On the Feasibility of Extending Oblivious Transfer*

Yehuda Lindell Hila Zarosim

Dept. of Computer Science
Bar-Ilan University, Israel
lindell@biu.ac.il,zarosih@cs.biu.ac.il

January 23, 2013

Abstract

Oblivious transfer is one of the most basic and important building blocks in cryptography. As such, understanding its cost is of prime importance. Beaver (STOC 1996) showed that it is possible to obtain $\mathsf{poly}(n)$ oblivious transfers given only n actual oblivious transfer calls and using one-way functions, where n is the security parameter. In addition, he showed that it is impossible to extend oblivious transfer information theoretically. The notion of extending oblivious transfer is important theoretically (to understand the complexity of computing this primitive) and practically (since oblivious transfers can be expensive and thus extending them using only one-way functions is very attractive).

Despite its importance, very little is known about the feasibility of extending oblivious transfer, beyond the fact that it is impossible information theoretically. Specifically, it is not known whether or not one-way functions are actually necessary for extending oblivious transfer, whether or not it is possible to extend oblivious transfers with adaptive security, and whether or not it is possible to extend oblivious transfers when starting with $O(\log n)$ oblivious transfers. In this paper, we address these questions and provide almost complete answers to all of them. We show that the existence of any oblivious transfer extension protocol with security for static semi-honest adversaries implies one-way functions, that an oblivious transfer extension protocol with adaptive security implies oblivious transfer with static security, and that the existence of an oblivious transfer extension protocol from only $O(\log n)$ oblivious transfers implies oblivious transfer itself.

*This research was supported by the israel science foundation (grant No. 189/11). Hila Zarosim is grateful to the Azrieli Foundation for the award of an Azrieli Fellowship.

1 Introduction

Background – extending oblivious transfer. In the oblivious transfer problem [16, 5], a sender holds a pair of input bits (b_0, b_1) and enables a receiver to obtain one of them at its choice. The security requirements are that the sender learns nothing about which input is obtained by the receiver, while the receiver learns only one bit.

Oblivious transfer is one of the most basic and important primitives in cryptography in general, and in secure computation in particular. Oblivious transfer is used in almost all general protocols for secure computation with no honest majority (e.g., see [18, 7]), and has been shown to imply essentially all basic cryptographic tasks [14]. Due to its importance, the complexity of computing oblivious transfer is of great importance. Oblivious transfer can be constructed from enhanced trapdoor permutations [5, 10] and from homomorphic encryption [1]. In addition, it is known that it is not possible to construct oblivious transfer from public-key encryption (or one-way functions and permutations) in a black-box manner [6]. Thus, oblivious transfer requires quite strong hardness assumptions (at least when considering black-box constructions, and no nonblack-box constructions from weaker assumptions are known).

Due to the importance of oblivious transfer and its cost, Beaver asked whether or not it is possible to use a small number of oblivious transfers and a weaker assumption like one-way functions in order to obtain many oblivious transfers [3]; such a construction is called an OT extension. Beaver answered this question in the affirmative and in a beautiful construction showed how to obtain $\mathsf{poly}(n)$ oblivious transfers given ideal calls to $O(n)$ oblivious transfers and using a pseudorandom generator and symmetric encryption, which can both be constructed from any one-way function. In addition, he showed that OT extensions cannot be achieved information theoretically. These results of [3] are of great importance theoretically since they deepen our understanding of the complexity of oblivious transfer. In addition, OT extensions are of interest practically, since oblivious transfer is much more expensive than symmetric primitives. Thus, OT extensions can potentially be used to speed up protocols that rely on many oblivious transfers. In this direction, efficient OT extensions (based on a stronger assumption than one-way functions) were presented in [12].

This paper – a feasibility study of OT extensions. In this paper, we ask the following questions:

1. *What is the minimal assumption required for constructing OT extensions?* It has been shown that one-way functions suffice, and that OT extensions cannot be carried out information theoretically [3]. However, it is theoretically possible that OT extensions can be achieved under a weaker assumption than that of the existence of one-way functions. Admittedly, it is hard to conceive of a cryptographic construction that is not information theoretic and does not require one-way functions. However, a proof that one-way functions really are necessary is highly desired.

2. *Can oblivious transfer be extended with adaptive security?* The known constructions of OT extensions maintain security only in the presence of static corruptions, where the set of corrupted parties is fixed before the protocol begins. This is because the messages sent by the sender in the constructions of [3, 12] are binding with respect to the sender's input strings, and so an adaptive simulator cannot explain a transcript in multiple ways. Nothing is known

about whether or not adaptively secure OT extensions exist without assuming erasures[1].

3. *How many oblivious transfers are needed for extensions?* In the constructions of [3, 12], one must start with $O(n)$ oblivious transfers where n is the security parameter. These constructions can also be made to work when a superlogarithmic number $\omega(\log n)$ of oblivious transfers are given. However, they completely break down if $O(\log n)$ oblivious transfers only are available. We ask whether or not it is possible to extend a logarithmic number of oblivious transfers.

We prove the following theorems:

Theorem 1.1 *If there exists an OT extension protocol from n to $n+1$ (with security in the presence of static semi-honest adversaries), then there exist one-way functions.*

Thus, one-way functions are *necessary and sufficient* for OT extensions.

Theorem 1.2 *If there exists an OT extension protocol from n to $n+1$ that is secure in the presence of* adaptive *semi-honest adversaries, then there exists an oblivious transfer protocol that is secure in the presence of static semi-honest adversaries.*

This means that the construction of an adaptive OT extension protocol involves constructing statically secure oblivious transfer from scratch. This can still be meaningful, since adaptive oblivious transfer cannot be constructed from static oblivious transfer in a black-box manner [15]. However, it does demonstrate that adaptive OT extensions based on weaker assumptions than those necessary for static oblivious transfer do not exist.

Theorem 1.3 *If there exists an OT extension protocol from $f(n) = \mathcal{O}(\log n)$ to $f(n) + 1$ that is secure in the presence of static malicious adversaries, then there exists an oblivious transfer protocol that is secure in the presence of static malicious adversaries.*

This demonstrates that in order to extend only a logarithmic number of oblivious transfers (with security for *malicious* adversaries), one has to construct an oblivious transfer protocol from scratch. Thus, meaningful OT extensions exist only if one starts with a superlogarithmic number of oblivious transfers.

We stress that all of our results are unconditional, and are not black-box separations. Rather, we construct concrete one-way functions and OT protocols in order to prove our results.

Our results provide quite a complete picture regarding the feasibility of constructing OT extensions. The construction of [3] is optimal in terms of the computational assumption, and the constructions of [3, 12] are optimal in terms of the number of oblivious transfers one starts with. Finally, the fact that no OT extensions are known for the setting of adaptive corruptions is somewhat explained by Theorem 2.

[1] Note that in the erasures model, an OT extension can be constructed from one-way functions using the original construction of Beaver and the two-party computation protocol of [**?**] that is adaptively secure with erasures and is based on Yao's protocol.

Open questions. Theorem 2 shows that there do not exist adaptively secure OT extensions based on weaker assumptions than what is needed for *statically secure* OT. However, we do not know how to construct an adaptively secure OT extension even from statically secure OT. Thus, the question of whether or not it is possible to construct an adaptively secure OT extension from an assumption weaker than adaptive OT is still open.

Theorem 3 holds only with respect to OT-extensions that are secure against *malicious* adversaries. For the case of semi-honest adversaries, the question of whether one can construct an an OT-extension from $f(n) = \mathcal{O}(\log n)$ to $f(n) + 1$ from an assumption weaker than statically secure OT protocol is open.

In this paper, we have investigated OT extensions. However, the basic question of extending a cryptographic primitive using a weaker assumption than that needed for obtaining the primitive from scratch is of interest in other contexts as well. For example, hybrid encryption (where one encrypts a symmetric key using an asymmetric scheme, and then encrypts the message using a symmetric scheme) is actually an extension of public-key encryption that requires one-way functions only.

A primitive that could certainly benefit from a study such as this one is *key agreement*. In this context, the question is whether it is possible for two parties to agree on an $m + 1$-bit long key, given an m-bit key, under assumptions that are weaker than those required for constructing a secure key-agreement from scratch. In the basic case, it is clear that OWFs are necessary and sufficient for any nontrivial KA extension that starts with n bits (where n is the security parameter). A more interesting question regarding this problem relates to the adaptive setting. Specifically, since adaptive key agreement is very expensive, it would be very beneficial if one could extend this primitive more efficiently and/or under weaker assumptions.

2 Definitions and Notations

We denote the security parameter by n, and we denote by U_n a random variable uniformly distributed over $\{0,1\}^n$. We say that a function $\mu : \mathbb{N} \to \mathbb{N}$ is negligible if for every positive polynomial $p(\cdot)$ and all sufficiently large n it holds that $\mu(n) < \frac{1}{p(n)}$. We use the abbreviation PPT to denote probabilistic polynomial-time. We denote the bits of a string $x \in \{0,1\}^n$ by $x_1, \ldots, x_n$; for a subscripted string x_b, we denote the bits by $x_b^1, \ldots, x_b^n$. In addition, for strings $x_0, x_1, \sigma \in \{0,1\}^n$ we denote by x_σ the string $x_{\sigma_1}^1, \ldots, x_{\sigma_n}^n$.

Definition 2.1 *Let $X = \{X(a,n)\}_{a\in\{0,1\}^*, n\in\mathbb{N}}$ and $Y = \{Y(a,n)\}_{a\in\{0,1\}^*, n\in\mathbb{N}}$ be two distribution ensembles. We say that X and Y are* computationally indistinguishable*, denoted $X \stackrel{c}{\equiv} Y$, if for every PPT machine D, every $a \in \{0,1\}^*$, every positive polynomial $p(\cdot)$ and all sufficiently large n:*

$$\left|\Pr\left[D(X(a,n), 1^n) = 1\right] - \Pr\left[D(Y(a,n), 1^n) = 1\right]\right| < \frac{1}{p(n)}.$$

We say that X and Y are statistically close*, denoted $X \stackrel{s}{\equiv} Y$, if for every $a \in \{0,1\}^*$, every positive polynomial $p(\cdot)$ and all sufficiently large n:*

$$SD(X,Y) \stackrel{\text{def}}{=} \frac{1}{2} \cdot \sum_{\alpha} \left|\Pr[X(a,n) = \alpha] - \Pr[Y(a,n) = \alpha]\right| < \frac{1}{p(n)}.$$

Interactive Protocols. Let $\pi = \langle A, B \rangle$ be an interactive protocol for computing a functionality f. We denote $f = (f_A, f_B)$, where f_A is the first output of f (for party A) and f_B is the second output of f (for party B). For inputs x_A and x_B of A and B (respectively) and random tapes r_A and r_B, we denote by $\text{TRANS}^\pi(x_A, x_B, r_A, r_B)$ the transcript obtained by running π on inputs x_A and x_B and random tapes r_A and r_B, and by $\text{TRANS}^\pi(x_A, x_B)$ the random variable describing $\text{TRANS}^\pi(x_A, x_B; r_A, r_B)$ where r_A and r_B are uniformly chosen.

The random variable $\text{VIEW}^\pi_A(x_A, x_B)$ denotes the view of the party A in an execution of π with inputs x_A for A and x_B for B, where the random tapes of the parties are uniformly chosen. Note that a view of a party contains its input, randomness and the messages it has received during the execution.

The random variable $\text{OUTPUT}^\pi_A(x_A, x_B)$ denotes the output of the party A in an execution of π with inputs x_A for A and x_B for B, where the random tapes of the parties are uniformly chosen.

Definition 2.2 *Let $f(\cdot,\cdot)$ be a deterministic binary functionality, let $\pi = \langle A, B \rangle$ be an interactive protocol and let n be the security parameter. We say that π* computes the functionality *f if there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that for all n, x_A and x_B:*

$$\Pr\left[\langle A(1^n, x_A), B(1^n, x_B)\rangle = (f_A(x_A, x_B), f_B(x_A, x_B))\right] \geq 1 - \mathsf{negl}(n).$$

Definition 2.3 *Let $\pi = \langle A, b \rangle$ be a protocol that computes a deterministic functionality $f = (f_A, f_B)$. We say that π* securely computes *f* in the presence of static semi-honest adversaries *if there exist two probabilistic polynomial-time algorithms $\mathcal{S}_A$ and $\mathcal{S}_B$ such that:*

$$\left\{\mathcal{S}_A(1^n, x_A, f_A(x_A, x_B))\right\}_{x_A, x_B \in \{0,1\}^*, n \in \mathbb{N}} \stackrel{\text{c}}{\equiv} \left\{\text{VIEW}^\pi_A(1^n, x_A, x_B)\right\}_{x_A, x_B \in \{0,1\}^*, n \in \mathbb{N}}$$

and

$$\left\{\mathcal{S}_B(1^n, x_B, f_B(x_A, x_B))\right\}_{x_A, x_B \in \{0,1\}^*, n \in \mathbb{N}} \stackrel{\text{c}}{\equiv} \left\{\text{VIEW}^\pi_B(1^n, x_A, x_B)\right\}_{x_A, x_B \in \{0,1\}^*, n \in \mathbb{N}}$$

Security in the presence of malicious adversaries. To define security in the presence of malicious adversaries, we use the ideal/real framework as defined by Canetti in [4]. Loosely speaking, in this approach we formalize the real-life computation as a setting where the parties, given their private inputs, interact according to the protocol in the presence of a real-life adversary that controls a set of corrupted parties. The real-life adversary can be either static (where the set of corrupted parties is fixed before the protocol begins) or adaptive (where the adversary can choose to corrupt parties during the protocol execution based on what it sees). At the end of the computation, the honest parties output what is specified by the protocol and the adversary outputs some arbitrary function of its view. If the adversary is adaptive, there is an additional entity $\mathcal{Z}$, called the environment, who sees the output of all of the parties. In addition, there is a "postexecution phase", where $\mathcal{Z}$ can instruct the adversary to also corrupt parties after the execution of the protocol ends (and the transcript is fixed, implying that "rewinding" is no longer allowed). At the end of the postexecution phase, $\mathcal{Z}$ outputs some function of its view.

