On Constructing One-Way Permutations from Indistinguishability Obfuscation*

Gilad Asharov and Gil Segev

Hebrew University of Jerusalem, Jerusalem 91904, Israel.
{asharov,segev}@cs.huji.ac.il

Abstract. We prove that there is no black-box construction of a one-way permutation family from a one-way function and an indistinguishability obfuscator for the class of all oracle-aided circuits, where the construction is "domain invariant" (i.e., where each permutation may have its own domain, but these domains are independent of the underlying building blocks).

Following the framework of Asharov and Segev (FOCS '15), 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. For example, we fully capture the construction of a trapdoor permutation family from a one-way function and an indistinguishability obfuscator due to Bitansky, Paneth and Wichs (TCC '16). Their construction is *not* domain invariant and our result shows that this, somewhat undesirable property, is unavoidable using the common techniques.

In fact, we observe that constructions which are not domain invariant circumvent all known negative results for constructing one-way permutations based on one-way functions, starting with Rudich's seminal work (PhD thesis '88). We revisit this classic and fundamental problem, and resolve this somewhat surprising gap by ruling out *all* such black-box constructions – even those that are not domain invariant.

1 Introduction

One-way permutations are among the most fundamental primitives in cryptography, enabling elegant constructions of a wide variety of central cryptographic primitives. Although various primitives, such as universal one-way hash functions and pseudorandom generators, can be constructed based on any one-way

* This work was supported by the European Union's 7th Framework Program (FP7) via a Marie Curie Career Integration Grant, by the Israel Science Foundation (Grant No. 483/13), by the Israeli Centers of Research Excellence (I-CORE) Program (Center No. 4/11), by the US-Israel Binational Science Foundation (Grant No. 2014632), and by a Google Faculty Research Award.

function [57, 41], their constructions based on one-way permutations are much simpler and significantly more efficient [15, 52].

Despite the key role of one-way permutations in the foundations of cryptography, only very few candidates have been suggested over the years. Whereas one-way functions can be based on an extremely wide variety of assumptions, candidate one-way permutation families are significantly more scarce. Up until recently, one-way permutation families were known to exist only based on the hardness of problems related to discrete logarithms and factoring [56, 54]. Moreover, the seminal work by Rudich [58], within the framework of Impagliazzo and Rudich [44], initiated a line of research showing that a one-way permutation cannot be constructed in a black-box manner from a one-way function or from various other cryptographic primitives [24, 45, 51, 50].

Very recently, a one-way (trapdoor!) permutation family was constructed by Bitansky, Paneth and Wichs [13] based on indistinguishability obfuscation [6, 32] and one-way functions. Their breakthrough result provides the first trapdoor permutation family that is not based on the hardness of factoring, and motivates the task of studying the extent to which indistinguishability obfuscation can be used for constructing one-way permutations. Specifically, their work leaves completely unresolved the following question, representing to a large extent the "holy grail" of constructing one-way permutations:

Is there a construction of a one-way permutation over $\{0,1\}^n$
based on indistinguishability obfuscation and one-way functions?

While exploring this intriguing question, one immediately identifies two somewhat undesirable properties in the construction of Bitansky, Paneth and Wichs:

- Even when not aiming for trapdoor invertibility, their approach seems limited to providing a *family* of permutations instead of a *single* permutation[1].
- Their construction provides permutations that are defined over domains which both depend on the underlying building blocks and are extremely sparse[2].

From the theoretical perspective, one-way permutation families with these two properties are typically still useful for most constructions that are based on one-way permutations. However, such families lack the elegant structure that makes constructions based on one-way permutations more simple and significantly more efficient when compared to constructions based on one-way functions.

[1] Moreover, Bitansky et al. note that their permutations do not seem certifiable. That is, they were not able to provide an efficient method for certifying that a key is well-formed and describes a valid permutation. In contrast, a single permutation is certifiable by its nature.

[2] Each permutation in their construction is defined over a domain of elements of the form $(x, \mathsf{PRF}_K(x))$, where PRF is a pseudorandom function, and each permutation is associated with a different key K. This domain depends on the underlying building block, i.e., the pseudorandom function (equivalently, one-way function).

1.1 Our Contributions

Motivated by the recent construction of Bitansky et al. [13], we study the limitations of using indistinguishability obfuscation for constructing one-way permutations. Following the framework of Asharov and Segev [3], we consider indistinguishability obfuscation for *oracle-aided* circuits, and thus 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 [59] and its variants, as well as sub-exponential security assumptions. For example, we fully capture the construction of a trapdoor permutation family from a one-way function and an indistinguishability obfuscator due to Bitansky et al. [13]. We refer the reader to Section 1.3 for an overview of our framework and of the type of constructions that it captures.

Our work considers three progressively weaker one-way permutation primitives: (1) a *domain-invariant* one-way permutation, (2) a domain-invariant one-way permutation *family*, and (3) a one-way permutation family (which may or may not be domain invariant). Roughly speaking, we say that a construction of a one-way permutation (or a one-way permutation family) is domain invariant if the domain of the permutation is independent of the underlying building blocks (in the case of a permutation family we allow each permutation to have its own domain, but these domains have to be independent of the underlying building blocks).

Within our framework we prove the following two impossibility results, providing a tight characterization of the feasibility of constructing these three progressively weaker one-way permutation primitives based on one-way functions and indistinguishability obfuscation using the common techniques (we summarize this characterization in Figure 1).

$i\mathcal{O}$+OWF $\not\Rightarrow$ domain-invariant OWP family. Bitansky et al. [13] showed that any sub-exponentially-secure indistinguishability obfuscator and one-way function imply a one-way permutation family which is *not* domain invariant. We show that using the common techniques (as discussed above) one cannot construct the stronger primitive of a *domain-invariant* one-way permutation family (even when assuming sub-exponential security). In particular, we show that the above-described undesirable properties of their construction are unavoidable unless new non-black-box techniques are introduced.[3]

Theorem 1.1 *There is no fully black-box construction of a domain-invariant one-way permutation family from a one-way function f and an indistinguishability obfuscator for the class of all oracle-aided circuits C^f.*

[3] In addition to the above-described undesirable properties, our impossibility result holds even for constructions of one-way permutation families that have a "pseudo" input-sampling procedure instead of an "exact" input-sampling procedure (as in [13]), as well as to constructions that are not necessarily certifiable (again, as in [13]).

OWF $\not\Rightarrow$ OWP family. In fact, we observe that constructions which are not domain invariant circumvent the known negative results for constructing one-way permutations based on one-way functions, starting with Rudich's seminal work [58, 45, 51, 53]. We revisit this classic and fundamental problem, and resolve this surprising gap by ruling out *all* black-box constructions of one-way permutation *families* from one-way functions – even those that are *not* domain invariant.

Theorem 1.2 *There is no fully black-box construction of a one-way permutation family (even a non-domain-invariant one) from a one-way function.*

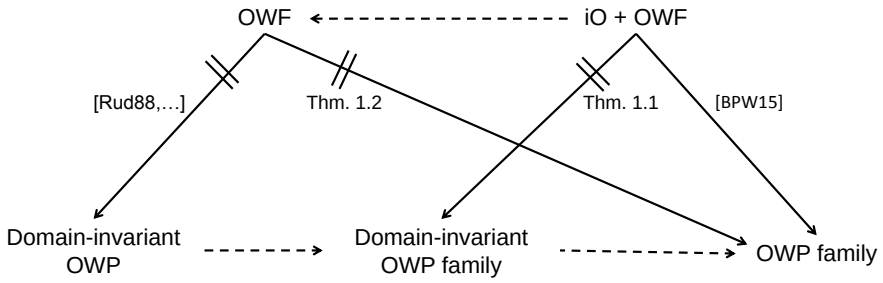

Fig. 1. A dashed arrow from a primitive A to a primitive B indicates that A implies B by definition.

Bitansky et al. [13] showed that any sub-exponentially-secure indistinguishability obfuscator and one-way function imply a one-way permutation family (which is not domain invariant), and we show that one cannot construct the stronger primitive of a *domain-invariant* one-way permutation family unless new non-black-box techniques are introduced (even when assuming sub-exponential security).

The line of research starting with Rudich [58] showed that one cannot construct a *domain-invariant* one-way permutation from a one-way function in a black-box manner. We improve this result, showing that one cannot construct the weaker primitive of a one-way permutation *family* (even one that is *not* domain invariant) from a one-way function in a black-box manner (again, even when assuming sub-exponential security).

1.2 Related Work

The recent line of research focusing on new constructions based on indistinguishability obfuscation has been extremely fruitful so far (e.g., [32, 8–10, 17, 19, 26, 39, 31, 21, 27, 33, 42, 46, 59, 1, 2, 11, 13, 12, 14, 23, 22, 61] and the references therein). However, the extent to which indistinguishability obfuscation can be used as a building block has been insufficiently explored. Our approach for proving meaningful impossibility results for constructions based on indistinguishability obfuscation is based on that of Asharov and Segev [3] (which, in turn,

was inspired by that of Brakerski, Katz, Segev and Yerukhimovich [18]). They showed that the common techniques (including non-black-box ones) that are used in constructions based on indistinguishability obfuscation can be captured by considering the stronger notion of indistinguishability obfuscation for oracle-aided circuits (see Section 1.3 for an elaborate discussion). Generalizing the work of Simon [60] and Haitner et al. [40], they showed that using these common techniques one cannot construct a collision-resistant hash function family from a general-purpose indistinguishability obfuscator (even when assuming sub-exponential security). In addition, generalizing the work of Impagliazzo and Rudich [44] and Brakerski et al. [18], they showed a similar result from constructing a perfectly-complete key-agreement protocol from a private-key functional encryption scheme (again, even when assuming sub-exponential security).

It is far beyond the scope of this paper to provide an overview of the lines of research on black-box impossibility results in cryptography (see, for example, [44, 60, 36, 43, 55, 34, 35, 62, 7, 28, 49, 5, 25, 29, 48, 16, 20] and the references therein). Impossibility results for constructing one-way permutations start with the seminal work of Rudich [58]. This line of research has successfully shown that one-way permutations cannot be based on a variety of fundamental cryptographic primitives (e.g., [24, 45, 51, 50]). However, these impossibility results capture only constructions of a *single* permutation that is *domain invariant*, and do not seem to capture more general constructions (such as the construction of Bitansky et al. [13] producing a permutation *family* which is *not* domain invariant).

The notion of "domain invariance" that we consider in this work for black-box constructions is somewhat related to that of "function obliviousness" that was introduced by Dachman-Soled, Mahmoody and Malkin [30] for coin-flipping protocols. They proved an impossibility result for constructing an optimally-fair coin-flipping protocol based on any one-way function, as long as the outcome of the protocol is completely independent of the specific one-way function that is used.

1.3 Overview of Our Results

In this section we provide a high-level overview of our two results. First, we describe the framework that enables us to prove a meaningful impossibility result for constructions that are based on indistinguishability obfuscation. Next, we describe Rudich's attack for inverting any domain-invariant permutation relative to a random oracle. Extending Rudich's approach, we then discuss the main technical ideas underlying our results: We present an attack on any domain-invariant permutation family relative to our, significantly more structured, oracle, and we generalize Rudich's attack to non-domain-invariant permutation families in the random-oracle model.

