2015 IEEE 56th Annual Symposium on Foundations of Computer Science

Limits on the Power of Indistinguishability Obfuscation and Functional Encryption

Gilad Asharov
Hebrew University of Jerusalem
Jerusalem 91904, Israel
asharov@cs.huji.ac.il

Gil Segev
Hebrew University of Jerusalem
Jerusalem 91904, Israel
segev@cs.huji.ac.il

Abstract

Recent breakthroughs in cryptography have positioned indistinguishability obfuscation as a "central hub" for almost all known cryptographic tasks, and as an extremely powerful building block for new cryptographic tasks resolving long-standing and foundational open problems. However, constructions based on indistinguishability obfuscation almost always rely on non-black-box techniques, and thus the extent to which it can be used as a building block in cryptographic constructions has been completely unexplored so far.

We present a framework for proving meaningful negative results on the power of indistinguishability obfuscation. By considering indistinguishability obfuscation for *oracle-aided* circuits, we capture the common techniques that have been used so far in constructions based on indistinguishability obfuscation. These include, in particular, *non-black-box* techniques such as the punctured programming approach of Sahai and Waters (STOC '14) and its variants, as well as sub-exponential security assumptions.

Within our framework we prove the first negative results on the power of indistinguishability obfuscation and of the tightly related notion of functional encryption. Our results are as follows:

- There is no fully black-box construction of a collision-resistant function family from an indistinguishability obfuscator for oracle-aided circuits.
- There is no fully black-box construction of a key-agreement protocol with perfect completeness from a private-key functional encryption scheme for oracle-aided circuits.

Specifically, we prove that any such potential constructions must suffer from an exponential security loss, and thus our results cannot be circumvented using sub-exponential security assumptions. Our framework captures constructions that may rely on a wide variety of primitives in a non-black-box manner (e.g., obfuscating or generating a functional key for a function that uses the evaluation circuit of a puncturable pseudorandom function), and we only assume that the underlying indistinguishability obfuscator or functional encryption scheme themselves are used in a black-box manner.

I. Introduction

The study of program obfuscation within the cryptography community was initiated by Barak, Goldreich, Impagliazzo, Rudich, Sahai, Vadhan, and Yang [2], [3] with the goal of understanding the extent to which programs can be made "unintelligible" while preserving their functionality. Barak et al. introduced two notions of obfuscation, virtual black-box obfuscation and indistinguishability obfuscation, which offer two such extents: Virtual black-box obfuscation asks that an obfuscated program reveals no more information than a black box implementing the program, and indistinguishability obfuscation asks, somewhat more modestly, that obfuscations of any two functionally-equivalent programs be computationally indistinguishable.

Barak et al. showed that general-purpose virtual black-box obfuscation is impossible to achieve, and left open the problem of whether or not indistinguishability obfuscation exists[1]. This has rooted the cryptography community with a somewhat pessimistic view, as it was not at all clear whether indistinguishability obfuscation (if even exists) is nearly as useful as virtual black-box obfuscation. In a recent breakthrough result, Garg, Gentry, Halevi, Raykova, Sahai, and Waters [6] proposed the first candidate construction of a general-purpose indistinguishability obfuscator, and showed how to apply indistinguishability obfuscation for constructing the first general-purpose functional-encryption scheme.

The power of indistinguishability obfuscation. The initial breakthrough by Garg et al. [6] has motivated Sahai and Waters [7] to carry out a systematic study of the applicability of indistinguishability obfuscation. Despite the somewhat modest flavor of obfuscation that is offered by indistinguishability obfuscation, Sahai and Waters presented strikingly-surprising constructions, positioning indistinguishability obfuscation as a "central hub" in cryptography. Specifically, relying

Due to space limitations we refer the reader to the full version of this work [1].

Supported by the European Union's Seventh Framework Programme (FP7) via a Marie Curie Career Integration Grant, by the Israel Science Foundation (Grant No. 483/13), and by the Israeli Centers of Research Excellence (I-CORE) Program (Center No. 4/11).

[1]Nevertheless, there are examples of function families that can be obfuscated in a virtual black-box manner under strong assumptions (e.g., [4], [5]).

0272-5428/15 $31.00 © 2015 IEEE
DOI 10.1109/FOCS.2015.21

on indistinguishability obfuscation Sahai and Waters resolved the long-standing foundational open problem of constructing a deniable encryption scheme [8], as well as constructed a variety of core cryptographic objects such as public-key encryption, short "hash-and-sign" signatures, CCA-secure public-key encryption, non-interactive zero-knowledge proofs, injective trapdoor functions, and oblivious transfer.

Following the work of Sahai and Waters, extensive research has been recently devoted to the applicability of indistinguishability obfuscation, positioning it as a central hub for almost all known cryptographic tasks, and as an extremely powerful building block for new cryptographic tasks resolving long-standing and foundational open problems. Applications of indistinguishability obfuscations range from fundamental building blocks such as one-way functions [9] and trapdoor permutations [10], to more complex tasks such as full-domain hash without random oracles [11], multiparty key exchange [12], efficient traitor tracing [12], multi-input functional encryption [13], [14], functional encryption for randomized functionalities [15], adaptively-secure multiparty computation [16], [17], [18], [19], communication-efficient secure computation [20], adaptively-secure functional encryption [21], polynomially-many hardcore bits for any one-way function [22], ZAPs and non-interactive witness-indistinguishable proofs [23], constant-round zero-knowledge proofs [24], fully-homomorphic encryption [25], and even cryptographic hardness for the complexity class PPAD [26].

Despite the extensive recent research on indistinguishability obfuscation, the following fundamental question has remained completely open:

Is there a natural task that cannot be solved using indistinguishability obfuscation?

The power of functional encryption. Functional encryption [27], [28] allows tremendous flexibility when accessing encrypted data: it supports restricted decryption keys that allow users to learn specific functions of the encrypted data and nothing else. The notion of functional encryption is tightly related to that of indistinguishability obfuscation. On one hand, Garg et al. [6] showed that indistinguishability obfuscation implies general-purpose public-key functional encryption. On the other hand, Ananth and Jain [14] and Bitansky and Vaikuntanathan [29] showed that general-purpose public-key functional encryption with sub-exponential security (and some additional requirements) implies general-purpose indistinguishability obfuscation.

Recent research has explored the private-key variant of functional encryption [30], [31], [14], [32], [33], [34], showing that it is a very useful building block that can essentially be used instead of indistinguishability obfuscation in various cases based on the notion of function privacy [32]. Although private-key functional encryption seems significantly weaker than its public-key variant, constructions of private-key functional encryption schemes are currently known based only on public-key functional encryption [6], [35], [21]. At the same time, however, we currently do not know how to construct any public-key primitive based on private-key functional encryption, and this raises the following fundamental question[2]:

Does general-purpose private-key functional encryption imply public-key cryptography?

The challenge of dealing with non-black-box constructions. The above-stated open questions clearly call for a study of the limitations of using indistinguishability obfuscation and functional encryption as building blocks in cryptographic constructions[3]. In general, one cannot argue about limitations of cryptographic constructions without placing any restrictions (as long as we believe that the cryptographic primitives under consideration exist – since a construction may simply ignore its building blocks). In an effort to capture what is meant by a "natural" construction of one primitive from another, Impagliazzo and Rudich [38] introduced the framework of black-box constructions, which has been extensively and successfully used over the years for studying the limitations of cryptographic constructions.[4]

The main challenge in the setting of indistinguishability obfuscation and functional encryption is that constructions that are based on either one of these two primitives almost always have a non-black-box ingredient. Specifically, such constructions obfuscate or generate a functional key for a function that uses the *description* of a cryptographic primitive (e.g., the evaluation circuit of a puncturable pseudorandom function or of a pseudorandom generator). Thus, any study of the limitations of indistinguishability obfuscation and functional encryption must face the challenge of dealing with such non-black-box ingredients.

[2]We note that for other forms of encryption, such as homomorphic encryption or re-randomizable encryption, it is known that any private-key scheme implies a public-key one [36], [25].

[3]Whereas various efforts have already been devoted to studying the *existence* of indistinguishability obfuscation (see, for example, [37] and the references therein), our goal is to explore the *limitations* of indistinguishability obfuscation and functional encryption.

[4]Informally, a black-box construction is a construction that "ignores the internal structure" of its underlying building blocks.

A. Our Contributions

In this work we prove the first limitations on the power of indistinguishability obfuscation and functional encryption as building blocks in cryptographic constructions. Our results are obtained within a subtle framework capturing constructions that may rely on a wide variety of primitives in a non-black-box manner (e.g., obfuscating or generating a functional key for a function that uses the evaluation circuit of a puncturable pseudorandom function), and we only assume that the underlying indistinguishability obfuscator or functional encryption scheme themselves are used in a black-box manner. Specifically, we observe that by considering indistinguishability obfuscation and functional encryption for *oracle-aided* circuits, we can capture the common techniques that have been used so far in constructions based on indistinguishability obfuscation. These include, in particular, *non-black-box* techniques such as the punctured programming approach of Sahai and Waters [7] and its variants, as well as sub-exponential security assumptions.

Within our framework we begin by exploring the limitations of indistinguishability obfuscation and prove the following theorem:

Theorem (informal) I.1. *There is no fully black-box construction of a collision-resistant function family from a general-purpose indistinguishability obfuscator and a one-way permutation (and even a trapdoor permutation).*

Specifically, we prove that any such potential construction must suffer from an exponential security loss, and thus sub-exponential security assumptions (e.g., [14], [26], [25], [10], [29]) are not useful for circumventing our negative result.

This shows that not only there exists a cryptographic task that cannot be solved using indistinguishability obfuscation using the common techniques, but in fact identifies collision-resistant hashing (one of the most fundamental primitives) as a specific such task. In turn, this implies a similar result for any primitive that is known to imply collision-resistant hashing in a fully black-box manner, such as homomorphic encryption, homomorphic commitments, two-message private information retrieval, and more (see [39]).[5]

We then explore the limitations of private-key functional encryption and prove the following theorem:

Theorem (informal) I.2. *There is no fully black-box construction of a key-agreement protocol with perfect completeness from a general-purpose private-key functional encryption scheme and a one-way permutation.*

As with Theorem I.1, we prove that any such potential construction must suffer from an exponential security loss. Moreover, Theorem I.2 holds even when the underlying functional encryption scheme is "compact" [14], [29] in the sense that the efficiency of its encryption algorithm depends only on the security parameter and on length of the plaintext (in particular, it is independent of the complexity of the function family supported by the scheme). This result positions private-key functional encryption as a private-key primitive unless additional non-black-box techniques are used.

Finally, we show that Theorem I.2 can be extended for separating indistinguishability obfuscation for oracle-aided circuits from private-key functional encryption for oracle-aided circuits (this does not follow immediately from Theorem I.2 since indistinguishability obfuscation and one-way permutations are not known to imply key agreement in a black-box manner). We prove the following theorem:

Theorem (informal) I.3. *There is no fully black-box construction of an indistinguishability obfuscator for oracle-aided circuits from a private-key functional encryption scheme for oracle-aided circuits.*

As with Theorems I.1 and I.2, we prove that any such potential construction must suffer from an exponential security loss. An impossibility result of a somewhat similar flavor was proved by Goldwasser and Rothblum [40] who showed that in the *programmable* random-oracle model there exists a class of oracle-aided circuits for which there is no indistinguishability obfuscator[6]. Our result is obtained by considering a more structured oracle relative to which we prove that there exists a private-key functional encryption scheme for polynomial-size oracle-aided circuits, but there is no indistinguishability obfuscator for polynomial-size oracle-aided circuits. Our result can be viewed as strengthening that of Goldwasser and Rothblum, both by avoiding any flavor of *oracle programmability*, and by considering a seemingly stronger building block (specifically, a private-key functional encryption scheme for oracle-aided circuits in our case vs. a random function in their case).

Finally, we note that unlike Theorems I.1 and I.2 that show separation results in the *standard model*, Theorem I.3 does not necessarily imply a separation result in the standard model (we refer the reader to [40] for a discussion of the possible

[5] A construction of a (fully-)homomorphic encryption scheme based on indistinguishability obfuscation with sub-exponential security was recently given by Canetti et al. [25]. Their construction assumes, in addition, re-randomizable encryption – which is currently known to exists based only on assumptions that already imply collision-resistant hashing.

[6] A different impossibility result for obfuscation considered the significantly stronger notion of virtual black-box obfuscation – see [41], [42] who extended the work of Barak et al. [3] to the random-oracle model.

implications of results of this flavor). Specifically, Theorem I.3 does not necessarily imply a separation in the standard model, since it may be that there exists an indistinguishability obfuscator for circuits, but there does not exist such an obfuscator for *oracle-aided* circuits. Nevertheless, this provides substantial evidence that *private-key* functional encryption is somewhat unlikely to imply indistinguishability obfuscation using the common techniques.