Next we consider an ideal process, where an ideal-world adversary controls a set of corrupted parties. Then, in the computation phase, all parties send their inputs to some incorruptible trusted party. The ideal-world adversary sends inputs on behalf of the corrupted parties. The trusted party evaluates the function and hands each party its output. The honest parties then output whatever

they received from the trusted party and the ideal-world adversary outputs some arbitrary value. Similarly to the real-life setting, in the case of adaptive security, there is an environment $\mathcal{Z}$ who sees all outputs and can instruct the adversary to also corrupt parties in the postexecution phase. At the end of the postexecution phase, $\mathcal{Z}$ outputs some function of its view.

Loosely speaking, a protocol π is secure in the presence of static malicious adversaries, if for every static malicious real-life adversary $\mathcal{A}$, there exists a static malicious ideal-world adversary $\mathcal{SIM}$ such that the distribution obtained in a real-life execution of π with adversary $\mathcal{A}$ is indistinguishable from the distribution obtained in a ideal-world with adversary $\mathcal{SIM}$. Likewise, a protocol π is secure in the presence of adaptive malicious adversaries, if for every adaptive malicious real-life adversary $\mathcal{A}$ and environment $\mathcal{Z}$, there exists an adaptive malicious ideal-world adversary $\mathcal{SIM}$ such that the output of $\mathcal{Z}$ in a real-life execution of π with adversary $\mathcal{A}$ is indistinguishable from its output in a ideal-world with adversary $\mathcal{SIM}$.

Security in the presence of adaptive semi-honest adversaries is defined in the same way as adaptive malicious adversaries, except that the adversary only sees the internal state of a corrupted party but cannot instruct it to deviate from the protocol specification. For full definitions see [4].

The hybrid model. Let ϕ be a functionality. The ϕ-hybrid model is defined as follows. The real-life model for protocol π is augmented with an incorruptible trusted party T for evaluating the functionality ϕ, and the parties are allowed to make calls to the ideal functionality ϕ by sending their ϕ-inputs to T. If we consider malicious adversaries, the adversary specifies the inputs of all parties under its control. If the adversary is semi-honest, then even the corrupted parties hand T inputs as specified by the protocol π. At each invocation of ϕ, the trusted party T sends the parties their respective outputs.

We stress that if π is in the ϕ-hybrid model, then a view of a party A contains also the inputs sent by A to the functionality ϕ and the outputs sent to A by T computing ϕ.

Oblivious transfer and extensions. We are now ready to define oblivious transfer and OT extensions.

Definition 2.4 *The* bit oblivious transfer functionality *OT is defined by $OT((b_0, b_1), \sigma) = (\lambda, b_\sigma)$. The* parallel oblivious transfer functionality *$m \times OT$ is defined for strings $x_0, x_1, \sigma \in \{0,1\}^m$ as follows: $m \times OT((x_0, x_1), \sigma) = (\lambda, (x^1_{\sigma_1}, \ldots, x^m_{\sigma_m})) = (\lambda, x_\sigma)$ (recall that x_σ denotes the string $x^1_{\sigma_1}, \ldots, x^n_{\sigma_n}$).*

We denote by OT^k the ideal functionality of k independent OT computations. We stress that OT^k is not the same as $k \times OT$, since in the latter all of the inputs are given at once whereas in OT^k the inputs can be chosen over time (in particular, the receiver can choose its inputs as a function of the previous outputs it received). Using this notation, we have that an OT extension protocol is a protocol that securely computes $m \times OT$ given access to OT^k, where $k < m$. Formally:

Definition 2.5 (OT-extension) *Let π be a protocol and let $k, m : \mathbb{N} \to \mathbb{N}$ be two functions where $k(n) < m(n)$ for all n. We say that π is an* OT-extension from $k = k(n)$ to $m = m(n)$ *if π securely computes the $m \times OT$ functionality in the OT^k-hybrid model.*

OT extensions – two technical propositions. We present two propositions that we use throughout the paper. Beaver showed that OT can be precomputed [2]. That is, it is possible to first compute OT on random inputs and then use the result to later compute an OT on any input. Stated formally:

Proposition 2.6 (Beaver [2]) *Let $m = m(n)$ be a polynomial. If there exists a protocol that securely computes the $m \times OT$ functionality, then there exists a protocol that securely computes the OT^m ideal functionality.*

Proposition 2.6 shows that Definition 2.5 could have been stated as a protocol that securely computes OT^m in the OT^k (or even the $k \times OT$) hybrid model.

The fact that a single extension implies many has been stated many times in the literature (e.g., [3]) and is well accepted folklore, but has not been formally proved. We sketch a proof of this here. We stress that this holds irrespectively of how many oblivious transfers you start with (even if only a *constant number*), as long as only a polynomial number of transfers are derived. We state the proposition for adaptive malicious adversaries and observe that it holds for all four combinations of static/adaptive and semi-honest/malicious adversaries.

Proposition 2.7 *Let $f : \mathbb{N} \to \mathbb{N}$ be* any *polynomially-bounded function, and let n be the security parameter. If there exists a protocol π that is an OT-extension from $f(n)$ to $f(n)+1$ that is secure in the presence of adaptive malicious adversaries, then for every polynomial $p(\cdot)$ there exists an OT-extension protocol from $f(n)$ to $p(n)$ that is secure in the presence of adaptive malicious adversaries.*

Proof Sketch: First, we remark that any secure extension protocol π can be converted into a secure extension protocol π' with the property that all of the $f(n)$ calls to the ideal OT are made at the beginning of the protocol. We actually divide the execution of π' into two phases: in the first phase the parties make $f(n)$ calls to an ideal OT, and in the second phase they use the results of the first phase to compute the OT calls in the original extension protocol π. This transformation follows easily from the fact that OT can be precomputed [2].

We now use π' to construct a new protocol $\tilde{\pi}$ that is an OT-extension from $f(n)$ to $p(n)$. Protocol $\tilde{\pi}$ iteratively invokes π' in the following way. First, $f(n)$ calls are made to an ideal OT. Then, invoke phase 2 of π' to obtain $f(n)+1$ new OT's using the result of the $f(n)$ OT's from the previous iteration. The first $f(n)$ of these OT's are used to once again obtain $f(n)+1$ OT's by invoking phase 2 of π'. Repeating this process $p(n)$ times, and noting that there is one "spare" OT in each iteration, we have that $p(n)$ OTs remain and can be used for actual transfers.

This is the same methodology as that used to prove that the existence of pseudorandom generators that stretch the input by a single bit implies the existence of pseudorandom generators that stretch the input by any polynomial amount (see [9, Sec. 3.3.2]). The proof of security also follows a hybrid argument in the same way. We stress that since we use a hybrid argument on the number of times the original extension is applied, it makes no difference how many OT calls are used in the original extension protocol. Thus, this holds also for small $f(n)$. ■

2.1 A Lemma on Statistical Distance

Lemma 2.8 *Let D_1 and D_2 be two distributions over a set U and let E be an event such that $\Pr_{D_1}[E] = \Pr_{D_2}[E]$. Then, it holds that*

$$SD(D_1, D_2) \leq SD(D_1 \mid E, D_2 \mid E) + \Pr_{D_1}[\neg E]$$

Proof:

$$\begin{aligned}
SD(D_1, D_2) &= \frac{1}{2}\sum_{x \in U} \Big|\Pr_{D_1}[x] - \Pr_{D_2}[x]\Big| \\
&= \frac{1}{2}\sum_{x \in U} \Big|\Pr_{D_1}[x \mid E] \cdot \Pr_{D_1}[E] + \Pr_{D_1}[x \mid \neg E] \cdot \Pr_{D_1}[\neg E] \\
&\qquad - \Pr_{D_2}[x \mid E] \cdot \Pr_{D_2}[E] + \Pr_{D_2}[x \mid \neg E] \cdot \Pr_{D_2}[\neg E]\Big| \\
&= \frac{1}{2}\sum_{x \in U} \Big|\Pr_{D_1}[x \mid E] \cdot \Pr_{D_1}[E] - \Pr_{D_2}[x \mid E] \cdot \Pr_{D_2}[E] \\
&\qquad + \Pr_{D_1}[x \mid \neg E] \cdot \Pr_{D_1}[\neg E] - \Pr_{D_2}[x \mid \neg E] \cdot \Pr_{D_2}[\neg E]\Big| \\
&\leq \frac{1}{2}\sum_{x \in U} \Big|\Pr_{D_1}[x \mid E] \cdot \Pr_{D_1}[E] - \Pr_{D_2}[x \mid E] \cdot \Pr_{D_1}[E]\Big| \\
&\qquad + \frac{1}{2}\sum_{x \in U} \Big|\Pr_{D_1}[x \mid \neg E] \cdot \Pr_{D_1}[\neg E] - \Pr_{D_2}[x \mid \neg E] \cdot \Pr_{D_1}[\neg E]\Big| \\
&= \Pr_{D_1}[E] \cdot SD(D_1 \mid E, D_2 \mid E) + \Pr_{D_1}[\neg E] \cdot SD(D_1 \mid \neg E, D_2 \mid \neg E) \\
&\leq SD(D_1 \mid E, D_2 \mid E) + \Pr_{D_1}[\neg E]
\end{aligned}$$

■

3 OT Extensions Imply One-Way Functions

In this section we show that the existence of an OT extension protocol implies the existence of one-way functions. We prove the theorem for any OT extension that is secure in the presence of static semi-honest adversaries (thus the theorem also holds when the OT extension is secure in the presence of adaptive and/or malicious adversaries).

Theorem 3.1 *If there exists a protocol that is an OT-extension from n to $n+1$ (where n is the security parameter) that is secure for static semi-honest adversaries, then there exist one-way functions.*

Proof: By Proposition 2.7, if there exists an OT extension protocol from n to $n+1$ then there exists an OT extension protocol from n to $2n+1$. We therefore prove the theorem by showing that the existence of a protocol π that is an OT-extension from n to $2n+1$ implies the existence of two polynomial-time constructible probability ensembles that are computationally indistinguishable and yet their statistical distance is noticeable. The fact that this implies one-way functions was shown in [8]. We begin by defining the probability ensembles and then provide intuition as to why they fulfill the above property.

Let $X_0, X_1, X'_0, X'_1, \Sigma$ be (dependent) random variables chosen as follows:

1. $\Sigma \in_R \{0,1\}^{2n+1}$ is a uniformly distributed string (representing the receiver's input)

2. $X_0, X_1, X'_0, X'_1 \in \{0,1\}^{2n+1}$ (representing the sender's possible inputs) are uniformly distributed under the constraint that for every $i = 1, \ldots, 2n+1$ it holds that $X^i_{\Sigma^i} = X'^i_{\Sigma^i}$, where $\Sigma = \Sigma_1, \ldots, \Sigma_{2n+1}$ and $X_0 = X^1_0, \ldots, X^{2n+1}_0$ (likewise for X_1, X'_0, X'_1). (Thus, the pairs (X_0, X_1) and (X'_0, X'_1) "agree" on the bits chosen by Σ and are independent otherwise.)

Let $\text{TRANS}^\pi(x_0, x_1, \sigma)$ be a random variable over the transcript of π on sender-inputs (x_0, x_1) and receiver-input σ. We stress that the transcript contains all of the messages sent between the parties, but does *not* contain the n input/output values sent by the parties to the ideal OT functionality within the extension protocol. We are now ready to define the two probability ensembles $\mathcal{E}^1 = \{\mathcal{E}^1_n\}_{n\in\mathbb{N}}$ and $\mathcal{E}^2 = \{\mathcal{E}^2_n\}_{n\in\mathbb{N}}$:

$$\mathcal{E}^1_n = (X_0, X_1, \Sigma, \text{TRANS}^\pi(X_0, X_1, \Sigma)) \quad \text{and} \quad \mathcal{E}^2_n = (X'_0, X'_1, \overline{\Sigma}, \text{TRANS}^\pi(X_0, X_1, \Sigma)),$$

where $\overline{\Sigma}$ denotes the bitwise complement of Σ. Observe that in $\mathcal{E}^1$ the transcript is generated from the given inputs (X_0, X_1, Σ), whereas in $\mathcal{E}^2$ the given inputs are (X'_0, X'_1) and $\overline{\Sigma}$ (and (X'_0, X'_1) "agree" with (X_0, X_1) on Σ and are independent of each other on $\overline{\Sigma}$).

