

Indistinguishability Obfuscation versus Multi-Bit Point Obfuscation with Auxiliary Input

Christina Brzuska¹ and Arno Mittelbach²

¹ Tel Aviv University, Israel

² Darmstadt University of Technology, Germany

brzuska@post.tau.ac.il arno.mittelbach@cased.de

Abstract. In a recent celebrated breakthrough, Garg et al. (FOCS 2013) gave the first candidate for so-called indistinguishability obfuscation (iO) thereby reviving the interest in obfuscation for a general purpose. Since then, iO has been used to advance numerous sub-areas of cryptography. While indistinguishability obfuscation is a general purpose obfuscation scheme, several obfuscators for specific functionalities have been considered. In particular, special attention has been given to the obfuscation of so-called *point functions* that return zero everywhere, except for a single point x . A strong variant is point obfuscation with auxiliary input (AIPO), which allows an adversary to learn some non-trivial auxiliary information about the obfuscated point x (Goldwasser, Tauman-Kalai; FOCS, 2005).

Multi-bit point functions are a strengthening of point functions, where on x , the point function returns a string m instead of 1. Multi-bit point functions with auxiliary input (MB-AIPO) have been constructed from composable AIPO by Canetti and Dakdouk (Eurocrypt 2008) and have been used by Matsuda and Hanaoka (TCC 2014) to construct CCA-secure public-key encryption schemes and by Bitansky and Paneth (TCC 2012) to construct three-round weak zero-knowledge protocols for NP.

In this paper we present both positive and negative results. We show that if indistinguishability obfuscation exists, then MB-AIPO does not. Towards this goal, we build on techniques by Brzuska, Farshim and Mittelbach (Crypto 2014) who use indistinguishability obfuscation as a mean to attack a large class of assumptions from the Universal Computational Extractor framework (Bellare, Hoang and Keelveedhi; Crypto 2013). On the positive side we introduce a weak version of MB-AIPO which we deem to be outside the reach of our impossibility result. We build this weak version of MB-AIPO based on iO and AIPO and prove that it suffices to construct a public-key encryption scheme that is secure even if the adversary can learn an arbitrary leakage function of the secret key, as long as the secret key remains computationally hidden. Thereby, we strengthen a result by Canetti et al. (TCC 2010) that showed a similar connection in the symmetric-key setting.

Keywords Indistinguishability obfuscation, differing-inputs obfuscation, point function obfuscation, multi-bit point function obfuscation, auxiliary input obfuscation, leakage resilient PKE

1 Introduction

The obfuscation of a program should hide its inner workings while preserving the functionality of the program. Inspired by heuristic code-obfuscation techniques [28], obfuscation turned into a major research area of cryptography due to its manifold applications. The formal definition of Virtual Black-Box Obfuscation (VBB) demands that an obfuscated program is as good as a black-box that provides the same input-output behaviour as the program. Since the seminal paper of Barak et al. [5, 4], we know that this strong notion of obfuscation is generally not achievable.

Hence, research focused on special-purpose obfuscators and, in particular, there are various positive results for obfuscating so-called *point functions* p_x , that map all strings to 0, except for a single string x that they map to 1 [23, 27, 30, 52, 40, 24, 29, 26, 9, 13]. Other positive examples include obfuscating re-encryption [41] and encrypted signatures [38].

Point functions vs. point functions with multi-bit output. When considering point function obfuscation, we need to make a clear distinction between plain point functions such as p_x which map every input to 0 except for the single input x that is mapped to 1 and point functions with multi-bit output (MBPF) such as $p_{x,m}$ where input x is mapped to string m . Obfuscators for plain point functions are constructed in [23, 52, 40, 29].

Another important distinction is, whether the adversary is given some “leakage” about x , so-called *auxiliary information*, as introduced by Goldwasser and Tauman-Kalai [36]. We note that the obfuscator by Canetti [23] also allow for auxiliary information about the point x to leak and the obfuscator by Dodis et al. [29] allows for auxiliary information that hides the point statistically.

Although very similar, obfuscation schemes for MBPFs seem to be harder to construct than obfuscation schemes for plain point functions. Indeed, Canetti and Dakdouk initiated the study of obfuscation for MBPFs and showed that such obfuscation schemes are closely related to *composable* obfuscation schemes for plain point functions [24]. They show that obfuscators for MBPFs exist if composable obfuscators for plain point functions exist. Moreover, they show that composability is a non-trivial property. Both of these results carry over to obfuscation in the presence of auxiliary information, as long as the auxiliary information does not allow to recover the point. We refer to this type of auxiliary information as hard-to-invert or more specifically to computationally hard-to-invert.

Bitansky and Paneth [13] provide a clean treatment of auxiliary inputs and introduce the notion of *point obfuscation with auxiliary input* secure against unpredictable distributions (AIPO). Assuming composable AIPO they construct a three-round weak zero-knowledge protocol for \mathcal{NP} . Matsuda and Hanaoka [49] extend the notion of AIPO to the multi-bit point function case (MB-AIPO) and show how to use it to build CCA-secure public-key encryption. We adopt the notions AIPO and MB-AIPO in this paper.

Indistinguishability obfuscation. Simultaneously to constructing task-specific obfuscation schemes, the quest for general obfuscators continued, and in a celebrated breakthrough [32], Garg, Gentry, Halevi, Raykova, Sahai and Waters presented a candidate construction for indistinguishability obfuscation (iO). The notion of indistinguishability obfuscation is weaker than VBB-obfuscation and assures that, for any two circuits that compute the same function, their obfuscations are indistinguishable. As Goldwasser and Rothblum [37] establish, this seemingly weak notion of obfuscation is actually the *best possible* notion of obfuscation. And indeed, the work by Garg et al. [32] inspired simultaneous breakthroughs for hard problems in several sub-areas of cryptography [51, 16, 1, 31, 42, 15, 8, 22] such as functional encryption, deniable encryption, two-round secure multi-party computation, full-domain hash, poly-many hardcore bits for any one-way function and more.

Contribution. In this paper we give both positive and negative results. We show that the existence of indistinguishability obfuscation contradicts the existence of multi-bit point function obfuscation in the presence of computationally hard-to-invert auxiliary information (MB-AIPO), a notion which was built upon in [13, 49]. That is, if indistinguishability obfuscation exists, then MB-AIPO does not exist and some of the results in [13, 49] are based on a false assumption. (We discuss the precise implications shortly.) Or, equivalently, if MB-AIPO exists, then indistinguishability obfuscation does not exist and all candidate assumptions are false [32, 50, 34]. However, we do not have a candidate construction for MB-AIPO¹, but we do have a candidate construction for iO. Therefore, given the current advancements in the understanding of indistinguishability obfuscation—for example, Gentry et al. [34] show in a very recent work that iO can be based on the Multilinear Subgroup Elimination Assumption thereby giving the first construction based on an instance-independent assumption—we consider the existence of iO to be more likely.

In summary, we derive the following negative results.

Theorem (informal). *If indistinguishability obfuscation exists, then MB-AIPO and hence composable AIPO do not exist.*

Our proof is inspired by the result by Barak et al. [5, 4]. Technically, they show that multi-bit output point functions cannot be VBB-obfuscated when “coupled” with a particularly chosen second function. Let $p_{x,m}$ be a multi-bit output point function that maps all strings to 0, except for the single point x which the function maps to the string m . Now, the second function is a *test function* $\mathcal{T}_{x,m}$ that takes as input a circuit C and tests whether $C(x)$ is equal to m . Now, if an adversary is given access to two oracles that compute $p_{x,m}$ and $\mathcal{T}_{x',m'}$ then it cannot check whether the two functions “match”, i.e., whether $(x', m') = (x, m)$.

