

On Symmetric Encryption and Point Obfuscation

Ran Canetti* Yael Tauman Kalai† Mayank Varia‡ Daniel Wichs§

January 30, 2010

Abstract

We show tight connections between several cryptographic primitives, namely encryption with weakly random keys, encryption with key-dependent messages (KDM), and obfuscation of point functions with multi-bit output (which we call multi-bit point functions, or MBPFs, for short). These primitives, which have been studied mostly separately in recent works, bear some apparent similarities, both in the flavor of their security requirements and in the flavor of their constructions and assumptions. Still, rigorous connections have not been drawn.

Our results can be interpreted as indicating that MBPF obfuscators imply a very strong form of encryption that *simultaneously* achieves security for weakly-random keys and key-dependent messages as special cases. Similarly, each one of the other primitives implies a certain restricted form of MBPF obfuscation. Our results carry both constructions and impossibility results from one primitive to others. In particular:

- The recent impossibility result for KDM security of Haitner and Holenstein (TCC '09) carries over to MBPF obfuscators.
- The Canetti-Dakdouk construction of MBPF obfuscators based on a strong variant of the DDH assumption (EC '08) gives an encryption scheme which is secure w.r.t. *any* weak key distribution of super-logarithmic min-entropy (and in particular, also has very strong leakage resilient properties).
- All the recent constructions of encryption schemes that are secure w.r.t. weak keys imply a weak form of MBPF obfuscators.

*School of Computer Science, Tel Aviv University (canetti@cs.tau.ac.il)

†Microsoft Research New England (yael@microsoft.com)

‡Massachusetts Institute of Technology (varia@csail.mit.edu)

§New York University (wichs@cs.nyu.edu)

1 Introduction

Symmetric encryption is an algorithmic tool that allows a pair of parties to communicate secret information over open communication media that are accessible to eavesdroppers. In order to achieve this goal, the communicating parties need to have some shared secret randomness (a *key*). The classic view of symmetric encryption allows the encryption scheme to determine the distribution of the key precisely (typically it is a uniformly random string). It also assumes that the encryption and decryption algorithms are executed in a completely sealed way, so no information about the key is leaked to the eavesdroppers. Finally, the classic model assumes that the parties only use the key in the encryption and decryption routines and not for any other purpose. In particular, their messages are never related to the key.

In recent years, much research has been done to investigate various relaxations of this classic (and somewhat naïve) model. One relaxation is to consider the case where the key is chosen using a “defective” source of randomness that does not generate uniform and independent random bits. (See e.g. [1, 2, 13, 20, 24] and the references therein). Namely, the key is assumed to be taken from a distribution that is adversarially chosen under some restriction. Typically the restriction is that the min-entropy of the distribution of the secret key is at least α , for some value of α . In this case the scheme is said to be secure w.r.t. α -weak keys.

A different relaxation of the classic model considers the case where the key is chosen uniformly but some *arbitrary* information on the key is leaked to the adversary (see e.g. [1, 24]). This models both direct attacks where the adversary gains access to the internal storage of the parties, such as the cold-boot attack of [17], and indirect information leakage that occurs when the shared key is derived from the communication between the parties, such as the information exchange used to agree on the key. Of course, all security is lost if the adversary learns the key in its entirety, and therefore some restriction needs to be imposed on the *amount* of information that the adversary can get. One possibility is to require that the key has some significant statistical entropy left, even given the leakage. We call this the *entropic* setting. Another, stronger, security notion only insists that it is *computationally* infeasible to compute the secret key from the leaked information, but allows the leakage to completely determine the key statistically. We call this type of leakage *auxiliary input*.¹ It turns out that encryption resilient to weak keys is also resilient to a comparable amount of leakage in the entropic setting. Conversely, in some settings there is a simple transformation from leakage resilient encryption to one that withstands comparably weak keys.²

Yet another relaxation of the classic model considers the case where the messages may depend on the shared key. Security in this more demanding setting was termed *key-dependent message security* (KDM security) by Black, Rogaway and Shrimpton in [7]. In the last few years, the notion of KDM security has been extensively studied [3, 4, 5, 8, 9, 16, 18, 19], and several positive results emerged, most notably the results of [3, 8] who showed how to obtain KDM security w.r.t. the class of affine functions (the former under the DDH assumption and the latter under the LWE assumption). In contrast, [16] show that there exist no black-box reductions from the KDM security of any encryption scheme w.r.t. all efficient functions to “any standard cryptographic assumption.”

While the constructions for KDM-secure schemes and the constructions of schemes that are secure w.r.t. α -weak keys bear significant similarities to each other (eg., see [8, 24], [3, 13], and [1, 3]), no formal connections between the problems have been made so far.

Another recently studied primitive, which may seem unrelated at a cursory look, is obfuscation of point functions (programs) with multi-bit output. Obfuscation is the task of constructing an algorithm, called an *obfuscator* \mathcal{O} , that takes as input a program p from a family P of programs and outputs a program $q = \mathcal{O}(p)$ that has essentially the same functionality as p , but where the code of q only gives information that can also be determined with oracle access to p . A central point here is that \mathcal{O} should work correctly and securely for *every* program in P .

A point function with multi-bit output (or a MBPF) is a function $I_{(k,m)}$ which, on input x , outputs m if $x = k$ and \perp otherwise. In the special case of point functions, the value m is fixed to some constant, say 1. Obfuscators for point functions are constructed in [10, 25] under strong assumptions, and in [21] in the random oracle model.

¹Many other models of leakage-resilience, such as the “only computation leaks information” model [14, 22], place further restrictions on the type of information that may be leaked, and are not considered in this work.

²In the case of semantic security for symmetric-key encryption (without chosen-plaintext attacks), we can use the following transformation: Given a scheme (Enc, Dec) that’s secure against key leakage, construct the weak-key scheme $(Enc'_k(m) = (r, Enc_{k+r}(m)))$ for a random $|k|$ -bit r , $Dec'_k(r, c) = Dec_{k+r}(c)$.

Obfuscators for MBPFs are only known based on very strong and specific assumptions, such as the existence of fully-composable point function obfuscators [11]. Different constructions exist for restricted settings, such as the case where m is shorter than k , or the case where m and k are distributed independently from each other [11, 13]. In all of these constructions the obfuscator is given the values k and m explicitly.

The applicability of MBPF obfuscation to symmetric encryption has been pointed out in [11], who proposed to encrypt a message m with key k by letting $\mathcal{O}(I_{(k,m)})$ be the ciphertext. The fact that security holds for any k was used to suggest that m remains hidden even when k is taken from a distribution which is not uniform, as long as it has sufficient min-entropy that it cannot be guessed in polynomial time. Also, [13] show that their construction of leakage resilient encryption can be used as a restricted variant of MBPF obfuscation.

1.1 Our Results

We show tight relations between the primitives of symmetric-key encryption and MBPF obfuscation. Specifically, we show that symmetric-key encryption with weak key resilience, leakage resilience, and KDM security, each with its own variants, can all be viewed as natural special cases of the MBPF obfuscation problem. In fact, MBPF obfuscation incorporates weak key resilience and KDM security simultaneously. In addition to providing some insight and intuition to these primitives, the drawn connections provide new results – both constructions and hardness results – for the primitives considered.

As a preliminary step towards drawing these connections, we set up a framework for relaxing the standard notion of security of MBPF obfuscation. This notion, called virtual black-box (VBB) security [6], essentially requires that for any adversary with binary output there exists a simulator such that, for any k, m , the output of the adversary given $\mathcal{O}(I_{(k,m)})$ is indistinguishable from the output of the simulator given oracle access to $I_{(k,m)}$. We wish to consider the relaxed case where k and m are taken from an unknown *distribution* from a given class. We capture this relaxation by replacing the “for any k, m ” requirement in the VBB definition with “for any distribution on k, m from a given class of distributions.” Note that the simulator knows the class of distributions, but not the distribution itself. We relate the different notions of encryption to MBPF obfuscation for different classes of distributions.

Obfuscation vs. Weak-Key and Leakage Resilient Encryption. We say that an MBPF obfuscator is α -entropic with independent messages if it satisfies the above definition for product distributions on k, m where the distribution of k has min-entropy at least α . Note that the product distribution ensures that m is drawn *independently* of k , and that we impose no restriction on the entropy of m .

Our first result is that α -entropic MBPF obfuscators with independent messages are equivalent to symmetric key encryption with α -weak keys.³ We describe both directions of the equivalence.

From obfuscation to encryption. Given an obfuscator \mathcal{O} , we construct an encryption scheme by the transformation $\text{Enc}_k(m) = \mathcal{O}(I_{(k,m)})$ and $\text{Dec}_k(c) = c(k)$, where c is interpreted as the description of a circuit.

From encryption to obfuscation. Conversely, given an encryption scheme, we construct an obfuscator as follows. On input the pair k, m , simply encrypt m with key k to obtain a ciphertext c . Then, the obfuscated program simply has c hard-coded, and on input x , runs $\text{Dec}_x(c)$ and outputs the result. Here, for the correctness of obfuscation, we require that the encryption scheme can *detect* if it is decrypting a ciphertext with an incorrect secret key. We show that this property can be added generically to any semantically secure encryption scheme.