Capturing non-black-box constructions via $i\mathcal{O}$ for oracle-aided circuits The fact that constructions that are based on indistinguishability obfuscation 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 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, as observed by Asharov and Segev [3], most of the non-black-box techniques that are used on such constructions have essentially the same flavor: The obfuscator is applied to functions that can be constructed in a fully black-box manner from a low-level primitive, such as a one-way function. In particular, the vast majority of constructions rely on the obfuscator itself in a black-box manner. By considering the stronger primitive of an indistinguishability obfuscator for *oracle-aided* circuits (see Definition 2.4), Asharov and Segev showed that such 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 [59] 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 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 [59]. 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*, 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 [3, Sec. 4.6] for an in-depth technical treatment).

This reasoning extends to various variants of the punctured programming approach by Sahai and Waters [59], and in particular fully captures the construction of a trapdoor permutation family from a one-way function and an indistinguishability obfuscator due to Bitansky, Paneth and Wichs [13]. As noted in [3], this approach does not capture constructions that rely on the obfuscator itself in a non-black-box manner (e.g., [11])[4], or constructions that rely on zero-knowledge techniques and require using NP reductions[5].

[4] 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 [3, Sec. 3.1].

[5] Such techniques are captured by the work of Brakerski et al. [18], and we leave it as an intriguing open problem to see whether the two approaches for capturing non-black-box techniques can be unified.

The oracle. Our first result is obtained by presenting an oracle Γ relative to which the following two properties hold: (1) there is no domain-invariant one-way permutation family, and (2) there exist an *exponentially-secure* one-way function f and an *exponentially-secure* indistinguishability obfuscator $i\mathcal{O}$ for the class of all polynomial-size oracle-aided circuits C^f. Our oracle is quite intuitive and consists of three functions: (1) a random function f that will serve as the one-way function, (2) a random injective length-increasing function $\mathcal{O}$ that will serve as the obfuscator (an obfuscation of an oracle-aided circuit C is a "handle" $\mathcal{O}(C, r)$ for a uniformly-chosen string r), and (3) a function Eval 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 $C^f(x)$. We refer the reader to Section 3.2 for more details.

The vast majority of our effort is in showing that relative to Γ there is no domain-invariant one-way permutation family. Specifically, as for the second part, our oracle Γ is somewhat similar to the oracle introduced by [3], relative to which they proved the existence of an exponentially-secure one-way function and an exponentially-secure indistinguishability obfuscator (see Section 3.2 for the differences between the oracles).

In the remainder of this section we first provide a high-level overview of Rudich's attack on any single domain-invariant permutation in the random-oracle model. Inspired by this attack, we explain the main challenges in extending Rudich's attack to domain invariant constructions relative to our oracle, and to non-domain invariant constructions in the random-oracle model. We again refer the reader to Figure 1 which summarizes our characterization of the feasible constructions.

Warm-up: Rudich's attack in the random-oracle model Following [58, 45, 51] we show that for any oracle-aided polynomial-time algorithm P, if P^f implements a permutation over the same domain $\mathcal{D}$ for all functions f (i.e., P is domain invariant), then there exists an oracle-aided algorithm $\mathcal{A}$ that for any function f inverts P^f with probability 1 by querying f for only a polynomial number of times. The algorithm $\mathcal{A}$ is given some string $y^* \in \mathcal{D}$ and oracle access to f, and is required to find the unique $x^* \in \mathcal{D}$ such that $P^f(x^*) = y^*$. It first initializes a set of queries/answers Q, which will contain the actual queries made by $\mathcal{A}$ to the true oracle f. It repeats the following steps polynomially many times:

1. **Simulation:** $\mathcal{A}$ finds an input $x' \in \mathcal{D}$ and a set of oracle queries/answers f' that is consistent with Q (i.e., $f'(w) = f(w)$ for every $w \in Q$) such that $P^{f'}(x') = y^*$.
2. **Evaluation:** $\mathcal{A}$ evaluates $P^f(x')$ (i.e., evaluation with respect to the true oracle f). If the output is y^*, it terminates and outputs x'.
3. **Update:** $\mathcal{A}$ asks f for all queries in f' that are not in Q, and updates the set Q.

The proof relies on the following observation: In each iteration, either (1) $\mathcal{A}$ finds the pre-image x^* such that $P^f(x^*) = y^*$ or (2) in the update phase,

$\mathcal{A}$ queries f with at least one new query that is also made by P during the computation of $P^f(x^*) = y^*$.

Intuitively, if neither of the above holds, then we can construct a "hybrid" oracle $\widetilde{f}$ that behaves like f in the evaluation of $P^f(x^*) = y^*$ and behaves like f' in the evaluation of $P^{f'}(x') = y^*$. This hybrid oracle can be constructed since the two evaluations $P^{f'}(x')$ and $P^f(x^*)$ have no further intersection queries rather than the queries which are already in Q. According to this hybrid oracle $\widetilde{f}$ it holds that $P^{\widetilde{f}}(x') = P^{\widetilde{f}}(x^*) = y^*$ but yet $x^* \neq x'$, and thus relative to $\widetilde{f}$ the value y^* has two pre-images, in contradiction to the fact that P always implements a permutation. Using this claim, since there are only polynomially many f-queries in the evaluation of $P^f(x^*) = y^*$, the algorithm $\mathcal{A}$ must output x^* after a polynomial number of iterations (more specifically, after at most $q+1$ iterations, where q is the number of oracle gates in the circuit P).

Attacking domain-invariant permutation families relative to our oracle. We extend the attack described above in two different aspects. First, we rule out constructions of domain-invariant permutation *families* and not just a single permutation. Second, we extend the attack to work relative to our oracle, which is a significantly more structured oracle than a random oracle and therefore raises new technical challenges. Indeed, by the discussion in Section 1.3, relative to our oracle *there exists a non-domain-invariant construction of one-way permutation family* [13]. This mere fact represents the subtleties we have to deal with in our setting. In the following overview we focus our attention on the challenges that arise due to the structure of our oracle, as these are the most important and technically challenging ones.

Recall that our oracle Γ consists of three oracles: A length-preserving function f, an *injective* length-increasing function $\mathcal{O}$, and an "evaluation" oracle Eval that depends on both f and $\mathcal{O}$. We now sketch the challenges that these oracles introduce. The first challenge is that the evaluation oracle Eval is not just a "simple" function. This oracle performs (by definition) exponential time computations (e.g., an exponential number of queries to f and $\mathcal{O}$) which may give immense power to the construction P. Specifically, unlike in Rudich's case, here it is no longer true that the computation $P^\Gamma(x^*)$ performs a polynomial number of oracle queries (although P itself is of polynomial size). The second challenge is that since the oracle Eval depends on both f and $\mathcal{O}$, each query to Eval determines many other queries to f and $\mathcal{O}$ implicitly, which we need to make sure that they are considered in the attack. Specifically, given the structured dependencies between f, $\mathcal{O}$ and Eval, in some cases it may not be possible to construct a hybrid oracle even if there are no more intersection queries (in Rudich's case a hybrid oracle always exists).

Finally, the third challenge is the fact that $\mathcal{O}$ is *injective*, which causes the following problem (somewhat similar to [51]). In our case, we are forced to assume that P^Γ is a permutation only when $\mathcal{O}$ is an *injective length-increasing* function and not just any arbitrary function as in Rudich's case (as otherwise our obfuscator may not preserve functionality). Therefore, when constructing

the hybrid oracle $\widetilde{\mathcal{O}}$, we must ensure that it is also *injective* in order to reach a contradiction. However, the hybrid oracle $\widetilde{\mathcal{O}}$ might be non-injective when there is some overlap between the images of the true oracle $\mathcal{O}$ and the sampled oracle $\mathcal{O}'$ on elements that are not in Q.

We revise the attack and its analysis to deal with the above obstacles. As in Rudich's attack, the algorithm $\mathcal{A}$ considers the collection of all oracles that are consistent with Q. However, for dealing with the third challenge, it then chooses one of these oracles *uniformly at random* and does not pick just an arbitrarily one as in Rudich's attack. We then show that with all but an exponentially-small probability, there is no overlap between the range of the sampled oracle $\mathcal{O}'$ and the true oracle $\mathcal{O}$, and therefore the hybrid oracle $\widetilde{\mathcal{O}}$ can almost always be constructed in an injective manner. Then, dealing with the first challenge, we show that Eval does not give P a significant capability as one may imagine. Intuitively, this is due to the fact that $\mathcal{O}$ is length increasing, and therefore its range is very sparse. As a result, it is hard to sample a valid image of $\mathcal{O}$ without first querying it, and almost any Eval query can be simulated by the construction P itself. Finally, due to the dependencies between the oracles, for dealing with the second challenge, the algorithm $\mathcal{A}$ will have to sample additional, carefully-chosen, queries that do not necessarily appear in the evaluations $P^{\Gamma}(x^*) = y^*$ or $P^{\Gamma'}(x') = y^*$, but are related to the set of queries that appears in these evaluations. This results in a rather involved proof, where we carefully define this set of queries, and extend the analysis accordingly.

As expected, our proof does not extend to constructions that are not domain invariant. For example, in such constructions for two distinct (injective) functions Γ and Γ', the domain of the permutations P^{Γ} and $P^{\Gamma'}$ may be completely distinct, and this forces additional restrictions on the number of oracles Γ' that are "valid" (i.e., can be used to construct the hybrid oracle $\widetilde{\Gamma}$ as above). As a result, while in the original proof of Rudich all of the oracles Γ' that the adversary may pick are valid, and while in our case all but some exponentially-small amount of oracles Γ' are valid, here the number of valid oracles may be significantly smaller and therefore the attack may succeed with only a negligibly small probability.

Attacking non-domain-invariant permutation families in the random-oracle model. At a first sight, it seems that a natural approach towards ruling out non-domain-invariant families relative to a random oracle, is to reduce them to the case of a single permutation. That is, the adversary receives some index α of some permutation in the family, together with the challenge element $y^* \in \mathcal{D}^f_\alpha$ which it needs to invert (note that now the respective domain $\mathcal{D}^f_\alpha$ may depend on both f and α). A natural approach is to apply Rudich's attack to the single permutation $P^f(\alpha, \cdot)$.

However, this approach seems somewhat insufficient due to the following reasons. First, since the construction is not domain invariant, the set of valid indices depends on the underlying primitive, and the set of valid indices for the true oracle f may be completely different than the set of valid indices for the

oracle f' that will be sampled by $\mathcal{A}$ in each iteration (e.g., α might even not be a valid index with respect to the sampled f').

Second, when $\mathcal{A}$ inverts y^* relative to f', it may be that the pre-image x' that it finds is not even in the domain $\mathcal{D}_\alpha^f$ of the permutation $P^f(\alpha,\cdot)$ that it needs to invert. That is, it may be that even when the index α is valid relatively to both f and f', the domain of the permutation indexed by α relative to f is completely different than the domain relative to f'. One can try restricting $\mathcal{A}$ to sampling x' from the domain $\mathcal{D}_\alpha^f$, but conditioning on $P^{f'}(\alpha,x')=y^*$ it is not clear that such an x' even exists (and, even if it exists, $\mathcal{A}$ would typically need an exponential number of queries to f for finding it – since $\mathcal{A}$ has no "simple" representation of the sets $\mathcal{D}_\alpha^f$ and $\mathcal{D}_\alpha^{f'}$).