¹ Note that the construction by Canetti and Dakdouk [24] is from composable AIPO for which we do not have a candidate construction. The construction by Bitansky and Canetti [9, 10] achieves composable point obfuscation in the virtual grey-box setting (VGB) which implies MB-AIPO, but only for statistically hard-to-invert leakage [49].

In turn, when given a circuit C that computes $p_{x,m}$, the adversary can run $\mathcal{T}_{x,m}$ on C and simply check whether $\mathcal{T}_{x,m}(C)$ returns 1. Hence, the obfuscation of $p_{x,m}$ and the obfuscation of $\mathcal{T}_{x,m}$ leak more information than two oracles for $p_{x,m}$ and $\mathcal{T}_{x,m}$ thus establishing a counterexample for VBB obfuscation.

Although the starting point of Barak et al.’s result is a point function $p_{x,m}$, they actually construct an unobfuscatable function that is a combination of the point function $p_{x,m}$ together with test function $\mathcal{T}_{x,m}$ and thus their result is an impossibility result for general VBB obfuscation rather than an impossibility result for point function obfuscation.

In order to obtain a result for point function obfuscation based on the above idea, we proceed in two steps. Firstly, we think of the test circuit $\mathcal{T}_{x,m}$ as “auxiliary information” [36] about the point function $p_{x,m}$. Secondly, we do not use the “plain” test function $\mathcal{T}_{x,m}$ but rather, based on indistinguishability obfuscation, we construct an obfuscated circuit that approximates the behaviour of $\mathcal{T}_{x,m}$.

Matsuda and Hanoaka [49] introduce MB-AIPO as follows. A first stage of the adversary \mathcal{B}_1 defines a distribution over a point address x , a message m and auxiliary input z —we sometimes refer to the auxiliary input as “leakage”.

Now, a second stage of the adversary \mathcal{B}_2 gets the leakage z as well as an obfuscation of the point function $p_{x,m}$ or an obfuscation of the point function $p_{x,m'}$, where m' is drawn at random. The distinguisher \mathcal{B}_2 tries to guess which of the two it received.

A multi-bit point function obfuscator is called secure, if for all efficiently computable distributions² \mathcal{B}_1 , for the second stage of the adversary \mathcal{B}_2 , given z , obfuscations of $p_{x,m}$ and obfuscations of $p_{x,m'}$ are indistinguishable.

As such, the definition is not satisfiable, because \mathcal{B}_1 can leak the pair (x, m) so that \mathcal{B}_2 can check whether this pair “matches” the point function that \mathcal{B}_1 received. Hence, we additionally require that \mathcal{B}_1 be computationally *unpredictable*, that is, for all efficient predictors Pred , it holds that with high probability over $(z, x, m) \leftarrow \mathcal{B}_1$, given z , the algorithm Pred outputs x at most with negligible probability.

To recap, \mathcal{B}_1 outputs a point address x , a point value m and some leakage z such that z hides the value x . Then, the second stage of the adversary \mathcal{B}_2 receives z as well as an obfuscation of $p_{x,m}$ or an obfuscation of $p_{x,m'}$ and needs to distinguish between the two. See Definition 5 for a formal definition.

Hence, to attack MB-AIPO, we need to define an adversarial distribution \mathcal{B}_1 that is unpredictable and that returns some leakage z that allows \mathcal{B}_2 to distinguish between obfuscations of $p_{x,m}$ and obfuscations of $p_{x,m'}$. Our adversarial distribution \mathcal{B}_1 draws a random value x and a random value m . Moreover, as auxiliary information z , it will output a specially devised obfuscation that approximates the behaviour of the test function $\mathcal{T}_{x,m}$.

Given the circuit z and a multi-bit point function p , the second stage of the adversary \mathcal{B}_2 outputs whatever the circuit z outputs when run on p . It distinguishes successfully between an obfuscation p of the “matching” multi-bit

² We add the condition of unpredictability in the next paragraph.

point function $p_{x,m}$ and the obfuscation p of a non-matching multi-bit point function $p_{x,m'}$.

We now explain how adversary \mathcal{B}_1 constructs z . The hardness resides in constructing an obfuscation of the test function $\mathcal{T}_{x,m}$ such that indeed, x is unpredictable given the description of the obfuscated test function. Towards this goal, we build on techniques developed by Brzuska, Farshim and Mittelbach [20] who show a similar 1-out-of-2 result, namely that indistinguishability obfuscation and a large class of assumptions of the Universal Computational Extractor framework (UCE) [6] are mutually exclusive. We obfuscate the test function via indistinguishability obfuscation and prove that it is indistinguishable from an obfuscation of the zero circuit $\mathbf{0}$, the circuit that returns 0 on all inputs. As the zero circuit does not contain any information about x , indistinguishability obfuscation guarantees that likewise, an obfuscation of the test function $\mathcal{T}_{x,m}$ hides x computationally.

In detail, let y be the output of a pseudo-random generator G when applied to m . The circuit z is an indistinguishability obfuscation of the following circuit $C[x, y]$ with parameters x and y hard-coded. Circuit $C[x, y]$ gets as input a circuit p , runs p on x and checks whether $G(p(x))$ is equal to y . If yes, it outputs 1. Else, it outputs 0.

For simplicity, let us assume that the G is injective. Then, $C[x, y]$ behaves exactly like the test function $T_{x,m}$. Interestingly, and that is the key idea, we do not actually use m to compute the circuit $C[x, y]$, we only need $y = G(m)$. In particular, as G is a one-way function, y does not leak m . Moreover, as G is a pseudo-random generator, y does not even leak whether a pre-image m exists.

We will now use the PRG property to argue that an indistinguishability obfuscation of $C[x, y]$ does not leak anything about x . Namely, if y is in the image of the PRG, then $C[x, y]$ is equal to the test function $T_{x,m}$. In turn, when y is not in the image of the PRG, then $C[x, y]$ is the all-zero function. Due to the PRG security, these two distributions— $C[x, y]$ when y is drawn as an output from the G and $C[x, y]$ when y is drawn at random—are computationally indistinguishable. Moreover, when the PRG has enough stretch, then with overwhelming probability, a random y is not in the image of the PRG, and hence, with overwhelming probability over a random y , the circuit $C[x, y]$ is the all-zero circuit $\mathbf{0}$. For the two functionally equivalent circuits $C[x, y]$ and $\mathbf{0}$, it holds that $\text{iO}(C[x, y])$ is computationally indistinguishable from $\text{iO}(\mathbf{0})$. As $\mathbf{0}$ leaks nothing about x , we can argue that also $\text{iO}(C[x, y])$ leaks nothing about x and hence, x is unpredictable from the leakage of \mathcal{B}_1 as required by the definition of MB-AIPO.

We note that our usage of the PRG is somewhat similar to the use by Sahai and Waters in their construction of a CCA-secure PKE scheme from iO [51] as well as the range-extension of Matsuda and Hanaoka [49] of a multi-bit point function to obtain shorter point values and the range-extension of a UCE1-secure hash-function by Bellare et al. [7] used to strengthen the impossibility result by Brzuska et al. [20].

