

Indistinguishability Obfuscation versus Multi-Bit Point Obfuscation with Auxiliary Input

Christina Brzuska¹

Arno Mittelbach²

¹Tel Aviv University, Israel

²Darmstadt 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

Contents

1	Introduction	3
2	Preliminaries	10
3	IO Implies the Impossibility of MB-AIPO	15
3.1	IO and MB-AIPO are Mutually Exclusive	15
3.2	Implications	19
3.3	On Circumventing our Impossibility Result	20
4	Weak MB-AIPO from iO and AIPO	20
4.1	Point-independent Point Function Obfuscation	20
4.2	Weak MB-AIPO from AIPO and iO	21
5	Leakage Resilient Public-key Encryption	26
A	Constructions of Point Obfuscation Schemes	34

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 [CTL97], 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. [BGI⁺01, BGI⁺12], 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 [Can97, CMR98, Fis99, Wee05, HMLS07, CD08, DKL09, CKVW10, BC10, BP12]. Other positive examples include obfuscating re-encryption [HRsV07] and encrypted signatures [Had10].

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 [Can97, Wee05, HMLS07, DKL09].

Another important distinction is, whether the adversary is given some “leakage” about x , so-called *auxiliary information*, as introduced by Goldwasser and Tauman-Kalai [GK05]. We note that the obfuscator by Canetti [Can97] also allow for auxiliary information about the point x to leak as long as x remains computationally hidden and the obfuscator by Dodis et al. [DKL09] 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 [CD08]. 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 [BP12] 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 [MH14] 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 [GGH⁺13], 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 [GR07] establish, this seemingly weak notion of obfuscation is actually the *best possible* notion of obfuscation. And indeed, the work

by Garg et al. [GGH⁺13] inspired simultaneous breakthroughs for hard problems in several sub-areas of cryptography [SW14, BCP14, ABG⁺13, GGHR14, HSW14, BZ14, BST14, BM14] 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 [BP12, MH14]. That is, if indistinguishability obfuscation exists, then MB-AIPO does not exist and some of the results in [BP12, MH14] 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 [GGH⁺13, PST14, GLSW14]. 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. [GLSW14] 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. [BGI⁺01, BGI⁺12]. 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 $(x', m') = (x, m)$. 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” [GK05] 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 [MH14] 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”.

¹Note that the construction by Canetti and Dakdouk [CD08] is from composable AIPO for which we do not have a candidate construction. The construction by Bitansky and Canetti [BC10, BC14] achieves composable point obfuscation in the virtual grey-box setting (VGB) which implies MB-AIPO, but only for statistically hard-to-invert leakage [MH14].

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 2.9 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 $T_{x,m}$.

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 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 [BFM14] 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) [BHK13a] 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 PRG 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 $\text{PRG}(p(x))$ is equal to y . If yes, it outputs 1. Else, it outputs 0.

For simplicity, let us assume that the PRG 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 $y = \text{PRG}(m)$. In particular, as PRG is a one-way function, y does not leak m . Moreover, as PRG 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

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

function. Due to the PRG security, these two distributions— $C[x, y]$ when y is drawn as an output from the PRG 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 [SW14] as well as the range-extension of Matsuda and Hanaoka [MH14] 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. [BHK13b] used to strengthen the impossibility result by Brzuska et al. [BFM14].

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 z . 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. [BGI⁺01, BGI⁺12] 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. [CKVW10] 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 $y = \text{PRG}(m)$. 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

³Canetti and Dakdouk [CD08] show that composable AIPO already implies MB-AIPO.

⁴Access 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.

result was only known for symmetric-key encryption [CKVW10]. We next discuss existing notions of multi-bit point obfuscation.

NOTIONS OF MULTI-BIT POINT-OBFUSCATION. Lynn et al. [LPS04] initiate the study of obfuscators for point functions with multi-bit output (MBPF) in the idealized 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 [CD08] 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 [BC10, BC14] show that the point obfuscator by Canetti [Can97] 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 [BC10, BC14] and allows the simulator to run in unbounded time—and not the stronger notion of VBB obfuscation. In [CKVW10] 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 [BP12] introduce a clean treatment of a form of auxiliary information which hides the obfuscated point computationally (AIPO) and Matsuda and Hanaoka [MH14] 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 $\mathcal{N}\mathcal{P}$ based on composable AIPO [BP12] thereby circumventing a black-box impossibility result [GK96]. Matsuda and Hanaoka (MH, [MH14]) 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. [CFPR14] 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 [MH14], 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 [MH14] as well as to the construction of a three-round weak zero-knowledge protocol for $\mathcal{N}\mathcal{P}$ by Bitansky and Paneth [BP12].⁵ We 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

⁵Bitanski 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 protocols from AIPO [BP12] 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 [MH14] are not affected by our result.

can show that the VBB obfuscator given by Lynn et al. [LPS04] 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, [BM14]) 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 [Hof14] 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 [BW13, BGI14, KPTZ13]. 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 [Can97].

FURTHER 1-OUT-OF-2 RESULTS. Indistinguishability obfuscation has led to many surprising breakthroughs in a number of sub-areas of cryptography [SW14, BCP14, ABG⁺13, GGHR14, HSW14, BZ14, BST14, BM14]. 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).

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. [BFM14] who show that iO is mutually exclusive with a large class of assumptions from the UCE framework [BHK13a].

Interestingly, several notions of obfuscation are mutually exclusive with iO. Bitansky et al. [BCC⁺14] 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. [BCPR14] 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 [BP13] 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.

Also here, iO gives insights into what is (im)possible. Drawing from the insights of the impossibility result, Bitansky et al. [BCPR14] show how to construct extractable one-way functions with *bounded* auxiliary input under relatively standard assumptions. Finally, Marcedone et al. [MO14], as well as Koppula et al. [KRW13] 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. [BGI⁺01, BGI⁺12] 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. [KMN⁺14] show that (even imperfect) indistinguishability obfuscation does not exist in Pessiland [Imp95], 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. [GGHW14] show that differing-inputs obfuscation—a stronger form of indistinguishably obfuscation that was also introduced in the seminal paper by Barak et al. [BGI⁺01, BGI⁺12]—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 [PST14, BR14, BGK⁺14, AGIS14, GLSW14] supporting its plausibility.

AUXILIARY INPUT. Auxiliary input (AI) has been introduced by Goldwasser and Tauman-Kalai [GK05] 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. [BCPR14] 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. [BCC⁺14] 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 [BM14] and CCA secure PKE schemes [MH14].

While in the VBB-setting AI is a strong notion that corresponds to the existence of a universal simulator [BCC⁺14], 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 [BC10]. Secure AIPOs in the VGB-setting imply AIPOs with statistically hard-to-invert leakage [MH14]. 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. [CFPR14] 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 [BP12] and [MH14] 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 [BP12, MH14] and with computationally hard-to-invert auxiliary information—are mutually exclusive. It remains to investigate whether the positive results in [BP12, MH14] can be salvaged through weaker notions of MB-AIPO or, perhaps, when combining AIPO and iO in a similar way as we did in Section 4 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 [BP12, MH14] 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. [GGHW14] 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.

2 Preliminaries

NOTATION. We denote by $\lambda \in \mathbb{N}$ the security parameter, which all algorithms get implicitly and in unary representation 1^λ . By $\{0, 1\}^\ell$ we denote the set of all bit-strings of length ℓ , and by $\{0, 1\}^*$ the set of all bit-strings of finite length. The length of x is denoted by $|x|$. If $x, y \in \{0, 1\}^*$ are two bit strings of the same length, then we denote their inner product over $\mathbb{GF}(2)$ by $\langle x, y \rangle$. For a finite set X , we denote the action of sampling x uniformly at random from X by $x \leftarrow_s X$, and denote the cardinality of X by $|X|$. We call a randomized algorithm efficient or PPT if it runs in time polynomial in the security parameter. We assume algorithms to be randomized, unless explicitly

⁶Note that [BP12] 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 [MH14] 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.

stated differently. If \mathcal{A} is randomized then by $y \leftarrow \mathcal{A}(x; r)$ we denote that \mathcal{A} is run on input x and with random coins r and produced output y . If no randomness is specified, then we assume that \mathcal{A} is run with freshly sampled uniform random coins, and write this as $y \leftarrow_s \mathcal{A}(x)$. We often refer to algorithms, or tuples of algorithms, as adversaries. If E is an event then we denote by $\Pr[E]$ its probability and if X is a random variable, we denote its expectation by $\mathbb{E}[X]$. We say a function $\text{negl}(\lambda)$ is negligible if $\text{negl}(\lambda) \in \lambda^{-\omega(1)}$. We say a function poly is polynomial if $\text{poly} \in \lambda^{\mathcal{O}(1)}$.

If we speak of an ensemble or a family $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ of circuits, denoted by a calligraphic letter such as \mathcal{C} , we mean that \mathcal{C}_λ contains a set of circuits for each security parameter $\lambda \in \mathbb{N}$. We speak of a sequence of circuits $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ to denote a non-uniform circuit, that is, one circuit for every security parameter. By a distribution or an ensemble of distributions $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ we identify a function ensemble $\{f_\lambda : S_\lambda \rightarrow [0, 1]\}_{\lambda \in \mathbb{N}}$ with corresponding set S_λ that assigns each element in S_λ a probability weight in $[0, 1]$ such that $\sum_{x \in S_\lambda} f_\lambda(x) = 1$ for all $\lambda \in \mathbb{N}$. We consider only efficiently sampleable distributions \mathcal{D} by which we mean that a (possibly non-uniform) algorithm Sam_λ exists that on input a uniformly random string r outputs a value in S_λ according to distribution D_λ , that is, such that for all $\lambda \in \mathbb{N}$ and $x \in S_\lambda$

$$\Pr_r[\text{Sam}_\lambda(r) = x] = f_\lambda(x).$$

We often say we “run” a distribution or we simply write $D_\lambda(1^\lambda)$ to denote that the corresponding sample algorithm is invoked on fresh random coins.

OBFUSCATION. Obfuscation has been extensively used within cryptography and it comes in many different flavors. We now present the various definitions that we use in this paper. We start by recalling the strongest definition of virtual black-box (VBB) obfuscation with auxiliary inputs due to [BGI⁺01, GK05, BGI⁺12].

