Full Characterization of Completeness for Two-party Randomized Function Evaluation

Hemanta K. Maji [*] Manoj Prabhakaran [†]

November 5, 2012

Abstract

We close a long line of work that has pursued a full characterization of completeness of (possibly randomized) finite functions for 2-party computation that is secure against active adversaries. The first such function was discovered almost quarter of a century back (Kilian, FOCS 1988). Since then the question of which all finite 2-party functions are complete has been studied extensively, leading to characterizations in many special cases. In this work, we answer this problem in full.

The main tools in our solution include:

– a linear-algebraic characterization of *redundancy* in a 2-party function,

– the notion of statistical testability (a version of interactive proofs in the information-theoretic setting) and a result that evaluation of a 2-party function is statistically testable if and only if the function is redundancy free, and

– an extension to the (weak) converse of Shannon's channel coding theorem.

Our main construction also gives a generalization of a line of work on obtaining protocols secure against active adversaries from protocols secure against passive adversaries in a black-box manner.

[*] University of California, Los Angeles. Supported by NSF CI Postdoctoral Fellowship. hmaji@cs.ucla.edu.

[†] University of Illinois, Urbana-Champaign. Supported by NSF grant CNS 07-47027. mmp@illinois.edu.

Contents

1 Introduction

Understanding the *complexity* of functions is central to theoretical computer science. While the most studied notion of complexity in this literature is that of computational complexity, there have also been other important aspects explored, most notably, *communication complexity* [Yao79]. Another aspect of complexity of a (distributed) function is its *cryptographic complexity* which seeks to understand the cryptographic utility of a function, stemming from how it hides and reveals information. While it is only recently that the term has been explicitly used, cryptographic complexity theory has been vigorously pursued at least since Kilian introduced the notion of *completeness* of cryptographic primitives [Kil88].

Completeness has been the first and most important question of cryptographic complexity: what properties of a function let all other cryptographic tasks (in the context of secure multi-party computation) be *reduced* to it. This question has been asked and answered several times [Kil88, CK88, Kil91, Kil00, CMW04, KM11, MPR12] each time for a different class of functions, or restricted to different kinds of reductions. These works produced several exciting ideas and advances, and brought together concepts from different fields. For instance, [Kil00] used the Nash equilibrium in a zero-sum game defined using the function to obtain a secure protocol; earlier [CK88] identified the binary symmetric channel (noisy channel) as a complete function, paving the way to a connection with information-theory literature that has been going strong.

However, these works left open what is arguably the hardest part of the characterization: completeness of general *randomized* functions without any restrictions on which parties have inputs and which parties have outputs, under reductions that are secure against an active adversary. Indeed, even with a (usually simplifying) restriction that only one of the two parties receives an output from the function, it was not known which *randomized* functions are complete. In this work, we finally provide a full characterization of completeness of 2-party functions. This brings to close this rich line of investigation, but also throws out several new ideas and questions regarding cryptographic complexity.

We remark that while this result subsumes all the prior results on this front, it does rely on some of them (in particular results from [Kil00, IPS08, MPR12]).

Also, along the way to our main construction, we generalize a result in another line of work, on black-box constructions [IKLP06, Hai08, CDMW09]. We give a black-box transformation from a passive-secure OT protocol *in a hybrid setting* (wherein the protocol has access to an ideal functionality) to a UC-secure OT protocol in the same hybrid setting, with access to the commitment functionality. This is significant to the cryptographic complexity theory in the computationally bounded setting: it throws light on the question of *which computational assumptions are "distinct" and which ones are not* (in the sense of Impagliazzo's worlds [Imp95]).

Finally, our tools for analysis are novel in this line of work. One of our constructions crucially relies on the converse of Shannon's Channel Coding theorem to obtain a hiding property from a "channel." This is perhaps an unusual (but in hindsight, natural) use of the converse of the channel coding theorem, which was originally used to establish the optimality of the channel coding theorem. Another important contribution is the introduction of *statistically testable games*, which brings the notion of interactive proof systems to the information-theoretic setting. We discuss these in more detail in Section 1.1 and in subsequent sections.

Our Results. We close the line of work that has been pursuing the characterization of general 2-party functions that are complete for secure function evaluation. Our main result can be summarized as follows. As is customary in this line of work, we restrict ourselves to finite functions.

(Below, the "core" of a 2-party function f stands for a "redundancy free" function which is equivalent to f as far as security against active adversaries is concerned; we formally define these notions in Section 3.1 and Section 5).

Theorem 1. *Suppose f is a (possibly randomized) finite 2-party function. Then the following are equivalent.*

1. *f is UC-complete.*
2. *f is standalone-complete.*
3. *f has a core that is passive-complete.*

Our result also completes the following elegant characterization that was conjectured in [MPR12] (using the terminology from [MPR12], which is explained in Section 5).

Theorem 2. *A finite 2-party function is passive-complete if and only if it is not simple. A finite 2-party function is UC-complete (or equivalently, standalone-complete) if and only if it has a core that is not simple.*

Related Work. We briefly summarize the results on completeness from prior work. The function oblivious transfer (OT) was identified independently by Wiesner and Rabin [Rab81, Wie83]. Several years later Kilian identified OT as the first active-complete function [Kil88]. Prior to this Goldreich and Vainish, and independently Micali and Haber, showed that OT is passive-complete [HM86, GV87]. Crépeau and Kilian then showed that the noisy channel is also active-complete [CK88]. The first characterization of completeness appeared in [Kil91] where it was shown that among deterministic "symmetric" functions (in which both parties get the same output) a function f is active-complete if and only if there is an "OR minor" in the matrix representing f. Beimel, Malkin and Micali showed that among "asymmetric" functions (in which only one party gets the output), a function if passive-complete if and only if it is not "trivial" [BMM99]. ([BMM99] also concerned itself with the computational setting and asked cryptographic complexity questions regarding computational assumptions.) Kilian vastly generalized this by giving several completeness characterizations: active-complete deterministic asymmetric functions, passive-complete symmetric functions and passive-complete asymmetric functions [Kil00]. Kilian's result for active-completeness was extended in two different directions by subsequent work: Crépeau, Morozov and Wolf [CMW04] considered "channel functions" which are randomized asymmetric functions (only one party has output), but with the additional restriction that only one party has input; Kraschewski and Müller-Quade [KM11] considered functions in which both parties can have inputs and outputs, but restricted to deterministic functions. Kilian's result for passive-completeness was extended to all functions in a recent work [MPR12], which also presented a unification of all the prior characterizations and posed the question of completing the characterization. The full characterization we obtain matches the unified conjecture from [MPR12].

1.1 Technical Overview

An important ingredient of our result is a combinatorial/linear-algebraic characterization of "redundancy" in a general 2-party function. The importance of redundancy is two fold:

- Any function f is "equivalent" (or *weakly isomorphic*, as defined in [MPR12]) to a "core" function $\widehat{f}$ which is redundancy free, so that f is complete against active adversaries if and only if $\widehat{f}$ is. Thus it is enough to characterize only redundancy free functions.
- Our various protocols rely on being given access to a redundancy free function. Redundancy makes it possible for an adversary to deviate from a prescribed interaction with a function without any chance of being detected. Thus the statistical checks used to enforce that the adversary does not deviate from its behavior crucially rely on the protocol using only redundancy free functions.

While redundancy of special classes of 2-party functions have appeared in the literature previously, it turns out that for general 2-party functions, the nature of redundancy is significantly more intricate. Redundancy, in this case, stands for the possibility that a party could replace its prescribed input to a function by a randomized mixture of inputs, and further, probabilistically alter the output it receives from the function before reporting it, while resulting in the same *distribution* of the view for the rest of the system (environment). Here the deviation from the protocol is defined not by a single probability distribution (as in prior instances of redundancy in the literature) but by a distribution over conditional distributions. The more intricate nature of redundancy raises the possibility that even if a function does not have exact redundancy, it may be arbitrarily close to having redundancy. We provide a non-trivial linear algebraic analysis of redundancy and show that this is not the case. We show that for any finite function, the *irredundancy parameter* — which measures the *ratio* between the extent of difference in the environment's view and the extent of deviation by the adversary — is either equal to 0 (showing that the function has redundancy) or bounded away from 0 by a constant. This proves crucial in being able to detect deviation from honest behavior, when using a redundancy free function.

Our main construction shows that any redundancy free function f which is *passive-complete* (i.e., complete with respect to reductions which are secure against passive adversaries) is also UC-complete. This construction separates into two parts:

- A protocol to UC-securely reduce the commitment functionality $\mathcal{F}_{\text{COM}}$ to f.
- A protocol in the $\mathcal{F}_{\text{COM}}$-hybrid model that UC-securely reduces OT to f, starting from a passive-secure reduction of OT to f (since f is passive-complete, such a protocol exists). That is, we compile (in a black-box manner) a passive-secure OT protocol using f, to a UC-secure OT protocol using f (and $\mathcal{F}_{\text{COM}}$).

In building the commitment functionality we rely on a well-known result from information theory, namely the (weak) *converse of Shannon's Channel Coding Theorem*, extended to the case of adaptively chosen channel characteristics. We also rely on statistical testability of redundancy free functions (see below) to enforce binding, and further on a weak statistical test that uses a carefully chosen input distribution to ensure that a malicious receiver, even if it may deviate without being detected, does not completely avoid the hiding property of f.

The second part, which gives a compiler, is similar in spirit to protocols that established that a passive-secure OT protocol (in the plain model) can be converted to an active-secure OT protocol *in a black-box manner* [IKLP06, Hai08, CDMW09]. In particular, it resembles the protocol in [CDMW09]. However, the key contribution in our protocol compared to these earlier protocols is that the passive-secure OT protocol that we are given is not in the plain model (as was in the case of prior work, which restricted itself to the setting with computational assumptions), but is in the f-hybrid model. The technical difficulty in our case is in ensuring that a cut-and-choose technique can be used to (mostly) verify an adversary's claims about what inputs it sent to a 2-party function and what outputs it received, when the verifier has access to only the other end of the function.

We mention that in contrast with prior work, we do not use "OT reversal" [WW06] (and two uses of the compiler) to obtain security against active corruption of the receiver and then that of the sender. Instead, we directly obtain a somewhat good OT protocol that is secure against active corruption of either player, and use the OT combiner from [IPS08] to obtain the final protocol.

Statistically Testable Games. We introduce a formal notion of statistically testable game: two parties interact with a (possibly probabilistic) system, providing local inputs and obtaining local outputs. After repeating this for a large number of times, one party must declare and prove to the other what sequence of (input, output) pairs it obtained from the execution (or more generally, a function of this sequence). The game is said to be statistically testable if there is a sound and complete proof system which ensures that the sequence declared by the prover has a $o(1)$ fraction hamming distance from the actual sequence (or more generally, some other measure of distance). We point out that our notion of a proof here is information-theoretic (unlike the traditional notion of an interactive proof system, which is meaningless in the information theoretic setting). The verifier's uncertainty arises from having obtained only partial information about the prover's view in the game.

The game we consider is of 2-party function evaluation. A major technical ingredient used both in our commitment protocol and in our compiler is the following theorem: *evaluation of a 2-party function is statistically testable if and only if the function is redundancy free.*

To illustrate the power of statistical testing, consider the (simpler) question of using an *asymmetric* function f (in which only one party, say Bob, obtains an output) to obtain a commitment functionality (we only discuss binding here). A natural approach would be to have the sender in the commitment protocol play the role of Alice in interacting with the function f; Alice's inputs to f can encode her message, and Bob will have some uncertainty about it (which can then be amplified). Since Alice does not obtain any output from f, she has no idea what Bob received and later it would be hard for her to claim to have fed a different set of inputs to f. *However, it is much less intuitive to have a commitment protocol in which the sender plays the role of Bob and the receiver play the role of Alice.* Here, the sender would learn information about receiver's input to f and the receiver will not learn what the sender learned. This makes it, intuitively, much more likely for the sender to be able to equivocate. What statistical testability provides us with is a means for Alice to catch Bob if he lies about his inputs to f, even though all she obtained from the interaction with f was her own inputs (as long as f is redundancy free).