To recap, we use the pseudo-random generator to hide m , and we use the indistinguishability obfuscation to hide x . Note that unpredictability in the

MB-AIPO definition only requires that x is unpredictable from the leakage. Therefore, hiding m might seem unnecessary. Interestingly, it turns out that this is not merely an artefact of our proof. Namely, we define a strong notion of unpredictability where x needs to be unpredictable from the pair (z, m) , and we show that MB-AIPO can be achieved under this definition, assuming plain AIPO in conjunction with iO.

Indeed, our negative results do not carry over to the setting of obfuscating plain point functions in the presence of auxiliary information, that is, to plain AIPO (assuming they are not composable³). This is due to the fact that we cannot apply the PRG to a function that only outputs a single bit.

Analogously, it looks unlikely that the result of Barak et al. [5, 4] carries over to plain point functions, because it seems crucial that the point function $p_{x,m}$ has a multi-bit output m . Imagine that \mathcal{T}_x takes the circuit C as input and returns 1 if and only if $C(x) = 1$. Then, an adversary can perform binary search and recover x , even when only given access to \mathcal{T}_x and p_x as oracles.⁴ Hence, also their result does not carry over to standard point functions.

On the positive side, as hinted above, we show ways to work around our impossibility result. Firstly, note that Canetti et al. [26] introduce weaker versions of MB-AIPO that are not affected by our negative results. In particular, they use these weaker notions to build a symmetric-key encryption scheme that is secure in the presence of hard-to-invert leakage about the key. We strengthen their result insofar, as we present a notion that lies between their weaker versions of MB-AIPO and full MB-AIPO.

Our weak notion of MB-AIPO requires that the auxiliary information L computationally hides the point x even when given the corresponding point value m for some multi-bit point function $p_{x,m}$.

This definition circumvents our impossibility result because we cannot use the security of the PRG anymore. In the proof of the impossibility result, we used that the circuit $C[x, y]$ does not need m as a parameter and only needs $y = G(m)$. In the presence of the value m , the reduction to the PRG-security does not carry through.

This argument merely shows that our proof fails. However, we provide positive evidence for the new security notion. Assuming AIPO and iO, we give a construction that achieves MB-AIPO for strongly unpredictable distributions. We show that this weaker notion of MB-AIPO is useful for applications. Based on our weak MB-AIPO construction, we build a public-key encryption scheme which is leakage resilient in the presence of hard-to-invert leakage of the key. Previously, such a result was only known for symmetric-key encryption [26]. We next discuss existing notions of multi-bit point obfuscation.

Notions of Multi-Bit Point-Obfuscation. Lynn et al. [47] initiate the study of obfuscators for point functions with multi-bit output (MBPF) in the idealized

³ Canetti and Dakdouk [24] show that composable AIPO already implies MB-AIPO.

⁴ Access to the testing function \mathcal{T}_x suffices to recover x , even when not given access to p_x neither as a circuit nor as an oracle.

random oracle model (ROM) and give a construction of a VBB obfuscator in the ROM. Though they do not explicitly introduce auxiliary information, it is easily seen that their construction allows for computationally hard-to-invert auxiliary information. Canetti and Dakdouk [24] initiated the study of MBPF-obfuscators in the standard model and showed that these exist if so-called t -composable obfuscators exist for plain point functions. Building on these results Canetti and Bitansky [9, 10] show that the point obfuscator by Canetti [23] meets the requirements of a t -composable point function obfuscator down to a strong variant of the decisional Diffie–Hellman assumption (DDH), namely the t -strong vector DDH assumption. Note that the notion they achieve is the so-called notion of Virtual Grey-Box obfuscation (VGB)—the virtual grey box notion was introduced by Bitansky and Canetti [9, 10] and allows the simulator to run in unbounded time—and not the stronger notion of VBB obfuscation. In [26] Canetti et al. show that obfuscators for MBPFs are closely related to symmetric encryption and that obfuscators for MBPFs secure in the presence of (certain types of) auxiliary inputs imply the existence of (certain types of) leakage resilient symmetric encryption schemes. Bitansky and Paneth [13] introduce a clean treatment of a form of auxiliary information which hides the obfuscated point computationally (AIPO) and Matsuda and Hanaoka [49] extend their notion to multi-bit output functions which is also the notion considered in this paper (MB-AIPO). Using composable AIPOs Bitansky and Paneth construct a three-round weak zero-knowledge protocol for \mathcal{NP} based on composable AIPO [13] thereby circumventing a black-box impossibility result [35]. Matsuda and Hanaoka (MH, [49]) introduce also an average case variant of MB-AIPO and a more restricted version of MB-AIPO which requires the auxiliary input to statistically hide the obfuscated point. They further study the relation between these average case MB-AIPO notions and the worst-case notions of point obfuscation, that is, virtual black-box and virtual grey-box. MH show how to construct CCA secure public-key encryption schemes from an IND-CPA secure encryption scheme using MB-AIPO with computationally hard-to-invert auxiliary information, as well as, how to achieve CCA security starting from a CPA-secure lossy encryption scheme and using MB-AIPO with statistically hard-to-invert auxiliary information. In a very recent work, Canetti et al. [25] show how to build fuzzy extractors using t -composable point obfuscation secure in the presence of auxiliary information in the virtual grey-box setting. MH show that this form of point obfuscation implies MB-AIPO with respect to statistically hard-to-invert auxiliary information [49], it is, however, not known if it can be shown to also imply MB-AIPO with computationally hard-to-invert auxiliary information.

Our negative result shows that if indistinguishability obfuscation exists that MB-AIPO with computationally hard-to-invert auxiliary information does not exist. This applies to the first of the two constructions of CCA secure PKE schemes by Matsuda and Hanaoka [49] as well as to the construction of a three-round weak zero-knowledge protocol for \mathcal{NP} by Bitansky and Paneth [13].⁵ We

⁵ Bitanski and Paneth actually consider the stronger notion of composable AIPO which implies MB-AIPO. We also note that the construction of 3-message witness-hiding

leave as open problems, whether our negative results can be strengthened to encompass further uses of MBPF obfuscation or, whether the above constructions can be based on weaker notions of MBPF obfuscation not ruled out by our result.

Finally, we note that our result can be regarded as a random oracle uninstantiability result. One can show that the VBB obfuscator given by Lynn et al. [47] is a secure MB-AIPO in the random oracle model, even if hard-to-invert leakage is allowed. Our results shows that, if indistinguishability obfuscation exists, then there is no hash-function that instantiates the random oracle securely according to this notion of security.

Point obfuscation and indistinguishability obfuscation. For our positive result, a construction of weak MB-AIPO and subsequently a construction of a leakage resilient PKE scheme, we combine AIPOs and indistinguishability obfuscation. In a recent work Brzuska and Mittelbach (BM, [22]) show that combining these techniques allows to build powerful primitives and they give the first construction of a standard model hash function which is UCE secure for a non-trivial UCE notion which implies universal hardcore-functions and q -query correlated input secure hash functions. Furthermore, we note that our notion of weak MB-AIPO is inspired by the UCE notion introduced by BM: UCE security with respect to strongly unpredictable sources.

In a recent and independent work, Hofheinz constructs fully secure constrained pseudorandom functions [39] in the random oracle model. A constrained PRF allows for the generation of keys that enable the holder to evaluate the PRF on a set of points but not on all points, and various forms have been suggested [14, 17, 44]. In contrast to previous works Hofheinz uses point obfuscation and an extension he calls *extensible testers*—an extensible tester can be regarded as an obfuscation of a set of points Z which can be combined with a known set Z' into a tester for set $(Z \cup Z')$ —in conjunction with indistinguishability obfuscation to hide which points a given key allows to honestly evaluate. This allows him to achieve full security without relying on complexity leveraging which was used in previous constructions entailing a superpolynomial loss of security in the adaptive setting. We note that unlike this work (and the work by BM) Hofheinz relies on the simpler assumption of plain point obfuscation (that is, obfuscation without auxiliary inputs) and shows how to build extensible testers based on the DDH-based point obfuscator by Canetti [23].

