Y}$, such that $f(\hat{x}_0, \hat{y}_0) = f(\hat{x}_1, \hat{y}_0)$, and $\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1]$ is comparable to $\mathsf{P}[\widehat{F_X}]$ (with uniformly random inputs (x,y)), and for every $u \in \widehat{F_X}$, $\hat{y}_1$ sufficiently distinguishes $\hat{x}_0$ and $\hat{x}_1$. More precisely,

$$\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1] \geq \frac{\theta|\mathcal{Y}|}{(1+\delta)^N} \mathsf{P}[\widehat{F_X}], \tag{4}$$

and for all $u \in \widehat{F_X}$, if $w = \mathsf{Apred}(u)$, then $\mathsf{P}[u|w; \hat{x}_0, \hat{y}_1] \geq (1+\delta)^{1/(|\mathcal{X}|-1)} \mathsf{P}[u|w; \hat{x}_1, \hat{y}_1]$, and hence

$$\mathsf{P}[u|w; \hat{x}_0, \hat{y}_1] - \mathsf{P}[u|w; \hat{x}_1, \hat{y}_1] \geq \frac{\delta'}{1+\delta'} \mathsf{P}[u|w; \hat{x}_0, \hat{y}_1]. \tag{5}$$

where $\delta' = (1+\delta)^{1/(|\mathcal{X}|-1)} - 1$.

2. **Part 2.** We shall also show that $\mathsf{P}[\widehat{S_X}|\hat{x}_0, \hat{y}_1]$, where $\widehat{S_X} = \widehat{F_X} \cap \mathsf{Achildren}$, must be "small" if the protocol is secure. (For a node $u \in \widehat{S_X}$, $\mathsf{Apred}(u) \in \mathsf{Anodes}$.)

3. **Part 3.** Then we shall show that $\mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1]$, where $\widehat{R_X} = \widehat{F_X} \backslash \mathsf{Achildren}$, must be small if the protocol is secure. (For a node $u \in \widehat{R_X}$, $\mathsf{Apred}(u) \in \mathsf{Achildren}$.)

   Since $\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1] = \mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1] + \mathsf{P}[\widehat{S_X}|\hat{x}_0, \hat{y}_1]$, Parts 2 and 3 imply that $\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1]$ is small as well. Further, by Part 1, $\mathsf{P}[\widehat{F_X}]$ and, thus, $\mathsf{P}[\widehat{F_X}]$ is small as well.

The error terms $\widehat{\varepsilon}_0$ and $\widehat{\varepsilon}_1$ appear in Parts 2 and 3 respectively, from Claim 5.6 and Claim 5.7. The claims are consequences of the independence properties obtained by Eve of Lemma A.1. Below we state the former claim (and show how it follows from Lemma A.2 proven in Appendix A), which states that Alice's message is almost independent of Bob's input, conditioned on Eve's view thus far.

Claim 5.6. *For all $x \in \mathcal{X}$, $y, y' \in \mathcal{Y}$, if $W \subseteq \mathsf{Anodes}$ is such that for all $w \in W$, $\mathsf{P}[y|w; x], \mathsf{P}[y'|w; x] \geq \sigma$ for $\sigma = \frac{1}{\mathrm{poly}(|\mathcal{X}||\mathcal{Y}|)}$, then, for $\varepsilon \leq 1/\mathrm{poly}(\kappa|\mathcal{X}||\mathcal{Y}|)$ (for some polynomial) and an error parameter $\widehat{\varepsilon}_0 = \varepsilon^{\Omega(1)} \mathrm{poly}(\kappa|\mathcal{X}||\mathcal{Y}|)$, we have*

$$\sum_{w \in W} \mathsf{P}[w|x, y] \cdot \mathsf{SD}\left(\{\mathsf{chldrn}(w)|w; x, y\}, \{\mathsf{chldrn}(w)|w; x, y'\}\right) \leq N\widehat{\varepsilon}_0, \tag{6}$$

where $\{\mathsf{chldrn}(w)|w; x, y\}$ and $\{\mathsf{chldrn}(w)|w; x, y'\}$ stand for the distribution of the next node after w (i.e., Alice's message at w) in $\mathbb{T}^+$ when Π is executed with inputs (x, y) and (x, y') respectively.

Proof. Lemma A.2, stated in terms of a traversal of the tree $\mathbb{T}^+$, partitions the nodes at each level in the tree into three sets, a low-probability set W_0^i such that $\mathsf{P}[W_0^i|x, y] \leq \varepsilon'$, W_1^i such that for $w \in W_1^i$, $\mathsf{P}[y|w; x] < \varepsilon'$ or $\mathsf{P}[y'|w; x] < \varepsilon'$ and W_2^i such that for $w \in W_2^i$, $\mathsf{SD}\left(\{\mathsf{chldrn}(w)|w; x, y\}, \{\mathsf{chldrn}(w)|w; x, y'\}\right) \leq \varepsilon'$. Note that $W_1^i \cap W = \emptyset$ because (for sufficiently small values of ε), $\varepsilon' = \varepsilon^{\Omega(1)} \mathrm{poly}(\kappa|\mathcal{X}||\mathcal{Y}|) < \sigma$. So,

$$\begin{aligned}\sum_{w \in W} &\mathsf{P}[w|x, y] \cdot \mathsf{SD}\left(\{\mathsf{chldrn}(w)|w; x, y\}, \{\mathsf{chldrn}(w)|w; x, y'\}\right) \\ &\leq \sum_i \sum_{w \in W_0^i} \mathsf{P}[w|x, y] + \sum_i \sum_{w \in W_2^i} \mathsf{P}[w|x, y]\varepsilon' \leq N\varepsilon' + N\varepsilon' \leq N\widehat{\varepsilon}_0,\end{aligned}$$

where $\widehat{\varepsilon}_0 = 2\varepsilon'$ □

We mention a few other technical inequalities that are useful in the proof.

For $u \in \widetilde{F_X}$, if $w = \mathsf{Apred}(u)$, then w is strictly above the frontier F_Y^0, and hence

$$\mathsf{P}[w|\hat{x}_0, \hat{y}_0] \geq \frac{1}{(1+\delta)^N} \mathsf{P}[w|\hat{x}_0, \hat{y}_1]. \tag{7}$$

For any subset W of nodes,

$$-2N\nu_0 \leq \sum_{w \in W} \mathsf{P}[w|\hat{x}_0, \hat{y}_0] - \mathsf{P}[w|\hat{x}_1, \hat{y}_0] \leq 2N\nu_0, \tag{8}$$

because $f(\hat{x}_0, \hat{y}_0) = f(\hat{x}_1, \hat{y}_0)$ and by the security guarantee of Π, restricted to the intersection of W with the frontier corresponding to a fixed round number, this summation is at most $2\nu_0$ (since in the ideal world, the simulated views are identical, and for each execution, the error from the simulated distribution is at most ν_0).

It will be useful to relate $\sum_{w\in W}\left(\mathsf{P}[w|\hat{x}_0,\hat{y}_0]\sum_{u\in S_w} g(u,w)\right)$ to $\sum_{w\in W}\left(\mathsf{P}[w|\hat{x}_1,\hat{y}_0]\sum_{u\in S_w} g(u,w)\right)$, where for all $w \in W$, $\sum_{u\in S_w} g(u,w) \leq 1$. This arises for us when S_w forms part of a frontier, and $g(u,w)$ is a probability distribution (possibly conditioned on w) or statistical distance between two probability distributions.

$$\begin{aligned}
&\sum_{w\in W} \mathsf{P}[w|\hat{x}_1,\hat{y}_0] \sum_{u\in S_w} g(u,w) \\
&\quad = \sum_{w\in W} \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{u\in S_w} g(u,w) - \sum_{w\in W} \left(\mathsf{P}[w|\hat{x}_0,\hat{y}_0] - \mathsf{P}[w|\hat{x}_1,\hat{y}_0]\right) \sum_{u\in S_w} g(u,w) \\
&\quad = \sum_{w\in W} \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{u\in S_w} g(u,w) \pm 2N\nu_0 \qquad \text{(By Eq. 8.)}
\end{aligned} \tag{9}$$

Here, we applied Eq. 8 to two subsets of W (where $(\mathsf{P}[w|\hat{x}_0,\hat{y}_0] - \mathsf{P}[w|\hat{x}_1,\hat{y}_0])$ is positive and negative, respectively.) and also used the fact that $\sum_{u\in S_w} g(u,w) \leq 1$.

Part 1. We define $\widehat{F_X}$ and $(\hat{x}_0, \hat{x}_1, \hat{y}_0, \hat{y}_1)$.

For any node $u \in F_X^\theta$, there exists $y_u^* \in \mathcal{Y}$ and some $x, x' \in \mathcal{X}$ such that $\mathsf{P}[u|w;x,y_u^*] > (1+\delta)\mathsf{P}[u|w;x',y_u^*]$, where $w = \mathsf{Apred}(u)$. W.l.o.g, we consider x which maximizes $\mathsf{P}[u|w;x,y_u^*]$; we call the maximum value $\alpha(u,y_u^*)$. Since f is not row-decomposable at the top-level, there exist a sequence of $t+1 \leq |\mathcal{X}|$ values $x_0, \cdots, x_t$ such that

- $x_0 = x$, $x_t = x'$ (and hence $\mathsf{P}[u|w;x_0,y_u^*] > (1+\delta)\mathsf{P}[u|w;x_t,y_u^*]$);
- for every $i = 0, \cdots, t-1$, there exists $y_i \in \mathcal{Y}$ such that $f(x_i, y_i) = f(x_{i+1}, y_i)$.

Then, there exists an i such that $\mathsf{P}[u|w;x_i,y_u^*] > (1+\delta)^{1/t}\mathsf{P}[u|w;x_{i+1},y_u^*]$ and $\mathsf{P}[u|w;x_i,y_u^*] > \mathsf{P}[u|w;x_0,y_u^*]/(1+\delta)$. We will denote the nodes (x_i, x_{i+1}, y_i) by (x_u, x'_u, y_u). Thus, for every node $u \in F_X^\theta$, there are nodes (x_u, x'_u, y_u, y_u^*) such that

- $f(x_u, y_u) = f(x'_u, y_u)$, and
- $\mathsf{P}[u|w;x_u,y_u^*] > (1+\delta)^{1/t}\mathsf{P}[u|w;x'_u,y_u^*]$ and $\mathsf{P}[u|w;x_u,y_u^*] > \alpha(u,y_u^*)/(1+\delta)$.

Suppose that $\mathsf{P}[\widetilde{F_X}] = p$; i.e., when the protocol is executed with a random input pair (x,y), with probability p, the transcript passes through some $u \in \widetilde{F_X}$. Since there are at most $|\mathcal{X}|^2|\mathcal{Y}|^2$ values for the tuples (x_u, x'_u, y_u, y_u^*), we can find a tuple $(\hat{x}_0, \hat{x}_1, \hat{y}_0, \hat{y}_1)$ such that the transcript passes through $u \in \widetilde{F_X}$ with $(x_u, x'_u, y_u, y_u^*) = (\hat{x}_0, \hat{x}_1, \hat{y}_0, \hat{y}_1)$ with probability at least $p' = p/(|\mathcal{X}|^2|\mathcal{Y}|^2)$. We define $\widehat{F_X} \subseteq \widetilde{F_X}$ as containing those u with $(x_u, x'_u, y_u, y_u^*) = (\hat{x}_0, \hat{x}_1, \hat{y}_0, \hat{y}_1)$. Then $\mathsf{P}[\widehat{F_X}] \geq p'$.

For $w = \mathsf{Apred}(u)$ for $u \in \widehat{F_X}$, w is strictly above F_X^θ, and hence $\mathsf{P}[w|\hat{x}_0,\hat{y}_1] \geq \mathsf{P}[w|\hat{y}_1]/(1+\delta)^{N-1}$. (Since w has a child u, we upper-bound its depth by $N-1$.)

Also, since for $u \in \widehat{F_X}$ we have $\mathsf{P}[u|w;\hat{x}_0,\hat{y}_1] \geq \alpha(u,\hat{y}_1)/(1+\delta) \geq \mathsf{P}[u|w;\hat{y}_1]/(1+\delta)$, we get that

$$\begin{aligned}
\mathsf{P}[\widehat{F_X}|\hat{x}_0,\hat{y}_1] &= \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{F_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_0,\hat{y}_1] \\
&\geq \frac{1}{(1+\delta)^N} \sum_w \mathsf{P}[w|\hat{y}_1] \sum_{\substack{u\in\widehat{F_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{y}_1] \\
&\geq \frac{1}{(1+\delta)^N} \mathsf{P}[\widehat{F_X}|\hat{y}_1].
\end{aligned}$$

Finally, note that for $u \in \widehat{F_X}$, $\mathsf{P}[\hat{y}_1|u] \geq \theta$ and hence

$$\mathsf{P}[\widehat{F_X}|\hat{y}_1] = \sum_{u\in\widehat{F_X}} \mathsf{P}[u|\hat{y}_1] = \sum_{u\in\widehat{F_X}} |\mathcal{Y}|\mathsf{P}[\hat{y}_1|u]\mathsf{P}[u] \geq \theta|\mathcal{Y}|\mathsf{P}[\widehat{F_X}].$$

Hence,

$$\mathsf{P}[\widehat{F_X}|\hat{x}_0,\hat{y}_1] \geq \frac{\theta|\mathcal{Y}|}{(1+\delta)^N}\mathsf{P}[\widehat{F_X}].$$

Part 2. This part is in fact similar to the argument in [MPR09], except that we need to rely on the independence guarantee from Claim 5.6 to say that Alice's message is (almost) independent of Bob's input, conditioned on the (augmented) transcript so far. We shall show that $|\mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_0] - \mathsf{P}[\widehat{S_X}|\hat{x}_1,\hat{y}_0]|$ is significant if $\mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_1]$ is significant. However, since $f(\hat{x}_0,\hat{y}_0) = f(\hat{x}_1,\hat{y}_0)$, the former must be "small", and hence the latter too must be small.