CPA security vs. self-composability. If we start with a *CPA secure* encryption scheme, then the resulting MBPF obfuscator \mathcal{O} is *self-composable*, in the sense that security is preserved even if \mathcal{O} is run multiple times with the *same* input k and (possibly) different inputs m_i . As was shown by [11], this property is not, in general, implied by obfuscation alone. The converse is true as well.

³The traditional notion of encryption without leakage resilience, where the secret key is chosen uniformly at random, is captured by the $\alpha(n) = 2^n$ case of the equivalence.

Semantically secure encryption with:	Is equivalent to MBPF obfuscator for:
α -weak keys	α -entropic sec. for indep. messages
auxiliary input	auxiliary input
CPA security	self-composability
fully weak keys	fully-entropic sec.
KDM security	dependent messages

Table 1: Equivalence between symmetric-key encryption (left) and obfuscation (right) terminology. The five rows can be combined arbitrarily, except that row 1 is required along with rows 3-5, and we do not consider the combination of rows 2 and 5.

Fully-entropic obfuscation and fully-weak key security. We say that an MBPF obfuscator for independent messages is *fully entropic* if it satisfies α -entropic security for all super-logarithmic α . If we start with such an obfuscator, the transformation produces an encryption scheme with semantic security for fully-weak keys (i.e. security for any key distribution with super-logarithmic entropy).

To connect our new α -entropic definition to previous works, we show that any MBPF obfuscator that is fully-entropic also satisfies the virtual black-box property, i.e., it works for *any* k, m . We note that the proof of this result is trickier than it might seem, the main difficulty being that in the case of α -entropic security the simulator has the bound α , whereas in the virtual black-box case no such bound exists.

Auxiliary input. If we start from a leakage-resilient encryption w.r.t. auxiliary input, then the resulting MBPF obfuscator is secure with respect to auxiliary input as well, as defined in [15]. The converse is true as well.

KDM security. All of the above equivalence results were stated with respect to the restricted notion of obfuscation to *independent* messages. Interestingly, the standard notion of MBPF obfuscation provides the additional (and very powerful) security guarantee for encryption with *key-dependent messages* (KDM).

We say that \mathcal{O} is an α -entropic (*dependent message*) MBPF obfuscator if it withstands any *joint* distribution on k, m where the projection distribution on k has min-entropy at least α , and the message m may depend on k . Typically, we view m as a function of k . Such an obfuscator is equivalent to an α -KDM semantically secure encryption scheme, via the same transformations as before.

Multiple extensions. Finally, we note that the four extensions of the original connection between obfuscation and encryption can be achieved concurrently, with two caveats. First, when combining CPA and KDM security, we require that the function connecting the message to the key be chosen non-adaptively prior to viewing any ciphertexts. Second, we do not consider KDM security with auxiliary input.

1.2 Implications

We describe some implications of the above correspondence results. See Section 6 for more details.

Secure encryption w.r.t. (fully) weak keys. The known constructions of α -weak key secure encryption schemes require that the bound α be chosen in advance, and then the scheme is constructed based on α . Using our transformations, we get that, under the strong DDH assumption in [10], the [10, 11] MBPF obfuscator provides an encryption scheme that simultaneously achieves α -weak key security for all super-logarithmic functions α . The main advantage of this scheme is that the min-entropy α does not need to be chosen in advance.

We remark that the hardness assumption we use has a similar flavor – it explicitly makes an assumption for every distribution with super logarithmic min-entropy. The crucial point is however that the construction does *not* depend on α and so it provides a tradeoff between the strength of the assumption and the strength of the obtained guarantee.

Impossibility for MBPF obfuscators and fully composable point function obfuscators. Using our transformations, the negative result due to Haitner and Holenstein [16] implies that MBPF obfuscators cannot be proven secure via a “black box reduction to standard cryptographic primitives.” Since full MBPF obfuscators can be constructed in a black-box way from fully composable point function obfuscators [11], the impossibility carries over to this primitive as well.

Constructing self-composable MBPF obfuscators with independent messages. Using our transformations, we can use constructions of encryption schemes that are secure w.r.t. α -weak keys, to get self composable MBPF obfuscators with independent messages. More specifically, we construct self composable obfuscators for MBPFs $\{I_{(k,m)}\}$ as long as the distribution of m is independent of the distribution of k , both distributions are efficiently sampleable, and the distribution of k has min-entropy α .

1.3 Organization

Section 2 contains some basic definitions for obfuscation and encryption. Section 3 draws connections between obfuscation and weak key encryption, for both semantic and CPA security. Section 4 extends the connection to the auxiliary input setting. Section 5 draws connections between obfuscation and KDM encryption. Section 6 states the implications that we draw from the general connections. Some proofs are located in the appendices.

2 Definitions

2.1 Obfuscation of Point Functions with Multi-bit Output

Let $I_{(k,m)} : \{0, 1\}^* \cup \{\perp\} \rightarrow \{0, 1\}^* \cup \perp$ denote the function

$$I_{(k,m)}(x) = \begin{cases} m & \text{if } x = k \\ \perp & \text{otherwise} \end{cases}$$

which outputs the *message* m given the *key* k , and \perp otherwise. Let $\mathcal{I} = \{I_{(k,m)} \mid k, m \in \{0, 1\}^*\}$ be the family of all such functions, which we call the family of *point functions with multi-bit output* or just *multi-bit point functions (MBPF)* for short.

Definition 2.1 (Obfuscation of Point Functions with Multi-bit Output). *A multi-bit point function (MBPF) obfuscator is a PPT algorithm \mathcal{O} which takes as input values (k, m) describing a function $I_{(k,m)} \in \mathcal{I}$ and outputs a circuit C . We will abuse notation and write $\mathcal{O}(I_{(k,m)})$, but will always assume that \mathcal{O} gets k and m as clearly delineated inputs.*

Correctness: For all $(k, m) \in \{0, 1\}^*$ with $|k| = n, |m| = \text{poly}(n)$, all $x \in \{0, 1\}^n$,

$$\Pr[C(x) \neq I_{(k,m)}(x) \mid C \leftarrow \mathcal{O}(I_{(k,m)})] \leq \text{negl}(n)$$

where the probability is taken over the randomness of the obfuscator algorithm.

Polynomial Slowdown: For any k, m , the size of the circuit $C = \mathcal{O}(I_{(k,m)})$ is polynomial in $|k| + |m|$.

Entropic Security: We say that the scheme has $\alpha(n)$ -**entropic security** if for any PPT adversary \mathcal{A} with 1 bit output, any polynomial $\ell(\cdot)$, there exists a PPT simulator \mathcal{S} such that for all jointly-distributed $\{X_n, Y_n\}_{n \in \mathbb{N}}$ where X_n takes values in $\{0, 1\}^n$, Y_n takes values in $\{0, 1\}^{\ell(n)}$ and $H_\infty(X_n) \geq \alpha(n)$, we have:

$$\left| \Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr[\mathcal{S}^{I_{(k,m)}(\cdot)}(1^n) = 1] \right| \leq \text{negl}(n)$$

where the probability is taken over the randomness of $(k, m) \leftarrow (X_n, Y_n)$, the randomness of the obfuscator \mathcal{O} and the randomness of \mathcal{A}, \mathcal{S} . We say that a scheme has **fully-entropic security** if it has $\alpha(n)$ -entropic security for all $\alpha(n) \in \omega(\log(n))$.

We relate the notion of fully-entropic security, defined above, to the standard security guarantee provided by obfuscation called the *virtual black-box property*:

Definition 2.2 (Virtual black-box property [6, 10, 25]). For any PPT adversary \mathcal{A} with 1 bit output and any polynomials $p(\cdot)$, $\ell(\cdot)$, there exists a PPT simulator \mathcal{S} such that for all distributions $\{X_n, Y_n\}_{n \in \mathbb{N}}$ with X_n taking values in $\{0, 1\}^n$ and Y_n taking values in $\{0, 1\}^{\ell(n)}$, we have:

$$|\Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr[\mathcal{S}^{\mathcal{I}_{(k,m)}}(1^n) = 1]| \leq \frac{1}{p(n)}.$$

The probability is taken over the randomness of $(k, m) \leftarrow (X_n, Y_n)$, \mathcal{A} , \mathcal{S} , and \mathcal{O} .

Note the differences between the fully-entropic definition and the VBB definition: the former allows a different simulator for each entropy threshold $\alpha(\cdot)$, but requires a negligible error in simulation, while the latter allows a different simulator for each simulation-error $p(\cdot)$, but requires the simulator to work for all distributions regardless of entropy. Interestingly, we show that the fully-entropic definition implies VBB (but don't know whether the converse holds as well).

Theorem 2.1. If \mathcal{O} is a MBPF obfuscator that satisfies fully-entropic security (as in Definition 2.1) then \mathcal{O} also satisfies virtual black-box obfuscation (as in Definition 2.2).