Further 1-out-of-2 results. Indistinguishability obfuscation has led to many surprising breakthroughs in a number of sub-areas of cryptography [51, 16, 1, 31, 42, 15, 8, 22]. Interestingly, the existence of indistinguishability obfuscation collides with the existence of other desirable primitives. If indistinguishability obfuscation exists, then it draws a fine line between what is possible and what is impossible, e.g., MB-AIPO and iO are mutually exclusive, but weak MB-AIPO can be build from iO (and AIPO).

protocols from AIPO [13] as well as the construction of a CCA secure PKE scheme from a lossy encryption scheme and MB-AIPO with statistically hard-to-invert information [49] are not affected by our result.

If indistinguishability obfuscations does not exist, then 1-out-of-2 results are a promising way to prove such an impossibility result. In particular, it would be highly interesting to show a 1-out-of-2 result for iO and some other primitive for which we have a candidate construction, e.g., AIPO. Whether indistinguishability obfuscation exists or not, 1-out-of-2 results for iO help us explore the boundaries of what is possible. Either, they increase our understanding of iO , or they increase our understanding of other primitives.

Before our result, several 1-out-of-2 results have been established for iO . We already discussed the result by Brzuska et al. [20] who show that iO is mutually exclusive with a large class of assumptions from the UCE framework [6].

Interestingly, several notions of obfuscation are mutually exclusive with iO . Bitansky et al. [11] show that iO implies the non-existence of average-case virtual black-box obfuscation with auxiliary input (AI-VBB) for circuit families with super-polynomial pseudo-entropy. In particular, AI-VBB obfuscation is impossible for all pseudo-random function families. Moreover, they show that indistinguishability obfuscation implies the non-existence of average-case virtual black-box obfuscation with a universal simulator for circuit families with a superpolynomial amount of pseudo-entropy. Bitansky et al. [12] show that if indistinguishability obfuscation exists, then for every extractable one-way function family there is an (unbounded polynomial-length) auxiliary input distribution \mathcal{L} and an adversary \mathcal{A} such that all extractors fail for \mathcal{A} . Similar to our result for MB-AIPO, they embed an attack circuit into the auxiliary input. Boyle and Pass [18] strengthen this result under the assumption of differing-input obfuscation (diO). If diO exists, then the quantifiers can be reversed so that \mathcal{L} does not depend on the one-way function family.

Moreover, Bitansky et al. [12] show how to construct extractable one-way functions with *bounded* auxiliary input under relatively standard assumptions. Finally, Marcedone et al. [48], as well as Koppula et al. [46] show that if indistinguishability obfuscation exists, then IND-CPA-security of an encryption scheme does not imply its circular security, even if the cycles are of arbitrary polynomial-length.

On the plausibility of iO . Barak et al. [5, 4] introduce Indistinguishability Obfuscation as a notion of obfuscation that is not ruled out by their impossibility result for virtual black-box obfuscation. The amount and quality of positive results based on iO as well as the number of 1-out-of-2 results indicate that indeed, indistinguishability obfuscation is a strong assumption and Komargodski et al. [45] show that (even imperfect) indistinguishability obfuscation does not exist in Pessiland [43], a world where NP is hard but one-way functions do not exist. Their result does not carry over to a world where one-way functions exist.

Garg et al. [33] show that differing-inputs obfuscation—a stronger form of indistinguishably obfuscation that was also introduced in the seminal paper by Barak et al. [5, 4]—is mutually exclusive with some special-purpose obfuscator. As the particular special-purpose obfuscator that they consider seems to be a relatively mild assumption, we interpret their result as a conditional impossibility result for differing-inputs obfuscation. However, their result does not apply to

indistinguishability obfuscation. In particular, recent results show how to improve the assumptions that underly indistinguishability obfuscation [50, 19, 3, 2, 34] supporting its plausibility.

Auxiliary Input. Auxiliary input (AI) has been introduced by Goldwasser and Tauman-Kalai [36] and the specifics of how AI is modeled are very important when it comes to the (im)possibility of notions of obfuscation. Notably, for extractable one-way functions, the aforementioned results by Bitansky et al. [12] show that, assuming iO, this notion of security is impossible under unbounded AI, but possible when the length of the AI is bounded by a fixed polynomial that is known a priori. Potentially, bounded AI—for example, if the amount of AI is restricted to be less than the size of an MB-AIPO—could also be used to circumvent our iO-based impossibility result while preserving a reasonably wide range of applications.

Moreover, one can consider independent AI rather than dependent AI, which would also help to circumvent our impossibility result. However, AI is usually useful for composition where partial information about the obfuscated circuit/point is leaked to the outside and thus, dependent AI is often quite powerful in applications. However, even security under independent AI is non-trivial to achieve. Assuming iO, Bitansky et al. [11] show that a large class of functions cannot be VBB-obfuscated in the presence of independent AI.

A further possibility to circumvent our impossibility result is to consider a statistical notion of unpredictability rather than computational unpredictability. Statistical unpredictability has already proved useful for the construction of q -query secure correlation-secure hash functions [22] and CCA secure PKE schemes [49].

While in the VBB-setting AI is a strong notion that corresponds to the existence of a universal simulator [11], in the VGB-setting AI is trivial. That is, it is equivalent whether one considers VGB security with AI or without AI. The reason is that the VGB simulator is unbounded and hence able to compute the best AI itself [9]. Secure AIPOs in the VGB-setting imply AIPOs with statistically hard-to-invert leakage [49]. Our result does not rule out composable AIPOs in the VGB-setting, and indeed, this assumption has been used very recently by Canetti et al. [25] to build computationally secure fuzzy extractors that work for classes of sources that have more errors than entropy.

In light of the subtle modeling of AI, it remains to investigate whether those results in [13] and [49] that use an assumption which is mutually exclusive with iO can be based on an alternative assumption that is compatible with iO. Towards this goal, one might consider our weakened notion of MB-AIPO or model AI in a way that circumvents our result. Finally, it would then be interesting to come up with candidate assumptions for such a notion of security.

Conclusion and future work. We show that indistinguishability obfuscation and MB-AIPO—that is, MB-AIPO as used in [13, 49] and with computationally hard-to-invert auxiliary information—are mutually exclusive. It remains to investigate whether the positive results in [13, 49] can be salvaged through weaker notions of

MB-AIPO or, perhaps, when combining AIPO and iO in a similar way as we do in the full version [21] to receive our positive result for weak MB-AIPO. We note, however, that, at a first glance, it is not straightforward to base the applications in [13, 49] on our weakened notion of MB-AIPO.⁶

On the other hand, one might ask whether our negative result can be extended to showing that AIPO and iO are mutually exclusive. Currently, we do not know whether this is possible. We consider such a result to be a highly interesting finding and suspect that it would require different techniques than the ones we use. Our result implies directly that differing-inputs obfuscation (diO) and MB-AIPO are mutually exclusive. Perhaps, using different techniques, one might be able to first show that diO and AIPO are mutually exclusive, for example, by showing that we can instantiate the special-purpose obfuscator by Garg et al. [33] using AIPO.