Intuitively, these ensembles are computationally indistinguishable by the privacy properties of oblivious transfer (the change from (X_0, X_1) to (X'_0, X'_1) cannot be distinguished or a receiver with input Σ could learn more than allowed, and the change from Σ to $\overline{\Sigma}$ cannot be distinguished or the sender could learn something about the receiver's input). Furthermore, they are statistically far apart because the transcript must contain meaningful information about the inputs being used (in which case, the transcript will be consistent with the inputs in $\mathcal{E}^1$ but not in $\mathcal{E}^2$). In order to see why this is the case, observe that since the number of calls made to the ideal *OT* functionality is only n, it cannot be the case that all information regarding the inputs is transferred via the use of the ideal *OT* calls. Thus the transcript itself must contain some meaningful information, and this information will not be consistent in $\mathcal{E}^2$.

We begin by proving that $\mathcal{E}^1$ and $\mathcal{E}^2$ are computationally indistinguishable. Intuitively, this follows from the privacy property of secure oblivious transfer.

Lemma 3.2 *The ensembles $\mathcal{E}^1$ and $\mathcal{E}^2$ are computationally indistinguishable.*

Proof: We prove the lemma by separately considering the privacy guarantees with respect to the receiver's input and the sender's inputs. Towards this goal, consider the following hybrid ensemble: Let $\mathcal{E}^h = \{\mathcal{E}^h_n\}_{n\in\mathbb{N}}$ be the following probability ensemble:

$$\mathcal{E}^h_n = (X'_0, X'_1, \Sigma, \text{TRANS}^\pi(X_0, X_1, \Sigma)).$$

Note that in $\mathcal{E}^h_n$ we change only the inputs of the sender, whereas in $\mathcal{E}^2_n$ both the inputs of the sender and the receiver are changed (and in $\mathcal{E}^1_n$ none of the inputs is changed). We prove the claim by proving that $\mathcal{E}^1$ and $\mathcal{E}^h$ are computationally indistinguishable and $\mathcal{E}^h$ and $\mathcal{E}^2$ are computationally indistinguishable. We sketch the proof of computational indistinguishability:

1. The only difference between $\mathcal{E}^1$ and $\mathcal{E}^h$ is that $\mathcal{E}^1$ contains the actual input used by the sender whereas $\mathcal{E}^h$ outputs a pair of strings that are random in the locations that are not part of the receiver's output. Intuitively, these are indistinguishable since otherwise a corrupted receiver could obtain information about the sender's inputs that it did not choose, in contradiction

to the security of oblivious transfer. This can be formalized by defining an experiment in which the receiver's input σ is chosen at random, and then two sets of sender inputs are chosen randomly under the constraint that they are the same for the bits to be received for the receiver input σ. The oblivious transfers are run using one of the two sender inputs, and an adversary receiving the receiver's view attempts to guess which one was used. It is easy to show that the privacy of oblivious transfer implies that no adversary can succeed in guessing correctly with probability non-negligibly greater than 1/2.

2. The only difference between $\mathcal{E}^h$ and $\mathcal{E}^2$ is that in $\mathcal{E}^h$ the receiver's actual input appears whereas in $\mathcal{E}^2$ the complement of the receiver's input appears. As above, these are indistinguishable since otherwise a corrupted sender could obtain some information about the receiver's input, in contradiction to the security of oblivious transfer. Again, this can be formalized by defining an experiment where a string σ is chosen at random and given to the sender. Then, the oblivious transfer implies that no adversary can succeed in guessing if the receiver input was σ or $\overline{\sigma}$ with probability non-negligibly greater than 1/2.

The formal proofs of the above are straightforward and are therefore omitted. ∎

We now prove that the ensembles are statistically far apart.

Lemma 3.3 *There exists a polynomial $p(\cdot)$ such that for all large enough n's, $SD(\mathcal{E}^1_n, \mathcal{E}^2_n) \geq \frac{1}{p(n)}$.*

Proof: Given the input $\sigma \in \{0,1\}^{2n+1}$ of the receiver and a transcript t, let $\{(\tau_i, \omega_i)\}_{i=1}^n$ denote a sequence of size n containing the inputs $\{\tau_i\}_{i=1}^n$ sent by the receiver in the n calls to the ideal OT and the respective outputs $\{\omega_i\}_{i=1}^n$ obtained from these calls. We use the following notation:

- For every sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$, let $\mathsf{R}_{\mathsf{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ denote the set of all random tapes of the receiver that are consistent with σ, t and $\{(\tau_i, \omega_i)\}_{i=1}^n$. Moreover, for every string $x \in \{0,1\}^{2n+1}$, let $\mathsf{R}_{\mathsf{out}}(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ denote the set of all random tapes of the receiver that are consistent with σ, t and $\{(\tau_i, \omega_i)\}_{i=1}^n$ and lead the receiver to output x. Note that for every x, it holds that $\mathsf{R}_{\mathsf{out}}(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) \subseteq \mathsf{R}_{\mathsf{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$. Let $p_\pi(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ denote the ratio between the size of these two sets; that is:
$$p_\pi(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) = \frac{|\mathsf{R}_{\mathsf{out}}(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{R}_{\mathsf{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}$$
- Let $\textsc{LikelySet}(\sigma, t)$ denote the set of all strings $x \in \{0,1\}^{2n+1}$ for which there *exists* a sequence of n pairs $\{(\tau_i, \omega_i)\}_{i=1}^n$ such that
$$p_\pi(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) > \frac{1}{2}$$
($\textsc{LikelySet}(\sigma, t)$ is empty if no such x exists). From the definition, for a given receiver-input σ and transcript t, the set $\textsc{LikelySet}(\sigma, t)$ contains all of the strings x for which there exists a sequence $\{(\tau_i, \omega_i)\}$ so that the receiver outputs x after the execution of π with probability greater than 1/2.

To prove the statistical distance, we construct an unbounded distinguisher $\mathcal{A}$ and show the existence of a polynomial $p(\cdot)$ such that for all sufficiently large n's:
$$\left|\Pr[\mathcal{A}(\mathcal{E}^1_n) = 1] - \Pr[\mathcal{A}(\mathcal{E}^2_n) = 1]\right| \geq \frac{1}{p(n)}$$

We define our (computationally unbounded) distinguisher $\mathcal{A}$ as follows: $\mathcal{A}$ receives as input a tuple $(\tilde{x}_0, \tilde{x}_1, \tilde{\sigma}, t)$ that was chosen from either $\mathcal{E}^1$ or $\mathcal{E}^2$ and outputs 1 if and only if $\tilde{x}_{\tilde{\sigma}} \in \textsc{LikelySet}(\tilde{\sigma}, t)$. Observe that $\tilde{x}_{\tilde{\sigma}}$ is the correct receiver output in the case that the parties' inputs were $\tilde{x}_0, \tilde{x}_1, \tilde{\sigma}$.

The intuition behind this construction is as follows. If $(\tilde{x}_0, \tilde{x}_1, \tilde{\sigma}, t)$ was sampled from $\mathcal{E}^1$, then $\tilde{x}_0, \tilde{x}_1$ and $\tilde{\sigma}$ are the inputs used to generate the transcript t, and by the correctness of the protocol the receiver should output $\tilde{x}_{\tilde{\sigma}}$ with probability close to 1. Thus, with high probability $\tilde{x}_{\tilde{\sigma}} \in \textsc{LikelySet}(\tilde{\sigma}, t)$. In contrast, if $(\tilde{x}_0, \tilde{x}_1, \tilde{\sigma}, t)$ was sampled from $\mathcal{E}^2 = (X'_0, X'_1, \overline{\Sigma}, \textsc{Trans}(X_0, X_1, \Sigma))$, then t is a transcript generated from (x_0, x_1, σ), where x_0, x_1 are uniform and independent of $(\tilde{x}_0, \tilde{x}_1)$ on the bits chosen by $\tilde{\sigma}$, and $\tilde{\sigma} = \overline{\sigma}$. This implies that $\tilde{x}_{\tilde{\sigma}} = \tilde{x}_{\overline{\sigma}}$ is a random string of size $2n+1$ that is independent of t and so the probability that $\tilde{x}_{\tilde{\sigma}} \in \textsc{LikelySet}(\tilde{\sigma}, t)$ cannot be too large.

We show that $\mathcal{A}$ distinguishes $\mathcal{E}^1$ from $\mathcal{E}^2$ with probability close to $1/2$. Surprisingly, the main challenge is actually to show that $\mathcal{A}$ outputs 1 when receiving a sample from $\mathcal{E}^1$ with probability close to 1. We explain the difficulty involved at the beginning of the proof of Claim 3.5.

Claim 3.4 *For every n, it holds that $\Pr[\mathcal{A}(\mathcal{E}^2_n) = 1] \leq \frac{1}{2}$.*

Proof: Recall that upon input $(\tilde{x}_0, \tilde{x}_1, \tilde{\sigma}, t)$, distinguisher $\mathcal{A}$ outputs 1 if and only if $\tilde{x}_{\tilde{\sigma}} \in \textsc{LikelySet}(\tilde{\sigma}, t)$; that is, if and only if there exists a sequence of pairs $\{(\tau_i, \omega_i)\}_{i=1}^n$ such that $p_\pi(\tilde{x}_{\tilde{\sigma}}, \tilde{\sigma}, t, \{(\tau_i, \omega_i)\}_{i=1}^n) > \frac{1}{2}$. As we have described, in this case of ensemble $\mathcal{E}^2$, the string $\tilde{x}_{\tilde{\sigma}}$ is *independent* of t. To stress this point, the distribution $\mathcal{E}^2$ can be generated by choosing X_0, X_1, Σ and generating t, and only then choosing the bits of X'_0, X'_1 corresponding to $\overline{\Sigma}$ (observe that $\tilde{x}_{\tilde{\sigma}}$ corresponds exactly to these bits chosen last). Now, for every given $(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ there exists at most one x such that $p_\pi(x, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) > \frac{1}{2}$ (since it is required that the probability be strictly greater than $1/2$). Since t depends only on random coins generated before the remaining bits of X'_0, X'_1 and so $\tilde{x}_{\tilde{\sigma}}$ are chosen, this implies that for every series $\{(\tau_i, \omega_i)\}_{i=1}^n$,

$$\Pr\left[p_\pi(\tilde{x}_{\tilde{\sigma}}, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) > \frac{1}{2}\right] = \frac{1}{2^{2n+1}}.$$

We therefore have that for every n,

$$\begin{aligned}
\Pr\left[\mathcal{A}(\mathcal{E}^2_n) = 1\right] &= \Pr\left[\exists \{(\tau_i, \omega_i)\}_{i=1}^n \text{ s.t. } p_\pi(\tilde{x}_{\tilde{\sigma}}, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) > \frac{1}{2}\right] \\
&\leq \sum_{\{(\tau_i, \omega_i)\}_{i=1}^n} \Pr\left[p_\pi(\tilde{x}_{\tilde{\sigma}}, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) > \frac{1}{2}\right] \\
&\leq 2^{2n} \cdot \frac{1}{2^{2n+1}} = \frac{1}{2}\ .
\end{aligned}$$

■

Denote by $\textsc{output}^\pi_R(x_0, x_1, \sigma; 1^n)$ the output of the receiver R after an execution with sender-inputs (x_0, x_1), receiver-input σ, and security parameter n. We prove:

Claim 3.5 *Let $\mu(\cdot)$ be the negligible function so that $\Pr\left[\textsc{output}^\pi_R(x_0, x_1, \sigma; 1^n) = x_\sigma)\right] \geq 1 - \mu(n)$ (from the correctness requirement). Then, for every n it holds that $\Pr[\mathcal{A}(\mathcal{E}^1_n) = 1)] \geq 1 - 2\mu(n)$.*

Proof: Recall that $\mathcal{E}^1$ samples tuples (x_0, x_1, σ, t) such that t is a transcript of π on inputs x_0, x_1 and σ, where x_0, x_1 and σ are uniformly chosen. Intuitively, this claim follows from the correctness of the oblivious transfer protocol. That is, if $x_\sigma \notin \text{LIKELYSET}(\sigma, t)$ then the receiver would output the correct output x_σ with probability less than $1/2$, contradicting the correctness requirement. Unfortunately, this intuitive argument is far more involved to prove. The reason for this is that the correctness requirement is based on the probability over the random coins of *both parties*. In contrast, LIKELYSET is defined based on the random coins of the receiver only. In order to demonstrate why this could be problematic, consider the situation where for any given transcript t and sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$, the majority of receiver coins r_R result in an *incorrect* output. However, there are only very few sender coins r_S that are consistent with t and the bad receiver coins r_R. Therefore, when taking the probability over both the sender and receiver coins, the incorrect output is received with only very small probability. However, when considering the receiver's coins only, the incorrect output is obtained very often. We stress that such an event is easily shown to not be possible in a standard protocol where the transcript contains all information. This is because there is no dependence between the sender's coins and the receiver's coins, for all possible coins that are consistent with the transcript. However, in our scenario where ideal OT calls are included (and the inputs and outputs to these calls are not part of the transcript), such dependence may be introduced via the ideal OT calls. Proving that such a case cannot occur constitutes the majority of the proof of this claim.