We use statistical testing also in our compiler to ensure that a party cannot freely lie about how it behaved in an execution of a passive-secure protocol: just from the view at one of the function f, the other party can detect lying (up to small errors).

2 Preliminaries

Matrix Definitions. In the following we shall refer to the following matrix norms: $\|A\|_\infty = \max_i \sum_j |a_{ij}|$ (maximum absolute row sum norm), and $\|A\|_{\text{sum}} = \sum_{i,j} |a_{ij}|$ (absolute sum norm). We shall also use the function $\max(A) = \max_{i,j} a_{ij}$ (maximum value among all entries); not that here we do not consider the absolute value of the entries in A. For a probability distribution $\mathfrak{p}^X$over a space X (denoted as vectors), we define $\min(\mathfrak{p}^X) = \min_{x \in X} \mathfrak{p}^X[x]$, the minimum probability it assigns to an element in X. The norm $\|\cdot\|_\infty$ when applied to a column vector simply equals the largest absolute value entry in the vector. We say that a matrix P is a *probability matrix* if its entries are all in the range $[0,1]$ and $\|P\|_{\text{sum}} = 1$. We say that a matrix is a *stochastic matrix* (or row-stochastic matrix) if all its entries are in the range $[0,1]$ and every row sums up to 1. For convenience, we define the notation $D(M)$ for a square matrix M to be the diagonal matrix derived from M by replacing all non-diagonal entries by 0.

2-Party Secure Function Evaluation. A two-party randomized function (also called a secure function evaluation (SFE) functionality) is specified by a single randomized function denoted as $f : X \times Y \to W \times Z$. Despite the notation, the range of f is, more accurately, the space of probability distributions over $W \times Z$. The functionality takes an input $x \in X$ from Alice and an input $y \in Y$ from Bob and samples $(w,z) \in W \times Z$ according to the distribution $f(x,y)$; then it delivers w to Alice and z to Bob. Through out, we shall denote the probability of outputs being (w,z) when Alice and Bob use inputs x and y respectively is represented by $\mathfrak{p}^f[w,z|x,y]$. We use the following variables for the sizes of the sets W, X, Y, Z:

$$|X| = m \qquad |Y| = n \qquad |W| = q \qquad |Z| = r.$$

In this paper we shall restrict to function evaluations where m, n, q and r are constants, i.e. as the security parameter increases the domains do not expand. (But the efficiency and security of our reductions are only polynomially dependent on m, n, q, r, so one could let them grow polynomially with the security parameter. We have made no attempt to optimize this dependency.) W.l.o.g., we shall assume that $X = [m]$ (i.e., the set of first m positive integers), $Y = [n]$, $W = [q]$ and $Z = [r]$.

We consider standard security notions in the information-theoretic setting: UC-security, standalone-security and passive-security against computationally unbounded adversaries (and with computationally unbounded simulators).

Complete Functionalities. A two-party randomized function evaluation f is *standalone-complete* (respectively, *UC-complete*) against information theoretic adversaries if any functionality g can be standalone securely (respectively, UC securely) computed in f hybrid. We shall also consider passive-complete functions where we consider security against passive (semi-honest) adversaries.

3 Main Tools

In this section we introduce the main tools used in our construction.

3.1 Characterizing Irredundancy

Redundancy in a function allows at least one party to deviate in its behavior in the ideal world and not be detected (with significant probability) by an environment. In our protocol, which are designed to detect deviation, it is important to use a function in a form in which redundancy has been removed. We define irredundancy in an explicit linear algebraic fashion, and introduce a parameter to measure the extent of irredundancy.

Irredundancy of a System of Stochastic Matrices. Let P_i, $i = 1, \ldots, m$ be a collection of $s \times q$ probability matrices (i.e., entries in the range $[0,1]$, with $\|P_i\|_{\text{sum}} = 1$). Consider tuples of the form $(j, \{M_i, \alpha_i\}_{i=1}^{n})$, where $j \in [n]$, M_i are $q \times q$ stochastic matrices, and $\alpha_i \in [0,1]$ are such that $\sum_i \alpha_i = 1$. Then we define the irredundancy of this system as

$$\mathfrak{D}(P_1, \ldots, P_m) = \inf_{(j, \{\alpha_i, M_i\}_{i=1}^m)} \frac{\|(\sum_{i=1}^m \alpha_i P_i M_i) - P_j\|_\infty}{1 - \alpha_j \|P_j \cdot D(M_j)\|_{\text{sum}}}$$

where the infimum is over tuples of the above form. (Recall that $D(M_j)$ refers to the diagonal matrix with the diagonal entries of M_j.)

Intuitively, consider the rows of P_i to be probability distributions over a q-ary alphabet produced as the outcome of a process with the row index corresponding to a hidden part of the outcome, and the column index being an observable outcome. Then, irredundancy measures how well a P_j can (or rather, cannot) be approximated by a convex combination of all the matrices P_i, possibly with the observable outcome transformed using a stochastic matrix (corresponding to a probabilistic mapping of the observable outcomes); the denominator normalizes the approximability by how much overall *deviation* (probability of changing the process or changing the outcome) is involved. This excludes the trivial possibility of perfectly matching P_j by employing zero deviation (i.e., taking $\alpha_j = 1$ and $M_j = I$).

Irredundancy of a 2-Party Secure Function Evaluation Function. Recall that a 2-party SFE function f with input domains, $X \times Y$ and output domain $W \times Z$ is defined by probabilities $\mathfrak{p}^f[w, z|x, y]$. We define left and right redundancy of f as follows. Below, $|X| = m, |Y| = n, |W| = q, |Z| = r$.

To define left-redundancy, consider representing f by the matrices $\{P^x\}_{x \in X}$ where each P^x is an $nr \times q$ matrix with $P^x_{(y,z),w} = \mathfrak{p}^f[w, y, z|x]$. Here, $\mathfrak{p}^f[w, y, z|x] \triangleq \frac{1}{n}\mathfrak{p}^f[w, z|x, y]$ (where we pick y independent of x, with uniform probability $\mathfrak{p}^f[y|x] = \frac{1}{n}$).

Definition 1. *For an SFE function $f : X \times Y \to W \times Z$, represented by matrices $\{P^x\}_{x \in X}$, with $P^x_{(y,z),w} = \Pr[w, y, z|x]$, we say that an input $\hat{x} \in X$ is* left-redundant *if there is a set $\{(\alpha_x, M_x)|x \in X\}$, where $0 \leq \alpha_x \leq 1$ with $\sum_x \alpha_x = 1$, and each M_x is a $q \times q$ stochastic matrix such that if $\alpha_{\hat{x}} = 1$ then $M_{\hat{x}} \neq I$, and $P^{\hat{x}} = \sum_{x \in X} \alpha_x P^x M_x$.*

We say $\hat{x}$ is strictly left-redundant *if it is left-redundant as above, but $\alpha_{\hat{x}} = 0$. We say $\hat{x}$ is* self left-redundant *if it is left-redundant as above, but $\alpha_{\hat{x}} = 1$ (and hence $M_{\hat{x}} \neq I$).*

f is said to be left-redundancy free if there is no $x \in X$ that is left-redundant.

Right-redundancy notions are defined analogously. f is said to be *redundancy-free* if it is left-redundancy free and right-redundancy free.

Lemma 1. *For an SFE function f, if $\hat{x}$ is left-redundant, then it is either strictly left-redundant or self left-redundant.*

Proof. Suppose $\hat{x}$ is left-redundant. Then $P^{\hat{x}} = \sum_{x \in X} \alpha_x P^x M_x$ as in the definition of left-redundancy. If $\alpha_{\hat{x}} = 1$, then by definition it is self left-redundant. If $\alpha_{\hat{x}} < 1$, we shall show that it is strictly left-redundant. We can write $P^{\hat{x}}(I - \alpha_{\hat{x}} M_x) = \sum_{x \neq \hat{x}} \alpha_x P^x M_x$. To rewrite $P^{\hat{x}}$ as required by strict left-redundancy we depend on the following the observation.

Claim 1. *If M is a $q \times q$ stochastic matrix and $0 \leq \alpha < 1$, then $I - \alpha M$ is invertible and $(1-\alpha)(I - \alpha M)^{-1}$ is a stochastic matrix.*

Proof. Consider the series $D = I + \alpha M + \alpha^2 M^2 + \cdots$. Since $|\alpha| < 1$ and M is stochastic (and in particular, $\|\alpha M\|_\infty < 1$), this series converges, and then since, $D = I + \alpha M \cdot D$, we have $(I - \alpha M)D = I$. Further, $D \cdot \mathbf{1}^T = \frac{1}{1-\alpha} \cdot \mathbf{1}^T$ (where $\mathbf{1}$ is the row matrix of all 1's), because M^t is stochastic for all t and $\alpha^t M^t \cdot \mathbf{1}^T = \alpha^t \cdot \mathbf{1}^T$. □

Using the above claim, let $M = (1 - \alpha_{\hat{x}})(I - \alpha_{\hat{x}} M_{\hat{x}})^{-1}$ be a stochastic matrix. Then

$$\begin{aligned} P^{\hat{x}} &= \frac{1}{1-\alpha_{\hat{x}}} \cdot P^{\hat{x}} \cdot (I - \alpha_{\hat{x}} M_x) \cdot M = \frac{1}{1-\alpha_{\hat{x}}} \cdot \sum_{x \neq \hat{x}} \alpha_x P^x \cdot M_x M \\ &= \sum_{x \neq \hat{x}} \alpha'_x \cdot P^x \cdot M'_x \end{aligned}$$

where $\alpha'_x = \frac{\alpha_x}{1-\alpha_{\hat{x}}}$ (except $\alpha'_{\hat{x}} = 0$) and $M'_x = M_x \cdot M$ satisfy the conditions required for strict left-redundancy. □

We shall see that for an SFE function $f : X \times Y \to W \times Z$ defined by the probability matrices $\{P^x\}_{x \in X}$, if there is a self left-redundant input $\hat{x} \in X$, then $P^{\hat{x}}$ has two columns (neither all zero) which are scalar multiples of each other. (Similarly, if $\hat{y} \in Y$ is a self right-redundant input, then there must be two columns in $P^{\hat{y}}$ that are scalar multiples of each other.) In fact, we shall show the following quantitative version, which shows that if no two columns of $P^{\hat{x}}$ are close to being scalar multiples of each other, then $\hat{x}$ is not close to being self left-redundant.

Claim 2. *Suppose M is a $q \times q$ stochastic matrix such that $\max(M - I) \geq \delta > 0$. Also, suppose P is an $s \times q$ matrix such that for any two columns c_i and c_j of P, $\inf_\gamma \|c_i - \gamma c_j\|_\infty \geq \epsilon > 0$. Then $\|PM - P\|_\infty \geq \delta\epsilon$.*

Proof. Firstly, since we require that $\inf_\gamma \|c_i - \gamma c_j\|_\infty \geq \epsilon$, $\|c_i\|_\infty \geq \epsilon$, for any column c_i in P. Also note that $M \neq I$ since $\max(M - I) > 0$. We need to establish a lowerbound on $\|PN\|_\infty$, where $N = M - I$.

For this, we prove the following by induction, for all integers $t \geq 1$. Suppose N is a $q \times q$ matrix with t non-zero rows, such that all diagonal entries of N are at most 0 and all non-diagonal entries of N are at least 0; also every row of N sums up to 0. Then $\|PN\|_\infty \geq \delta\epsilon$ (where P is as given).

Base case: $t = 1$. Consider (i, j) such that $N_{ij} = \max(N)$. Since there is only one non-zero row, the j^{th} column in N has this as the only non-zero entry. Hence the j^{th} column in PN is $\max(N) \cdot c_i$. Since $\max(N) \geq \delta$, $\|\max(N) \cdot c_i\|_\infty \geq \delta\|c_i\|_\infty \geq \delta\epsilon$. Hence, $\|PN\|_\infty \geq \delta\epsilon$.