We hope that our work sparks further interest in studying the connections between iO/diO on the one hand and notions of (multi-bit) point obfuscation on the other hand. More generally, we believe that it is an interesting question to identify notions of security that collide with indistinguishability obfuscation and we expect more results of that flavor in the future.

Full version. Due to space restrictions, this version should be regarded as an extended abstract as we defer many details and all proofs to the full version [21]. In the remainder of this extended abstract we present our main impossibility result and give some intuition for the underlying proof. For details as well as for our positive results (our weak MB-AIPO notion and the construction of a leakage resilient PKE scheme) we refer to the full version [21].

2 Preliminaries

Indistinguishability obfuscation. Virtual black-box (VBB) obfuscation [5, 36, 4] requires that for any PPT adversary given the code of some functionality (and some auxiliary input) there exists a PPT simulator that given only black-box access to the functionality (and as input the same auxiliary input) produces a computationally indistinguishable distribution. While VBB obfuscation provably does not exist for all circuits [5, 4], weaker notions such as *indistinguishability obfuscation* may well do. An indistinguishability obfuscation (iO) scheme, on the other hand, only ensures that the obfuscations of any two functionally equivalent circuits are computationally indistinguishable. Indistinguishability obfuscation

⁶ Note that [13] use composable point functions which is a stronger security notion than MB-AIPO for showing the existence of 3-round protocols that are weakly zero-knowledge. Also note, that their second result, a 3-round witness-hiding protocol, is not affected by our result. Likewise, our result only affects the CCA-encryption scheme in [49] that is based on CPA-security and MB-AIPO. They also build a CCA-secure encryption scheme based on *lossy* IND-CPA secure encryption and MB-AIPO with statistically hard-to-invert auxiliary input. The latter result is not affected by our result.

was originally proposed by Barak et al. [5] as a potential weakening of virtual-black-box obfuscation. We recall the definition from [32].

Definition 1. A PPT algorithm iO is called an indistinguishability obfuscator for a circuit ensemble $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ if the following conditions are satisfied:

- **Correctness.** For all security parameters $\lambda \in \mathbb{N}$, for all $C \in \mathcal{C}_\lambda$, and for all inputs x we have that $\Pr[C'(x) = C(x) : C' \leftarrow_{\$} \text{iO}(1^\lambda, C)] = 1$.
- **Security.** For any PPT distinguisher \mathcal{D} , for all pairs of circuits $C_0, C_1 \in \mathcal{C}_\lambda$ such that $C_0(x) = C_1(x)$ on all inputs x the following distinguishing advantage is negligible:

$$|\Pr[\mathcal{D}(1^\lambda, \text{iO}(1^\lambda, C_1)) = 1] - \Pr[\mathcal{D}(1^\lambda, \text{iO}(1^\lambda, C_0)) = 1]| \leq \text{negl}(\lambda).$$

Differing-inputs obfuscation. Differing-inputs obfuscation is closely related to indistinguishability obfuscation and also goes back to the seminal paper of Barak et al. [5, 4]. Building on a theorem by Boyle, Chung and Pass [16], we are able to avoid diO as an assumption and only use it as an intermediary concept in our proof. We refer for details to the full version of this work [21].

Point obfuscation. Besides the general purpose indistinguishability obfuscator we consider obfuscators for the specific class of so-called point functions. A point function p_x for some value $x \in \{0, 1\}^*$ is defined as outputting \perp on all inputs except for x where it outputs 1. In this paper, we consider a variant of point function obfuscators under auxiliary input which was first formalized by Canetti [23]. We here give the definition from [13] presented in a game based formulation. The first definition formalizes unpredictable distributions which are in turn used to define obfuscators for point functions.

Definition 2 (Unpredictable distribution). A distribution ensemble $\mathcal{D} = \{D_\lambda = (Z_\lambda, X_\lambda)\}_{\lambda \in \mathbb{N}}$, on pairs of strings is unpredictable if no poly-size (non-uniform) circuit can predict X_λ from Z_λ . That is, for every poly-size circuit sequence $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ and for all large enough λ :

$$\Pr_{(z,x) \leftarrow_{\$} D_\lambda}[C_\lambda(z) = x] \leq \text{negl}(\lambda)$$

Definition 3 (Auxiliary input point obfuscation for unpredictable distributions (AIPO)). A PPT algorithm AIPO is a point obfuscator for unpredictable distributions if it satisfies the functionality and polynomial slowdown requirements as in VBB-obfuscation [5, 4], and the following secrecy property: for any (efficiently sampleable) unpredictable distribution \mathcal{B}_1 over $\{0, 1\}^{\text{poly}(\lambda)} \times \{0, 1\}^\lambda$ it holds for any PPT algorithm \mathcal{B}_2 that the probability that the following experiment outputs true for $(\mathcal{B}_1, \mathcal{B}_2)$ is negligibly close to $\frac{1}{2}$:

```

b  $\leftarrow_{\$}$   $\{0, 1\}$ 
 $(z, x_0) \leftarrow_{\$}$   $\mathcal{B}_1(1^\lambda)$ 
 $x_1 \leftarrow_{\$}$   $\{0, 1\}^\lambda$ 
 $p \leftarrow_{\$}$  AIPO( $x_b$ )
 $b' \leftarrow_{\$}$   $\mathcal{B}_2(1^\lambda, p, z)$ 
return  $b = b'$ 

```

The probability is over the coins of adversary $(\mathcal{B}_1, \mathcal{B}_2)$, the coins of AIPO and the choices of x_1 and b .

Obfuscation for point functions with multi-bit output. While point functions only return a single bit, a point function with multi-bit output (MBPF) $p_{x,m}$ for values $x, m \in \{0, 1\}^*$ is defined as \perp on any input except for input x which is mapped to m . For an MBPF $p_{x,m}$ we call x the point address and m the point value. Similar to AIPO we can define MB-AIPO via an unpredictable distribution—the notion was introduced by Matsuda and Hanaoka [49] in an average case formulation called AIND- δ -cPUAI—where the distribution outputs a tuple (x, m) (defining a point function $p_{x,m}$) together with auxiliary information z . We require that it be computationally infeasible to recover the point address x given auxiliary information z . Thus, in the MBPF setting we define the unpredictable distribution as $\mathcal{D} = \{D_\lambda = (Z_\lambda, X_\lambda, M_\lambda)\}_{\lambda \in \mathbb{N}}$ but still require that the point address x remains hidden given the auxiliary input. An MB-AIPO assures that the obfuscation of $p_{x,m}$ is indistinguishable from an obfuscation with a changed point value m' that is chosen uniformly at random, which captures that the obfuscation does not reveal any information about the point value m .