B. Related Work

Our approach in this paper is inspired by a combination of various approaches and ideas that were developed in previous work. Our framework for capturing certain non-black-box techniques in constructions that are based on indistinguishability obfuscation and functional encryption is inspired by the work of Brakerski, Katz, Yerukhimovich and Segev [43]. Their work addressed the question of whether non-black-box techniques that are originated from zero-knowledge proofs can be used for circumventing known impossibility results for black-box constructions. Although our work shares the same theme of capturing non-black-box techniques in a black-box manner, we capture a different class of non-black-box techniques, which raises many additional difficulties when compared to their work.

Our impossibility result for collision-resistant hashing generalizes the approach of Haitner et al. [44], who in turn generalized the ideas of Simon [45], Gennaro et al. [46], and Wee [47]. Specifically, we rely on Simon's collision-finding oracle, but since we also use additional (inefficient) oracles for implementing indistinguishability obfuscation (and these oracles can interact with Simon's oracle), we again deal with many additional difficulties when compared to previous work.

Our impossibility result for key-agreement protocols is inspired by ideas that were developed in the early work of Impagliazzo and Rudich [38], in its improvement by Barak and Mahmoody-Ghidary [48], and in the work of Brakerski, Katz, Yerukhimovich and Segev [43] who focused on the case of perfectly-complete protocols. We refer the reader to the work of Reingold, Trevisan and Vadhan [49] and to various recent impossibility results (see, for example, [50], [47], [51], [48], [52], [53], [54], [55], [56] and the references therein) for an overview of black-box reductions and the known impossibility results.

C. Overview of Our Approach

In this section we provide a high-level overview of our approach. First, we describe our framework that enables us to prove meaningful impossibility results for constructions that are based on indistinguishability obfuscation and functional encryption. Then, within our framework, we describe the main ideas underlying our impossibility results.

Capturing non-black-box constructions via oracle-aided circuits. The fact that constructions that are based on indistinguishability obfuscation or functional encryption are almost always *non-black-box* makes it extremely challenging to prove any impossibility results. For example, a typical such construction would apply the obfuscator or the key-generation algorithm to a function that uses the evaluation circuit of a pseudorandom generator or a pseudorandom function, and this requires *specific implementations* of its underlying building blocks.

However, most of the non-black-box techniques that are used on such constructions have essentially the same flavor: The obfuscator or the key-generation algorithm are applied to functions that can be constructed in a fully black-box manner from a low-level primitive, such as a one-way function or a trapdoor permutation. In particular, the vast majority of constructions rely on the obfuscator or the functional encryption scheme themselves in a black-box manner. Thus, when considering the stronger primitives of indistinguishability obfuscation and functional encryption for *oracle-aided* circuits (see Sections II-A and II-B), these non-black-box techniques in fact directly translate into black-box ones. These include, in particular, non-black-box techniques such as the punctured programming approach of Sahai and Waters [7] and its variants (as well as sub-exponential security assumptions – which are already captured by most frameworks for black-box impossibility results).

Example: The Sahai-Waters punctured programming approach. Consider, for example, the construction of a public-key encryption scheme from a one-way function and a general-purpose indistinguishability obfuscator by Sahai and Waters [7]. Their construction relies on the underlying one-way function in a non-black-box manner. However, relative to an oracle that allows the existence of a one-way function f and indistinguishability obfuscation $i\mathcal{O}$ for *oracle-aided* circuits C^f, it is in fact a fully black-box construction. Specifically, Sahai and Waters use the underlying indistinguishability obfuscator for obfuscating a circuit that invokes a puncturable pseudorandom function and a pseudorandom generator as sub-routines. Given that puncturable pseudorandom functions and pseudorandom generators can be based on any one-way function in a fully black-box manner, from our perspective such a circuit is a polynomial-size oracle-aided circuit C^f – which can be obfuscated using $i\mathcal{O}$ (we refer to reader to Section IV-C for more details).

This reasoning naturally extends to various variants of the punctured programming approach by Sahai and Waters [7]. However, it does not extend to constructions that may rely on the obfuscator or the functional encryption scheme themselves

in a non-black-box manner (e.g., [23])[7], or constructions that may rely on zero-knowledge techniques and require using NP reductions[8].

Our separation results. Within our framework we obtain our results via two oracle separations on which we now elaborate. For ruling out a fully black-box construction of a primitive X from a primitive Y, it suffices to construct an oracle relative to which there exists an implementation of Y but any implementation of X can be efficiently broken by an oracle-aided algorithm (see, for example, [38], [49]).

We prove Theorem I.1 by presenting an oracle Γ relative to which the following three properties hold: (1) there exists an exponentially-secure one-way permutation (and even a trapdoor permutation family), (2) there exists an exponentially-secure general-purpose obfuscator for oracle-aided circuits, and (3) there does not exist a collision-resistant function family. Our oracle Γ is quite structured and consists of three main ingredients. The first ingredient is a random permutation f that will serve as a one-way permutation. The second ingredient is a pair $(\mathcal{O}, \mathsf{Eval}^{f,\mathcal{O}})$ of functions, where $\mathcal{O}$ is a random permutation that will serve as an obfuscator (obfuscating a circuit C is done by computing a "handle" $\mathcal{O}(C, r)$ for a randomly-chosen string r), and Eval is a function that enables evaluations of obfuscated circuits (Eval has access to both f and $\mathcal{O}$): Given a handle $\mathcal{O}(C, r)$ and an input x, it "finds" C and returns $C(x)$. The third and final ingredient is the "collision-finding" oracle of Simon [45] and Haitner et al. [44] that takes as input any oracle-aided circuit (which may access f, $\mathcal{O}$, and Eval), and returns a random collision for the circuit. It is straightforward to prove that relative to Γ there are no collision-resistant function families, and the vast majority of our effort is in showing that relative to Γ there exist a one-way permutation and an indistinguishability obfuscator. We refer the reader to Section III for more information.

We prove Theorems I.2 and I.3 by presenting an oracle Ψ relative to which the following three properties hold: (1) there exists an exponentially-secure one-way permutation, (2) there exists an exponentially-secure private-key functional encryption scheme for oracle-aided circuits, and (3) there does not exist a perfectly-complete bit-agreement protocol. Our oracle Ψ is again quite structured and consists of two main ingredients. The first ingredient is again a random permutation f that will serve as a one-way permutation. The second ingredient is a triplet $(\mathcal{K}, \mathcal{E}, \mathcal{D}^{f,\mathcal{K},\mathcal{E}})$ of functions that we use for implementing a functional encryption scheme as follows. A functional key for an oracle-aided circuit C is a handle $\mathcal{K}(\mathsf{msk}, C)$, and an encryption of a message m is a handle $\mathcal{E}(\mathsf{msk}, m, r)$, where msk is the master-secret key, $\mathcal{K}$ and $\mathcal{E}$ are randomly-chosen functions, and r is a randomly-chosen string for each encryption. The decryption algorithm is implemented using $\mathcal{D}^{f,\mathcal{K},\mathcal{E}}$ which on input a functional key $\mathsf{sk}_C = \mathcal{K}(\mathsf{msk}, C)$ and an encryption $\mathsf{ct} = \mathcal{E}(\mathsf{msk}, m, r)$ "finds" C and m, and then returns $C^f(m)$ if and only if sk_C and ct were indeed computed using the same master secret key msk. Here, as opposed to the proof of Theorem I.1, each of the above-specified three properties is quite challenging to prove, and we refer the reader to Section IV for more information.

D. Paper Organization

The remainder of this paper is organized as follows. In Section II we introduce the cryptographic primitives under consideration in this paper. In Sections III and IV we present our results on the limitations of indistinguishability obfuscation and functional encryption, respectively. In each of these sections, we first define the class of constructions to which our result applies, and then provide an overview of both the main ideas underling the proof and of the structure of the proof. Due to space limitations we refer the reader to the full version of this work [1] for the formal proofs.

II. Preliminaries

In this section we present the notation and basic definitions that are used in this work. For a distribution X we denote by $x \leftarrow X$ the process of sampling a value x from the distribution X. Similarly, for a set $\mathcal{X}$ we denote by $x \leftarrow \mathcal{X}$ the process of sampling a value x from the uniform distribution over $\mathcal{X}$. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \ldots, n\}$. A function $\mathsf{negl} : \mathbb{N} \to \mathbb{R}^+$ is *negligible* if for every constant $c > 0$ there exists an integer N_c such that $\mathsf{negl}(n) < n^{-c}$ for all $n > N_c$. Throughout the paper, we denote by n the security parameter.

A. Indistinguishability Obfuscation for Oracle-Aided Circuits

We consider the standard notion of indistinguishability obfuscation [2], [3], [6] when naturally generalized to the setting of oracle-aided computations. We first define the notion of functional equivalence relative to a specific function (provided as an oracle), and then we define the notion of an indistinguishability obfuscation for a class of oracle-aided circuits. In

[7]With the exception of obfuscating a function that may invoke an indistinguishability obfuscator in a black-box manner. This is captured by our approach – see Section III-A.

[8]Such techniques are captured by the work of Brakerski et al. [43], and we leave it as an intriguing open problem to see whether the two approaches for capturing non-black-box techniques can be unified.

what follows, when considering a class $\mathcal{C} = \{\mathcal{C}_n\}_{n\in\mathbb{N}}$ of oracle-aided circuits, we assume that each $\mathcal{C}_n$ consists of circuits of size at most n.

Definition II.1. *Let C_1 and C_2 be two oracle-aided circuits, and let f be a function. We say that C_1 and C_2 are* functionally equivalent relative to *f, denoted $C_1^f \equiv C_2^f$, if for any input x it holds that $C_1^f(x) = C_2^f(x)$.*

Definition II.2. *A probabilistic polynomial-time algorithm $i\mathcal{O}$ is an* indistinguishability obfuscator *relative to an oracle Γ for a class $\mathcal{C} = \{\mathcal{C}_n\}_{n\in\mathbb{N}}$ of oracle-aided circuits if the following conditions are satisfied:*

- **Functionality.** *For all $n \in \mathbb{N}$ and for all $C \in \mathcal{C}_n$ it holds that*

$$\Pr\left[C^\Gamma \equiv \widehat{C}^\Gamma \ : \ \widehat{C} \leftarrow i\mathcal{O}^\Gamma(1^n, C)\right] = 1.$$

- **Indistinguishability.** *For any probabilistic polynomial-time distinguisher $D = (D_1, D_2)$ there exists a negligible function $\mathsf{negl}(\cdot)$ such that*

$$\mathsf{Adv}^{\mathsf{iO}}_{\Gamma,i\mathcal{O},D,\mathcal{C}}(n) \stackrel{\mathsf{def}}{=} \left|\Pr\left[\mathsf{Exp}^{\mathsf{iO}}_{\Gamma,i\mathcal{O},D,\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \leq \mathsf{negl}(n)$$

for all sufficiently large $n \in \mathbb{N}$, where the random variable $\mathsf{Exp}^{\mathsf{iO}}_{\Gamma,i\mathcal{O},D,\mathcal{C}}(n)$ is defined via the following experiment:

1) $b \leftarrow \{0,1\}$.
2) $(C_0, C_1, \mathsf{state}) \leftarrow D_1^\Gamma(1^n)$ *where* $C_0, C_1 \in \mathcal{C}_n$ *and* $C_0^\Gamma \equiv C_1^\Gamma$.
3) $\widehat{C} \leftarrow i\mathcal{O}^\Gamma(1^n, C_b)$.
4) $b' \leftarrow D_2^\Gamma(\mathsf{state}, \widehat{C})$.
5) *If $b' = b$ then output 1, and otherwise output 0.*

Throughout this paper we sometimes find it convenient to denote by $\mathsf{Exp}^{\mathsf{iO}}_{\Gamma,i\mathcal{O},D,\mathcal{C}}(n; b, r^*)$ the above experiment when using specific values b and r^*, where b is the bit chosen in the first step of the experiment and r^* is the randomness used by the algorithm $i\mathcal{O}$ to obfuscating the challenge circuit in Step 3 of the experiment.