This theorem is proved in Appendix A. The idea is to extend the technique used in [10] to show that a distribution-based definition implies the virtual black box property in the case of point functions. At a high level, the distributional definition says that if a user chooses a key from a well-spread distribution, then an adversary cannot learn anything from an obfuscated point function beyond the fact that the key is from this distribution, so in particular the key is hard to determine. We show how to extend the distributional definition to the MBPF setting and prove that fully-entropic security implies this distributional requirement, and therefore the virtual black-box property as well.

Fully entropic security, as well as virtual black box security, are quite strong, and difficult to satisfy. The notion of $\alpha(n)$ -entropic security, for some particular $\alpha(n) \in \omega(\log(n))$, corresponds to a meaningful weakening of that notion where security is only provided when the input comes from a reasonably random source. A similar weakening of obfuscation, in the special case of point functions, was also considered by Canetti, Micciancio and Reingold [12] in the context of perfectly one-way hash functions.

Instead of restricting attention to distribution with $\alpha(n)$ min-entropy, one might instead give the simulator the ability to ask its oracle more queries, by a factor of $2^{\alpha(n)}$ (i.e. the simulator is no longer polynomial time). In Appendix B, we show that this alternative relaxed notion is actually implied by α -entropic security.

We consider several additional variants of obfuscation throughout the paper. First, we propose an additional weakening of the definition, which we call security for *independent messages*, and where we require that the distribution on the output m is independent from that of the input k for a point function $I_{(k,m)}$.

Definition 2.3 (Independent Messages). We say that an obfuscator \mathcal{O} is $\alpha(n)$ -entropically secure for independent messages if we restrict the definition of $\alpha(n)$ -entropic security only to distributions $\{X_n, Y_n\}$ where X_n and Y_n are independently distributed. We define the notion of fully-entropic security for independent messages analogously.

We also define a stronger variant of plain obfuscation, which provides some *composability* guarantees. There are two variants: For *full composition* we require that the security of obfuscation is preserved even if the adversary gets (freshly and independently) obfuscated circuits for many functions, where the various obfuscated functions are related in arbitrary ways (i.e., both the keys and the messages may differ). For *self composition* we require that all the obfuscated functions have the same value of the key k . That is, one should obfuscate the functions $I_{(k,m_1)}, \dots, I_{(k,m_t)}$ with the *same key* k but *potentially different messages* m_1, \dots, m_t . (For point functions, self composition boils down to the case of many obfuscated versions of the same function.)

Definition 2.4 (Composability). A multi-bit point function obfuscator \mathcal{O} with $\alpha(n)$ -entropic security is said to be **fully-composable** if for any PPT adversary \mathcal{A} with 1 bit output, any polynomials $t(\cdot), \ell(\cdot)$, there exists a PPT

simulator \mathcal{S} such that for all distributions $\{(X_n, Y_n)\}_{n \in \mathbb{N}}$, where $X_n = X_n^{(1)}, \dots, X_n^{(t)}$, $Y_n = Y_n^{(1)}, \dots, Y_n^{(t)}$, and $X_n^{(i)}$ taking values in $\{0, 1\}^n$, $Y_n^{(i)}$ taking values in $\{0, 1\}^{\ell(n)}$ and $H_\infty(X_n) \geq \alpha(n)$, we have:

$$|\Pr[\mathcal{A}(\mathcal{O}(I_{k_1, m_1}), \dots, \mathcal{O}(I_{k_t, m_t})) = 1] - \Pr[\mathcal{S}^{I_{k_1, m_1}, \dots, I_{k_t, m_t}}(1^n) = 1]| \leq \text{negl}(n),$$

where the probabilities are over $(k_1, \dots, k_t, m_1, \dots, m_t) \leftarrow (X_n, Y_n)$ and over the randomness of $\mathcal{A}, \mathcal{S}, \mathcal{O}$.

If the above holds only for the distributions X_n where $\Pr[k_1 = k_2 \dots = k_t] = 1$, then we say that \mathcal{O} is **self-composable**.

The notions of composability extend naturally to obfuscators with fully-entropic security, where we require that the above definition holds for all $\alpha(n) \in \omega(\log(n))$. It also extends to obfuscators for independent messages, where we restrict the definition to the case where X_n and Y_n are independent. (It is stressed that there is no independence assumption among the coordinates within X_n or Y_n .)

2.2 Definitions for Encryption with Weak Keys

A symmetric encryption scheme consists of efficient algorithms (Enc, Dec) .⁴ We say that the encryption scheme is semantically secure for $\alpha(n)$ -weak keys if the usual notion of semantic security holds even when the key comes from any weak-source of entropy $\alpha(n)$. We propose the following definition of symmetric key encryption with weak keys.

Definition 2.5 (Symmetric Encryption with Weak Keys). We say that an encryption scheme has **CPA security** for $\alpha(n)$ -weak keys if there exists an efficient algorithm $D(n, \ell)$ running in time $\text{poly}(n, \ell)$, such that, for all PPT adversaries \mathcal{A} and all distribution-ensembles $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$, we have:

$$|\Pr[\text{CPA}_0^{X, D}(\mathcal{A}, n) = 1] - \Pr[\text{CPA}_1^{X, D}(\mathcal{A}, n) = 1]| \leq \text{negl}(n)$$

where the games $\text{CPA}_b^{X, D}(\mathcal{A}, n)$ for $b = 0, 1$ are defined via the following experiment:

1. $k \leftarrow X_n$
2. Repeat: \mathcal{A} submits a query m and receives a ciphertext c where:
 - In game $\text{CPA}_0^{X, D}$, the challenger sets $c \leftarrow \text{Enc}_k(m)$.
 - In game $\text{CPA}_1^{X, D}$, the challenger sets $c \leftarrow D(n, |m|)$.
3. The output of the game is the output of \mathcal{A} .

The algorithm $D(n, \ell)$ can keep persistent state during stage 2. We define **semantic security** with $\alpha(n)$ -weak keys via the games $\text{SEM}_0^{X, D}, \text{SEM}_1^{X, D}$, which are equivalent to the CPA games except that step (2) is performed only once.

We say that an encryption scheme is CPA-secure (resp. semantically-secure) for **fully weak keys** if it is CPA-secure (resp. semantically-secure) secure for $\alpha(n)$ -weak keys for all $\alpha(n) \in \omega(\log(n))$.

Note that, in case of $\alpha(n) = n$ (i.e. uniformly random secret keys), the above definition is equivalent to the standard notion of CPA/semantic security, since we can always simply define $D(n, \ell)$ to always output fresh encryptions $\text{Enc}_k(0^\ell)$, where k is initially chosen uniformly at random and re-used for all queries. On the other hand, when considering $\alpha(n)$ -weak keys, the above definition is somewhat stronger than just requiring that the adversary cannot distinguish between an encryption of m and that of some set message, such as 0^ℓ . In particular, it requires that there is a single *universal* distribution D on ciphertexts, which is indistinguishable from encryption with *any* key distribution X_n of sufficient entropy. For example, consider an encryption scheme which, along with the ciphertext, always outputs the first bit of the secret key. Although such scheme might satisfy a natural definition where encryption of m_0 and m_1 are indistinguishable, it could never satisfy the above definition, even for $\alpha(n) = n - 1$. The reason is that the ciphertext distribution is now different depending on whether the keys come from a distribution that fixes the first bit at 0 versus one which fixes the first bit at 1. Although our definition is

⁴That is, the key generation algorithm is implicit and is assumed to always generate a uniform n -bit string.

stronger than one may need, we will show that it is necessary and sufficient for our equivalence with obfuscation to hold. Moreover, all natural constructions of encryption schemes with weak-keys that we know of achieve the above definition.

We also define a “wrong-key detection” property, which will be needed to achieve correctness in obfuscation.

Definition 2.6 (Wrong-Key Detection). *We say that an encryption scheme satisfies the wrong-key detection property if for all $k \neq k' \in \{0, 1\}^n$, all $m \in \{0, 1\}^{\text{poly}(n)}$, $\Pr[\text{Dec}_{k'}(\text{Enc}_k(m)) \neq \perp] \leq \text{negl}(n)$.*

We note that a similar, but weaker, property called confusion freeness, was defined in [23]. For confusion freeness, the keys k, k' are random and independent, while we consider a worst-case choice of k, k' and the probability above is only over the randomness of the encryption scheme.

Lemma 2.1 shows that, in the case of semantic security, wrong-key detection can always be achieved via a simple transformation. We note, however, that this transformation no longer works in the case of CPA security.

Lemma 2.1. *Let (Enc, Dec) be a semantically-secure encryption scheme for $\alpha(n)$ -weak keys and let \mathcal{H} be a pairwise-independent permutation family. Define an encryption scheme $(\text{Enc}', \text{Dec}')$ by:*

$$\begin{aligned} \text{Enc}'_k(m) &\triangleq \begin{cases} \text{Choose: } h \leftarrow \mathcal{H}, r \leftarrow U_n \\ \text{Output: } \langle r, h, c = \text{Enc}_{h(k)}(r||m) \rangle \end{cases} \\ \text{Dec}'_k(\langle r, h, c \rangle) &\triangleq \begin{cases} \text{Compute: } (r' || m') = \text{Dec}_{h(k)}(c) \\ \text{Output: } m' \text{ if } r' = r \text{ and } \perp \text{ otherwise} \end{cases} \end{aligned}$$

Then $(\text{Enc}', \text{Dec}')$ is a semantically-secure encryption scheme for $\alpha(n)$ -weak keys, with wrong-key detection. The above also holds if we replace “ $\alpha(n)$ ” with “fully”.