Definition 2.1 (Worst-case obfuscator with auxiliary input (VBB-AI)). *A PPT \mathcal{O} is a worst-case obfuscator with auxiliary input for an ensemble $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ of poly-size circuits if it satisfies:*

- **Functionality.** *For any $\lambda \in \mathbb{N}$ and $C \in \mathcal{C}_\lambda$, $\mathcal{O}(C)$ is a circuit which computes the same function as C , that is, for all x it holds*

$$\Pr[C'(x) = C(x) \mid C' \leftarrow_s \mathcal{O}(C)] = 1.$$

- **Polynomial slowdown.** *For any $\lambda \in \mathbb{N}$ and $C \in \mathcal{C}_\lambda$, $\Pr[|C'| \leq \text{poly}(|C|) \mid C' \leftarrow_s \mathcal{O}(C)] = 1$.*
- **Virtual black-box.** *For any PPT adversary \mathcal{A} there is a PPT simulator Sim such that for all sufficiently large $\lambda \in \mathbb{N}$, $C \in \mathcal{C}_\lambda$ and $z \in \{0, 1\}^{\text{poly}(\lambda)}$:*

$$\left| \Pr[\mathcal{A}(z, \mathcal{O}(C)) = 1] - \Pr[\text{Sim}^C(z, 1^{|C|}) = 1] \right| \leq \text{negl}(\lambda)$$

where the probability is taken over the coins of \mathcal{A} , Sim and \mathcal{O} .

INDISTINGUISHABILITY OBFUSCATION. While VBB obfuscation as defined above provably does not exist [BGI⁺01] for all circuits, weaker notions such as *indistinguishability obfuscation* may well do. VBB obfuscation 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. 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 was originally proposed by Barak et al. [BGI⁺01] as a potential weakening of virtual-black-box obfuscation. We recall the definition from [GGH⁺13].

Definition 2.2. 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

$$\Pr \left[C'(x) = C(x) : C' \leftarrow_{\S} \text{iO}(1^\lambda, C) \right] = 1.$$

- **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:

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

DIFFERING-INPUTS OBFUSCATION. Differing-inputs obfuscation is closely related to indistinguishability obfuscation and also goes back to the seminal paper of Barak et al. [BGI⁺01, BGI⁺12]. While indistinguishability obfuscation requires circuits to be identical on all inputs, differing-inputs obfuscation intuitively says that if a distinguisher can tell apart two obfuscated circuits then one can efficiently extract a value on which the circuits differ. We here follow the definition of Ananth et al. [ABG⁺13] and Boyle et al. [BCP14] and first define the notion of *differing-inputs distributions* which in turn are then used to define differing-inputs obfuscation.

Definition 2.3 (Differing-inputs circuits). A sample algorithm $(C_0, C_1, z) \leftarrow_{\S} \text{Sam}(1^\lambda)$ that samples circuits from a circuit ensemble $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is said to be a differing-inputs distribution if for all PPT algorithms \mathcal{A} there is a negligible function negl such that:

$$\Pr \left[C_0(x) \neq C_1(x) : (C_0, C_1, z) \leftarrow_{\S} \text{Sam}(1^\lambda), x \leftarrow_{\S} \mathcal{A}(1^\lambda, C_0, C_1, z) \right] \leq \text{negl}(\lambda)$$

Definition 2.4 (Differing-inputs obfuscation). A PPT algorithm diO is a differing-inputs obfuscator for a differing-inputs distribution Sam (for circuit ensemble $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$) if the following holds:

- **Correctness.** For all security parameters $\lambda \in \mathbb{N}$, for all $C \in \mathcal{C}_\lambda$, and for all inputs x we have that

$$\Pr \left[C'(x) = C(x) : C' \leftarrow_{\S} \text{diO}(1^\lambda, C) \right] = 1.$$

- **Security.** For any PPT distinguisher \mathcal{D} , for any $(C_0, C_1, z) \leftarrow_{\S} \text{Sam}(1^\lambda)$ the following distinguishing advantage is negligible:

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

Boyle, Chung and Pass [BCP14] show that any general indistinguishability obfuscator is also a differing-inputs obfuscator for certain classes of circuits. That is, any indistinguishability obfuscator for all circuits in \mathcal{P}/poly is also a differing-inputs obfuscator for distributions over pairs of circuits that differ on at most polynomially many inputs. We recall their Theorem:

Theorem 2.5 ([BCP14]). Let iO be an indistinguishability obfuscator for \mathcal{P}/poly . Let $(\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}, \text{Sam})$ be a differing-inputs distribution for which there exists a polynomial $d : \mathbb{N} \rightarrow \mathbb{N}$, such that

$$\Pr \left[|\{x : C_0(x) \neq C_1(x)\}| \leq d(\lambda) \mid (C_0, C_1, z) \leftarrow_{\S} \text{Sam}(1^\lambda) \right] \geq 1 - \text{negl}(\lambda).$$

Then iO is a differing-inputs obfuscator for $(\{\mathcal{C}_\lambda\}, \text{Sam})$.

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

$$p_x(s) := \begin{cases} 1 & \text{if } s = x \\ \perp & \text{o/w} \end{cases}$$

In this paper, we consider a variant of point function obfuscators under auxiliary input which was first formalized by Canetti [Can97]. We here give the definition from [BP12] presented in a game based formulation. The first definition formalizes unpredictable distributions which are in turn used to define obfuscators for point functions.

Definition 2.6 (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 :*

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

Remark. Alternatively, we could use a variant of Definition 2.6 for *uniform* distributions \mathcal{D} . Jumping ahead, we note that our negative result, Theorem 3.1 also carries over to uniform distributions. Theorem 2.5 by Boyle et al. [BCP14] holds for both, uniform and non-uniform samplers. For our construction of weak MB-AIPO in Proposition 4.5, AIPO for non-uniform samplers yields weak MB-AIPO for non-uniform samplers, and AIPO for uniform samplers yields weak MB-AIPO for uniform samplers. Likewise, for Proposition 5.3, weak MB-AIPO for non-uniform samplers yields a leakage-resilient PKE secure against non-uniform adversaries, while weak MB-AIPO for uniform samplers yields a leakage-resilient PKE secure against uniform adversaries. For ease of presentation, we omit the explicit treatment of uniform and non-uniform adversaries.

Definition 2.7 (Auxiliary input point obfuscation for unpredictable distributions (AIPO)). *A PPT algorithm AIPO is a point obfuscator for unpredictable distributions if on input (z, x) it outputs a polynomial-size circuit that returns 1 on x and 0 everywhere else and satisfies 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

$$p_{x,m}(s) := \begin{cases} m & \text{if } s = x \\ \perp & \text{o/w} \end{cases}$$

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 [MH14] 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 (aka., x) remains hidden given the auxiliary input. From an MB-AIPO obfuscator we now require that the obfuscation of $p_{x,m}$ is indistinguishable from an obfuscation with a changed point value m' where m' is chosen uniformly at random. Intuitively this captures that the obfuscation does not reveal any information about the point value m . We note that also other definitional choices are possible here, which we discuss after presenting the formal definition. We now state unpredictability for a distribution over triples (z, x, m) and then give the formal definition of MB-AIPOs.

Definition 2.8 (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 :*

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

Definition 2.9 (Auxiliary input point obfuscation for unpredictable distributions (MB-AIPO)). *A PPT algorithm MB-AIPO is a multi-bit point obfuscator for unpredictable distributions if on input (z, x, m) it outputs a polynomial-size circuit that returns m on x and 0 everywhere else and satisfies 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  $\leftarrow_{\S} \{0, 1\}$ 
 $(z, x, m_0) \leftarrow_{\S} \mathcal{B}_1(1^\lambda)$ 
 $m_1 \leftarrow_{\S} \{0, 1\}^\lambda$ 
 $p \leftarrow_{\S} \text{MB-AIPO}(x, m_b)$ 
 $b' \leftarrow_{\S} \mathcal{B}_2(1^\lambda, p, z)$ 
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 .

Definition 2.9 was used in [MH14] and is implied by the notion of composable AIPO [CD08] that is used in [BP12]⁷. Interestingly, an MB-AIPO as by the above definition is not necessarily also AIPO secure. For this note that given an obfuscation scheme MB-AIPO we can construct an MB-AIPO' such that MB-AIPO'(x, m) works like MB-AIPO, but additionally, it leaks a hardcore bit of x by outputting a string r and the inner product of x and r . The hardcore bit does not hurt the security of MB-AIPO' as an MB-AIPO, yet given this single bit of information on x an adversary in the AIPO security game can trivially win.

There are two immediate alternative definitions of MB-AIPO that come to mind. One could modify the MB-AIPO definition by requiring that the obfuscation of point function $p_{x,m}$ is indistinguishable from an obfuscation of $p_{x',m'}$. That is, instead of only adapting the point value, we could

⁷Canetti and Dakdouk [CD08] show that composable point functions (without auxiliary input) imply multi-bit point functions (without auxiliary input). Their result carries over to obfuscation in the presence of auxiliary input

conceive a notion in which the honest obfuscation should be indistinguishable from one where both the point address and the point value are chosen uniformly at random. A second alternative would be to require that the obfuscation of point function $p_{x,m}$ is indistinguishable from an obfuscation of $p_{x',m}$. This gives us three alternatives:

$$(x, m) \text{ vs. } (x, m') \quad (\text{Definition 2.9}) \quad (1)$$

$$(x, m) \text{ vs. } (x', m') \quad (2)$$

$$(x, m) \text{ vs. } (x', m) \quad (3)$$

where the first alternative corresponds to our Definition 2.9. We note that the second and third alternative are easily seen to imply also AIPO. However, the third alternative allows the obfuscation to leak large portions of m which goes against the intuition of what MB-AIPO wants to model.

In Section 3, we show that, assuming indistinguishability obfuscation exists, neither of the three definitions can be met. Throughout the paper, the term MB-AIPO refers to Definition 2.9 unless explicitly stated otherwise.

AVERAGE-CASE POINT OBFUSCATION AND STATISTICAL UNPREDICTABILITY. The notions for point obfuscation as defined above is over arbitrary high-entropy distributions over the point address. Instead, one could consider a slightly weaker variant where the point address is sampled according to the uniform distribution. Indeed, Matsuda and Hanaoka [MH14] 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 AIND- δ -cPUAI. Our impossibility result also applies to AIND- δ -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 2.9 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 [MH14] and to circumvent black-box impossibility results for three-round weak zero-knowledge protocols for \mathcal{NP} [BP12]. Our following result says that, if indistinguishability obfuscation and pseudo-random generators exist, then MB-AIPOs (as defined in Definition 2.9) 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; [BFM14]). 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 [BHK13a]—cannot

exist [BFM14]. In the UCE-framework, a hash function H gets a hash key hk and an input x and outputs y . BFM obfuscate the circuit

$$(H(\cdot, x) = y),$$

that given a hash-key hk checks whether hk “matches” the pair (x, y) , that is, whether $H(hk, \cdot)$ maps x to y . They show that, if $|hk| < 2|y|$, 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 $H(hk, \cdot)$ 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 is given as Lemma 3.2.