For inputs x_0, x_1, and σ, let $\mathsf{Good}(x_0, x_1, \sigma)$ denote the set of all transcripts t such that $x_\sigma \in \text{LIKELYSET}(\sigma, t)$; i.e., $\mathsf{Good}(x_0, x_1, \sigma) = \{t \mid x_\sigma \in \text{LIKELYSET}(\sigma, t)\}$. Intuitively, this is the set of all transcripts that are "good" in the sense that in those executions the receiver (may) output the correct output with a good probability (it won't necessarily output the correct output because this just means that there exists a sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$ for which it outputs the correct output with probability greater than $1/2$). Recall that $\mathcal{A}$ on input (x_0, x_1, σ, t) returns 1 if and only if $x_\sigma \in \text{LIKELYSET}(\sigma, t)$ and hence $\mathcal{A}$ outputs 1 if and only if $t \in \mathsf{Good}(x_0, x_1, \sigma)$. Thus, it suffices to prove that $\Pr[t \in \mathsf{Good}(x_0, x_1, \sigma)] > 1 - 2\mu(n)$, when (x_0, x_1, σ, t) are sampled from $\mathcal{E}^1$.

In order to prove this, we use the fact that

$$\begin{aligned}
&\Pr[\text{OUTPUT}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma] \\
&\quad = \Pr[\text{OUTPUT}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \in \mathsf{Good}(x_0, x_1, \sigma)] \cdot \Pr[t \in \mathsf{Good}(x_0, x_1, \sigma)] \\
&\qquad + \Pr[\text{OUTPUT}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \notin \mathsf{Good}(x_0, x_1, \sigma)] \cdot \Pr[t \notin \mathsf{Good}(x_0, x_1, \sigma)] \\
&\quad \leq \Pr[t \in \mathsf{Good}(x_0, x_1, \sigma)] \\
&\qquad + \Pr[\text{OUTPUT}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \notin \mathsf{Good}(x_0, x_1, \sigma)] \cdot \Pr[t \notin \mathsf{Good}(x_0, x_1, \sigma)]
\end{aligned}$$

Below, we will prove that

$$\Pr[\text{OUTPUT}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \notin \mathsf{Good}(x_0, x_1, \sigma)] \leq \frac{1}{2}. \tag{1}$$

Combining the above calculation with Eq. (1) and with the correctness requirement of the protocol stating that $\Pr[\text{OUTPUT}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma] \geq 1 - \mu(n)$, we have:

$$\begin{aligned}
1 - \mu(n) &\leq \Pr[t \in \mathsf{Good}(x_0, x_1, \sigma)] + \frac{1}{2} \cdot \Pr[t \notin \mathsf{Good}(x_0, x_1, \sigma)] \\
&= 1 - \frac{1}{2} \cdot \Pr[t \notin \mathsf{Good}(x_0, x_1, \sigma)]
\end{aligned}$$

and so $\Pr[t \notin \mathsf{Good}(x_0, x_1, \sigma)] \leq 2\mu(n)$. Thus, $\Pr[\mathcal{A}(\mathcal{E}_n^1) = 1] = \Pr[t \in \mathsf{Good}(x_0, x_1, \sigma)] > 1 - 2\mu(n)$ as required.

It therefore remains to prove Eq. (1) in order to prove Claim 3.5. By the definition of Good, for every $t \notin \mathsf{Good}(x_0, x_1, \sigma)$ we have that $x_\sigma \notin \textsc{LikelySet}(\sigma, t)$, which by the definition of $\textsc{LikelySet}(\sigma, t)$ implies that for *every* sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$, it holds that

$$p_\pi(x_\sigma, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) = \frac{|\mathsf{R}_{\mathsf{out}}(x_\sigma, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{R}_{\mathsf{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|} \leq \frac{1}{2}. \tag{2}$$

Fix x_0, x_1, σ and fix $t \notin \mathsf{Good}(x_0, x_1, \sigma)$. We prove Eq. (1) by showing that *for all* $\{(\tau_i, \omega_i)\}_{i=1}^n$

$$\Pr[\textsc{output}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \wedge \{(\tau_i, \omega_i)\}_{i=1}^n] \leq \frac{1}{2}.$$

For every $t \notin \mathsf{Good}(x_0, x_1, \sigma)$ and $\{(\tau_i, \omega_i)\}_{i=1}^n$ we define the following two sets (recall that x_0, x_1 and σ are fixed):

1. Let $\mathsf{RS}_{\mathsf{All}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ contain all pairs of random tapes (r_R, r_S) for which the execution $\langle S(x_0, x_1; r_S), R(\sigma; r_R)\rangle$ results in transcript t and the sequence of input/output ideal calls $\{(\tau_i, \omega_i)\}_{i=1}^n$.
2. Let $\mathsf{RS}_{\mathsf{good}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ contain all pairs of random tapes (r_R, r_S) for which the execution $\langle S(x_0, x_1; r_S), R(\sigma; r_R)\rangle$ results in transcript t, sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$ and receiver-output x_σ.

It follows immediately from the definition of these sets that

$$\Pr[\textsc{output}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \wedge \{(\tau_i, \omega_i)\}_{i=1}^n] = \frac{|\mathsf{RS}_{\mathsf{good}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{RS}_{\mathsf{All}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}. \tag{3}$$

In order to see this, denote by All the set of all pairs of random tapes, and observe that

$$\Pr[\textsc{output}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \wedge t \wedge \{(\tau_i, \omega_i)\}_{i=1}^n] = \frac{|\mathsf{RS}_{\mathsf{good}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{All}|}$$

and

$$\Pr[t \wedge \{(\tau_i, \omega_i)\}_{i=1}^n] = \frac{|\mathsf{RS}_{\mathsf{All}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{All}|}\ .$$

Observe that this is very similar to Eq. (2), except that Eq. (2) refers to $\mathsf{R}_{\mathsf{All}}$ and $\mathsf{R}_{\mathsf{out}}$ which are based on the receiver's random tape only, and here we refer to $\mathsf{RS}_{\mathsf{All}}$ and $\mathsf{RS}_{\mathsf{good}}$ which refer to both the receiver and sender's random tapes. Thus, it remains to show that they have the same ratio, and this will imply that $\Pr[\textsc{output}_R^\pi(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \wedge \{(\tau_i, \omega_i)\}_{i=1}^n] \leq 1/2$.

Let $\mathsf{S}_{\mathsf{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ be the set of all random tapes of the sender that are consistent with x_0, x_1, t and $\{(\tau_i, \omega_i)\}_{i=1}^n$. We prove:

$$\mathsf{RS}_{\mathsf{All}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n) = \mathsf{S}_{\mathsf{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n) \times \mathsf{R}_{\mathsf{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) \tag{4}$$

$$\mathsf{RS}_{\mathsf{good}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n) = \mathsf{S}_{\mathsf{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n) \times \mathsf{R}_{\mathsf{out}}(x_\sigma, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) \tag{5}$$

(Recall that this is trivial in the case that there are no ideal calls to a functionality. However, in this case, it is conceivable that the ideal calls may introduce dependence and thus it requires a proof; see Footnote 2 below.) We begin by proving Eq. (4). Let $r_S \in \mathsf{S}_{\mathsf{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ and let $r_R \in \mathsf{R}_{\mathsf{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$. We show that $(r_R, r_S) \in \mathsf{RS}_{\mathsf{All}}(t, \{(\tau_i, \omega_i)\}_{i=1}^n)$ by showing that the execution $\langle S(x_0, x_1; r_S), R(\sigma; r_R)\rangle$ results in transcript t and sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$.

This can be proved by a simple induction on the round number k. Assume that up to the k^{th} round, the execution $\langle S(x_0, x_1; r_S), R(\sigma; r_R)\rangle$ is consistent with t and the n pairs $\{(\tau_i, \omega_i)\}_{i=1}^n$; we show that this argument holds also after the $k+1^{th}$ round. There are three cases for the $k+1^{th}$ round:

- *The sender sends a message:* By the induction hypothesis, all the information that S has up to this point is consistent with t and $\{(\tau_i, \omega_i)\}_{i=1}^n$. Since $r_S \in \mathsf{S_{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$, it follows that the message sent by the sender in this round is consistent with t.
- *The receiver sends a message:* Exactly as above, using the fact that $r_R \in \mathsf{R_{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)$.
- *The parties make the j^{th} call to the ideal OT functionality:* By a similar argument to the previous cases, we deduce that the input sent by the sender to the ideal OT functionality is consistent with (τ_j, ω_j) and the input sent by the receiver is consistent with (τ_j, ω_j). Hence, letting m_0, m_1 be the input of the sender to the OT functionality, we have that $m_{\tau_j} = \omega_j$ and the input of the receiver is τ_j. This implies that the output of the receiver is ω_j and hence (r_R, r_S) remains consistent after this call to the OT functionality.[2]

We therefore conclude that Eq. (4) holds; the proof of Eq. (5) is almost identical (with the addition that the output remains the same). Combining Equations (2) to (5), we obtain that for every *fixed* x_0, x_1, σ, $t \notin \mathsf{Good}(x_0, x_1, \sigma)$ and for *every* sequence $\{(\tau_i, \omega_i)\}_{i=1}^n$,

$$\begin{aligned}
&\Pr[\text{OUTPUT}_R^{\pi}(x_0, x_1, \sigma; 1^n) = x_\sigma \mid t \wedge \{(\tau_i, \omega_i)\}_{i=1}^n] \\
&= \frac{|\mathsf{S_{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n)| \cdot |\mathsf{R_{out}}(x_\sigma, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{S_{All}}(x_0, x_1, t, \{(\tau_i, \omega_i)\}_{i=1}^n)| \cdot |\mathsf{R_{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|} = \frac{|\mathsf{R_{out}}(x_\sigma, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|}{|\mathsf{R_{All}}(\sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n)|} \\
&= p_\pi(x_\sigma, \sigma, t, \{(\tau_i, \omega_i)\}_{i=1}^n) \leq \frac{1}{2}.
\end{aligned}$$

This completes the proof of Eq. (1), thereby implying Claim 3.5. ■

Combining Claims 3.5 and 3.4, we obtain that the statistical distance of $\mathcal{E}^1$ and $\mathcal{E}^2$ is greater than $1/2 - 2\mu(n)$, completing the proof of Lemma 3.3. ■

We have demonstrated that the existence of an OT extension protocol implies the existence of two ensembles that are computationally indistinguishable and yet statistically far apart, which in turn implies the existence of one-way functions, by [8]. ■

4 Adaptive Security

In this section we consider the feasibility of constructing OT-extension protocols that are secure in the presence of adaptive adversaries. It is easy to see that the OT-extension protocols of Beaver [3]

[2] We stress that this argument does not hold if we considered only the outputs ω_j of the ideal OT calls, and not *both the input τ_j and output ω_j*. This is because the consistency of r_S with t and $\{\omega_i\}_{i=1}^n$ just guarantees that one of the inputs sent by S is ω_j; it does not guarantee that this is the output received by R. For example, consider the case that R inputs a random bit, and the sender inputs $(b, \bar{b})$ for a random b. The sender's tape r_S is consistent with t and any $\omega_j \in \{0,1\}$ since there exists a receiver's tape r_R for which R receives ω_j. However, there also exists a receiver's tape r'_R that is in $\mathsf{R_{All}}$ (because there exists a sender tape providing consistency), but the pair (r_S, r'_R) is *not consistent*. Thus, although seemingly trivial, this argument requires care and only holds since we consider both the inputs and outputs to the ideal OT calls.

and Ishai et al. [12] are not secure when considering adaptive security. This is because the receiver's view is essentially a binding commitment to all of the sender's inputs.[3] This raises the question as to whether there exists an OT extension protocol at all in the presence of adaptive adversaries. Of course, if the existence of an OT extension protocol (that is secure for adaptive adversaries) implies OT that is secure for adaptive adversaries, then this means that only a trivial OT extension that constructs OT from scratch exists. We provide a partial answer to this question and show that a protocol for OT-extension that is secure in the presence of adaptive adversaries implies the existence of an OT protocol that is secure in the presence of *static* adversaries. Thus, any protocol for extending OT that maintains adaptive security needs to assume, at the very least, the existence of a statically secure protocol for OT. We state and prove this for semi-honest adversaries; an analogous theorem for malicious adversaries can be obtained by applying a GMW-type compiler. Formally, we prove the following theorem (the intuition appears immediately after Protocol 4.2 below):

Theorem 4.1 *Let n be the security parameter. If there exists an OT-extension protocol from n to $n+1$ that is secure in the presence of adaptive semi-honest adversaries, then there exists an OT protocol that is secure in the presence of static semi-honest adversaries.*

Proof: We prove the theorem by building an OT protocol that is secure in the presence of static adversaries from any OT extension from n to $4n$ that is secure in the presence of adaptive adversaries. (Note that by Proposition 2.7, an OT extension from n to $4n$ exists if there exists an extension from n to $n+1$.) We first present the construction of the OT protocol for static adversaries and then provide intuition as to why it is secure.