B. Private-Key Functional Encryption for Oracle-Aided Circuits

A private-key functional encryption scheme over a message space $\mathcal{M} = \{\mathcal{M}_n\}_{n\in\mathbb{N}}$ and a function space $\mathcal{C} = \{\mathcal{C}_n\}_{n\in\mathbb{N}}$ is a quadruple $(\mathsf{Setup}, \mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ of probabilistic polynomial-time oracle-aided algorithms. The setup algorithm Setup takes as input the unary representation 1^n of the security parameter $n \in \mathbb{N}$ and outputs a master secret key msk. The key-generation algorithm KG takes as input a master secret key msk and a circuit $C \in \mathcal{C}_n$, and outputs a functional key sk_C. The encryption algorithm Enc takes as input a master secret key msk and a message $m \in \mathcal{M}_n$, and outputs a ciphertext ct. In terms of correctness we require that for all sufficiently large $n \in \mathbb{N}$, for every circuit $C \in \mathcal{C}_n$ and message $m \in \mathcal{M}_n$ it holds that $\mathsf{Dec}(\mathsf{KG}(\mathsf{msk}, C), \mathsf{Enc}(\mathsf{msk}, m)) = C(m)$ with all but a negligible probability over the internal randomness of the algorithms Setup, KG, and Enc.

In terms of security, we rely on the standard private-key indistinguishability-based notion of adaptive security (see, for example, [30], [32]) when naturally generalized to the setting of oracle-aided computations. For simplicity, we consider a *single*-challenge security notion, and we note that it is in fact equivalent to a *multiple*-challenge one (via a standard hybrid argument – as we allow adversaries to access an encryption oracle).

Definition II.3 (Valid adversary)**.** *A probabilistic polynomial-time algorithm $\mathcal{A}$ is a* valid adversary relative to an oracle Ψ *if for all private-key functional encryption schemes $\Pi = (\mathsf{Setup}, \mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$, for all $n \in \mathbb{N}$ and $b \in \{0,1\}$, and for all oracle-aided circuits C with which $\mathcal{A}$ queries the key-generation oracle KG it holds that $C^\Psi(m_0) = C^\Psi(m_1)$, where m_0 and m_1 are the challenge messages produced by $\mathcal{A}$.*

Definition II.4 (Adaptive security)**.** *A private-key functional encryption scheme $\Pi = (\mathsf{Setup}, \mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ over a message space $\mathcal{M} = \{\mathcal{M}_n\}_{n\in\mathbb{N}}$ and an oracle-aided function space $\mathcal{C} = \{\mathcal{C}_n\}_{n\in\mathbb{N}}$ is* adaptively secure relative to an oracle Ψ *if for any probabilistic polynomial-time valid adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function $\mathsf{negl}(\cdot)$ such that*

$$\mathsf{Adv}^{\mathsf{FE}}_{\Psi,\Pi,\mathcal{A},\mathcal{C}}(n) \stackrel{\mathsf{def}}{=} \left|\Pr\left[\mathsf{Exp}^{\mathsf{FE}}_{\Psi,\Pi,\mathcal{A},\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \leq \mathsf{negl}(n)$$

for all sufficiently large $n \in \mathbb{N}$, where the random variable $\mathsf{Exp}^{\mathsf{FE}}_{\Psi,\Pi,\mathcal{A},\mathcal{C}}(n)$ is defined via the following experiment:

1) $\mathsf{msk} \leftarrow \mathsf{Setup}^\Psi(1^n)$, $b \leftarrow \{0,1\}$.
2) $(m_0, m_1, \mathsf{state}) \leftarrow \mathcal{A}_1^{\Psi, \mathsf{KG}^\Psi(\mathsf{msk},\cdot), \mathsf{Enc}^\Psi(\mathsf{msk},\cdot)}(1^n)$.

3) $c^* \leftarrow \mathsf{Enc}^{\Psi}(\mathsf{msk}, m_b; r^*)$, *where* $r^* \leftarrow \{0,1\}^*$.
4) $b' \leftarrow \mathcal{A}_2^{\Psi, \mathsf{KG}^{\Psi}(\mathsf{msk},\cdot), \mathsf{Enc}^{\Psi}(\mathsf{msk},\cdot)}(\mathsf{state}, c^*)$.
5) *If* $b' = b$ *then output 1, and otherwise output 0.*

Throughout the paper we sometimes find it convenient to denote by $\mathsf{Exp}^{\mathsf{FE}}_{\Psi,\Pi,\mathcal{A},\mathcal{C}}(n; \mathsf{msk}, b, r^*)$ the above experiment when using specific values msk, b, and r^*.

C. Oracle-Aided Collision-Resistant Function Families

We rely on the standard notion of a collision-resistant function family when generalized to the setting of oracle-aided computations. For our purposes in this paper it suffices to consider functions that compress their input by a single bit.

Definition II.5. *A pair* $(\mathsf{Gen}, \mathsf{Eval})$ *of polynomial-time oracle-aided algorithms is a collision-resistant function family relative to an oracle* Γ *if it satisfies the following properties:*

- *The index-generation algorithm* Gen *is a probabilistic algorithm that on input* 1^n *and oracle access to* Γ *outputs a function index* $\sigma \in \{0,1\}^{m(n)}$.
- *The evaluation algorithm* Eval *is a deterministic algorithm that takes as input a function index* $\sigma \in \{0,1\}^{m(n)}$ *and a string* $x \in \{0,1\}^n$*, has oracle access to* Γ*, and outputs a string* $y = \mathsf{Eval}^{\Gamma}(\sigma, x) \in \{0,1\}^{n-1}$.
- *For any probabilistic polynomial-time oracle-aided algorithm* $\mathcal{A}$ *there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\begin{array}{c} x \neq x' \text{ and} \\ \mathsf{Eval}^{\Gamma}(\sigma, x) = \mathsf{Eval}^{\Gamma}(\sigma, x') \end{array} \middle| \begin{array}{c} \sigma \leftarrow \mathsf{Gen}^{\Gamma}(1^n) \\ (x, x') \leftarrow \mathcal{A}^{\Gamma}(1^n, \sigma) \end{array}\right] \leq \mathsf{negl}(n)$$

for all sufficiently large $n \in \mathbb{N}$.

D. Oracle-Aided Key-Agreement Protocols

We rely on the standard notion of a key-agreement protocol when generalized to the setting of oracle-aided computations. For our purposes in this paper it suffices to consider key-agreement protocols in which the parties agree on a single bit, and we refer to such protocols as bit-agreement protocols.

A bit-agreement protocol consists of a pair $(\mathcal{A}, \mathcal{B})$ of probabilistic polynomial-time oracle-aided algorithms. We denote by $(k_{\mathcal{A}}, k_{\mathcal{B}}, T) \leftarrow \langle \mathcal{A}^{\Psi}(1^n; r_{\mathcal{A}}), \mathcal{B}^{\Psi}(1^n; r_{\mathcal{B}}) \rangle$ the random process of executing the protocol relative to an oracle Ψ, where $r_{\mathcal{A}}$ and $r_{\mathcal{B}}$ are the random tapes of $\mathcal{A}$ and $\mathcal{B}$, respectively, $k_{\mathcal{A}}$ and $k_{\mathcal{B}}$ are the output bits of $\mathcal{A}$ and $\mathcal{B}$, respectively, and T is the transcript of the protocol (i.e., the messages exchanged by the parties; this does not include the oracle queries/answers of the parties during the execution). In this paper we consider bit-agreement protocol that are perfectly complete (i.e., the parties always output the same bit).

Definition II.6. *A pair* $\Pi = (\mathcal{A}, \mathcal{B})$ *of probabilistic polynomial-time oracle-aided algorithms is a perfectly-complete bit-agreement protocol relative to an oracle* Ψ *if the following two conditions hold:*

- **Perfect completeness.** *For any* $n \in \mathbb{N}$ *it holds that*

$$\Pr_{r_{\mathcal{A}}, r_{\mathcal{B}}}\left[k_{\mathcal{A}} = k_{\mathcal{B}} \,\middle|\, (k_{\mathcal{A}}, k_{\mathcal{B}}, T) \leftarrow \langle \mathcal{A}^{\Psi}(1^n; r_{\mathcal{A}}), \mathcal{B}^{\Psi}(1^n; r_{\mathcal{B}}) \rangle\right] = 1.$$

- **Security.** *For any probabilistic polynomial-time oracle-aided algorithm* E*, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\mathsf{Adv}^{\mathsf{KA}}_{\Psi,\Pi,E}(n) \stackrel{\mathsf{def}}{=} \left|\Pr\left[\mathsf{Exp}^{\mathsf{KA}}_{\Psi,\Pi,E}(n) = 1\right] - \frac{1}{2}\right| \leq \mathsf{negl}(n)$$

for all sufficiently large $n \in \mathbb{N}$*, where the random variable* $\mathsf{Exp}^{\mathsf{KA}}_{\Psi,\Pi,E}(n)$ *is defined via the following experiment:*

1) $(k_{\mathcal{A}}, k_{\mathcal{B}}, T) \leftarrow \langle \mathcal{A}^{\Psi}(1^n), \mathcal{B}^{\Psi}(1^n) \rangle$.
2) $k' \leftarrow E^{\Psi}(1^n, T)$.
3) *If* $k' = k_{\mathcal{A}}$ *then output 1, and otherwise output 0.*

III. Limits on the Power of Indistinguishability Obfuscation

In this section we present our negative result for constructing a collision-resistant function family from a general-purpose indistinguishability obfuscator and a one-way permutation (and even an enhanced trapdoor permutation family). First, in Section III-A we formally define the class of constructions to which our negative result applies. Then, in Section III-B we present the structure of our proof, which is provided in the full version of this work [1].

A. The Class of Reductions

We consider fully black-box constructions of a collision-resistant function family from a general-purpose indistinguishability obfuscator and a one-way permutation (and even an enhanced trapdoor permutation family). we model these primitives as two independent building blocks due to the following four reasons. First, although indistinguishability obfuscation is known to imply one-way functions under reasonable assumptions [9], it is not known to imply one-way permutations unless one assumes sub-exponential hardness [10]. Second, this enables us to prove an *unconditional* result. Third, as we show in the full version of this work [1], our proof extends to the case where the one-way permutation is replaced by an enhanced trapdoor permutations family. Finally, and most importantly, this enables us to capture constructions that may apply the indistinguishability obfuscator to any primitive that can be constructed in a fully black-box manner from a one-way permutation (and, more general, from trapdoor permutations). For example, this enables us to capture constructions that may apply the indistinguishability obfuscator to any circuit that uses a puncturable pseudorandom function or a pseudorandom generator as a sub-routine[9].

Moreover, this also enables us to capture constructions that may apply the indistinguishability obfuscator to any primitive that can be constructed in a fully black-box manner from a one-way permutation (or trapdoor permutations) *and indistinguishability obfuscation*: Whenever an "outer" indistinguishability obfuscator is applied to a circuit that uses an "inner" indistinguishability obfuscator as a sub-routine, the functionality and security of the outer obfuscator imply that the inner obfuscator can be simply replaced by the identity function (with a suitable padding of its output). For example, this enables us to capture constructions that may apply the indistinguishability obfuscator to the decryption circuit of a general-purpose functional encryption scheme (such a scheme was recently constructed by Waters [21], where in our setting we view his construction as relying on indistinguishability obfuscation for circuits that use one-way functions in a fully black-box manner). We emphasize that *black-box* access in our setting is only required with respect to the indistinguishability obfuscator and the one-way permutation themselves (e.g., applying an indistinguishability obfuscator to a circuit that uses a pseudorandom generator as a sub-routine is considered black-box access).

We now formally define the class of constructions considered in this section, tailoring our definitions to the specific primitives under consideration. We consider any implementation of a one-way permutation f and an indistinguishability obfuscator $i\mathcal{O}$ for the class of all polynomial-size oracle-aided circuits C^f. Before providing the formal definition we remind the reader that two oracle-aided circuits, C_0 and C_1, are functionally equivalent relative to a function f if for any input x it holds that $C_1^f(x) = C_2^f(x)$ (see Definition II.1). The following definition is directly inspired by those of [62], [63], [47], [44].