Proof. Let us first show that the modification preserves semantic security for $\alpha(n)$ -weak keys. Let $D(n, \ell)$ be the distribution for which the semantic security of (Enc, Dec) is satisfied, and define $D'(n, \ell) = D(n, \ell + n)$. Then, for any adversary \mathcal{A} attacking the modified scheme $(\text{Enc}', \text{Dec}')$, and any distribution-ensemble $\{X_n\}_{n \in \mathbb{N}}$ we have

$$\begin{aligned} &|\Pr[\text{SEM}_0(\mathcal{A}, n) = 1] - \Pr[\text{SEM}_1(\mathcal{A}, n) = 1]| \\ &= \left| \Pr \left[\begin{array}{c} m \leftarrow \mathcal{A}(1^n) \\ r \leftarrow U_n, h \leftarrow \mathcal{H}, k \leftarrow X_n, c \leftarrow \text{Enc}_{h(k)}(r||m_0) \\ \mathcal{A}(1^n, c) = 1 \end{array} \right] - \Pr \left[\begin{array}{c} m \leftarrow \mathcal{A}(1^n) \\ c \leftarrow D'(n, m) = D(n, |m| + n) \\ \mathcal{A}(1^n, c) = 1 \end{array} \right] \right| \\ &\leq \max_{r \in \{0,1\}^n, h \in \mathcal{H}} \left| \Pr \left[\begin{array}{c} m, \leftarrow \mathcal{A}(1^n) \\ k \leftarrow h(X_n), c \leftarrow \text{Enc}_k(r||m) \\ \mathcal{A}(1^n, c) = 1 \end{array} \right] - \Pr \left[\begin{array}{c} m \leftarrow \mathcal{A}(1^n) \\ c \leftarrow D(n, |m| + n) \\ \mathcal{A}(1^n, c) = 1 \end{array} \right] \right| \\ &\leq \text{negl}(n) \end{aligned} \tag{1}$$

where the last inequality simply follows from the semantic-security of the original (Enc, Dec) scheme and noting that, for any fixed permutation h , the distribution $h(X_n)$ has the same entropy as X_n .

Now we show that modified scheme has *wrong-key detection*. Assume otherwise, that there is some polynomial $p(\cdot)$ and infinitely many values n for which there exists $k \neq k' \in \{0, 1\}^n$, $m \in \{0, 1\}^{\text{poly}(n)}$ such that

$$\begin{aligned} \Pr[\text{Dec}'_{k'}(\text{Enc}_k(m)) \neq \perp] &= \Pr_{h \leftarrow \mathcal{H}, r \leftarrow U_n} [\text{Dec}_{h(k')}(\text{Enc}_{h(k)}(r||m)) \text{ has a prefix } r] \\ &= \Pr_{k \leftarrow U_n, k' \leftarrow U_n, r \leftarrow U_n} [\text{Dec}_{k'}(\text{Enc}_k(r||m)) \text{ has a prefix } r] \geq p(n) \end{aligned} \tag{2}$$

We now show that (2) contradicts the semantic security of (Enc, Dec) (even for uniform keys). In particular, consider an adversary \mathcal{A} which queries the challenger on the messages $m^* = r||m$, for a random r and for the m which contradicts correctness and satisfies (2). On input c , the adv. \mathcal{A} picks a random $k' \leftarrow U_n$ and outputs 1 iff $\text{Dec}_{k'}(c)$ begins with r . By (2), we see that in the semantic-security game SEM_0 (where the encryption is of the message m) \mathcal{A} outputs 1 with probability $p(n)$. On the other hand, in SEM_1 , no matter what distribution D the challenger

samples from, the result is independent of r and therefore, the probability that \mathcal{A} outputs 1 is the probability that $\text{Dec}_{k'}(\text{Enc}_k(m_0))$ begins with r , which is at most $1/2^n$. Therefore the adversary \mathcal{A} has a non-negligible advantage in the semantic-security game (for any D), which gives a contradiction. \square

3 Encryption with Weak Keys and MBPF Obfuscation

3.1 Sem. Sec. Encryption and Obfuscation with Independent Messages

In this section, we show equivalence between semantically secure encryption with weak keys and MBPF obfuscators for *independent messages*.

Theorem 3.1. *Let $\alpha(n) \in \omega(\log(n))$. There exist MBPF obfuscators with $\alpha(n)$ -entropic security for independent messages if and only if there exist semantically secure encryption schemes with wrong key detection for $\alpha(n)$ -weak keys. Furthermore, the above also holds if we replace “ $\alpha(n)$ ” with “fully”.*

We prove the “if” and “only if” directions in Lemmas 3.1 and 3.2, respectively.

Lemma 3.1. *Let $\alpha(n) \in \omega(\log(n))$ and let \mathcal{O} be a MBPF obfuscator with $\alpha(n)$ -entropic security for independent messages. Let $\text{Enc}_k(m) \triangleq \mathcal{O}(I_{(k,m)})$, $\text{Dec}_k(C) \triangleq C(k)$ where the ciphertext C is interpreted as a circuit. Then the encryption scheme (Enc, Dec) is semantically secure with $\alpha(n)$ -weak keys and has the wrong-key detection property.*

Proof. The correctness of decryption follows from the correctness of obfuscation. For the security of the encryption scheme with $\alpha(n)$ -weak keys. Fix any adversary \mathcal{A} and any distribution $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$. The distribution $\{Y_n\}$ is defined by running $\mathcal{A}(1^n)$ and outputting the message m that \mathcal{A} gives to its challenger. Define the distribution $D(n, \ell) = \mathcal{O}(I_{(k,m)})$ where $(k, m) \leftarrow (U_n, U_\ell)$. Then, by the $\alpha(n)$ -entropic security of obfuscation, there must be a simulator \mathcal{S} such that

$$\begin{aligned}
& \left| \Pr[\text{SEM}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{SEM}_1^{X,D}(\mathcal{A}, n) = 1] \right| \\
&= \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (U_n, U_\ell)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] \right| \\
&\leq \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}(\cdot)}(1^n) = 1] \right| \tag{3} \\
&\quad + \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}}(1^n) = 1] - \Pr_{(k,m) \leftarrow (U_n, U_\ell)} [\mathcal{S}^{I_{(k,m)}}(1^n) = 1] \right| \tag{4} \\
&\quad + \left| \Pr_{(k,m) \leftarrow (U_n, U_\ell)} [\mathcal{S}^{I_{(k,m)}}(1^n) = 1] - \Pr_{(k,m) \leftarrow (U_n, U_\ell)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] \right| \tag{5} \\
&\leq \text{negl}(n)
\end{aligned}$$

where (3),(5) follow by the definition of entropic security of obfuscation, and (4) follows since the only way that a PPT simulator can get anything from its oracle is by querying it on the input k , which happens with negligible probability when k comes from a source of super-logarithmic entropy $\alpha(n)$. \square

Lemma 3.2. *Let (Enc, Dec) be an encryption scheme with semantic security for $\alpha(n)$ -weak keys and with the wrong-key detection property. We define the obfuscator \mathcal{O} which, on input $I_{(k,m)}$, computes a ciphertext $c = \text{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ defined by $C_c(x) \triangleq \text{Dec}_x(c)$. Then the obfuscator \mathcal{O} has $\alpha(n)$ -entropic security for independent messages.*

Proof. First, we show the correctness property of the obfuscator. Fix $k, x \in \{0, 1\}^n$ and $m \in \{0, 1\}^{\text{poly}(n)}$. If $k = x$ then

$$\Pr[C(x) \neq I_{(k,m)}(x) \mid C \leftarrow \mathcal{O}(I_{(k,m)})] = \Pr[\text{Dec}_k(\text{Enc}_k(m)) \neq m] \leq \text{negl}(n)$$

by the correctness of encryption. On the other hand, if $k \neq x$ then

$$\Pr[C(x) \neq I_{(k,m)}(x) \mid C \leftarrow \mathcal{O}(I_{(k,m)})] = \Pr[\text{Dec}_x(\text{Enc}_k(m)) \neq \perp] \leq \text{negl}(n)$$

by the *wrong-key detection* of encryption.

The polynomial slowdown property of the obfuscator follows from the fact that the size of the circuit is only proportional to the ciphertext size and the size of the decryption circuit, which are polynomial in $|k|, |m|$.