Theorem 3.1. *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 Definition 2.9 and all the three variants that we discuss below that definition. It also applies to the average-case version where the point address is sampled uniformly, because our adversary samples both, x and m uniformly at random.

To prove Theorem 3.1 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 the unpredictable distribution 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 PRG and finally compares this result with some hard-coded value y . That is, we consider the circuit

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

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 $\text{PRG} : \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{2n(\lambda)}$ 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.⁸ 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 unpredictable

⁸ Canetti et al. [CKVW10] 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.

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

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

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

$$\text{PRG}(p_{x,m}(x)) = y$$

where y was computed as $\text{PRG}(m)$. Thus, it will always output 1. If, however, p is an obfuscation of $p_{x,u}$, then with overwhelming probability over the choice of u , adversary \mathcal{B}_2 returns 0, because PRG-security implies that for any fixed value, the probability that $\text{PRG}(u)$ is equal to that value for a random u , is negligible. It follows that $(\mathcal{B}_1, \mathcal{B}_2)$ break the security of the obfuscator with overwhelming probability. We now prove that $(\mathcal{B}_1, \mathcal{B}_2)$ is also a valid pair of adversaries, that is, that \mathcal{B}_1 is an unpredictable distribution. In the following lemma, we show that, 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 have the same length as the (unobfuscated) leaked circuit, that is, the circuit $(\text{PRG}(\text{uC}(\cdot, x)) = y)$. As the zero circuit does not leak any information about y , the leakage is unpredictable.

Lemma 3.2. *Let n, ℓ, t be polynomials. For $x \in \{0, 1\}^{\ell(\lambda)}$, let $\text{uC}(\cdot, x)$ be a universal circuit that on input a description of a circuit $C \in \{0, 1\}^{t(\lambda)}$ and value x outputs $C(x)$.*

If $\text{PRG} : \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{2n(\lambda)}$ is a pseudo-random generator and iO is a secure indistinguishability obfuscator for all circuits in \mathcal{P}/poly , then for all efficient PPT distinguishers Dist there exists a negligible function negl such that

$$\left| \Pr \left[\text{Dist} \left(1^\lambda, \text{iO} \left(\text{PRG}(\text{uC}(\cdot, x)) = \text{PRG}(m) \right) \right) = 1 \right] - \Pr \left[\text{Dist}(1^\lambda, \text{iO}(\mathbf{0})) = 1 \right] \right| \leq \text{negl}(\lambda)$$

where $\mathbf{0}$ denotes the constant zero circuit padded to the same length as circuit $(\text{PRG}(\text{uC}(\cdot, x)) = \text{PRG}(m))$. The first probability is over the random choice of x and m and the coins of iO and Dist and the second probability is over the coins of iO and Dist .

Proof. We bound the distinguishing probability in Lemma 3.2 with the security of the PRG and the indistinguishability obfuscator iO . We consider the following hybrids (which are also depicted in Figure 1):

Game₁: The game chooses a random value m and computes $y \leftarrow \text{PRG}(m)$. It constructs the predicate circuit $C_{x,y} \leftarrow (\text{PRG}(\text{uC}(\cdot, x)) = y)$ and an obfuscation $\tilde{C} \leftarrow_{\$} \text{iO}(C_{x,y})$. It then calls distinguisher Dist on input \tilde{C} and outputs whatever Dist outputs.

Game₂: As before, except that the game chooses a uniformly random value y . It constructs the predicate circuit $C_{x,y} \leftarrow (\text{PRG}(\text{uC}(\cdot, x)) = y)$ and an obfuscation $\tilde{C} \leftarrow_{\$} \text{iO}(C_{x,y})$. It then calls distinguisher Dist on input \tilde{C} and outputs whatever Dist outputs.

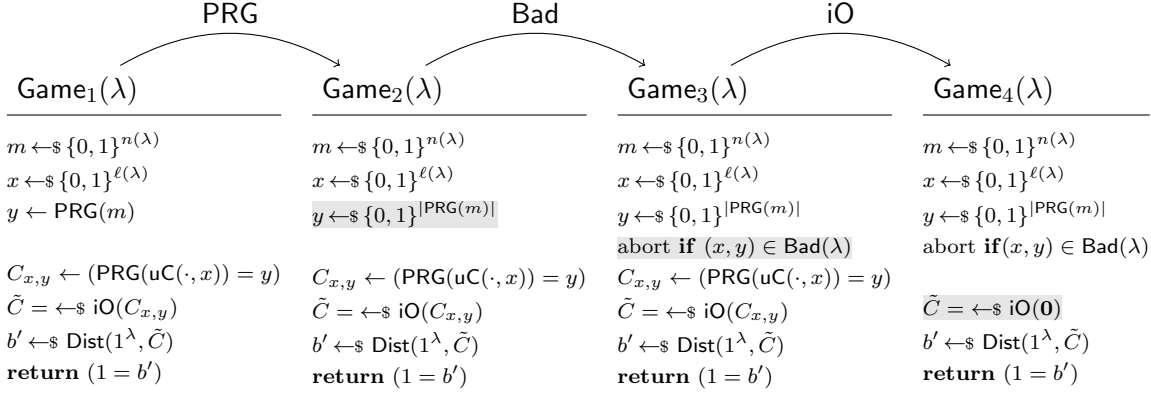


Figure 1: The hybrids for the proof of Lemma 3.2. We have highlighted the changes between the games with a light-grey background.

Game₃: As before, except that the game terminates if there exists a circuit description $C \in \{0, 1\}^{t(\lambda)}$ such that $C_{x,y}(C) = 1$. We denote this event by $(x, y) \in \text{Bad}(\lambda)$ if for values (x, y) such a circuit description C exists.

Game₄: As before, except that now an obfuscation of the constant zero-circuit (padded to appropriate length) $\tilde{C} \leftarrow \text{iO}(\mathbf{0})$ is constructed.

We bound the difference between **Game₁** and **Game₂** by the security of the PRG. **Game₂** and **Game₃** are, by the fundamental lemma of the game-playing technique [BR06], identical until event $\text{Bad}(\lambda)$ occurs. We now prove that $\text{Bad}(\lambda)$ only happens with negligible probability. For a uniformly random pair of values (x, y) we have that the probability that there exists circuit C such that $\text{PRG}(\text{uC}(C, x)) = y$ is upper bounded by the probability that y is in the image of the PRG. As the PRG maps $n(\lambda)$ bits to $2n(\lambda)$ bits, the probability that a random y is in the image of the PRG is upper bounded by $2^{-n(\lambda)}$.

Finally, we bound the difference between **Game₃** and **Game₄** by the security of the indistinguishability obfuscator iO by noting that if $(x, y) \notin \text{Bad}(\lambda)$ then circuit $C_{x,y}$ encodes the constant zero circuit. We now explain this formally.

Firstly, let us externalize some of the variables that the games use and introduce a unified notation for **Game₃** and **Game₄**. For $i \in \{3, 4\}$, let $\text{Game}_i[x, y](\lambda)$ be equal to the game **Game_i**(λ) where the game chooses values x and y . We define $\mathcal{A}[x, y](C)$ to be an adversary against the indistinguishability obfuscator iO that gets a circuit C as input, where C is either an obfuscation of circuit $C_{x,y}$ or an obfuscation of the constant zero circuit $\mathbf{0}$ padded to the length of circuit $(\text{PRG}(\text{uC}(\cdot, x)) = y)$. Adversary $\mathcal{A}[x, y](C)$ runs distinguisher D on input $(1^\lambda, C)$ and outputs whatever D outputs.

If $C = C_{x,y}$ then adversary $\mathcal{A}[x, y](C)$ perfectly simulates game **Game₃** $[x, y](\lambda)$ and if $C = \mathbf{0}$ then the adversary simulates **Game₄** $[x, y](\lambda)$. Thus, we can rewrite the difference between the game's distributions

$$\Pr[\text{Game}_3(\lambda)] - \Pr[\text{Game}_4(\lambda)]$$

as

$$\mathbb{E}_{x,y} \left[\Pr[\text{Game}_3[x, y](\lambda)] \right] - \mathbb{E}_{x,y} \left[\Pr[\text{Game}_4[x, y](\lambda)] \right]$$

Note that x and y are chosen such that $(x, y) \notin \text{Bad}(\lambda)$. Due to the linearity of expectation we can further rearrange this as

$$\begin{aligned}
&= \mathbb{E}_{x,y} \left[\Pr[\text{Game}_3[x, y](\lambda)] - \Pr[\text{Game}_4[x, y](\lambda)] \right] \\
&= \mathbb{E}_{x,y} \left[\Pr \left[\mathcal{A}[x, y](1^\lambda, \text{iO}(C_{x,y})) = 1 \right] - \Pr \left[\mathcal{A}[x, y](1^\lambda, \text{iO}(\mathbf{0})) = 1 \right] \right] \\
&= \mathbb{E}_{x,y} \left[\text{Adv}_{\text{iO}, \mathcal{A}[x,y], C_{x,y}, \mathbf{0}}^{\text{iO}}(\lambda) \right] \\
&\leq \max_{x,y} \text{Adv}_{\text{iO}, \mathcal{A}[x,y], C_{x,y}, \mathbf{0}}^{\text{iO}}(\lambda)
\end{aligned}$$

By the security of the indistinguishability obfuscator, the advantage of any efficient adversary is negligible and, hence, also $\max_{x,y} \text{Adv}_{\text{iO}, \mathcal{A}[x,y], C_{x,y}, \mathbf{0}}^{\text{iO}}(\lambda)$ is negligible. \square

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 [MH14]. 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 3.3. *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 [CD08] 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 3.4. *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 [MH14] and Bitansky and Paneth give a three-round weak zero-knowledge protocol for \mathcal{NP} based on composable AIPO [BP12].⁹ In Section 4 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 [BP12, MH14] and such a proof is not straightforward, so it remains to study whether other weak variants of MB-AIPO could be used.

A RANDOM ORACLE UNINSTANTIABILITY. Lynn et al. [LPS04] 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 indistinguishability obfuscation exists our result rules out the existence of a standard model hash function that can instantiate the random oracle in their construction.

Corollary 3.5. *If indistinguishability obfuscation exists, then the multi-bit output point function obfuscator by Lynn et al. [LPS04] cannot be instantiated in the standard model so that it achieves VBB security with auxiliary input.*

⁹We note that the construction of 3-message witness-hiding protocols from AIPO [BP12] as well as the construction of a CCA secure PKE scheme from lossy encryption schemes and MB-AIPO with statistically hard-to-invert information [MH14] are not affected by our result.