Definition 4 (Unpredictable distribution). A distribution ensemble $\mathcal{D} = \{D_\lambda = (Z_\lambda, X_\lambda, M_\lambda)\}_{\lambda \in \mathbb{N}}$, on triples of strings is unpredictable if no poly-size (non-uniform) circuit can predict X_λ from Z_λ . That is, for every poly-size circuit sequence $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ and for all large enough λ :

$$\Pr_{(z,x,m) \leftarrow_{\$} D_\lambda} [C_\lambda(z) = x] \leq \text{negl}(\lambda)$$

Definition 5 (Auxiliary input point obfuscation for unpredictable distributions (MB-AIPO)). A PPT algorithm MB-AIPO is a multi-bit point obfuscator for unpredictable distributions if it satisfies the functionality and polynomial slowdown requirements as in VBB-obfuscation [5, 4], and the following secrecy property: for any (efficiently sampleable) unpredictable distribution \mathcal{B}_1 over $\{0, 1\}^{\text{poly}(\lambda)} \times \{0, 1\}^\lambda \times \{0, 1\}^{\text{poly}(\lambda)}$ it holds for any PPT algorithm \mathcal{B}_2 that the probability that the following experiment outputs true for $(\mathcal{B}_1, \mathcal{B}_2)$ is negligibly close to $\frac{1}{2}$:

```

b  $\leftarrow_{\$}$   $\{0, 1\}$ 
 $(z, x, m_0) \leftarrow_{\$}$   $\mathcal{B}_1(1^\lambda)$ 
 $m_1 \leftarrow_{\$}$   $\{0, 1\}^\lambda$ 
 $p \leftarrow_{\$}$  MB-AIPO( $x, m_b$ )
 $b' \leftarrow_{\$}$   $\mathcal{B}_2(1^\lambda, p, z)$ 
return  $b = b'$ 

```

The probability is over the coins of adversary $(\mathcal{B}_1, \mathcal{B}_2)$, the coins of AIPO and the choices of x, m_0, m_1 and b .

We note that also different definitional choices are possible and we discuss various choices in the full version of this work [21]

Average-case point obfuscation and statistical unpredictability. The above notions for point obfuscation are for arbitrary high-entropy distributions over the point address. Instead, we can consider a weaker variant where the point address is sampled according to the uniform distribution. Indeed, Matsuda and Hanaoka [49] recently presented constructions of CCA-secure public-key encryption schemes based on this version of point obfuscation. They call AIPO with arbitrary high-entropy samplers a worst-case notion, and AIPO with the uniform distribution an average-case notion and denote it by $\text{AIND-}\delta\text{-cPUAI}$. Our impossibility result also applies to $\text{AIND-}\delta\text{-cPUAI}$ which we refer to as *average case MB-AIPO*.

A second avenue to weaken the security requirements of point obfuscators is to require that the auxiliary input needs to hide the point address statistically. We call unpredictable distributions for which this is the case *statistically unpredictable*. Our impossibility result does not carry over to this notion.

3 IO Implies the Impossibility of MB-AIPO

In the following we present our negative result, namely that indistinguishability obfuscation and multi-bit point function obfuscation in the presence of auxiliary information (MB-AIPO) are mutually exclusive. This holds for MB-AIPO as defined in Definition 5 as well as for the two alternative definitions discussed below the definition. We discuss implications of our result in Section 3.2.

3.1 IO and MB-AIPO are Mutually Exclusive

Multi-bit point obfuscation with auxiliary inputs is a powerful primitive and has, for example, been used to construct CCA-secure encryption schemes [49] and to circumvent black-box impossibility results for three-round weak zero-knowledge protocols for \mathcal{NP} [13]. Our following result says that, if indistinguishability obfuscation and pseudo-random generators exist, then MB-AIPOs (as defined in Definition 5) cannot exist. The result remains valid even if we consider average case MB-AIPOs (where point address x is chosen uniformly at random). Technically our result builds on techniques used by Brzuska, Farshim and Mittelbach (BFM; [20]). BFM show a similar 1-out-of-2 result, namely that if indistinguishability obfuscation exists, then certain kinds of UCE-secure hash functions—a hash function security notion recently introduced in [6]—cannot exist [20]. In the UCE-framework, a hash function \mathbf{H} gets a hash key \mathbf{hk} and an input x and outputs y . BFM obfuscate the circuit $(\mathbf{H}(\cdot, x) = y)$, that given a hash-key \mathbf{hk} checks whether \mathbf{hk} “matches” the pair (x, y) , that is, whether $\mathbf{H}(\mathbf{hk}, \cdot)$ maps x to y . They show that, if $|\mathbf{hk}| < 2|y|$, then it is likely (in the corresponding experiment) that the circuit is the $\mathbf{0}$ -circuit that outputs 0 on all inputs and hence, the indistinguishability obfuscation of this circuit does not leak x .

We will use a similar technique to hide the point address. In order to break AIPO with indistinguishability obfuscation, we need to show that, given the auxiliary input, it is hard to recover the point address, but that, given the auxiliary input and the point function, one can distinguish. Similarly, for UCEs,

one needs to show that, given some leakage about x and y , it is hard to recover x , but that, given the leakage and the hash-key hk , one can distinguish whether y was generated by applying $H(\text{hk}, \cdot)$ to x or whether y was drawn at random.

Showing that an indistinguishability obfuscation hides a certain value is usually the crux in proofs involving iO . For this, we construct a new technique which may be of independent interest and which we discuss in further detail in the full version [21].

Theorem 1. *If indistinguishability obfuscation exists for all circuits in \mathcal{P}/poly , then average-case obfuscation for multi-bit point functions secure under auxiliary input (MB-AIPO) does not exist.*

This theorem applies to the average-case version where the point address is sampled uniformly, because our adversary samples both, x and m uniformly at random. It also applies to other variants of the MB-AIPO definition which we discuss in the full version of this work [21].

To prove Theorem 1 we use indistinguishability obfuscation to construct an unpredictable distribution \mathcal{B}_1 together with an adversary \mathcal{B}_2 that, given leakage from the unpredictable distribution can distinguish between point obfuscations from \mathcal{B}_1 and point obfuscations from the uniform distribution.

We first give the unpredictable distribution \mathcal{B}_1 which takes as input the security parameter 1^λ and outputs two values x, m together with some auxiliary information (resp. leakage) z . Here leakage z will be the indistinguishability obfuscation of a predicate circuit that takes as input a description of a circuit C , evaluates the circuit on a hard-coded value x , runs the result through a pseudo-random generator G and finally compares this result with some hard-coded value y . That is, we consider the circuit

$$C[x, y](\cdot) := \text{iO} \left(G(\text{uC}(\cdot, x)) = y \right),$$

where uC denotes a universal circuit taking as input a circuit description C of a fixed length and a value x and which outputs $C(x)$. This use of a PRG allows us later to argue that if value y is chosen uniformly at random that with high probability it falls outside the image of the PRG and thus the circuit is 0 on all inputs, that is, it implements the zero-circuit $\mathbf{0}$.

We next formally define the unpredictable distribution. For this let n and ℓ be two polynomials and let $G : \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{2n(\lambda)}$ be a pseudo-random generator with stretch 2. Note that we do not need to additionally assume the existence of PRGs as AIPOs (and in particular MB-AIPOs) already imply one-way functions.⁷ Let, furthermore, $\text{uC}(\cdot, x)$ be a universal circuit that on input a description of a circuit C and value x outputs $C(x)$. Adversary \mathcal{B}_1 computes an

⁷ Canetti et al. [26] show that multi-bit point function obfuscation is tightly related to symmetric encryption and that MB-AIPO implies the existence of (leakage-resilient) IND-CPA symmetric encryption schemes.