Lastly, we show $\alpha(n)$ -entropic security for independent messages. Let $D(n, \ell)$ be the distribution defined by the semantic-security of the encryption scheme. For any polynomial $\ell(n)$ any PPT adversary \mathcal{A} which attacks the obfuscation scheme, we define the simulator \mathcal{S} which chooses a random ciphertext c from the distribution $D(n, \ell(n))$ and runs \mathcal{A} on a circuit C_c constructed using the ciphertext c . Then

$$\begin{aligned} & \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}}(1^n, 1^\ell) = 1] \right| & (6) \\ & = \left| \Pr \left[\mathcal{A}(C_c) = 1 \mid \begin{array}{c} (k, m) \leftarrow (X_n, Y_n) \\ c \leftarrow \text{Enc}_k(m) \end{array} \right] - \Pr [\mathcal{A}(C_c) = 1 \mid c \leftarrow D(n, \ell)] \right| \\ & \leq \text{negl}(n) & (7) \end{aligned}$$

Where (7) follows by semantic-security. \square

3.2 CPA Encryption and Composable Obfuscation for Indep. Messages

In this section, we show equivalence between CPA secure encryption with weak keys and self-composable MBPF obfuscators for *independent messages*.

Theorem 3.2. *Let $\alpha(n) \in \omega(\log(n))$. There exist **self-composable MBPF obfuscators** with $\alpha(n)$ -entropic security for independent messages if and only if there exist **CPA secure encryption schemes** for $\alpha(n)$ -weak keys and the *wrong-key detection property*. The above also holds if we replace “ $\alpha(n)$ ” with “fully”.*

We prove the two sides of the “if and only if” separately. First we show that composable obfuscation implies encryption (Lemma 3.3) and then we show that encryption implies obfuscation (Lemma 3.4).

In the next lemma, going from obfuscation to encryption, it would be natural to define $\text{Enc}_k(m) = \mathcal{O}(I_{(k,m)})$. However, we instead define $\text{Enc}_k(m) = (\mathcal{O}(I_{(k,r)}), m \oplus r)$ for a uniform r . The reason for this is that the messages m chosen by the adversary in the CPA game can depend adaptively on prior ciphertexts. However, for composable obfuscation, the distributions Y_i of the messages m_i are independent of prior obfuscated circuits. We get around this by making sure that the obfuscation is applied to a random value.

Lemma 3.3. *Let $\alpha(n) \in \omega(\log(n))$ be an arbitrary function. Let \mathcal{O} be a **self-composable MBPF obfuscator** with $\alpha(n)$ -entropic security for independent messages. We define (Enc, Dec) by*

$$\text{Enc}_k(m) \triangleq (\mathcal{O}(I_{(k,r)}), m \oplus r) \quad , \quad \text{Dec}_k(C, y) \triangleq C(k) \oplus y$$

where r is uniformly random, and C is interpreted as a circuit. The resulting encryption scheme is **CPA secure** with $\alpha(n)$ -weak keys.

Proof. The correctness of decryption, and the wrong-key detection property, follow from the correctness of obfuscation. For the CPA security of the encryption scheme with $\alpha(n)$ -weak keys, we define the distribution $D(n, \ell)$ which chooses a uniformly random $k \leftarrow U_n$ in the beginning, and then, on each invocation, outputs $(r, \mathcal{O}(I_{(k,r)}))$ for uniformly random and independent $r, r' \leftarrow_R U_\ell$. We need to show that for all PPT adversaries \mathcal{A} and all distribution-ensembles $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$, we have:

$$\left| \Pr[\text{CPA}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{CPA}_1^{X,D}(\mathcal{A}, n) = 1] \right| \leq \text{negl}(n) \quad (8)$$

for the CPA attack game defined in Definition 2.5.

Fix a PPT adversary \mathcal{A} and let t be an upper-bound on the number of queries that \mathcal{A} sends to its encryption oracle (including the challenge query). Then there are some values of the random coins (r_1, \dots, r_t) used by the encryption algorithm during the computation of ciphertexts $(r_i \oplus m, I_{(k,m)})$ that maximizes the difference in equation (8). For this value, set the distributions $Y_1^{(0)}, \dots, Y_{t+1}^{(0)}$ to the point-values r_1, \dots, r_t . Set the distributions $Y_1^{(1)}, \dots, Y_t^{(1)}$ to be uniform on $\{0, 1\}^\ell$. We define an adversary $\mathcal{B}_{(r_1, \dots, r_t)}(C_1, \dots, C_t)$ that attacks the obfuscation scheme. Namely, \mathcal{B} simulates the CPA game with \mathcal{A} so that, whenever \mathcal{A} queries its oracle on messages m_i or asks for a challenge ciphertext (for $i = 1, \dots, t$), the adversary \mathcal{B} responds with $(C_i, r_i \oplus m_i)$. Notice that, when C_1, \dots, C_t are obfuscations of points $r_1, \dots, r_t \leftarrow Y_1^{(0)}, \dots, Y_{t+1}^{(0)}$ under a random $k \leftarrow X_n$ then the above simulation is equivalent to $\text{CPA}_0^{X,D}$. On the other hand when C_1, \dots, C_t are obfuscation of uniformly random $r_1, \dots, r_t \leftarrow Y_1^{(1)}, \dots, Y_{t+1}^{(1)}$ under a uniformly random key $k \leftarrow U_\ell$, then the above is equivalent to $\text{CPA}_1^{X,D}$.

Therefore,

$$\begin{aligned} & \left| \Pr[\text{CPA}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{CPA}_1^{X,D}(\mathcal{A}, n) = 1] \right| \\ & \leq \left| \Pr \left[\mathcal{B}(C_1, \dots, C_t) = 1 \mid \begin{array}{l} (k, r_1, \dots, r_t) \leftarrow X_n, Y_1^{(0)}, \dots, Y_{t+1}^{(0)} \\ C_i \leftarrow \mathcal{O}(I_{(k,r_i)}) \end{array} \right] \right. \\ & \quad \left. - \Pr \left[\mathcal{B}(C_1, \dots, C_t) = 1 \mid \begin{array}{l} (k, r_1, \dots, r_t) \leftarrow U_n, Y_1^{(1)}, \dots, Y_{t+1}^{(1)} \\ C_i \leftarrow \mathcal{O}(I_{(k,r_i)}) \end{array} \right] \right| \\ & \leq \text{negl}(n) \end{aligned}$$

where, the last inequality follows since, by the definition of self-composable obfuscation, there is a simulator that simulates both sides of the difference equivalently. \square

The other direction is shown via the same construction as in the case of semantic security.

Lemma 3.4. *Let (Enc, Dec) be an encryption scheme with CPA security for $\alpha(n)$ -weak keys and having the wrong-key detection property. We define the obfuscator \mathcal{O} which, on input $I_{(k,m)}$, computes a ciphertext $c = \text{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ defined by $C_c(x) = \text{Dec}_x(c)$. Then, \mathcal{O} is a **self-composable MBPF obfuscator** with $\alpha(n)$ -entropic security for independent messages.*

Proof. The correctness and polynomial slowdown properties follow from the same argument as that in the proof of Lemma 3.2.

We show that \mathcal{O} is self-composable with $\alpha(n)$ -entropic security for independent messages. For any PPT adversary \mathcal{A} and for any $t = \text{poly}(n), \ell = \text{poly}(n)$, we define the simulator \mathcal{S} which, on input 1^n chooses t random ciphertexts c_1, \dots, c_t from the distribution $D(n, \ell)$ as defined by CPA encryption, and runs \mathcal{A} on a circuits $(C_{c_1}, \dots, C_{c_t})$ constructed using the ciphertexts c_1, \dots, c_t . Then, for any distribution ensemble $\{X_n\}_{n \in \mathbb{N}}$ where X_n is distributed over $\{0, 1\}^n$ with $H_\infty(X_n) \geq \alpha(n)$, and t messages $m_1, \dots, m_t \in \{0, 1\}^{\ell(n)}$ we have

$$\begin{aligned} & \left| \Pr[\mathcal{A}(\mathcal{O}(I_{k,m_1}), \dots, \mathcal{O}(I_{k,m_t})) = 1 \mid k \leftarrow X_n] - \Pr[\mathcal{S}^{I_{k,m_1}, \dots, I_{k,m_t}}(1^n, 1^\ell) = 1] \right| \\ & \leq \left| \Pr \left[\mathcal{A}(\{C_{c_i}\}_{i=1}^t) = 1 \mid \begin{array}{l} k \leftarrow X_n \\ \{c_i \leftarrow \text{Enc}_k(m_i)\}_{i=1}^t \end{array} \right] - \Pr[\mathcal{A}(\{C_{c_i}\}_{i=1}^t) = 1 \mid \{c_i \leftarrow D(n, \ell)\}_{i=1}^t] \right| \\ & \leq \text{negl}(n) \end{aligned}$$

where the last inequality simply follows from CPA security. \square

4 Encryption/Obfuscation with Auxiliary Input

In this section we define semantic/CPA secure encryption *with auxiliary input family* \mathcal{F} , where the adversary gets to learn $f(k)$ for any $f \in \mathcal{F}$.⁵ Similarly, we define (self-composable) MBPF obfuscation with auxiliary input family \mathcal{F} , where the adversary and simulator both get $f(k)$ for some $f \in \mathcal{F}$ and the obfuscated point k (we only consider this notion for obfuscation with independent messages). Both notions can be defined for $\alpha(n)$ -weak keys as well as fully weak keys. Then, we show that all of the results of Section 3 extend naturally to the auxiliary input setting.