Base case: $t = 2$. Again, consider (i, j) such that $N_{ij} = \max(N)$. If this is the only non-zero entry in the j^{th} column in N, then the same analysis as before holds. Otherwise, there is one more non-zero entry, say $N_{i'j}$ in that column. Then the j^{th} column in PN equals $N_{ij}c_i + N_{i'j}c_j = N_{ij}(c_i + \gamma c_j)$ where $\gamma = N_{i'j}/N_{ij}$. Hence, $\|PN\|_\infty \geq N_{ij}\|c_i + \gamma c_j\|_\infty \geq \delta\epsilon$.

Induction step. Suppose N has $t \geq 3$ non-zero rows. We shall construct N' with non-negative non-dagonal entries, with each row summing up to 0, with $t' < t$ non-zero rows and with $\max(N') \geq \max(N)$, such that $\|PN\|_\infty \geq \|PN'\|_\infty$. Then by the induction hypothesis, it follows that $\|PN'\|_\infty \geq \delta\epsilon$.

To construct N' from N, consider (i, j) such that $N_{ij} = \max(N)$. Let k be a non-zero row in N such that $k \neq i$ and $k \neq j$. (This is possible since $t' > 2$.) For $j' \neq k$ and all i', we set $N'_{i'j'} = N_{i'j'} - \frac{N_{kj'}}{N_{kk}}N_{i'k}$. Also we set $N'_{i'k} = 0$ for all i'. This zeroes out the k^{th} row and k^{th} column of N'. Note that $N_{kk} < 0$ and for $i' \neq k$ and $j' \neq k$ we have $N_{kj'}, N_{i'k} \geq 0$; so $N'_{i'j'} \geq N_{i'j'}$ for all elements except those in the k^{th} row or column (which are 0 in N'). In particular, $N'_{ij} \geq N_{ij}$. So $\max(N') \geq \max(N)$.

We can write $N' = NT$, where T is a $q \times q$ matrix defined as follows:

$$T_{i'j'} = \begin{cases} 1 & \text{if } i' = j' \neq k. \\ -\frac{N_{kj'}}{N_{kk}} & \text{if } i' = k \neq j'. \\ 0 & \text{otherwise.} \end{cases}$$

Hence $\|PN'\|_\infty = \|PNT\|_\infty \leq \|PN\|_\infty\|T\|_\infty$ by the sub-mutiplicativity of the $\|\cdot\|_\infty$ norm. But $\|T\|_\infty = 1$ (since the k^{th} row has positive numbers that sum up to 1, and the other rows have a single non-zero entry equal to 1). Thus, $\|PN'\|_\infty \leq \|PN\|_\infty$ as required. □

Lemma 2. *Suppose a 2-party function $f : X \times Y \to W \times Z$ is left redundancy free. Let $\mathfrak{p}^Y$ be a probability distribution over Y. Let the probability matrices $\{P^x\}_{x \in X}$, be defined by $P^x_{(y,z),w} = \mathfrak{p}^f[w, z|x, y]\mathfrak{p}^Y[y]$. Then there is a constant $\epsilon_f > 0$ (depending only on f) such that $\mathfrak{D}(P^1, \cdots, P^m) \geq \epsilon_f \min(\mathfrak{p}^Y)$.*

Proof. Consider any tuple $(\hat{x}, \{M_i, \alpha_i\}_{i=1}^n)$ as in the definition of irredundancy such that it is not the case that $\alpha_{\hat{x}} = 1$ and $M_{\hat{x}} = I$ (so that the denominator is non-zero). We need to show that such tuples cannot achieve arbitrarily low values of the irredundancy parameter.

We consider two cases: firstly, if $\alpha_{\hat{x}} = 1$ (but $M_{\hat{x}} \neq I$) and secondly, if $\alpha_{\hat{x}} < 1$, and give a lower-bound in both cases. Below, we define the matrices Q^x to be similar to P^x but with using the uniform distribution over y rather than the distribution $\mathfrak{p}^Y$. That is, $Q^x_{(y,z),w} = \mathfrak{p}^f[w, z|x, y] \cdot \frac{1}{n}$.

For the case of $\alpha_{\hat{x}} = 1$, $M_{\hat{x}} \neq I$, firstly note that if the denominator in the irredundancy parameter (i.e., the probability of deviation) is $1 - \|P^{\hat{x}} \cdot D(M)\|_{\text{sum}} = \delta_0 > 0$, then there must be $w \in W$ such that $(M_{\hat{x}})_{ww} \leq (1 - \delta_0)$. Hence there must be $w' \neq w$ such that $(M_{\hat{x}})_{ww'} \geq \frac{\delta_0}{q}$. Also, note that since $\hat{x}$ is not a self left-redundant input, for any two non-zero columns c_i and c_j of $Q^{\hat{x}}$, it is not the case that c_i is proportional to c_j. That is, the point $c_i \in \mathbb{R}^q$ lies outside the line through origin and c_j. Note that the matrix $Q^{\hat{x}}$ depends only of f, and so the infimum, $\inf_{i,j,\gamma} \|c_i - \gamma c_j\|_\infty$ is lowerbounded by some constant ϵ that depends only on f. Then by considering P in Claim 2 to be $Q^{\hat{x}}$ without the zero columns, we have $\|Q^{\hat{x}}M - Q^{\hat{x}}\|_\infty \geq \frac{\delta_0}{q}\epsilon$. Since each row of $P^{\hat{x}}$ is obtained by multiplying a row of $Q^{\hat{x}}$ by $n\mathfrak{p}^Y[y]$ for some y, $\|P^{\hat{x}}M - P^{\hat{x}}\|_\infty \geq n \cdot \min(\mathfrak{p}^Y) \cdot \frac{\delta_0}{q}\epsilon$, and hence the ratio $\frac{\|P^{\hat{x}}M - P^{\hat{x}}\|_\infty}{1 - \|P^{\hat{x}} \cdot D(M)\|_{\text{sum}}} \geq nq\epsilon \cdot \min(\mathfrak{p}^Y)$.

For the case when $\alpha_{\hat{x}} < 1$, first we use the argument from Lemma 1, to write the matrix $K := \sum_x \alpha_x P^x M_x - P^{\hat{x}}$ as $\frac{1}{1-\alpha^{\hat{x}}} K \cdot M = (\sum_{x \neq \hat{x}} \alpha'_x P^x M'_x) - P^{\hat{x}}$, where $\sum_{x \neq \hat{x}} \alpha'_x = 1$ and M'_x are stochastic. Now, we note that the set of points $\{(\sum_{x \neq \hat{x}} \alpha'_x P^x M'_x) | \alpha'_x \geq 0, \sum_{x \neq \hat{x}} \alpha'_x = 1 \text{ and } M'_x \text{ stochastic}\}$ is a *closed* region (with $nr \times q$ matrices treated as points in $\mathbb{R}^{nrq}$). Since $P^{\hat{x}}$ is not in this region, there is a lowerbound on the distance between $P^{\hat{x}}$ and this region (under various norms, including the $\|\cdot\|_\infty$ norm we use in the numerator of the irredundancy parameter). Note that the denominator is at most 1. Hence for this setting of $\{M_x, \alpha_x\}_{x \in X}$ we again have the irredundancy parameter lowerbounded by a positive constant that depends only on f.

□

3.2 Statistically Testable Function Evaluation

In this section we consider the notion of a statistically testable function evaluation game. (The notion is more general and could be extended to reactive systems, or multi-player settings; for simplicity we define it only for the relevant setting of 2-party functions.) We informally defined a statistical test in Section 1.1. As mentioned there, we shall show that *evaluation of a 2-party function is statistically testable if and only if the function is redundancy free*. For simplicity, we define a particular test and show that it is sound and complete for redundancy free functions (without formally defining statistical tests in general). (It is easy to see that functions with redundancy cannot have a sound and complete test. Since this is not relevant to our proof, we omit the details.)

Let f be redundancy free. Consider the following statistical test, formulated as a game between an honest challenger (verifier) and an adversary (prover) in the f-hybrid.

Left-Statistical-Test$(f, \mathfrak{p}^Y; N)$:

1. The adversary picks $\tilde{\mathbf{x}} = (\tilde{x}_1, \ldots, \tilde{x}_N) \in X^N$, and for each $i \in [N]$ the challenger (secretly) picks uniform i.i.d $y_i \in Y$, according to the distribution $\mathfrak{p}^Y$.
2. For each $i \in [N]$, the adversary picks $x_i \in X$ (possibly adaptively), and they invoke f with inputs x_i and y_i respectively; the adversary receives w_i and the challenger receives z_i, where $(w_i, z_i) \overset{\$}{\leftarrow} f(x_i, y_i)$.
3. The adversary then outputs $\tilde{\mathbf{w}} = (\tilde{w}_1, \ldots, \tilde{w}_N) \in W^N$.

The adversary wins this game (breaks the soundness) if the following conditions hold:

1. Consistency: Let $\mu_{\tilde{w},\tilde{x},y,z}$ be the number of indices $i \in [N]$ such that $\tilde{w}_i = \tilde{w}, \tilde{x}_i = \tilde{x}$, $y_i = y$ and $z_i = z$. Also, let $\mu_{\tilde{x},y}$ be the number of indices $i \in [N]$ such that $\tilde{x}_i = \tilde{x}$ and $y_i = y$. The

consistency condition requires that $\forall (w,x,y,z) \in W \times X \times Y \times Z$,

$$\mu_{\tilde{w},\tilde{x},y,z} = \mu_{\tilde{x},y} \times \mathfrak{p}^f[\tilde{w},z|\tilde{x},y] \pm N^{2/3}.$$

2. Separation: Let vectors $\mathbf{A}, \tilde{\mathbf{A}} \in (W \times X)^N$ be defined by $A_i := (w_i, x_i)$ and $\tilde{A}_i = (\tilde{w}_i, \tilde{x}_i)$. The separation condition requires that the hamming distance between the vectors $\mathbf{A}$ and $\tilde{\mathbf{A}}$ is $\Delta(\mathbf{A}, \tilde{\mathbf{A}}) \geq N^{7/8}$.

The *Right-Statistical-Test*$(f, \mathfrak{p}^X; N)$ is defined analogously. The experiment *Statistical-Test*$(f, \mathfrak{p}^X, \mathfrak{p}^Y; N)$ consists of Left-Statistical-Test$(f, \mathfrak{p}^Y; N)$ and Right-Statistical-Test$(f, \mathfrak{p}^X; N)$, and the adversary wins if it wins in either experiment.

Before proceeding, we note that the above statistical test is indeed "complete": if the prover plays "honestly" and uses $\tilde{\mathbf{x}} = \mathbf{x}$ and $\tilde{\mathbf{w}} = \mathbf{w}$, then the consistency condition will be satisfied with all but negligible probability (for any choice of $\mathbf{x}$).

Lemma 3. *If f is redundancy free, and $\mathfrak{p}^X$ and $\mathfrak{p}^Y$ have full support over X and Y respectively, then the probability that any adversary wins in Statistical-Test$(f, \mathfrak{p}^Y, \mathfrak{p}^X; N)$ is* $\mathrm{negl}(N)$.

Proof. We shall only argue that if f is left-redundancy free, then the probability of any adversary winning the Left-Statistical-Test$(f, \mathfrak{p}^Y; N)$ is negligible in N. The argument for the Right-Statistical-Test is similar. Then the result follows by union bound.

The experiment involves the adversary adaptively choosing x_i. To facilitate the analysis, instead we shall analyze *all* choices of $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w}, \tilde{\mathbf{w}})$, but restricted to $\mathbf{w}$ being "typical" for a randomly chosen $\mathbf{y}$ (for the given vector $\mathbf{x}$). Since this would hold except with negligible probability (over random choice of $\mathbf{y}$ and the randomness of f), this restriction will not affect the conclusion. Then, assuming that the adversary satisfies the sufficient-distance condition, we analyze the probability of the consistency condition holding. We shall argue that this probability is negligible if f is redundancy free.