Since $\widehat{S_X}$ is part of a frontier, for all x, y,

$$\mathsf{P}[\widehat{S_X}|x,y] = \sum_{u\in\widehat{S_X}} \mathsf{P}[u|x,y] = \sum_{w\in\mathsf{Anodes}} \mathsf{P}[w|x,y] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;x,y].$$

For $u \in \widehat{S_X}$, $w = \mathsf{Apred}(u)$ is u's parent, an Alice node which is strictly above F_Y^0.

$$\begin{aligned}
\mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_0] &= \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_0,\hat{y}_0] \\
&= \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_0,\hat{y}_1] \pm N\hat{\varepsilon}_0 \qquad \text{(By Eq. 6.)}
\end{aligned}$$

Note that Eq. 6 can be applied above, since the summation is over w strictly above F_Y^0 (since $w = \mathsf{Apred}(u)$ for $u \in \widehat{S_X}$), and for such w, $\mathsf{P}[y|w;x] > \frac{1}{(1+\delta)^N|\mathcal{Y}|} = \frac{1}{\text{poly}(|\mathcal{X}||\mathcal{Y}|)}$.

$$
\begin{aligned}
\mathsf{P}[\widehat{S_X}|\hat{x}_1,\hat{y}_0] &= \sum_w \mathsf{P}[w|\hat{x}_1,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_1,\hat{y}_0] \\
&= \sum_w \mathsf{P}[w|\hat{x}_1,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_1,\hat{y}_1] \pm N\widehat{\varepsilon}_0 && \text{(By Eq. 6.)} \\
&= \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_1,\hat{y}_1] \pm 2N\nu_0 \pm N\widehat{\varepsilon}_0 && \text{(By Eq. 9.)}
\end{aligned}
$$

The above expressions for $\mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_0]$ and $\mathsf{P}[\widehat{S_X}|\hat{x}_1,\hat{y}_0]$, combined with Eq. 5 and Eq. 7 let us relate their difference to $\mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_1]$, as follows.

$$
\begin{aligned}
&\mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_0] - \mathsf{P}[\widehat{S_X}|\hat{x}_1,\hat{y}_0] \\
&\quad\geq \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} (\mathsf{P}[u|w;\hat{x}_0,\hat{y}_1] - \mathsf{P}[u|w;\hat{x}_1,\hat{y}_1]) - 2N(\nu_0+\widehat{\varepsilon}_0) \\
&\quad\geq \left(\frac{\delta'}{1+\delta'}\right) \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_0,\hat{y}_1] - 2N(\nu_0+\widehat{\varepsilon}_0) \\
&\quad\geq \left(\frac{\delta'}{(1+\delta')(1+\delta)^N}\right) \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{S_X}\\ \mathsf{Apred}(u)=w}} \mathsf{P}[u|w;\hat{x}_0,\hat{y}_1] - 2N(\nu_0+\widehat{\varepsilon}_0) \\
&\quad\geq \left(\frac{\delta'}{(1+\delta')(1+\delta)^N}\right) \mathsf{P}[\widehat{S_X}|\hat{x}_0,\hat{y}_1] - 2N(\nu_0+\widehat{\varepsilon}_0)
\end{aligned}
$$

Part 3. We shall consider an attack when the protocol is run with inputs $(\hat{x}_0,\hat{y}_0)$ or $(\hat{x}_1,\hat{y}_0)$ (which must be indistinguishable for security). We shall show that if $\mathsf{P}[\widehat{R_X}|\hat{x}_0,\hat{y}_1]$ is significant, then the curious Bob's output is significantly correlated with Alice's input x (biased more towards 0 when $x=\hat{x}_0$). This will contradict the security of the protocol, since in the ideal world, Bob's input $\hat{y}_0$ cannot distinguish between Alice's input being $\hat{x}_0$ or $\hat{x}_1$.

The probability that the execution with input $(\hat{x}_0,\hat{y}_0)$ reaches a node $w=\mathsf{Apred}(u)$ for $u\in\widehat{R_X}$ is significant if this probability is significant in the execution with input $(\hat{x}_0,\hat{y}_1)$, since each such w falls above the F_Y^0 frontier, and replacing $\hat{y}_1$ with $\hat{y}_0$ causes only a constant factor change in the probabilities. In Figure 4 we describe a *curious Bob* who can, at such a point, mentally substitute its input $\hat{y}_0$ with $\hat{y}_1$ and simulate the augmented execution (inluding – and this is the non-trivial part – the answers from the oracle) till before the next Alice message. The probability that this simulated execution goes through $\widehat{R_X}$ remains significant when Alice's input is $\hat{x}_0$ (since the simulated execution will have input $(\hat{x}_0,\hat{y}_1)$). At the same time, the probability of the execution with $(\hat{x}_1,\hat{y}_1)$ hitting each node in $\widehat{R_X}$ differs by a significant *factor* from that when Alice's input is $\hat{x}_0$ (Eq. 5). This will let the curious Bob distinguish between when Alice's input is $\hat{x}_0$ and when it is $\hat{x}_1$, even though Bob's real input is $\hat{y}_0$, leading to a contradiction.

Curious Bob: Learning what Eve learns, with a different input

Bob is given $\hat{y}_0$ as input, and Alice is given a uniformly random element from $x \leftarrow \{\hat{x}_0, \hat{x}_1\}$ as input. Alice and Bob execute the protocol honestly, with access to a random oracle $\mathcal{O}$. But at the end Bob carries out the following computation.

For every Alice node w in the augmented transcript, which is strictly above F_Y^0, Bob carries out an *exploration* as follows. He samples a view $V_{B,\hat{y}_1}(w)$ for himself with input $\hat{y}_1$, conditioned on node w (and in particular Eve's view $V_E(w)$). Bob mentally carries out the execution with the hypothetical view $V_{B,\hat{y}_1}(w)$, till the next message from Alice (i.e., Eve queries, followed by Bob's own queries and his message in the protocol, and then further Eve queries) by simulating an oracle $\mathcal{O}'$ defined as follows. Below, $V_{B,\hat{y}_0}(w)$ denotes the actual view of Bob in the protocol at that point, $\mathcal{O}$ is the actual oracle and $\mathcal{O}''$ is a freshly sampled independent random oracle. On query q,

- if $q \in \mathcal{Q}(V_{B,\hat{y}_1}(w)) \cup \mathcal{Q}(V_E(w))$, answer according to $V_{B,\hat{y}_1}(w)$ or $V_E(w)$;[a]
- else, if $q \in \mathcal{Q}(V_{B,\hat{y}_0}(w))$, answer according to $\mathcal{O}''$;
- else, answer according to $\mathcal{O}$.

Let the set of nodes encountered by Bob during this exploration (over explorations from every Alice node w) be $\mathcal{E}^x_{\hat{y}_0,\hat{y}_1}$, where x is Alice's input, and Bob substitutes $\hat{y}_0$ with $\hat{y}_1$ for exploration. If $\mathcal{E}^x_{\hat{y}_0,\hat{y}_1} \cap \widehat{R_X} \neq \emptyset$, then Bob outputs 0; else he outputs 1.

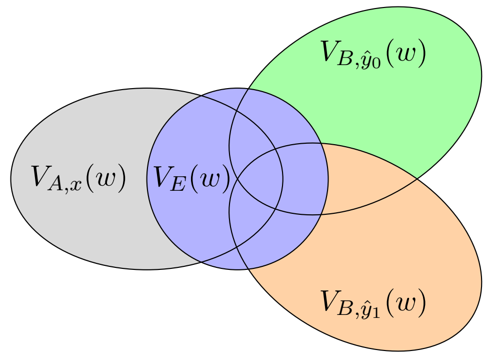

Figure 3: Simulating the oracle answers during exploration. The ovals represent the sets of queries in the views $V_{A,x}(w)$, $V_E(w)$, $V_{B,\hat{y}_0}(w)$ and $V_{B,\hat{y}_1}(w)$. Queries already answered in $V_E(w)$ (blue) or in the hypothetical Bob view $V_{B,\hat{y}_1}(w)$ (orange) are answered according to these views. Answers to the remaining queries in $\mathcal{Q}(V_{B,\hat{y}_0}(w))$ (green), are freshly sampled, i.e. answered according to $\mathcal{O}''$. All other queries are answered using the actual random oracle $\mathcal{O}$. When the "safety" condition Eq. 10 holds, i.e., the orange and green regions (which have "edited" answers) do not intersect the gray region, this yields a perfect simulation (see Eq. 12).

[a] As $V_{B,\hat{y}_1}(w)$ is conditioned on $V_E(w)$, if $q \in \mathcal{Q}(V_{B,\hat{y}_1}(w)) \cap \mathcal{Q}(V_E(w))$, both views will have the same answer for q.

Figure 4: Curious Bob strategy to show that $\mathsf{P}[\widehat{R_X} | \hat{x}_0, \hat{y}_1]$ is small.

Before we prove this, we define a game and state a sub-claim, which will help us with the analysis (for which we will derive yet another game based on this).

Game $\mathsf{G}^{y'}(x, y)$. An oracle $\mathcal{O}$, and random tapes for Alice and Bob are picked at random. Then, Alice and Bob execute the protocol Π using oracle $\mathcal{O}$ and the chosen random tapes, with inputs x and y respectively; at each node $w \in \mathsf{Achildren}$ in the transcript path, we define $V_{A,x}(w)$, $V_{B,y}(w)$ and $V_E(w)$ as the views of Alice, Bob and Eve respectively. Further, at each such node w that is strictly above the frontier F_Y^0 (i.e., $\nexists z \in F_Y^0$, $z \preceq w$.), we pick a random view for Bob conditioned on w and input y'. Let $V_{B,y'}(w)$ represent that Bob view.

We define the event $\mathsf{safe}(w)$ to occur in this game for a node w where $V_{B,y'}(w)$ is sampled (i.e., child of an Alice node that is strictly above F_Y^0), if

$$\mathcal{Q}\left(V_{A,x}(w)\right) \ \cap\ \left(\mathcal{Q}(V_{B,y}(w)) \ \cup\ \mathcal{Q}(V_{B,y'}(w))\right) \ \subseteq\ \mathcal{Q}(V_E(w)) \tag{10}$$

Claim 5.7. *For any $x \in \mathcal{X}$, $y, y' \in \mathcal{Y}$, in the game $\mathsf{G}^{y'}(x, y)$, for an error parameter $\widehat{\varepsilon}_1 = \varepsilon^{\Omega(1)} \operatorname{poly}(\kappa|\mathcal{X}||\mathcal{Y}|)$, we have*

$$\sum_w \mathsf{P}[w|x, y']\mathsf{P}'_x[\overline{\mathsf{safe}(w)}|w] \leq N\widehat{\varepsilon}_1 \tag{11}$$

(where the summation is over w for which $\mathsf{safe}(w)$ is defined: i.e., $w \in \mathsf{Achildren}$ such that $w \prec F_Y^0$).

Proof. This follows from Lemma A.3. At every level L_i, Lemma A.3 guarantees that $\sum_{\substack{w \in L_i \\ \mathsf{P}[y|w;x] \geq \varepsilon'}} \mathsf{P}[w \wedge \overline{\mathsf{safe}(W)}] \leq \varepsilon'$ for $\varepsilon' = \varepsilon^{\Omega(1)} \operatorname{poly}(\kappa|\mathcal{X}||\mathcal{Y}|)$. For $w \prec F_Y^0$, we have $\mathsf{P}[y|w; x] \geq \frac{1}{(1+\delta)^N|\mathcal{Y}|} > \varepsilon'$. Hence the sum in the claim is bounded by $N\varepsilon'$. We set $\widehat{\varepsilon}_1 = \varepsilon'$. □

Two Experiments: $\widehat{\mathsf{G}}(x)$ and $\mathsf{G}'(x)$. Now, we define two experiments $\widehat{\mathsf{G}}(x)$ and $\mathsf{G}'(x)$ as follows:

$\widehat{\mathsf{G}}(x)$ (which corresponds to the curious attack above) is the same as $\mathsf{G}^{\hat{y}_1}(x, \hat{y}_0)$, but with the following addition. At each node $w \in \mathsf{Achildren}$ above F_Y^0 in the transcript path, we carry out an "exploration" of Eve's steps and Bob's step till the next message from Alice (Eve, Bob, Eve), using the view sampled for $\hat{y}_1$. This exploration is carried out as defined above for the curious Bob strategy (Figure 4). For simplifying notation, we make the following definition. For a node u and $w = \mathsf{Apred}(u)$, we define the probability of the *exploration* starting at w visiting u as

$$\widehat{\mathsf{P}}_x[u|w] = \Pr_{\widehat{\mathsf{G}}(x)}[u \text{ reached in exploration from } w | w \text{ reached in execution}].$$

We also define $\widehat{\mathsf{P}}_x[u] = \mathsf{P}[w|x, \hat{y}_0]\widehat{\mathsf{P}}_x[u|w]$ to be the probability of the exploration reaching $u \in \widehat{R_X}$ (not conditioned on visiting $w = \mathsf{Apred}(u)$).

Note that in the right-hand side of the equation, we have $\mathsf{P}[w|x, \hat{y}_0]$, i.e. the node w is generated with Alice interacting with her input x and Bob with his input $\hat{y}_0$. After reaching w, Bob samples a new view conditioned on his input being $\hat{y}_1$ and proceed to explore till Alice is supposed to send the next message. This part of the probability, i.e. probability of reaching a node u *conditioned* on reaching w is expressed by the term $\widehat{\mathsf{P}}_x[u|w]$. We point out that $\widehat{\mathsf{P}}_x[u|w]$ is *not necessarily equal to* $\mathsf{P}[u|w, x, \hat{y}_1]$ since the exploration uses a simulated oracle that is simulated without knowing x. (However, as we shall see, it will be closely related to the latter.)

$\mathsf{G}'(x)$ is in fact, the same as $\mathsf{G}^{\hat{y}_0}(x, \hat{y}_1)$ (note the reversal of roles for $\hat{y}_0$ and $\hat{y}_1$): i.e., an execution with inputs $(x, \hat{y}_1)$, along with sampling Bob's view for input $\hat{y}_0$ at each node $w \in \mathsf{Achildren}$ encountered that is strictly above F_Y^0. This experiment involves no exploration. Now, for a node u and $w = \mathsf{Apred}(u)$, we define the probability of the execution visiting u, conditioned on it having visited w, as

$$\mathsf{P}'_x[u|w] = \Pr_{\mathsf{G}'(x)}[u \text{ reached in the execution } | w \text{ reached in execution}].$$

We also define $\mathsf{P}'_x[u] = \mathsf{P}[w|x, \hat{y}_1]\mathsf{P}'_x[u|w]$ to be the probability of the execution reaching u (not conditioned on visiting $w = \mathsf{Apred}(u)$). Note that in this experiment the only significance of $\hat{y}_0$ is in defining the event $\mathsf{safe}(w)$. In particular, there is no exploration phase or switching of inputs, and the execution considered for defining the probability $\mathsf{P}'_x[u|w]$ is simply the same as a faithful execution of the original augmented protocol. Thus, $\mathsf{P}'_x[u|w] = \mathsf{P}[u|w, x, \hat{y}_1]$.

A priori, there is no direct relation between the probability terms $\widehat{\mathsf{P}}_x[u|w]$ and $\mathsf{P}'_x[u|w]$. This is because the sampling of the Bob view in $\widehat{\mathsf{G}}(x)$ is not correlated with the view of Alice given Eve's view; while, on the other hand, Bob's view in $\mathsf{G}'(x)$ could possibly be correlated with Alice's view even when Eve view is given. But, by additionally conditioning on the event $\mathsf{safe}(w)$, these two probabilities are identical. More formally, we have the following key observation:[13] For all x, for all u, w such that $w = \mathsf{Apred}(u)$, $\widehat{\mathsf{P}}_x[w, \mathsf{safe}(w)] > 0$ and $\mathsf{P}'_x[w, \mathsf{safe}(w)] > 0$,

$$\widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)] = \mathsf{P}'_x[u|w, \mathsf{safe}(w)]. \tag{12}$$

This is because, given a node w, in either experiment, the set of Alice views, the set of Bob views with input $\hat{y}_0$ and the set of Bob views with input $\hat{y}_1$ each compatible with the view in w (individually) are determined. On conditioning on $\mathsf{safe}(w)$, the distribution over triplets of views (one from each of the three sets) is the same in both experiments: they correspond to pairs of edges in the "views graph" at w, with both edges incident on the same Alice view, and the probability of a pair is (before conditioning) product of the probabilities on the two edges (according to distributions obtained by conditioning on $\hat{y}_0$ and $\hat{y}_1$), and the conditioning removes all those pairs of edges that violate the safety condition; these operations (multiplication and safety condition) are symmetric in $\hat{y}_0$,$\hat{y}_1$ and hence, both the distributions are the same. Now, conditioned on $\mathsf{safe}(w)$, the exploration in $\widehat{\mathsf{G}}(x)$ for a triplet of views is identical to the execution in $\mathsf{G}'(x)$ for the same triplet.