4.1 Definitions

Definition 4.1 (Symmetric Encryption with Weak Keys and Auxiliary Inputs). *We say that an encryption scheme has CPA security for $\alpha(n)$ -weak keys and auxiliary inputs in \mathcal{F} if there exists an efficient algorithm $D(n, \ell)$ running in time $\text{poly}(n, \ell)$, such that, for all PPT adversaries \mathcal{A} and all distribution-ensembles $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$, we have:*

$$|\Pr[\text{CPA}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{CPA}_1^{X,D}(\mathcal{A}, n) = 1]| \leq \text{negl}(n)$$

where the games $\text{CPA}_b^{X,D}(\mathcal{A}, n)$ for $b = 0, 1$ are defined via the following experiment:

1. $k \leftarrow X_n$
2. \mathcal{A} submits a function $f \in \mathcal{F}$ and gets back $f(k)$.
3. Repeat: \mathcal{A} submits a query m . Set $c_0 \leftarrow \text{Enc}_k(m)$, $c_1 \leftarrow D(n, |m|)$ and give c_b to \mathcal{A} .
4. The output of the game is the output of \mathcal{A} .

We define **semantic security** with $\alpha(n)$ -weak keys and auxiliary inputs in \mathcal{F} via the games $\text{SEM}_0^{X,D}, \text{SEM}_1^{X,D}$, which are equivalent to the CPA games except that step (3) is performed only once.

Definition 4.2 (Self composability). *A multi-bit point function obfuscator \mathcal{O} with $\alpha(n)$ -entropic security with auxiliary inputs from \mathcal{F} is said to be **self-composable** if for any PPT adversary \mathcal{A} with 1 bit output, any polynomials $t(\cdot), \ell(\cdot)$, there exists a PPT simulator \mathcal{S} such that for every $f \in \mathcal{F}$ and all distributions $\{(X_n, Y_n)\}_{n \in \mathbb{N}}$, where $Y_n = Y_n^{(1)}, \dots, Y_n^{(t)}$, and X_n taking values in $\{0, 1\}^n$, $Y_n^{(i)}$ taking values in $\{0, 1\}^{\ell(n)}$ and $H_\infty(X_n) \geq \alpha(n)$, we have:*

$$|\Pr[\mathcal{A}(f(k), \mathcal{O}(I_{(k, m_1)}), \dots, \mathcal{O}(I_{(k, m_t)})) = 1] - \Pr[\mathcal{S}^{I_{(k, m_1)}(\cdot), \dots, I_{(k, m_t)}(\cdot)}(f(k), 1^n) = 1]| \leq \text{negl}(n),$$

where the probabilities are over $(k_1, \dots, k_t, m_1, \dots, m_t) \leftarrow (X_n, Y_n)$ and over the randomness of $\mathcal{A}, \mathcal{S}, \mathcal{O}$.

The notion of self composability extends naturally to obfuscators with fully-entropic security, where we require that the above definition holds for all $\alpha(n) \in \omega(\log(n))$. It also extends to obfuscators for independent messages, where we restrict the definition to the case where X_n and Y_n are independent. (It is stressed that there is no independence assumption among the coordinates within X_n or Y_n .)

4.2 Sem. Sec. Encryption and Obfuscation with Independent Messages

In this section, we show equivalence between semantically secure encryption with weak keys and auxiliary inputs and MBPF obfuscators with auxiliary inputs for *independent messages*.

Theorem 4.1. *Let $\alpha(n) \in \omega(\log(n))$ and let \mathcal{F} be a family of efficiently computable functions. There exist MBPF obfuscators that are $\alpha(n)$ -entropic secure with auxiliary inputs from \mathcal{F} , for independent messages, if and only if there exist semantically secure encryption schemes for $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} , that also have the wrong key detection property. Furthermore, the above also holds if we replace “ $\alpha(n)$ ” with “fully”.*

⁵This is only interesting for families \mathcal{F} where each $f \in \mathcal{F}$ is *hard* to invert, as otherwise $f(k)$ completely reveals k and no security is possible. Often, it makes sense to restrict \mathcal{F} much further, such as requiring that $f(k)$ is exponentially-hard to invert.

The proof of the above theorem is very similar to that of Theorem 3.1. As before, we prove the “if” and “only if” directions separately, in Lemmas 4.1 and 4.2, respectively.

Lemma 4.1. *Let $\alpha(n) \in \omega(\log(n))$ and let \mathcal{F} be a family of efficiently computable functions. Let \mathcal{O} be a MBPF obfuscator that is $\alpha(n)$ -entropic secure with auxiliary inputs from \mathcal{F} , for independent messages. Let $\text{Enc}_k(m) \triangleq \mathcal{O}(I_{(k,m)})$, $\text{Dec}_k(C) \triangleq C(k)$ where the ciphertext C is interpreted as a circuit. Then the encryption scheme (Enc, Dec) is semantically secure with $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} , and has the wrong-key detection property.*

Proof. The correctness of decryption follows from the correctness of obfuscation. For the security of the encryption scheme with $\alpha(n)$ -weak keys and auxiliary inputs in \mathcal{F} : Fix any adversary \mathcal{A} and any distribution $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$. The distribution $\{Y_n\}$ is defined by running $\mathcal{A}(1^n)$, and outputting the message m that \mathcal{A} gives to its challenger, after obtaining some auxiliary input $f(k)$.

Define the distribution $D(n, \ell) = \mathcal{O}(I_{(k',m')})$ where $(k', m') \leftarrow (U_n, U_\ell)$. Then, by the $\alpha(n)$ -entropic security of obfuscation w.r.t. auxiliary inputs from \mathcal{F} , there must be a simulator \mathcal{S} such that

$$\begin{aligned} & \left| \Pr[\text{SEM}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{SEM}_1^{X,D}(\mathcal{A}, n) = 1] \right| \\ &= \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(f(k), \mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k',m') \leftarrow (U_n, U_\ell)} [\mathcal{A}(f(k), \mathcal{O}(I_{(k',m')})) = 1] \right| \\ &\leq \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(f(k), \mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}}(f(k)) = 1] \right| \end{aligned} \quad (9)$$

$$+ \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}}(f(k)) = 1] - \Pr_{(k',m') \leftarrow (U_n, U_\ell)} [\mathcal{S}^{I_{(k',m')}}(f(k)) = 1] \right| \quad (10)$$

$$\begin{aligned} &+ \left| \Pr_{(k',m') \leftarrow (U_n, U_\ell)} [\mathcal{S}^{I_{(k',m')}}(f(k)) = 1] - \Pr_{(k',m') \leftarrow (U_n, U_\ell)} [\mathcal{A}(f(k), \mathcal{O}(I_{(k',m')})) = 1] \right| \\ &\leq \text{negl}(n) \end{aligned} \quad (11)$$

where (9) follows by the definition of entropic security of obfuscation with auxiliary inputs; (10) follows since the only way that a PPT simulator can get anything from its oracle is by querying it on the input k , which happens with negligible probability when k comes from a source of super-logarithmic entropy $\alpha(n)$; and (11) follows by the definition of entropic security of obfuscation. \square

Conversely, we show that an encryption scheme with the wrong-key detection property implies obfuscation.

Lemma 4.2. *Let \mathcal{F} be a family of efficiently computable functions. Let (Enc, Dec) be an encryption scheme with semantic security for $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} , that has the wrong-key detection property. We define the obfuscator \mathcal{O} which, on input $I_{(k,m)}$, computes a ciphertext $c = \text{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ defined by $C_c(x) \triangleq \text{Dec}_x(c)$. Then the obfuscator \mathcal{O} with auxiliary inputs from \mathcal{F} has $\alpha(n)$ -entropic security for independent messages.*

Proof. The proof of correctness and polynomial slowdown of the obfuscator follows exactly the proof of Lemma 3.2. We next show that the obfuscator is $\alpha(n)$ -entropic secure with auxiliary inputs from \mathcal{F} , for independent messages. Let $D(n, \ell)$ be the distribution defined by the semantic-security of the encryption scheme. For any polynomial $\ell(n)$ any PPT adversary \mathcal{A} which gets an auxiliary input $f(k)$ (for some $f \in \mathcal{F}$) and attacks the obfuscation scheme, we define the simulator \mathcal{S} which chooses a random ciphertext c from the distribution $D(n, \ell(n))$ and runs \mathcal{A} on a circuit C_c constructed using the ciphertext c . Then,

$$\begin{aligned} & \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(f(k), \mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}(\cdot)}(f(k), 1^n, 1^\ell) = 1] \right| \quad (12) \\ &= \left| \Pr \left[\mathcal{A}(f(k), C_c) = 1 \mid \begin{array}{c} (k, m) \leftarrow (X_n, Y_n) \\ c \leftarrow \text{Enc}_k(m) \end{array} \right] - \Pr [\mathcal{A}(f(k), C_c) = 1 \mid c \leftarrow D(n, \ell)] \right| \leq \text{negl}(n) \end{aligned} \quad (13)$$

Where (13) follows by semantic-security. \square

4.3 CPA Encryption and Composable Obfuscation for Indep. Messages

In this section, we show equivalence between CPA secure encryption with weak keys and auxiliary inputs, and self-composable MBPF obfuscators with auxiliary inputs for *independent messages*.