3.3 On Circumventing our Impossibility Result

Matsuda and Hanaoka [MH14] present a CCA-secure PKE scheme that is based on MB-AIPO and, thus, is ruled out by our impossibility result, if indistinguishability obfuscation exists. However, they also present a version of a CCA-secure PKE scheme based on the weaker assumption of MB-AIPO that is secure only with respect to statistically unpredictable distributions. Indeed, our techniques do not carry over to ruling out MB-AIPO for statistically unpredictable distributions, because the way in which we use indistinguishability obfuscation, inherently relies on computational security. Switching to a statistical notion of security was also proposed for UCEs in order to salvage a large number of applications [BFM14, BHK13a].

Moreover, Canetti et al. [CKVW10] present notions of MB-AIPO in the setting of computational unpredictability that are not affected by our impossibility result. In the following section, we present a variant of MB-AIPO that is stronger than theirs but weaker than the definition that we show to be impossible. Namely, we strengthen the assumption on unpredictable distributions to remain unpredictable even in case the point value m is given. We call this notion strong unpredictability and, indeed, give a construction based on AIPO and indistinguishability obfuscation. An analogous notion of unpredictability has recently been introduced by Brzuska and Mittelbach in the context of UCE security [BM14].

4 Weak MB-AIPO from iO and AIPO

In this section we show that, despite the negative results from the previous section, point obfuscation—in particular AIPO—and indistinguishability obfuscation (iO) together make a powerful team. Our first observation is that we can use iO to construct a point obfuscation scheme which can securely obfuscate point functions given as input the point *function*, rather than the point *address*. This is inherently different from all previous point obfuscation schemes [Can97, CMR98, Fis99, Wee05, CD08, CKVW10, BC10, BP12] which always take the point address as input. We call this sort of obfuscator point-independent as the obfuscator works independently of the actual point (in fact it may be computationally infeasible for the obfuscator to recover the point from its input). Using a similar technique we then construct a mild form of a multi-bit point function obfuscation scheme secure in the presence of auxiliary input that we call *weak MB-AIPO*. Later, in Section 5, we use our construction of weak MB-AIPOs to construct a public-key encryption scheme that is leakage resilient with respect to any computationally hard-to-invert leakage of the secret-key.

4.1 Point-independent Point Function Obfuscation

Goldwasser and Rothblum [GR07] introduce the notion of best-possible obfuscation. Intuitively an obfuscator is a best-possible obfuscator if an obfuscation leaks as little information about the original program as any functionally equivalent circuit. In the context of point functions this means that a best-possible obfuscator for point functions (or for more general function classes) on input a point function p_x would need to output an obfuscated point function which is at least as good as an obfuscation produced by, for example, an AIPO obfuscator on input the point address x . Goldwasser and Rothblum show that if we consider PPT obfuscators, then the notions of best-possible obfuscation and indistinguishability obfuscation are equivalent [GR07]. Let us recall the definition of best-possible obfuscation:

Definition 4.1 ([GR07]: Best-possible obfuscation). *A PPT \mathcal{O} is a best-possible obfuscator for an ensemble $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ of families of poly-size circuits if it satisfies the preserving functionality and*

polynomial slowdown properties as in Definition 2.1, and also has the following property (instead of the virtual black-box property).

- **Computational best-possible obfuscation.** For any polynomial size learner \mathcal{L} , there exists a polynomial size simulator Sim such that for every large enough input length λ , for any circuit $C_1 \in \mathcal{C}_\lambda$ and for any circuit $C_2 \in \mathcal{C}_\lambda$ that computes the same function as C_1 and such that $|C_1| = |C_2|$ it holds for any PPT adversary \mathcal{A} that

$$\left| \Pr \left[\mathcal{A}(1^\lambda, \mathcal{L}(\mathcal{O}(C_1))) = 1 \right] - \Pr \left[\mathcal{A}(1^\lambda, \text{Sim}(C_2)) = 1 \right] \right| \leq \text{negl}(\lambda)$$

Thus, if iO is an indistinguishability obfuscator for all circuits in \mathcal{P}/poly , then the mapping $C \mapsto \text{iO}(C)$ yields a point-independent point function obfuscator since the indistinguishability obfuscation of point function p_x must not leak more about x than what can be extracted from the “best possible” point-obfuscation scheme on input x . Furthermore, if AIPO exists, then a best possible point obfuscation is at least as good as AIPO and hence, our construction yields a point-independent obfuscation scheme that is secure in the presence of auxiliary inputs.

4.2 Weak MB-AIPO from AIPO and iO

In the following we give a relaxed definition of MB-AIPO and subsequently give a construction based on plain AIPO and indistinguishability obfuscation. We weaken the original MB-AIPO definition (Definition 2.9) by requiring that an obfuscation must only be secure for unpredictable distributions that hide the point address even given the the point value m . We call this notion of unpredictability *strong unpredictability* and note that an analogous notion has recently been introduced in the context of UCE security [BM14]. If we restrict adversaries to strong unpredictability we yield an MB-AIPO notion that we call *weak MB-AIPO*.

Definition 4.2 (Strongly unpredictable distribution). *We say that a distribution ensemble $\mathcal{D} = \{D_\lambda = (Z_\lambda, X_\lambda, M_\lambda)\}_{\lambda \in \mathbb{N}}$, on triples of strings is strongly unpredictable if no poly-size (non-uniform) circuit family can predict X_λ from (Z_λ, M_λ) . That is, for every sequence of poly-sized circuits $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ and for all large enough λ :*

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

Definition 4.3 (Weak MB-AIPO). *A PPT algorithm AIPO is a weak multi-bit point obfuscator if it is a MB-AIPO for strongly unpredictable distributions.*

Next, we present our construction of a weak MB-AIPO scheme. The idea will be to use a plain AIPO (for the point address x). We then construct a circuit which evaluates the AIPO and outputs m if, and only if, this evaluation returns 1. The construction will be the indistinguishability obfuscation of that circuit.

Construction 4.4. *Let AIPO be a secure AIPO and iO be a secure indistinguishability obfuscator for all circuits in \mathcal{P}/poly . We construct a weak MB-AIPO obfuscator MB-AIPO as follows. On input a point address x and value m MB-AIPO constructs a point obfuscation $p_x \leftarrow_{\S} \text{AIPO}(x)$. It then constructs the following circuit*

$$C[p_x, m](x^*) := \mathbf{if} (p_x(x^*) = 1) \mathbf{then return} m \mathbf{else return} \perp$$

and outputs an indistinguishability obfuscation of $C[p_x, m]$.

Proposition 4.5. *If AIPO exists and if iO is a secure indistinguishability obfuscator for all circuits in \mathcal{P}/poly then the above construction is a weak MB-AIPO.*

Proof. Consider an adversary $(\mathcal{B}_1, \mathcal{B}_2)$ against the weak MB-AIPO property of MB-AIPO where \mathcal{B}_1 implements an unpredictable distribution, that is, on input the security parameter it outputs a point function description (x, m) together with some auxiliary information z . We proceed to proof the result in three hybrid steps where the first game is identical to the MB-AIPO game where adversary \mathcal{B}_2 receives an honest obfuscation of point function $p_{x,m}$ and the last game is identical to the MB-AIPO game where the adversary receives an obfuscation of point function $p_{x,m'}$ for a uniformly random m' :

Game₁: Is the original MB-AIPO game where \mathcal{B}_1 outputs $(z, (x, m))$, that is, a point description (x, m) and auxiliary information z and where \mathcal{B}_2 receives as input $(\text{iO}(C[p_x, m]), z)$.

Game₂: Is identical to before but now \mathcal{B}_2 receives an indistinguishability obfuscation of the constant zero circuit $\mathbf{0}$.

Game₃: Is the original MB-AIPO game where \mathcal{B}_1 outputs $(z, (x, m))$, that is, a point description (x, m) and auxiliary information z and where \mathcal{B}_2 receives as input $(\text{iO}(C[p_x, m']), z)$ where m' is a uniformly random string.

We, thus, need to show that the difference between the above games is negligible. For the difference between games **Game₁** and **Game₂** we consider the following claim.

Claim 4.6. *Let Sam be the following sample algorithm. It runs adversary \mathcal{B}_1 to receive point function description (x, m) and auxiliary information z . It constructs an AIPO $p_x \leftarrow_{\$} \text{AIPO}(x)$ and then builds circuit $C[p_x, m]$ as in the construction. Additionally it pads the constant zero circuit $\mathbf{0}$ to the same length as $C[p_x, m]$. It outputs $(C[p_x, m], \mathbf{0}, z)$. If AIPO is secure then this distribution of pairs of circuits is differing-inputs.*

Assume this is not the case. Then there exists an extractor Ext that on input $(C[p_x, m], \mathbf{0}, z)$ outputs a target value τ (with noticeable probability) such that $C[p_x, m](\tau) \neq \mathbf{0}(\tau)$ and, hence, $\tau = x$. We will use this extractor to break the security of the AIPO scheme, that is, we construct $(\mathcal{A}_1, \mathcal{A}_2)$ as an adversary against the security of AIPO. Adversary \mathcal{A}_1 runs \mathcal{B}_1 to receive point function description (x, m) and auxiliary information z . It chooses a random value r and computes $b \leftarrow \langle r, x \rangle$. It outputs $(x, (z, m, r, b))$. Adversary \mathcal{A}_2 gets as input a (z, m, r, b) and a point function obfuscation p which is either an obfuscation for p_x or for p_u for a uniformly random u . It constructs circuits $C[p, m]$ and the constant zero circuit $\mathbf{0}$ and runs extractor $\tau \leftarrow_{\$} \text{Ext}(C[p, m], \mathbf{0}, z)$. If $\tau = \perp$ it flips a bit and outputs it. Else, if $\tau \neq \perp$, then $p(\tau) = 1$ and, hence, value τ is either equal to x or to u . In this case adversary \mathcal{A}_2 outputs 1 if $\langle \tau, r \rangle = b$ and 0 otherwise.

ANALYSIS. Let us denote by ϵ the probability that Ext outputs a value $\tau \neq \perp$ in the differing-inputs game. Then, if $\tau = x$ adversary \mathcal{A}_2 will always output 1 as, by construction, $\langle x, r \rangle = b$. If, on the other hand, $\tau = u$, then adversary \mathcal{A}_2 will output 1 only with probability $\frac{1}{2}$ as u and r are randomly chosen values and r remains hidden from Ext . Thus, our adversary has a distinguishing advantage of $\frac{1}{2}\epsilon$. In the following we make this intuition formal. We note that our simulation technique is inspired by Brzuska and Mittelbach [BM14] who build variants of UCE security based on puncturable PRFs, iO and AIPO and that the formal analysis is almost taken verbatim.