Assuming that $\mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1]$ is significant, we are interested in lower-bounding $\widehat{\mathsf{P}}_{\hat{x}_0}[\widehat{R_X}] - \widehat{\mathsf{P}}_{\hat{x}_1}[\widehat{R_X}]$.

For $x \in \{\hat{x}_0, \hat{x}_1\}$, we have:

$$\widehat{\mathsf{P}}_x[\widehat{R_X}] = \sum_{u \in \widehat{R_X}} \widehat{\mathsf{P}}_x[u] = \sum_w \left(\mathsf{P}[w|x, \hat{y}_0] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u) = w}} \widehat{\mathsf{P}}_x[u|w] \right)$$

Note that the last summation will be over $w \in \mathsf{Achildren}$ that are strictly above F_Y^0, since we consider only those w for which there exists some $u \in \widehat{R_X}$ with $\mathsf{Apred}(u) = w$.

[13]We shall use this claim for w strictly above F_Y^0. It can be seen that if only one of $\widehat{\mathsf{P}}_x[w, \mathsf{safe}(w)]$ and $\mathsf{P}'_x[w, \mathsf{safe}(w)]$ is positive, then by the convention in Footnote 10, the node w cannot be strictly above F_Y^0. Hence the claim will be applicable. Alternately, similar to the normal form for protocols mentioned in Footnote 10, we can assume w.l.o.g that for all $w \in \mathsf{Achildren}$, $\widehat{\mathsf{P}}_x[w, \mathsf{safe}(w)] > 0$ and $\mathsf{P}'_x[w, \mathsf{safe}(w)] > 0$, so that the claim holds for all w.

Fix a node w and consider $u \in \widehat{R_X}$ such that $\mathsf{Apred}(u) = w$. Then, (using the convention in Footnote 10),

$$\begin{aligned}
\widehat{\mathsf{P}}_x[u|w] &= \widehat{\mathsf{P}}_x[u, \mathsf{safe}(w)|w] + \widehat{\mathsf{P}}_x[u, \overline{\mathsf{safe}(w)}|w] \\
&= \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)]\widehat{\mathsf{P}}_x[\mathsf{safe}(w)|w] + \widehat{\mathsf{P}}_x[u, \overline{\mathsf{safe}(w)}|w] \\
&= \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)] - \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)]\widehat{\mathsf{P}}_x[\overline{\mathsf{safe}(w)}|w] + \widehat{\mathsf{P}}_x[u, \overline{\mathsf{safe}(w)}|w] \\
&= \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)] + (\widehat{\mathsf{P}}_x[u|w, \overline{\mathsf{safe}(w)}] - \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)])\widehat{\mathsf{P}}_x[\overline{\mathsf{safe}(w)}|w]
\end{aligned}$$

The sums $\sum_{u \in \widehat{R_X}, \mathsf{Apred}(u)=w} \widehat{\mathsf{P}}_x[u|w, \overline{\mathsf{safe}(w)}]\widehat{\mathsf{P}}_x[\overline{\mathsf{safe}(w)}|w]$ and $\sum_{u \in \widehat{R_X}, \mathsf{Apred}(u)=w} \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)]\widehat{\mathsf{P}}_x[\overline{\mathsf{safe}(w)}|w]$ are both bounded by $\widehat{\mathsf{P}}_x[\overline{\mathsf{safe}(w)}|w]$. Thus we can write

$$\begin{aligned}
\widehat{\mathsf{P}}_x[\widehat{R_X}] &= \sum_w \left(\mathsf{P}[w|x, \hat{y}_0] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \widehat{\mathsf{P}}_x[u|w] \right) \\
&= \sum_w \left(\mathsf{P}[w|x, \hat{y}_0] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)] \right) \pm \sum_w \mathsf{P}[w|x, \hat{y}_0]\widehat{\mathsf{P}}_x[\overline{\mathsf{safe}(w)}|w] \\
&= \sum_w \left(\mathsf{P}[w|x, \hat{y}_0] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \widehat{\mathsf{P}}_x[u|w, \mathsf{safe}(w)] \right) \pm N\widehat{\varepsilon}_1 \qquad (13)
\end{aligned}$$

where the last step follows by Claim 5.7. Note that $\mathsf{P}'_x[\widehat{R_X}] = \mathsf{P}[\widehat{R_X}|x, \hat{y}_1]$.

In our derivation below, we shall rely on conditioning the experiments $\widehat{\mathsf{G}}(x)$ and $\mathsf{G}'(x)$ on the event $\mathsf{safe}(\cdot)$. To facilitate our arguments we relate certain probabilities when conditioned on $\mathsf{safe}(\cdot)$ and otherwise.

Claim 5.8. *The following two inequalities hold:*

$$\sum_w \mathsf{P}[w|\hat{x}_0, \hat{y}_1] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_0}[u|w, \mathsf{safe}(w)] = \left(\sum_w \mathsf{P}[w|\hat{x}_0, \hat{y}_1] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_0}[u|w] \right) \pm N\widehat{\varepsilon}_1 \qquad (14)$$

$$\sum_w \mathsf{P}[w|\hat{x}_0, \hat{y}_1] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_1}[u|w, \mathsf{safe}(w)] = \left(\sum_w \mathsf{P}[w|\hat{x}_0, \hat{y}_1] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_1}[u|w] \right) \pm (1+\delta)^N (N\widehat{\varepsilon}_1 + 2N\nu_0). \qquad (15)$$

Proof. Firstly,

$$\sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_x[u|w, \mathsf{safe}(w)] = \left(\sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_x[u|w] \right) \pm \mathsf{P}'_x[\overline{\mathsf{safe}(w)}|w] \qquad (16)$$

We get Eq. 14 as follows:

$$\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_0}[u|w,\mathsf{safe}(w)]$$

$$= \left(\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_0}[u|w]\right) \pm \left(\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1]\mathsf{P}'_{\hat{x}_0}[\overline{\mathsf{safe}(w)}|w]\right) \quad \text{By Eq. 16.}$$

$$= \left(\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_0}[u|w]\right) \pm N\widehat{\varepsilon}_1 \quad \text{By Eq. 11.}$$

To prove Eq. 15, first we note the following:

$$\begin{aligned}\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0]\mathsf{P}'_{\hat{x}_1}[\overline{\mathsf{safe}(w)}|w] &\le \sum_w \mathsf{P}[w|\hat{x}_1,\hat{y}_0]\mathsf{P}'_{\hat{x}_1}[\overline{\mathsf{safe}(w)}|w] + 2N\nu_0 && \text{By Eq. 9.}\\ &\le N\widehat{\varepsilon}_1 + 2N\nu_0 && \text{By Eq. 11.}\end{aligned} \tag{17}$$

$$\begin{aligned}\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1]\mathsf{P}'_{\hat{x}_1}[\overline{\mathsf{safe}(w)}|w] &\le (1+\delta)^N \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_0]\mathsf{P}'_{\hat{x}_1}[\overline{\mathsf{safe}(w)}|w]\\ &\le (1+\delta)^N(N\widehat{\varepsilon}_1 + 2N\nu_0) && \text{By Eq. 17.}\end{aligned} \tag{18}$$

Hence, we conclude

$$\sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_1}[u|w,\mathsf{safe}(w)]$$

$$= \sum_w \left(\mathsf{P}[w|\hat{x}_0,\hat{y}_1] \sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_1}[u|w]\right) \pm \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1]\mathsf{P}'_{\hat{x}_1}[\overline{\mathsf{safe}(w)}|w] \quad \text{By Eq. 16.}$$

$$= \sum_w \mathsf{P}[w|\hat{x}_0,\hat{y}_1]\left(\sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_1}[u|w]\right) \pm (1+\delta)^N(N\widehat{\varepsilon}_1 + 2N\nu_0) \quad \text{By Eq. 18.}\square$$

To lower bound $\widehat{\mathsf{P}}_{\hat{x}_0}[\widehat{R_X}] - \widehat{\mathsf{P}}_{\hat{x}_1}[\widehat{R_X}]$ we proceed as follows:

$$\widehat{\mathsf{P}}_{\hat{x}_0}[\widehat{R_X}] - \widehat{\mathsf{P}}_{\hat{x}_1}[\widehat{R_X}]$$

$$\ge \sum_w \Bigg[\mathsf{P}[w|\hat{x}_0,\hat{y}_0]\left(\sum_{\substack{u\in\widehat{R_X},\\ \mathsf{Apred}(u)=w}} \widehat{\mathsf{P}}_{\hat{x}_0}[u|w,\mathsf{safe}(w)]\right)$$

$$- \mathsf{P}[w|\hat{x}_1, \hat{y}_0]\Bigg(\sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \widehat{\mathsf{P}}_{\hat{x}_1}[u|w, \mathsf{safe}(w)] \Bigg) \Bigg] - 2N\widehat{\varepsilon}_1 \qquad \text{By Eq. 13.}$$

$$= \sum_w \Bigg[\mathsf{P}[w|\hat{x}_0, \hat{y}_0]\Bigg(\sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_0}[u|w, \mathsf{safe}(w)] \Bigg)$$

$$- \mathsf{P}[w|\hat{x}_1, \hat{y}_0]\Bigg(\sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \mathsf{P}'_{\hat{x}_1}[u|w, \mathsf{safe}(w)] \Bigg) \Bigg] - 2N\widehat{\varepsilon}_1 \qquad \text{By Eq. 12 (and Footnote 13).}$$

$$\geq \sum_w \Bigg(\mathsf{P}[w|\hat{x}_0, \hat{y}_0] \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \Bigg(\mathsf{P}'_{\hat{x}_0}[u|w, \mathsf{safe}(w)] - \mathsf{P}'_{\hat{x}_1}[u|w, \mathsf{safe}(w)] \Bigg) \Bigg)$$

$$- 2N(\nu_0 + \widehat{\varepsilon}_1) \qquad \text{By Eq. 9.}$$

$$\geq \sum_w \Bigg(\frac{\mathsf{P}[w|\hat{x}_0, \hat{y}_1]}{(1+\delta)^N} \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \Bigg(\mathsf{P}'_{\hat{x}_0}[u|w, \mathsf{safe}(w)] - \mathsf{P}'_{\hat{x}_1}[u|w, \mathsf{safe}(w)] \Bigg) \Bigg)$$

$$- 2N(\nu_0 + \widehat{\varepsilon}_1) \qquad \text{By Eq. 7.}$$

$$\geq \sum_w \Bigg(\frac{\mathsf{P}[w|\hat{x}_0, \hat{y}_1]}{(1+\delta)^N} \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \Bigg(\mathsf{P}'_{\hat{x}_0}[u|w] - \mathsf{P}'_{\hat{x}_1}[u|w] \Bigg) \Bigg)$$

$$- 2N(\nu_0 + \widehat{\varepsilon}_1) - \frac{N\widehat{\varepsilon}_1}{(1+\delta)^N} - (2N\nu_0 + N\widehat{\varepsilon}_1) \qquad \text{By Eq. 14 and Eq. 15.}$$

$$= \sum_w \Bigg(\frac{\mathsf{P}[w|\hat{x}_0, \hat{y}_1]}{(1+\delta)^N} \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \Bigg(\mathsf{P}[u|w; \hat{x}_0, \hat{y}_1] - \mathsf{P}[u|w; \hat{x}_1, \hat{y}_1] \Bigg) \Bigg)$$

$$- 2N(\nu_0 + \widehat{\varepsilon}_1) - \frac{N\widehat{\varepsilon}_1}{(1+\delta)^N} - (2N\nu_0 + N\widehat{\varepsilon}_1) \qquad \text{By definition of } \mathsf{P}'_x[u|w].$$

$$\geq \sum_w \Bigg(\frac{\mathsf{P}[w|\hat{x}_0, \hat{y}_1]}{(1+\delta)^N} \sum_{\substack{u \in \widehat{R_X}, \\ \mathsf{Apred}(u)=w}} \Big(\frac{\delta'}{1+\delta'} \Big) \mathsf{P}[u|w; \hat{x}_0, \hat{y}_1] \Bigg) - 4N(\nu_0 + \widehat{\varepsilon}_1) \qquad \text{By Eq. 5.}$$

$$\geq \frac{\delta'}{(1+\delta')(1+\delta)^N} \mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1] - 4N(\nu_0 + \widehat{\varepsilon}_1)$$

Putting things Together Let us define $\mu = (1+\delta)^N$ and recall that $\delta' = (1+\delta)^{1/(|\mathcal{X}|-1)} - 1$. From Part 2 and 3, we obtain a lower-bound on the distinguishing advantage obtained in terms of $\mathsf{P}[\widehat{S_X}|\hat{x}_0, \hat{y}_1]$ and $\mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1]$. We can assume that this advantages are ρ_S and ρ_R respectively. But we know that simulation error is ν_0, so $\rho_S + \rho_R \leq 2\nu_0$. Thus, we obtain the following bounds:

$$\mathsf{P}[\widehat{S_X}|\hat{x}_0, \hat{y}_1] \leq \frac{(1+\delta')\mu}{\delta'} \left(\rho_S + 2N(\nu_0 + \widehat{\varepsilon}_0)\right)$$
$$\mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1] \leq \frac{(1+\delta')\mu}{\delta'} \left(\rho_R + 4N(\nu_0 + \widehat{\varepsilon}_1)\right)$$

Finally, we can obtain a bound on the overall bad event $\mathsf{P}[\widetilde{F_X}]$:

$$\begin{aligned}
\mathsf{P}[\widetilde{F_X}] &\leq |\mathcal{X}|^2|\mathcal{Y}|^2\mathsf{P}[\widehat{F_X}] \\
&\leq \frac{\mu|\mathcal{X}|^2|\mathcal{Y}|}{\theta}\mathsf{P}[\widehat{F_X}|\hat{x}_0, \hat{y}_1] \\
&= \frac{\mu|\mathcal{X}|^2|\mathcal{Y}|}{\theta}\left(\mathsf{P}[\widehat{S_X}|\hat{x}_0, \hat{y}_1] + \mathsf{P}[\widehat{R_X}|\hat{x}_0, \hat{y}_1]\right) \\
&\leq \frac{2\mu^2|\mathcal{X}|^2|\mathcal{Y}|(1+\delta')N}{\theta\delta'}\left(4\nu_0 + \widehat{\varepsilon}_0 + 2\widehat{\varepsilon}_1\right)
\end{aligned}$$

This completes the proof of Claim 5.5, and in turn that of Claim 5.1. As discussed in Section 5.2, this (combined with Claim 5.3), is used to prove Theorem 1.1.