We shall consider the expectation of the quantity $\mu_{\tilde{w},\tilde{x},y,z} - \mathfrak{p}^f[\tilde{w},z|\tilde{x},y]\mu_{\tilde{x},y}$ and argue that for some value of $x, \tilde{y}, \tilde{z}$, the absolute value of this expectation should be large, say, $\Omega(N^{7/8})$. Note that, once we fix $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w}, \tilde{\mathbf{w}})$, then for any quadruple $(\tilde{x}, x, w, \tilde{w})$, $\mu_{\tilde{w},\tilde{x},y,z}$ and $\mu_{\tilde{x},y}$ can both be written as the sum of i.i.d indicator random variables. This is because the random experiment we consider consists only of picking y_i, z_i, for each i independently: if $x_i = x$ and $w_i = w$, then $\Pr[y_i = y, z_i = z] = \mathfrak{p}^{f,Y}[y,z|x,w] := \frac{\mathfrak{p}^Y[y]\cdot\mathfrak{p}^f[w,z|x,y]}{\sum_{z',y'}\mathfrak{p}^Y[y']\cdot\mathfrak{p}^f[w,z'|x,y']}$. Then by Chernoff bounds, we obtain that except with negligible probability, the consistency condition will be violated.

We shall define the set Good of "good" $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w})$ in which, for each $\tilde{x}, x, w$, the number of positions i with $w_i = w$ among the positions i with $\tilde{x}_i = \tilde{x}, x_i = x$ is as expected (over uniformly random i.i.d y_i and randomness of f) up to an additive error of $N^{2/3}$. (Note that this assumption is non-trivial only when there are at least $N^{2/3}$ positions with $\tilde{x}_i = \tilde{x}, x_i = x$.) The analysis below would be for every tuple $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w}) \in \mathsf{Good}$. W.l.o.g we assume that for each $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w})$ the adversary chooses $\tilde{\mathbf{w}}$ deterministically.

Fix $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w}) \in \mathsf{Good}$ and an arbitrary $\tilde{\mathbf{w}}$. Let $\tilde{I}_{\tilde{x}\tilde{w}}$ denote the subset of indices $i \in [N]$ such that $(\tilde{x}_i, \tilde{w}_i) = (\tilde{x}, \tilde{w})$, and $I_{y,z}$ denote the set of i such that $(y_i, z_i) = (y, z)$. We also write $\tilde{I}_{\tilde{x}}$ to denote the set of all indices i with $\tilde{x}_i = \tilde{x}$.

Let $\tilde{J}_{\tilde{x}} = \tilde{I}_{\tilde{x}} \setminus (\cup_{w \in W} \tilde{I}_{\tilde{x},w} \cap I_{\tilde{x},w})$. By the sufficient-distance condition, we know that there is some value $\hat{x} \in X$ such that $|\tilde{J}_{\hat{x}}| \geq \frac{1}{m} N^{7/8}$. Henceforth, we restrict our attention to $\tilde{I}_{\hat{x}}$.

The probabilities in the expressions below are conditioned on $(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w})$, where the random choices made are of $\mathbf{y}$ and $(\mathbf{w}, \mathbf{z})$. (We do not assume any distribution over $\tilde{\mathbf{x}}$ and $\mathbf{x}$ which are chosen by the adversary.) For any $y \in Y$, we have:

$$\begin{aligned}
\mathrm{E}\left[\mu_{\tilde{w},\hat{x},y,z}\right] = \mathrm{E}\left[|\tilde{I}_{\hat{x},\tilde{w}} \cap I_{y,z}|\right] &= \sum_{\substack{x \in X, w \in W \\ \mathfrak{p}^{f,Y}[w|x] > 0}} \mathrm{E}\left[|\tilde{I}_{\hat{x}\tilde{w}} \cap I_{w,x,y,z}|\right] \\
&= \sum_{x,w} |\tilde{I}_{\hat{x}\tilde{w}} \cap I_{xw}| \cdot \mathfrak{p}^{f,Y}[y,z|x,w] = \sum_{x,w} |\tilde{I}_{\hat{x}\tilde{w}} \cap I_{xw}| \cdot \frac{\mathfrak{p}^{f,Y}[w,y,z|x]}{\mathfrak{p}^{f,Y}[w|x]}
\end{aligned}$$

Here, $\mathfrak{p}^{f,Y}[w,y,z|x] \triangleq \mathfrak{p}^{Y}[y]\mathfrak{p}^{f}[w,z|x,y]$ (since we pick y independent of x, with probability $\mathfrak{p}^{Y}[y|x] = \mathfrak{p}^{Y}[y]$) and $\mathfrak{p}^{f,Y}[w|x] = \sum_{y,z} \mathfrak{p}^{f,Y}[w,y,z|x]$. Also, we define $\beta^{x}_{w\tilde{w}}$ to be the fraction of indices $i \in \tilde{I}_{\hat{x}}$ in which the adversary obtained $w_i = w$, for which it reported $\tilde{w}_i = \tilde{w}$.

$$\begin{aligned}
\beta^{x}_{w\tilde{w}} &= \begin{cases} \frac{|\tilde{I}_{\hat{x}\tilde{w}} \cap I_{xw}|}{|\tilde{I}_{\hat{x}} \cap I_{xw}|} & \text{if } |\tilde{I}_{\hat{x}} \cap I_{xw}| \neq 0 \\ 0 & \text{otherwise.} \end{cases} \\
|\tilde{I}_{\hat{x}\tilde{w}} \cap I_{xw}| &= |\tilde{I}_{\hat{x}} \cap I_{xw}| \cdot \beta^{x}_{w\tilde{w}} && \text{by definition of } \beta^{x}_{w\tilde{w}} \qquad (1) \\
&= \left(|\tilde{I}_{\hat{x}} \cap I_{x}| \cdot \mathfrak{p}^{f,Y}[w|x] \pm N^{2/3}\right) \cdot \beta^{x}_{w\tilde{w}} && \text{since } (\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w}) \in \mathsf{Good}. \qquad (2)
\end{aligned}$$

We substitute this into the above expression for $\mathrm{E}\left[\mu_{\tilde{w},\hat{x},y,z}\right]$. Note that $\mathfrak{p}^{f,Y}[w|x] > 0$ implies that it is lower-bounded by a positive constant (depending on f, independent of N), and so $\frac{N^{2/3}}{\mathfrak{p}^{f,Y}[w|x]} = O(N^{2/3})$. Thus,

$$\begin{aligned}
\mathrm{E}\left[\mu_{\tilde{w},\hat{x},y,z}\right] &= \sum_{x,w} |\tilde{I}_{\hat{x}\tilde{w}} \cap I_{xw}| \cdot \mathfrak{p}^{f,Y}[w,y,z|x] \cdot \beta^{x}_{w\tilde{w}} \pm O(N^{2/3}) \\
&= |\tilde{I}_{\hat{x}}| \cdot \sum_{x} \alpha^{x} \left(P^{x} \cdot B^{x}\right)_{(y,z),\tilde{w}} \pm O(N^{2/3})
\end{aligned}$$

where $\alpha^{x} = \frac{|\tilde{I}_{\hat{x}} \cap I_{x}|}{|\tilde{I}_{\hat{x}}|}$, P^{x} is an $nr \times q$ matrix with $P^{x}_{(y,z),w} = \mathfrak{p}^{f,Y}[w,y,z|x]$ and B^{x} is a $q \times q$ matrix with $B^{x}_{w\tilde{w}} = \beta^{x}_{w\tilde{w}}$. Note that the sum of all the entries in P^{x} is 1; also, $\sum_{x} \alpha^{x} = 1$ and for each x, B^{x} is a stochastic matrix.

Next we consider the following:

$$\begin{aligned}
\mathrm{E}\left[\mu_{\hat{x},y}\right] &= \sum_{x,w} \mathfrak{p}^{f,Y}[y|x,w]|\tilde{I}_{\hat{x}} \cap I_{xw}| \\
&= \sum_{x,w} \mathfrak{p}^{f,Y}[y|x,w]\mathfrak{p}^{f,Y}[w|x]|\tilde{I}_{\hat{x}} \cap I_{x}| \pm O(N^{2/3}) && \text{since } (\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{w}) \in \mathsf{Good} \\
&= |\tilde{I}_{\hat{x}}| \sum_{x,w} \alpha^x \mathfrak{p}^{f,Y}[w,y|x] \pm O(N^{2/3}) \\
&= |\tilde{I}_{\hat{x}}|\mathfrak{p}^Y[y] \sum_{x} \alpha^x \pm O(N^{2/3}) && \text{since } \mathfrak{p}^{f,Y}[y|x] = \mathfrak{p}^Y[y] \\
&= |\tilde{I}_{\hat{x}}|\mathfrak{p}^Y[y] \pm O(N^{2/3})
\end{aligned}$$

So, $\mathfrak{p}^f[\tilde{w}, z|\hat{x}, y] \cdot \mathrm{E}\left[\mu_{\hat{x},y}\right] = |\tilde{I}_{\hat{x}}|P^{\hat{x}}_{(y,z),\tilde{w}} \pm O(N^{2/3})$ since $P^{\hat{x}}_{(y,z),\tilde{w}} = \mathfrak{p}^{f,Y}[\tilde{w}, y, z|\hat{x}] = \mathfrak{p}^f[\tilde{w}, z|\hat{x}, y]\mathfrak{p}^{f,Y}[y|\hat{x}] = \mathfrak{p}^f[\tilde{w}, z|\hat{x}, y]\mathfrak{p}^Y[y]$. Thus,

$$\mathrm{E}\left[\mu_{\tilde{w},\hat{x},y,z} - \mathfrak{p}^f[\tilde{w}, z|\hat{x}, y] \cdot \mu_{\hat{x},y}\right] = |\tilde{I}_{\hat{x}}| \left((\sum_x \alpha^x P^x \cdot B^x) - P^{\hat{x}} \right)_{(y,z),\tilde{w}} \pm O(N^{2/3})$$

Finally, we can rewrite $|\tilde{I}_{\hat{x}}|$ in terms of $|\tilde{J}_{\hat{x}}|$ as follows:

$$\begin{aligned}
|\tilde{J}_{\hat{x}}| &= |\tilde{I}_{\hat{x}}| - \sum_{w} |\tilde{I}_{\hat{x}w} \cap I_{\hat{x}w}| \\
&= |\tilde{I}_{\hat{x}}| - \left(|\tilde{I}_{\hat{x}} \cap I_{\hat{x}}| \sum_{w} \mathfrak{p}^{f,Y}[w|\hat{x}] \cdot \beta^{\hat{x}}_{ww} \right) \pm N^{2/3} && \text{by Equation 2} \\
&= |\tilde{I}_{\hat{x}}| \left(1 - \alpha^{\hat{x}} \cdot \sum_{w,y,z} \mathfrak{p}^{f,Y}[w, y, z|\hat{x}] \cdot \beta^{\hat{x}}_{ww} \right) \pm N^{2/3}
\end{aligned}$$

So,

$$\begin{aligned}
\mathrm{E}\left[\mu_{\tilde{w},\hat{x},y,z} - \mathfrak{p}^f[\tilde{w}, z|\hat{x}, y] \cdot \mu_{\hat{x},y}\right] &= \left(|\tilde{J}_{\hat{x}}| \pm O(N^{2/3}) \right) \left(\frac{((\sum_x \alpha^x P^x \cdot B^x) - P^{\hat{x}})_{(y,z),\tilde{w}}}{1 - \alpha^{\hat{x}} \cdot \sum_{w,y,z} \mathfrak{p}^{f,Y}[w, y, z|\hat{x}] \cdot \beta^{\hat{x}}_{ww}} \right) \pm O(N^{2/3}) \\
&= |\tilde{J}_{\hat{x}}| \left(\frac{((\sum_x \alpha^x P^x \cdot B^x) - P^{\hat{x}})_{(y,z),\tilde{w}}}{1 - \alpha^{\hat{x}} \cdot \sum_{w,y,z} \mathfrak{p}^{f,Y}[w, y, z|\hat{x}] \cdot \beta^{\hat{x}}_{ww}} \right) \pm o(N^{7/8})
\end{aligned}$$

where in the last step we used that fact that $|\tilde{J}_{\hat{x}}| = \Omega(N^{7/8})$ and $|\tilde{J}_{\hat{x}}| \leq |\tilde{I}_{\hat{x}}| \leq N$.