Finally, even when x' is the pre-image of y^* relative to f' and x^* is the pre-image of y^* relative to f, we have no guarantee that neither x' or x^* are even in the domain of the permutation indexed by α when considering the hybrid oracle $\widetilde{f}$. Therefore, the fact that $P^f(\alpha,x^*)=P^{f'}(\alpha,x')$ and $x^*\neq x'$ may not indicate any contradiction.

In Section 4 we show how to overcome these obstacles. Intuitively, when sampling some function f' and the element x', the algorithm $\mathcal{A}$ samples in addition two "certificates" that ensure that α is a valid index relative to f', and that x' is in the respective domain. These certificates include the randomness used by the index sampling and input sampling procedures of the permutation family, as well as all oracle queries and answers that are involved in the execution of these two procedures. We later use these certificates when defining the hybrid function $\widetilde{f}$, and thus ensure that α is a valid index relative to $\widetilde{f}$ and that x' is in the respective domain. Similarly, relative to the true oracle f, there exist some other certificates (which are unknown to $\mathcal{A}$), that ensure that α and x^* are valid, and are considered as well when defining the hybrid $\widetilde{f}$. Only then we can conclude the existence of a hybrid oracle $\widetilde{f}$ relative to which there exist an index α and two distinct inputs x^* and x' in the domain of α such that $P^{\widetilde{f}}(\alpha,x^*)=P^{\widetilde{f}}(\alpha,x')$. L

1.4 Paper Organization

The remainder of this paper is organized as follows. In Section 2 we introduce the cryptographic primitives under consideration in this paper, oracle-aided one-way permutation families and indistinguishability obfuscation for oracle-aided circuits, as well as some standard notation. In Section 3 we present our negative result for constructing domain-invariant one-way permutation families from indistinguishability obfuscation and one-way functions. Then, in Section 4 we present our negative result for constructing one-way permutation families from one-way functions.

2 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.

2.1 Oracle-Aided One-Way Permutation Families

We consider the standard notion of a one-way permutation family (see, for example, [38]) when naturally generalized to the setting of oracle-aided algorithms (as required within the context of black-box reductions [44, 55]). We start by formalizing the notion of an oracle-aided permutation family, and then introduce the standard one-wayness requirement.

Definition 2.1 *Let* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ *be a triplet of oracle-aided polynomial-time algorithms. We say that* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ *is* an oracle-aided permutation family relative to an oracle Γ *if the following properties are satisfied:*

- **Index sampling:** $\mathsf{Gen}^{\Gamma}(\cdot)$ *is a probabilistic algorithm that takes as input the security parameter* 1^n *and produces a distribution over indices* α*. For every* $n \in \mathbb{N}$ *we denote by* $\mathcal{I}_n^{\Gamma}$ *the support of the distribution* $\mathsf{Gen}^{\Gamma}(1^n)$*, and we let* $\mathcal{I}^{\Gamma} \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} \mathcal{I}_n^{\Gamma}$*.*
- **Input sampling:** $\mathsf{Samp}^{\Gamma}(\cdot)$ *is a probabilistic algorithm that takes as input an index* $\alpha \in \mathcal{I}^{\Gamma}$*, and produces a uniform distribution over a set denoted* $\mathcal{D}_{\alpha}^{\Gamma}$*.*
- **Permutation evaluation:** *For any index* $\alpha \in \mathcal{I}^{\Gamma}$*,* $\mathsf{P}^{\Gamma}(\alpha, \cdot)$ *is a deterministic algorithm that computes a permutation over the set* $\mathcal{D}_{\alpha}^{\Gamma}$*.*

Definition 2.2 *An oracle-aided permutation family* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ *is* one way relative to an oracle Γ *if for any probabilistic polynomial-time algorithm* $\mathcal{A}$ *there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathcal{A}^{\Gamma}(\alpha, \mathsf{P}^{\Gamma}(\alpha, x)) = x\right] \leq \mathsf{negl}(n)$$

for all sufficiently large $n \in \mathbb{N}$*, where the probability is taken over the choice of* $\alpha \leftarrow \mathsf{Gen}^{\Gamma}(1^n)$*,* $x \leftarrow \mathsf{Samp}^{\Gamma}(\alpha)$*, and over the internal randomness of* $\mathcal{A}$*.*

2.2 Indistinguishability Obfuscation for Oracle-Aided Circuits

We consider the standard notion of indistinguishability obfuscation [6, 32] when naturally generalized to oracle-aided circuits (i.e., circuits that may contain oracle gates in addition to standard gates). 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 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 2.3 *Let C_0 and C_1 be two oracle-aided circuits, and let f be a function. We say that C_0 and C_1 are* functionally equivalent relative to *f, denoted $C_0^f \equiv C_1^f$, if for any input x it holds that $C_0^f(x) = C_1^f(x)$.*

Definition 2.4 *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.*

3 Impossibility for Constructions Based on $i\mathcal{O}$ and One-Way Functions

In this section we present our negative result for domain-invariant constructions of a one-way permutation family from from a one-way function and an indistinguishability obfuscator. In section 3.1 we formally define the class of constructions to which our negative result applies. Then, in section 3.2 we present the structure of our proof, which is provided in Sections 3.3–3.4.

3.1 The Class of Constructions

We consider fully black-box constructions of a one-way permutation family from a one-way function f and an indistinguishability obfuscator for all oracle-aided circuits C^f. Following [3], we model these primitives as two independent building blocks due to the following reasons. First, although indistinguishability obfuscation is known to imply one-way functions under reasonable assumptions [46], this enables us to prove an unconditional result. Second, and more importantly, this enables us to capture the common techniques that have been used so far in constructions based on indistinguishability obfuscation. As discussed in Section

1.3, these include, in particular, *non-black-box* techniques such as the punctured programming approach of Sahai and Waters [59] and its variants.

We now formally define the class of constructions considered in this section, tailoring our definitions to the specific primitives under consideration. We remind the reader that two oracle-aided circuits, C_0 and C_1, are functionally equivalent relative to a function f, denoted $C_0^f \equiv C_1^f$, if for any input x it holds that $C_0^f(x) = C_1^f(x)$ (see Definition 2.3). The following definition is based on those of [3] (which, in turn, are motivated by [47, 37, 55]).

Definition 3.1 *A fully black-box construction of a one-way permutation family from a one-way function and an indistinguishability obfuscator for the class $\mathcal{C} = \{C_n\}_{n\in\mathbb{N}}$ of all polynomial-size oracle-aided circuits, consists of a triplet of oracle-aided probabilistic polynomial-time algorithms $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$, 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 conditions hold:*

- **Correctness:** *For any functions f $i\mathcal{O}$ such that $i\mathcal{O}(C;r)^f \equiv C^f$ for all $C \in \mathcal{C}$ and $r \in \{0,1\}^*$, the triplet $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ is a permutation family relative to the oracle $(f, i\mathcal{O})$ (as in Definition 2.1).*
- **Black-box proof of security:** *For any function 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 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*

$$\Pr\left[\mathcal{A}^{f,i\mathcal{O}}(\alpha, \mathsf{P}^{f,i\mathcal{O}}(\alpha, x)) = x\right] \geq \epsilon_{\mathcal{A}}(n)$$

for infinitely many values of $n \in \mathbb{N}$, where the probability is taken over the choice of $\alpha \leftarrow \mathsf{Gen}^{f,i\mathcal{O}}(1^n)$, $x \leftarrow \mathsf{Samp}^{f,i\mathcal{O}}(\alpha)$, and over the internal randomness of $\mathcal{A}$, then either

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

for infinitely many values of $n \in \mathbb{N}$, where the probability is taken over the choice of $x \leftarrow \{0,1\}^n$ and over the internal randomness of M, 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 $n \in \mathbb{N}$ (see Definition 2.4 for the description of the experiment $\mathsf{Exp}^{\mathsf{iO}}_{(f,i\mathcal{O}),i\mathcal{O},M^{\mathcal{A}},\mathcal{C}}(n)$).

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 [47] and Goldreich [37], 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. [13] (and in various other recent constructions based on indistinguishability obfuscation), the adversary-dependent loss is polynomial whereas the adversary-independent loss is sub-exponential[6].

Domain-invariant constructions. We now define the notion of *domain invariance* which allows us to refine the above class of constructions. Recall that for an oracle-aided permutation family $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ and for any oracle Γ, we denote by $\mathcal{I}_n^\Gamma$ the support of the distribution $\mathsf{Gen}^\Gamma(1^n)$ for every $n \in \mathbb{N}$, and we let $\mathcal{I}^\Gamma \stackrel{\text{def}}{=} \bigcup_{n\in\mathbb{N}} \mathcal{I}_n^\Gamma$ (i.e., $\mathcal{I}^\Gamma$ is the set of all permutation indices). In addition, for any permutation index $\alpha \in \mathcal{I}^\Gamma$ we denote by $\mathcal{D}_\alpha^\Gamma$ the domain of the permutation $\mathsf{P}^\Gamma(\alpha, \cdot)$.

Definition 3.2 *An oracle-aided one-way permutation family* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ *is* domain invariant relative to a set $\mathfrak{S}$ of oracles *if there exist sequences* $\{\mathcal{I}_n\}_{n\in\mathbb{N}}$ *and* $\{\mathcal{D}_\alpha\}_{\alpha\in\mathcal{I}}$ *such that for every oracle* $\Gamma \in \mathfrak{S}$ *the following conditions hold:*

1. $\mathcal{I}_n^\Gamma = \mathcal{I}_n$ *for every* $n \in \mathbb{N}$ *(i.e., a permutation index* α *is either valid with respect to all oracles in* $\mathfrak{S}$ *or invalid with respect to all oracles in* $\mathfrak{S}$*).*
2. $\mathcal{D}_\alpha^\Gamma = \mathcal{D}_\alpha$ *for every* $\alpha \in \bigcup_{n\in\mathbb{N}} \mathcal{I}_n$ *(i.e., the domain of* $\mathsf{P}^\Gamma(\alpha, \cdot)$ *is the same for all* $\Gamma \in \mathfrak{S}$*).*

3.2 Proof Overview and the Oracle Γ

Our result in this section is obtained by presenting a distribution over oracles Γ relative to which the following two properties hold: (1) there is no domain-invariant one-way permutation family $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$, and (2) there exist an *exponentially-secure* one-way function f and an *exponentially-secure* indistinguishability obfuscator $i\mathcal{O}$ for the class of all polynomial-size oracle-aided circuits C^f. Equipped with the notation and terminology introduced in Section 3.1, we prove the following theorem:

Theorem 3.3 *Let* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P}, M, T_M, \epsilon_{M,1}, \epsilon_{M,2})$ *be a fully black-box domain-invariant construction of a one-way permutation family from a one-way function* f *and an indistinguishability obfuscator for the class of all polynomial-size oracle-aided circuits* C^f*. Then, at least one of the following propertied holds:*

[6] 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.