6 Beyond Semi-Honest Security

In this section we prove Theorem 1.2, which tells us that in the context of building 2-party SFE protocols secure against active adversaries, a random oracle is only useful as a means for securely realizing the commitment functionality, denoted by $\mathcal{F}_{\text{COM}}$. This holds true for both UC and standalone security.

Theorem 1.2 (Restated.) *For a deterministic finite 2-party function f, the following statements are equivalent:*

1. *f has a statistically UC-secure SFE protocol in the random oracle model.*
2. *f has a statistically standalone-secure SFE protocol in the random oracle model.*
3. *f has a statistically UC-secure SFE protocol in the $\mathcal{F}_{\text{COM}}$-hybrid model.*
4. *f has a statistically standalone-secure SFE protocol in the $\mathcal{F}_{\text{COM}}$-hybrid model.*

Proof. Clearly, $(1) \Rightarrow (2)$ and $(3) \Rightarrow (4)$. That $(3) \Rightarrow (1)$ and $(4) \Rightarrow (2)$ follow from the fact that in the random oracle model, we can UC-securely implement the $\mathcal{F}_{\text{COM}}$ functionality. (This implication holds not only for deterministic SFE, but also for reactive or randomized functionalities as well.)

To complete the proof we shall show that $(2) \Rightarrow (3)$. So suppose f has a standalone secure protocol using a random oracle. Let f' be a redundancy free function obtained by removing redundant inputs one by one from f (see Footnote 4). Then, it is enough to show $(2') \Rightarrow (3')$ where $(2')$ and $(3')$ are identical to (2) and (3), but with f replaced by f' (because, $(2) \Leftrightarrow (2')$ and $(3) \Leftrightarrow (3')$ [MPR10, KM11]). Now, if f' has a standalone secure protocol (possibly in the random oracle model), then the same protocol is semi-honest secure as well. Then, by Claim 5.2, f' must be isomorphic to a symmetric function f'' which has a semi-honest secure protocol in the plain model. That is f'' must be decomposable. Then, by a result in [MPR09], f'' has a UC secure protocol in the $\mathcal{F}_{\text{COM}}$-hybrid model. Since f'' is isomorphic to f', the latter also has UC secure protocol in the $\mathcal{F}_{\text{COM}}$-hybrid model, proving $(3')$ as desired. □

7 Black-Box Separations

The random oracle model is of interest not only as an abstract theoretical framework, but also because it models a (strong) one-way function. Thus, informally, the impossibility results in the random oracle model translate to impossibility of constructions that rely on a one-way function as its sole computational primitive. This intuition can be formalized as black-box separation results, following [IR89, RTV04].

For our black-box separation results, we shall follow the definitions as introduced by [RTV04] with minor modifications. Following [RTV04], we consider primitives to be specified as pairs of the form $(F_{\mathcal{Q}}, R_{\mathcal{Q}})$. The set $F_{\mathcal{Q}}$ is a set of functions that are candidate implementations of primitive $\mathcal{Q}$. For example, for the one-way function primitive (represented by OWF) the set F_{OWF} consists of all functions defined over $\{0,1\}^*$. The set $R_{\mathcal{Q}}$ is a set of pairs (Q, M), where Q is a candidate implementation of $\mathcal{Q}$ and M is an adversary which breaks the security of Q. (Sometimes we shall abuse the notation and write $(\Pi, M) \in R_{\mathcal{Q}}$ if Π implements a function Q such that $(Q, M) \in R_{\mathcal{Q}}$.) Continuing our example of OWF, (Q, M) would consist of one-way functions Q where the inverter M inverts non-negligible fraction of outputs of Q.

Next, we recall the definition of fully black-box reductions (or as presented below, fully black-box constructions) as introduced in [RTV04]. Below, we say that a (possibly non-uniform) algorithm is *efficient* if it is probabilistic polynomial time (PPT).

Definition 7.1 (Fully Black-box Constructions)**.** A *fully black-box construction* of a primitive $\mathcal{P}$ from another primitive $\mathcal{Q}$ consists of a pair of efficient oracle algorithms (Π, S), such that the following two conditions hold:

1. **Correct Implementation:** For any $Q \in F_{\mathcal{Q}}$, Π^Q implements a function $P \in F_{\mathcal{P}}$.
2. **Security:** For any $Q \in F_{\mathcal{Q}}$ and any (possibly inefficient) adversary A that breaks the security of Π^Q, the reduction $S^{Q,A}$ breaks the security of Q as an implementation of $\mathcal{Q}$. That is, $\forall A, \forall Q \in F_{\mathcal{Q}}$, $(\Pi^Q, A) \in R_{\mathcal{P}} \Rightarrow (Q, S^{Q,A}) \in R_{\mathcal{Q}}$.

We emphasize that the construction Π and the reduction S are efficient.

Constructions that Preserve the Security Parameter. As is standard in cryptographic constructions, we shall associate a security parameter with primitives and state security condition in terms of it. Formally, we shall consider that any primitive $\mathcal{P}$, the input to any $P \in F_{\mathcal{P}}$ has a security parameter encoded as part of its inputs. We prove our separation results with a technical restriction on blackbox constructions, namely that the constructions respect the security parameter: that is, in a black-box construction of $\mathcal{P}$ from $\mathcal{Q}$, when the implementation Π^Q, for $Q \in F_{\mathcal{Q}}$, is given an input with security parameter κ, it always invokes Q with the same security parameter κ. However, there is no such restriction on the security reduction S.

For $Q \in F_{\mathcal{Q}}$, we denote by Q_i the restriction of Q to inputs which have security parameter i. We will often identify Q with the infinite tuple $(Q_1, Q_2, \dots)$. For a security parameter respecting construction (Π, S), when invoked with security parameter κ, Π^Q will access only Q_κ. There is no such restriction of the security reduction S. When invoked with security parameter κ, $S^{Q,A}$ (for an adversary A attacking Π) is expected to invert points in the range of Q_κ. To perform this inversion, $S^{Q,A}$ is permitted access $Q_{\kappa'}$, for all values of $\kappa' \in \mathbb{N}$ (including $\kappa' \neq \kappa$), and in particular can invoke Π^Q and A^Q with different security parameters κ'.

This restriction is not as limiting as it may appear at first, since we can define primitives like one-way function to allow access to a range of input lengths for a single value of the security parameter. (See the definition of OWF_ζ below.)

Below we define the various primitives used to formalize our results. The primitives are formally specified by the F and R sets as mentioned above. We shall specify the functions in F separately for each value of the security parameter. We consider the machines M in all the definitions below as non-uniform machines (with non-uniform advice for each security parameter); however, one could relax the security definition of any of the primitives to consider only uniform M, and by requiring the fully black-box construction to also be uniform, our results hold unchanged.

One-Way Function Primitive OWF. First, for simplicity, we consider a one-way function primitive OWF which considers the security parameter as the input length itself.[14]

- F_{OWF} consists of all functions from $\{0,1\}^*$ to $\{0,1\}^*$, and the security parameter is the length of the input.
- $(Q, M) \in R_{\mathsf{OWF}}$ if there is a non-negligible function δ such that for infinitely many $\kappa \in \mathbb{N}$, $\Pr[Q(M(y)) = y : x \xleftarrow{\$} \{0,1\}^\kappa, y = Q(x)] > \delta(\kappa)$.

Primitive for Semi-honest Secure SFE Protocol. For a 2-party function $f\colon \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}_A \times \mathcal{Z}_B$, we define the primitive SFE_f corresponding to a semi-honest secure protocol for evaluating f. For simplicity, we consider the domain and range of f itself to be finite and fixed (independent of the security parameter).[15] A protocol Π will be identified with the next message function of the protocol. One of its inputs is the security parameter κ.

- $\Pi \in F_{\mathsf{SFE}_f}$ if the protocol defined by Π is "correct", i.e. for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the pair of outputs from Alice and Bob when they execute Π with security parameter κ and inputs (x, y), is $(a, b) = f(x, y)$ except with probability negligible in κ.
- An adversary Adv breaks Π, i.e. $(\Pi, \mathsf{Adv}) \in R_{\mathsf{SFE}_f}$ if there exists (x, x', y, y') such that
  1. $f(x, y) = f(x, y')$ and $|\Pr[\mathsf{Adv}(V_B^{\Pi(x,y)}) = 1] - \Pr[\mathsf{Adv}(V_B^{\Pi(x,y')}) = 1]| > \delta(\kappa)$, or
  2. $f(x, y) = f(x', y)$ and $|\Pr[\mathsf{Adv}(V_A^{\Pi(x,y)}) = 1] - \Pr[\mathsf{Adv}(V_A^{\Pi(x',y)}) = 1]| > \delta(\kappa)$,

  where $V_A^{\Pi(x,y)}$ and $V_B^{\Pi(x,y)}$ stand for Alice's and Bob's views after executing Π with inputs (x, y) and the advantage $\delta(\kappa)$ is non-negligible in κ.

Note that we used a game-based definition of semi-honest security. This is in general weaker than the standard simulation based definition of semi-honest security (unless simulation with unbounded computational power is considered, in which case they are identical). Since we are ruling out blackbox constructions of SFE_f, using a weaker definition of security for SFE_f makes our result only stronger.

One-Way Function Primitive OWF_ζ. Since we consider only security-parameter preserving constructions, a construction using the primitive OWF above can access the one-way function on inputs of length exactly equal to the security-parameter. This limits the implications of a separation result, as it leaves open the possibility that a construction that uses a one-way function on more than one input length could be secure. To rule out this possibility as well, we consider a more elaborate primitive and rule out fully black-box

[14] This is the same one-way function primitive as considered in [RTV04]. However, in the case of security parameter preserving constructions, this primitive prevents the construction from using the one-way function with any other input length other than the security parameter. Later we remove this restriction by considering the primitives OWF_ζ defined below.

[15] One could consider f to have infinite domains and range, and define restrictions of f, $f_\kappa : \mathcal{X}_\kappa \times \mathcal{Y}_\kappa \to \mathcal{Z}_{A,\kappa} \times \mathcal{Z}_{B,\kappa}$, where $\mathcal{X}_1 \subseteq \mathcal{X}_2 \subseteq \cdots \mathcal{X}$ etc., with efficient representations for the subdomains and subrange. Our results hold as long as $|\mathcal{X}_\kappa|, |\mathcal{Y}_\kappa| \leq \mathrm{poly}(\kappa)$. We omit such a formalization for the sake of simplicity.

construction of SFE_f from this primitive as well. Formally, we define a primitive OWF_ζ for each polynomial ζ as follows.

For any function $g : \{0,1\}^* \to \{0,1\}^*$, let g^ζ be defined as follows: $g^\zeta(\kappa, x) = g(x)$ if $|x| \leq \zeta(\kappa)$ and $g^\zeta(\kappa, x) = 0$ otherwise. Let $\mathcal{W}_i^\zeta = \{g^\zeta(i, \cdot) | g : \{0,1\}^* \to \{0,1\}^*\}$.

- $F_{\mathsf{OWF}_\zeta} = \mathcal{W}_1^\zeta \times \mathcal{W}_2^\zeta \times \cdots$. That is, for $Q \in F_{\mathsf{OWF}_\zeta}$, $Q = (Q_1, Q_2, \dots)$, the function Q_κ is of the form $g^\zeta(\kappa, \cdot)$ for some function g.
- $(Q, M) \in R_{\mathsf{OWF}_\zeta}$ if there is a non-negligible function δ such that for infinitely many $\kappa \in \mathbb{N}$, $\Pr[Q(\kappa, M(\kappa, y)) = y : x \xleftarrow{\$} \{0,1\}^1 \cup \cdots \cup \{0,1\}^{\zeta(\kappa)}, y = Q(\kappa, x)] > \delta(\kappa)$.

Theorem 7.2. *For a deterministic two-party function f, SFE_f (semi-honest secure protocol for f) the following statements are equivalent:*

(1) *f has a perfectly semi-honest secure protocol (in the plain model).*
(2) *SFE_f has a security-parameter preserving fully black-box construction from* OWF.
(3) *SFE_f has a security-parameter preserving fully black-box construction from OWF_ζ, for some polynomial ζ.*

We prove this theorem in Appendix C.

8 Open Problems and Future Work

We have shown a black-box separation between one-way functions and semi-honest SFE protocols for 2-party secure function evaluation for any function which does not have already have a semi-honest SFE protocol in the plain model. Intuitively, this introduces new worlds between "minicrypt" and "cryptomania" [Imp95], corresponding to where these functions have semi-honest SFE protocols. There are several interesting questions that this gives rise to. We mention a few directions below.

1. Our result relies on the combinatorial characterization of undecomposable function evaluations. In particular, our strategy is not able to "compile out" the random oracle completely in the context of 2-party deterministic semi-honest function evaluation, i.e., we are not able to rule out that access to a random oracle could enable secure computation (of say, a randomized functionality) that cannot be achieved by a protocol in the plain model. Understanding the precise power of random oracles in the context of secure computation in its full generality (especially, for randomized functions) remains open.

2. The separation of OT from one-way functions (implicit) in [IR89] was strengthened to separate OT from public-key encryption in [GKM+00]. In on going work, we give a similar strengthening of our results, separating every function which does not have a semi-honest SFE protocol in the plain model (undecomposable functions, among symmetric functions) from public-key encryption. This, in particular, would give an alternate proof for the result in [GKM+00].

3. In this work we *do not* show that (semi-honest) SFE for the various functions we separate from one-way functions really correspond to *new* worlds in Impagliazzo's universe. In particular, we do not separate them from the "OT protocol" primitive. Indeed, one could hope to prove our current results by simply showing that SFE for all the functions we considered can, in a fully black-box manner, yield an OT protocol. But we conjecture that such a construction simply does not exist. We leave it open to fully understand the relationship between the worlds corresponding to (semi-honest) SFE protocols for the different functions, and in particular, find out if there is an infinite hierarchy of such distinct worlds.