Let us denote by $d = 0$ the event that in the AIPO-game, the honest point function p_x gets obfuscated, and let $d = 1$ describe the event that in the AIPO-game, p_u gets obfuscated for a random

value u . Let further ϵ be the probability that Ext returns a value $\tau \neq \perp$ in the differing-inputs game, that is, $\epsilon := \Pr[\perp \neq \text{Ext} \mid d = 0]$ (note that in the differing-inputs game the obfuscated point function is always p_x). For readability we will drop the inputs given to adversaries Ext and \mathcal{A}_2 in the following formal treatment. We can now consider the distinguishing probability of our adversary \mathcal{A}_2

$$\begin{aligned}
& \Pr[\mathcal{A}_2 = 1 \mid d = 0] - \Pr[\mathcal{A}_2 = 1 \mid d = 1] \\
&= \Pr[\mathcal{A}_2 = 1 \mid d = 0, \text{Ext} \neq \perp] \cdot \Pr[\text{Ext} \neq \perp \mid d = 0] + \\
& \quad \Pr[\mathcal{A}_2 = 1 \mid d = 0, \text{Ext} = \perp] \cdot \Pr[\text{Ext} = \perp \mid d = 0] - \Pr[\mathcal{A}_2 = 1 \mid d = 1] \\
&= \Pr[\text{Ext} \neq \perp \mid d = 0] + \frac{1}{2} \cdot \Pr[\text{Ext} = \perp \mid d = 0] - \Pr[\mathcal{A}_2 = 1 \mid d = 1] \\
&= \Pr[\text{Ext} \neq \perp \mid d = 0] + \frac{1}{2} \cdot \left(1 - \Pr[\text{Ext} \neq \perp \mid d = 0]\right) - \Pr[\mathcal{A}_2 = 1 \mid d = 1] \\
&= \frac{1}{2} \cdot \Pr[\text{Ext} \neq \perp \mid d = 0] + \frac{1}{2} - \Pr[\mathcal{A}_2 = 1 \mid d = 1] = \frac{1}{2}\epsilon + \frac{1}{2} - \Pr[\mathcal{A}_2 = 1 \mid d = 1]
\end{aligned}$$

Let U denote a random variable describing the choice of point function p_u (in case $d = 1$) and note that u is chosen from $\{0, 1\}^\lambda$.

$$\begin{aligned}
&= \frac{1}{2}\epsilon + \frac{1}{2} - \Pr[\mathcal{A}_2 = 1 \mid d = 1, \text{Ext} \neq \perp] \cdot \Pr[\text{Ext} \neq \perp \mid d = 1] + \\
& \quad \Pr[\mathcal{A}_2 = 1 \mid d = 1, \text{Ext} = \perp] \cdot \Pr[\text{Ext} = \perp \mid d = 1] \\
&= \frac{1}{2}\epsilon + \frac{1}{2} - \\
& \quad \frac{1}{2^\lambda} \sum_{u \in \{0, 1\}^\lambda} \left(\Pr[\mathcal{A}_2 = 1 \mid d = 1, U = u, \text{Ext} \neq \perp] \cdot \Pr[\text{Ext} \neq \perp \mid U = u, d = 1] + \right. \\
& \quad \left. \Pr[\mathcal{A}_2 = 1 \mid d = 1, U = u, \text{Ext} = \perp] \cdot \Pr[\text{Ext} = \perp \mid U = u, d = 1] \right)
\end{aligned}$$

If extractor Ext outputs a value u (given that $d = 1$), then the probability that \mathcal{A}_2 outputs 1 ($\Pr[\mathcal{A}_2 = 1 \mid d = 1, U = u, \text{Ext} \neq \perp]$) is equivalent to $\Pr_{R,b}[\langle R, u \rangle = b]$ where random variable R denotes the choice of value r by \mathcal{A}_1 to compute $b = \langle r, x^* \rangle$. Note that extractor Ext is independent of R and b and, thus, $\Pr_{R,b}[\langle R, u \rangle = b] = \frac{1}{2}$. It follows

$$\begin{aligned}
&= \frac{1}{2}\epsilon + \frac{1}{2} - \\
& \quad \frac{1}{2^\lambda} \sum_{u \in \{0, 1\}^\lambda} \left(\Pr_{R,b}[\langle R, u \rangle = b] \cdot \Pr[\text{Ext} \neq \perp \mid U = u, d = 1] + \frac{1}{2} \cdot \Pr[\text{Ext} = \perp \mid U = u, d = 1] \right) \\
&= \frac{1}{2}\epsilon + \frac{1}{2} - \frac{1}{2^\lambda} \sum_{u \in \{0, 1\}^\lambda} \left(\frac{1}{2} \cdot \left(\Pr[\text{Ext} \neq \perp \mid U = u, d = 1] + \Pr[\text{Ext} = \perp \mid U = u, d = 1] \right) \right) \\
&= \frac{1}{2}\epsilon + \frac{1}{2} - \frac{1}{2^\lambda} \sum_{u \in \{0, 1\}^\lambda} \frac{1}{2} \cdot 1 = \frac{1}{2}\epsilon
\end{aligned}$$

This establishes that adversary \mathcal{A}_2 is able to distinguish with noticeable probability since, by assumption, the success probability ϵ of extractor Ext is noticeable.

It remains to show that \mathcal{A}_1 implements an unpredictable distribution. Note that \mathcal{B}_1 is strongly unpredictable (cp. Definition 4.2) and hence point x remains computationally hidden given values z

and m . As r is a uniformly random value chosen independently of x and b is a single bit which can be guessed it follows that, indeed, \mathcal{A}_1 implements an unpredictable distribution.

This concludes the proof of Claim 4.6. \diamond

We next show that with Claim 4.6 it follows that games Game_1 and Game_2 are negligibly close. Boyle et al. [BCP14] show that every indistinguishability obfuscator is also a differing-inputs obfuscator for circuit families that differ on at most polynomially many points (we give their result as Theorem 2.5 on page 12). As the circuits considered in Claim 4.6 differ only on a single point it follows that their obfuscation under an indistinguishability obfuscator are computationally indistinguishable and hence games Game_1 and Game_2 are negligibly close.

For games Game_2 and Game_3 , the analysis is analogous. The only difference consists in m' being chosen at random. Thus, we can make use of Claim 4.6 with the sample adapted to output circuit $C[p_x, m']$. This concludes the proof. \square

A SECOND CONSTRUCTION. In the next section we will construct a leakage resilient public-key encryption scheme, for which we will need an extended construction of MB-AIPO that we present next. Namely, we prove that the construction is still secure if, additionally, we return the AIPO of x . Intuitively, that should not harm security, because, given an MB-AIPO for (x, m) , it is easy to construct an AIPO for x . However, making this statement formal requires some care.

Construction 4.7. *Let AIPO be a secure AIPO and iO be a secure indistinguishability obfuscator for all circuits in \mathcal{P}/poly . We construct a weak MB-AIPO obfuscator MB-AIPO as follows. On input a point address x and value m MB-AIPO constructs a point obfuscation $p_x \leftarrow_{\$} \text{AIPO}(x)$. It then constructs the following circuit*

$$C[p_x, m](x^*) := \mathbf{if} (p_x(x^*) = 1) \mathbf{then return } m \mathbf{ else return } \perp$$

and outputs an indistinguishability obfuscation of $C[p_x, m]$ together with p_x .

We next show that also this adapted construction fulfills the security properties of a weak MB-AIPO scheme.

Proof. We proceed by the following game hops where the first is identical to the MB-AIPO setting where the adversary $(\mathcal{B}_1, \mathcal{B}_2)$ gets an honest obfuscation of point function $p_{x,m}$ and the last is identical to the dual setting where it gets as input an obfuscation of $p_{x,m'}$ for a uniformly random point value m' .

Game₁: Is the original MB-AIPO game with where adversary \mathcal{B}_1 outputs $(z, (x, m))$ and where \mathcal{B}_2 gets as input $((\text{iO}(C[p_x, m]), p_x), z)$. Note that p_x is an obfuscation of the point function for point address x .

Game₂: As before, but instead of returning p_x , that is the obfuscation of coming from the construction, we construct a fresh point obfuscation of x . Hence, the adversary \mathcal{B}_2 gets as input $((\text{iO}(C[p_x, m]), \text{AIPO}(x)), z)$.

Game₃: Instead of returning $\text{AIPO}(x)$, we return $\text{AIPO}(u)$ for a random point u , that is, the adversary \mathcal{B}_2 gets $((\text{iO}(C[p_x, m]), \text{AIPO}(u)), z)$.

Game₄: Instead of returning $\text{iO}(C[p_x, m])$, we return $\text{iO}(C[p_x, m'])$ for a random m' , that is, the adversary \mathcal{B}_2 gets $((\text{iO}(C[p_x, m']), \text{AIPO}(u)), z)$.

Game₅: Instead of returning $\text{AIPO}(u)$ for a random point u , we return $\text{AIPO}(x)$, that is, the adversary \mathcal{B}_2 gets $((\text{iO}(C[p_x, m']), \text{AIPO}(x)), z)$.

Game₆: Instead of returning a fresh point obfuscation of x , we return p_x , that is, the adversary \mathcal{B}_2 gets $((\text{iO}(C[p_x, m']), p_x), z)$.

Note that the last game corresponds to the MB-AIPO-game where the point value is chosen uniformly at random. Hence, it suffices to show that the six games are computationally indistinguishable.

Game₁ TO Game₂. We reduce to the security of the indistinguishability obfuscator iO . Note that the two circuits p_x (given to the adversary in **Game₁**) and $\text{AIPO}(x)$ (given to the adversary in **Game₂**) are functionally equivalent as they are two independently generated obfuscations of point function for point x . Let $(\mathcal{B}_1, \mathcal{B}_2)$ be a distinguisher between **Game₁** and **Game₂**. Then, we construct an adversary against the security property of obfuscator iO analogously to the final game hop of the proof of Lemma 3.2.

Game₂ TO Game₃. We reduce to the distinguishing advantage between games **Game₂** and **Game₃** to the security of the point obfuscation scheme AIPO . Let $(\mathcal{B}_1, \mathcal{B}_2)$ be a distinguisher between **Game₂** and **Game₃**. Then, we construct an adversary $(\mathcal{D}_1, \mathcal{D}_2)$ against AIPO as follows: \mathcal{D}_1 runs \mathcal{B}_1 to obtain $(z, (x, m))$. It runs $p_x \leftarrow \text{AIPO}(x)$, computes $C \leftarrow \text{iO}(C[p_x, m])$ and returns $((z, C), x)$. That is, the auxiliary input returned by \mathcal{D}_1 is (z, C) . \mathcal{D}_2 receives $(p, (z, C))$ as input and runs \mathcal{B}_2 on $((C, p), z)$, that is, on multi-bit point function (C, p) and auxiliary information z . \mathcal{D}_2 outputs whatever \mathcal{B}_2 outputs. If p is $\text{AIPO}(x)$, then the input distribution to $(\mathcal{B}_1, \mathcal{B}_2)$ is as in **Game₂**. If p is $\text{AIPO}(u)$, then the input distribution to $(\mathcal{B}_1, \mathcal{B}_2)$ is as in **Game₃**. Hence, if $(\mathcal{B}_1, \mathcal{B}_2)$ is successful, then so is $(\mathcal{D}_1, \mathcal{D}_2)$.