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 (and, in particular, as in the construction of Bitansky et al. [13]), 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 – as discussed in Section 3.1).

In what follows we describe the oracle Γ (more accurately, the distribution over such oracles), and then explain the structure of our proof.

The oracle Γ. The oracle Γ is a triplet $\left(f, \mathcal{O}, \mathsf{Eval}^{f,\mathcal{O}}\right)$ that is sampled from a distribution $\mathfrak{S}$ 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 function $f_n : \{0,1\}^n \to \{0,1\}^n$.
  Looking ahead, we will prove that f is a one-way function relative to Γ.
- **The functions $\mathcal{O} = \{\mathcal{O}_n\}_{n\in\mathbb{N}}$ and $\mathsf{Eval}^{f,\mathcal{O}} = \{\mathsf{Eval}_n^{f,\mathcal{O}}\}_{n\in\mathbb{N}}$.** For every $n \in \mathbb{N}$ the function $\mathcal{O}_n$ is an injective function $\mathcal{O}_n : \{0,1\}^{2n} \to \{0,1\}^{10n}$ chosen uniformly at random. The function $\mathsf{Eval}_n^{f,\mathcal{O}}$ on input $(\widehat{C}, x) \in \{0,1\}^{10n} \times \{0,1\}^n$ finds the unique pair $(C, r) \in \{0,1\}^n \times \{0,1\}^n$ such that $\mathcal{O}_n(C, r) = \widehat{C}$, where C is an oracle-aided circuit and r is a string (uniqueness is guaranteed since $\mathcal{O}_n$ is injective). If such a pair exists, it evaluates and outputs $C^f(x)$, and otherwise it outputs $\perp$.
  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 structure of our proof. Our proof consists of two parts: (1) showing that relative to Γ there is no domain-invariant one-way permutation family, and (2) showing that relative to Γ the function f is an *exponentially-secure* one-way function and that the pair $(\mathcal{O}, \mathsf{Eval})$ can be used for implementing an *exponentially-secure* indistinguishability obfuscator for oracle-aided circuits C^f.

The vast majority of our effort in this proof is in showing that relative to Γ there is no domain-invariant one-way permutation family. Specifically, as for the second part, our oracle Γ is somewhat similar to the oracle introduced by [3], relative to which they proved the existence of an exponentially-secure one-way function and an exponentially-secure indistinguishability obfuscator. The main difference between the oracles is that the function $\mathcal{O}$ in their case is a *permutation*, whereas in our case it is an *injective length-increasing* function. Since our aim here is to rule out constructions of one-way permutations, then clearly we cannot allow $\mathcal{O}$ to be a permutation. This requires us to revisit the proof of [3] and generalize it to the case where $\mathcal{O}$ is injective and length increasing.

In what follows, we say that an algorithm $\mathcal{A}$ that has oracle access to Γ is a q-query algorithm if it makes at most q queries to Γ, and each of its queries to Eval consists of a circuit of size at most q.

Part 1: Inverting any domain-invariant construction. Building upon and generalizing the work of Rudich [58], we show that relative to the oracle Γ there are no domain-invariant one-way permutations families. As discussed in Section 1.3, Rudich presented an attacker that inverts any *single* domain-invariant permutation that has oracle access to a random function. Here we need to deal with constructions that have oracle access to a significantly more structured functionality[7], and that are permutation *families*. Nevertheless, inspired by the main ideas underlying Rudich's attacker we prove the following theorem in Section 3.3:

Theorem 3.4 (simplified) *Let* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ *be an oracle-aided domain-invariant permutation family. Then, there exist a polynomial* $q(\cdot)$ *and a* q*-query algorithm* $\mathcal{A}$ *such that*

$$\Pr\left[\mathcal{A}^{\Gamma}(\alpha, \mathsf{P}^{\Gamma}(\alpha, x)) = x\right] \geq 1 - 2^{-10}$$

for any $n \in \mathbb{N}$*, where the probability is taken over the choice of* $\Gamma \leftarrow \mathfrak{S}$*,* $\alpha \leftarrow \mathsf{Gen}^{\Gamma}(1^n)$*,* $x \leftarrow \mathsf{Samp}^{\Gamma}(\alpha)$*, and over the internal randomness of* $\mathcal{A}$*. Moreover, the algorithm* $\mathcal{A}$ *can be implemented in polynomial time given access to a* PSPACE*-complete oracle.*

Part 2: The existence of a one-way function and an indistinguishability obfuscator. As discussed above, by refining the proof of [3] we prove that f is an exponentially-secure one-way function relative to Γ, and we construct an exponentially-secure indistinguishability obfuscator $i\mathcal{O}$. Our obfuscator is defined as follows: For obfuscating an oracle-aided circuit $C \in \{0,1\}^n$ (i.e., we denote by $n = n(C)$ the bit length of C's representation), the obfuscator $i\mathcal{O}$ samples $r \leftarrow \{0,1\}^n$ uniformly at random, computes $\widehat{C} = \mathcal{O}_n(C, r)$, and outputs the circuit $\mathsf{Eval}(\widehat{C}, \cdot)$. That is, the obfuscated circuit consists of a single Eval gate with hardwired input $\widehat{C}$. We prove the following theorem in the full version of this paper [4]:

Theorem 3.5 (simplified) *For any oracle-aided* $2^{n/4}$*-query algorithm* $\mathcal{A}$ *it hold that*

$$\Pr\left[\mathcal{A}^{\Gamma}(f(x)) \in f^{-1}(f(x))\right] \leq 2^{-n/2}$$

and

$$\left|\Pr\left[\mathsf{Exp}^{\mathsf{iO}}_{\Gamma, i\mathcal{O}, \mathcal{A}, \mathcal{C}}(n) = 1\right] = 1 - \frac{1}{2}\right| \leq 2^{-n/4}$$

[7] For example, there are dependencies between $\mathcal{O}$, Eval and f which allow Eval to query $\mathcal{O}$ for a exponential number of times.

for all sufficiently large $n \in \mathbb{N}$, where the probability is taken over the choice of $\Gamma \leftarrow \mathfrak{S}$ and internal randomness of $\mathcal{A}$ for both cases, in addition to the choice of $x \leftarrow \{0,1\}^n$ in the former case and to the internal randomness of the challenger in the latter case.

3.3 Attacking Domain-Invariant Permutation Families Relative to Γ

We show that relative to the oracle Γ there are no domain-invariant one-way permutations families. As discussed in Section 1.3, Rudich presented an attacker that inverts any *single* domain-invariant permutation that has oracle access to a random function. Here we need to deal with constructions that have oracle access to a significantly more structured functionality. We prove the following theorem:

Theorem 3.6 *Let $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ be an oracle-aided permutation family that is domain invariant relative to the support of the distribution $\mathfrak{S}$. Then, there exist a polynomial $q(\cdot)$ and a q-query algorithm $\mathcal{A}$ such that*

$$\Pr\left[\mathcal{A}^{\Gamma}(\alpha, \mathsf{P}^{\Gamma}(\alpha, x^*)) = x^*\right] \geq 1 - 2^{-10}$$

for any $n \in \mathbb{N}$, where the probability is taken over the choice of $\Gamma \leftarrow \mathfrak{S}$, $\alpha \leftarrow \mathsf{Gen}^{\Gamma}(1^n)$, $x^* \leftarrow \mathsf{Samp}^{\Gamma}(\alpha)$, and over the internal randomness of $\mathcal{A}$. Moreover, the algorithm $\mathcal{A}$ can be implemented in polynomial time given access to a PSPACE*-complete oracle.*

We first provide additional notation definitions that we require for the proof of the above theorem, and then we provide its formal proof.

The event spoof. The event spoof will help up show that the oracle Eval does not provide the construction with any significant capabilities. We formally define this event and then state an important claim that will help up to prove our theorem.

Definition 3.7 *For any oracle-aided algorithm M, consider the following event spoof_n that may occur during an execution of $M^{\Gamma}(1^n)$: The algorithm makes a query $\mathsf{Eval}_n(\widehat{C}, a)$ with $|\widehat{C}| = 10n$ whose output is not $\perp$, yet $\widehat{C}$ was not an output of a previous $\mathcal{O}_n$-query.*

In the full version of this paper [4] we prove the following claim:

Claim 3.8 *For any $n \in \mathbb{N}$, for any f and $\mathcal{O}_{-n} = \{\mathcal{O}_m\}_{m \in \mathbb{N}, m \neq n} m$ and for any q-query algorithm M, the probability that spoof_n occurs in an execution of $M^{\Gamma}(1^n)$ satisfies*

$$\Pr_{\mathcal{O}_n}[\ \mathsf{spoof}_n\ ] \leq q \cdot 2^{-8n}\ .$$

Notation. Denote by $\mathcal{T}$ the support of the distribution $\mathfrak{S}$ from which our oracle $\Gamma = (f, \mathcal{O}, \mathsf{Eval}^{f,\mathcal{O}})$ is sampled. Note that the oracle Eval is fully determined given f and $\mathcal{O}$, and therefore it is enough to consider the choice of the latter only. For every $n \in \mathbb{N}$ we let $\mathcal{I}_n$ denote the support of $\mathsf{Gen}^{\Gamma}(1^n)$, which is the same for every $\Gamma \in \mathcal{T}$ due to the domain invariant assumption, and we let $\mathcal{I} = \bigcup_{n\in\mathbb{N}} \mathcal{I}_n$. In addition, we let $\mathcal{D} = \{D_\alpha\}_{\alpha\in\mathcal{I}}$ be the set of domains (which is again the same for any $\Gamma \in \mathcal{T}$).

We let $\mathsf{Partial}(\Gamma')$ denote the set of oracle queries that our adversary $\mathcal{A}$ will sample in each iteration. We let Q denote the set of actual queries that made by $\mathcal{A}$ to the true oracle Γ. We write, e.g., $[\mathcal{O}_n(C, r) = \widehat{C}] \in Q$ to denote that Q contains an $\mathcal{O}_n$-query with input (C, r) and output $\widehat{C}$. Likewise, $[f_n(x) = y] \in \mathsf{Partial}(\Gamma')$ denotes that there is some f_n query in $\mathsf{Partial}(\Gamma')$ with input x and output y. We also use the symbol $\star$ to indicate an arbitrary value, for instance $[\mathsf{Eval}(\widehat{C}, a) = \star] \in Q$ denotes that $\mathcal{A}$ made an Eval call to Γ on the pair $(\widehat{C}, a)$, but we are not interested in the value that was returned by the oracle.