Finally, by Lemma 2, since f is redundancy free, $\mathfrak{D}(P^1, \ldots, P^m) \geq \epsilon_f \cdot \min(\mathfrak{p}^Y)$, where $\epsilon_f > 0$ is a constant. Since $\mathfrak{p}^Y$ has full support (and is independent of N), $\min(\mathfrak{p}^Y) > 0$ is also a constant.

Thus,

$$\begin{aligned}
\max_{(\tilde{w},y,z)} \left| \mathrm{E}\left[\mu_{\tilde{w},\hat{x},y,z} - \mathfrak{p}^f[\tilde{w},z|\hat{x},y]\cdot\mu_{\hat{x},y}\right]\right| &\geq |\tilde{J}_{\hat{x}}| \left(\frac{\|(\sum_x \alpha^x P^x \cdot B^x) - P^{\hat{x}}\|_{\max}}{1-\alpha^{\hat{x}}\cdot\sum_{w,y,z}\mathfrak{p}^{f,Y}[w,y,z|\hat{x}]\cdot\beta^{\hat{x}}_{ww}} \right) \pm o(N^{7/8}) \\
&\geq \frac{|\tilde{J}_{\hat{x}}|}{q} \left(\frac{\|(\sum_x \alpha^x P^x \cdot B^x) - P^{\hat{x}}\|_{\infty}}{1-\alpha^{\hat{x}}\cdot\sum_{w,y,z}\mathfrak{p}^{f,Y}[w,y,z|\hat{x}]\cdot\beta^{\hat{x}}_{ww}} \right) \pm o(N^{7/8}) \\
&\geq \frac{|\tilde{J}_{\hat{x}}|}{q}\mathfrak{D}(P^1,\ldots,P^m) \pm o(N^{7/8}) = \Omega(N^{7/8}).
\end{aligned}$$

To complete the proof we use Chernoff bounds to argue that with all but negligible probability, for $(\tilde{w},y,z)$ which maximizes the above expectation, $|\mu_{\tilde{w},\hat{x},y,z} - \mathfrak{p}^f[\tilde{w},z|\hat{x},y]\cdot\mu_{\hat{x},y}| > N^{2/3}$ (when N is sufficiently large). □

3.3 A Converse of The Channel Coding Theorem

We first prove a generalization of the weak converse of channel coding theorem where the receiver can adaptively choose the channel based on its current view. If it uses μ fraction of channels which have low rate, then we lower bound its error probability of predicting the input codeword as a function of μ, an upperbound on the channel capacities, and rate of the code.

Lemma 4 (Weak Converse of Channel Coding Theorem, Generalization)**.** *Let $\mathcal{F} = \{\mathcal{F}_1,\ldots,\mathcal{F}_K\}$ be a set of K channels which take as input alphabets from a set A, with $|A| = 2^\lambda$. Let $\mathcal{G} \subseteq [K]$ be such that for all $i \in \mathcal{G}$, the capacity of the channel $\mathcal{F}_i$ is at most $\lambda - c$, for a constant $c > 0$.*

Let $\mathcal{C} \subseteq A^M$ be a rate $R \in [0,1]$ code. Consider the following experiment: a random codeword $c_1 \ldots c_M \equiv \mathbf{c} \overset{\$}{\leftarrow} \mathcal{C}$ is drawn and each symbol $c_1 \ldots c_M$ is transmitted sequentially; the channel used for transmitting each symbol is chosen (possibly adaptively) from the set $\mathcal{F}$ by the receiver.

Let S denote the set of indices $j \in [M]$ for which the receiver chose a channel in $\mathcal{G}$ for receiving c_j. If the receiver always chooses the channels such that $|S|/M \geq \mu$, then the probability of error of the receiver in predicting $\mathbf{c}$ is

$$P_e \geq 1 - \frac{1}{MR\lambda} - \frac{1-c\mu/\lambda}{R}.$$

Proof. Let the codeword $\mathbf{c}$ be chosen uniformly from the code; and $\mathbf{d}$ represent the outputs $d_1 \ldots d_M$ received by the receiver and $\mathbf{y}$ represent $y_1 \ldots y_M$ the inputs used by the receiver. First note that:

$$\begin{aligned}
MR\lambda = H(\mathbf{c}) &= H(\mathbf{c}|\mathbf{y},\mathbf{d}) + I(\mathbf{c};\mathbf{y},\mathbf{d}) \\
&\leq 1 + P_e MR\lambda + I(\mathbf{c};\mathbf{y},\mathbf{d}) && \text{By Fano's Inequality}
\end{aligned}$$

Now, we shall upper bound the mutual information $I(\mathbf{c};\mathbf{y},\mathbf{d})$. We use $\mathbf{c}^{(j)}$ to denote $c_1 \ldots c_j$; and similarly define $\mathbf{y}^{(j)}$ and $\mathbf{d}^{(j)}$. We can write:

$$\begin{aligned}
I(\mathbf{c};\mathbf{y},\mathbf{d}) &= \sum_{j\in[M]} I(\mathbf{c};y_j,d_j|\mathbf{y}^{(j-1)},\mathbf{d}^{(j-1)}) \\
&= \sum_{j\in[M]} I(\mathbf{c};y_j|\mathbf{y}^{(j-1)},\mathbf{d}^{(j-1)}) + I(\mathbf{c};d_j|\mathbf{y}^{(j)},\mathbf{d}^{(j-1)})
\end{aligned}$$

Note that $\mathbf{c} \to (\mathbf{y}^{(j-1)}, \mathbf{d}^{(j-1)}) \to y_j$, so we have $I(\mathbf{c}; y_j | \mathbf{y}^{(j-1)}, \mathbf{d}^{(j-1)}) = 0$. Therefore, we get:

$$\begin{aligned} I(\mathbf{c}; \mathbf{y}, \mathbf{d}) &= \sum_{j\in[M]} I(\mathbf{c}; d_j | \mathbf{y}^{(j)}, \mathbf{d}^{(j-1)}) \\ &= \sum_{j\in[M]} H(d_j | \mathbf{y}^{(j)}, \mathbf{d}^{(j-1)}) - H(d_j | \mathbf{y}^{(j)}, \mathbf{d}^{(j-1)}, \mathbf{c}) \\ &\leq \sum_{j\in[M]} H(d_j | y_j) - H(d_j | \mathbf{y}^{(j)}, \mathbf{d}^{(j-1)}, \mathbf{c}) \\ &= \sum_{j\in[M]} H(d_j | y_j) - H(d_j | y_j, c_j) \\ &= \sum_{j\in[M]} I(d_j; c_j | y_j) \end{aligned}$$

Let E_j be the indicator variable for the event that the j-th index $i_j \in \mathcal{G}$, i.e. the adversary chooses a channel with capacity $\leq \lambda - c$. Define p_j as the probability of $E_j = 1$. We know that $\sum_{j\in[M]} p_j \geq \mu M$. Now, we know that $I(d_j; c_j | y_j) \leq p_j(\lambda - c) + (1 - p_j)\lambda = \lambda - cp_j$ Therefore, $I(\mathbf{c}; \mathbf{y}, \mathbf{d}) \leq M(\lambda - c\mu)$.

Combining this with the previous result, we get:

$$\begin{aligned} & MR\lambda \leq 1 + P_e MR\lambda + M(1 - c\mu/\lambda)\lambda \\ & \Rightarrow P_e \geq 1 - \frac{1}{MR\lambda} - \frac{1 - c\mu/\lambda}{R} \end{aligned}$$

This completes the proof of the lemma. □

4 Main Construction

In this section we prove the following theorem, which forms the main ingredient for the proof of Theorem 1.

Theorem 3. *If f is a redundancy free 2-party function and f is passive-complete, then f is UC-complete.*

To prove this theorem we show that the OT function (which is known to be UC-complete [Kil88, IPS08]) reduces to an f as in the theorem. Since f is passive-complete we know that OT does reduce to f against passive adversaries. We shall take such a passive-secure OT protocol in the f-hybrid, and convert it into a UC-secure protocol. For this we need two ingredients: first a UC-secure commitment protocol in the f-hybrid model, and secondly a compiler to turn the passive secure OT protocol in the f-hybrid model to a UC-secure protocol in the commitment-hybrid model.

4.1 A UC Secure Commitment Protocol

Before presenting the protocol, we define a distribution over the inputs of f which has a special property, and show that such a distribution always exists if f is redundancy free and passive-complete. We say that an input $y \in Y$ is a *fully-revealing* input, if $\nexists z \in Z, x, x' \in X, w, w' \in W$

such that $(x, w) \neq (x', w')$ and $\mathfrak{p}^f[w, z|x, y] > 0, \mathfrak{p}^f[w', z|x', y] > 0$. That is, from z (and y) Bob can exactly find out what (w, z) is. Note that if all of Bob's inputs are fully revealing, then f is passive-trivial (it is passive-secure for Alice to send x to Bob, who computes w and sends it back to her), and hence f is not passive-complete.

For each $y \in Y$ we can define a point $\mathbf{d}_y \in \mathbb{R}^{mq}$ denoting the probability distribution over $X \times W$ of the view (x, w) of Alice, when Bob's input to f is y, and Alice's input is uniformly randomly chosen. Consider the convex hull of the set of points $D = \{\mathbf{d}_y | y \text{fully-revealing}\}$ (which may or may not be empty). We claim that if f is redundancy-free and has an input $y \in Y$ that is not fully-revealing, then $\mathbf{d}_y$ is outside the region D. Otherwise, y would be a right-redundant input: a convex combination of fully revealing inputs produces (x, w) distributed identically to what y produces; further, due to these inputs being fully revealing, Bob can sample $\tilde{z}$ to be consistent with (x, w).

We define an *unrevealing distribution* over Y (with respect to the function f) to be a distribution with full support over Y such that the corresponding weighted average of the vectors $\mathbf{d}_y$ is outside the region D. If f is redundancy free and is passive-complete, such a distribution exists since we saw that for at least one y, $\mathbf{d}_y$ is outside of D, and hence a weighted average between (say) the uniform distribution (whose point may or may not be within D) and the point distribution that puts all its weight on y (whose point is outside D), will result in a point that is outside D.

Bit-Commitment(b, f, M, N):
The protocol is presented in terms of a code $\mathcal{C}$ over the alphabet $\mathbb{Z}_{mq}$ (or equivalently $X \times W$) with block-length M, rate 1 and distance $\omega(M^{7/8})$. (An explicit code is not necessary: Bob can pick at random $\omega(M^{7/8})$ "parity checks" to construct the code and announce it to Alice.)

1. Commit Phase:

   (a) The sender picks $x_1 \dots x_N \xleftarrow{\$} X^N$. The receiver picks $y_1 \dots y_N \in Y^N$, where each y_i is i.i.d., according to *an unrevealing distribution* $\mathfrak{p}^Y$. Both parties invoke f with respective inputs x_i and y_i; the function computes $(w_i, z_i) \xleftarrow{\$} f(x_i, y_i)$, and sends w_i to the sender and z_i to the receiver.

   (b) The sender carries out a consistency check on $\{(x_i, w_i)\}_{i=1}^{N/M}$: it checks that for each value of $(x, w) \in X \times W$ the number of indices i with $(x_i, w_i) = (x, w)$ is $(\sum_{y \in Y} \mathfrak{p}^Y[y]\mathfrak{p}^f[x, w|y] \pm N^{2/3}$. If not, it aborts the protocol.

   (c) The sender picks a string $c_1 \dots c_N \in \mathbb{Z}_{mq}^N$ in the form of N/M codewords $\mathbf{c}_1 \dots \mathbf{c}_{N/M} \xleftarrow{\$} \mathcal{C}^{N/M}$. Let $r_i = c_i + \phi(x_i, w_i)$. The sender sends $r_1 \dots r_N$ to the receiver.

   (d) The sender picks $h \leftarrow \mathcal{H}$, a universal hash function family mapping X^N to $\{0, 1\}$ and sends $(h, b \oplus h(c_1 \dots c_N))$ to the receiver.