It remains to show that $(\mathcal{D}_1, \mathcal{D}_2)$ is a valid adversary against AIPO , that is, that \mathcal{D}_1 is unpredictable. This follows from the strong unpredictability of \mathcal{B}_1 and the security of AIPO , that is, given a predictor \mathcal{P} against the strong unpredictability of \mathcal{D}_1 , we will either construct an adversary against the security of AIPO or a predictor \mathcal{R} against the strong unpredictability of \mathcal{B}_1 .

Let \mathcal{P} be a predictor against the unpredictability of \mathcal{D}_1 . Then, we construct \mathcal{R} against the strong unpredictability of \mathcal{B}_1 as follows: Adversary \mathcal{R} gets as input (z, m) where z is the auxiliary information generated by \mathcal{B}_1 and m is the point value output by \mathcal{B}_1 . It draws a random point v and runs $p_v \leftarrow \text{AIPO}(v)$, $C \leftarrow \text{iO}(C[p_v, m])$. It then runs predictor \mathcal{P} on input (C, z) and outputs whatever \mathcal{P} returns. We now need to argue that \mathcal{P} produces the right output, although v is used in the generation of C and not x as expected by \mathcal{P} . We reduce the difference in \mathcal{P} 's behavior to the AIPO security.

Assume that \mathcal{P} has non-negligible probability of returning x when getting $(\text{iO}(C[p_x, m]), z)$, but not when getting $(\text{iO}(C[p_v, m]), z)$. Then, we construct an adversary $(\mathcal{D}_1, \mathcal{D}_2)$ against the AIPO security as follows. Adversary \mathcal{D}_1 runs \mathcal{B}_1 to get $(z, (x, m))$. It draws a random string r and sets bit b to be the inner product of r and x , that is, $b \leftarrow \langle r, x \rangle$. Adversary \mathcal{D}_1 outputs $((z, m, r, b), x)$, that is, its leakage is (z, m, r, b) . Now, \mathcal{D}_2 gets (z, m, r, b) as well as a point function p as input. It runs the predictor $\mathcal{P}(\text{iO}(C[p, m]), z)$ to obtain some value x' . It tests whether $p(x') = 1$ and if this is not the case it returns a random bit. If $p(x') = 1$, then it returns 1 if and only if the inner product of r and x' is equal to b , that is, it returns $\langle r, x' \rangle$.

Now, to see that $(\mathcal{D}_1, \mathcal{D}_2)$ breaks the security of AIPO , we first show that it is a valid adversary, i.e., that \mathcal{D}_1 is unpredictable and then analyse the success probability. Adversary \mathcal{D}_1 is unpredictable because \mathcal{B}_1 is strongly unpredictable and because b is a single bit that can be guessed. Let us turn to the success probability of $(\mathcal{D}_1, \mathcal{D}_2)$. The formal treatment is identical to the advantage computation

within Claim 4.6. We now give the intuition behind the success probability of $(\mathcal{D}_1, \mathcal{D}_2)$. If p is a point obfuscation of p_x , then \mathcal{P} returns x with non-negligible probability ν and in these cases, \mathcal{D}_2 returns 1 with probability 1, because the bit always matches. Thus, if p is a point function of p_x , then \mathcal{D}_2 returns 1 with probability $\nu + (1 - \nu) \cdot \frac{1}{2} = \frac{1}{2} + \frac{\nu}{2}$. If p is a point obfuscation of a random point p_v , then, independently of the behaviour of \mathcal{P} , \mathcal{D}_2 returns 1 with probability $\frac{1}{2}$. Hence, the success probability is non-negligible.

Game₃ TO Game₄. We reduce to the MB-AIPO security of Construction 4.4 that we established in Proposition 4.5. Let $(\mathcal{B}_1, \mathcal{B}_2)$ be an adversary that distinguishes between Game₃ and Game₄, where \mathcal{B}_1 is strongly unpredictable. We construct an adversary $(\mathcal{D}_1, \mathcal{D}_2)$ against the weak MB-AIPO property of Construction 4.4 as follows. \mathcal{D}_1 runs \mathcal{B}_1 and outputs whatever \mathcal{B}_1 outputs. \mathcal{D}_2 gets (C, z) . It draws a random value u and computes $p \leftarrow_{\$} \text{AIPO}(u)$. It then runs \mathcal{B}_2 on $((C, p), z)$ and outputs whatever \mathcal{B}_2 outputs. \mathcal{D}_1 is strongly unpredictable, because \mathcal{B}_1 is. Moreover, the simulation is perfect. Hence, the advantage of $(\mathcal{D}_1, \mathcal{D}_2)$ is identical to the advantage of $(\mathcal{B}_1, \mathcal{B}_2)$.

Game₄ TO Game₅. Analogous to the game hop from Game₂ to Game₃.

Game₅ TO Game₆. Analogous to the game hop from Game₁ to Game₂. □

5 Leakage Resilient Public-key Encryption

In this section, we will use the construction of weak MB-AIPO to build a leakage resilient public-key encryption scheme. Our result is inspired by Canetti et al. who show that multi-bit point obfuscation is tightly connected to symmetric encryption [CKVW10]. They give an intriguingly simple construction of a symmetric encryption scheme from an MB-AIPO as follows. Encryption under key k is defined as $\text{Enc}_k(m) := \text{MB-AIPO}(k, m)$. Correspondingly, decryption interprets the ciphertext as a circuit and runs it on the key, that is, $\text{Dec}_k(c) := c(k)$. Furthermore, they show how to build an MB-AIPO scheme from symmetric encryption. They classify the relationships between the two primitives depending on the strength of the MB-AIPO and encryption scheme, respectively. In particular, they show that a version of MB-AIPO obfuscation implies the existence of a symmetric key encryption scheme secure in the presence of leakage (of the key) with the only requirement that the leakage computationally hides the secret key.

Let us recall a somewhat simplified version of their notion of semantic security of a symmetric encryption scheme with weak keys and auxiliary inputs:

Definition 5.1 ([CKVW10]: Symmetric Encryption with Weak Keys and Auxiliary Inputs). *Let $\mathcal{D} = \{D_\lambda = (Z_\lambda, K_\lambda)\}_{\lambda \in \mathbb{N}}$ be an unpredictable distribution ensemble. We say that an encryption scheme has semantic security with keys chosen from $\{K_\lambda\}_{\lambda \in \mathbb{N}}$ and auxiliary inputs from $\{Z_\lambda\}_{\lambda \in \mathbb{N}}$ if there exists a PPT algorithm $\text{Sim}(1^\lambda, \ell)$ such that, for all PPT adversaries \mathcal{A} we have:*

$$\left| \Pr \left[\text{SEM}_0^{\mathcal{D}, \text{Sim}}(\mathcal{A}, \lambda) = 1 \right] - \Pr \left[\text{SEM}_1^{\mathcal{D}, \text{Sim}}(\mathcal{A}, \lambda) = 1 \right] \right| \leq \text{negl}(\lambda)$$

where the games $\text{SEM}_b^{\mathcal{D}, \text{Sim}}$ for $b = 0, 1$ are defined via the following experiment:

1. $(z, k) \leftarrow_{\$} D_\lambda$ and give z to \mathcal{A}
2. Adversary \mathcal{A} submits a query m . Set $c_0 \leftarrow_{\$} \text{Enc}_k(m)$, $c_1 \leftarrow_{\$} \text{Sim}(1^\lambda, |m|)$ and give c_b to \mathcal{A} .
3. The output of the game is the output of \mathcal{A} .

Figure 2: The IND-CPA game for public-key encryption schemes with hard-to-invert key leakage. An adversary is deemed admissible if it is PPT and if the output of \mathcal{A}_0 computationally hides the key, that is, secret key sk has super-logarithmic min-entropy given z .

```

IND-CPA $_{\mathcal{E}}^{\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2}(\lambda)$ 
-----
 $b \leftarrow_{\$} \{0, 1\}; (pk, sk) \leftarrow_{\$} \mathcal{E}.\text{KGen}(1^\lambda)$ 
 $z \leftarrow_{\$} \mathcal{A}_0(1^\lambda, sk); (m, \text{st}) \leftarrow_{\$} \mathcal{A}_1(1^\lambda, pk, z)$ 
 $c_0 \leftarrow_{\$} \mathcal{E}.\text{Enc}(m)$ 
 $r \leftarrow_{\$} \{0, 1\}^{|m|}; c_1 \leftarrow_{\$} \mathcal{E}.\text{Enc}(r)$ 
 $b' \leftarrow_{\$} \mathcal{A}_2(1^\lambda, c_b, \text{st})$ 
return  $(1 = b')$ 

```

Canetti et al. show that such a strong form of symmetric encryption exists if, and only if, certain types of MB-AIPOs exist. Their type of MB-AIPO requires that the message m (point value) is drawn independently from the point x (point address). Their notion of MB-AIPO for *independent messages* is weaker than our notion of MB-AIPO against strongly computationally unpredictable distributions as presented in Definition 4.3 and, in particular, not affected by our impossibility result.

We improve the result by Canetti et al. by building a leakage-resilient encryption scheme, that is public key rather than symmetric key. We first present the variant of IND-CPA security with hard-to-invert key-leakage of a public-key encryption scheme \mathcal{E} that we consider (we give the pseudocode in Figure 2). In the IND-CPA game with adversary $(\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$ an initial adversary \mathcal{A}_0 takes as input the secret key and outputs some leakage z . Adversary \mathcal{A}_1 is run on input the public key pk and leakage z and outputs a single messages m together with some state st . Then, according to a secret bit b either message m or a uniformly random message m' of the same length is encrypted yielding ciphertext c which is given together with state st to the final adversary \mathcal{A}_2 which needs to guess bit b .¹⁰

Construction 5.2. *Let λ be the security parameter, let AIPO denote a point obfuscator and iO an indistinguishability obfuscator. Key generation picks a secret key $x \leftarrow_{\$} \{0, 1\}^\lambda$ as a uniformly random bit string of length λ . As public key it outputs a point obfuscation of x , that is, it outputs $pk \leftarrow_{\$} \text{AIPO}(x)$. To encrypt a message m one constructs the circuit*

$$C[pk, m](x^*) := \text{if } (pk(x^*) = 1) \text{ then return } m \text{ else return } \perp$$

and computes an indistinguishability obfuscation $c \leftarrow_{\$} \text{iO}(C[pk, m])$ which yields the ciphertext c . For decryption one computes $m \leftarrow c(x)$.