The set of queries/answers that the adversary samples. Our adversary $\mathcal{A}$ will sample in each iteration some oracle queries/answers $\mathsf{Partial}(\Gamma') = (f', \mathcal{O}', \mathsf{Eval}')$ that are consistent with the actual queries Q it made so far. However, since the oracles $(f, \mathcal{O}, \mathsf{Eval})$ have some dependencies, we want that these dependencies will appear explicitly in the set of queries/answers that the adversary samples (looking ahead, by doing so, we will be able to construct a hybrid oracle $\widetilde{\Gamma}$). Formally, we define:

Definition 3.9 (Consistent oracle queries/answers) *Let* $\mathsf{Partial}(\Gamma') = (f', \mathcal{O}', \mathsf{Eval}')$ *be a set of queries/answers. We say it is* consistent *if for every* $m \in \mathbb{N}$ *it holds that:*

1. *For every query* $\left[\mathsf{Eval}_m(\widehat{C}, \star) = \star\right] \in \mathsf{Eval}'$*, there exists a query* $\left[\mathcal{O}_m(\star) = \widehat{C}\right] \in \mathcal{O}'$*.*
2. *For every query* $\left[\mathsf{Eval}_m(\widehat{C}, a) = \beta\right] \in \mathsf{Eval}'$ *with* $|\widehat{C}| = 10m$ *and* $|a| = m$*, let* $\left[\mathcal{O}_m(C, r) = \widehat{C}\right] \in \mathcal{O}'$ *that is guaranteed to exist by the previous requirement. Then, the oracle* f' *contains also queries/answers sufficient for the evaluation of* $C^{f'}(a)$*, and the value of this evaluation is indeed* β*.*

Augmented oracle queries. For the analysis, we consider the queries that are associated with the execution of $\mathsf{P}^{\Gamma}(\alpha, x^*) = y^*$, for some $\alpha \in \mathcal{I}$. In fact, the set that we consider may contain some additional queries that do not necessarily appear in the execution of $\mathsf{P}^{\Gamma}(\alpha, x^*)$, but are still associated with this execution. Let $\mathsf{RealQ}(\Pi, \Gamma, \alpha, x^*)$ denote the set of actual queries to Γ in the evaluation of $\mathsf{P}^{\Gamma}(\alpha, x^*)$. We define:

Definition 3.10 (Augmented oracle queries) *The set of extended queries, denoted* $\mathsf{AugQ}(\Pi, \Gamma, x^*)$*, consists of the following queries:*

1. *All the queries in* $\mathsf{RealQ}(\Pi, \Gamma, \alpha, x^*)$*.*

2. *For every query* $[\mathsf{Eval}_m(\widehat{C}, a) = \beta] \in \mathsf{RealQ}(\Pi, \Gamma, \alpha, x^*)$ *with* $|\widehat{C}| = 10m$, $|a| = m$ *and* $b \neq \perp$, *let* $C, r \in \{0,1\}^m$ *be the unique pair such that* $\mathcal{O}_m(C, r) = \widehat{C}$. *Then, the set* $\mathsf{AugQ}(\Pi, \Gamma, x^*)$ *contains all the* f*-queries/answers sufficient to for the evaluation of* $C^f(a)$.

Note that these additional queries correspond to the consistent queries/answers that the adversary samples in the attack, as in Definition 3.9. We do not explicitly require the first requirement of Definition 3.9 here. This is because our analysis focuses on the case where there is no Eval query on an obfuscated circuit $\widehat{C}$ that is not an output of a previous $\mathcal{O}$-query.

Looking ahead, all the circuits that will be evaluated by the oracle Eval are of some polynomial size in the security parameter,and therefore each evaluation adds some polynomial number of oracle queries to f. Therefore, the overall size of $\mathsf{AugQ}(\Pi, \Gamma, x^*)$ is some polynomial. Let $\ell = \ell(n) > n$ be an upper bound of $|\mathsf{AugQ}(P, \widetilde{\Gamma}, x)|$ for all possible $\widetilde{\Gamma} \in \mathcal{T}$ and all $x \in D_\alpha$.

Equipped with the above notation and definitions, we are now ready to prove Theorem 3.6.

Proof of Theorem 3.6. Let $\Pi = (\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ be an oracle-aided permutation family that is domain invariant relative to the support of the distribution $\mathfrak{S}$. Consider the following oracle-aided algorithm $\mathcal{A}$:

The algorithm $\mathcal{A}$.

- **Input:** *An index* $\alpha \in \mathcal{I}$ *and a value* $y^* \in D_\alpha$.
- **Oracle access:** *The oracle* Γ.
- **The algorithm:**
  1. *Initialize an empty list* Q *of oracle queries/answers to* Γ *(looking ahead, the list* Q *will always be consistent with the true oracle* Γ*).*
  2. **Avoiding spoof$_m$ for small m.** *Let* $t = \log(16\ell)$. *The adversary* $\mathcal{A}$ *queries the oracle* f_m *on all inputs* $|x| = m$ *for all* $m \leq t$. *It queries* $\mathcal{O}_m(C, r)$ *for all* $|C| = |r| = m \leq t$*; and queries* $\mathsf{Eval}_m(\widehat{C}, a)$ *on all* $m \leq t$ *with* $|\widehat{C}| = m/10$ *and* $|a| = m$. *Denote this set of queries by* Q^*.
  3. *Run the following for* $\ell + 1$ *iterations:*
     (a) **Simulation phase:** $\mathcal{A}$ *finds a value* $x' \in D_\alpha$ *and a set* $\mathsf{Partial}(\Gamma')$ *of consistent oracle queries/answers that is consistent with the list of queries/answers* Q*, such that* $\mathsf{P}^{\mathsf{Partial}(\Gamma')}(\alpha, x') = y^*$ *as follows:* [8]
        i. $\mathcal{A}$ *samples an oracle* $\Gamma' = (f', \mathcal{O}', \mathsf{Eval}')$ *uniformly at random from the set of all oracles that are consistent with* Q. *That is,* f' *and* $\mathcal{O}'$ *are sampled uniformly at random conditioned on* Q*, and then* Eval' *is defined accordingly.*
        ii. $\mathcal{A}$ *inverts* y^* *relative to* Γ'. *Specifically,* $\mathcal{A}$ *enumerates over* $\mathcal{D}_\alpha$ *and find the unique input* $x' \in D_\alpha$ *for which* $\mathsf{P}^{\Gamma'}(\alpha, x') = y^*$.

[8] Note that the set of queries/answers $\mathsf{Partial}(\Gamma')$ may be inconsistent with the true oracle Γ on all queries $\mathsf{Partial}(\Gamma') \setminus Q$.

iii. $\mathcal{A}$ *sets* $\mathsf{Partial}(\Gamma')$ *to be all the queries in* Q*, and all the queries included in the evaluation of* $\mathsf{P}^{\Gamma'}(\alpha, x')$*.*

(b) **Evaluation phase:** *The adversary evaluates* $\mathsf{P}^{\Gamma}(\alpha, x')$*. If the output of the evaluation is* y^**, it halts and outputs* x'*.*

(c) **Update phase:** *Otherwise,* $\mathcal{A}$ *makes all the queries in* $\mathsf{Partial}(\Gamma')\backslash Q$ *to the true oracle* Γ*, and continues to the next iteration.*

4. In case the adversary has not halted yet, it outputs $\perp$*.*

Analysis. We show that in each iteration the adversary either finds x^* or learns some query associated with the evaluation $\mathsf{P}^{\Gamma}(\alpha, x^*)$. We now define these two "bad" events and show that they occur with small probability. We then proceed to the analysis conditioned that these two bad events do not occur.

The event spoof. For any $m \in \mathbb{N}$, define spoof_m to be the event where

$$\left[\mathsf{Eval}_m(\widehat{C}, a) \neq \perp\right] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$$

but

$$\left[\mathcal{O}_m(\star, \star) = \widehat{C}\right] \notin \mathsf{AugQ}(\Pi, \Gamma, x^*) \cup Q^* \ .$$

Let $\mathsf{spoof}_\Gamma = \bigvee_m \mathsf{spoof}_m$. By construction, Q^* contains all possible $\mathcal{O}_m$-queries for every $m \leq t$, and therefore spoof_m cannot occur for $m \leq t$. Moreover, by Claim 3.8, we have that

$$\begin{aligned}\Pr\left[\ \mathsf{spoof}_\Gamma\ \right] &\leq \Pr\left[\ \bigvee_m \mathsf{spoof}_m\ \right] \leq \sum_{m=t}^{\infty} \Pr\left[\ \mathsf{spoof}_m\ \right] \\ &\leq \sum_{m=\log 16\ell}^{\infty} \ell \cdot 2^{-8m} \leq 2 \cdot \ell \cdot 2^{-8 \log 16\ell} \leq 2^{-31}\end{aligned}$$

Let spoof'_m be the event where the adversary $\mathcal{A}$ queries the real oracle Γ some query $[\mathsf{Eval}_m(\widehat{C}, \star)]$, receives a value differ than $\perp$, but $\widehat{C}$ was not an output of Γ on some previous query of $\mathcal{A}$ to $\mathcal{O}_m$. Let $\mathsf{spoof}_{\mathcal{A}} = \bigvee_m \mathsf{spoof}'_m$. Similarly to the above, the probability of $\mathsf{spoof}_{\mathcal{A}}$ is bounded by 2^{-31}. Finally, we let $\mathsf{spoof} = \mathsf{spoof}_\Gamma \vee \mathsf{spoof}_{\mathcal{A}}$, and this probability is bounded by 2^{-30}.

The event fail. The second bad event that we consider is the event fail. This event occurs whenever $\mathcal{A}$ samples an oracle Γ' that has some contradiction with the oracle Γ, and therefore the hybrid oracle $\widetilde{\Gamma}$ cannot be constructed.

Let $\mathcal{T}(Q)$ be the set of all oracles Γ' that are consistent with Q (namely, each query in Q is answered the same for all $\Gamma' \in \mathcal{T}(Q)$, with the same answer as Γ). In each iteration, the adversary $\mathcal{A}$ samples the oracle Γ' which is consistent with the true oracle queries Q. Let $\mathcal{T}\text{-}\mathsf{admissible}$ denote the set of "valid" oracles that $\mathcal{A}$ may sample; the set $\mathcal{T}\text{-}\mathsf{admissible}$ contains all oracles $\Gamma' = (f', \mathcal{O}', \mathsf{Eval}')$ such that:

- Γ' *is consistent with* Q.
- Γ' *avoids the outputs of* $\mathcal{O}$. For every $m \in \mathbb{N}$, the true oracle $\mathcal{O}_m$ and the sampled oracle $\mathcal{O}'_m$ should have disjoint outputs (except for the queries in Q). Formally, let $Q_m^{\mathcal{O}} = \{x \in \{0,1\}^{2m} \mid [\mathcal{O}_m(x) = \star] \in Q\}$. Then, we require that for every $x, y \notin Q_m^{\mathcal{O}}$ it holds that $\mathcal{O}_m(x) \neq \mathcal{O}'_m(y)$.
- Γ' *avoids invalid* Eval*-queries.* That is, for every $\left[\mathsf{Eval}_m(\widehat{C}, a) = \bot\right] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$, with $|\widehat{C}| = 10m$, for every $C, r \in \{0,1\}^m$ it holds that $\mathcal{O}'_m(C, r) \neq \widehat{C}$.

Notice that the first two conditions relate to the set of queries Q, whereas the third condition relates to the set $\mathsf{AugQ}(\Pi, \Gamma, x^*)$. Moreover, note that the second condition defines $2^{2m} - |Q|$ outputs of $\mathcal{O}'_m$ that are invalid, and the third condition defines at most q invalid outputs. Therefore, there are overall at most 2^{2m} outputs of $\mathcal{O}'_m$ that are invalid.