Theorem 4.2. *Let $\alpha(n) \in \omega(\log(n))$ and let \mathcal{F} be a family of efficiently computable functions. . There exist **self-composable MBPF obfuscators** that are $\alpha(n)$ -entropic secure with auxiliary inputs from \mathcal{F} , for independent messages, if and only if there exist **CPA secure encryption schemes** for $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} , that have the wrong-key detection property. The above also holds if we replace “ $\alpha(n)$ ” with “fully.”*

The proof of the above theorem is very similar to that of Theorem 3.2. We prove the two sides of the “if and only if” separately. First we show that composable obfuscation with auxiliary inputs implies encryption with auxiliary inputs, (Lemma 4.3) and then we show that encryption with auxiliary inputs implies obfuscation with auxiliary inputs (Lemma 4.4).

Lemma 4.3. *Let $\alpha(n) \in \omega(\log(n))$ be an arbitrary function and let \mathcal{F} be any family of efficiently computable functions. Let \mathcal{O} be a **self-composable MBPF obfuscator** that is $\alpha(n)$ -entropic secure with auxiliary inputs from \mathcal{F} , for independent messages. We define the encryption function by*

$$\text{Enc}_k(m) \triangleq (\mathcal{O}(I_{(k,r)}), m \oplus r),$$

where r is a random message. We define the decryption function by

$$\text{Dec}_k(C, y) \triangleq C(k) \oplus y,$$

where C is interpreted as a circuit. Then, the resulting encryption scheme (Enc, Dec) is **CPA secure** with $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} .

Proof. The correctness of decryption, and the wrong-key detection property, follow from the correctness of obfuscation. For the CPA security of the encryption scheme with $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} , we define the distribution $D(n, \ell)$ which chooses a uniformly random $k' \leftarrow U_n$ in the beginning, and then, on each invocation, outputs $(\mathcal{O}(I_{(k',r')}), r)$ for uniformly random and independent $r, r' \leftarrow_R U_\ell$. We need to show that for all PPT adversaries \mathcal{A} and all distribution-ensembles $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$, we have:

$$|\Pr[\text{CPA}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{CPA}_1^{X,D}(\mathcal{A}, n) = 1]| \leq \text{negl}(n) \quad (14)$$

for the CPA attack game defined in Definition 4.1.

Fix a PPT adversary \mathcal{A} and let t be an upper-bound on the number of queries that \mathcal{A} sends to its encryption oracle (including the challenge query), and let $f \in \mathcal{F}$ be the auxiliary input that \mathcal{A} takes. Then there are some values of the random coins (r_1, \dots, r_t) used by the encryption algorithm during the computation of ciphertexts $(I_{(k,m)}, r_i \oplus m)$ that maximizes the difference in equation (14). For this value, set the distributions $Y_1^{(0)}, \dots, Y_t^{(0)}$, to the point-values r_1, \dots, r_t . Set the distributions $Y_1^{(1)}, \dots, Y_t^{(1)}$ to be uniform on $\{0, 1\}^\ell$. We define an adversary $\mathcal{B}_{(r_1, \dots, r_t)}(f(k), C_1, \dots, C_t)$ that attacks the obfuscation scheme. Namely, \mathcal{B} simulates the CPA game with \mathcal{A} so that, whenever \mathcal{A} queries its oracle on messages m_i or asks for a challenge ciphertext (for $i = 1, \dots, t$), the adversary \mathcal{B} responds with $(C_i, r_i \oplus m_i)$. Notice that, when C_1, \dots, C_t are obfuscations of points $r_1, \dots, r_t \leftarrow Y_1^{(0)}, \dots, Y_t^{(0)}$ under the key k (which was chosen according to X_n) then the above simulation is equivalent to $\text{CPA}_0^{X,D}$. On the other hand when C_1, \dots, C_t are obfuscation of uniformly random $r'_1, \dots, r'_t \leftarrow Y_1^{(1)}, \dots, Y_t^{(1)}$ under a uniformly random key $k' \leftarrow U_\ell$, then the above is equivalent to $\text{CPA}_1^{X,D}$.

Therefore,

$$\begin{aligned}
& \left| \Pr[\text{CPA}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{CPA}_1^{X,D}(\mathcal{A}, n) = 1] \right| \\
\leq & \left| \Pr \left[\mathcal{B}(f(k), C_1, \dots, C_t) = 1 \mid \begin{array}{l} (k, r_1, \dots, r_t) \leftarrow X_n, Y_1^{(0)}, \dots, Y_t^{(0)} \\ C_i \leftarrow \mathcal{O}(I_{(k, r_i)}) \end{array} \right] \right. \\
& \left. - \Pr \left[\mathcal{B}(f(k), C_1, \dots, C_t) = 1 \mid \begin{array}{l} k \leftarrow X_n \\ (k', r'_1, \dots, r'_t) \leftarrow U_n, Y_1^{(1)}, \dots, Y_t^{(1)} \\ C_i \leftarrow \mathcal{O}(I_{(k', r'_i)}) \end{array} \right] \right| \\
\leq & \text{negl}(n)
\end{aligned}$$

where the last inequality follows from the definition of self-composable obfuscation with auxiliary inputs, since there exists a simulator for both cases, and these simulators output the same thing, since they cannot tell their oracles apart. \square

The other direction is shown via the same construction as in the case of semantic security.

Lemma 4.4. *Let \mathcal{F} be a family of efficiently computable functions. Let (Enc, Dec) be an encryption scheme with CPA security for $\alpha(n)$ -weak keys and auxiliary inputs from \mathcal{F} . In addition assume that (Enc, Dec) has the wrong-key detection property. We define the obfuscator \mathcal{O} which, on input $I_{(k,m)}$, computes a ciphertext $c = \text{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ defined by $C_c(x) = \text{Dec}_x(c)$. Then, \mathcal{O} is a **self-composable MBPF obfuscator** for $\alpha(n)$ -entropic secure and auxiliary inputs in \mathcal{F} , for independent messages.*

Proof. The correctness and polynomial slowdown properties follow from the same argument as that in the proof of Lemma 3.2.

We show that \mathcal{O} is self-composable for $\alpha(n)$ -entropic security and auxiliary inputs from \mathcal{F} , for independent messages. For any PPT adversary \mathcal{A} and for any $t = \text{poly}(n), \ell = \text{poly}(n)$, we define the simulator \mathcal{S} which, on input 1^n chooses t random ciphertexts c_1, \dots, c_t from the distribution $D(n, \ell)$ as defined by CPA encryption, and runs \mathcal{A} on a circuits $(C_{c_1}, \dots, C_{c_t})$ constructed using the ciphertexts c_1, \dots, c_t . Then, for any distribution ensemble $\{X_n\}_{n \in \mathbb{N}}$ where X_n is distributed over $\{0, 1\}^n$ with $H_\infty(X_n) \geq \alpha(n)$, and t messages $m_1, \dots, m_t \in \{0, 1\}^{\ell(n)}$ we have

$$\begin{aligned}
& \left| \Pr[\mathcal{A}(f(k), \mathcal{O}(I_{k, m_1}), \dots, \mathcal{O}(I_{k, m_t})) = 1 \mid k \leftarrow X_n] - \Pr[\mathcal{S}^{I_{k, m_1}, \dots, I_{k, m_t}}(f(k), 1^n, 1^\ell) = 1] \right| \\
\leq & \left| \Pr \left[\mathcal{A}(f(k), \{C_{c_i}\}_{i=1}^t) = 1 \mid \begin{array}{l} k \leftarrow X_n \\ \{c_i \leftarrow \text{Enc}_k(m_i)\}_{i=1}^t \end{array} \right] - \Pr[\mathcal{A}(f(k), \{C_{c_i}\}_{i=1}^t) = 1 \mid \{c_i \leftarrow D(n, \ell)\}_{i=1}^t] \right| \\
\leq & \text{negl}(n)
\end{aligned}$$

where the last inequality simply follows from CPA security. \square

5 KDM Encryption and MBPF Obfuscation

5.1 Semantically Secure KDM Encryption and Obfuscation

In this section, we show equivalence between encryption with *key dependent messages* (KDM) and obfuscation with dependent messages. First, we define the notion of semantically-secure KDM encryption with $\alpha(n)$ -weak keys.

Definition 5.1 (Semantic KDM Encryption with Weak Keys). *A symmetric encryption scheme (Enc, Dec) is semantically secure for **key dependent messages (KDM)** and $\alpha(n)$ -weak keys if there exists a distribution $D(n, \ell)$, which is efficiently sampleable in time $\text{poly}(n, \ell)$, such that for all functions f , all PPT adversaries \mathcal{A} , and all distribution-ensembles $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$, we have:*

$$\left| \Pr[\text{KDM}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\text{KDM}_1^{X,D}(\mathcal{A}, n) = 1] \right| \leq \text{negl}(n), \quad (15)$$

where $\text{KDM}_b^{X,D}(\mathcal{A}, n)$ is defined via the following experiment:

$$\begin{aligned} k &\leftarrow X_n \\ c_0 &\leftarrow \text{Enc}_k(f(k)), c_1 \leftarrow D(n, \ell) \text{ where } \ell \text{ is the output size of } f. \\ \text{Output: } &\mathcal{A}(c_b) \end{aligned}$$

Note that, unlike standard definitions of KDM security, our definition is stronger in that we do not necessarily insist that f is an efficient function. We now show that semantically secure encryption with KDM and weak key security is equivalent to MBPF obfuscation.