References

[BBCR10] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In Leonard J. Schulman, editor, *STOC*, pages 67–76. ACM, 2010. 1

[Bea89] Donald Beaver. Perfect privacy for two-party protocols. In Joan Feigenbaum and Michael Merritt, editors, *Proceedings of DIMACS Workshop on Distributed Computing and Cryptography*, volume 2, pages 65–77. American Mathematical Society, 1989. 1, 4, 8

[BM07] Boaz Barak and Mohammad Mahmoody. Lower bounds on signatures from symmetric primitives. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 2007. 4

[BM09] Boaz Barak and Mohammad Mahmoody. Merkle puzzles are optimal - an $O(n^2)$-query attack on any key exchange from a random oracle. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390. Springer, 2009. 3, 4, 5, 7, 8, 39

[BMM99] Amos Beimel, Tal Malkin, and Silvio Micali. The all-or-nothing nature of two-party secure computation. In *CRYPTO*, pages 80–97, 1999. 1

[BPR$^+$08] Dan Boneh, Periklis A. Papakonstantinou, Charles Rackoff, Yevgeniy Vahlis, and Brent Waters. On the impossibility of basing identity based encryption on trapdoor permutations. In *FOCS*, pages 283–292. IEEE Computer Society, 2008. 1, 4

[CI93] Richard Cleve and Russell Impagliazzo. Martingales, collective coin flipping and discrete control processes (extended abstract), 1993. 1, 4

[CK89] Benny Chor and Eyal Kushilevitz. A zero-one law for boolean privacy (extended abstract). In Johnson [Joh89], pages 62–72. 1, 4

[DBL89] *30th Annual Symposium on Foundations of Computer Science, 30 October-1 November 1989, Research Triangle Park, North Carolina, USA*. IEEE, 1989. 36, 37

[DLMM11] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On black-box complexity of optimally-fair coin-tossing. In *Theory of Cryptography Conference - TCC 2011*, 2011. 4, 7, 39

[GGKT05] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM J. Comput.*, 35(1):217–246, 2005. 4

[GKM$^+$00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The relationship between public key encryption and oblivious transfer. In *FOCS*, pages 325–335. IEEE Computer Society, 2000. 1, 4, 34

[GMR01] Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *FOCS*, pages 126–135, 2001. 1, 4

[Hai08] Iftach Haitner. Semi-honest to malicious oblivious transfer - the black-box way. In *TCC*, pages 412–426, 2008. 4

[HHRS07] Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols - a tight lower bound on the round complexity of statistically-hiding commitments. In *FOCS*, pages 669–679. IEEE Computer Society, 2007. 4

[HNO+09] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009. 4

[HOZ12] Iftach Haitner, Eran Omri, and Hila Zarosim. On the power of random oracles. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:129, 2012. 4

[HOZ13] Iftach Haitner, Eran Omri, and Hila Zarosim. Limits on the usefulness of random oracles. *Theory of Cryptography Conference (TCC, to appear)*, 2013. 4

[IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *FOCS* [DBL89], pages 230–235. 4

[Imp95] Russell Impagliazzo. A personal view of average-case complexity. In *Structure in Complexity Theory Conference*, pages 134–147, 1995. 4, 34

[IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Johnson [Joh89], pages 44–61. 1, 3, 4, 5, 7, 8, 10, 32, 34

[Ish11] Yuval Ishai, editor. *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings*, volume 6597 of *Lecture Notes in Computer Science*. Springer, 2011. 36, 37

[Joh89] David S. Johnson, editor. *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing, 15-17 May 1989, Seattle, Washington, USA*. ACM, 1989. 35, 36

[Kil91] Joe Kilian. A general completeness theorem for two-party games. In *STOC*, pages 553–560. ACM, 1991. 1

[KKMO00] Joe Kilian, Eyal Kushilevitz, Silvio Micali, and Rafail Ostrovsky. Reducibility and completeness in private computations. *SIAM J. Comput.*, 29(4):1189–1208, 2000. 1

[KM11] Daniel Kraschewski and Jörn Müller-Quade. Completeness theorems with constructive proofs for finite deterministic 2-party functions. In *TCC*, 2011. 1, 31

[KMR09] Robin Künzler, Jörn Müller-Quade, and Dominik Raub. Secure computability of functions in the it setting with dishonest majority and applications to long-term security. In Reingold [Rei09], pages 238–255. 4

[Kre11] Gunnar Kreitz. A zero-one law for secure multi-party computation with ternary outputs. In Ishai [Ish11], pages 382–399. 1

[KST99] Jeong Han Kim, Daniel R. Simon, and Prasad Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *FOCS*, pages 535–542, 1999. 4

[KSY11] Jonathan Katz, Dominique Schröder, and Arkady Yerukhimovich. Impossibility of blind signatures from one-way permutations. In Ishai [Ish11], pages 615–629. 1, 4

[Kus89] Eyal Kushilevitz. Privacy and communication complexity. In *FOCS* [DBL89], pages 416–421. 1, 2, 4, 8

[LTW05] Henry C. Lin, Luca Trevisan, and Hoeteck Wee. On hardness amplification of one-way functions. In Joe Kilian, editor, *TCC*, volume 3378 of *Lecture Notes in Computer Science*, pages 34–49. Springer, 2005. 4

[Maj11] Hemanta K. Maji. *On Computational Intractability Assumptions in Cryptography*. PhD thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, 2011. 4, 6

[MM11] Takahiro Matsuda and Kanta Matsuura. On black-box separations among injective one-way functions. In Ishai [Ish11], pages 597–614. 1, 4

[MOPR11] Hemanta K. Maji, Pichayoot Ouppaphan, Manoj Prabhakaran, and Mike Rosulek. Exploring the limits of common coins using frontier analysis of protocols. In *TCC*, 2011. 1, 2, 4, 18

[MPR09] Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. Complexity of multi-party computation problems: The case of 2-party symmetric secure function evaluation. In Reingold [Rei09], pages 256–273. Full version available from IACR Eprint Archive: `http://eprint.iacr.org`. 1, 2, 4, 5, 23, 31

[MPR10] Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. A zero-one law for cryptographic complexity with respect to computational uc security. In *CRYPTO*, pages 595–612, 2010. 1, 31

[Ost91] Rafail Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *Structure in Complexity Theory Conference*, pages 133–138, 1991. 4

[OW93] Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. Technical Report TR-93-073, International Computer Science Institute, Berkeley, CA, November 1993. Preliminary version in Proc. 2nd Israeli Symp. on Theory of Computing and Systems, 1993, pp. 3–17. 4

[Rei09] Omer Reingold, editor. *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*. Springer, 2009. 36, 37

[RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *TCC*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2004. 3, 4, 9, 32, 33

[Sim98] Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *EUROCRYPT*, pages 334–345, 1998. 1, 4

A Independence Learners

A.1 Some Notations

Before, we proceed, we introduce some notations used in this section.

A.1.1 Random Variables.

We use bold letters to emphasize the nature of a random variable (e.g., $\mathbf{x}$). By $\mathrm{Supp}(\mathbf{x})$ we denote $\{x \mid \Pr[\mathbf{x} = x] > 0\}$. By $x \xleftarrow{\$} \mathbf{x}$ we mean that x is sampled according to the distribution of the random variable $\mathbf{x}$. We usually use the same letter to denote a sample from a random variable. When we say an event occurs with negligible probability denoted by $\mathrm{negl}(\kappa)$, we mean it occurs with probability $\kappa^{-\omega(1)}$. We call two random variables $\mathbf{x}, \mathbf{y}$ (or their corresponding distributions) ε-close if their statistical distance, defined as $\mathsf{SD}(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \cdot \sum_{s \in \mathrm{Supp}(\mathbf{x}) \cup \mathrm{Supp}(\mathbf{y})} |\Pr[\mathbf{x} = s] - \Pr[\mathbf{y} = s]|$ is at most ε. By $\mathbf{x} \equiv \mathbf{y}$ we denote that the random variables $\mathbf{x}$ and $\mathbf{y}$ are distributed identically.

By $(\mathbf{x}_1, \mathbf{x}_2, \dots)$ we denote a (perhaps infinite) sequence of *correlated* random variables where $\mathbf{x}_i$ is the random variable of the i^{th} coordinate. For correlated random variables $(\mathbf{x}, \mathbf{y})$, by $(\mathbf{x} \times \mathbf{y})$ we refer to a new random variable that samples *independent* copies for $\mathbf{x}$ and $\mathbf{y}$ (i.e., sample two pairs $(x_1, y_1) \leftarrow (\mathbf{x}, \mathbf{y}), (x_2, y_2) \xleftarrow{\$} (\mathbf{x}, \mathbf{y})$ and output (x_1, y_2)). For correlated random variables $(\mathbf{x}, \mathbf{y})$ and $y \in \mathrm{Supp}(\mathbf{y})$ by $(\mathbf{x} \mid \mathbf{y} = y)$ we denote the random variable $\mathbf{x}$ conditioned on $\mathbf{y} = y$. When it is clear from the context we simply write $(\mathbf{x} \mid y)$ instead of $(\mathbf{x} \mid \mathbf{y} = y)$.

A.1.2 Two Party Protocols

In the proofs in this section, we mostly analyze the protocols by rounds rather than frontiers. Hence it will be convenient to introduce notation involving round numbers (rather than nodes in the transcript tree). Below we describe notation associated with Alice; similar notation is associated with Bob as well.

1. $\mu = (\mu_1, \mu_2, \dots)$ denotes the transcript generated during the interaction where the i^{th} message is sent by Alice, if i is odd, and it is sent by Bob, if i is even. By $\mu^{(i)}$ we denote $(\mu_1, \dots, \mu_i)$.

2. Let P_A denote the set of oracle query-answer pairs obtained by Alice from the oracle. By $P_A^{(i)}$ we denote the set of query-answers obtained by Alice before $\mu^{(i)}$ is sent. Q_A and $Q_A^{(i)}$ are defined similarly to P_A and $P_A^{(i)}$ while only containing the queries. Namely, using the query-operator $\mathcal{Q}$ defined earlier it, holds that $Q_A = \mathcal{Q}(P_A)$ and $Q_A^{(i)} = \mathcal{Q}(P_A^{(i)})$.

3. V_A denotes the view of Alice which is equal to (x, r_A, P_A, μ), where r_A denotes the private random-tape of Alice. By $V_A^{(i)}$ we denote the view of Alice till the message $\mu^{(i)}$ is sent which is equal to $(x, r_A, P_A^{(i)}, \mu^{(i-1)})$

A Public Query Strategy. For a 2-party protocol Π in the $\mathcal{O}$ model, we define a *public query strategy* Eve as a deterministic algorithm which takes as input a prefix $\mu^{(i)}$ of the messages of an execution of Π and a set $P_E^{(i-1)}$ of query-answer pairs from $\mathcal{O}$ (standing for the queries that she has asked previously), and then adaptively queries $\mathcal{O}$ multiple times. The view of Eve, denoted by V_E, is equal to (μ, P_E). We also define $V^{(i+1)} = (P_E^{(i+1)}, \mu^{(i)})$ as the view of Eve before μ_{i+1} is sent.

We define the query complexity of a public query strategy Eve for a protocol Π to be the maximum number of queries Eve makes to $\mathcal{O}$ over an entire augmented execution of Π and Eve.

A Round. For an odd i, the i^{th} round starts right after the $(i-1)^{\text{st}}$ message is received by Alice and starts asking its oracle queries (which are contained in $Q_A^{(i)}$). When Alice sends μ_i the i^{th} round continues when Eve asks its oracle queries (contained in $P_E^{(i)}$). This round ends when Eve is done with asking her oracle queries. For an even i, the definition of the i^{th} round is similar (switching between Alice and Bob).

A.2 Independence Learner

The following lemma was implicit in the work of [BM09] and was proved explicitly in [DLMM11] (here, for simplicity, we use this lemma with more relaxed parameters).

Lemma A.1 (Independence Learner for Protocols with No Input [BM09, DLMM11])**.** *Let Π be an N-round input-less randomized two-party protocol using a random oracle $\mathcal{O}$, with m query complexity. Then, for any threshold $0 \le \varepsilon \le 1$, Π has a public query strategy* Eve *(who only observes the public messages) with query complexity* $\text{poly}(m/\varepsilon)$*, such that with probability at least $1-\varepsilon$ over the choice of the view of* Eve*: $V_E \xleftarrow{\$} \mathbf{V}_E$ the following holds. (Recall that $V_E = (P_E, \mu)$ and $V_E^{(i)}$ is the part of V_E that corresponds to the first i rounds).*

1. $(1-\varepsilon)$**-Independence:** *For every $i \in [N]$ the following distributions are ε-close:*

$$\left((\mathbf{V}_A^{(i)} \mid V_E^{(i)}) \times (\mathbf{V}_B^{(i)} \mid V_E^{(i)})\right) \ \text{ and } \ \left((\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)}) \mid V_E^{(i)}\right).$$

   Namely, if we sample the views of Alice and Bob jointly conditioned on $V_E^{(i)}$, this joint distribution is ε-close to the product distribution in which Alice and Bob's views are sampled independently (each conditioned on the same $V_E^{(i)}$).

2. ε**-Lightness:** *For every $q \notin Q_E^{(i)}$ (where $Q_E^{(i)} = \mathcal{Q}(V_E^{(i)})$) it holds that*

$$\mathsf{P}_{V_A^{(i)} \xleftarrow{\$} (\mathbf{V}_A^{(i)} | V_E^{(i)})}[q \in \mathcal{Q}(V_A^{(i)})] \ \le \varepsilon \ \text{ and } \ \mathsf{P}_{V_B^{(i)} \xleftarrow{\$} (\mathbf{V}_B^{(i)} | V_E^{(i)})}[q \in \mathcal{Q}(V_B^{(i)})] \ \le \varepsilon.$$

A.3 Using the Independence Learner

The Independence Learner of Lemma A.1 is not directly useful in our context. We need two additional technical properties ensured by the independence learner which are mentioned below.

The first lemma formalizes the intuition that a curious eavesdropper when run with appropriate parameters can ensure that whenever Alice sends a message in the protocol she can only add information about her input and not Bob's input.

Lemma A.2 (Independence Learner for Likely Inputs)**.** *Let Π be a secure protocol for some secure function evaluation relative to a random oracle $\mathcal{O}$ and Alice asks m queries to the random oracle. Suppose $\mathcal{X}$ and $\mathcal{Y}$ are, respectively, the set of inputs for Alice and Bob. We run* Eve *with input parameter $\varepsilon < 1$ over Π assuming that Π is run with $\tilde{x} \xleftarrow{\$} \mathcal{X}$ and $\tilde{y} \xleftarrow{\$} \mathcal{Y}$. Let $x \in \mathcal{X}$ and $y, y' \in \mathcal{Y}$ be fixed inputs. Then for some $\varepsilon' = \varepsilon^{\Omega(1)}(m \cdot |\mathcal{X}|)^{O(1)}$, if we run the protocol Π with inputs x and y together the curious eavesdropper* Eve*, for every even $i \in [N]$ (i.e., Bob sends μ_i), with probability at least $1-\varepsilon'$ over the choice of the view of* Eve *$V_E^{(i)} \xleftarrow{\$} \mathbf{V}_E^{(i)}$ at least one of the following holds:*

1. $\mathsf{P}[y \mid V_E^{(i)}, x] < \varepsilon'$,
2. $\mathsf{P}[y' \mid V_E^{(i)}, x] < \varepsilon'$, *or*
3. $\mathsf{SD}((\boldsymbol{\mu}_{i+1} \mid V_E^{(i)}, x, y), (\boldsymbol{\mu}_{i+1} \mid V_E^{(i)}, x, y')) \le \varepsilon'$.