Correctness of the scheme follows from the correctness criteria of indistinguishability obfuscation and AIPO. We reduce IND-CPA security of the scheme to the security of weak MB-AIPO (note, that an encryption is nothing but an obfuscation following Construction 4.4).

Proposition 5.3. *If Construction 4.7 is a weak MB-AIPO, then Construction 5.2 is IND-CPA secure in the presence of computationally uninvertible leakage on the secret-key.*

Proof. Assume that there exists a successful adversary $\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2$. We are going to construct an adversary $\mathcal{B}_1, \mathcal{B}_2$ against MB-AIPO.

Adversary \mathcal{B}_1 selects a random point x and runs adversary \mathcal{A}_0 on input the security parameter and x to receive some leakage z . It constructs an AIPO obfuscation $pk \leftarrow_{\$} \text{AIPO}(x)$ and runs adversary \mathcal{A}_1 on input $(1^\lambda, pk, z)$ to receive a message m and state st . It outputs $((x, m), \text{st})$. Adversary \mathcal{B}_2 gets as input an obfuscation c that is either equal to $\text{MB-AIPO}(x, m)$ or equal to $\text{MB-AIPO}(x, m')$ as well as state st . It runs adversary \mathcal{A}_2 on input $(1^\lambda, c, \text{st})$ and outputs whatever \mathcal{A}_2 outputs.

¹⁰The described real-or-random notion of IND-CPA can be shown to be equivalent upto a factor of 2 in the reduction to the more frequently used left-or-right security notion [BDJR97].

ANALYSIS. Our analysis proceeds in two steps. First, we show that if $\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2$ is successful, then so is $\mathcal{B}_1, \mathcal{B}_2$. Then, we prove that \mathcal{B}_1 implements a strongly unpredictable distribution as required for weak MB-AIPO.

To see that $\mathcal{B}_1, \mathcal{B}_2$ are successful, we observe that the simulation is perfect. It remains to show that \mathcal{B}_1 implements a strongly unpredictable distribution. We reduce to the unpredictability of \mathcal{A}_0 . Let \mathcal{P} be a predictor against the strong unpredictability of \mathcal{B}_1 , then we construct a predictor \mathcal{R} against the unpredictability of \mathcal{A}_0 .

\mathcal{R} receives the leakage z that was created by \mathcal{A}_0 . It draws a random point u and sets $pk \leftarrow \text{AIPO}(u)$. Then, the predictor \mathcal{R} runs \mathcal{A}_1 on (z, pk) to obtain a message m and some state st . It runs \mathcal{P} on (m, st) to get a value x' and returns x' . Now, \mathcal{A}_1 gets as input $\text{AIPO}(u)$ instead of $\text{AIPO}(x)$. We argue that, assuming the security of AIPO, \mathcal{P} is also successful on this distribution. Assume that \mathcal{P} has non-negligible probability ν in returning x when \mathcal{A}_1 is run on $(z, \text{AIPO}(x))$, but negligible probability ν in returning x when \mathcal{A}_1 is run on $(z, \text{AIPO}(u))$ for a random u . Then, we construct an adversary $(\mathcal{C}_1, \mathcal{C}_2)$ against AIPO as follows. Adversary \mathcal{C}_1 runs \mathcal{A}_0 to create (z, x) . Then, \mathcal{C}_1 draws a random string r and sets b to be the inner product of r and x , that is, $b = \langle r, x \rangle$. Adversary \mathcal{C}_1 returns $((z, r, b), x)$. The second stage \mathcal{C}_2 gets $((z, r, b), p)$ and runs \mathcal{A}_1 on (z, p) to obtain a message m and some state st . It runs \mathcal{P} on (m, st) to get a value x' . It checks whether $p(x') = 1$. If no, it returns a random bit. If yes, then it returns 1 if and only if the inner product of x' and r is equal to b .

Firstly, the first stage \mathcal{C}_1 is unpredictable because \mathcal{A}_0 is and b is only a single bit of x . Now, let us see that $(\mathcal{C}_1, \mathcal{C}_2)$ are also successful. If p is a point obfuscation of x , then \mathcal{P} returns x with probability ν and thus, \mathcal{C}_2 returns 1 with probability $\nu + (1 - \nu) \cdot \frac{1}{2}$. If p is not a point obfuscation of x , then, independently of the behaviour of \mathcal{P} , \mathcal{C}_2 returns 1 with probability $\frac{1}{2}$ thus yielding an overall advantage of $\frac{\nu}{2}$. The formal treatment is identical to the advantage computation within Claim 4.6. This concludes the proof. \square

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.

References

- [ABG⁺13] Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. <http://eprint.iacr.org/2013/689>. (Cited on pages 4, 8, and 12.)
- [AGIS14] Prabhanjan Ananth, Divya Gupta, Yuval Ishai, and Amit Sahai. Optimizing obfuscation: Avoiding barrington’s theorem. Cryptology ePrint Archive, Report 2014/222, 2014. <http://eprint.iacr.org/>. (Cited on page 9.)
- [BC10] Nir Bitansky and Ran Canetti. On strong simulation and composable point obfuscation. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture*

- Notes in Computer Science*, pages 520–537, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Berlin, Germany. (Cited on pages 3, 4, 7, 10, and 20.)
- [BC14] Nir Bitansky and Ran Canetti. On strong simulation and composable point obfuscation. *J. Cryptology*, Volume 27(2):317–357, 2014. (Cited on pages 4 and 7.)
- [BCC⁺14] Nir Bitansky, Ran Canetti, Henry Cohn, Shafi Goldwasser, Yael Tauman Kalai, Omer Paneth, and Alon Rosen. The impossibility of obfuscation with auxiliary input or a universal simulator. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 71–89, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Berlin, Germany. (Cited on pages 8, 9, and 10.)
- [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 52–73, San Diego, CA, USA, February 24–26, 2014. Springer, Berlin, Germany. (Cited on pages 4, 8, 12, 13, and 24.)
- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 505–514, New York, NY, USA, May 31 – June 3, 2014. ACM Press. (Cited on page 9.)
- [BDJR97] Mihir Bellare, Anand Desai, Eric Jokipii, and Phillip Rogaway. A concrete security treatment of symmetric encryption. In *38th Annual Symposium on Foundations of Computer Science*, pages 394–403, Miami Beach, Florida, October 19–22, 1997. IEEE Computer Society Press. (Cited on page 27.)
- [BFM14] Christina Brzuska, Pooya Farshim, and Arno Mittelbach. Indistinguishability obfuscation and UCES: The case of computationally unpredictable sources. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 188–205, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Berlin, Germany. (Cited on pages 5, 6, 8, 15, 16, and 20.)
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Berlin, Germany. (Cited on pages 3, 4, 6, 9, 11, and 12.)
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, May 2012. (Cited on pages 3, 4, 6, 9, 11, and 12.)
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014: 17th International Workshop on Theory and Practice in Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 501–519, Buenos Aires, Argentina, March 26–28, 2014. Springer, Berlin, Germany. (Cited on page 8.)

- [BGK⁺14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 221–238, Copenhagen, Denmark, May 11–15, 2014. Springer, Berlin, Germany. (Cited on page 9.)
- [BHK13a] Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Instantiating random oracles via UCEs. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 398–415, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Berlin, Germany. (Cited on pages 5, 8, 15, and 20.)
- [BHK13b] Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Personal communication. Sep, 2013. (Cited on page 6.)
- [BM14] Christina Brzuska and Arno Mittelbach. Using indistinguishability obfuscation via uces. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014*, *Lecture Notes in Computer Science*, pages ??–??, Kaohsiung, Taiwan, December 7–11, 2014. Springer, Berlin, Germany. (Cited on pages 4, 8, 9, 20, 21, and 22.)
- [BP12] Nir Bitansky and Omer Paneth. Point obfuscation and 3-round zero-knowledge. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 190–208, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Berlin, Germany. (Cited on pages 3, 4, 7, 10, 13, 14, 15, 19, 20, and 34.)
- [BP13] Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. *Cryptology ePrint Archive*, Report 2013/703, 2013. <http://eprint.iacr.org/2013/703>. (Cited on page 9.)
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Berlin, Germany. (Cited on page 18.)
- [BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 1–25, San Diego, CA, USA, February 24–26, 2014. Springer, Berlin, Germany. (Cited on page 9.)
- [BST14] Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Poly-many hardcore bits for any one-way function. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014*, *Lecture Notes in Computer Science*, pages ??–??, Kaohsiung, Taiwan, December 7–11, 2014. Springer, Berlin, Germany. (Cited on pages 4 and 8.)
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology – ASIACRYPT 2013, Part II*, volume 8270 of *Lecture Notes in Computer Science*, pages 280–300, Bangalore, India, December 1–5, 2013. Springer, Berlin, Germany. (Cited on page 8.)

- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 480–499, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Berlin, Germany. (Cited on pages 4 and 8.)
- [Can97] Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In Burton S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO’97*, volume 1294 of *Lecture Notes in Computer Science*, pages 455–469, Santa Barbara, CA, USA, August 17–21, 1997. Springer, Berlin, Germany. (Cited on pages 3, 7, 8, 13, 20, and 34.)
- [CD08] Ran Canetti and Ronny Ramzi Dakdouk. Obfuscating point functions with multibit output. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 489–508, Istanbul, Turkey, April 13–17, 2008. Springer, Berlin, Germany. (Cited on pages 3, 4, 6, 7, 14, 19, 20, and 35.)
- [CFPR14] Ran Canetti, Benjamin Fuller, Omer Paneth, and Leonid Reyzin. Key derivation from noisy sources with more errors than entropy. Cryptology ePrint Archive, Report 2014/243, 2014. <http://eprint.iacr.org/>. (Cited on pages 7 and 10.)
- [CKVW10] Ran Canetti, Yael Tauman Kalai, Mayank Varia, and Daniel Wichs. On symmetric encryption and point obfuscation. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 52–71, Zurich, Switzerland, February 9–11, 2010. Springer, Berlin, Germany. (Cited on pages 3, 6, 7, 16, 20, and 26.)
- [CMR98] Ran Canetti, Daniele Micciancio, and Omer Reingold. Perfectly one-way probabilistic hash functions (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 131–140, Dallas, Texas, USA, May 23–26, 1998. ACM Press. (Cited on pages 3 and 20.)
- [CTL97] Christian Collberg, Clark Thomborson, and Douglas Low. A taxonomy of obfuscating transformations. Technical Report 148, Department of Computer Science, University of Auckland, July 1997. (Cited on page 3.)
- [DKL09] Yevgeniy Dodis, Yael Tauman Kalai, and Shachar Lovett. On cryptography with auxiliary input. In Michael Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 621–630, Bethesda, Maryland, USA, May 31 – June 2, 2009. ACM Press. (Cited on page 3.)
- [Fis99] Marc Fischlin. Pseudorandom function tribe ensembles based on one-way permutations: Improvements and applications. In Jacques Stern, editor, *Advances in Cryptology – EUROCRYPT’99*, volume 1592 of *Lecture Notes in Computer Science*, pages 432–445, Prague, Czech Republic, May 2–6, 1999. Springer, Berlin, Germany. (Cited on pages 3 and 20.)
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press. (Cited on pages 3, 4, and 11.)