Let $\pi = \langle S, R \rangle$ be a protocol that securely computes the $4n \times OT$ functionality in the OT^n-hybrid model in the presence of adaptive semi-honest adversaries. Without loss of generality, we assume that all of the ideal calls to OT in π are such that S plays the sender and R plays the receiver. This is without loss of generality since the roles in OT can always be reversed [17]. We construct an OT protocol $\hat{\pi}$ in the plain model (i.e., with no calls to an ideal OT functionality), as follows:

Protocol 4.2 (OT protocol $\hat{\pi} = \langle \hat{S}, \hat{R} \rangle$ for Static Adversaries)

- **Inputs:** *The input of the sender $\hat{S}$ is $b_0, b_1 \in \{0,1\}$ and the input of the $\hat{R}$ is $\sigma \in \{0,1\}$.*
- **The protocol:**
  1. *$\hat{S}$ chooses two random strings $\alpha_0, \alpha_1 \in \{0,1\}^{4n}$.*
  2. *$\hat{S}$ and $\hat{R}$ run the extension protocol π as follows:*
     (a) *$\hat{S}$ plays the sender S in π with inputs (α_0, α_1).*
     (b) *$\hat{R}$ plays the receiver R in π with input σ^{4n} (i.e., the string of length $4n$ with all bits set to σ)*
     (c) *The parties follow the instructions of π exactly except that whenever π instructs them to make an ideal call to the OT functionality with input (β_0, β_1) for S and input τ for R, the sender $\hat{S}$ sends the pair (β_0, β_1) to $\hat{R}$, and $\hat{R}$ proceeds to run R with output β_τ from the simulated ideal call.*

[3] In [3] a Yao garbled circuit is used which is binding when instantiated with known encryption methods. Likewise, [12] uses correlation-robust hash functions for which it is hard to find collisions, which is exactly what is needed in order to "explain the transcript" in different ways as is needed for proving adaptive security.

(d) Let $\gamma \in \{0,1\}^{4n}$ denote the output of R in the execution of π.

3. *$\hat{S}$ chooses two random strings $r_0, r_1 \in_R \{0,1\}^{4n}$ and sets:*

$$z_0 = \langle \alpha_0, r_0 \rangle \oplus b_0 \quad \text{and} \quad z_1 = \langle \alpha_1, r_1 \rangle \oplus b_1.$$

$\hat{S}$ sends (r_0, z_0) and (r_1, z_1) to $\hat{R}$.

- **Output:** *$\hat{R}$ outputs $z_\sigma \oplus \langle \gamma, r_\sigma \rangle$.*

It is clear that $\hat{\pi}$ correctly computes the OT functionality. This is because by the correctness of the OT extension protocol, R will output $\gamma = \alpha_\sigma$ in Step 2d, except with negligible probability. Thus, $z_\sigma \oplus \langle \gamma, r_\sigma \rangle = z_\sigma \oplus \langle \alpha_\sigma, r_\sigma \rangle = b_\sigma$, as required.

We proceed to prove that π securely computes the OT functionality in the presence of semi-honest adversaries. We begin with the intuition. If $\hat{S}$ and $\hat{R}$ were to run the original extension protocol π with the ideal calls, then it is clear that $\hat{\pi}$ is a secure OT protocol. This is because $\hat{S}$ learns nothing about σ, and $\hat{R}$ learns α_σ but nothing about $\alpha_{1-\sigma}$. Thus, $\hat{R}$ learns b_σ but nothing about $b_{1-\sigma}$ (observe that $\langle \alpha_{1-\sigma}, r_{1-\sigma} \rangle$ hides $b_{1-\sigma}$ by the fact that $\alpha_{1-\sigma}$ is random). Now, in $\hat{\pi}$ the difference is that $\hat{S}$ sends both inputs to $\hat{R}$ in every ideal OT call within the execution of π. Clearly, $\hat{S}$'s view can be simulated since its view is identical to the case that π with the ideal OT calls is used. In contrast, $\hat{R}$ learns more information since it obtains both sender inputs in all ideal OT calls. Since the inputs to each ideal call are a single bit, we have that $\hat{R}$ obtains n more bits of information than in the original extension protocol using ideal OT calls. However, $\alpha_{1-\sigma}$ is $4n$ bits long and so still must have high entropy even given the n additional bits of information learned. This entropy is enough to hide $b_{1-\sigma}$ since $\langle \alpha_{1-\sigma}, r_{1-\sigma} \rangle$ is a perfect universal hash function, and so a good randomness extractor.

The above seems to have nothing to do with the fact that the extension protocol π is secure in the presence of *adaptive adversaries*. However, the argument that just n more bits of information are obtained is valid only in this case. Specifically, by the definition of security in the presence of adaptive adversaries, the simulator must be able to simulate in the case that the receiver is corrupted at the onset, and the sender is corrupted at the end after the protocol concludes (formally, in the "post-execution corruption phase"). This means that the simulator must first generate a receiver-view (given the receiver's input and output), and must then later generate a sender-view (given the sender's input) that is consistent with the *already fixed* receiver-view that it previously generated. This sender-view contains, amongst other things, the inputs that the sender uses in all of the n ideal calls to the OT functionality within the extension protocol π. Thus, it is possible to add these inputs of the sender to the previously generated receiver-view (we call this the extended receiver view) and the result is the receiver-view in the modified extension protocol used in Step 2 of $\hat{\pi}$; in particular, both sender's inputs to all ideal OT calls appear. Observe that only n bits of additional information are added to the receiver view in order to obtain the extended view, and so there are at most 2^n extended views for any given receiver view. However, there are 2^{4n} different possible strings $\alpha_{1-\sigma}$. The crucial point here is that the above implies that many different possible strings $\alpha_{1-\sigma}$ must be consistent with any given extended view (except with negligible probability). This relies critically on the fact that the receiver-view is fixed before the sender corruption and so the same extended receiver-view must be consistent with many different sender inputs to the ideal OT calls. Now, once we have that many different possible $\alpha_{1-\sigma}$ strings are consistent, we can use the fact that $\alpha_{1-\sigma}$ is randomly chosen to apply the leftover hash lemma and conclude that $\langle \alpha_{1-\sigma}, r_{1-\sigma} \rangle$ is a bit that is statistically close to uniform. We now proceed to the formal proof.

Corrupted sender: The case of a corrupted sender is straightforward since the sender $\hat{S}$ receives no information in Step 2 of $\hat{\pi}$ beyond what it receives in a real execution of π with ideal OT calls. Thus the simulator that is assumed to exist for the sender S in π can be used to generate the exact view of $\hat{S}$ in Step 2 of $\hat{\pi}$. Since $\hat{S}$ receives no messages beyond in Step 2, there is nothing more to be added to the view of $\hat{S}$.

Corrupted receiver: In order to construct our simulator $\mathcal{S}_{\hat{R}}$ for the corrupted receiver $\hat{R}$ in $\hat{\pi}$, we first define a specific simulator $\mathcal{SIM}$ for the extension protocol π for the adaptive setting. Let $\mathcal{A}$ and $\mathcal{Z}$ be the following real-life semi-honest adversary and environment for π; see Section 2 for a brief overview of the definition of adaptive security and [4] for full definitions. At the beginning of the execution of π, the adversary $\mathcal{A}$ corrupts the receiver and learns its input $\sigma \in \{0,1\}^{4n}$. It then follows the honest strategy for R and at the end of the execution, outputs its entire view. In the post-execution phase, $\mathcal{Z}$ generates a "corrupt S" message, sends it to $\mathcal{A}$ who corrupts S and hands $\mathcal{Z}$ the internal view of S. $\mathcal{Z}$ then outputs its internal view (note that it contains views of both R and S). Let $\mathcal{SIM}$ be the ideal-process adversary that is guaranteed to exist for this $\mathcal{A}$ and $\mathcal{Z}$ by the security of π. We remark that $\mathcal{SIM}$ generates a view of an execution of π in the OT-hybrid model, where ideal calls are used for the n invocations of OT. We use $\mathcal{SIM}$ to construct the simulator $\mathcal{S}_{\hat{R}}$ for the case of a corrupted receiver in $\hat{\pi}$.

Construction 4.3 ($\mathcal{S}_{\hat{R}}$) *$\mathcal{S}_{\hat{R}}$ receives σ and b_σ as input and works in three stages as follows:*

1. Stage 1 – obtain simulated receiver-view in π:
   (a) *Choose a random string $\alpha_\sigma \in_R \{0,1\}^{4n}$ as the "output of π" and a random tape $r_{\mathcal{SIM}}$ for $\mathcal{SIM}$ of the appropriate length.*
   (b) *Start an execution of $\mathcal{SIM}$ with random-tape $r_{\mathcal{SIM}}$. When $\mathcal{SIM}$ corrupts the receiver, hand σ^{4n} to $\mathcal{SIM}$ as the input of R.*
   (c) *In the computation stage, play the role of the trusted party and send α_σ to $\mathcal{SIM}$ as the output of R from $4n \times OT$. (Since we are in the semi-honest setting, R always sends its specified input σ^{4n} and so the output that it would receive is always α_σ.)*
   (d) *Let v_R be the output of $\mathcal{SIM}$ at the end of the execution phase (this consists of a view for the receiver). If v_R is not consistent with σ^{4n} and α_σ,[4] return $\perp$ and abort. Otherwise, proceed to the next stage.*
2. Stage 2 – obtain extended receiver-view:
   (a) *Choose a random string $\alpha_{1-\sigma} \in \{0,1\}^{4n}$.*
   (b) *Send a "corrupt S" message to $\mathcal{SIM}$ on behalf of $\mathcal{Z}$. When $\mathcal{SIM}$ corrupts the sender, hand (α_0, α_1) to $\mathcal{SIM}$ as the input of S.*
   (c) *Let v_S be the view of the sender sent by $\mathcal{SIM}$ to $\mathcal{Z}$. If v_S is not consistent with v_R and the inputs, output $\perp$ and abort. If v_S is consistent with v_R and the inputs, then for each of the n calls for the ideal OT functionality, extend v_R by appending the other input used by the sender (as appear in v_S) into the view v_R (note that v_R already contains one of the inputs used by the sender in each call since the receiver receives one output in each ideal call). Let v'_R be the extended view.*

[4] We say that a view is consistent with inputs and outputs if when running the party on the given view and input, it outputs the correct output.

3. Stage 3 – complete simulation:

(a) Choose two random strings $r_0, r_1 \in \{0,1\}^{4n}$; let $z_\sigma = \langle \alpha_\sigma, r_\sigma \rangle \oplus b_\sigma$ (where b_σ is from the input of $\mathcal{S}_{\hat{R}}$) and let $z_{1-\sigma}$ be a random bit.

(b) Output $v'_{\hat{R}}, r_0, r_1, z_0, z_1$.

We now prove that $\mathcal{S}_{\hat{R}}$ is a good simulator. That is, we prove that:

$$\left\{\mathcal{S}_{\hat{R}}(1^n, \sigma, b_\sigma)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{\text{c}}{\equiv} \left\{\text{VIEW}^{\hat{\pi}}_{\hat{R}}(1^n, b_0, b_1, \sigma)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \tag{6}$$

To prove Eq. (6), we consider a hybrid simulator $\mathcal{S}^h$ that receives as input $b_{1-\sigma}$ in addition to the input (σ, b_σ) of $\mathcal{S}_{\hat{R}}$. It then works exactly as $\mathcal{S}_{\hat{R}}$ except that in Stage 3 of the simulation it sets $z_{1-\sigma} = \langle \alpha_{1-\sigma}, r_{1-\sigma} \rangle \oplus b_{1-\sigma}$ (instead of setting $z_{1-\sigma}$ to a random bit as $\mathcal{S}_{\hat{R}}$ does).

We first prove that the output of the hybrid simulator is indistinguishable from the receiver view in a real execution. That is, we prove that:

$$\left\{\mathcal{S}^h(1^n, \sigma, b_0, b_1)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{\text{c}}{\equiv} \left\{\text{VIEW}^{\hat{\pi}}_{\hat{R}}(1^n, b_0, b_1, \sigma)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \tag{7}$$

Note that the only difference between the two distributions is that in $\text{VIEW}^{\hat{\pi}}_{\hat{R}}(1^n, b_0, b_1, \sigma)$, the "extended view of R" (including both inputs used by the sender in each ideal OT call) is generated in a real execution of π, whereas in $\mathcal{S}^h(1^n, \sigma, b_0, b_1)$ the extended view is generated by $\mathcal{SIM}$ after the corruption at the end. So intuitively the guarantee that $\mathcal{SIM}$ is a good simulator implies that the two ensembles are computationally indistinguishable. Formally, we define a machine $\mathcal{D}$ that receives the output of $\mathcal{Z}$ after an execution of π in the adaptive setting, and attempts to determine whether it obtained a pair of receiver/sender views from a real or ideal execution. $\mathcal{D}$ generates an extended receiver-view from the pair of receiver/sender views that it received, and in addition computes the messages $(r_0, z_0), (r_1, z_1)$ using the correct sender inputs b_0, b_1 (that it's given as auxiliary input) and using the strings α_0, α_1 that appear in $\mathcal{Z}$'s output. Finally, $\mathcal{D}$ outputs the extended receiver-view together with the last message; this constitutes a view of the receiver $\hat{R}$ in $\hat{\pi}$. It is immediate that if $\mathcal{D}$ received a pair of views from a real execution of π then it outputs a view which is *identical* to $\text{VIEW}^{\hat{\pi}}_{\hat{R}}(1^n, b_0, b_1, \sigma)$. In contrast, if $\mathcal{D}$ received a pair of views generated by $\mathcal{SIM}$ in an ideal execution, then it outputs a view which is *identical* to $\mathcal{S}^h(1^n, \sigma, b_0, b_1)$. Thus, Eq. (7) follows from the security of π with simulator $\mathcal{SIM}$.