The second lemma is slightly more technical. The curious eavesdropper of Lemma A.1 ensures that all intersection queries are covered with high probability when Alice and Bob execute the protocol with actual inputs x and y. We need a stronger version of this result. We want to claim that even if Bob pretends to change his input to y' and samples a corresponding local view, the intersection queries of this "hypothetical view" are also covered by the actual Eve view with high probability. This ensures that we can sample a consistent random oracle even without the knowledge of actual Alice input x while simulating the hypothetical view. Looking ahead, this lemma shall be useful when Bob launches a curious attack by changing his private input appropriately.

Lemma A.3 (Bounding Collisions of Queries for Likely Inputs)**.** *Let Π be a secure protocol for some secure function evaluation relative to a random oracle $\mathcal{O}$ in which Alice asks m queries to the random oracle. Suppose $\mathcal{X}$ and $\mathcal{Y}$ are, respectively, the set of inputs for Alice and Bob. We run* Eve *with input parameter $\varepsilon < 1$ over Π assuming that Π is run with $\tilde{x} \xleftarrow{\$} \mathcal{X}$ and $\tilde{y} \xleftarrow{\$} \mathcal{Y}$. Let $x \in \mathcal{X}, y, y' \in \mathcal{Y}$ be some fixed inputs. Suppose we perform the following samplings:*

$$(V_E^{(i)}, Q_B^{(i)}, Q_A^{(i+1)}) \xleftarrow{\$} (\mathbf{V}_E^{(i)}, \mathbf{Q}_B^{(i)}, \mathbf{Q}_A^{(i+1)} \mid x, y) \text{ and } Q'^{(i)}_B \xleftarrow{\$} (\mathbf{Q}_B^{(i)} \mid V_E^{(i)}, \mu_{i+1}, y').$$

In the second sampling: the protocol is executed with inputs x, y and Alice's message μ_{i+1} is generated, and after that we sample a view of Bob for the first i rounds conditioned on $V_E^{(i)}$, μ_{i+1} and Bob's input being y'. Then for some $\varepsilon' = \varepsilon^{\Omega(1)}(m \cdot |\mathcal{X}| \cdot |\mathcal{Y}|)^{O(1)}$ with probability at least $1 - \varepsilon'$ it holds that either

1. $\mathsf{P}[y' \mid V_E^{(i)}, x] < \varepsilon'$, *or*
2. $Q_A^{(i+1)} \cap (Q_B^{(i)} \cup Q'^{(i)}_B) \subseteq \mathcal{Q}(V_E^{(i)})$.

Before proving Lemma A.2 and Lemma A.3 we need to develop some general tools of probability.

A.4 General Useful Lemmas

A corollary to Lemma 2.1 is that the actual inputs of Alice and Bob will not become "unlikely" conditioned on Eve's view, except with small probability.

Corollary A.4. *Suppose Alice and Bob run a two party protocol with inputs x, y chosen from an arbitrary distribution and suppose Eve is some public query strategy. Then the probability that at some point during the protocol it holds that $\mathsf{P}[(x, y)|u] < \theta$ where u is the view of Eve so far, is at most $\theta / \mathsf{P}[(x, y)]$ (if the inputs are chosen uniformly at random from the sets $\mathcal{X}, \mathcal{Y}$, this probability is at most $\theta|\mathcal{X}||\mathcal{X}|$).*

Proof. Corollary A.4 follows by a direct application of Lemma 2.1 by using the event X corresponds to the case that (x, y) are the inputs, and the sequence of random variables $(\mathbf{m}_1, \mathbf{m}_2, \dots)$ corresponds to the sequence of the bits representing the view of Eve. □

The following lemma states that if two random variables $\mathbf{a}, \mathbf{b}$ are statistically close, they will "remain close" even if we condition on a "likely event" defined over their supports.

Lemma A.5. *Let $\mathbf{a}, \mathbf{b}$ be two random variables such that $\mathsf{SD}(\mathbf{a}, \mathbf{b}) \leq \varepsilon$. Suppose $E \subseteq \text{Supp}(\mathbf{a}) \cup \text{Supp}(\mathbf{b})$ be an event such that $\mathsf{P}[\mathbf{a} \in E] \geq \delta > 0$ and $\mathsf{P}[\mathbf{b} \in E] > 0$. Define $\mathbf{a}_E \equiv (\mathbf{a} \mid E)$ and $\mathbf{b}_E \equiv (\mathbf{b} \mid E)$. Then, $\mathsf{SD}(\mathbf{a}_E, \mathbf{b}_E) \leq \varepsilon/\delta$.*

Proof. First, we prove a weaker bound of $3\varepsilon/2\delta$ and then will sharpen the analysis to obtain the optimal bound of ε/δ.

Let $\alpha = \mathsf{P}[\mathbf{a} \in E]$ and $\beta = \mathsf{P}[\mathbf{b} \in E]$. Recall, we are guaranteed that $\alpha \geq \delta > 0$ and $\beta > 0$. Moreover, $\sum_{s \in E} |\mathsf{P}[\mathbf{a} = s] - \mathsf{P}[\mathbf{b} = s]| \leq 2\varepsilon$ and $|\alpha - \beta| \leq \varepsilon$, because $\mathsf{SD}(\mathbf{a}, \mathbf{b}) \leq \varepsilon$. Observe that $\mathsf{P}[\mathbf{a}_E = s] = \mathsf{P}[\mathbf{a} = s]/\alpha$ and $\mathsf{P}[\mathbf{b}_E = s] = \mathsf{P}[\mathbf{b} = s]/\beta$, for $s \in E$. Therefore, we can perform the following simplification:

$$\begin{aligned}
\mathsf{SD}(\mathbf{a}_E, \mathbf{b}_E) &= \frac{1}{2} \sum_{s \in E} \left| \frac{\mathsf{P}[\mathbf{a} = s]}{\alpha} - \frac{\mathsf{P}[\mathbf{b} = s]}{\beta} \right| \\
&\leq \left(\frac{1}{2} \sum_{s \in E} \left| \frac{\mathsf{P}[\mathbf{a} = s]}{\alpha} - \frac{\mathsf{P}[\mathbf{b} = s]}{\alpha} \right| \right) + \left(\frac{1}{2} \sum_{s \in E} \left| \frac{\mathsf{P}[\mathbf{b} = s]}{\alpha} - \frac{\mathsf{P}[\mathbf{b} = s]}{\beta} \right| \right) \\
&\leq \left(\frac{2\varepsilon}{2\alpha} \right) + \left(\frac{|\alpha - \beta|}{2\alpha} \right) \quad \leq \frac{3\varepsilon}{2\alpha} \quad \leq \frac{3\varepsilon}{2\delta}
\end{aligned}$$

With a more careful case analysis, the upper bound can be improved to ε/δ (which is tight). Consider these two cases:

1. Case $\alpha \geq \beta$: We shall partition the set E into three sets E_1, E_2 and E_3 as follows:

$$\begin{aligned}
E_1 &= \{s | s \in E, \mathsf{P}[\mathbf{a} = s]/\alpha \geq \mathsf{P}[\mathbf{b} = s]/\beta\} \\
E_2 &= \{s | s \in E, \mathsf{P}[\mathbf{a} = s] \geq \mathsf{P}[\mathbf{b} = s] \text{ but } \mathsf{P}[\mathbf{a} = s]/\alpha < \mathsf{P}[\mathbf{b} = s]/\beta\} \\
E_3 &= \{s | s \in E, \mathsf{P}[\mathbf{a} = s] < \mathsf{P}[\mathbf{b} = s]\}
\end{aligned}$$

Let $u_i = \mathsf{P}[\mathbf{a} \in E_i]$ and $v_i = \mathsf{P}[\mathbf{b} \in E_i]$, where $i \in \{1, 2, 3\}$. We shall use the following constraints: $v_2 \leq u_2$, $v_3 \leq u_3 + \varepsilon$ and $v_1 \geq u_1 - \varepsilon$. Now, consider the following manipulation:

$$\begin{aligned}
\mathsf{SD}(\mathbf{a}_E, \mathbf{b}_E) &= \frac{1}{2} \sum_{s \in E} \left| \frac{\mathsf{P}[\mathbf{a} = s]}{\alpha} - \frac{\mathbf{b} = s}{\beta} \right| \\
&= \frac{1}{2} \left[\left(\frac{u_1}{\alpha} - \frac{v_1}{\beta} \right) - \left(\frac{u_2}{\alpha} - \frac{v_2}{\beta} \right) - \left(\frac{u_3}{\alpha} - \frac{v_3}{\beta} \right) \right] \\
&= \frac{u_1}{\alpha} - \frac{v_1}{v_1 + v_2 + v_3} \\
&\leq \frac{u_1}{\alpha} - \frac{v_1}{v_1 + u_2 + v_3} \qquad (\because v_2 \leq u_2) \\
&\leq \frac{u_1}{\alpha} - \frac{v_1}{v_1 + u_2 + u_3 + \varepsilon} \qquad (\because v_3 \leq u_3 + \varepsilon) \\
&\leq \frac{u_1}{\alpha} - \frac{u_1 - \varepsilon}{u_1 + u_2 + u_3} \qquad (\because v_1 \geq u_1 - \varepsilon) \\
&= \frac{\varepsilon}{\alpha} \qquad \leq \frac{\varepsilon}{\delta}
\end{aligned}$$

2. Case $\alpha < \beta$: We shall partition the set E into three sets E_1, E_2 and E_3 as follows:

$$E_1 = \{s | s \in E, \mathsf{P}[\mathbf{a} = s] \geq \mathsf{P}[\mathbf{b} = s]\}$$
$$E_2 = \{s | s \in E, \mathsf{P}[\mathbf{a} = s] < \mathsf{P}[\mathbf{b} = s] \text{ but } \mathsf{P}[\mathbf{a} = s]/\alpha \geq \mathsf{P}[\mathbf{b} = s]/\beta\}$$
$$E_3 = \{s | s \in E, \mathsf{P}[\mathbf{a} = s]/\alpha < \mathsf{P}[\mathbf{b} = s]/\beta\}$$

Let $u_i = \mathsf{P}[\mathbf{a} \in E_i]$ and $v_i = \mathsf{P}[\mathbf{b} \in E_i]$, where $i \in \{1, 2, 3\}$. We shall use the following constraints: $v_2 > u_2$, $v_3 \leq u_3 + \varepsilon$ and $v_1 \geq u_1 - \varepsilon$. Now, consider the following manipulation:

$$\begin{aligned}
\mathsf{SD}(\mathbf{a}_E, \mathbf{b}_E) &= \frac{1}{2} \sum_{s \in E} \left| \frac{\mathsf{P}[\mathbf{a} = s]}{\alpha} - \frac{\mathbf{b} = s}{\beta} \right| \\
&= \frac{1}{2} \left[\left(\frac{u_1}{\alpha} - \frac{v_1}{\beta} \right) + \left(\frac{u_2}{\alpha} - \frac{v_2}{\beta} \right) - \left(\frac{u_3}{\alpha} - \frac{v_3}{\beta} \right) \right] \\
&= \frac{v_3}{v_1 + v_2 + v_3} - \frac{u_3}{\alpha} \\
&< \frac{v_3}{v_1 + u_2 + v_3} - \frac{u_3}{\alpha} \qquad (\because v_2 > u_2) \\
&\leq \frac{u_3 + \varepsilon}{v_1 + u_2 + u_3 + \varepsilon} - \frac{u_3}{\alpha} \qquad (\because v_3 \leq u_3 + \varepsilon) \\
&\leq \frac{u_3 + \varepsilon}{u_1 + u_2 + u_3} - \frac{u_3}{\alpha} \qquad (\because v_1 \geq u_1 - \varepsilon) \\
&= \frac{\varepsilon}{\alpha} \qquad \leq \frac{\varepsilon}{\delta}
\end{aligned}$$

This completes the proof that $\mathsf{SD}(\mathbf{a}_E, \mathbf{b}_E) \leq \varepsilon/\delta$. Equality holds if and only if, $\{s | \mathsf{P}[\mathbf{a} = s] \neq \mathsf{P}[\mathbf{b} = s]\} \subseteq E$ and $\mathsf{P}[\mathbf{a} \in E] = \delta$. ☐

The following lemma states that if two random are close to being independent iff they are close to the product of their marginal distribution.

Lemma A.6. *Let* $(\mathbf{a}, \mathbf{b})$ *be jointly distributed random variables such that* $\mathsf{SD}((\mathbf{a}, \mathbf{b}), (\mathbf{u} \times \mathbf{v})) \leq \varepsilon$ *for some random variables* $\mathbf{u}$ *and* $\mathbf{v}$*. Then it holds that* $\mathsf{SD}((\mathbf{a}, \mathbf{b}), (\mathbf{a} \times \mathbf{b})) \leq 3\varepsilon$*.*

Proof. $\mathsf{SD}((\mathbf{a}, \mathbf{b}), (\mathbf{u} \times \mathbf{v})) \leq \varepsilon$ implies that $\mathsf{SD}(\mathbf{a}, \mathbf{u}) \leq \varepsilon$ and $\mathsf{SD}(\mathbf{b}, \mathbf{v}) \leq \varepsilon$. Therefore, by two applications of triangle inequality it holds that: $\mathsf{SD}((\mathbf{a}, \mathbf{b}), (\mathbf{a} \times \mathbf{b})) \leq \mathsf{SD}((\mathbf{a}, \mathbf{b}), (\mathbf{u} \times \mathbf{v})) + \mathsf{SD}((\mathbf{u} \times \mathbf{v}), (\mathbf{a} \times \mathbf{v})) + \mathsf{SD}((\mathbf{a} \times \mathbf{v}), (\mathbf{a} \times \mathbf{b})) \leq 3\varepsilon$. ☐

The following lemma states that whenever two random variables $(\mathbf{a}, \mathbf{b})$ are close to being independent, then they will remain so, even if we sample $\mathbf{a}$ conditioned on some partial leakage c as a function of $\mathbf{b}$.