Definition III.1. *A fully black-box construction of a collision-resistant function family from a one-way permutation and an indistinguishability obfuscator for the class $\mathcal{C} = \{\mathcal{C}_n\}_{n\in\mathbb{N}}$ of all polynomial-size oracle-aided circuits consists of a pair of probabilistic polynomial-time oracle-aided algorithms $(\mathsf{Gen}, \mathsf{Eval})$, an oracle-aided algorithm M that runs in time $T_M(\cdot)$, and functions $\epsilon_{M,1}(\cdot)$ and $\epsilon_{M,2}(\cdot)$, such that the following two conditions hold:*

- **Correctness:** *For any $n \in \mathbb{N}$, for any permutation f, for any function $i\mathcal{O}$ such that $i\mathcal{O}(C;r)^f \equiv C^f$ for all $C \in \mathcal{C}$ and $r \in \{0,1\}^*$, and for any function index σ produced by $\mathsf{Gen}^{f,i\mathcal{O}}(1^n)$, it holds that $\mathsf{Eval}^{f,i\mathcal{O}}(\sigma,\cdot) : \{0,1\}^n \to \{0,1\}^{n-1}$.*
- **Black-box proof of security:** *For any permutation f, for any function $i\mathcal{O}$ such that $i\mathcal{O}(C;r)^f \equiv C^f$ for all $C \in \mathcal{C}$ and $r \in \{0,1\}^*$, for any probabilistic oracle-aided algorithm $\mathcal{A}$ that runs in time $T_{\mathcal{A}} = T_{\mathcal{A}}(n)$, and for any function $\epsilon_{\mathcal{A}} = \epsilon_{\mathcal{A}}(n)$, if*

$$\Pr\left[\begin{array}{c} x \neq x' \text{ and} \\ \mathsf{Eval}^{f,i\mathcal{O}}(\sigma,x) = \mathsf{Eval}^{f,i\mathcal{O}}(\sigma,x') \end{array} \middle| \begin{array}{c} \sigma \leftarrow \mathsf{Gen}^{f,i\mathcal{O}}(1^n) \\ (x,x') \leftarrow \mathcal{A}^{f,i\mathcal{O}}(1^n,\sigma) \end{array}\right] \geq \epsilon_{\mathcal{A}}(n)$$

for infinitely many values of n, then either

$$\Pr_{x\leftarrow\{0,1\}^n}\left[M^{f,i\mathcal{O},\mathcal{A}}(f(x)) = x\right] \geq \epsilon_{M,1}\left(T_{\mathcal{A}}(n)\cdot\epsilon_{\mathcal{A}}^{-1}(n)\right)\cdot\epsilon_{M,2}(n)$$

or

$$\left|\Pr\left[\mathsf{Exp}^{\mathsf{iO}}_{(f,i\mathcal{O}),i\mathcal{O},M^{\mathcal{A}},\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \geq \epsilon_{M,1}\left(T_{\mathcal{A}}(n)\cdot\epsilon_{\mathcal{A}}^{-1}(n)\right)\cdot\epsilon_{M,2}(n)$$

for infinitely many values of n (see Definition II.2 for the description of the experiment $\mathsf{Exp}^{\mathsf{iO}}_{(f,i\mathcal{O}),i\mathcal{O},M^{\mathcal{A}},\mathcal{C}}$).

At this point we would like to highlight the following aspects of the above definition:

[9]We note that both puncturable pseudorandom functions and pseudorandom generators can be built from any one-way function in a fully black-box manner [57], [58], [59], [60], [7], [61].

- **$i\mathcal{O}$ can obfuscate *oracle-aided* circuits C^f.** As discussed above, our definition captures an indistinguishability obfuscator $i\mathcal{O}$ for the class of all polynomial-size oracle-aided circuits C^f, and thus *black-box* access in our setting is only required with respect to the indistinguishability obfuscator and the one-way permutation themselves. For example, this enables us to capture constructions that may apply the indistinguishability obfuscator to the evaluation circuit of a puncturable pseudorandom function, or to any circuit that uses a pseudorandom generator as a sub-routine.
  We note that our definition does not capture, for example, constructions that use non-interactive zero-knowledge (or witness-indistinguishable) proofs for languages that are defined relative to a one-way function. We leave it as an open problem to extend our framework to such constructions, noting that the approach of Brakerski et al. [43] seems quite promising in this direction.
- **The "security loss" functions.** Black-box constructions are typically formulated with a reduction algorithm M that runs in *polynomial* time and offers a *polynomial* security loss. In our setting, as we are interested in capturing constructions that may be based on super-polynomial security assumptions, we allow the algorithm M to run in arbitrary time $T_M(n)$ and to have an arbitrary security loss.
  In general, the security loss of a reduction is a function of the adversary's running time $T_{\mathcal{A}}(n)$, of its success probability $\epsilon_{\mathcal{A}}(n)$, and of the security parameter $n \in \mathbb{N}$. Following Luby [62] and Goldreich [63], we simplify the presentation by considering Levin's unified security measure $T_{\mathcal{A}}(n) \cdot \epsilon_{\mathcal{A}}^{-1}(n)$. Specifically, our definition captures the security loss of a reduction by considering an "adversary-dependent" security loss $\epsilon_{M,1}(T_{\mathcal{A}}(n) \cdot \epsilon_{\mathcal{A}}^{-1}(n))$, and an "adversary-independent" security loss $\epsilon_{M,2}(n)$. By considering arbitrary security loss functions, we are indeed able to capture constructions that rely on super-polynomial security assumptions. For example, in the recent construction of Bitansky et al. [10] (and in various other recent constructions based on indistinguishability obfuscation), the adversary-dependent loss is polynomial whereas the adversary-independent loss is sub-exponential[10].

B. Proof Overview and the Oracle Γ

Our result is obtained by presenting an oracle Γ relative to which there exist an exponentially-secure one-way permutation f and an exponentially-secure indistinguishability obfuscator $i\mathcal{O}$ for the class of all polynomial-size oracle-aided circuits C^f, but any collision-resistant function family can be easily broken. We prove the following theorem:

Theorem III.2. *Let* $(\mathsf{Gen}, \mathsf{Eval}, M, T_M, \epsilon_{M_1}, \epsilon_{M_2})$ *be a fully black-box construction of a collision-resistant function family from a one-way permutation* f *and an indistinguishability obfuscator for the class of all polynomial-size oracle-aided circuits* C^f *(see Definition III.1). Then, at least one of the following propertied holds:*

1) $T_M(n) \geq 2^{\zeta n}$ *for some constant* $\zeta > 0$ *(i.e., the reduction runs in exponential time).*
2) $\epsilon_{M,1}(n^c) \cdot \epsilon_{M,2}(n) \leq 2^{-n/50}$ *for some constant* $c > 1$ *(i.e., the security loss is exponential).*

In particular, the theorem implies that if the running time $T_M(\cdot)$ of the reduction is sub-exponential and the adversary-dependent security loss $\epsilon_{M,1}(\cdot)$ is polynomial as in the vast majority of constructions, then the adversary-independent security loss $\epsilon_{M,2}(\cdot)$ must be exponential (thus ruling out even constructions that rely on sub-exponential security assumptions).

In what follows we describe the oracle Γ (which is in fact a distribution over oracles), and then explain the structure of our proof. The proof is inspired by a combination of ideas that were developed in the work of Haitner et al. [44] (generalizing ideas of Simon [45], Gennaro et al. [46], and Wee [47]) and in the work of Brakerski et al. [43], as discussed below.

The oracle Γ. The oracle Γ is a quadruple $\left(f, \mathcal{O}, \mathsf{Eval}^{f,\mathcal{O}}, \mathsf{CollFinder}^{f,\mathcal{O},\mathsf{Eval},\mathcal{R}}\right)$ that is defined as follows:

- **The function $f = \{f_n\}_{n\in\mathbb{N}}$.** For every $n \in \mathbb{N}$ the function f_n is a uniformly chosen permutation over $\{0,1\}^n$.
  Looking ahead, we will prove that f is a one-way permutation relative to Γ (as detailed in the full version of this work [1], our result extends to the case where f is replaced by a family of trapdoor permutations).
- **The functions $\mathcal{O} = \{\mathcal{O}_n\}_{n\in\mathbb{N}}$ and $\mathsf{Eval}^{f,\mathcal{O}}$.** For every $n \in \mathbb{N}$ the function $\mathcal{O}_n$ is a uniformly chosen permutation over $\{0,1\}^{2n}$. The function $\mathsf{Eval}^{f,\mathcal{O}}$ on input $(y, x) \in \{0,1\}^*$ finds the unique oracle-aided circuit $C \in \{0,1\}^{|y|/2}$ (i.e., C is encoded as a $|y|/2$-bit string) and the unique string $r \in \{0,1\}^{|y|/2}$ such that $\mathcal{O}_{|y|/2}(C, r) = y$ (note that the uniqueness is guaranteed by the fact that $\mathcal{O}_{|y|/2}$ is a permutation over $\{0,1\}^{|y|}$). Then, it computes and outputs $C^f(x)$.
  Looking ahead, we will use $\mathcal{O}$ and Eval for realizing an indistinguishability obfuscator $i\mathcal{O}$ relative to Γ for the class of all polynomial-size oracle-aided circuits C^f.
- **The function $\mathsf{CollFinder}^{f,\mathcal{O},\mathsf{Eval},\mathcal{R}}$.** On input an encoding of an oracle-aided circuit C that may access f, $\mathcal{O}$ and Eval, the function $\mathsf{CollFinder}^{f,\mathcal{O},\mathsf{Eval},\mathcal{R}}$ outputs a uniform pair (w, w') such that $C^{f,\mathcal{O},\mathsf{Eval}}(w) = C^{f,\mathcal{O},\mathsf{Eval}}(w')$.

[10]This is also the situation, for example, when using "complexity leveraging" for arguing that any selectively-secure identity-based encryption scheme is in fact adaptively secure.

Looking ahead, we will use $\mathsf{CollFinder}$ for finding a non-trivial collision in any circuit that compresses its input. More specifically, $\mathsf{CollFinder}$ is provided with an infinite collection $\mathcal{R}$ of permutations, where for every possible circuit $C^{f,\mathcal{O},\mathsf{Eval}}:\{0,1\}^m \to \{0,1\}^{m'}$ the collection $\mathcal{R}$ contains two uniformly and independently sampled permutations, π^1_C and π^2_C, over $\{0,1\}^m$. Now, given a circuit $C^{f,\mathcal{O},\mathsf{Eval}}$, the oracle $\mathsf{CollFinder}$ sets $w = \pi^1_C(0^m)$, and then computes $w' = \pi^2_C(t)$ for the lexicographically smallest $t \in \{0,1\}^m$ such that $C(\pi^2_C(t)) = C(w)$.[11]

Equipped with the oracle Γ, our proof consists of the following three parts.

Part 1: Finding collisions in any compressing circuit. As in the work of Simon [45], it is straightforward to observe that there are no collision-resistant function families relative to Γ. Specifically, for any $n \in \mathbb{N}$, for any functions f and $\mathcal{O}$ as above, and for any oracle-aided circuit $C = C^{f,\mathcal{O},\mathsf{Eval}^{f,\mathcal{O}}}:\{0,1\}^n \to \{0,1\}^{n-1}$, querying $\mathsf{CollFinder}$ with C results in a non-trivial collision with probability at least $1/4$ over the choice of the relevant permutations in $\mathcal{R}$. More generally, n independent such queries result in at least one non-trivial collision with probability at least $1-(3/4)^n$.[12]

Part 2: The existence of an indistinguishability obfuscator. The most challenging part in this section is in proving that there exists a general-purpose indistinguishability obfuscator relative to Γ. Our construction of such an obfuscator $i\mathcal{O}^{\mathcal{O}}$ is quite intuitive: For obfuscating an oracle-aided circuit $C = C^f \in \{0,1\}^n$ (i.e., a circuit that is encoded as an n-bit string for some $n \in \mathbb{N}$), the obfuscator $i\mathcal{O}^{\mathcal{O}}$ samples $r \leftarrow \{0,1\}^n$ uniformly at random, computes $\widehat{C} = \mathcal{O}_n(C,r)$, and outputs the circuit $C'(\cdot) = \mathsf{Eval}(\widehat{C}^f,\cdot)$ (i.e., the obfuscated circuit C' consists of a single Eval-gate with hardwired input $\widehat{C}^f$). The definition of the function Eval guarantees that $i\mathcal{O}$ preserves the functionality of the obfuscated circuit C.

The vast majority of our effort is focused on showing that obfuscations of any two circuits that are functionally equivalent relative to f are computationally indistinguishable even for algorithms that can access the oracle Γ. Essentially, the technical challenge here is that the oracle $\mathsf{CollFinder}$ may perform an *exponential* number of queries to the oracles $\mathcal{O}$ and Eval, and thus most of the standard arguments that are typically used in black-box impossibility results are not applicable here. By generalizing the approach of Gennaro et al. [46] and its extensions by Haitner et al. [44] (for dealing with Simon's oracle $\mathsf{CollFinder}$) and by Brakerski et al. [43] (for dealing with indistinguishability experiments), we prove that it is hard to distinguish between the obfuscations of any two functionally-equivalent circuits.