We now proceed to prove that the output of $\mathcal{S}_{\hat{R}}$ is statistically close to the output of the hybrid simulator S^h. That is:

$$\left\{\mathcal{S}_{\hat{R}}(1^n, \sigma, b_\sigma)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{\text{s}}{\equiv} \left\{\mathcal{S}^h(1^n, \sigma, b_0, b_1)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \tag{8}$$

First note that $\mathcal{S}_{\hat{R}}$ and $\mathcal{S}^h$ work identically in the first two stages of the simulation and differ only in how $z_{1-\sigma}$ is computed. In particular, the distributions over the extended views generated by $\mathcal{S}_{\hat{R}}$ and by $\mathcal{S}^h$ are identical; let $V'_{\hat{R}}(1^n, \sigma)$ denote this distribution.

The first step is to show that with probability negligibly close to 1, there are exponentially many strings $\alpha_{1-\sigma}$ that are consistent with an extended view generated by $\mathcal{SIM}$ (as run by $\mathcal{S}^h$ or equivalently $\mathcal{S}_{\hat{R}}$). Fix $\sigma \in \{0,1\}$ and b_σ (the following holds for all σ, b_σ and we fix them here for clarity). For a given random tape $r_{\mathcal{SIM}}$ of $\mathcal{SIM}$ and a given α_σ, let v_R be the (regular, non-extended) view generated by $\mathcal{SIM}$ with random tape $r_{\mathcal{SIM}}$ and α_σ in the execution phase. Let

$\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)$ be the set of all strings $\alpha_{1-\sigma}$ of size $4n$ for which the views v_R, v_S generated by $\mathcal{SIM}$ with random tape $r_{\mathcal{SIM}}$ and inputs α_σ and $\alpha_{1-\sigma}$ in the computation and post-execution phases, respectively, are all *consistent* (we have already fixed σ and b_σ so consistency is also with respect to these values; see Footnote 4). Note that if $\mathcal{S}^h$ or $\mathcal{S}_{\hat{R}}$ would output $\perp$ in the first stage (i.e., if v_R is not consistent with the input and output) when choosing $r_{\mathcal{SIM}}, \alpha_\sigma$ then $\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)$ is *empty*.

We now prove that for every $\sigma, b_\sigma \in \{0,1\}$, there exists a negligible function μ such that

$$\Pr_{r_{\mathcal{SIM}}, \alpha_\sigma}\left[|\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)| \geq 2^{3n}\right] \geq 1 - \mu(n).$$

Intuitively, this holds because if $\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)$ is "small", then this means that $\mathcal{SIM}$ would fail with high probability. Formally, assume that $\Pr_{r_{\mathcal{SIM}}, \alpha_\sigma}[|\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)| \geq 2^{3n}]$ is non-negligibly smaller than 1. We consider two cases:

1. With non-negligible probability, the view v_R generated by $\mathcal{SIM}$ with random tape $r_{\mathcal{SIM}}$ and α_σ cause $\mathcal{S}^h$ and $\mathcal{S}_{\hat{R}}$ to output $\perp$ (i.e., it is not consistent with the inputs/outputs): In this case, a distinguisher $\mathcal{Z}$ easily distinguishes the output of $\mathcal{SIM}$ from the views of v_R, v_S in a real execution of π since in a real execution the views are consistent except with negligible probability.

2. With non-negligible probability, the view v_R *is* consistent but $|\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)| < 2^{3n}$: In this case, it is possible to distinguish a real execution of π from an ideal execution with $\mathcal{SIM}$ because the probability that a random $\alpha_{1-\sigma}$ is in $\Delta(r_{\mathcal{SIM}}, \alpha_\sigma)$ is less than $\frac{2^{3n}}{2^{4n}} = 2^{-n}$. Thus, the environment $\mathcal{Z}$ can just supply a random $\alpha_{1-\sigma}$ and see if in the post-execution corruption it receives a consistent view. In the real execution it will always receive a consistent view. However, in the ideal (simulated) execution, it will receive a consistent view with probability less than 2^{-n}. This is due to the fact that when $\alpha_{1-\sigma} \notin \Delta(r_{\mathcal{SIM}}, \alpha_\sigma)$ the view is *not* consistent. Thus, $\mathcal{Z}$ distinguishes with probability $(1 - 2^{-n})$ times the probability that this case occurs, which is non-negligible.

We stress that the calculation in the second case holds since the view of the receiver v_R is fixed before the post-execution phase and thus is fixed before $\alpha_{1-\sigma}$ is essentially chosen.

We now fix $r^*_{\mathcal{SIM}}$ and α^*_σ for which $|\Delta(r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| \geq 2^{3n}$ and prove that the outputs of $\mathcal{S}^h$ and $\mathcal{S}_{\hat{R}}$ are statistically close for such $r^*_{\mathcal{SIM}}$ and α^*_σ. First, recall that an extended view v'_R is obtained by concatenating the other (previously not received) input of the sender in the n calls to the ideal OT to the view v_R. Since there are 2^n possible "other sender inputs" in the n ideal OT calls, it follows that for any given receiver-view v_R (which is fully determined by $r^*_{\mathcal{SIM}}$ and α^*_σ; recall that σ, b_σ are already fixed) there are at most 2^n possible associated extended views. (Again, this relies on the fact that the receiver-view is fixed before the post-execution corruption phase.)

Now, since there are 2^n possible extended views, we can partition the at least 2^{3n} consistent strings $\alpha_{1-\sigma} \in \Delta(r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$ so that each partition contains the set of strings $\alpha_{1-\sigma}$ that yield the extended view v'_R. Equivalently, we associate $\alpha_{1-\sigma}$ with v'_R if $\mathcal{SIM}$ with $r^*_{\mathcal{SIM}}$ and α^*_σ outputs the extended view v'_R when given $\alpha_{1-\sigma}$ in the post-execution corruption phase. We denote by $\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$ the set of all strings $\alpha_{1-\sigma} \in \Delta(r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$ which are associated with v'_R, as described above.

We argue that the probability of obtaining an extended view v'_R for which $|\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| < 2^n$ is at most 2^{-n} (i.e., an extended view for which the set of associated strings $\alpha_{1-\sigma}$ is small is obtained

with probability at most 2^{-n}). We stress that the probability is over the choice of $\alpha_{1-\sigma}$ (all other randomness is fixed).

In order to see this, observe that the fact that $|\Delta(r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| \geq 2^{3n}$ implies that there are at least 2^{3n} strings $\alpha_{1-\sigma}$ that are associated with *some* extended view v'_R. Now, for every v'_R for which $|\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| < 2^n$, we have that v'_R is generated by less than 2^n of those 2^{3n} strings. Thus, such a v'_R is obtained with probability less than $2^n/2^{3n} = 2^{-2n}$. By union bound over the 2^n possible extended views v'_R (which also bounds the number of extended views for which $|\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| < 2^n$) we conclude that

$$\Pr\Big[\,\big|\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)\big| < 2^n\Big] < 2^n \cdot \frac{1}{2^{2n}} = \frac{1}{2^n} \tag{9}$$

where the probability is over the choice of $\alpha_{1-\sigma}$.

From Eq. (9), we know that when choosing $\alpha_{1-\sigma}$ at random, the probability that we will obtain an extended view v'_R such that $\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$ is small (with less than 2^n strings $\alpha_{1-\sigma}$ associated with it) is less than 2^{-n}. We therefore proceed by conditioning further over views v'_R for which $|\Gamma(v'_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| \geq 2^n$. Specifically, we argue that the distributions generated by $\mathcal{S}_{\hat{R}}$ and $\mathcal{S}^h$ are statistically close, conditioned on $r^*_{\mathcal{SIM}}, \alpha^*_\sigma$ such that $|\Delta(r^*_{\mathcal{SIM}}, \alpha^*_\sigma)| \geq 2^{3n}$ and conditioned on the extended view being a specific v'^*_R for which $\big|\Gamma(v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)\big| \geq 2^n$.

First, observe that since $\alpha_{1-\sigma}$ is chosen uniformly and independently of $r^*_{\mathcal{SIM}}, \alpha_\sigma$, it is uniformly distributed in $\Gamma(v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$, when conditioning on all of the above. (The conditioning over v'^*_R is equivalent to saying that $\alpha_{1-\sigma}$ is uniform in $\Gamma(v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$ instead of being uniform in $\{0,1\}^{4n}$.) Second, recall that $\Gamma(v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma)$ is a set of size at least 2^n. Third, note that $H_{r_{1-\sigma}}(x) = \langle r_{1-\sigma}, x\rangle)$ is a universal hash function from $\{0,1\}^{4n}$ to $\{0,1\}$. Thus, by the Leftover Hash Lemma (the version given in [13]), it holds that:

$$SD\Big((r_{1-\sigma}, \langle r_{1-\sigma}, \alpha_{1-\sigma}\rangle), (r_{1-\sigma}, U_1)\Big) \leq \frac{1}{2^{(n-1)/2}}$$

where SD denotes statistical distance and U_1 denotes the uniform distribution over $\{0,1\}$ (as above, this statistical distance is computed when conditioned over $v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma$). Thus, these random variables are statistically close, conditioned on $v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma$ as above. Noting that in the output of $\mathcal{S}_{\hat{R}}$ we have $(r_{1-\sigma}, z_{1-\sigma}) = (r_{1-\sigma}, U_1)$, and in the output of $\mathcal{S}^h$ we have that $(r_{1-\sigma}, z_{1-\sigma}) = (r_{1-\sigma}, \langle r_{1-\sigma}, \alpha_{1-\sigma}\rangle)$, we conclude that

$$\Big\{\mathcal{S}_{\hat{R}}(1^n, \sigma, b_\sigma) \mid v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma\Big\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{\text{s}}{\equiv} \Big\{\mathcal{S}^h(1^n, \sigma, b_0, b_1) \mid v'^*_R, r^*_{\mathcal{SIM}}, \alpha^*_\sigma\Big\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}}$$

where the conditioning is as described above. We reiterate that this holds since the extended views and the pair (r_σ, z_σ) are generated in an identical way by $\mathcal{S}_{\hat{R}}$ and $\mathcal{S}^h$, and the only difference is with respect to $(r_{1-\sigma}, z_{1-\sigma})$. Eq. (8) follows from the fact that we condition here on events that occur with all but negligible probability (and the events have identical probability with $\mathcal{S}_{\hat{R}}$ and $\mathcal{S}^h$; see Lemma 2.8). Combining Eq. (7) with Eq. (8), we derive Eq. (6), thereby completing the proof of Theorem 4.1. ■

Corollary – lengthening string OT. Observe that in our proof above the receiver always uses σ^{4n} for input. Thus, it follows that the theorem holds even if the receiver is interested in only obtaining the string of all of the "0 inputs" or the string of all of the "1 inputs". Stated differently, our proof holds also for the problem of lengthening string OT; i.e., for the problem of obtaining a *single string OT* for strings of length $n+1$ or more, given a *single string OT* for strings of length n.

5 OT Extensions Require Super-Logarithmic Calls

Theorem 5.1 *Let $f : \mathbb{N} \to \mathbb{N}$ be a function such that $f(n) \in \mathcal{O}(\log n)$, and let n be the security parameter. Then, if there exists a protocol π that is an OT-extension from $f(n)$ to $f(n)+1$ that is secure in the presence of malicious adversaries, then there exists a protocol for the OT functionality that is secure in the presence of malicious adversaries.*

Proof: Intuitively, in an OT extension protocol using only $\mathcal{O}(\log n)$ ideal OT calls, it is possible for the receiver to guess the bits that it would receive as output from these calls instead of actually running them. Since there are only $\mathcal{O}(\log n)$ calls, the probability that the receiver guesses correctly is $2^{-\mathcal{O}(\log n)} = 1/\mathsf{poly}(n)$. This idea can be used to construct an OT protocol that is weak in the sense that full privacy is maintained, but correctness only holds with probability $1/2 + 1/\mathsf{poly}(n)$. We stress that a naive attempt to implement the above idea will not work since it is necessary to ensure that if the receiver's guesses are incorrect then it still outputs the correct output of the protocol with probability almost 1/2. Otherwise, the "advantage" in obtaining the correct output when the receiver guesses correctly can be canceled out by the "disadvantage" when the receiver guesses incorrectly. We therefore use a similar technique as in the proof regarding adaptive adversaries above. Specifically, we use the fact that an extension from $f(n)$ to $f(n)+1$ implies an extension from $f(n)$ to n, and then use this to obliviously transfer n random bits. The actual oblivious transfer is carried out by applying a universal hash function to the random strings and using the result to mask the actual bits being transferred. This ensures that we obtain correctness that is noticeable greater than 1/2 and so can be amplified. However, in addition, we also have to claim that privacy is maintained. This is not immediate since the receiver does not follow the specified protocol (rather, it chooses the outputs from the ideal OT calls at random, and this may effect the other messages that it sends). By requiring that the extension protocol be secure for malicious adversaries, this ensures that the receiver cannot learn more by behaving in this way. In addition, we show that a malicious sender can also achieve the same affect by inputting a random bit (for both sender inputs) in each ideal OT call. This implies that a malicious sender can also not learn anything by the receiver behaving in this way. We now proceed to the formal proof.