Note that between iterations, the set Q varies. We define by $\mathsf{Invalid\text{-}Im}_m^{(i)}$ the set of all invalid outputs for $\mathcal{O}'_m$, in the ith iteration. In all iterations, the set $\mathsf{Invalid\text{-}Im}_m^{(i)}$ is bounded by 2^{2m}.

Let $\mathsf{fail}_m^{(i)}$ denote the event where $\mathcal{A}$ samples an invalid oracle $\mathcal{O}'_m$ in some iteration i. Let $\mathsf{fail}^{(i)} = \bigvee_m \mathsf{fail}_m^{(i)}$, and let $\mathsf{fail} = \bigvee_i \mathsf{fail}^{(i)}$. For every m, we have that:

$$\begin{aligned}\Pr_{\mathcal{O}'_m}\left[\ \mathsf{fail}_m^{(i)}\ \right] &= \Pr_{\mathcal{O}'_m}\left[\ \exists x \in \{0,1\}^{2m} \text{ s.t. } \mathcal{O}'_m(x) \in \mathsf{Invalid\text{-}Im}_m^{(i)}\ \right] \\ &\leq 2^{2m} \cdot \frac{\left|\mathsf{Invalid\text{-}Im}_m^{(i)}\right|}{2^{10m} - 2^{2m}} \leq 2^{-5m}\ .\end{aligned}$$

As a result, we get that the probability that sampling $\mathcal{O}$ fails for some length $m > t$ is bounded by

$$\Pr_{\mathcal{O}'}\left[\ \mathsf{fail}^{(i)}\ \right] \leq \sum_{m=t}^{\infty} 2^{-5m} \leq 2 \cdot 2^{-5t}\ .$$

We therefore conclude that the probability that in some of the $\ell + 1$ iterations, the adversary $\mathcal{A}$ samples some oracle $\Gamma' \notin \mathcal{T}\text{-}\mathsf{admissible}$ is bounded by

$$\Pr[\ \mathsf{fail}\ ] \leq \sum_{i=1}^{\ell+1} \Pr\left[\ \mathsf{fail}^{(i)}\ \right] \leq (\ell+1) \cdot 2 \cdot 2^{-5t} = 2(\ell+1) \cdot \left(2^{-4} \cdot \ell^{-1}\right)^5 \leq 2^{-19}\ ,$$

where recall that $t = \log(16\ell)$. We are now ready for the main claim of the analysis.

Claim 3.11 *Assume that* fail *and* spoof *do not occur. Then, in every iteration at least one of the following occurs:*

1. $\mathcal{A}$ *finds the pre-image* x^* *such that* $\mathsf{P}^{\Gamma}(\alpha, x^*) = y^*$.
2. *During the update phase* $\mathcal{A}$ *queries* Γ *with at least one of the queries in* $\mathsf{AugQ}(\Pi, \Gamma, x^*)$.

Proof. Assume that neither one of the above conditions hold. Then, we show that there exists an oracle $\widetilde{\Gamma} \in \mathcal{T}$ that behaves like the true oracle Γ on $\mathsf{P}^{\widetilde{\Gamma}}(\alpha, x^*) = \mathsf{P}^{\Gamma}(\alpha, x^*) = y^*$, and on the other hand, it behaves like Γ' in the evaluation of $\mathsf{P}^{\widetilde{\Gamma}}(\alpha, x') = \mathsf{P}^{\mathsf{Partial}(\Gamma')}(\alpha, x') = y^*$. According to this oracle $\widetilde{\Gamma}$, the following hold:

1. Since Π is a domain-invariant construction, and since $\widetilde{\Gamma} \in \mathcal{T}$, there exists some randomness $r \in \{0,1\}^*$ such that $\mathsf{Gen}^{\widetilde{\Gamma}}(1^n; r) = \alpha$.
2. Since Π is a domain-invariant construction, it holds that $\mathrm{Im}(\mathsf{Samp}^{\widetilde{\Gamma}}(\alpha)) = \mathrm{Im}(\mathsf{Samp}^{\Gamma}(\alpha)) = \mathrm{Im}(\mathsf{Samp}^{\mathsf{Partial}(\Gamma')}(\alpha)) = D_\alpha$. As a result, there exists some randomness $r' \in \{0,1\}^*$ such that $\mathsf{Samp}^{\widetilde{\Gamma}}(\alpha; r') = x'$ and $\mathsf{Samp}^{\widetilde{\Gamma}}(\alpha; r^*) = x^*$.
3. As mentioned above, $\mathsf{P}^{\widetilde{\Gamma}}(\alpha, x') = y^*$ and $\mathsf{P}^{\widetilde{\Gamma}}(\alpha, x^*) = y^*$.

Since the first condition in the statement does not hold, we conclude that $x' \neq x^*$ but still $\mathsf{P}^{\widetilde{\Gamma}}(\alpha, x') = \mathsf{P}^{\widetilde{\Gamma}}(\alpha, x^*)$, in contradiction to the assumption that $\mathsf{P}^{\widetilde{\Gamma}}(\alpha, \cdot)$ defines a permutation.

We now show that the oracle $\widetilde{\Gamma} = (\widetilde{f}, \widetilde{\mathcal{O}}, \widetilde{\mathsf{Eval}})$ as above can be constructed. Recall that we assume that the both conditions of the statement of the claim do not hold, and therefore in particular it holds that $\mathsf{AugQ}(\Pi, \Gamma, x^*) \cap \mathsf{Partial}(\Gamma') \subseteq Q$.

The oracle $\widetilde{f}$. Note that for every $m \leq t$, the set of queries Q^* contains all the functions $\{f_m\}_{m \leq t}$ and thus agrees completely with f (i.e., also with f'). We therefore set $\widetilde{f}_m = f_m$.

For every $m > t$, we define the function $\widetilde{f}_m$ as follows. For every x such that $[f_m(x) = y'] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$, we set $\widetilde{f}_m(x) = y'$. For every $[f_m(x) = y] \in \mathsf{Partial}(\Gamma')$, we set $\widetilde{f}_m(x) = y$. Since $\mathsf{AugQ}(\Pi, \Gamma, x^*) \cap \mathsf{Partial}(\Gamma') \subseteq Q$, we have that there is no contradiction, i.e, there are no input x and outputs y, y' such that $y \neq y'$ and $[f_m(x) = y'] \in f'$ and $[f_m(x) = y] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$. For any other value $x \notin \mathsf{Partial}(\Gamma') \cap \mathsf{AugQ}(\Pi, \Gamma, x^*)$, we set $\widetilde{f}_m(x) = 0^m$.

Before we continue to define the oracle $\widetilde{\mathcal{O}}$, we first define some set of output values that $\widetilde{\mathcal{O}}$ will have to avoid. For every $m > t$, we define the set $\mathsf{avoid}\text{-}\mathcal{O}_m$ as

$$\mathsf{avoid}\text{-}\mathcal{O}_m \stackrel{\text{def}}{=} \left\{ \widehat{C} \in \{0,1\}^{10m} \;\middle|\; \exists\, [\mathsf{Eval}_m(\widehat{C}, \star) = \star] \in \mathsf{AugQ}(\Pi, \Gamma, x^*) \cup \mathsf{Partial}(\Gamma') \right\} .$$

The oracle $\widetilde{\mathcal{O}}$. The oracle is already defined for every $m \leq t$. For every $m > t$, we define the function $\widetilde{\mathcal{O}}_m$ as follows. For every $[\mathcal{O}_m(x) = y] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$, we set $\widetilde{\mathcal{O}}_m(x) = y$. Likewise, for every $[\mathcal{O}_m(x) = y] \in \mathsf{Partial}(\Gamma')$, we set $\widetilde{\mathcal{O}}_m(x) = y$. Since $\mathsf{AugQ}(\Pi, \Gamma, x^*) \cap \mathsf{Partial}(\Gamma') \subseteq Q$, we have that there is no contradiction, that is, there is no pre-image that has two possible outputs. Moreover, since fail does not occur, it holds that $\Gamma' \in \mathcal{T}\text{-}\mathsf{admissible}$, the two

functions $\mathcal{O}_m$ and (the partially defined function) $\mathcal{O}'_m$ do not evaluate to the same output, and so the partially defined function $\widetilde{\mathcal{O}}_m$ is injective. We continue to define $\widetilde{\mathcal{O}}_m$ on the additional values, such that $\mathcal{O}_m$ is injective and avoids the set $\mathsf{avoid}\text{-}\mathcal{O}_m$.

The oracle $\widetilde{\mathsf{Eval}}$. We define the oracle $\widetilde{\mathsf{Eval}}$ using the oracles $\widetilde{f}$ and $\widetilde{\mathcal{O}}$ exactly as the true oracle Eval is defined using the true oracles f and $\mathcal{O}$. We now show that $\widetilde{\mathsf{Eval}}$ is consistent with $\mathsf{AugQ}(\Pi, \Gamma, x^*)$ and $\mathsf{Partial}(\Gamma')$. That is, that every query $[\mathsf{Eval}_m(\star,\star)] \in \mathsf{AugQ}(\Pi, \Gamma, x^*) \cup \mathsf{Partial}(\Gamma')$ has the same answer with $\widetilde{\mathsf{Eval}}$, and therefore $\mathsf{P}^{\Gamma}(\alpha, x^*) = \mathsf{P}^{\widetilde{\Gamma}}(\alpha, x^*)$ and $\mathsf{P}^{\Gamma'}(x\alpha, x') = \mathsf{P}^{\widetilde{\Gamma}}(\alpha, x')$. We have:

1. Assume that there exists $[\mathsf{Eval}(\widehat{C}, a) = \beta] \in \mathsf{Eval}'$ for some $\beta \neq \bot$. Since the oracle $\mathsf{Partial}(\Gamma') = (f', \mathcal{O}', \mathsf{Eval}')$ is consistent (recall Definition 3.9), then there exists a query $\left[\mathcal{O}_m(C, r) = \widehat{C}\right] \in \mathsf{Partial}(\Gamma')$ and f' contains all the necessary queries/answers for the evaluation of $C^{f'}(a)$, and it also holds that $C^{f'}(a) = \beta$. However, since any $(f', \mathcal{O}')$-queries in $\mathsf{Partial}(\Gamma')$ has the exact same answer with $(\widetilde{f}, \widetilde{\mathcal{O}})$, it holds that $C^{\widetilde{f}}(a) = \beta$ and $\widetilde{\mathcal{O}}(C, r) = \widehat{C}$, and so, from the definition of $\widetilde{\mathsf{Eval}}$ it holds that $\widetilde{\mathsf{Eval}}(\widehat{C}, a) = \beta$ as well.
2. Assume that there exists $[\mathsf{Eval}(\widehat{C}, a) = \beta] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$ for some $\beta \neq \bot$. Since the event spoof does not occur, there exists a query $[\mathcal{O}(C, r) = \widehat{C}] \in \mathsf{AugQ}(\Pi, \Gamma, x^*)$ as well, and $\mathsf{AugQ}(\Pi, \Gamma, x^*)$ contains also all the f-queries necessary for the evaluation $C^f(a)$. Since these queries appear in $\mathsf{AugQ}(\Pi, \Gamma, x^*)$, it holds that $\widetilde{f}$ and $\widetilde{\mathcal{O}}$ agree on the same queries, and therefore $\widetilde{\mathsf{Eval}}(\widehat{C}, a) = \beta$, as well.
3. For every query $[\mathsf{Eval}(\widehat{C}, a) = \bot] \in \mathsf{Partial}(\Gamma') \cup \mathsf{AugQ}(\Pi, \Gamma, x^*)$ we show that $\widetilde{\mathsf{Eval}}(\widehat{C}, a) = \bot$ as well. Specifically, it suffices to show that there do not exist C and r for which $\widetilde{\mathcal{O}}(C, r) = \widehat{C}$. Assume towards a contradiction that there exist such C and r, then there is inconsistency only if $\widetilde{\mathcal{O}}(C, r) = \widehat{C}$ but $[\mathsf{Eval}(\widehat{C}, a) = \bot] \in \mathsf{Partial}(\Gamma') \cup \mathsf{AugQ}(\Pi, \Gamma, x^*)$. However, this cannot occur since the oracles $\mathcal{O}$ and $\mathcal{O}'$ do not contradict, and $\widetilde{\mathcal{O}}$ avoids all Eval-queries in both $\mathsf{Partial}(\Gamma')$ and $\mathsf{AugQ}(\Pi, \Gamma, x^*)$, since it avoids the set $\mathsf{avoid}\text{-}\mathcal{O}$.