We note that our approach heavily relies on the specific setting of *indistinguishability* obfuscation, and it is currently not clear whether or not it can be extended to stronger notions of indistinguishability such as *differing-input* obfuscation [64], [65], [66]. The following is a simplified variant of the theorem that we prove in the full version [1]:

Theorem III.3 (Simplified variant)**.** *For any probabilistic oracle-aided algorithm $\mathcal{A}$ that runs in time at most $2^{n/15}$, and for any permutation f, it holds that*

$$\left|\Pr_{\mathcal{O},\mathcal{R}}\left[\mathsf{Exp}^{\mathsf{iO}}_{\Gamma,i\mathcal{O},\mathcal{A},\mathcal{C}}(n)=1\right]-\frac{1}{2}\right| \le 2^{-n/40}.$$

Part 3: The existence of a one-way permutation. Finally, we prove that f is a one-way permutation relative to Γ. The following is a simplified variant of the theorem that we prove in the full version [1], and also here our inspiration is rooted at ideas originated in [46], [44], [43] that we generalize to our oracle Γ:

Theorem III.4 (Simplified variant)**.** *For any probabilistic oracle-aided algorithm $\mathcal{A}$ that runs in time at most $2^{n/50}$, and for any function $\mathcal{O}$, it holds that*

$$\Pr_{\substack{f,\mathcal{R}\\ x\leftarrow\{0,1\}^n}}\left[\mathcal{A}^{\Gamma}(f(x))=x\right] \le 2^{-n/50}.$$

Given the above discussion we are now ready to prove Theorem III.2.

Proof of Theorem III.2. Let $(\mathsf{Gen},\mathsf{Eval},M,T_M,\epsilon_{M,1},\epsilon_{M,2})$ be a fully black-box construction of a collision-resistant function family from a one-way permutation f and an indistinguishability obfuscator $i\mathcal{O}$ for the class $\mathcal{C}$ of all polynomial-size oracle-aided circuits C^f (recall Definition III.1). Note that in our setting, relative to the oracle Γ, this means we allow the algorithms Gen and Eval to access f, $\mathcal{O}$ and Eval (but we do not allow them to access $\mathsf{CollFinder}$).

Then, as discussed above, there exists an oracle-aided algorithm $\mathcal{A}$ that runs in polynomial time $T_{\mathcal{A}}(n)$, such that for any f and $\mathcal{O}$, it holds that

$$\Pr_{\mathcal{R}}\left[\begin{array}{c} x\neq x' \text{ and} \\ \mathsf{Eval}^{f,\mathcal{O},\mathsf{Eval}^{f,\mathcal{O}}}(\sigma,x)=\mathsf{Eval}^{f,\mathcal{O},\mathsf{Eval}^{f,\mathcal{O}}}(\sigma,x')\end{array}\middle|\begin{array}{c}\sigma\leftarrow\mathsf{Gen}^{f,\mathcal{O},\mathsf{Eval}^{f,\mathcal{O}}}(1^n)\\ (x,x')\leftarrow\mathcal{A}^{\Gamma}(1^n,\sigma)\end{array}\right]\ge\epsilon_{\mathcal{A}}(n), \tag{1}$$

[11] Note that w is uniformly distributed over $\{0,1\}^m$, and that w' is uniformly distributed over $\{0,1\}^m$ subject to forming a collision with w.

[12] $\mathsf{CollFinder}$ is deterministic, and therefore n independent responses for same input circuit can be obtained, for example, by appending to the description of the circuit a "dummy" counter (thus having n different circuits with the same functionality).

where $\epsilon_{\mathcal{A}}(n) = 1/4$ for all values of $n \in \mathbb{N}$ (as noted above, we can in fact use $\epsilon_{\mathcal{A}}(n) = 1 - (3/4)^n$). Definition III.1 then states that there are two possible cases to consider: $\mathcal{A}$ can be used either for breaking the indistinguishability obfuscator $i\mathcal{O}$, or for inverting the one-way permutation f.

In the first case, noting that Eq. (1) holds for any f and $\mathcal{O}$, we obtain from Definition III.1 that for any f it holds that

$$\left| \Pr_{\mathcal{O},\mathcal{R}} \left[\mathsf{Exp}^{\mathsf{iO}}_{\Gamma,i\mathcal{O},M^{\mathcal{A}},\mathcal{C}}(n) = 1 \right] - \frac{1}{2} \right| \geq \epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot 4\right) \cdot \epsilon_{M,2}(n),$$

where M runs in time $T_M(n)$. The algorithm M may invoke $\mathcal{A}$ on various security parameters (i.e., in general M is not restricted to invoking $\mathcal{A}$ only on security parameter n), and we denote by $\ell(n)$ the maximal security parameter on which M invokes $\mathcal{A}$ (when M itself is invoked on security parameter n). Thus, viewing $M^{\mathcal{A}}$ as a single algorithm, its running time $T_{M^{\mathcal{A}}}(n)$ satisfies $T_{M^{\mathcal{A}}}(n) \leq T_M(n) \cdot T_{\mathcal{A}}(\ell(n))$ (this follows since M may invoke $\mathcal{A}$ at most $T_M(n)$ times, and the running time of $\mathcal{A}$ on each such invocation is at most $T_{\mathcal{A}}(\ell(n))$). Theorem III.3 then implies that either $2^{n/15} \leq T_{M^{\mathcal{A}}}(n)$ or $\epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot 4\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/40}$. In the first sub-case, noting that $\ell(n) \leq T_M(n)$, we obtain that

$$2^{n/15} \leq T_{M^{\mathcal{A}}}(n) \leq T_M(n) \cdot T_{\mathcal{A}}(\ell(n)) \leq T_M(n) \cdot T_{\mathcal{A}}(T_M(n)).$$

The running time $T_{\mathcal{A}}(n)$ of the adversary $\mathcal{A}$ is some fixed polynomial in n, and therefore $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$. In the second sub-case, we have that $\epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot 4\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/40}$, and since $T_{\mathcal{A}}(n)$ is some fixed polynomial in n we obtain that $\epsilon_{M,1}\left(n^c\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/40}$ for some constant $c > 1$.

In the second case, noting again that Eq. (1) holds for any f and $\mathcal{O}$, we obtain from Definition III.1 that for any $\mathcal{O}$ it holds that

$$\Pr_{\substack{f,\mathcal{R} \\ x \leftarrow \{0,1\}^n}} \left[\left(M^{\mathcal{A}}\right)^{\Gamma}\left(f(x)\right) = x \right] \geq \epsilon_{M_1}\left(T_{\mathcal{A}}(n) \cdot 4\right) \cdot \epsilon_{M_2}(n),$$

where M runs in time $T_M(n)$. As in the first case, viewing $M^{\mathcal{A}}$ as a single algorithm, its running time $T_{M^{\mathcal{A}}}(n)$ satisfies $T_{M^{\mathcal{A}}}(n) \leq T_M(n) \cdot T_{\mathcal{A}}(\ell(n))$. Theorem III.4 then implies that either $2^{n/50} \leq T_{M^{\mathcal{A}}}(n)$ or $\epsilon_{M_1}\left(T_{\mathcal{A}}(n) \cdot 4\right) \cdot \epsilon_{M_2}(n) \leq 2^{-n/50}$. As in the first case, this implies that either $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$ or $\epsilon_{M,1}\left(n^c\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/50}$ for some constant $c > 1$. ∎

IV. Limits on the Power of Private-Key Functional Encryption

In this section we present our negative result for constructing a perfectly-complete key-agreement protocol from a general-purpose private-key functional encryption scheme and a one-way permutation. First, in Section IV-A we formally define the class of constructions to which our negative result applies. Then, in Section IV-B we present the structure of our proof, which is provided in the full version of this work [1]. Finally, in Section IV-C we show that our result can be extended for separating indistinguishability obfuscation for oracle-aided circuits from private-key functional encryption for oracle-aided circuits.

A. The Class of Reductions

We consider fully black-box constructions of a perfectly-complete bit-agreement protocol from a general-purpose private-key functional encryption scheme and a one-way permutation. Similarly to our approach from Section III, we model these primitives as two independent building blocks due to the following two reasons. First, although any private-key functional encryption scheme clearly implies the existence of a one-way function, it is not known whether any such scheme implies the existence of a one-way permutation. Second, this enables us to capture constructions that may use the underlying functional encryption scheme for generating functional keys to any circuit that can be constructed in a fully black-box manner from a one-way permutation. For example, as in Section III, this enables us to capture constructions that may generate functional keys to any circuit that uses a puncturable pseudorandom function or a pseudorandom generator as a sub-routine.

We now formally define the class of constructions considered in this section, tailoring our definitions to the specific primitives under consideration. We consider key-agreement protocols in which the parties agree on a single bit, and we refer to such protocols as bit-agreement protocols. We consider any implementation of a one-way permutation f and a private-key functional encryption scheme Π for the class of all polynomial-size oracle-aided circuits C^f.

Definition IV.1. *A fully black-box construction of a perfectly-correct bit-agreement protocol from a one-way permutation and a private-key functional encryption scheme for the class $\mathcal{C} = \{\mathcal{C}_n\}_{n\in\mathbb{N}}$ of all polynomial-size oracle-aided circuits consists of a pair of probabilistic polynomial-time oracle-aided algorithms $(\mathcal{A}, \mathcal{B})$, an oracle-aided algorithm M that runs in time $T_M(\cdot)$, and functions $\epsilon_{M,1}(\cdot)$ and $\epsilon_{M,2}(\cdot)$, such that the following two conditions hold:*

- **Correctness:** *For any $n \in \mathbb{N}$, for any permutation f and for any correct private-key functional encryption scheme Π, it holds that*
$$\Pr\left[k_{\mathcal{A}} = k_{\mathcal{B}} \left| (k_{\mathcal{A}}, k_{\mathcal{B}}, T) \leftarrow \left\langle \mathcal{A}^{f,\Pi}(1^n), \mathcal{B}^{f,\Pi}(1^n) \right\rangle \right.\right] = 1,$$
where the probability is taken over the internal randomness of $\mathcal{A}$ and $\mathcal{B}$.
- **Black-box proof of security:** *For any permutation f, for any correct private-key functional encryption scheme Π, for any probabilistic oracle-aided algorithm E that runs in time $T_E = T_E(n)$, and for any function $\epsilon_E = \epsilon_E(n)$, if*
$$\left| \Pr\left[\mathsf{Exp}^{\mathsf{KA}}_{(f,\Pi),(\mathcal{A},\mathcal{B}),E}(n) = 1 \right] - \frac{1}{2} \right| \geq \epsilon_E(n)$$
for infinitely many values of n (see Definition II.6 for the description of the experiment $\mathsf{Exp}^{\mathsf{KA}}_{(f,\Pi),(\mathcal{A},\mathcal{B}),E}$), then either
$$\Pr_{x \leftarrow \{0,1\}^n}\left[M^{f,\Pi,E}(f(x)) = x \right] \geq \epsilon_{M,1}\left(T_E(n) \cdot \epsilon_E^{-1}(n)\right) \cdot \epsilon_{M,2}(n)$$
or
$$\left| \Pr\left[\mathsf{Exp}^{\mathsf{FE}}_{(f,\Pi),\Pi,M^E,\mathcal{C}}(n) = 1 \right] - \frac{1}{2} \right| \geq \epsilon_{M_1}\left(T_E(n) \cdot \epsilon_E^{-1}(n)\right) \cdot \epsilon_{M,2}(n)$$
for infinitely many values of n (see Definition II.4 for the description of the experiment $\mathsf{Exp}^{\mathsf{FE}}_{(f,\Pi),\Pi,M^E,\mathcal{C}}$).

Similarly to the discussion in Section III-A, we emphasize the fact that our definition captures an underlying functional encryption scheme Π for the class of all polynomial-size oracle-aided circuits C^f, and thus *black-box* access in our setting is only required with respect to the functional encryption scheme Π and the one-way permutation f themselves. We refer the reader to Section III-A for a more elaborated discussion of this issue, as well as a discussion on the roles of the "security loss" functions T_M, $\epsilon_{M,1}$ and $\epsilon_{M,2}$.