Throughout the proof, we will construct protocols that are secure for *semi-honest adversaries* only. This suffices since semi-honest OT implies malicious OT [7, 11]. Let $f : \mathbb{N} \to \mathbb{N}$ be a function such that $f(n) \in \mathcal{O}(\log n)$ and let $\pi = \langle S, R \rangle$ be a protocol such that on security parameter n and inputs $x_0, x_1 \in \{0,1\}^{f(n)+1}$ and $\sigma \in \{0,1\}^{f(n)+1}$ securely computes the $(f(n)+1) \times OT$ functionality in the $OT^{f(n)}$-hybrid model (that is, making at most $f(n)$ calls to an ideal OT). We assume that π is secure in the presence of malicious adversaries. We assume that in all of these calls, R is the one to receive output (this is without loss of generality since oblivious transfer is symmetric [17] and so the roles can be reversed by adding additional messages in π). We show how to construct a protocol for computing the OT functionality without any further assumptions other than the existence of an extension protocol π with the parameters in the theorem statement. This is achieved in two steps. First, we use the OT-extension from $f(n) = \mathcal{O}(\log n)$ to n to construct a protocol $\tilde{\pi}$ which is simulatable and therefore fully secure, but whose error might be large. Then we amplify the correctness of the protocol using multiple execution. As we show, this can be done once the basic protocol is fully secure.

Step 1 – constructing a weak-OT. We begin by formally defining weak-OT, which is an oblivious transfer for semi-honest adversaries that has weak correctness but full simulation security.[5] We then show how to construct a weak-OT protocol $\tilde{\pi} = \langle \tilde{S}, \tilde{R} \rangle$ from an OT-extension from $f(n)$ to n. Note that by Proposition 2.7, if there exists an extension protocol from $f(n)$ to $f(n)+1$, then there exists an extension protocol from $f(n)$ to n.

Definition 5.2 (Weak-OT) *A two-party protocol $\pi = \langle S, R \rangle$ is a* weak-OT *if the following hold:*

- **Weak-correctness:** *There exists a polynomial $p(\cdot)$ such that for all $b_0, b_1, \sigma \in \{0,1\}$ and all sufficiently large n's, it holds that* $\Pr[\text{OUTPUT}_R^{\pi}(1^n, b_0, b_1, \sigma) = b_\sigma] \geq \frac{1}{2} + \frac{1}{p(n)}$.
- **Privacy:** *There exists PPT machines $\mathcal{S}_R$ and $\mathcal{S}_S$ such that*

$$\{\mathcal{S}_R(1^n, \sigma, b_\sigma)\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{c}{\equiv} \{\text{VIEW}_R^{\pi}(1^n, b_0, b_1, \sigma)\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}}$$
$$\{\mathcal{S}_S(1^n, b_0, b_1)\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{c}{\equiv} \{\text{VIEW}_S^{\pi}(1^n, b_0, b_1, \sigma)\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}}$$

Let $\alpha_0, \alpha_1, c \in \{0,1\}^n$ be n-bit strings. Let $\alpha_0 = \alpha_0^1, \ldots, \alpha_0^n$, $\alpha_1 = \alpha_1^1, \ldots, \alpha_1^n$, and $c = c_1, \ldots, c_n$. Recall that $\alpha_c = \alpha_{c_1}^1, \alpha_{c_2}^2, \ldots, \alpha_{c_n}^n$; that is, the ith bit of α_c is either α_0^i or α_1^i, depending on the value of c_i.

Let $\pi = \langle S, R \rangle$ be an OT-extension protocol from $f(n) = \mathcal{O}(\log n)$ to n. We construct a weak OT protocol $\tilde{\pi} = \langle \tilde{S}, \tilde{R} \rangle$ as follows:

Protocol 5.3 (A weak-OT with no ideal OT calls)

- **Inputs:** *The sender $\tilde{S}$ has two bits $b_0, b_1 \in \{0,1\}$ and the receiver $\tilde{R}$ has $\sigma \in \{0,1\}$.*
- **The protocol:**
  1. *$\tilde{S}$ chooses two random strings $\alpha_0, \alpha_1 \in_R \{0,1\}^n$.*
  2. *$\tilde{R}$ chooses a random string $c \in_R \{0,1\}^n$.*
  3. *$\tilde{S}$ and $\tilde{R}$ simulate an execution of the extension protocol π, as follows:*
     (a) *$\tilde{S}$ plays the role of the sender S with input $\alpha_0, \alpha_1 \in \{0,1\}^n$ and $\tilde{R}$ plays the role of the receiver R with input $c \in \{0,1\}^n$.*
     (b) *Whenever π instructs the parties to make an OT call, the parties make no call and $\tilde{R}$ chooses a random bit as its output from the call. We denote by $\beta_1, \ldots, \beta_{f(n)}$ the random bits chosen by $\tilde{R}$ as the OT outputs.*
     (c) *Let $\gamma \in \{0,1\}^n$ denote the receiver-output of the simulation of π received by $\tilde{R}$.*
  4. *$\tilde{R}$ chooses a random $c' \in_R \{0,1\}^n$ and sends (c_0, c_1) to $\tilde{S}$, where $c_\sigma = c$ and $c_{1-\sigma} = c'$.*
  5. *$\tilde{S}$ chooses two random strings $r_0, r_1 \in_R \{0,1\}^n$, computes $z_0 = \langle r_0, \alpha_{c_0} \rangle \oplus b_0$ and $z_1 = \langle r_1, \alpha_{c_1} \rangle \oplus b_1$, and sends $(r_0, z_0), (r_1, z_1)$ to $\tilde{R}$.*
- **Output:** *$\tilde{S}$ outputs nothing and $\tilde{R}$ outputs* $\mathsf{out} = z_\sigma \oplus \langle r_\sigma, \gamma \rangle$.

[5] Note that we cannot cast this as a special case of Definition 2.3 since full correctness is required there by stating that π computes f.

We now prove that Protocol 5.3, also denoted $\tilde{\pi}$, is a weak-OT protocol. We begin by showing the weak-correctness of $\tilde{\pi}$; that is we show that the receiver $\tilde{R}$ outputs the correct bit b_σ with probability at least $\frac{1}{2}+\frac{1}{2^{f(n)+2}}$. This suffices since $f(n)=\mathcal{O}(\log n)$ and thus $\frac{1}{2}+\frac{1}{2^{f(n)+2}}=\frac{1}{2}+\frac{1}{2^{c\cdot\log n+2}}=\frac{1}{2}+\frac{1}{4n^c}$ for some constant c. Intuitively, weak correctness holds because $\tilde{R}$ correctly guesses the outputs of the OT calls with probability $1/2^{f(n)}$ in which case $\gamma=\alpha_c$ (except with negligible probability) by the correctness of π and thus $\langle r_\sigma,\gamma\rangle=\langle r_\sigma,\alpha_c\rangle$ and $\mathsf{out}=b_\sigma$. In addition, when the guesses made by $\tilde{R}$ are not correct, it still outputs b_σ with probability $1/2$.

Let b_0,b_1 and σ be the inputs of $\tilde{S}$ and $\tilde{R}$. Note that $\mathsf{out}=z_\sigma\oplus\langle\gamma,r_\sigma\rangle=\langle\alpha_{c_\sigma},r_\sigma\rangle\oplus b_\sigma\oplus\langle\gamma,r_\sigma\rangle=\langle\alpha_c,r_\sigma\rangle\oplus b_\sigma\oplus\langle\gamma,r_\sigma\rangle$ and thus $\mathsf{out}=b_\sigma$ if and only if $\langle\gamma,r_\sigma\rangle=\langle\alpha_c,r_\sigma\rangle$, where r_σ is a random string. Thus,

$$\begin{aligned}\Pr[\mathsf{out}=b_\sigma] &= \Pr[\langle\alpha_c,r_\sigma\rangle=\langle\gamma,r_\sigma\rangle]\\ &= \Pr[\langle\alpha_c,r_\sigma\rangle=\langle\gamma,r_\sigma\rangle\mid\gamma=\alpha_c]\cdot\Pr[\gamma=\alpha_c]\\ &\quad+\Pr[\langle\alpha_c,r_\sigma\rangle=\langle\gamma,r_\sigma\rangle\mid\gamma\neq\alpha_c]\cdot\Pr[\gamma\neq\alpha_c]\\ &= 1\cdot\Pr[\gamma=\alpha_c]+\Pr[\langle\alpha_c,r_\sigma\rangle=\langle\gamma,r_\sigma\rangle\mid\gamma\neq\alpha_c]\cdot(1-\Pr[\gamma=\alpha_c])\end{aligned}$$

Now, let $\mathsf{Correct}$ denote the event that the guesses made by $\tilde{R}$ for the outputs of the $f(n)$ ideal-OT's are the correct outputs. Then, by the correctness of protocol π, $\Pr[\gamma=\alpha_c\mid\mathsf{Correct}]\geq 1-\mathsf{negl}(n)$. This is because when all the outputs from the ideal calls are correct, the execution is exactly the same as in a real execution of π. We therefore have:

$$\Pr[\gamma=\alpha_c]\geq\Pr[\gamma=\alpha_c\mid\mathsf{Correct}]\cdot\Pr[\mathsf{Correct}]\geq(1-\mathsf{negl}(n))\cdot\Pr[\mathsf{Correct}].$$

Noting that π makes $f(n)$ calls to the ideal OT and thus $\Pr[\mathsf{Correct}]=\frac{1}{2^{f(n)}}$, we have that

$$\Pr[\gamma=\alpha_c]\geq\frac{1}{2^{f(n)}}\cdot(1-\mathsf{negl}(n))\geq\frac{1}{2^{f(n)}}-\mathsf{negl}(n). \tag{10}$$

In addition, since the inner-product function $H_{r_\sigma}(x)=\langle x,r_\sigma\rangle$ is a universal hash function (for randomly chosen r_σ) it holds that $\Pr[\langle\alpha_c,r_\sigma\rangle=\langle\gamma,r_\sigma\rangle\mid\gamma\neq\alpha_c]=\frac{1}{2}$. Combining the above, we conclude that:

$$\begin{aligned}\Pr[\mathsf{out}=b_\sigma] &= \Pr[\gamma=\alpha_c]+\Pr[\langle\alpha_c,r_\sigma\rangle=\langle\gamma,r_\sigma\rangle\mid\gamma\neq\alpha_c]\cdot(1-\Pr[\gamma=\alpha_c])\\ &= \Pr[\gamma=\alpha_c]+\frac{1}{2}\cdot(1-\Pr[\gamma=\alpha_c])\\ &= \frac{1}{2}+\frac{1}{2}\cdot\Pr[\gamma=\alpha_c]\\ &= \frac{1}{2}+\frac{1}{2}\cdot\frac{1}{2^{f(n)}}-\mathsf{negl}'(n)\geq\frac{1}{2}+\frac{1}{2^{f(n)+2}}\end{aligned}$$

for all large enough n's (the last inequality may not hold for small values of n).

We proceed to prove *privacy*, by constructing $\mathcal{S}_{\tilde{S}}$ and $\mathcal{S}_{\tilde{R}}$ as required. We start by constructing the simulator $\mathcal{S}_{\tilde{S}}$ for the case that the sender is corrupted. To prove this we use the fact that the original protocol π is secure in the presence of malicious adversaries. Consider a malicious adversary $\mathcal{A}$ for π that controls the sender and learns its input $\alpha_0,\alpha_1\in\{0,1\}^n$. $\mathcal{A}$ follows the honest strategy for S except that it chooses random bits $\beta_1,\ldots,\beta_n$ and then in the jth call to the ideal OT functionality, it uses β_j as both sender inputs to the OT call (ensuring that R receives

β_j). We stress that in the rest of the execution, it behaves as if it has used the correct inputs that were supposed to be sent to the OT calls. Observe that the view of $\mathcal{A}$ in an execution of π is *identically distributed* to the view of $\tilde{S}$ in the simulation of π run in Step 3 of Protocol 5.3. Let $\mathcal{SIM}$ be the simulator that is guaranteed to exist for $\mathcal{A}$ by the security of π. We construct the simulator $\mathcal{S}_{\tilde{S}}$ using $\mathcal{SIM}$:

Construction 5.4 ($\mathcal{S}_{\tilde{S}}$) *: Upon input $b_0, b_1 \in \{0,1\}$, simulator $\mathcal{S}_{\tilde{S}}$ works as follows:*

1. *$\mathcal{S}_{\tilde{S}}$ chooses two random strings $\alpha_0, \alpha_1 \in_R \{0,1\}^n$ and runs $\mathcal{SIM}$ with sender-inputs α_0, α_1. Let v_S be the sender-view output by $\mathcal{SIM}$ at the end of its execution ($\mathcal{SIM}$ also sends input to the trusted party, but this is ignored by $\mathcal{S}_{\tilde{S}}$).*

2. *$\mathcal{S}_{\tilde{S}}$ chooses two random strings $c_0, c_1 \in_R \{0,1\}^n$ as the message received from $\tilde{R}$ in Step 4 of Protocol 5.3, and outputs $v_{\tilde{S}} = (v_S, c_0, c_1)$.*

The fact that $\mathcal{S}_{\tilde{S}}$ is a good simulator follows immediately from the fact that $\mathcal{SIM}$ generates a sender-view that is indistinguishable from what $\mathcal{A}$ would see in a real execution of π. Since we have already observed that the view of $\tilde{S}$ in Step 3 of Protocol 5.3 is identical to the view of $\mathcal{A}$ above in π, it follows that v_S is indistinguishable from $\tilde{S}$'s view in Step 3 of Protocol 5.3. Next observe that a distinguisher $\mathcal{D}$ for $\mathcal{SIM}$ and π obtains the input/output used (α_0, α_1, c) and thus can extend the view of the sender to include c_0, c_1 where $c_\sigma = c$, and c is the input of R into the execution of π with $\mathcal{A}$ (we can assume that D knows σ as auxiliary input). Thus, the view of $\tilde{S}$ in Protocol 5.3 (resp., as generated by simulator $\mathcal{S}_{\tilde{S}}$) can be perfectly constructed by $\mathcal{D}$ from the real view v_S of S in π (resp., from a simulated view v_S of S as generated by $\mathcal{SIM}$). This implies that if the output of $\mathcal{S}_{\tilde{S}}$ can be distinguished from the view of $\tilde{S}$ in a real execution of Protocol 5.3, then the output of $\mathcal{SIM}$ can be distinguished from the view of $\mathcal{A}$ in a real execution of π, in contradiction to the security of π with simulator $\mathcal{SIM}$. The formal reduction is straightforward.