- [GGHR14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure MPC from indistinguishability obfuscation. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 74–94, San Diego, CA, USA, February 24–26, 2014. Springer, Berlin, Germany. (Cited on pages 4 and 8.)
- [GGHW14] Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 518–535, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Berlin, Germany. (Cited on pages 9 and 10.)
- [GK96] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, February 1996. (Cited on page 7.)
- [GK05] Shafi Goldwasser and Yael Tauman Kalai. On the impossibility of obfuscation with auxiliary input. In *46th Annual Symposium on Foundations of Computer Science*, pages 553–562, Pittsburgh, PA, USA, October 23–25, 2005. IEEE Computer Society Press. (Cited on pages 3, 4, 9, and 11.)
- [GLSW14] Craig Gentry, Allison Lewko, Amit Sahai, and Brent Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. Cryptology ePrint Archive, Report 2014/309, 2014. <http://eprint.iacr.org/2014/309>. (Cited on pages 4 and 9.)
- [GR07] Shafi Goldwasser and Guy N. Rothblum. On best-possible obfuscation. In Salil P. Vadhan, editor, *TCC 2007: 4th Theory of Cryptography Conference*, volume 4392 of *Lecture Notes in Computer Science*, pages 194–213, Amsterdam, The Netherlands, February 21–24, 2007. Springer, Berlin, Germany. (Cited on pages 3 and 20.)
- [Had10] Satoshi Hada. Secure obfuscation for encrypted signatures. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 92–112, French Riviera, May 30 – June 3, 2010. Springer, Berlin, Germany. (Cited on page 3.)
- [HMLS07] Dennis Hofheinz, John Malone-Lee, and Martijn Stam. Obfuscation for cryptographic purposes. In Salil P. Vadhan, editor, *TCC 2007: 4th Theory of Cryptography Conference*, volume 4392 of *Lecture Notes in Computer Science*, pages 214–232, Amsterdam, The Netherlands, February 21–24, 2007. Springer, Berlin, Germany. (Cited on page 3.)
- [Hof14] Dennis Hofheinz. Fully secure constrained pseudorandom functions using random oracles. Cryptology ePrint Archive, Report 2014/372, 2014. <http://eprint.iacr.org/>. (Cited on page 8.)
- [HRsV07] Susan Hohenberger, Guy N. Rothblum, abhi shelat, and Vinod Vaikuntanathan. Securely obfuscating re-encryption. In Salil P. Vadhan, editor, *TCC 2007: 4th Theory of Cryptography Conference*, volume 4392 of *Lecture Notes in Computer Science*, pages 233–252, Amsterdam, The Netherlands, February 21–24, 2007. Springer, Berlin, Germany. (Cited on page 3.)

- [HSW14] Susan Hohenberger, Amit Sahai, and Brent Waters. Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 201–220, Copenhagen, Denmark, May 11–15, 2014. Springer, Berlin, Germany. (Cited on pages 4 and 8.)
- [Imp95] R. Impagliazzo. A personal view of average-case complexity. In *Proceedings of the 10th Annual Structure in Complexity Theory Conference (SCT’95)*, SCT ’95, pages 134–, Washington, DC, USA, 1995. IEEE Computer Society. (Cited on page 9.)
- [KMN⁺14] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. Cryptology ePrint Archive, Report 2014/347, 2014. <http://eprint.iacr.org/>. (Cited on page 9.)
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13: 20th Conference on Computer and Communications Security*, pages 669–684, Berlin, Germany, November 4–8, 2013. ACM Press. (Cited on page 8.)
- [KRW13] Venkata Koppula, Kim Ramchen, and Brent Waters. Separations in circular security for arbitrary length key cycles. Cryptology ePrint Archive, Report 2013/683, 2013. <http://eprint.iacr.org/2013/683>. (Cited on page 9.)
- [LPS04] Ben Lynn, Manoj Prabhakaran, and Amit Sahai. Positive results and techniques for obfuscation. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 20–39, Interlaken, Switzerland, May 2–6, 2004. Springer, Berlin, Germany. (Cited on pages 7, 8, and 19.)
- [MH14] Takahiro Matsuda and Goichiro Hanaoka. Chosen ciphertext security via point obfuscation. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 95–120, San Diego, CA, USA, February 24–26, 2014. Springer, Berlin, Germany. (Cited on pages 3, 4, 6, 7, 9, 10, 14, 15, 19, and 20.)
- [MO14] Antonio Marcedone and Claudio Orlandi. Obfuscation \implies (ind-cpa security $= / \implies$ circular security). In Michel Abdalla and Roberto De Prisco, editors, *Security and Cryptography for Networks*, volume 8642 of *Lecture Notes in Computer Science*, pages 77–90. Springer International Publishing, 2014. (Cited on page 9.)
- [PST14] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 500–517, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Berlin, Germany. (Cited on pages 4 and 9.)
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press. (Cited on pages 4, 6, and 8.)

[Wee05] Hoeteck Wee. On obfuscating point functions. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 523–532, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press. (Cited on pages 3, 20, and 34.)

A Constructions of Point Obfuscation Schemes

In the following we present the two constructions (and their underlying assumptions) of AIPOs secure with respect to Definition 2.7. These constructions are, to the best of our knowledge, the only two candidates that achieve AIPO security. Other constructions, such as the construction by Wee [Wee05] either do not consider auxiliary information or put additional restrictions on the auxiliary information.

The first construction is due to Canetti [Can97] who bases his construction on a strong variant of the DDH assumption. We here present the construction in the formulation of [BP12] and then present the assumption it is based on.

Construction A.1 (AIPO obfuscator due to [Can97]). *Let $\mathcal{G} := \{\mathbb{G}_\lambda\}_{\lambda \in \mathbb{N}}$ be a group ensemble, where each \mathbb{G}_λ is a group of prime order $p_\lambda \in (2^{\lambda-1}, 2^\lambda)$. We define an obfuscator \mathcal{O} for points in the domain \mathbb{Z}_{p_λ} as follows: $x \mapsto C(r, r^x)$, where $r \leftarrow_s \mathbb{G}_\lambda$ is a random generator of \mathbb{G}_λ , and $C(r, r^x)$ is a circuit which on input i , checks whether $r^x = r^i$.*

Assumption A.2 ([Can97],[BP12]). *There exists an ensemble of prime order groups $\mathcal{G} := \{\mathbb{G}_\lambda\}_{\lambda \in \mathbb{N}}$ such that for any unpredictable distribution $\mathcal{D} = \{D_\lambda = (Z_\lambda, Y_\lambda)\}_{\lambda \in \mathbb{N}}$ with support $\{0, 1\}^{\text{poly}(\lambda)} \times \mathbb{Z}_{p_\lambda}$, it holds that for all PPT algorithms \mathcal{A} there exists a negligible function negl such that*

$$\left| \Pr_{r \leftarrow_s \mathbb{G}_\lambda, z \leftarrow_s D_\lambda} [\mathcal{A}(z, r, r^x) = 1] - \Pr_{r \leftarrow_s \mathbb{G}_\lambda, z \leftarrow_s Z_\lambda, u \leftarrow_s \mathbb{Z}_{p_\lambda}} [\mathcal{A}(z, r, r^u) = 1] \right| \leq \text{negl}(\lambda)$$

The second candidate construction for AIPO is due to Bitansky and Paneth [BP12] who adapt the point obfuscation scheme of Wee [Wee05] to allow for auxiliary input. Their construction is based on an assumption on the existence of strong pseudorandom permutations. Let us recall the underlying assumption (which generalizes the original assumption due to Wee [Wee05]) before recalling the construction.

Assumption A.3 ([BP12]). *There exists an ensemble of permutation families $\mathcal{F} = \{\mathcal{F}_\lambda = \{f\}\}$ such that for any unpredictable distribution ensemble $\mathcal{D} = \{D_\lambda = (Z_\lambda, Y_\lambda)\}_{\lambda \in \mathbb{N}}$, the following two distribution ensembles are also unpredictable:*

- $((Z_\lambda, f(Y_\lambda), f); Y_\lambda)$
- $((Z_\lambda, f); f(Y_\lambda)),$

where in both $f \leftarrow_s \mathcal{F}_\lambda$ (independently of D_λ).

Based on Assumption A.3, Bitansky and Paneth show that the following construction yields an AIPO obfuscator satisfying Definition 2.7 [BP12].

Construction A.4 ([BP12]). *Let \mathcal{F} be a family of permutations as given by Assumption A.3. AIPO obfuscator \mathcal{O} works as follows: given a point $y \in \{0, 1\}^\lambda$, \mathcal{O} samples 3λ permutations $\{f_i\}_{i \in [3\lambda]}$ from \mathcal{F}_λ and 3λ strings $\{r_i\}_{i \in [3\lambda]}$ from $\{0, 1\}^\lambda$. For every $i \in [3\lambda]$, let $f^i := f_i \circ f_{i-1} \circ \dots \circ f_1$ (where \circ denotes composition). Obfuscator \mathcal{O} outputs a circuit C_y that has hardcoded into it the randomness of \mathcal{O} , $\{f_i, r_i\}_{i \in [3\lambda]}$ and the bits $\{b_i := \langle r_i, f^i(y) \rangle\}_{i \in [3\lambda]}$, where $\langle \cdot, \cdot \rangle$ denotes the inner product over \mathbb{GF}_2 . Circuit C_y outputs 1 on a point x if for all $i \in [3\lambda] : b_i = \langle r_i, f^i(x) \rangle$; and 0 otherwise.*

FROM AIPO TO MB-AIPO. Constructions of point obfuscation schemes for point functions with multi-bit output have first been studied by Canetti and Dakdouk [CD08] who show that composability of plain AIPOs is a sufficient condition for the existence of MB-AIPOs. To the best of our knowledge no direct constructions of MB-AIPOs have been proposed in the literature.