B. Proof Overview and the Oracle Ψ

Our result is obtained by presenting an oracle Ψ relative to which there exist an exponentially-secure one-way permutation f and an exponentially-secure private-key functional encryption scheme Π for the class of all polynomial-size oracle-aided circuits C^f, but any bit-agreement protocol with perfect completeness can be broken. We prove the following theorem:

Theorem IV.2. *Let $(\mathcal{A}, \mathcal{B}, M, T_M, \epsilon_{M,1}, \epsilon_{M,2})$ be a fully black-box construction of a perfectly-correct bit-agreement protocol from a one-way permutation f and a private-key functional encryption scheme Π for the class of all polynomial-size oracle-aided circuits C^f (see Definition IV.1). Then, at least one of the following propertied holds:*

1) *$T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$ (i.e., the reduction runs in exponential time).*
2) *$\epsilon_{M,1}(n^c) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$ (i.e., the security loss is exponential).*

In particular, the theorem implies that if the running time $T_M(\cdot)$ of the reduction is sub-exponential and the adversary-dependent security loss $\epsilon_{M,1}(\cdot)$ is polynomial as in the vast majority of constructions, then the adversary-independent security loss $\epsilon_{M,2}(\cdot)$ must be exponential (thus ruling out even constructions that rely on sub-exponential security assumptions).

In what follows we describe the oracle Ψ and then explain the structure of our proof.

The oracle Ψ. The oracle Ψ is a quadruple $\left(f, \mathcal{K}, \mathcal{E}, \mathcal{D}^{f,\mathcal{K},\mathcal{E}}\right)$ that is defined as follows:

- **The function $f = \{f_n\}_{n\in\mathbb{N}}$.** For every $n \in \mathbb{N}$ the function f_n is a uniformly chosen permutation over $\{0,1\}^n$. Looking ahead, we will prove that f is one way relative to Ψ.
- **The functions $\mathcal{K} = \{\mathcal{K}_n\}_{n\in\mathbb{N}}$ and $\mathcal{E} = \{\mathcal{E}_n\}_{n\in\mathbb{N}}$.** For every $n \in \mathbb{N}$ the functions $\mathcal{K}_n$ and $\mathcal{E}_n$ are uniformly chosen functions $\mathcal{K}_n : \{0,1\}^{2n} \to \{0,1\}^{10n}$ and $\mathcal{E}_n : \{0,1\}^{3n} \to \{0,1\}^{10n}$. Looking ahead, we will use $\mathcal{K}$ and $\mathcal{E}$ for implementing the (deterministic) key-generation algorithm and the (randomized) encryption algorithm, respectively, of the functional encryption scheme Π.
- **The function $\mathcal{D}^{f,\mathcal{K},\mathcal{E}} = \{\mathcal{D}_n^{f_n,\mathcal{K}_n,\mathcal{E}_n}\}_{n\in\mathbb{N}}$.** For every $n \in \mathbb{N}$ the function $\mathcal{D}_n^{f,\mathcal{K},\mathcal{E}} : \{0,1\}^{20n} \to \{0,1\}^n$ parses its input as a pair $(\mathsf{sk}, c) \in \{0,1\}^{10n} \times \{0,1\}^{10n}$, and is defined as follows: If there exist $\mathsf{msk}, C, m, r \in \{0,1\}^n$ such that $\mathsf{sk} = \mathcal{K}_n(\mathsf{msk}, C)$ and $c = \mathcal{E}_n(\mathsf{msk}, m, r)$ then it outputs $C^f(m)$ for the lexicographically-first such quadruple, and otherwise it outputs $\perp$. Looking ahead, we will use $\mathcal{D}$ for implementing the decryption algorithm of the functional encryption scheme Π.

Equipped with the oracle Ψ, our proof consists of the following two parts.

Part 1: The existence of a one-way permutation and a functional encryption scheme. We first prove that f is one way relative to Ψ, and this is rather standard at least when observing that each query to $\mathcal{D}$ requires only a bounded (and not too large) number of queries to f. Then, we show that relative to Ψ there exists a functional encryption scheme Π that is naturally defined given the way we set up Ψ (as discussed in Section I-C). Moreover, the functional encryption scheme Π is even "compact" [14], [29] in the sense that the efficiency of its encryption algorithm depends only on the security parameter and on length of the plaintext (in particular, it is independent of the complexity of the function family supported by the scheme). Proving that Π is secure is more tricky, since each query to $\mathcal{D}$ may require an exponential number of queries to $\mathcal{K}$ and $\mathcal{E}$. The following are simplified variants of the theorems that we prove in the full version of this work [1]:

Theorem IV.3 (Simplified variant)**.** *For any oracle-aided algorithm $\mathcal{A}$ that makes at most $2^{n/4}$ oracle queries, and for any functions $\mathcal{K}$ and $\mathcal{E}$, it holds that*

$$\Pr_{\substack{f \\ x \leftarrow \{0,1\}^n}} \left[\mathcal{A}^{\Psi}(f(x)) = x\right] \leq 2^{-n/4}.$$

Theorem IV.4 (Simplified variant)**.** *For any oracle-aided valid adversary $\mathcal{A}$ that makes at most $2^{n/4}$ oracle queries, and for any permutation f, it holds that*

$$\left|\Pr_{\mathcal{K},\mathcal{E}}\left[\mathsf{Exp}^{\mathsf{FE}}_{\Psi,\Pi,\mathcal{A},\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \leq 2^{-n/4}.$$

Part 2: Breaking any perfectly-complete bit-agreement protocol. The most challenging part in this section is in proving that any perfectly-complete bit-agreement protocol can be broken using a polynomial number of oracle queries. Our proof is inspired by a combination of ideas that were developed in the early work of Impagliazzo and Rudich [38] and its improvement by Barak and Mahmoody-Ghidary [48]. Specifically, for the case of perfectly-complete bit-agreement protocols, our proof generalizes the approach of Brakerski, Katz, Yerukhimovich and Segev [43] to the setting of our oracle Ψ. The following is a simplified variant of the theorem that we prove in the full version of this work [1]:

Theorem IV.5 (Simplified variant)**.** *For any polynomial-time oracle-aided perfectly-complete bit-agreement protocol $(\mathcal{A},\mathcal{B})$, there exists an oracle-aided algorithm E making a polynomial number of oracle queries such that*

$$\left|\Pr_{\Psi}\left[\mathsf{Exp}^{\mathsf{KA}}_{\Psi,(\mathcal{A},\mathcal{B}),E}(n) = 1\right] - \frac{1}{2}\right| \geq 1/4.$$

Finally, we now show that the above two parts imply Theorem IV.2 via a proof that is essentially identical to that of Theorem III.2.

Proof of Theorem IV.2. Let $(\mathcal{A},\mathcal{B},M,T_M,\epsilon_{M,1},\epsilon_{M,2})$ be a fully black-box construction of a perfectly-complete bit-agreement protocol from a one-way permutation f and a private-key functional encryption scheme Π for the class $\mathcal{C}$ of all polynomial-size oracle-aided circuits C^f (recall Definition IV.1). Note that in our setting, relative to the oracle Ψ, this means we allow the algorithms $\mathcal{A}$ and $\mathcal{B}$ to access f, $\mathcal{K}$, $\mathcal{E}$ and $\mathcal{D}$ (i.e., to access Ψ).

Theorem IV.5 guarantees the existence of an oracle-aided algorithm E that makes a polynomial number $T_E(n)$ of queries to the oracle Ψ, such that

$$\left|\Pr_{\Psi}\left[\mathsf{Exp}^{\mathsf{KA}}_{\Psi,(\mathcal{A},\mathcal{B}),E}(n) = 1\right] - \frac{1}{2}\right| \geq \epsilon_E(n),$$

where $\epsilon_E(n) = 1/4$ for all values of $n \in \mathbb{N}$. Definition IV.1 then states that there are two possible cases to consider: E can be used either for breaking the functional encryption scheme Π, or for inverting the one-way permutation f.

In the first case, we obtain from Definition IV.1 that

$$\left|\Pr_{\Psi}\left[\mathsf{Exp}^{\mathsf{FE}}_{\Psi,\Pi,M^E,\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \geq \epsilon_{M,1}\left(T_E(n)\cdot 4\right)\cdot\epsilon_{M,2}(n),$$

where M runs in time $T_M(n)$. The algorithm M may invoke E on various security parameters (i.e., in general M is not restricted to invoking E only on security parameter n), and we denote by $\ell(n)$ the maximal security parameter on which M invokes E (when M itself is invoked on security parameter n). Thus, viewing M^E as a single algorithm, its number of queries $T_{M^E}(n)$ to the oracle Ψ satisfies $T_{M^E}(n) \leq T_M(n)\cdot T_E(\ell(n))$ (this follows since M runs in time $T_M(n)$ and in each step of its execution it may query Ψ directly at most once or invoke E at most once where each such invocation results in at most $T_E(\ell(n))$ queries to Ψ). Theorem IV.4 then implies that either $2^{n/4} \leq T_{M^E}(n)$ or $\epsilon_{M,1}\left(T_E(n)\cdot 4\right)\cdot\epsilon_{M,2}(n) \leq 2^{-n/4}$. In the first sub-case, noting that $\ell(n) \leq T_M(n)$, we obtain that

$$2^{n/4} \leq T_{M^E}(n) \leq T_M(n)\cdot T_E(\ell(n)) \leq T_M(n)\cdot T_E(T_M(n)).$$

The number of queries $T_E(n)$ made by the adversary E is some fixed polynomial in n, and therefore $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$. In the second sub-case, we have that $\epsilon_{M,1}\left(T_E(n) \cdot 4\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$, and since $T_E(n)$ is some fixed polynomial in n we obtain that $\epsilon_{M,1}\left(n^c\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$.

In the second case, we obtain from Definition IV.1 that

$$\Pr_{\substack{\Psi \\ x \leftarrow \{0,1\}^n}}\left[\left(M^E\right)^{\Psi}(f(x)) = x\right] \geq \epsilon_{M,1}\left(T_E(n) \cdot 4\right) \cdot \epsilon_{M,2}(n),$$

where M runs in time $T_M(n)$. As in the first case, viewing M^E as a single algorithm, its number of queries $T_{M^E}(n)$ to the oracle Ψ satisfies $T_{M^E}(n) \leq T_M(n) \cdot T_E(\ell(n))$. Theorem IV.3 then implies that either $2^{n/4} \leq T_{M^E}(n)$ or $\epsilon_{M,1}\left(T_E(n) \cdot 4\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$. As in the first case, this implies that either $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$ or $\epsilon_{M,1}\left(n^c\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$. ∎

C. Extending the Result to Indistinguishability Obfuscation

In this section we show that Theorem IV.2 can be extended for separating indistinguishability obfuscation for oracle-aided circuits from private-key functional encryption for oracle-aided circuits. As discussed in Section I-A, this does not necessarily imply a separation in the standard model, since it may be that there exists an indistinguishability obfuscator for all polynomial-size circuits, but there does not exist such an obfuscator for polynomial-size *oracle-aided* circuits. Nevertheless, this provides substantial evidence that private-key functional encryption is somewhat unlikely to imply indistinguishability obfuscation using standard techniques.

We now formally define the class of constructions considered in this section, tailoring our definitions to the specific primitives under consideration, and then formally state our result. We consider any implementation of a one-way permutation f and a private-key functional encryption scheme Π for the class of all polynomial-size oracle-aided circuits C^f. As in Section III, we model these primitives as two independent building blocks.

Definition IV.6. *A fully black-box construction of an indistinguishability obfuscator for the class $\mathcal{C} = \{\mathcal{C}_n\}_{n \in \mathbb{N}}$ of all polynomial-size oracle-aided circuits from a one-way permutation and a private-key functional encryption scheme for the class $\mathcal{C}$ consists of a probabilistic polynomial-time oracle-aided algorithm $i\mathcal{O}$, an oracle-aided algorithm M that runs in time $T_M(\cdot)$, and functions $\epsilon_{M,1}(\cdot)$ and $\epsilon_{M,2}(\cdot)$, such that the following two conditions hold:*

- **Correctness:** *For any $n \in \mathbb{N}$, for any permutation f, for any correct private-key functional encryption scheme Π, and for any oracle-aided circuit $C \in \mathcal{C}_n$ it holds that*

$$\Pr\left[C^f \equiv \widehat{C}^f \;:\; \widehat{C} \leftarrow i\mathcal{O}^{f,\Pi}(1^n, C)\right] = 1.$$

- **Black-box proof of security:** *For any permutation f, for any correct private-key functional encryption scheme Π, for any (not necessarily uniform) probabilistic oracle-aided algorithm D that runs in time $T_D = T_D(n)$, and for any function $\epsilon_D = \epsilon_D(n)$, if*

$$\left|\Pr\left[\mathsf{Exp}^{\mathsf{iO}}_{(f,\Pi),i\mathcal{O},D,\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \geq \epsilon_D(n)$$

for infinitely many values of n (see Definition II.2 for the description of the experiment $\mathsf{Exp}^{\mathsf{iO}}_{(f,\Pi),i\mathcal{O},D,\mathcal{C}}$), then either

$$\Pr_{x \leftarrow \{0,1\}^n}\left[M^{f,\Pi,D}(f(x)) = x\right] \geq \epsilon_{M,1}\left(T_D(n) \cdot \epsilon_D^{-1}(n)\right) \cdot \epsilon_{M,2}(n)$$

or

$$\left|\Pr\left[\mathsf{Exp}^{\mathsf{FE}}_{(f,\Pi),\Pi,M^D,\mathcal{C}}(n) = 1\right] - \frac{1}{2}\right| \geq \epsilon_{M,1}\left(T_D(n) \cdot \epsilon_D^{-1}(n)\right) \cdot \epsilon_{M,2}(n)$$

for infinitely many values of n (see Definition II.4 for the description of the experiment $\mathsf{Exp}^{\mathsf{FE}}_{(f,\Pi),\Pi,M^E,\mathcal{C}}$).