2. Reveal Phase:

   (a) Sender sends $(x_1, w_1), \dots (x_N, w_N)$ to the receiver, who recovers $c_1 \dots c_N$ and b. The receiver accepts the opening if and only if the vector $\mathbf{c}_1 \dots \mathbf{c}_{N/M} \in \mathcal{C}^{N/M}$ and $(x_1, w_1), \dots (x_N, w_N)$ and $(y_1, z_1), \dots (y_N, z_N)$ satisfy the consistency checks in the Left-Statistical-Test.

We sketch the proof of security for this commitment protocol (with say $N/M = M = \kappa$). Since we are in the information-theoretic setting, with computationally unbounded adversaries and simulators, we focus on showing the statistical hiding property and statistical binding property separately. These can be easily turned into a simulation argument.

Hiding. To see the hiding property, consider the use of the function f as a "channel," which accepts c_i from Alice, y_i from Bob and samples (x_i, w_i, z_i) and outputs w_i to Alice and $a_i + c_i$ to Bob, where $a_i = \phi(x_i, w_i)$. The hiding property relies on the fact that Bob is forced to use f as channel with capacity strictly less than $\log nq$: as we shall see below, this is enforced by the sender's check in step (b). Then we appeal to (an extension) of the weak converse of Shannon's Channel Coding Theorem [Cov] to argue that since the code has rate 1, some information about the code remains hidden from the receiver. We need an extension of the (weak) converse of the channel coding theorem to handle that facts that (a) the receiver can adaptively choose the channel characteristic, by picking y_i adaptively, and (b) some of the channel characteristics that can be chosen include a noiseless channel, but the number of times such a characteristic can be used cannot be large (except with negligible probability). The reason this restriction can be enforced is because the $\mathfrak{p}^Y$ is an unrevealing distribution. The check carried out by the sender is simple and cannot *bind* the receiver to using $\mathfrak{p}^Y$, but it ensures that Bob cannot use almost always use only fully-revealing ys (because Alice's view induced by such y lies inside the region D whereas Alice checks it for being outside D).

Finally, the commitment is made hiding by masking the bit to be committed by a bit extracted from the codewords $\mathbf{c}_i$.

This argument can be easily turned into a (computationally unbounded) simulator.

Binding. Binding follows from the fact that the advantage any adversary has in the Left-Statistical-Test involving f is negligible in N. Note that to be able to equivocate the adversary has to give two explanations $(\tilde{\mathbf{x}}, \tilde{\mathbf{w}})$ and $(\mathbf{x}, \mathbf{w})$ both of which should result in the same value for $\mathbf{r}$, using two different codewords. Since the code $\mathcal{C}$ has minimum distance $N^{\omega 7/8}$

Note that a simulator, which sees $\mathbf{x}$ can decode the codeword closest to $\{x_i \oplus r_i\}_i$, and use it to extract a bit b.

4.2 Hiding of the Commitment Protocol

We use the converse of the channel coding theorem and the irredundancy of the function f to argue that our commitment protocol is statistically hiding.

First Game: Error in predicting each codeword. Note that if $\lambda = \Theta(1)$, $M = \omega(1)$, $R = 1 - o(1)$, $c = \Theta(1)$ and $\mu = \Theta(1)$, then $P_e = \Theta(1)$. As a direct application of this result, we get the following result:

Let $A = \mathbb{Z}_{mq}$ and establish a bijection between A and $X \times W$. Similarly, let $B = \mathbb{Z}_{nr}$ and establish a bijection between B and $Y \times Z$. Define $\lambda = \log |A|$. Consider the following game between an honest challenger and an adversary:

1. The challenger picks a codeword $c_1 \dots c_M \equiv \mathbf{c} \xleftarrow{\$} \mathcal{C}$, where $\mathcal{C} \subseteq A^M$, $|\mathcal{C}| = MR\lambda$ and $R = 1 - o(1)$.

2. For each $i \in [M]$, sequentially repeat these steps:

   (a) The challenger picks $x_i \xleftarrow{\$} X$, for $i \in [M]$ and feeds into f.

   (b) The adversary feeds $y_i \in Y$.

   (c) The output $(w_i, z_i) \xleftarrow{\$} f(x_i, y_i)$ is computed; and w_i is given to the challenger and z_i is given to the adversary.

   (d) The challenger sends $r_i = a_i + c_i$ to the adversary.

3. The adversary outputs $\tilde{\mathbf{c}} \in A^M$.

The adversary wins the game if $\tilde{\mathbf{c}} = \mathbf{c}$, i.e. it is able to correctly guess the codeword. We shall show that the probability that the adversary loses this game is at least a constant, if the adversary feeds $y_i \in Y$ which are not completely revealing for μM rounds, where μ is a constant.

To directly apply Lemma 4 consider the following alternate, but equivalent game:

1. The challenger picks a codeword $c_1 \dots c_M \equiv \mathbf{c} \xleftarrow{\$} \mathcal{C}$, where $\mathcal{C} \subseteq A^M$, $|\mathcal{C}| = MR\lambda$ and $R = 1 - o(1)$.

2. For each $i \in [M]$, sequentially repeat these steps:

   (a) The adversary feeds $y_i \in Y$ to the channel.

   (b) The challenger sends c_i to the channel.

   (c) The channel first samples $x_i \xleftarrow{\$} X$ and then computes $(w_i, z_i) \xleftarrow{\$} f(x_i, y_i)$. Next it computed $b_i \in B$ the mapping of (y_i, z_i) and $r_i = c_i + a_i$, where $a_i \in A$ is the mapping of (x_i, z_i). The channel sends b_i and r_i to the adversary.

3. The adversary outputs $\tilde{\mathbf{c}} \in A^M$.

The adversary wins the game if $\tilde{\mathbf{c}} = \mathbf{c}$. Now, this is formulated as a game where the adversary can pick channel, at least $\mu = \Theta(1)$ fraction of whom are not fully revealing. Applying Lemma 4, we directly get that the adversary loses the game with probability $P_e = \Theta(1)$, if $M = \omega(1)$, $R = 1 - o(1)$ and $\mu = \Theta(1)$.

Second Game: Negligible advantage in predicting the bit. Consider the following game between an honest challenge and an adversary:

1. The challenger chooses $\mathbf{c}_1 \dots \mathbf{c}_{N/M} \xleftarrow{\$} \mathcal{C}^{N/M}$.

2. The challenger and adversary perform N/M copies of the previous game and in the k-th game, the challenger uses $\mathbf{c_k}$.

3. Interpret $\mathbf{c}_1 \dots \mathbf{c}_{N/M} \equiv u_1 \dots u_N$, where $u_i \in A$ for all $i \in [N]$. The challenger draws $h \xleftarrow{\$} \mathcal{H}$, where cH is a family of universal hash functions mapping A^N to $\{0, 1\}$. The challenger computes $b = h(u_1 \dots u_N)$ and sends h to the adversary.

4. The adversary output $\tilde{b} \in \{0, 1\}$.

The adversary wins the game if $b = \tilde{b}$.

The following analysis is conditioned on the fact that among the inputs $\{y_{(k-1)M+1}, \ldots, y_{kM}\}$ used by the adversary, at least $\mu = \Theta(1)$ fraction of them are not fully revealing, for every $k \in [N/M]$. We shall be using the notion of average min-entropy (denoted by $\tilde{H}_\infty$) as introduced by [DORS08]. Let us denote the complete view of the adversary by V and $\mathbf{u}$ denote the random variable $u_1 \ldots u_N$. Note that for each $k \in [N/M]$, the codeword $\mathbf{c_k}$ is incorrectly predicted by the adversary with probability at least $P_e = \Theta(1)$. Therefore, $\tilde{H}_\infty(\mathbf{u}|V) \geq \Theta(N/M)$. Finally, using the result that universal hash functions are good strong extractors for sources with high average min-entropy [DORS08], we get that b is statistically hidden from the adversary, if $N/M = \omega(1)$. Formally, let U be the uniform bit. Then $\mathsf{SD}\left((b, V), (U, V)\right) \leq \frac{1}{\sqrt{2}} 2^{-\tilde{H}_\infty(\mathbf{u}|V)/2} = 2^{-\Theta(N/M)}$.

Combining these two results, we get the following lemma:

Lemma 5. *Let $M = \omega(1)$, $R = 1 - o(1)$ and $\mathcal{C} \subseteq A^M$ be a rate R code. Define $S_k = \{(k-1)M + 1, \ldots, kM\}$, for $k \in [N/M]$. If the adversary uses $\Theta(M)$ inputs which are not fully revealing in rounds indexed by S_k, for every $k \in [N/M]$, then the advantage of the adversary in the following game is at most $2^{-\Theta(N/M)}$.*
$\mathsf{Hiding\text{-}Game}(N, M, \mathcal{C})$*:*

1. *For $k \in [N/M]$ Repeat the following steps:*
   - *(a) The challenger picks a code word $c_{(k-1)M+1} \ldots c_{kM} \equiv \mathbf{c_k} \xleftarrow{\$} \mathcal{C}$.*
   - *(b) For $i \in [M]$ repeat the following steps:*
     - *i. The challenger picks $x_{(k-1)M+i} \xleftarrow{\$} X$ and the adversary picks $y_{(k-1)M+i} \in Y$. They invoke f with these inputs and receive respective outcomes $w_{(k-1)M+i}$ and $z_{(k-1)M+i}$ from the functionality. Let $a_{(k-1)M+i} \in A$ be the element corresponding to $(x_{(k-1)M+i}, w_{(k-1)M+i})$.*
2. *The challenger computes $r_i = c_i + a_i$, for every $i \in [N]$. It samples $h \leftarrow \mathcal{H}$ and computes $b = h(c_1 \ldots c_N)$. The challenger sends $(h, r_1 \ldots r_N)$ to the adversary.*
3. *The adversary finally outputs $\tilde{b}$.*

The adversary wins the game if $b = \tilde{b}$.

4.3 Passive-to-Active Security Compiler

For any redundancy free SFE f, we describe a "compiler" that takes a 2-party protocol π in the f-hybrid and produces another 2-party protocol $\Pi(\pi, f)$ in the commitment-hybrid such that if π is a semi-honest $\binom{2}{1}$-OT protocol, then $\Pi(\pi, f)$ is a UC secure $\binom{2}{1}$-OT protocol. For convenience, we shall place a requirement on π that it uses f with uniformly independent inputs chosen independently by the two parties.[1]

[1] This suffices for our main result, since we shall invoke this compiler with a protocol π that satisfies this requirement. However, we remark that a more tedious analysis could be used to remove this restriction: one can argue that the views of the two parties from the invocations of f produced during the execution of π should be "non-trivial" (conditioned on the rest of the view) and this suffices for variants of the binding lemmas to hold, and in turn for the compiled protocol to be secure, when the parameters are chosen suitably.

We present the compiled protocol in two steps. In the first step, we build a protocol $\rho_{\widetilde{\mathrm{OT}}}$ that UC-securely realizes the following functionality $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$.

Functionality $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$. Parametrized by a function $t(\kappa) = o(\kappa)$.

- Accept $I \subseteq [\kappa]$ from the adversary. Abort if $|I| > t(\kappa)$.
- (If not aborted) Provide κ instances of $\binom{2}{1}$-OT. Allow the adversary to control the instances indexed by I.

We implement OT in the $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$-hybrid using the OT "combiner" from [IPS08]. This protocol can tolerate a small constant fraction of corrupted OTs (we need to tolerate only a $o(1)$ fraction), and can also produce $\Omega(\kappa)$ OT instances from the κ instances provided by one instance of $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$ (we need to produce just one OT). In the rest of this section, we focus on how to implement $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$ in the commitment hybrid model.

Security of $\rho_{\widetilde{\mathrm{OT}}}$. We prove that the protocol $\rho_{\widetilde{\mathrm{OT}}}$ UC-securely realizes $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$ with parameter $t(\kappa) = \kappa^{15/16}$. For this we build a simulator interacting with the ideal functionality $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$. It simulates to the adversary an interaction of the protocol $\rho_{\widetilde{\mathrm{OT}}}$ in the f-hybrid as follows. Till Phase III it plays the part of the honest party faithfully. Note that the inputs to the protocol are not used until Phase IV, so this can be carried out faithfully. If the simulated honest party aborts its execution before entering Phase IV, the simulator completes the simulation. Otherwise it proceeds as follows.