This completes the proof of claim 3.11. ■

From the previous claim we conclude that:

$$\Pr_{\substack{\Gamma \leftarrow \mathfrak{S} \\ \alpha \leftarrow \mathsf{Gen}^{\Gamma}(1^n) \\ x^* \leftarrow \mathsf{Samp}^{\Gamma}(\alpha)}} \left[\mathcal{A}^{\Gamma}(\alpha, \mathsf{P}^{\Gamma}(\alpha, x^*)) = x^* \;\middle|\; \overline{\mathsf{fail}} \wedge \overline{\mathsf{spoof}} \right] = 1 \ .$$

Since $\Pr[\ \mathsf{fail}\ ] + \Pr[\ \mathsf{spoof}\ ] \leq 2^{-10}$, it holds that:

$$\Pr_{\substack{\Gamma \leftarrow \mathfrak{S} \\ \alpha \leftarrow \mathsf{Gen}^{\Gamma}(1^n) \\ x^* \leftarrow \mathsf{Samp}^{\Gamma}(\alpha)}} \left[\mathcal{A}^{\Gamma}(\alpha, \mathsf{P}^{\Gamma}(\alpha, x^*)) = x^* \right] \geq 1 - 2^{-10}.$$

Finally, we observe that $\mathcal{A}$ makes at most a polynomial number of oracle queries to Γ, and all other computations that are done by $\mathcal{A}$ can be done using a polynomial number of queries to a PSPACE-complete oracle (as in the work of Impagliazzo and Rudich [44]): In each iteration, sampling x' and $\mathsf{Partial}(\Gamma')$ can be done in polynomial space, requires access only to Q which is of polynomial size, and does not require access to Γ. ■

3.4 Proof of Theorem 3.3

Equipped with the proofs of Theorems 3.4 and 3.5, we are now ready to prove Theorem 3.3.

Proof of Theorem 3.3. Let $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P}, M, T_M, \epsilon_{M,1}, \epsilon_{M,2})$ be a fully black-box construction of a domain-invariant one-way permutation family from a one-way function f and an indistinguishability obfuscator $i\mathcal{O}$ for the class $\mathcal{C}$ of all oracle-aided polynomial-size circuits C^f (recall Definition 3.2). Theorem 3.4 guarantees the existence of an oracle-aided algorithm $\mathcal{A}$ that runs in polynomial time $T_{\mathcal{A}}(n)$ such that

$$\Pr\left[\mathcal{A}^{\mathsf{PSPACE},\Gamma}(\alpha, \mathsf{P}^{\Gamma}(\alpha, x)) = x\right] \geq \epsilon_{\mathcal{A}}(n)$$

for any $n \in \mathbb{N}$, where $\epsilon_{\mathcal{A}}(n) = 1 - 2^{-10}$, and the probability is taken over the choice of $\Gamma \leftarrow \mathfrak{S}$, $\alpha \leftarrow \mathsf{Gen}^{\Gamma}(1^n)$, $x \leftarrow \mathsf{Samp}^{\Gamma}(\alpha)$, and over the internal randomness of $\mathcal{A}$. Definition 3.1 then states that there are two possible cases to consider: $\mathcal{A}$ can be used either for inverting the one-way permutation f or for breaking the indistinguishability obfuscator $i\mathcal{O}$.

In the first case we obtain from Definition 3.1 that

$$\Pr\left[M^{\mathcal{A}^{\mathsf{PSPACE}},\Gamma}\left(f\left(x\right)\right) \in f^{-1}(f(x))\right] \geq \epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot \epsilon_{\mathcal{A}}^{-1}(n)\right) \cdot \epsilon_{M,2}(n)$$

for infinitely many values of $n \in \mathbb{N}$, where the probability is taken over the choice of $x \leftarrow \{0,1\}^n$ and over the internal randomness of M. 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 oracle-aided algorithm that has access to a PSPACE-complete oracle and to Γ, 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 $T_{\mathcal{A}}(\ell(n))$). In particular, viewing $M' \stackrel{\mathsf{def}}{=} M^{\mathcal{A}^{\mathsf{PSPACE}}}$ as a single oracle-aided algorithm that has oracle access to Γ, implies that M' is a q-query algorithm where $q(n) = T_{M^{\mathcal{A}}}(n)$.[9] Theorem 3.5 then implies that either $2^{n/4} \leq q(n)$ or $\epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot \epsilon_{\mathcal{A}}^{-1}(n)\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/2}$. In the first sub-case, noting that $\ell(n) \leq T_M(n)$, we obtain that

$$2^{n/4} \leq q(n) = T_{M^{\mathcal{A}}}(n) \leq T_M(n) \cdot T_{\mathcal{A}}(\ell(n)) \leq T_M(n) \cdot T_{\mathcal{A}}(T_M(n)).$$

[9] Recall that an algorithm that has oracle access to Γ is a q-query algorithm if it makes at most q queries to Γ, and each of its queries to Eval consists of a circuit of size at most q.

The running time $T_{\mathcal{A}}(n)$ of the adversary $\mathcal{A}$ (when given access to a PSPACE-complete oracle) 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 $\epsilon_{M,1}\left(T_{\mathcal{A}}(n)\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/2}$, and since $T_{\mathcal{A}}(n)$ is some fixed polynomial in n (and $\epsilon_{\mathcal{A}}(n)$ is a constant) we obtain that $\epsilon_{M,1}(n^c) \cdot \epsilon_{M,2}(n) \leq 2^{-n/2}$ for some constant $c > 1$.

In the second case we obtain from Definition 3.1 that

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

for infinitely many values of $n \in \mathbb{N}$, where $\Gamma \leftarrow \mathfrak{S}$. As in the first case, viewing $M' \stackrel{\mathsf{def}}{=} M^{\mathcal{A}^{\mathsf{PSPACE}}}$ as a single oracle-aided algorithm that has oracle access to Γ, implies that M' is a q-query algorithm where $q(n) = T_{M^{\mathcal{A}}}(n)$. Theorem 3.5 then implies that either $2^{n/4} \leq q(n)$ or $\epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot \epsilon^{-1}_{\mathcal{A}}(n)\right) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$. As in the first case, this implies that either $T_M(n) \geq 2^{\zeta n}$ for some constant $\zeta > 0$, or $\epsilon_{M,1}(n^c) \cdot \epsilon_{M,2}(n) \leq 2^{-n/4}$ for some constant $c > 1$. ■

4 Impossibility for Constructions Based on One-Way Functions

As discussed in Section 1.3, the known impossibility results for constructing one-way permutations based on one-way functions [58, 45, 51] fall short in two aspects. First, these results rule out constructions of a *single* one-way permutation, and do not rule out constructions of a one-way permutation *family*. Second, these results rule out constructions that are *domain invariant* (recall Definition 3.2), and do not rule out constructions that are *not* domain invariant (such as the construction of Bitansky et al. [13]).

In this section we resolve this surprising gap by ruling out *all* fully black-box constructions of one-way permutation *families* from one-way functions – even constructions that are *not* domain invariant. In what follows we first formally define this class of reductions, and then state and prove our result.

Definition 4.1 *A fully black-box construction of a one-way permutation family from a one-way function consists of a triplet of oracle-aided probabilistic polynomial-time algorithms* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$*, 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 conditions hold:*

- **Correctness:** *For any function* f *the triplet* $(\mathsf{Gen}, \mathsf{Samp}, \mathsf{P})$ *is a permutation family relative to* f *(as in Definition 2.1).*
- **Black-box proof of security:** *For any function* f*, for any 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*

$$\Pr\left[\mathcal{A}^f(\alpha, \mathsf{P}^f(\alpha, x)) = x\right] \geq \epsilon_{\mathcal{A}}(n)$$

for infinitely many values of $n \in \mathbb{N}$, where the probability is taken over the choice of $\alpha \leftarrow \mathsf{Gen}^f(1^n)$, $x \leftarrow \mathsf{Samp}^f(\alpha)$, and over the internal randomness of $\mathcal{A}$, then

$$\Pr\left[M^{f,\mathcal{A}}\left(f\left(x\right)\right) \in f^{-1}(f(x))\right] \geq \epsilon_{M,1}\left(T_{\mathcal{A}}(n) \cdot \epsilon_{\mathcal{A}}^{-1}(n)\right) \cdot \epsilon_{M,2}(n)$$

for infinitely many values of $n \in \mathbb{N}$, where the probability is taken over the choice of $x \leftarrow \{0,1\}^n$ and over the internal randomness of M.

The above definition clearly captures constructions that are not domain invariant. First, it allows the support of the distribution $\mathsf{Gen}^f(1^n)$ to depend on f. Second, for each permutation index α that is produced by $\mathsf{Gen}^f(1^n)$, it allows the domain of the permutation $\mathsf{P}^f(\alpha,\cdot)$ to depend on f. For this general class of reductions we prove the following theorem:

Theorem 4.2 *Let $(\mathsf{Gen},\mathsf{Samp},\mathsf{P},M,T_M,\epsilon_{M,1},\epsilon_{M,2})$ be a fully black-box construction of a one-way permutation family from a one-way function. 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. *$\epsilon_{M,1}(n^c) \cdot \epsilon_{M,2}(n) \leq 2^{-\beta n}$ for some constants $c > 1$ and $\beta > 0$ (i.e., the security loss is exponential).*[10]

Towards proving Theorem 4.2 we generalize the attack presented in Section 1.3 from inverting any *single* oracle-aided *domain-invariant* permutation to inverting any oracle-aided one-way permutation *family* – even such families that are not domain invariant. In the full version of this paper [4], we prove the following theorem:

Theorem 4.3 *Let $(\mathsf{Gen},\mathsf{Samp},\mathsf{P})$ be a triplet of oracle-aided probabilistic polynomial-time algorithms that is a permutation family relative to any oracle f. Then, there exists an oracle-aided algorithm $\mathcal{A}$ that makes a polynomial number of oracle queries such that for any function f it holds that*

$$\Pr\left[\mathcal{A}^f(\alpha,\mathsf{P}^f(\alpha,x)) = x\right] = 1$$

for any $n \in \mathbb{N}$, where the probability is taken over the choice of $\alpha \leftarrow \mathsf{Gen}^f(1^n)$ and $x \leftarrow \mathsf{Samp}^f(\alpha)$, and over the internal randomness of $\mathcal{A}$. Moreover, the algorithm $\mathcal{A}$ can be implemented in polynomial time given access to a PSPACE-complete oracle.