Theorem IV.7. *Let $(i\mathcal{O}, M, T_M, \epsilon_{M,1}, \epsilon_{M,2})$ be a fully black-box construction of an indistinguishability obfuscator for the class $\mathcal{C}$ of all polynomial-size oracle-aided circuits C^f from a one-way permutation f and a private-key functional encryption scheme Π for the class $\mathcal{C}$ (see Definition IV.6). Then, at least one of the following propertied holds:*

1) *$T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$ (i.e., the reduction runs in exponential time).*
2) *$\epsilon_{M,1}(n^c) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$ (i.e., the security loss is exponential).*

In particular, the theorem implies that if the running time $T_M(\cdot)$ of the reduction is sub-exponential and the adversary-dependent security loss $\epsilon_{M,1}(\cdot)$ is polynomial as in the vast majority of constructions, then the adversary-independent security loss $\epsilon_{M,2}(\cdot)$ must be exponential (thus ruling out even constructions that rely on sub-exponential security assumptions).

Theorem IV.7 follows from Theorem IV.2 and from the construction of a (perfectly-correct) public-key encryption scheme from a one-way function and a general-purpose indistinguishability obfuscator by Sahai and Waters [7]. Their construction, however, relies on the underlying one-way function in a non-black-box manner, and therefore Theorem IV.7 does not immediately follow from Theorem IV.2. However, although their construction relies the underlying one-way function in a non-black-box manner, relative to the oracle Ψ it is in fact a fully black-box construction (with a polynomial security loss). Specifically, Sahai and Waters use the underlying indistinguishability obfuscator for obfuscating a circuit that invokes a puncturable pseudorandom function and a pseudorandom generator as sub-routines. Given that puncturable pseudorandom functions and pseudorandom generators can be based on any one-way function in a fully black-box manner, in our setting such a circuit is a polynomial-size oracle-aided circuit C^f (and we note that black-box proofs of security with a polynomial security loss clearly compose nicely). Equipped with this view we now prove Theorem IV.7.

Proof of Theorem IV.7. Let $(i\mathcal{O}, M, T_M, \epsilon_{M,1}, \epsilon_{M,2})$ be a fully black-box construction of a general-purpose indistinguishability obfuscator for the class $\mathcal{C}$ of all polynomial-size oracle-aided circuits C^f from a one-way permutation f and a private-key functional encryption scheme Π for the class $\mathcal{C}$ (recall Definition IV.6). Note that in our setting, relative to the oracle Ψ, this means we allow the algorithm $i\mathcal{O}$ to access f, $\mathcal{K}$, $\mathcal{E}$ and $\mathcal{D}$ (i.e., to access Ψ). We now construct a perfectly-complete bit-agreement protocol by relying on the following two building blocks:

- A length-doubling pseudorandom generator that is constructed from the permutation f in a fully black-box manner [58] with a polynomial security loss (it in fact suffices for f to be a one-way function). This means that: (1) the pseudorandom generator is of the form G^f where $G \in \mathcal{C}$ is a polynomial-size oracle-aided circuit, and (2) any oracle-aided distinguisher that runs in time $T = T(n)$ and has an advantage $\epsilon = \epsilon(n)$ in breaking the pseudorandom generator, can be used in a black-box manner for inverting f in time that is polynomially related to T and with probability that is polynomially related to ϵ.
- A puncturable pseudorandom function that is constructed from the permutation f in a fully black-box manner [59], [60], [7], [61] with a polynomial security loss (it again suffices for f to be a one-way function). This means that: (1) the evaluation and puncturing algorithms of the family are of the form $\mathsf{PRF.Eval}^f$ and $\mathsf{PRF.Punc}^f$ where $\mathsf{PRF.Eval}, \mathsf{PRF.Punc} \in \mathcal{C}$ are polynomial-size oracle-aided circuits[13], and (2) any oracle-aided distinguisher that runs in time $T = T(n)$ and has an advantage $\epsilon = \epsilon(n)$ in breaking the puncturable pseudorandom function, can be used in a black-box manner for inverting f in time that is polynomially related to T and with probability that is polynomially related to ϵ.

Consider now the following perfectly-complete bit-agreement protocol $(\mathcal{A}, \mathcal{B})$ relative to the oracle Ψ:

- The algorithm $\mathcal{A}$, on input 1^n, samples $\mathsf{k} \leftarrow \{0,1\}^n$, computes $\widehat{C_{\mathsf{k}}} \leftarrow i\mathcal{O}^{\Psi}(C_{\mathsf{k}})$ and sends $\widehat{C_{\mathsf{k}}}$ to $\mathcal{B}$, where $C_{\mathsf{k}} \in \mathcal{C}$ is a polynomial-size oracle-aided circuit that is defined as follows: On input a value $r \in \{0,1\}^n$ and oracle access to f, it outputs $(G^f(r), \mathsf{PRF.Eval}^f(\mathsf{k}, G^f(r)))$.
- The algorithm $\mathcal{B}$, on input 1^n and $\widehat{C_{\mathsf{k}}}$, first samples $b \leftarrow \{0,1\}$ and $r \leftarrow \{0,1\}^n$. Then, it sends to $\mathcal{A}$ the value $(c_1, c_2 \oplus b)$ where $(c_1, c_2) = \widehat{C_{\mathsf{k}}}^f(r)$, and outputs $k_{\mathcal{B}} = b$.
- The algorithm $\mathcal{A}$, on input (c_1, c_2'), outputs $k_{\mathcal{A}} = c_2' \oplus \mathsf{PRF.Eval}^f(\mathsf{k}, c_1)$.

The protocol is directly based on the public-key encryption scheme of Sahai and Waters [7] by having $\mathcal{A}$ send $\mathcal{B}$ a public key, and then $\mathcal{B}$ samples a uniform bit b (which will serve as their output) and replies with its encryption. The protocol is clearly perfectly complete (based on the fact that $i\mathcal{O}^{\Psi}$ preserves functionality) and results is a uniformly-distributed key.

Theorem IV.5 guarantees the existence of an oracle-aided algorithm E that makes a polynomial number $T_E(n)$ of queries to the oracle Ψ, such that

$$\left| \Pr_{\Psi} \left[\mathsf{Exp}^{\mathsf{KA}}_{\Psi,(\mathcal{A},\mathcal{B}),E}(n) = 1 \right] - \frac{1}{2} \right| \geq \epsilon_E(n), \tag{2}$$

where $\epsilon_E(n) = 1/4$ for all values of $n \in \mathbb{N}$. Given that the protocol $(\mathcal{A}, \mathcal{B})$ directly corresponds to the public-key encryption scheme of Sahai and Waters, their proof of security states that there are two possible cases to consider:

Case 1. E can be used in a black-box manner by an oracle-aided algorithm E' for inverting the one-way permutation f. Since their construction guarantees a polynomial security loss, the algorithm E' makes $T_{E'}(n) = \mathsf{poly}(T_E(n)) = \mathsf{poly}(n)$ oracle queries and succeeds with probability $\epsilon_{E'}(n) = 1/\mathsf{poly}(T_E(n) \cdot 4) = 1/\mathsf{poly}(n)$. This range of parameters for $T_{E'}(n)$ and $\epsilon_{E'}(n)$ contradicts Theorem IV.3, stating that any algorithm that makes at most $2^{n/4}$ queries to Ψ inverts f with probability at most $2^{-n/4}$, and therefore this case is not possible.

Case 2. E can be used in a black-box manner by an oracle-aided algorithm E' for breaking the indistinguishability obfuscator $i\mathcal{O}$. As in the previous case, since their construction guarantees a polynomial security loss, the algorithm E' makes $T_{E'}(n) =$

[13]For simplicity we assume that the key-generation algorithm of the pseudorandom family on input 1^n outputs a uniform n-bit key k.

$\mathsf{poly}(T_E(n)) = \mathsf{poly}(n)$ oracle queries and succeeds with probability $\epsilon_{E'}(n) = 1/\mathsf{poly}(T_E(n) \cdot 4) = 1/\mathsf{poly}(n)$. That is, in this case we have an oracle-aided algorithm E' making $T_{E'}(n) = \mathsf{poly}(n)$ oracle queries such that

$$\left| \Pr_{\Psi} \left[\mathsf{Exp}^{\mathsf{iO}}_{\Psi, i\mathcal{O}, E', \mathcal{C}}(n) = 1 \right] - \frac{1}{2} \right| \geq \epsilon_{E'}(n)$$

for infinitely many values of n, where $\epsilon_{E'}(n) = 1/\mathsf{poly}(n)$.

The analysis of this case proceeds as in the proof of Theorem IV.2. Specifically, Definition IV.6 states that there are two possible sub-cases to consider: E' can be used in a black-box manner by M either for breaking the functional encryption scheme Π, or for inverting the one-way permutation f. In the first sub-case, we obtain from Definition IV.6 that

$$\left| \Pr_{\Psi} \left[\mathsf{Exp}^{\mathsf{FE}}_{\Psi, \Pi, M^{E'}, \mathcal{C}}(n) = 1 \right] - \frac{1}{2} \right| \geq \epsilon_{M,1}\left(T_{E'}(n) \cdot \epsilon^{-1}_{E'}(n)\right) \cdot \epsilon_{M,2}(n),$$

where M runs in time $T_M(n)$. The algorithm M may invoke E' on various security parameters (i.e., in general M is not restricted to invoking E' only on security parameter n), and we denote by $\ell(n)$ the maximal security parameter on which M invokes E' (when M itself is invoked on security parameter n). Thus, viewing $M^{E'}$ as a single algorithm, its number of queries $T_{M^{E'}}(n)$ to the oracle Ψ satisfies $T_{M^{E'}}(n) \leq T_M(n) \cdot T_{E'}(\ell(n))$ (this follows since M runs in time $T_M(n)$ and in each step of its execution it may query Ψ directly at most once or invoke E' at most once where each such invocation results in at most $T_{E'}(\ell(n))$ queries to Ψ). Theorem IV.4 then implies that either $2^{n/4} \leq T_{M^{E'}}(n)$ or $\epsilon_{M,1}\left(T_{E'}(n) \cdot \epsilon^{-1}_{E'}(n)\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$. Now, since $T_{E'}(n)$ and $\epsilon^{-1}_{E'}(n)$ are some fixed polynomials in n, and since $\ell(n) \leq T_M(n)$, as in the proof of Theorem IV.2 we obtain that either $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$ or $\epsilon_{M,1}\left(n^c\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$.

In the second sub-case, we obtain from Definition IV.6 that

$$\Pr_{\substack{\Psi \\ x \leftarrow \{0,1\}^n}} \left[\left(M^{E'}\right)^{\Psi}(f(x)) = x \right] \geq \epsilon_{M,1}\left(T_{E'}(n) \cdot \epsilon^{-1}_{E'}(n)\right) \cdot \epsilon_{M,2}(n),$$

where M runs in time $T_M(n)$. As in the first sub-case, viewing $M^{E'}$ as a single algorithm, its number of queries $T_{M^{E'}}$ to the oracle Ψ satisfies $T_{M^{E'}}(n) \leq T_M(n) \cdot T_{E'}(\ell(n))$. Theorem IV.3 then implies that either $2^{n/4} \leq T_{M^{E'}}(n)$ or $\epsilon_{M,1}\left(T_{E'}(n) \cdot \epsilon^{-1}_{E'}(n)\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$. Now, since $T_{E'}(n)$ and $\epsilon^{-1}_{E'}(n)$ are some fixed polynomials in n, and since $\ell(n) \leq T_M(n)$, then again as in the proof of Theorem IV.2 we obtain that either $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$ or $\epsilon_{M,1}\left(n^c\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$. ∎

ACKNOWLEDGMENT

We thank Benny Applebaum, Nir Bitansky, Zvika Brakerski, Ran Canetti, Ilan Komargodski and Alon Rosen for very insightful discussions in various stages of this work.