- If the simulated honest party does not abort its execution, but the adversary has deviated from the execution it has been committed to in more than $t(\kappa)$ of the executions or π indexed by $\overline{L}$, the simulator bails out. We shall use the binding property of f to argue that this happens with negligible probability.
- Else it requests $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$ to corrupt all those indices in $[\kappa]$ corresponding to the (at most $t(\kappa)$) instances of π in which the adversary deviated. In these instances, it will carry out Phase IV execution using the correct inputs of the honest player, and (if Bob is the honest player) takes its output from that execution and makes $\mathcal{F}^{\kappa}_{\widetilde{\mathrm{OT}}}$ provide that output for that instance of OT. For those instances where the adversary has not deviated, the simulator picks an arbitrary input for the honest player and completes the simulation of Phase IV. We remark that the simulator does not employ the simulation for π, but rather runs the protocol π itself. The security guarantee for π is used in arguing that the simulation is good.

We argue that this is a good simulation with only a negligible statistical difference with the real execution. Note that we can couple the real and ideal executions upto the end of Phase III. To prove that the entire simulation is good, we show:

(a) probability of the event bail-out is negligible in the coupled execution, and

(b) conditioned on the event bail-out not occuring in the coupled execution, the two executions have negligible statistical difference.

The first part follows from the binding property, from Lemma 3. Suppose the adversary deviates in t_0 instances of π, and t_1 of those instances were indexed in L during the cut-and-choose phase. With

Protocol $\rho_{\widetilde{\mathrm{OT}}}$. Alice's inputs are $\{(x_0^i, x_1^i)\}_{i=1}^{\kappa}$ and Bob's inputs are $\{b^i\}_{i=1}^{\kappa}$ (where x_0^i, x_1^i, b^i are all bits). In the following protocol, they will invoke several instances of π with security parameter $\kappa_\pi = \kappa^c$ for some constant $c > 0$; c chosen to be sufficiently small so that the total number of f invocations (in either direction) in one session of π is upperbounded by $\kappa^{1/8}$.

PHASE I: **Coin tossing in the well.** Alice and Bob commit to 2κ strings each (of poly(κ) length, corresponding to the length of the random tape and input (two bits) required in π with security parameter κ_π). Let Alice's strings be $\{\rho_i^A\}_{i=1}^{2\kappa}$, and Bob's strings be be $\{\rho_i^B\}_{i=1}^{2\kappa}$. Then Alice sends κ strings $\{\sigma_i^A\}_{i=1}^{2\kappa}$ to Bob and Bob sends $\{\sigma_i^B\}_{i=1}^{2\kappa}$ to Alice. Alice defines input/random-tapes $\{\tau_i^A\}_{i=1}^{2\kappa}$ where $\tau_i^A = \rho_i^A \oplus \sigma_i^B$. Similarly Bob defines input/random-tapes $\{\tau_i^B\}_{i=1}^{2\kappa}$ where $\tau_i^B = \rho_i^B \oplus \sigma_i^A$.

PHASE II: **Execution.** Alice and Bob engage in 2κ executions of protocol π in the f-hybrid model. The security parameter of these executions is set to $\kappa_\pi = \kappa^c$ for a sufficiently small constant $c > 0$ so that the total number of f invocations (in either direction) in one session of π is upperbounded by $\kappa^{1/8}$.

In the i^{th} instance, Alice and Bob use τ_i^A and τ_i^B respectively as their input/random tape.

PHASE III: **Cut and Choose.** Alice and Bob use a protocol in the $\mathcal{F}_{\text{COM}}$-hybrid to UC-securely generate random coins to randomly choose a subset $L \subseteq [2\kappa]$ with $|L| = \kappa$. For each $i \in L$, Alice and Bob must "open" their views in the i^{th} execution of π: that is, Alice and Bob should reveal $\{\rho_i^A\}_{i\in L}$ and $\{\rho_i^B\}_{i\in L}$ respectively. Further each party should also report the outputs it received from f in each invocation of f.

Then each party checks (a) if the messages received in the protocol are consistent with (i.e., has non-zero probability) the views opened/reported by the other party, and (b) if its own actual views in the invocations of f (over all the executions of π that are opened) is *statistically consistent* with the expected view, based on the views of f reported by the other party. (The statistical check is similar to that in Section 3.2; see below.) If either of the checks fail, then the party should abort.

PHASE IV: **Finalizing.** Let $\overline{L} = [2\kappa] \setminus L$. Note that $|\overline{L}| = \kappa$. Alice and Bob now perform a standard procedure for carrying out a fresh OT given a pre-computed OT instance. This is summarized below.

Let $\{(s_0^i, s_1^i)\}_{i\in[\kappa]}$ and $\{u^i\}_{i\in[\kappa]}$ denote the inputs for Alice and Bob in the κ instances of π with indices in $\overline{L}$; also let $\{v^i\}_{i\in[\kappa]}$ denote the outputs that Bob received from π in these instances. For each $i \in [\kappa]$, Bob sends $c^i := b^i \oplus u^i$ to Alice and Alice responds with $(r_0^i, r_1^i) := (x_0^i \oplus s_{c^i}^i, x_1^i \oplus s_{1-c^i}^i)$ to Bob. Bob outputs $\{r_{b^i}^i \oplus v^i\}_{i=1}^{\kappa}$.

Figure 1 Protocol $\rho_{\widetilde{\mathrm{OT}}}$

high probability t_1 is close to $t_0/2$. For the honest party to not abort, in all the t_1 instances in L, the adversary should pass parts (a) and (b) of the checks. Note that the only part not determined by the protocol, given the view of the honest party and the committed values, are the views of the adversary from f invocations: so for a deviation to be not caught by part (a) of the check, either the deviation should be that the adversary actually fed a different value as input to an instance of

f than it was supposed to, or it altered the output it received from f and continued the execution faithfully with this altered output (and reported the altered output). Thus there are at least t_1 executions of f from the t_1 executions of π in which the adversary deviated as above. Of these at least $t_1/2$ have f invoked in the same direction (with the adversary playing the role of the first party (with input domain X) or of the second party: w.l.o.g., assume that the adversary plays the role of the first party in $t_1/2$ instances of f in which it deviated. Let N denote the total number of instances of f invoked in this direction out of all the κ instances of π indexed by $\overline{L}$. Recall that π is invoked with a security parameter $\kappa_\pi = \kappa^c$ for a small enough constant $c > 0$ so that the number of invocations of f in each instance of π is at most $\kappa^{1/8}$; then $N \leq \kappa^{9/8}$. By the binding lemma, we know that if the consistency check is cleared then $t_1 \leq N^{7/8} \leq \kappa^{54/64} < 2t(\kappa)$ with all but negligible probability (since $t(\kappa) = \kappa^{15/16}$). Thus the probability of $t_0 \geq t(\kappa)$, which is the probability of the event bail-out, is negligible.

To prove the second part we observe that if an environment can distinguish between the two executions, then we can break the statistical (semi-honest) security of π. More formally, we consider the advantage of the adversary in the following experiment: a fail-stop adversary takes part in κ executions of π with randomly chosen inputs (for both players). The adversary follows the protocol honestly but it can adaptively choose to abort any number of these executions, and whenever it aborts an execution, it will be given the state of the honest party in that execution. When all the executions finish, for each execution that was not aborted, define the "hidden bit" to be (the part of) the input of the honest pary that is not revealed to the adversary by the ideal OT functionality (for the inputs). Then the adversary is given either the actual hidden bits in all the unaborted executions, or independently randomly chosen bits. The adversary's advantage is the difference in its probability of outputting 1 in these two cases. By a hybrid argument it is enough to consider a single execution. Then clearly, the adversary can be assumed to not abort the execution (its advantage remains the same by not aborting and instead making a random guess); that is, we can consider only semi-honest adversaries in this experiment. By the security guarantee of π, the advantage of semi-honest adversary in distinguishing the actual hidden bit from a random bit is negligible (in κ_π and hence in κ).

5 Full Characterization of Completeness

In this section we show how Theorem 1 follows from our main construction (Theorem 3) and other observations regarding redundancy free functions. First, we introduce some definitions, following [MPR12].

In a *local protocol for f which uses g as a setup*, each party probabilistically maps her f-input to a g-input, calls g once with that input and, based on her local view (i.e. her given f-input, the output of g, and possibly local randomness), computes her final output, without any further communication between the parties.

Definition 2 ((Weak) Isomorphism [MPR12]). *We say that f and g are* isomorphic *to each other if there exist two local protocols π_1 and π_2 such that:*

1. *π_1^g UC-securely realizes f and π_2^f UC-securely realizes g;*
2. *π_1^g passive-securely realizes f and π_2^f passive-securely realizes g.*

f and g are said to be weakly isomorphic *to each other if condition 1 is satisfied.*

Note that isomorphism and weak isomorphism are equivalence relations. Also note that if two functions are weakly isomorphic to each other then one is UC-complete if and only if the other is. Further, this holds for standalone-completeness as well, since a *local protocol* that is standalone-secure must be UC-secure as well.

A *core* of a 2-party function f is a redundancy free function $\widehat{f}$ which is weakly isomorphic to f. From Lemma 10, it follows that every finite 2-party function f has a core. By the above observation about weak isomorphism, to characterize standalone or UC completeness of finite 2-party functions, it is enough to characterize it for redundnacy free functions. Note that Section A.2 gives an explicit procedure for finding a core of a given function. While the core is not unique, all the cores of a function are weakly isomorphic with each other.

The *kernel* of a 2-party function f is a function which outputs to the two parties only the "common information" that f makes available to them. To formalize this, we define a weighted bipartite graph $\mathrm{G}(f)$ with partite sets are $X \times W$ and $Y \times Z$, and for every $(x, w) \in X \times W$ and $(y, z) \in Y \times Z$, the edge joining these two vertices is assigned weight $\mathsf{wt}\Big((x,w),(y,z)\Big) := \frac{\mathfrak{p}^f[w,z|x,y]}{|X \times Y|}$. The kernel of f is a randomized function which takes inputs $x \in X$ and $y \in Y$ from the parties, samples $(w, z) \xleftarrow{\$} f(x, y)$, and outputs to both parties the connected component of $\mathrm{G}(f)$ which contains the edge $\Big((x,w),(y,z)\Big)$.

Definition 3 (Simple Function [MPR12])**.** *A (possibly randomized) 2-party function f is said to be* simple *if it is isomorphic to its kernel.*

To prove Theorem 1 note that UC-completeness implies standalone-completeness. If f is standalone complete, its core is also standalone complete, and by Lemma 7, it is passive-complete. Our main work is in showing that if f has a core that is passive-complete then f is UC-complete. Since f is weakly isomorphic to its core, it is enough to show that any redundancy free function that is passive-complete is also UC-complete. This is precisely what Theorem 3 proves.

Now we show how Theorem 2 follows. The first part, characterizing passive completeness was shown in [MPR12], building on the results in [Kil00, MOPR11]. We shall see that a standalone-complete (or UC-complete) function must always be passive-complete as well (Lemma 7). To complete the proof, we need to show that if a finite 2-party function has a core that is not simple, then it is UC-complete. Suppose f is such a function with a core $\widehat{f}$ that is not simple. By the first part, $\widehat{f}$ is passive complete. Now by Theorem 1 f is UC-complete and standalone-complete.

5.1 A Special Case

In particular, for the class for asymmetric function evaluations, i.e. where only one of the parties receives outputs, we obtain the following dichotomy:

Theorem 4 (Special Case: Dichotomy for Asymmetric 2-party SFE)**.** *Any asymmetric 2-party SFE is either standalone/UC-trivial or standalone/UC-complete.*

This theorem was proven for the deterministic case in [BMM99]. Note that in the randomized case, if there exists an input for the receiver such that it can determine the sender's input with certainty, then the SFE is standalone-/UC-trivial; because the protocol where the sender sends her input to the receiver is a standalone-/UC-secure protocol. On the other hand, if all receiver inputs are such that the receiver input cannot be predicted with certainty then our construction provides a standalone-/UC-secure construction of OT from this SFE.