[10] In particular, if the adversary-dependent security loss $\epsilon_{M,1}(\cdot)$ is polynomial, then the adversary-independent security loss $\epsilon_{M,2}(\cdot)$ is exponential.

References

1. Ananth, P., Brakerski, Z., Segev, G., Vaikuntanathan, V.: From selective to adaptive security in functional encryption. In: Advances in Cryptology – CRYPTO '15. pp. 657–677 (2015)
2. Ananth, P., Jain, A.: Indistinguishability obfuscation from compact functional encryption. In: Advances in Cryptology – CRYPTO '15. pp. 308–326 (2015)
3. Asharov, G., Segev, G.: Limits on the power of indistinguishability obfuscation and functional encryption. To appear in *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science* (available at `https://eprint.iacr.org/2015/341.pdf`) (2015)
4. Asharov, G., Segev, G.: On constructing one-way permutations from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2015/752 (available at `http://eprint.iacr.org/2015/752.pdf`) (2015)
5. Baecher, P., Brzuska, C., Fischlin, M.: Notions of black-box reductions, revisited. In: Advances in Cryptology – ASIACRYPT '13. pp. 296–315 (2013)
6. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. Journal of the ACM 59(2), 6 (2012)
7. Barak, B., Mahmoody-Ghidary, M.: Merkle puzzles are optimal - An $O(n^2)$-query attack on any key exchange from a random oracle. In: Advances in Cryptology – CRYPTO '09. pp. 374–390 (2009)
8. Bellare, M., Stepanovs, I., Tessaro, S.: Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation. In: Advances in Cryptology – ASIACRYPT '14. pp. 102–121 (2014)
9. Bitansky, N., Canetti, R., Cohn, H., Goldwasser, S., Tauman Kalai, Y., Paneth, O., Rosen, A.: The impossibility of obfuscation with auxiliary input or a universal simulator. In: Advances in Cryptology – CRYPTO '14. pp. 71–89 (2014)
10. Bitansky, N., Canetti, R., Tauman Kalai, Y., Paneth, O.: On virtual grey box obfuscation for general circuits. In: Advances in Cryptology – CRYPTO '14. pp. 108–125 (2014)
11. Bitansky, N., Paneth, O.: ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 401–427 (2015)
12. Bitansky, N., Paneth, O., Rosen, A.: 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)
13. Bitansky, N., Paneth, O., Wichs, D.: Perfect structure on the edge of chaos. To appear in *Proceedings of the 13th Theory of Cryptography Conference* (2016)
14. Bitansky, N., Vaikuntanathan, V.: 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)
15. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo-random bits. SIAM Journal on Computing 13(4), 850–864 (1984)
16. Bogdanov, A., Brzuska, C.: On basing size-verifiable one-way functions on NP-hardness. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 1–6 (2015)

17. Boneh, D., Zhandry, M.: Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In: Advances in Cryptology – CRYPTO '14. pp. 480–499 (2014)
18. Brakerski, Z., Katz, J., Segev, G., Yerukhimovich, A.: Limits on the power of zero-knowledge proofs in cryptographic constructions. In: Proceedings of the 8th Theory of Cryptography Conference. pp. 559–578 (2011)
19. Brakerski, Z., Rothblum, G.N.: Virtual black-box obfuscation for all circuits via generic graded encoding. In: Proceedings of the 11th Theory of Cryptography Conference. pp. 1–25 (2014)
20. Brzuska, C., Farshim, P., Mittelbach, A.: Random-oracle uninstantiability from indistinguishability obfuscation. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 428–455 (2015)
21. Canetti, R., Goldwasser, S., Poburinnaya, O.: Adaptively secure two-party computation from indistinguishability obfuscation. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 557–585 (2015)
22. Canetti, R., Lin, H., Tessaro, S., Vaikuntanathan, V.: Obfuscation of probabilistic circuits and applications. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 468–497 (2015)
23. Canetti, R., Tauman Kalai, Y., Paneth, O.: On obfuscation with random oracles. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 456–467 (2015)
24. Chang, Y., Hsiao, C., Lu, C.: The impossibility of basing one-way permutations on central cryptographic primitives. Journal of Cryptology 19(1), 97–114 (2006)
25. Chung, K., Lin, H., Mahmoody, M., Pass, R.: On the power of nonuniformity in proofs of security. In: Proceedings of the 4th Innovations in Theoretical Computer Science Conference. pp. 389–400 (2013)
26. Chung, K., Lin, H., Pass, R.: Constant-round concurrent zero-knowledge from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2014/991 (2014)
27. Dachman-Soled, D., Katz, J., Rao, V.: Adaptively secure, universally composable, multiparty computation in constant rounds. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 586–613 (2015)
28. Dachman-Soled, D., Lindell, Y., Mahmoody, M., Malkin, T.: On the black-box complexity of optimally-fair coin tossing. In: Proceedings of the 8th Theory of Cryptography Conference. pp. 450–467 (2011)
29. Dachman-Soled, D., Mahmoody, M., Malkin, T.: Can optimally-fair coin tossing be based on one-way functions? In: Proceedings of the 11th Theory of Cryptography Conference. pp. 217–239 (2014)
30. Dachman-Soled, D., Mahmoody, M., Malkin, T.: Can optimally-fair coin tossing be based on one-way functions? In: Proceedings of the 11th Theory of Cryptography Conference. pp. 217–239 (2014)
31. Garg, S., Gentry, C., Halevi, S., Raykova, M.: Two-round secure MPC from indistinguishability obfuscation. In: Proceedings of the 11th Theory of Cryptography Conference. pp. 74–94 (2014)
32. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science. pp. 40–49 (2013)
33. Garg, S., Polychroniadou, A.: Two-round adaptively secure MPC from indistinguishability obfuscation. In: Proceedings of the 12th Theory of Cryptography Conference. pp. 614–637 (2015)

34. Gennaro, R., Gertner, Y., Katz, J., Trevisan, L.: Bounds on the efficiency of generic cryptographic constructions. SIAM Journal on Computing 35(1), 217–246 (2005)
35. Gertner, Y., Malkin, T., Myers, S.: Towards a separation of semantic and CCA security for public key encryption. In: Proceedings of the 4th Theory of Cryptography Conference. pp. 434–455 (2007)
36. Gertner, Y., Malkin, T., Reingold, O.: On the impossibility of basing trapdoor functions on trapdoor predicates. In: Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science. pp. 126–135 (2001)
37. Goldreich, O.: On security preserving reductions – revised terminology. Cryptology ePrint Archive, Report 2000/001 (2000)
38. Goldreich, O.: Foundations of Cryptography – Volume 1: Basic Techniques. Cambridge University Press (2001)
39. Goldwasser, S., Gordon, S.D., Goyal, V., Jain, A., Katz, J., Liu, F.H., Sahai, A., Shi, E., Zhou, H.S.: Multi-input functional encryption. In: Advances in Cryptology – EUROCRYPT '14. pp. 578–602 (2014)
40. Haitner, I., Hoch, J.J., Reingold, O., Segev, G.: Finding collisions in interactive protocols – Tight lower bounds on the round and communication complexities of statistically hiding commitments. SIAM Journal on Computing 44(1), 193–242 (2015)
41. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM Journal on Computing 28(4), 1364–1396 (1999)
42. Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In: Advances in Cryptology – EUROCRYPT '14. pp. 201–220 (2014)
43. Hsiao, C., Reyzin, L.: Finding collisions on a public road, or do secure hash functions need secret coins? In: Advances in Cryptology – CRYPTO '04. pp. 92–105 (2004)
44. Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: Proceedings of the 21st Annual ACM Symposium on Theory of Computing. pp. 44–61 (1989)
45. Kahn, J., Saks, M., Smyth, C.D.: The dual BKR inequality and Rudich's conjecture. Combinatorics, Probability & Computing 20(2), 257–266 (2011)
46. Komargodski, I., Moran, T., Naor, M., Pass, R., Rosen, A., Yogev, E.: One-way functions and (im)perfect obfuscation. In: Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science. pp. 374–383 (2014)
47. Luby, M.: Pseudorandomness and Cryptographic Applications. Princeton University Press (1996)
48. Mahmoody, M., Maji, H.K., Prabhakaran, M.: On the power of public-key encryption in secure computation. In: Proceedings of the 11th Theory of Cryptography Conference. pp. 240–264 (2014)
49. Mahmoody, M., Pass, R.: 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. pp. 701–718 (2012)
50. Matsuda, T.: On the impossibility of basing public-coin one-way permutations on trapdoor permutations. In: Proceedings of the 11th Theory of Cryptography Conference. pp. 265–290 (2014)
51. Matsuda, T., Matsuura, K.: On black-box separations among injective one-way functions. In: Proceedings of the 8th Theory of Cryptography Conference. pp. 597–614 (2011)

52. Naor, M., Yung, M.: Universal one-way hash functions and their cryptographic applications. In: Proceedings of the 21st Annual ACM Symposium on Theory of Computing. pp. 33–43 (1989)
53. Pass, R., Tseng, W.D., Venkitasubramaniam, M.: Towards non-black-box lower bounds in cryptography. In: Proceedings of the 8th Theory of Cryptography Conference. pp. 579–596 (2011)
54. Rabin, M.O.: Digitalized signatures and public-key functions as intractable as factorization. Technical report 212, Massachusetts Institute of Technology, Laboratory for Computer Science (1979)
55. Reingold, O., Trevisan, L., Vadhan, S.P.: Notions of reducibility between cryptographic primitives. In: Proceedings of the 1st Theory of Cryptography Conference. pp. 1–20 (2004)
56. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. Communication of the ACM 21(2), 120–126 (1978)
57. Rompel, J.: One-way functions are necessary and sufficient for secure signatures. In: Proceedings of the 22nd Annual ACM Symposium on Theory of Computing. pp. 387–394 (1990)
58. Rudich, S.: Limits on the Provable Consequences of One-way Functions. Ph.D. thesis, EECS Department, University of California, Berkeley (1988)
59. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: Deniable encryption, and more. In: Proceedings of the 46th Annual ACM Symposium on Theory of Computing. pp. 475–484 (2014)
60. Simon, D.R.: Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In: Advances in Cryptology – EUROCRYPT '98. pp. 334–345 (1998)
61. Waters, B.: A punctured programming approach to adaptively secure functional encryption. In: Advances in Cryptology – CRYPTO '15. pp. 678–697 (2015)
62. Wee, H.: One-way permutations, interactive hashing and statistically hiding commitments. In: Proceedings of the 4th Theory of Cryptography Conference. pp. 419–433 (2007)