We now proceed to construct a simulator $\mathcal{S}_{\tilde{R}}$ for the case that the receiver is corrupted. As above, we consider a malicious adversary $\mathcal{A}$ for π as follows. $\mathcal{A}$ receives the receiver's input $c \in \{0,1\}^n$ and follows the honest receiver strategy except that in each of the calls to the ideal OT functionality, it chooses a random bit β_j and proceeds with β_j as the output of the ideal OT. Let $\mathcal{SIM}$ be the simulator that is guaranteed to exist for $\mathcal{A}$ by the security of π. We use it construct the simulator $\mathcal{S}_{\tilde{R}}$ (recall that $\mathcal{SIM}$ works in the setting for malicious adversaries and thus interacts with a trusted party and sends a receiver-input which is not necessarily the prescribed receiver-input):

Construction 5.5 ($\mathcal{S}_{\tilde{R}}$) *: Upon input $\sigma, b_\sigma \in \{0,1\}$, simulator $\mathcal{S}_{\tilde{R}}$ works as follows:*

1. *$\mathcal{S}_{\tilde{R}}$ chooses three random strings $\alpha_0, \alpha_1, c \in_R \{0,1\}^n$.*

2. *$\mathcal{S}_{\tilde{R}}$ runs $\mathcal{SIM}$ with receiver input c.*

3. *When $\mathcal{SIM}$ sends some $c^* \in \{0,1\}^n$ to the trusted party, $\mathcal{S}_{\tilde{R}}$ hands α_{c^*} as the receiver-output to $\mathcal{SIM}$ from the trusted party. Let v_R be the output of $\mathcal{SIM}$.*

4. *$\mathcal{S}_{\tilde{R}}$ chooses random strings $c', r_0, r_1 \in_R \{0,1\}^n$, and sets $c_\sigma = c$ and $c_{1-\sigma} = c'$. Then, $\mathcal{S}_{\tilde{R}}$ computes $z_\sigma = \langle r_\sigma, \alpha_{c_\sigma} \rangle \oplus b_\sigma$ and sets $z_{1-\sigma} \in_R \{0,1\}$ to be a random bit.*

5. *$\mathcal{S}_{\tilde{R}}$ outputs a receiver view $(c_0, c_1, v_R, r_0, z_0, r_1, z_1)$. (Note that c_0, c_1 are actually part of $\tilde{R}$'s random tape, since they are chosen by $\tilde{R}$.)*

In order to show that $\mathcal{S}_{\tilde{R}}$ is a "good simulator", we construct a hybrid simulator $\mathcal{S}^h$ and show that its output is indistinguishable both from the output of the real simulator and the view of the receiver in the real execution of the protocol.

$\mathcal{S}^h$ receives as input σ and b_0, b_1 (in contrast to $\mathcal{S}_{\tilde{R}}$ which receives only σ and b_σ) and works exactly as $\mathcal{S}_{\tilde{R}}$ except that it lets $z_{1-\sigma} = \langle r_{1-\sigma}, \alpha_{c_{1-\sigma}} \rangle \oplus b_{1-\sigma}$ (rather than a random bit).

We begin by proving that the output of S^h is indistinguishable from the output of the receiver $\tilde{R}$'s view in a real execution of Protocol 5.3. That is:

$$\left\{\mathcal{S}^h(1^n, \sigma, b_0, b_1)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{\text{c}}{\equiv} \left\{\text{VIEW}^{\tilde{\pi}}_{\tilde{R}}(1^n, b_0, b_1, \sigma)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}}$$

The only difference between the two distributions is that in $\text{VIEW}^{\tilde{\pi}}_{\tilde{R}}(1^n, b_0, b_1, \sigma)$, π is actually executed and hence elements in $\text{VIEW}^{\tilde{\pi}}_{\tilde{R}}(1^n, b_0, b_1, \sigma)$ include a real view of the adversarial receiver $\mathcal{A}$ in π, whereas in $\left\{\mathcal{S}^h(1^n, \sigma, b_0, b_1)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}}$, π is not executed and hence elements in this distributions contain an output of $\mathcal{SIM}$. Hence intuitively the assumption that $\mathcal{SIM}$ is a good simulator implies that the two distributions are indistinguishable. The formal proof of this is almost identical to the proof of Eq. (7) in Theorem 4.1.

We now prove that the output of the hybrid simulator $\mathcal{S}^h$ is statistically close to the output of the actual simulator $\mathcal{S}_{\tilde{R}}$. That is,

$$\left\{\mathcal{S}_{\tilde{R}}(1^n, \sigma, b_\sigma)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}} \stackrel{\text{s}}{\equiv} \left\{\mathcal{S}^h(1^n, \sigma, b_0, b_1)\right\}_{b_0,b_1,\sigma\in\{0,1\},n\in\mathbb{N}}$$

The only difference between the two is that in $\mathcal{S}_{\tilde{R}}(1^n, \sigma, b_\sigma)$, it holds that $z_{1-\sigma}$ is a random bit whereas in $\mathcal{S}^h(1^n, \sigma, b_0, b_1)$, we have that $z_{1-\sigma} = \langle r_{1-\sigma}, \alpha_{c_{1-\sigma}} \rangle \oplus b_{1-\sigma}$. However, we show that since $c_{1-\sigma} = c'$ is chosen at random independently of the execution, and since $\mathcal{SIM}$ learns only the bits in the sender's input that correspond to c^*, with high probability there is enough uncertainty about $\langle \alpha_{c_{1-\sigma}}, r_{1-\sigma} \rangle$ and thus $z_{1-\sigma}$ is statistically close to a random bit.

To prove this formally, we first note that a receiver-view $v_{\tilde{R}}$ of Protocol 5.3 contains a receiver-view v_R in π, the strings c and c' and r_0, r_1, z_0, z_1. We note that it does not necessarily contain c^* and α_{c^*} (yet we include them anyway for clarity, and since this only strengthens the claim). Moreover, note that $v_R, c, c', c^*, \alpha_{c^*}, r_0, r_1$ and z_σ are generated exactly the same in both distributions and hence are identically distributed. We now restate what we want to prove. For every $\sigma, b_0, b_1 \in \{0,1\}$, we show that

$$\left\{v_R, c, c', c^*, \alpha_{c^*}, r_0, r_1, \langle r_\sigma, \alpha_c \rangle, \langle r_{1-\sigma}, \alpha_{c'} \rangle\right\} \stackrel{\text{s}}{\equiv} \left\{v_R, c, c', c^*, \alpha_{c^*}, r_0, r_1, \langle r_\sigma, \alpha_c \rangle, U_1\right\} \tag{11}$$

where U_1 is a random variable that is uniformly distributed over $\{0,1\}$. It suffices to show that, except with negligible probability, there exists an index $j \in \{1, \ldots, n\}$ such that $c'_j \neq c^*_j$, $r^j_{1-\sigma} = 1$ and $r^j_\sigma = 0$. This is due to the fact that if this holds then since $c'_j \neq c^*_j$ the receiver does not learn anything about $\alpha^j_{c'_j}$ (by the security of π). In addition, since the jth bit of α_c is zeroed by r^j_σ, the value $\langle r_\sigma, \alpha_c \rangle$ reveals nothing about $\alpha^j_{c'_j}$ (note that α^j_c may be correlated with $\alpha^j_{c'_j}$ and thus this is needed). Finally, since $r^j_{1-\sigma} = 1$, it follows that $r^j_{1-\sigma} \cdot \alpha^j_{c'_j} = \alpha^j_{c'_j}$ and so is uniformly distributed. This implies that $\langle r_{1-\sigma}, \alpha_{c'} \rangle$ is uniformly distributed since $\langle r_{1-\sigma}, \alpha_{c'} \rangle = (\sum_{i \neq j} r^i_{1-\sigma} \cdot \alpha^i_{c'_i}) + \alpha^j_{c'_j} \bmod 2$. Observing now that r_0, r_1, c' are all chosen at random and are of length n, a straightforward calculation yields that such a j exists except with at most negligible probability. This completes the proof of Eq. (11), demonstrating that Protocol 5.3 is a weak-OT protocol.

Step 2 – full-OT from weak-OT. It remains to show that any weak-OT protocol can be transformed into an OT that is fully correct and secure in the presence of semi-honest adversaries. This is achieved by simply running multiple executions of the weak-OT protocol and taking the majority result. By the Chernoff bound, if enough executions are run (say, $n \cdot p^2(n)$ where correctness is guaranteed with probability $\frac{1}{2} + \frac{1}{p(n)}$), then the majority result will be the correct one, except with negligible probability. Furthermore, the simulation is carried out by simply running the simulators of the weak-OT for each repetition; a standard hybrid argument (as used to prove sequential composition) shows that this yields a satisfactory simulation for the repeated protocol.

We conclude that the existence of an OT extension protocol that is secure for malicious adversaries and uses a logarithmic number of calls implies the existence of semi-honest OT. In order to show the existence of OT secure in the presence of malicious adversaries, one can simply apply the GMW compiler [7] (using the fact that OT implies one-way functions), or alternatively one could use the compilation of [11]). ■

References

[1] W. Aiello, Y. Ishai and O. Reingold. Priced Oblivious Transfer: How to Sell Digital Goods. In *EUROCRYPT 2001*, Springer-Verlag (LNCS 2045), pages 110–135, 2001.

[2] D. Beaver. Precomputing Oblivious Transfer. In *CRYPTO'95*, Springer-Verlag (LNCS 963), pages 97–109, 1995.

[3] D. Beaver. Correlated Pseudorandomness and the Complexity of Private Computations. In the 28*th STOC*, pages 479–488, 1996.

[4] R. Canetti. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology*, 13(1):143–202, 2000.

[5] S. Even, O. Goldreich and A. Lempel. A Randomized Protocol for Signing Contracts. In *Communications of the ACM,* 28(6):637–647, 1985.

[6] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The Relationship Between Public Key Encryption and Oblivious Transfer. In the 41*st FOCS*, page 325–335, 2000.

[7] O. Goldreich, S. Micali and A. Wigderson. How to Play any Mental Game – A Completeness Theorem for Protocols with Honest Majority. In 19*th STOC,* pages 218–229, 1987. For details see [10].

[8] O. Goldreich. A Note on Computational Indistinguishability. *Information Processing Letters*, 34(6):277–281, 1990.

[9] O. Goldreich. *Foundations of Cryptography: Volume 1 – Basic Tools.* Cambridge University Press, 2001.

[10] O. Goldreich. *Foundations of Cryptography: Volume 2 – Basic Applications.* Cambridge University Press, 2004.

[11] I. Haitner, Y. Ishai, E. Kushilevitz, Y. Lindell and E. Petrank. Black-Box Constructions of Protocols for Secure Computation. *SIAM Journal on Computing*, 40(2):225–266, 2011.

[12] Y. Ishai, J. Kilian, K. Nissim and E. Petrank. Extending Oblivious Transfer Efficiently. In *CRYPTO 2003*, Springer (LNCS 2729), pages 145–161, 2003.

[13] R. Impagliazzo and D. Zuckerman. How to Recycle Random Bits. In the 30*th FOCS*, 248–253, 1989.

[14] J. Kilian. Founding Cryptography on Oblivious Transfer. In the 20*th STOC*, pages 20–31, 1988.

[15] Y. Lindell and H. Zarosim. Adaptive Zero-Knowledge Proofs and Adaptively Secure Oblivious Transfer. In the *Journal of Cryptology,* 24(4):761-799, 2011. An extended abstract appeared in the 6*th TCC*, Springer (LNCS 5444), pages 183–201, 2009.

[16] M. Rabin. How to Exchange Secrets by Oblivious Transfer. *Tech. Memo TR-81,* Aiken Computation Laboratory, Harvard University, 1981.

[17] S. Wolf and J. Wullschleger. Oblivious Transfer is Symmetric. In *EUROCRYPT 2006*, Springer (LNCS 4004), pages 222–232, 2006.

[18] A. Yao. How to Generate and Exchange Secrets. In 27*th FOCS*, pages 162–167, 1986.