6 Conclusion

In this paper we resolve the completeness problem for general 2-party SFE. But not much is known about triviality characterizations. Only for the special case of asymmetric functions, we show that the dichotomy of functions into trivial and complete extends to the randomized case as well. For other function classes, such characterizations are unknown.

Our current construction does not efficiently convert instances of f into OT instances, if f is complete. We leave it as an open problem to achieve constant rate in such construction.

Acknowledgments. We thank Vinod Prabhakaran for helpful discussions on the converse of the Channel Coding Theorem. We also thank Amit Sahai for useful discussions.

References

[BMM99] Amos Beimel, Tal Malkin, and Silvio Micali. The all-or-nothing nature of two-party secure computation. In Michael J. Wiener, editor, *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 80–97. Springer, 1999. 2, 22

[CDMW09] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Simple, black-box constructions of adaptively secure protocols. In Omer Reingold, editor, *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 387–402. Springer, 2009. 1, 4

[CK88] Claude Crépeau and Joe Kilian. Achieving oblivious transfer using weakened security assumptions (extended abstract). In *FOCS*, pages 42–52. IEEE, 1988. 1, 2

[CMW04] Claude Crépeau, Kirill Morozov, and Stefan Wolf. Efficient unconditional oblivious transfer from almost any noisy channel. In Carlo Blundo and Stelvio Cimato, editors, *SCN*, volume 3352 of *Lecture Notes in Computer Science*, pages 47–59. Springer, 2004. 1, 2

[Cov] 16

[DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008. 18

[GV87] Oded Goldreich and Ronen Vainish. How to solve any protocol problem - an efficiency improvement. In Carl Pomerance, editor, *CRYPTO*, volume 293 of *Lecture Notes in Computer Science*, pages 73–86. Springer, 1987. 2

[Hai08] Iftach Haitner. Semi-honest to malicious oblivious transfer - the black-box way. In Ran Canetti, editor, *TCC*, volume 4948 of *Lecture Notes in Computer Science*, pages 412–426. Springer, 2008. 1, 4

[HM86] Stuart Haber and Silvio Micali. Unpublished manuscript, 1986. 2

[IKLP06] Yuval Ishai, Eyal Kushilevitz, Yehuda Lindell, and Erez Petrank. Black-box constructions for secure computation. In *STOC*, pages 99–108. ACM, 2006. 1, 4

[Imp95] Russell Impagliazzo. A personal view of average-case complexity. In *Structure in Complexity Theory Conference*, pages 134–147, 1995. 1

[IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, *CRYPTO*, volume 5157 of *Lecture Notes in Computer Science*, pages 572–591. Springer, 2008. 1, 4, 14, 19

[Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In Janos Simon, editor, *STOC*, pages 20–31. ACM, 1988. 1, 2, 14

[Kil91] Joe Kilian. A general completeness theorem for two-party games. In Cris Koutsougeras and Jeffrey Scott Vitter, editors, *STOC*, pages 553–560. ACM, 1991. 1, 2

[Kil00] Joe Kilian. More general completeness theorems for secure two-party computation. In F. Frances Yao and Eugene M. Luks, editors, *STOC*, pages 316–324. ACM, 2000. 1, 2, 22

[KM11] Daniel Kraschewski and Jörn Müller-Quade. Completeness theorems with constructive proofs for finite deterministic 2-party functions. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 364–381. Springer, 2011. 1, 2

[MOPR11] Hemanta K. Maji, Pichayoot Ouppaphan, Manoj Prabhakaran, and Mike Rosulek. Exploring the limits of common coins using frontier analysis of protocols. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2011. 22

[MPR12] Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. A unified characterization of completeness in secure function evaluation. To appear at INDOCRYPT, 2012. 1, 2, 3, 21, 22

[Rab81] M. Rabin. How to exchange secrets by oblivious transfer. Technical Report TR-81, Harvard Aiken Computation Laboratory, 1981. 2

[Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15:78–88, January 1983. 2

[WW06] Stefan Wolf and Jürg Wullschleger. Oblivious transfer is symmetric. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 222–232. Springer, 2006. 4

[Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC*, pages 209–213. ACM, 1979. 1

A Proofs

A.1 If Redundancy Free, Active-Secure ⇒ Passive-Secure

Lemma 6. *Let f be a redundancy free 2-party function. If f has a standalone (or UC) secure protocol in g-hybrid, then f also has a passive-secure protocol in g-hybrid.*

Proof. We will show that the same protocol that is a standalone secure realization of f in g-hybrid is also a passive-secure protocol for f in g-hybrid.

Consider the case when Alice is corrupt. We are given that there exists a simulator for corrupt Alice in the standalone or UC setting. We need to show that, conditioned on the existence of such a simulator, we get a semi-honest simulator for f. In fact, we shall leverage the left-redundancy of f to show this result.

For any input x, let N_x be the event that the simulator invokes the ideal functionality on inputs other than x or malicious Alice gets an output which was not the output sent by the ideal functionality to the simulator. If probability of N_x is negligible, then we consider a *semi-honest* simulator which faithfully simulates the standalone/UC simulator. If the input sent to the ideal functionality is different from x or the output obtained by malicious Alice is different from the output obtained from the ideal functionality then it aborts. For an external environment, interactions with these two simulator are statistically indistinguishable because the semi-honest simulation is statistically close to the original simulation. Hence, we can conclude that there exists a semi-honest simulator.

If the probability of the event N_x is non-negligible for some $x \in X$, then there exists an infinite set of security parameters κ where probability of N_x (represented by $p_x(\kappa)$) is significant, i.e. $1/\mathrm{poly}(\kappa)$, but the statistical distance between the real and simulated view of the environment is $\delta_x(\kappa) = \mathrm{negl}(\kappa)$ close to its real view. Now consider the set V of simulator views such that, on input x, the event N_x takes place. Define the following adversarial algorithm A: Randomly pick a view from V and follow its simulation strategy. Consider interaction of A in the Left-Statistical-Test. The separation condition is trivially satisfied, because the input fed to the simulator or the output received from the simulator does not match the input or the output given to the external environment.

Note that $p_x(\kappa)$ is significant, while the probability $\delta_x(\kappa)$ is negligible. Thus, restricted to the views V, the statistical distance between environment views can be at most $\delta_x(\kappa)/p_x(\kappa) = \mathrm{negl}(\kappa)$. This ensures that consistency check is also satisfied. So, we arrive at a contradiction (because for left redundancy free functionalities, it is impossible to win the binding experiment, except with negligible probability); thus, it is not possible that there exists $x \in X$ such that N_x is non-negligible.

Note that the whole argument is independent of the hybrid g being used. Further, considering the simulator for Bob and leveraging that f is right redundancy free, we can similarly conclude that there exists a semi-honest simulator for Bob. This concludes the proof. □

We can use this result to claim the following:

Lemma 7. *If a 2-party function g is standalone-complete (or UC-complete) then it is also passive-complete.*

Proof. Suppose g is standalone-complete (or UC-complete). Then there is a standalone-secure protocol for OT in g-hybrid. Since OT is redundancy free, by Lemma 6, this protocol is passive-secure as well. Since OT is passive-complete and passive-security admits secure composition, we conclude that g is passive-complete as well. □

A.2 An Algorithm to Find a Core

In this section we show that every function has a core and we give an explicit algorithm to find one. We begin by proving two results.

Lemma 8. *Suppose $x^* \in X$ is a strictly left-redundant input of a function $f : X \times Y \to W \times Z$. Let g be the function obtained by restricting f to the domain $(X \setminus \{x^*\}) \times Y$. Then, f and g are weakly isomorphic.*

Proof. Since x^* is strictly redundant, there exists $\{(\alpha_x, P^x, M_x) | x \in X\}$ and x^* such that $P^{x^*} = \sum_{x\in X} \alpha_x P^x M_x$, $\sum_{x\in X} \alpha_x = 1$, $\alpha_x \geq 0$ (for all $x \in X$) and $\alpha_{x^*} = 0$.

First, we show that there exists standalone/UC secure local protocol for f in the g-hybrid. Bob always feeds his input y to g. If Alice's input is $x \neq x^*$, simply feed x to g and both parties obtain the correct output distribution. If Alice's input is $x = x^*$, then sample sample x' from $X \setminus \{x^*\}$ according the the probability distribution $\{\alpha_x | x \in X \setminus \{x\}\}$. Alice invokes g with input x'. It receives outcome w' from the function. Sample an output w according to the distribution in $M_{x'}$ corresponding to output w' (i.e. the distribution represented by the row corresponding to output symbol w'). By definition of strict row redundancy, the protocol is correct. Simulation is trivial for both malicious Alice and Bob cases (the simulators just forward the input provided to the g-hybrid to the external ideal functionality and forward the output back to the party).

For the other direction, i.e. a secure protocol for g in f hybrid, the protocol is trivial. Both parties invoke f with their respective inputs and report their outputs. The simulator for malicious Bob is trivial (simply forward the input to f-hybrid to the external ideal functionality and report back the received outcome). The simulator for malicious Alice is as follows. If the f-hybrid is invoked with $x \neq x^*$, then simply forward that input to the ideal functionality g and report back the received output. If the f-hybrid is invoked with $x = x^*$, then sample x' according to the distribution $\{\alpha_x | x \in X \setminus \{x\}\}$. Invoke the ideal functionality g with input x' and receive the outcome w'. Translate w' into w by sampling according to the distribution in the row of $M_{x'}$ corresponding to the output symbol w'. The simulation is perfect due to strict left redundancy. □

Lemma 9. *Suppose $x \in X$ is a self left redundant input for f, and the two columns corresponding to w and w' in P^x are scalar multiples of each other. Suppose g is a function obtained by transferring all probability mass of w'-th column of P^x to w-th column: i.e., for all $y \in Y, z \in Z$, $\mathfrak{p}^g[w, z|x, y] = \mathfrak{p}^f[w, z|x, y] + \mathfrak{p}^f[w', z|x, y]$ and $\mathfrak{p}^g[w', z|z, y] = 0$. Then, f and g are weakly isomorphic.*

Proof. Protocol for f in g-hybrid is constructed as follows. Alice and Bob forwards their inputs to g. Alice, on input x, if she receives w as the output translates it into w' with probability $\mu/(1 + \mu)$, where the column corresponding to w' was μ times the column corresponding to w. Correctness is trivial. Simulator for malicious Bob simply forwards the input to g hybrid to the external ideal functionality and forwards the received output. Simulator for malicious Alice forwards the input g

hybrid to the external ideal functionality. If the input was x and the received outcome was w' then it forwards w to the adversary; otherwise it simply forwards the received outcome.

Protocol for g in the f hybrid is constructed as follows. Both parties forwards their inputs to f. If Alice receives output w' then it outputs w; otherwise she honestly reports the received output. Simulation for malicious Bob is trivial. Simulation for malicious Alice does the following: It forwards the input for f hybrid to the external ideal functionality and receives the output. If the input was x and the output was w, then it reports w' with probability $\mu/(\mu+1)$; otherwise it honestly reports w. □

Similar results also hold for strict right redundancy and self right redundancy. Now we give an algorithm that given a function f finds a core $\widehat{f}$. If f is redundancy free, then it is a core of itself. Otherwise, by Lemma 1, f is either strictly redundant or self redundant. In the former case, obtain g as in Lemma 8, and in the latter case obtain g as in Lemma 9. Note that in either case g is guaranteed to be well-defined (and in particular does not have an empty input or output domain). Then recursively apply this algorithm to g. Note that at every step the number of pairs $(x, w) \in X \times W$ or the number of pairs $(y, z) \in Y \times Z$ such that $\mathfrak{p}^f[w, z|x, y] > 0$ strictly reduces. Since we will never reach a situation where one of these sets become empty, the procedure must terminate with a well-defined function $\widehat{f}$ that is redundancy free. Since the function we chose at every step is weakly isomorphic to the previous function, $\widehat{f}$ is weakly isomorphic to f. Thus it is a core of f.

In particular, we have the following:

Lemma 10. *Every finite 2-party function has a core.*