Lemma A.7. *Let* $(\mathbf{a}, \mathbf{b})$ *be jointly distributed random variables such that* $\mathsf{SD}((\mathbf{a}, \mathbf{b}), (\mathbf{a} \times \mathbf{b})) \leq \varepsilon$*. Suppose* $c = f(b)$ *is a possibly randomized function of* b*, where the random tape for* $f(\cdot)$ *is chosen uniformly and independently at random. Given a sample for* $(b, f(b) = c)$*, let* $\mathbf{a}'$ *be another random variable sampled from the distribution* $(\mathbf{a} \mid c)$*. Then it holds that* $\mathsf{SD}((\mathbf{b}, \mathbf{a}'), (\mathbf{b} \times \mathbf{a})) \leq \varepsilon$*.*

Proof. Suppose $f(b;r)$ is the deterministic function where r is the random tape used to evaluate the randomized function f. This case reduces to the deterministic case as follows:

$$\mathsf{SD}((\mathbf{b},\mathbf{a}'),(\mathbf{b}\times\mathbf{a})) \le \mathsf{SD}((\mathbf{b},\mathbf{a}',\mathbf{r}),(\mathbf{b}\times\mathbf{a}\times\mathbf{r})) = \mathsf{SD}(((\mathbf{b}\times\mathbf{r}),\mathbf{a}'),((\mathbf{b}\times\mathbf{r})\times\mathbf{a}))$$

Henceforth, we can assume, without loss of generality, that f is a deterministic function. In this case:

$$\begin{aligned}
2\mathsf{SD}((\mathbf{b},\mathbf{a}'),(\mathbf{b}\times\mathbf{a})) &= \sum_a\sum_b \mathsf{P}[\mathbf{b}=b]\cdot|\mathsf{P}[\mathbf{a}=a|c=f(b)]-\mathsf{P}[\mathbf{a}=a]| \\
&= \sum_a\sum_c\sum_{b\in f^{-1}(c)} \mathsf{P}[\mathbf{b}=b]\cdot|\mathsf{P}[\mathbf{a}=a|c]-\mathsf{P}[\mathbf{a}=a]| \\
&= \sum_a\sum_c \mathsf{P}[\mathbf{c}=c]\cdot|\mathsf{P}[\mathbf{a}=a|c]-\mathsf{P}[\mathbf{a}=a]| \\
&= \sum_a\sum_c |\mathsf{P}[\mathbf{a}=a,\mathbf{c}=c]-\mathsf{P}[\mathbf{a}=a]\mathsf{P}[\mathbf{c}=c]| \\
&= \sum_a\sum_c \left|\sum_{b\in f^{-1}(c)} (\mathsf{P}[\mathbf{a}=a,\mathbf{b}=b]-\mathsf{P}[\mathbf{a}=a]\mathsf{P}[\mathbf{b}=b])\right| \\
&\le \sum_a\sum_c\sum_{b\in f^{-1}(c)} |\mathsf{P}[\mathbf{a}=a,\mathbf{b}=b]-\mathsf{P}[\mathbf{a}=a]\mathsf{P}[\mathbf{b}=b]| \\
&= 2\mathsf{SD}((\mathbf{b},\mathbf{a}),(\mathbf{b}\times\mathbf{a})).
\end{aligned}$$

□

A.5 Proving Lemma A.2 and Lemma A.3

We shall prove both Lemma A.2 and Lemma A.3 both using the following intermediate lemma.

Lemma A.8. *Suppose $V_E^{(i)}$ is the view of* Eve *by the end of the i^{th} round with respect to the two party protocol in which the inputs are chosen at random and is such that the $(1-\varepsilon)$-Independence and ε-Lightness properties hold conditioned on $V_E^{(i)}$. Suppose $x\in\mathcal{X}, y\in\mathcal{Y}$ are such that $\mathsf{P}[x,y\mid V_E^{(i)}]\ge\gamma$ and m is the total number Alice's queries. Then both of the following hold:*

1. $\mathsf{P}[Q_A^{(i+1)}\cap Q_B^{(i)}\not\subseteq Q_E^{(i)}\mid V_E^{(i)},x,y]\le O(m\varepsilon/\gamma)$.
2. *The following two are $O(m\varepsilon/\gamma)$-close:*

$$(\mathbf{V}_A^{(i+1)},\mathbf{V}_B^{(i)}\mid V_E^{(i)},x,y) \text{ and } ((\mathbf{V}_A^{(i+1)}\mid V_E^{(i)},x)\times(\mathbf{V}_B^{(i)}\mid V_E^{(i)},y)).$$

Before proving Lemma A.8 we shall see how it can be used to prove Lemma A.2 and Lemma A.3.

A.5.1 Proof of Lemma A.2

For simplicity, we shall use another parameter $0<\sigma<1$ and prove the following result: With probability at least $1-\varepsilon-\sigma|\mathcal{X}|$ over the choice of the Eve view $V_E^{(i)}\xleftarrow{\$}\mathbf{V}_E^{(i)}$ at least one of the following holds:

1. $\mathsf{P}[y\mid V_E^{(i)},x]<\sigma$,

2. $\mathsf{P}[y' \mid V_E^{(i)}, x] < \sigma$, or

3. $\mathsf{SD}((\boldsymbol{\mu}_{i+1} \mid V_E^{(i)}, x, y), (\boldsymbol{\mu}_{i+1} \mid V_E^{(i)}, x, y')) \leq \ O(m\varepsilon/\sigma^2)$ where m is the number of oracle queries asked by Alice during the protocol.

Then Lemma A.2 follows by setting $\varepsilon = \sigma^3$ and taking $\varepsilon' = \max(\sigma, \varepsilon + \sigma|\mathcal{X}|, m\varepsilon/\sigma^2)$ in the above mentioned statement.

By Lemma A.1, with probability at least $1 - \varepsilon$ over the choice of $V_E^{(i)}$, the $(1-\varepsilon)$-Independence and ε-Lightness properties both hold. Corollary A.4 implies that with probability at least $1 - \sigma|\mathcal{X}|$, we shall have $\mathsf{P}[x|V_E^{(i)}] \geq \sigma$. By union bound, both these events hold with probability at least $1 - \varepsilon - \sigma|\mathcal{X}|$. Henceforth, we shall assume that both these conditions hold for our choice of $V_E^{(i)}$.

For our choice of $V_E^{(i)}$, if one of the first two cases of Lemma A.2 holds then we are done. Suppose this is not the case. Then, we have $\mathsf{P}[y \mid V_E^{(i)}, x] \geq \sigma$ and $\mathsf{P}[y' \mid V_E^{(i)}, x] \geq \sigma$. Therefore we can conclude that *both pairs* of inputs (x, y) and (x, y') are "likely" conditioned on $V_E^{(i)}$. More formally:

$$\mathsf{P}[x, y \mid V_E^{(i)}] \geq \mathsf{P}[x \mid V_E^{(i)}] \cdot \mathsf{P}[y \mid V_E^{(i)}, x] \geq \sigma^2 \text{ and similarly } \mathsf{P}[x, y' \mid V_E^{(i)}] \geq \sigma^2$$

So, currently we are considering $V_E^{(i)}$ such that $\mathsf{P}[x, y|V_E^{(i)}] \geq \sigma^2$, $\mathsf{P}[x, y'|V_E^{(i)}] \geq \sigma^2$; and $(1-\varepsilon)$-Independence and ε-Lightness guarantees hold. Therefore Lemma A.2 follows by the second part of Lemma A.8 because $(\mathbf{V}_A^{(i+1)} \mid V_E^{(i)}, x)$ is independent of y and y' and μ_{i+1} is a function of $V_A^{(i+1)}$.

A.5.2 Proof of Lemma A.3

Similarly to the proof of Lemma A.2, we use another parameter $0 < \sigma < 1$ and prove the following statement: With probability $1 - O(\varepsilon + \sigma|\mathcal{X}| \cdot |\mathcal{Y}| + m\varepsilon/\sigma^2)$ over the samples at least one of the following is true:

1. $\mathsf{P}[y' \mid V_E^{(i)}, x] < \sigma$, or

2. $Q_A^{(i+1)} \cap (Q_B^{(i)} \cup Q_B'^{(i)}) \subseteq \mathcal{Q}(V_E^{(i)})$.

Lemma A.3 follows by setting $\sigma^3 = \varepsilon$ and $\varepsilon' = \varepsilon + \sigma|\mathcal{X}| \cdot |\mathcal{Y}| + m\varepsilon/\sigma^2$ in the above mentioned statement.

Recall that with probability at least $1 - \varepsilon$, the sampled Eve view $V_E^{(i)}$ has the $(1-\varepsilon)$-Independence and the ε-Lightness properties. Henceforth, we shall restrict ourselves to such $V_E^{(i)}$. By Corollary A.4 we conclude that with probability at least $1 - \sigma|\mathcal{X}|$ it holds that $\mathsf{P}[x \mid V_E^{(i)}] \geq \sigma$. If $\mathsf{P}[y'|V_E^{(i)}, x] < \sigma$ for this Eve view $V_E^{(i)}$, then we are done. So, assume on the contrary that $\mathsf{P}[y'|V_E^{(i)}, x] \geq \sigma$, which implies that:

$$\mathsf{P}[x, y' \mid V_E^{(i)}] \geq \mathsf{P}[x \mid V_E^{(i)}] \cdot \mathsf{P}[y' \mid V_E^{(i)}, x] \geq \sigma^2$$

Since $(x, y) \xleftarrow{\$} \mathcal{X} \times \mathcal{Y}$, we can apply Corollary A.4 directly to conclude that with probability at least $1 - \sigma^2|\mathcal{X}||\mathcal{Y}|$, $\mathsf{P}[x, y \mid V_E^{(i)}] \geq \sigma^2$. By union bound, we can assume that all these properties hold with probability $1 - O(\varepsilon + \sigma|\mathcal{X}| \cdot |\mathcal{Y}|)$. Henceforth, we shall assume that $V_E^{(i)}$ satisfies these conditions.

First, using $\mathsf{P}[x, y \mid V_E^{(i)}] \geq \sigma^2$ and by a direct application of Lemma A.8 we can conclude that with probability $1 - O(m\varepsilon/\sigma^2)$, it holds that $Q_A^{(i+1)} \cap Q_B^{(i)} \subseteq \mathcal{Q}(V_E^{(i)})$. Thus, it suffices to show that with probability $1 - O(m\varepsilon/\sigma^2)$, it holds that $Q_A^{(i+1)} \cap Q'^{(i)}_B \subseteq \mathcal{Q}(V_E^{(i)})$, in which case Lemma A.3 would trivially follow by a union bound from these two results.

Lets define $\widetilde{V}_B^{(i)}$ and $\widetilde{Q}_B^{(i)}$ similar to $V_B^{(i)}$ and $Q_B^{(i)}$ with the only difference that we sample them conditioned on the input y'. Then the same exact proof as for the case of likely input y, can be applied to the case of likely input y' and conclude that with probability $1 - O(m\varepsilon/\sigma^2)$, it holds that $Q_A^{(i+1)} \cap \widetilde{Q}_B^{(i)} \subseteq \mathcal{Q}(V_E^{(i)})$. We emphasize that the distributions $\widetilde{V}_B^{(i)}$ and $V'^{(i)}_B$ are not identical. Although, both are sampled based on Bob input being y' and Eve view being $V_E^{(i)}$, the latter is additionally conditioned on the next message μ_{i+1} of Alice. Here, we shall be leveraging Lemma A.7.

So, consider an Eve view $V_E^{(i)}$ with the following properties:

1. $(1-\varepsilon)$-Independence and ε-Lightness properties hold, and
2. $\mathsf{P}[x, y' | V_E^{(i)}, x, y'] \geq \sigma^2$.

Let $(\mathbf{V}_A^{(i+1)}, \widetilde{\mathbf{V}}_B^{(i)})$ represent the joint Alice-Bob views when Alice has input x and Bob has input y'. By Lemma A.8, we know that this distribution is $O(m\varepsilon/\sigma^2)$ close to the distribution $(\mathbf{V}_A^{(i+1)} | V_E^{(i)}, x) \times (\widetilde{\mathbf{V}}_B^{(i)} | V_E^{(i)}, y')$. Let $(\mathbf{V}_A^{(i+1)}, \mathbf{V'}^{(i)}_B)$ represent the joint Alice-Bob views when Alice has input x, Bob has input y' as picked in our experiment, i.e. Bob's view is additionally conditioned on the next message μ_{i+1}. Considering μ_{i+1} as a leakage on $V_A^{(i+1)}$, we can conclude that $(\mathbf{V}_A^{(i+1)}, \mathbf{V'}^{(i)}_B)$ is also $O(m\varepsilon/\sigma^2)$ close to $(\mathbf{V}_A^{(i+1)} | V_E^{(i)}, x) \times (\widetilde{\mathbf{V}}_B^{(i)} | V_E^{(i)}, y')$, by Lemma A.7. Consequently, the distributions $(\mathbf{V}_A^{(i+1)}, \widetilde{\mathbf{V}}_B^{(i)})$ and $(\mathbf{V}_A^{(i+1)}, \mathbf{V'}^{(i)}_B)$ are $O(m\varepsilon/\sigma^2)$ close.

Recall that the probability of the event $Q_A^{(i+1)} \cap \widetilde{Q}_B^{(i)} \subseteq \mathcal{Q}(V_E^{(i)})$ when Alice-Bob views are sampled according to $(\mathbf{V}_A^{(i+1)}, \widetilde{\mathbf{V}}_B^{(i)})$ is $1 - O(m\varepsilon/\sigma^2)$. So, the probability of the same event when Alice-Bob joint views are sampled according to $(\mathbf{V}_A^{(i+1)}, \mathbf{V'}^{(i)}_B)$ is also $1 - O(m\varepsilon/\sigma^2)$. This concludes the proof of Lemma A.3.

A.5.3 Proof of Lemma A.8

Finally we prove Lemma A.8. Recall that with respect to the Eve view $V_E^{(i)}$, $(1-\varepsilon)$-Independence and ε-Lightness hold, when the protocol is run with uniformly chosen $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Consider the space of all Alice and Bob private views and random oracles such that $V_E^{(i)}$ is produced as the view of Eve. We know by Lemma A.1 that the distribution of $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)})$ is ε close to a distribution $(\mathbf{U}_A \times \mathbf{U}_B) \equiv (\mathbf{V}_A^{(i)} | V_E^{(i)}) \times (\mathbf{V}_B^{(i)} | V_E^{(i)})$.

Additionally, we are also given that $\mathsf{P}[x, y | V_E^{(i)}] \geq \gamma$. Now, consider the event E such that x and y are actually the local inputs in sampled Alice and Bob views $V_A^{(i)}$ and $V_B^{(i)}$. By Lemma A.5, we can conclude that $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)}, x, y)$ is ε/γ close to the distribution $(\mathbf{U}_A \times \mathbf{U}_B | x, y) \equiv (\mathbf{U}_A | x) \times (\mathbf{U}_B | y)$. Now, observe that when Alice and Bob views are sampled according to $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)}, x, y)$, then they also satisfy ε/γ-Lightness property. Otherwise we can use the fact that $\mathsf{P}[x, y | V_E^{(i)}] \geq \gamma$ to show that $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)})$

does not satisfy the ε-Lightness property. Now, since the distribution $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)}, x, y)$ has ε/γ-Lightness property and is ε/γ close to the product distribution $(\mathbf{U}_A|x) \times (\mathbf{U}_B|y)$, this implies that the distribution $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)}, x, y)$ satisfies $(1-\varepsilon')$-Independence and ε'-Lightness properties, where $\varepsilon' = \varepsilon/\gamma$. Next, based on these properties, we shall first prove the first part of Lemma A.8 by showing that $\mathsf{P}[\mathcal{Q}(V_A^{(i+1)}) \cap \mathcal{Q}(V_B^{(i)}) \not\subseteq \mathcal{Q}(V_E^{(i)}) | V_E^{(i)}, x, y] \leq O(m\varepsilon')$.

We define several hybrid experiments where the distribution of $V_A^{(i+1)}$ and $V_B^{(i)}$ is defined differently in each of them. We are interested in comparing the probability p_i of the bad event B defined as $\mathcal{Q}(V_A^{(i+1)}) \cap \mathcal{Q}(V_B^{(i)}) \not\subseteq \mathcal{Q}(V_E^{(i)})$ in the game Game_i.