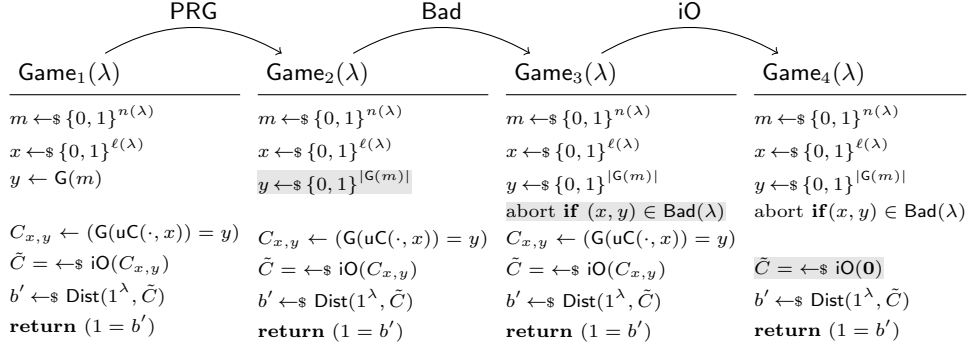


Fig. 1. The hybrids for the proof of Theorem 1. We have highlighted the changes between the games with a light-grey background.

unpredictable distribution over (z, x, m) as follows:

$$\begin{aligned}
& m \leftarrow_{\$} \{0, 1\}^{n(\lambda)} \\
& y \leftarrow G(m) \\
& x \leftarrow_{\$} \{0, 1\}^{\ell(\lambda)} \\
& z \leftarrow_{\$} \text{iO} \left(G(\text{uC}(\cdot, x)) = y \right) \\
& \mathbf{output:} (z, x, m)
\end{aligned}$$

We now present the adversary \mathcal{B}_2 that, given the leakage z from \mathcal{B}_1 , breaks the security of the multi-bit point obfuscator. We then argue that \mathcal{B}_1 , indeed, implements an unpredictable distribution. Adversary \mathcal{B}_2 gets values p and z as input, where p is either a point obfuscation of $p_{x,m}$ sampled according to \mathcal{B}_1 or an obfuscation for $p_{x,u}$ for a uniformly random value u . Adversary \mathcal{B}_2 computes $z(p)$ and outputs the result. If p is an obfuscation of $p_{x,m}$, then \mathcal{B}_2 computes the predicate function

$$G(p_{x,m}(x)) = y$$

where y is computed as $G(m)$ and outputs 1. In turn, if p is an obfuscation of $p_{x,u}$, then with overwhelming probability over the choice of u , adversary \mathcal{B}_2 returns 0. Thus, $(\mathcal{B}_1, \mathcal{B}_2)$ is a successful pair of adversaries. To prove that $(\mathcal{B}_1, \mathcal{B}_2)$ is also a valid pair of adversaries, we need to show that \mathcal{B}_1 is an unpredictable distribution. Under the assumption of indistinguishability obfuscation, the leakage computed by \mathcal{B}_1 is indistinguishable from an obfuscated zero circuit $\mathbf{0}$, the circuit that returns 0 on all inputs and which is padded to the same length as the (unobfuscated) leaked circuit, that is, the circuit $(G(\text{uC}(\cdot, x)) = y)$. As the zero circuit does not leak any information about y , the leakage is unpredictable. Formally we prove the unpredictability of \mathcal{B}_1 via a sequence of four hybrids depicted in Figure 1. We defer a formal proof to the full version [21].

3.2 Implications

Average case MB-AIPO is a relaxed notion of virtual-black-box point obfuscation in the presence of auxiliary input and in particular implied by it [49]. Consequently

our impossibility result also shows that VBB obfuscation of multi-bit point functions secure in the presence of auxiliary input cannot exist if indistinguishability obfuscation exist:

Corollary 1. *If indistinguishability obfuscation exists, then VBB multi-bit point obfuscation secure with auxiliary input does not exist.*

We note that VBB multi-bit point obfuscation is also often referred to as *Digital Lockers*. Canetti and Dakdouk [24] study the composition of point function obfuscation and show that composable AIPO implies the existence of composable MB-AIPO. And hence, applying our result we get the following corollary.

Corollary 2. *If indistinguishability obfuscation exists, then composable AIPO does not exist.*

Several results have been based on the existence of MB-AIPO (or composable AIPO). Matsuda and Hanaoka give a CCA secure public-key encryption scheme based on MB-AIPO [49] and Bitansky and Paneth give a three-round weak zero-knowledge protocol for \mathcal{NP} based on composable AIPO [13].⁸ In the full version [21] we present a weakened notion of MB-AIPO that we deem to fall outside our impossibility result. It is not clear whether this weaker notion suffices for the applications in [13, 49] and such a proof is not straightforward, so it remains to study whether one could use other weak variants of MB-AIPO.

A random oracle uninstantiability. Lynn et al. [47] construct VBB obfuscators for multi-bit point functions in the idealized random oracle model and their result can easily be seen to encompass auxiliary information. Thus, assuming iO exists our result rules out the existence of a standard model hash function that can instantiate the random oracle in their construction.

Corollary 3. *If indistinguishability obfuscation exists, then the multi-bit output point function obfuscator by Lynn et al. [47] cannot be instantiated in the standard model so that it achieves VBB security with auxiliary input.*

Acknowledgments

We thank the Asiacrypt 2014 reviewers for the many constructive comments. We especially thank Paul Baecher, Mihir Bellare, Nir Bitansky, Victoria Fehr, Peter Gazi, Dennis Hofheinz, Giorgia Azzurra Marson and Alon Rosen for many helpful comments and discussions throughout the various stages of this work. Christina Brzuska was supported by the Israel Science Foundation (grant 1076/11 and 1155/11), the Israel Ministry of Science and Technology grant 3-9094), and the German-Israeli Foundation for Scientific Research and Development (grant 1152/2011). Arno Mittelbach was supported by CASED (www.cased.de) and the German Research Foundation (DFG) SPP 1736.

⁸ We note that the construction of 3-message witness-hiding protocols from AIPO [13] as well as the construction of a CCA secure PKE scheme from lossy encryption schemes and MB-AIPO with statistically hard-to-invert information [49] are not affected by our result.

References

1. Ananth, P., Boneh, D., Garg, S., Sahai, A., Zhandry, M.: Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689 (2013), <http://eprint.iacr.org/2013/689>
2. Ananth, P., Gupta, D., Ishai, Y., Sahai, A.: Optimizing obfuscation: Avoiding barrington’s theorem. Cryptology ePrint Archive, Report 2014/222 (2014), <http://eprint.iacr.org/>
3. Barak, B., Garg, S., Kalai, Y.T., Paneth, O., Sahai, A.: Protecting obfuscation against algebraic attacks. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 221–238. Springer (May 2014)
4. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S., Yang, K.: On the (im)possibility of obfuscating programs. J. ACM 59(2), 6:1–6:48 (May 2012), <http://doi.acm.org/10.1145/2160158.2160159>
5. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer (Aug 2001)
6. Bellare, M., Hoang, V.T., Keelveedhi, S.: Instantiating random oracles via UCEs. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 398–415. Springer (Aug 2013)
7. Bellare, M., Hoang, V.T., Keelveedhi, S.: Personal communication (Sep, 2013)
8. Bellare, M., Stepanovs, I., Tessaro, S.: Poly-many hardcore bits for any one-way function. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. pp. ??–?? LNCS, Springer, Berlin, Germany, Kaohsiung, Taiwan (Dec 7–11, 2014)
9. Bitansky, N., Canetti, R.: On strong simulation and composable point obfuscation. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 520–537. Springer (Aug 2010)
10. Bitansky, N., Canetti, R.: On strong simulation and composable point obfuscation. J. Cryptology 27(2), 317–357 (2014)
11. Bitansky, N., Canetti, R., Cohn, H., Goldwasser, S., Kalai, Y.T., Paneth, O., Rosen, A.: The impossibility of obfuscation with auxiliary input or a universal simulator. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 71–89. Springer (Aug 2014)
12. Bitansky, N., Canetti, R., Paneth, O., Rosen, A.: On the existence of extractable one-way functions. In: Shmoys, D.B. (ed.) 46th ACM STOC. pp. 505–514. ACM Press (May / Jun 2014)
13. Bitansky, N., Paneth, O.: Point obfuscation and 3-round zero-knowledge. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 190–208. Springer (Mar 2012)
14. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 280–300. Springer (Dec 2013)
15. Boneh, D., Zhandry, M.: Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 480–499. Springer (Aug 2014)
16. Boyle, E., Chung, K.M., Pass, R.: On extractability obfuscation. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 52–73. Springer (Feb 2014)
17. Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer (Mar 2014)