REFERENCES

[1] G. Asharov and G. Segev, "Limits on the power of indistinguishability obfuscation and functional encryption," Cryptology ePrint Archive, Report 2015/341, 2015.

[2] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang, "On the (im)possibility of obfuscating programs," in *Advances in Cryptology – CRYPTO '01*, 2001, pp. 1–18.

[3] ——, "On the (im)possibility of obfuscating programs," *Journal of the ACM*, vol. 59, no. 2, p. 6, 2012.

[4] R. Canetti, "Towards realizing random oracles: Hash functions that hide all partial information," in *Advances in Cryptology – CRYPTO '97*, 1997, pp. 455–469.

[5] H. Wee, "On obfuscating point functions," in *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, 2005, pp. 523–532.

[6] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters, "Candidate indistinguishability obfuscation and functional encryption for all circuits," in *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, 2013, pp. 40–49.

[7] A. Sahai and B. Waters, "How to use indistinguishability obfuscation: Deniable encryption, and more," in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 2014, pp. 475–484.

[8] R. Canetti, C. Dwork, M. Naor, and R. Ostrovsky, "Deniable encryption," in *Advances in Cryptology – CRYPTO '97*, 1997, pp. 90–104.

[9] I. Komargodski, T. Moran, M. Naor, R. Pass, A. Rosen, and E. Yogev, "One-way functions and (im)perfect obfuscation," in *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science*, 2014, pp. 374–383.

[10] N. Bitansky, O. Paneth, and D. Wichs, "Perfect structure on the edge of chaos," Cryptology ePrint Archive, Report 2015/126, 2015.

[11] S. Hohenberger, A. Sahai, and B. Waters, "Replacing a random oracle: Full domain hash from indistinguishability obfuscation," in *Advances in Cryptology – EUROCRYPT '14*, 2014, pp. 201–220.

[12] D. Boneh and M. Zhandry, "Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation," in *Advances in Cryptology – CRYPTO '14*, 2014, pp. 480–499.

[13] S. Goldwasser, S. D. Gordon, V. Goyal, A. Jain, J. Katz, F.-H. Liu, A. Sahai, E. Shi, and H.-S. Zhou, "Multi-input functional encryption," in *Advances in Cryptology – EUROCRYPT '14*, 2014, pp. 578–602.

[14] P. Ananth and A. Jain, "Indistinguishability obfuscation from compact functional encryption," To appear in *Advances in Cryptology – CRYPTO '15* (available at https://eprint.iacr.org/2015/173.pdf), 2015.

[15] V. Goyal, A. Jain, V. Koppula, and A. Sahai, "Functional encryption for randomized functionalities," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 325–351.

[16] S. Garg, C. Gentry, S. Halevi, and M. Raykova, "Two-round secure MPC from indistinguishability obfuscation," in *Proceedings of the 11th Theory of Cryptography Conference*, 2014, pp. 74–94.

[17] R. Canetti, S. Goldwasser, and O. Poburinnaya, "Adaptively secure two-party computation from indistinguishability obfuscation," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 557–585.

[18] D. Dachman-Soled, J. Katz, and V. Rao, "Adaptively secure, universally composable, multiparty computation in constant rounds," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 586–613.

[19] S. Garg and A. Polychroniadou, "Two-round adaptively secure MPC from indistinguishability obfuscation," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 614–637.

[20] P. Hubacek and D. Wichs, "On the communication complexity of secure function evaluation with long output," in *Proceedings of the 6th Conference on Innovations in Theoretical Computer Science*, 2015, pp. 163–172.

[21] B. Waters, "A punctured programming approach to adaptively secure functional encryption," To appear in *Advances in Cryptology – CRYPTO '15* (available at https://eprint.iacr.org/2014/588.pdf), 2015.

[22] M. Bellare, I. Stepanovs, and S. Tessaro, "Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation," in *Advances in Cryptology – ASIACRYPT '14*, 2014, pp. 102–121.

[23] N. Bitansky and O. Paneth, "ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 401–427.

[24] K. Chung, H. Lin, and R. Pass, "Constant-round concurrent zero-knowledge from indistinguishability obfuscation," Cryptology ePrint Archive, Report 2014/991, 2014.

[25] R. Canetti, H. Lin, S. Tessaro, and V. Vaikuntanathan, "Obfuscation of probabilistic circuits and applications," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 468–497.

[26] N. Bitansky, O. Paneth, and A. Rosen, "On the cryptographic hardness of finding a nash equilibrium," To appear in *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science* (available at https://eprint.iacr.org/2014/1029.pdf), 2015.

[27] A. Sahai and B. Waters, "Slides on functional encryption," Available at http://www.cs.utexas.edu/~bwaters/presentations/files/functional.ppt, 2008.

[28] D. Boneh, A. Sahai, and B. Waters, "Functional encryption: Definitions and challenges," in *Proceedings of the 8th Theory of Cryptography Conference*, 2011, pp. 253–273.

[29] N. Bitansky and V. Vaikuntanathan, "Indistinguishability obfuscation from functional encryption," To appear in *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science* (available at https://eprint.iacr.org/2014/163.pdf), 2015.

[30] S. Agrawal, S. Agrawal, S. Badrinarayanan, A. Kumarasubramanian, M. Prabhakaran, and A. Sahai, "Function private functional encryption and property preserving encryption: New definitions and positive results," Cryptology ePrint Archive, Report 2013/744, 2013.

[31] P. Ananth, Z. Brakerski, G. Segev, and V. Vaikuntanathan, "From selective to adaptive security in functional encryption," To appear in *Advances in Cryptology – CRYPTO '15* (available at https://eprint.iacr.org/2014/917.pdf), 2015.

[32] Z. Brakerski and G. Segev, "Function-private functional encryption in the private-key setting," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 306–324.

[33] Z. Brakerski, I. Komargodski, and G. Segev, "From single-input to multi-input functional encryption in the private-key setting," Cryptology ePrint Archive, Report 2015/158, 2015.

[34] I. Komargodski, G. Segev, and E. Yogev, "Functional encryption for randomized functionalities in the private-key setting from minimal assumptions," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 352–377.

[35] S. Garg, C. Gentry, S. Halevi, and M. Zhandry, "Fully secure functional encryption without obfuscation," Cryptology ePrint Archive, Report 2014/666, 2014.

[36] R. Rothblum, "Homomorphic encryption: From private-key to public-key," in *Proceedings of the 8th Theory of Cryptography Conference*, 2011, pp. 219–234.

[37] N. Bitansky, R. Canetti, H. Cohn, S. Goldwasser, Y. Tauman Kalai, O. Paneth, and A. Rosen, "The impossibility of obfuscation with auxiliary input or a universal simulator," in *Advances in Cryptology – CRYPTO '14*, 2014, pp. 71–89.

[38] R. Impagliazzo and S. Rudich, "Limits on the provable consequences of one-way permutations," in *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, 1989, pp. 44–61.

[39] Y. Ishai, E. Kushilevitz, and R. Ostrovsky, "Sufficient conditions for collision-resistant hashing," in *Proceedings of the 2nd Theory of Cryptography Conference*, 2005, pp. 445–456.

[40] S. Goldwasser and G. N. Rothblum, "On best-possible obfuscation," *Journal of Cryptology*, vol. 27, no. 3, pp. 480–505, 2014.

[41] B. Lynn, M. Prabhakaran, and A. Sahai, "Positive results and techniques for obfuscation," in *Advances in Cryptology – EUROCRYPT '04*, 2004, pp. 20–39.

[42] R. Canetti, Y. Tauman Kalai, and O. Paneth, "On obfuscation with random oracles," in *Proceedings of the 12th Theory of Cryptography Conference*, 2015, pp. 456–467.

[43] Z. Brakerski, J. Katz, G. Segev, and A. Yerukhimovich, "Limits on the power of zero-knowledge proofs in cryptographic constructions," in *Proceedings of the 8th Theory of Cryptography Conference*, 2011, pp. 559–578.

[44] I. Haitner, J. J. Hoch, O. Reingold, and G. Segev, "Finding collisions in interactive protocols – Tight lower bounds on the round and communication complexities of statistically hiding commitments," *SIAM Journal on Computing*, vol. 44, no. 1, pp. 193–242, 2015.

[45] D. R. Simon, "Finding collisions on a one-way street: Can secure hash functions be based on general assumptions?" in *Advances in Cryptology – EUROCRYPT '98*, 1998, pp. 334–345.

[46] R. Gennaro, Y. Gertner, J. Katz, and L. Trevisan, "Bounds on the efficiency of generic cryptographic constructions," *SIAM Journal on Computing*, vol. 35, no. 1, pp. 217–246, 2005.

[47] H. Wee, "One-way permutations, interactive hashing and statistically hiding commitments," in *Proceedings of the 4th Theory of Cryptography Conference*, 2007, pp. 419–433.

[48] B. Barak and M. Mahmoody-Ghidary, "Merkle puzzles are optimal - An $O(n^2)$-query attack on any key exchange from a random oracle," in *Advances in Cryptology – CRYPTO '09*, 2009, pp. 374–390.

[49] O. Reingold, L. Trevisan, and S. P. Vadhan, "Notions of reducibility between cryptographic primitives," in *Proceedings of the 1st Theory of Cryptography Conference*, 2004, pp. 1–20.

[50] B. Barak and M. Mahmoody-Ghidary, "Lower bounds on signatures from symmetric primitives," in *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, 2007, pp. 680–688.

[51] I. Haitner, J. J. Hoch, and G. Segev, "A linear lower bound on the communication complexity of single-server private information retrieval," in *Proceedings of the 5th Theory of Cryptography Conference*, 2008, pp. 445–464.

[52] D. Dachman-Soled, Y. Lindell, M. Mahmoody, and T. Malkin, "On the black-box complexity of optimally-fair coin tossing," in *Proceedings of the 8th Theory of Cryptography Conference*, 2011, pp. 450–467.

[53] M. Mahmoody and R. Pass, "The curious case of non-interactive commitments – On the power of black-box vs. non-black-box use of primitives," in *Advances in Cryptology – CRYPTO '12*, 2012, pp. 701–718.

[54] K. Chung, H. Lin, M. Mahmoody, and R. Pass, "On the power of nonuniformity in proofs of security," in *Proceedings of the 4th Innovations in Theoretical Computer Science Conference*, 2013, pp. 389–400.

[55] D. Dachman-Soled, M. Mahmoody, and T. Malkin, "Can optimally-fair coin tossing be based on one-way functions?" in *Proceedings of the 11th Theory of Cryptography Conference*, 2014, pp. 217–239.

[56] M. Mahmoody, H. K. Maji, and M. Prabhakaran, "On the power of public-key encryption in secure computation," in *Proceedings of the 11th Theory of Cryptography Conference*, 2014, pp. 240–264.

[57] O. Goldreich, S. Goldwasser, and S. Micali, "How to construct random functions," *Journal of the ACM*, vol. 33, no. 4, pp. 792–807, 1986.

[58] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby, "A pseudorandom generator from any one-way function," *SIAM Journal on Computing*, vol. 28, no. 4, pp. 1364–1396, 1999.

[59] A. Kiayias, S. Papadopoulos, N. Triandopoulos, and T. Zacharias, "Delegatable pseudorandom functions and applications," in *Proceedings of the 20th Annual ACM Conference on Computer and Communications Security*, 2013, pp. 669–684.

[60] D. Boneh and B. Waters, "Constrained pseudorandom functions and their applications," in *Advances in Cryptology - ASIACRYPT '13*, 2013, pp. 280–300.

[61] E. Boyle, S. Goldwasser, and I. Ivan, "Functional signatures and pseudorandom functions," in *Proceedings of the 17th International Conference on Practice and Theory in Public-Key Cryptography*, 2014, pp. 501–519.

[62] M. Luby, *Pseudorandomness and Cryptographic Applications*. Princeton University Press, 1996.

[63] O. Goldreich, "On security preserving reductions – revised terminology," Cryptology ePrint Archive, Report 2000/001, 2000.

[64] P. Ananth, D. Boneh, S. Garg, A. Sahai, and M. Zhandry, "Differing-inputs obfuscation and applications," Cryptology ePrint Archive, Report 2013/689, 2013.

[65] E. Boyle, K. Chung, and R. Pass, "On extractability obfuscation," in *Proceedings of the 11th Theory of Cryptography Conference*, 2014, pp. 52–73.

[66] S. Garg, C. Gentry, S. Halevi, and D. Wichs, "On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input," in *Advances in Cryptology – CRYPTO '14*, 2014, pp. 518–535.