Game_0**:** In this game the views $V_A^{(i+1)}$ and $V_B^{(i)}$ are jointly sampled consistent with $V_E^{(i)}$ and local inputs x and y.

Game_1**:** This game is indeed a perfect lazy simulation of Game_0:

1. Sample $(V_A^{(i)}, V_B^{(i)})$ according to the distribution $(\mathbf{V}_A^{(i)}, \mathbf{V}_B^{(i)} | V_E^{(i)}, x, y)$.
2. Start the next message generation algorithm for Alice. If any query q asked by Alice is already contained in $\mathcal{Q}(V_A^{(i)}) \cup \mathcal{Q}(V_E^{(i)}) \cup \mathcal{Q}(V_B^{(i)})$, then it is consistently answered. Otherwise, a uniformly random answer is provided.

So, the probability p_1 of the bad event $\mathcal{Q}(V_A^{(i+1)}) \cap \mathcal{Q}(V_B^{(i)}) \not\subseteq \mathcal{Q}(V_E^{(i)})$ in this game is still equal to p_0.

Game_2**:** In this game

1. Alice and Bob views are drawn according to $(\mathbf{V}_A^{(i)} | V_E^{(i)}, x) \times (\mathbf{V}_B^{(i)} | V_E^{(i)}, y)$.
2. Start the next message generation algorithm for Alice with respecting the answers to Bob's private queries. Namely, if any query q asked by Alice is already contained in $\mathcal{Q}(V_A^{(i)}) \cup \mathcal{Q}(V_E^{(i)}) \cup \mathcal{Q}(V_B^{(i)})$, then it is consistently answered. Otherwise, an uniformly random answer is provided.

By $(1-\varepsilon')$-Independence we know that Game_1 and Game_2 are ε' close, so $p_1 \leq p_2 + \varepsilon'$.

Now, we shall bound p_2. Recall that ε'-Lightness of $(\mathbf{V}_B^{(i)} | V_E^{(i)}, y)$ implies that any query not already answered in $V_E^{(i)}$ occurs with probability at most ε' in a Bob view $V_B^{(i)} \xleftarrow{\$} (\mathbf{V}_B^{(i)} | V_E^{(i)}, y)$. So, the probability of m new queries of Alice hitting any query of Bob view $V_B^{(i)} \xleftarrow{\$} (\mathbf{V}_B^{(i)} | V_E^{(i)}, y)$ is at most $m\varepsilon'$, by union bound. So, $p_2 \leq m\varepsilon'$. This implies that $p_0 = p_1 \leq p_2 + \varepsilon' \leq (m+1)\varepsilon'$. This completes the proof of the first part of Lemma A.8.

Proving the Second part of Lemma A.8. In our previous hybrids, we showed that the joint distribution of views $(\mathbf{V}_A^{(i+1)}, \mathbf{V}_B^{(i)})$ in Game_0 and Game_2 are ε' far.

Consider the following Game_3 as the next hybrid following Game_2: In this game

1. Alice and Bob views are drawn according to $(\mathbf{V}_A^{(i)} | V_E^{(i)}, x) \times (\mathbf{V}_B^{(i)} | V_E^{(i)}, y)$.

2. Start the next message generation algorithm for Alice without respecting the answers to Bob's private queries. Namely, if any query q asked by Alice is already contained in $\mathcal{Q}(V_A^{(i)}) \cup \mathcal{Q}(V_E^{(i)})$, then it is consistently answered. Otherwise, an uniformly random answer is provided.

If the bad event $\mathcal{Q}(V_A^{(i+1)}) \cap \mathcal{Q}(V_B^{(i)}) \not\subseteq \mathcal{Q}(V_E^{(i)})$ does not occur, then the distribution of Alice-Bob joint views sampled in Game_2 and Game_3 are identical. Further, the distribution of Alice views in Game_3 is identical to $(\mathbf{V}_A^{(i+1)} | V_E^{(i)}, x)$. Note, that by the same argument at in Game_3, the probability p_3 of the bad event is at most $m\varepsilon'$, because the argument was independent of how Alice queries were answered. So, the joint distribution of views $(\mathbf{V}_A^{(i+1)}, \mathbf{V}_B^{(i)})$ in Game_2 and Game_3 are at most $\max\{p_2, p_3\} \leq m\varepsilon'$ far.

Therefore, the statistical distance between $(\mathbf{V}_A^{(i+1)}, \mathbf{V}_B^{(i)} | V_E^{(i)}, x, y)$ and $(\mathbf{V}_A^{(i+1)} | V_E^{(i)}, x) \times (\mathbf{V}_B^{(i)} | V_E, y)$ is at most $(m+1)\varepsilon'$. Thus, the second part of Lemma A.8 follows.

B Some Examples for Intuition

B.1 Undecomposable Functions

We give examples of some representative undecomposable functions in Figure 5, Figure 6 and Figure 7

$$\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$$

Figure 5: A Complete (and Undecomposable) Function.

$$\begin{pmatrix} 1 & 1 & 2 \\ 4 & 0 & 2 \\ 4 & 3 & 3 \end{pmatrix}$$

Figure 6: An Incomplete but Undecomposable Function (Minimum $|\mathcal{X}| + |\mathcal{Y}|$).

$$\begin{pmatrix} 1 & 1 & 3 & 4 \\ 3 & 2 & 2 & 4 \\ 3 & 4 & 1 & 1 \\ 2 & 4 & 3 & 2 \end{pmatrix}$$

Figure 7: An Incomplete but Undecomposable Function (Minimum $|\mathcal{Z}|$).

B.2 Decomposable Example

Let us consider the example of computing maximum of Alice and Bob inputs, where Alice's input set is $\{1, 3, 5\}$ and Bob's input set is $\{0, 2, 4\}$. This function is decomposable and its decomposition provides a perfectly semi-honest secure protocol, see Figure 8. The semi-honest protocol is as follows:

Protocol to compute maximum of Alice and Bob inputs:

1. If Alice's input is 5, then she announces the outcome to be 5; Otherwise she asks Bob to proceed.

2. If Bob's input is 4, then he announced the outcome to be 4; Otherwise he asks Alice to proceed.

3. If Alice's input is 3, then she announces the outcome to be 3; Otherwise she asks Bob to proceed.

4. Now, Alice's input is 1 for certain. If Bob's input is 2, then he announces the outcome to be 2; Otherwise the outcome is 1.

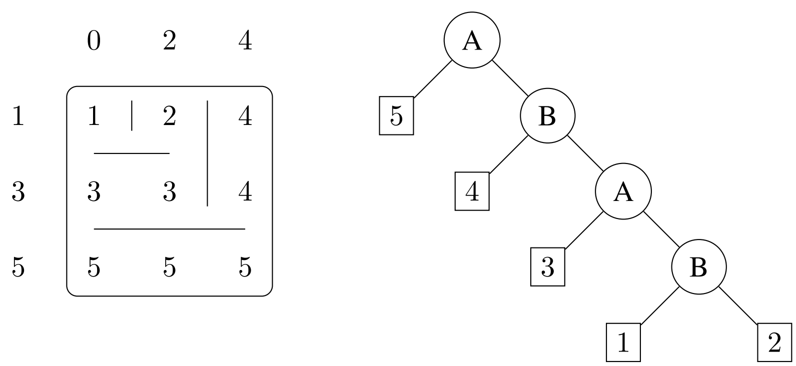

Figure 8: Decomposition of a Decomposable Function

C Black-box Separation Proof

Proof of Theorem 7.2. It is immediate that $(1) \Rightarrow (2) \Rightarrow (3)$. We shall show that $(3) \Rightarrow (1)$. In fact, for clarity, first we shall show $(2) \Rightarrow (1)$ before extending the argument to show $(3) \Rightarrow (1)$.

We rely on the following claim.

Claim C.1. *Let f be a a deterministic two-party function which does not have a perfectly semi-honest secure protocol. For any security-parameter preserving fully black-box construction (Π, S) of SFE_f from OWF, there exist $Q \in F_{\mathsf{OWF}}$ and an oracle algorithm Adv such that $(\Pi^Q, \mathsf{Adv}^Q) \in R_{\mathsf{SFE}_f}$ and $(Q, S^{Q,\mathsf{Adv}^Q}) \notin R_{\mathsf{OWF}}$.*

Before proving this claim, we note that it indeed shows $(2) \Rightarrow (1)$, as follows. Suppose, for the sake of contradiction, (Π, S) is a security-parameter preserving PPT-secure fully black-box construction of SFE_f from OWF, for some deterministic two-party function f which does not have a perfectly semi-honest secure protocol. For (Π, S), let $Q \in F_{\mathsf{OWF}}$ and Adv be as guaranteed in Claim C.1. Let A stand for Adv^Q. The claim guarantees that $(\Pi^Q, A) \in R_{\mathsf{SFE}_f}$. Consequently, by the security guarantee of fully black-box construction, we have $(Q, S^{Q,A}) \in R_{\mathsf{OWF}}$. But this contradicts the guarantee from Claim C.1.

Proof of Claim C.1. Let $\mathcal{U}_\kappa = \{g \colon \{0,1\}^\kappa \to \{0,1\}^\kappa\}$ denote the set of all length preserving functions over $\{0,1\}^\kappa$. Let $\mathcal{U} = \mathcal{U}_1 \times \mathcal{U}_2 \times \cdots$.

Note that since Π is security-parameter preserving, Π^Q accesses only Q_κ. Implicit in the proof of Theorem 1.1 is an adversary Adv such that, for $Q_\kappa \xleftarrow{\$} \mathcal{U}_\kappa$ the adversary Adv^{Q_κ} breaks the security of Π^{Q_κ} (as an implementation of SFE_f) with advantage $\delta(\kappa) > 1/\operatorname{poly}(\kappa)$, by asking $\operatorname{poly}(\kappa)$ queries to Q_κ. This will be the adversary Adv in the statement of the claim.

Next, we need to find a deterministic function Q such that Adv^Q breaks Π^Q, but there does not exist any efficient reduction S such that S^{Q,Adv^Q} breaks Q as a OWF implementation.

We show the existence of such a Q by the probabilistic method. For this, first we define $\mathcal{V}_\kappa \subseteq \mathcal{U}_\kappa$ for each $\kappa \in \mathbb{N}$ as follows. As mentioned above, Adv^{Q_κ} has an advantage of $\delta(\kappa) > 1/\operatorname{poly}(\kappa)$ in breaking Π^{Q_κ}, where $Q_\kappa \xleftarrow{\$} \mathcal{U}_\kappa$. Then, by an averaging argument, for a subset $\mathcal{V}_\kappa \subseteq \mathcal{U}_\kappa$ with $\frac{|\mathcal{V}_\kappa|}{|\mathcal{U}_\kappa|} \geq \delta(\kappa)$, it holds that for all $Q_\kappa \in \mathcal{V}_\kappa$, Adv^{Q_κ} has an advantage at least $\delta(\kappa)/2$ in the SFE_f security game for Π^{Q_κ}. Now, we pick $Q_\kappa \xleftarrow{\$} \mathcal{V}_\kappa$ independently for each security parameter κ. Q will be the composite oracle $(Q_1, Q_2, \dots)$.

By construction, $(\Pi^Q, \mathsf{Adv}^Q) \in R_{\mathsf{SFE}_f}$ with probability 1, since for all $Q_\kappa \in \mathcal{V}_\kappa$, Adv^{Q_κ} has a significant advantage (as a function of κ) in the security game. To complete the proof, we need to show that with positive probability Q is such that $(Q, S^{Q,\mathsf{Adv}^Q}) \notin R_{\mathsf{OWF}}$.

Consider again $Q_\kappa \xleftarrow{\$} \mathcal{U}_\kappa$ (rather than $Q_\kappa \xleftarrow{\$} \mathcal{V}_\kappa$, which we shall return to shortly). For each κ, for each choice of $Q_{\overline{\kappa}} = (Q_1, \dots, Q_{\kappa-1}, Q_{\kappa+1}, \dots)$, define the (inefficient) machine $T_{Q_{\overline{\kappa}}}$ such that $T_{Q_{\overline{\kappa}}}^{Q_\kappa}$ simulates S^{Q,Adv^Q}: for this, $T_{Q_{\overline{\kappa}}}$ internally simulates all of Q except Q_κ, which it accesses through oracle calls. Even though $T_{Q_{\overline{\kappa}}}$ is inefficient, since S is efficient, the number of oracle queries it makes is bounded by $\operatorname{poly}(\kappa)$. W.l.o.g, we can assume that a machine T^{Q_κ} can invert an input y with respect to its oracle, only if one of its oracle queries is answered by y (by adding a final query, in which it queries the oracle at its output). But when $Q_\kappa \xleftarrow{\$} \mathcal{U}_\kappa$ this happens with only negligible probability for a machine making polynomially many queries, because each distinct query is answered by a κ-bit string chosen uniformly at random which has a probability of $\frac{1}{2^\kappa}$ of being equal to y.

Thus, if $Q_\kappa \xleftarrow{\$} \mathcal{U}_\kappa$, then for each choice of $Q_{\overline{\kappa}}$, the probability that S^{Q,Adv^Q} has a non-negligible advantage in breaking Q at κ is $\nu(\kappa)$ for some negligible function ν. Then if $Q_\kappa \xleftarrow{\$} \mathcal{V}_\kappa$, this probability is at most $\nu(\kappa)\frac{|\mathcal{U}_\kappa|}{|\mathcal{V}_\kappa|}$ which is also negligible (since $\frac{|\mathcal{V}_\kappa|}{|\mathcal{U}_\kappa|} \geq \delta(\kappa)/2$).

Then, by a union bound over all $\kappa \geq \kappa_0$ for a sufficiently large value of κ_0, the probability that S^{Q,Adv^Q} has a non-negligible advantage in breaking Q at some $\kappa \geq \kappa_0$ is $\sum_{\kappa=\kappa_0}^{\infty} \nu(\kappa) < 1$ (and can in fact be made arbitrarily close to 0, by choosing κ_0 large enough). In particular, there exists Q such that S^{Q,Adv^Q} does not have a non-negligible advantage in breaking O at infinitely many values of κ. That is, $(Q, S^{Q,\mathsf{Adv}^Q}) \notin R_{\mathsf{OWF}}$. □

Extending to OWF_ζ. The above argument can be easily extended to show (3) $\Rightarrow$ (1), to complete the proof. Fix a polynomial ζ. Then, in the above argument consider the set $\mathcal{W}_\kappa := \{g^\zeta(\kappa, \cdot) | g : \{0,1\}^* \to \{0,1\}^\kappa\}$ (i.e., set of functions that map x, $|x| \leq \zeta(\kappa)$ to $y \in \{0,1\}^\kappa$), instead of $\mathcal{U}_\kappa$. We remark that for the adversary from the proof of Theorem 1.1 it was not crucial that the random oracle has input domain $\{0,1\}^\kappa$, or that the oracle is length-preserving, as long as the queries are answered independent of each other. The rest of the argument, including the fact that an inverter making polynomial queries to an oracle $W_\kappa \xleftarrow{\$} \mathcal{W}_\kappa$ can have only a negligible success probability, remains unchanged.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de