18. Boyle, E., Pass, R.: Limits of extractability assumptions with distributional auxiliary input. *Cryptology ePrint Archive*, Report 2013/703 (2013), <http://eprint.iacr.org/2013/703>
19. Brakerski, Z., Rothblum, G.N.: Virtual black-box obfuscation for all circuits via generic graded encoding. In: Lindell, Y. (ed.) *TCC 2014*. LNCS, vol. 8349, pp. 1–25. Springer (Feb 2014)
20. Brzuska, C., Farshim, P., Mittelbach, A.: Indistinguishability obfuscation and UCEs: The case of computationally unpredictable sources. In: Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part I*. LNCS, vol. 8616, pp. 188–205. Springer (Aug 2014)
21. Brzuska, C., Mittelbach, A.: Indistinguishability obfuscation versus multi-bit point obfuscation with auxiliary input. *Cryptology ePrint Archive*, Report 2014/405 (2014), <http://eprint.iacr.org/>
22. Brzuska, C., Mittelbach, A.: Using indistinguishability obfuscation via uces. In: Sarkar, P., Iwata, T. (eds.) *ASIACRYPT 2014*. pp. ??–?? LNCS, Springer, Berlin, Germany, Kaohsiung, Taiwan (Dec 7–11, 2014)
23. Canetti, R.: Towards realizing random oracles: Hash functions that hide all partial information. In: Kaliski Jr., B.S. (ed.) *CRYPTO’97*. LNCS, vol. 1294, pp. 455–469. Springer (Aug 1997)
24. Canetti, R., Dakdouk, R.R.: Obfuscating point functions with multibit output. In: Smart, N.P. (ed.) *EUROCRYPT 2008*. LNCS, vol. 4965, pp. 489–508. Springer (Apr 2008)
25. Canetti, R., Fuller, B., Paneth, O., Reyzin, L.: Key derivation from noisy sources with more errors than entropy. *Cryptology ePrint Archive*, Report 2014/243 (2014), <http://eprint.iacr.org/>
26. Canetti, R., Kalai, Y.T., Varia, M., Wichs, D.: On symmetric encryption and point obfuscation. In: Micciancio, D. (ed.) *TCC 2010*. LNCS, vol. 5978, pp. 52–71. Springer (Feb 2010)
27. Canetti, R., Micciancio, D., Reingold, O.: Perfectly one-way probabilistic hash functions (preliminary version). In: *30th ACM STOC*. pp. 131–140. ACM Press (May 1998)
28. Collberg, C., Thomborson, C., Low, D.: A taxonomy of obfuscating transformations. Technical Report 148, Department of Computer Science, University of Auckland (Jul 1997), <http://citeseer.ist.psu.edu/collberg97taxonomy.html>
29. Dodis, Y., Kalai, Y.T., Lovett, S.: On cryptography with auxiliary input. In: Mitzenmacher, M. (ed.) *41st ACM STOC*. pp. 621–630. ACM Press (May / Jun 2009)
30. Fischlin, M.: Pseudorandom function tribe ensembles based on one-way permutations: Improvements and applications. In: Stern, J. (ed.) *EUROCRYPT’99*. LNCS, vol. 1592, pp. 432–445. Springer (May 1999)
31. Garg, S., Gentry, C., Halevi, S., Raykova, M.: Two-round secure MPC from indistinguishability obfuscation. In: Lindell, Y. (ed.) *TCC 2014*. LNCS, vol. 8349, pp. 74–94. Springer (Feb 2014)
32. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: *54th FOCS*. pp. 40–49. IEEE Computer Society Press (Oct 2013)
33. Garg, S., Gentry, C., Halevi, S., Wichs, D.: On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In: Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part I*. LNCS, vol. 8616, pp. 518–535. Springer (Aug 2014)

34. Gentry, C., Lewko, A., Sahai, A., Waters, B.: Indistinguishability obfuscation from the multilinear subgroup elimination assumption. *Cryptology ePrint Archive, Report 2014/309* (2014), <http://eprint.iacr.org/2014/309>
35. Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. *SIAM J. Comput.* 25(1), 169–192 (Feb 1996), <http://dx.doi.org/10.1137/S0097539791220688>
36. Goldwasser, S., Kalai, Y.T.: On the impossibility of obfuscation with auxiliary input. In: 46th FOCS. pp. 553–562. IEEE Computer Society Press (Oct 2005)
37. Goldwasser, S., Rothblum, G.N.: On best-possible obfuscation. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 194–213. Springer (Feb 2007)
38. Hada, S.: Secure obfuscation for encrypted signatures. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 92–112. Springer (May 2010)
39. Hofheinz, D.: Fully secure constrained pseudorandom functions using random oracles. *Cryptology ePrint Archive, Report 2014/372* (2014), <http://eprint.iacr.org/>
40. Hofheinz, D., Malone-Lee, J., Stam, M.: Obfuscation for cryptographic purposes. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 214–232. Springer (Feb 2007)
41. Hohenberger, S., Rothblum, G.N., shelat, a., Vaikuntanathan, V.: Securely obfuscating re-encryption. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 233–252. Springer (Feb 2007)
42. Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 201–220. Springer (May 2014)
43. Impagliazzo, R.: A personal view of average-case complexity. In: Proceedings of the 10th Annual Structure in Complexity Theory Conference (SCT’95). pp. 134–. SCT ’95, IEEE Computer Society, Washington, DC, USA (1995), <http://dl.acm.org/citation.cfm?id=829497.829786>
44. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 13. pp. 669–684. ACM Press (Nov 2013)
45. Komargodski, I., Moran, T., Naor, M., Pass, R., Rosen, A., Yorgev, E.: One-way functions and (im)perfect obfuscation. *Cryptology ePrint Archive, Report 2014/347* (2014), <http://eprint.iacr.org/>
46. Koppula, V., Ramchen, K., Waters, B.: Separations in circular security for arbitrary length key cycles. *Cryptology ePrint Archive, Report 2013/683* (2013), <http://eprint.iacr.org/2013/683>
47. Lynn, B., Prabhakaran, M., Sahai, A.: Positive results and techniques for obfuscation. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 20–39. Springer (May 2004)
48. Marcedone, A., Orlandi, C.: Obfuscation \implies (ind-cpa security $\neq \implies$ circular security) (2014)
49. Matsuda, T., Hanaoka, G.: Chosen ciphertext security via point obfuscation. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 95–120. Springer (Feb 2014)
50. Pass, R., Seth, K., Telang, S.: Indistinguishability obfuscation from semantically-secure multilinear encodings. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 500–517. Springer (Aug 2014)
51. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) 46th ACM STOC. pp. 475–484. ACM Press (May / Jun 2014)
52. Wee, H.: On obfuscating point functions. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 523–532. ACM Press (May 2005)