Theorem 5.1. *Let $\alpha(n) \in \omega(\log(n))$. There exist MBPF obfuscators with $\alpha(n)$ -entropic security for the **standard notion of dependent messages** if and only if there exist semantically-secure **KDM** encryption schemes with $\alpha(n)$ -weak keys and the “wrong-key detection” property. In particular, the above also holds if we replace “ $\alpha(n)$ ” with “fully”.*

The proof of the above theorem is very similar to that of Theorem 3.2. We simply observe that allowing the adversary to get encryptions of values $f(k)$ corresponds to having a distribution Y_n that depends on X_n ; that is $Y_n = f(X_n)$. Conversely, for any joint distribution $\{X_n, Y_n\}$, we can define some (probabilistic, and possibly inefficient) function f so that $Y_n = f(X_n)$. We give the details below. Again, we prove the two sides of the “if and only if” separately. First we show that obfuscation implies KDM encryption in Lemma 5.1 and then we show that encryption implies obfuscation in Lemma 5.2.

Lemma 5.1. *Let $\alpha(n) \in \omega(\log(n))$ be an arbitrary function. Let \mathcal{O} be a MBPF obfuscator with $\alpha(n)$ -entropic security. We define the functions $\text{Enc}_k(m) = \mathcal{O}(I_{(k,m)})$, $\text{Dec}_k(C) = C(k)$ where the ciphertext C is interpreted as a circuit. The resulting encryption scheme (Enc, Dec) is semantically **KDM** secure with $\alpha(n)$ -weak keys and wrong-key detection.*

Proof. The correctness of decryption follows from the correctness of obfuscation. For the security of the resulting KDM encryption scheme with $\alpha(n)$ -weak keys, let us fix any adversary \mathcal{A} , function f , and distribution $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$. Define $Y_n = f(X_n)$. By the entropic-security of obfuscation with **dependent** messages, there exists a simulator \mathcal{S} such that:

$$\begin{aligned} &|\Pr[\text{KDM}_0(\mathcal{A}, n) = 1] - \Pr[\text{KDM}_1(\mathcal{A}, n) = 1]| \\ &= \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{k \leftarrow X_n} [\mathcal{A}(\mathcal{O}(I_{(k,0^\ell)})) = 1] \right| \\ &\leq \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}(\cdot)}(1^{|k|}, 1^\ell) = 1] \right| \quad (16) \\ &\quad + \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}(\cdot)}(1^{|k|}, 1^\ell) = 1] - \Pr_{k \leftarrow X_n} [\mathcal{S}^{I_{(k,0^\ell)}(\cdot)}(1^{|k|}, 1^\ell) = 1] \right| \quad (17) \\ &\quad + \left| \Pr_{k \leftarrow X_n} [\mathcal{S}^{I_{(k,0^\ell)}(\cdot)}(1^{|k|}, 1^\ell) = 1] - \Pr_{k \leftarrow X_n} [\mathcal{A}(\mathcal{O}(I_{(k,0^\ell)})) = 1] \right| \quad (18) \\ &\leq \text{negl}(n) \end{aligned}$$

where (16),(18) follow by the definition of entropic security of obfuscation with dependent messages, and (17) follows since the only way that a PPT simulator can get anything from its oracle is by querying it on the input k , which happens with negligible probability when k comes from a source of super-logarithmic entropy $\alpha(n)$. \square

Lemma 5.2. *Let (Enc, Dec) be an encryption scheme with **semantic KDM** security for $\alpha(n)$ -weak keys and with they wrong-key detection property. We define the obfuscator \mathcal{O} which, on input $I_{(k,m)}$, computes a ciphertext $c = \text{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ (with hard-coded ciphertext c) defined by $C_c(x) \triangleq \text{Dec}_x(c)$. Then the obfuscator \mathcal{O} has $\alpha(n)$ -entropic security for dependent messages.*

Proof. For any PPT adversary \mathcal{A} , we define the simulator \mathcal{S} which, on input 1^n , finds the message size ℓ by querying the $I_{(k,m)}$ oracle on \perp , chooses a random ciphertext c from the distribution $D(n, \ell)$ as defined by the *pseudorandom ciphertexts* property, and runs \mathcal{A} on a circuit C_c constructed using the ciphertext c . Then, for any distribution ensemble $\{X_n, Y_n\}_{n \in \mathbb{N}}$ where X_n is distributed over $\{0, 1\}^n$ with $H_\infty(X_n) \geq \alpha(n)$, define the function $f(k)$ which, on input k , outputs a random sample from the distribution $(Y_n | X_n = k)$ of Y_n conditioned on $X_n = k$. Note that this function may not be efficient (but our definition of KDM encryption allows this). Then

$$\begin{aligned} & \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)} [\mathcal{S}^{I_{(k,m)}(\cdot)}(1^n, 1^\ell) = 1] \right| \\ &= \left| \Pr \left[\mathcal{A}(C_c) = 1 \mid \begin{array}{l} k \leftarrow (X_n) \\ c \leftarrow \text{Enc}_k(f(k)) \end{array} \right] - \Pr[\mathcal{A}(C_c) = 1 \mid c \leftarrow D(n, \ell)] \right| \leq \text{negl}(n) \end{aligned}$$

□

5.2 Multi-KDM Encryption and Self-Composable Obfuscation

In this section, we explore a notion of CPA security with KDM and weak-keys. We essentially show results analogous to those in Section 3.2 connecting CPA encryption (without KDM) to obfuscation with independent messages, but only if we restrict ourselves to a non-adaptive attacker who chooses the function f of the secret key prior to seeing any ciphertexts.

Definition 5.2 (Multi-KDM Encryption with Weak Keys). *A symmetric encryption scheme (Enc, Dec) is **Multi-KDM secure** for $\alpha(n)$ -weak keys if there exists a distribution $D(n, \ell)$ such that for any $t \leq \text{poly}(n)$, any functions f_1, \dots, f_t , any PPT adversary \mathcal{A} , and any distribution-ensemble $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$, we have:*

$$\left| \Pr[\mathcal{A}(\text{Enc}_k(f_1(k)), \dots, \text{Enc}_k(f_t(k))) = 1] - \Pr[\mathcal{A}(c_0, \dots, c_t) = 1 \mid c_i \leftarrow D(n, \ell_i)] \right| \leq \text{negl}(n),$$

where ℓ_i is the output size of f_i .

Theorem 5.2. *Let $\alpha(n) \in \omega(\log(n))$ be an arbitrary function. Let \mathcal{O} be a self-composable MBPF obfuscators with $\alpha(n)$ -entropic security (for dependent messages). We define the encryption function by*

$$\text{Enc}_k(m) \triangleq \mathcal{O}(I_{k,m}),$$

and the decryption function by

$$\text{Dec}_k(C) \triangleq C(k),$$

where C is interpreted as a circuit. Then the resulting encryption scheme (Enc, Dec) is multi-KDM secure with $\alpha(n)$ -weak keys.

Proof. The correctness of decryption follows from the correctness of obfuscation. For the multi-KDM security of the encryption scheme, fix any $t \leq \text{poly}(n)$, any poly-size circuits f_1, \dots, f_t , any PPT adversary \mathcal{A} , and any distribution-ensemble $\{X_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \geq \alpha(n)$.

$$\begin{aligned} & \left| \Pr_{k \leftarrow X_n} [\mathcal{A}(\text{Enc}_k(f_1(k)), \dots, \text{Enc}_k(f_t(k))) = 1] - \Pr_{k \leftarrow X_n} [\mathcal{A}(\text{Enc}_k(0^{\ell_1}), \dots, \text{Enc}_k(0^{\ell_t})) = 1] \right| = \\ & \left| \Pr_{k \leftarrow X_n} [\mathcal{A}(\mathcal{O}(I_{(k,f_1(k))}), \dots, \mathcal{O}(I_{(k,f_t(k))})) = 1] - \Pr_{k \leftarrow X_n} [\mathcal{A}(\mathcal{O}(I_{(k,0^{\ell_1})}), \dots, \mathcal{O}(I_{(k,0^{\ell_t})})) = 1] \right| \leq \text{negl}(n) \end{aligned}$$

where the latter equation follows from the fact that \mathcal{O} is a self-composable obfuscator w.r.t. dependent messages and hence, if the probability of \mathcal{A} outputting 1 was non-negligibly different between the left and right-hand sides above, then there would be a PPT simulator that could distinguish the above distributions, but that cannot be the case. □

Theorem 5.3. *Let (Enc, Dec) be an encryption scheme with **multi-KDM** security for $\alpha(n)$ -weak keys and with the wrong-key detection property. We define the obfuscator \mathcal{O} which, on input $I_{(k,m)}$, computes a ciphertext $c = \text{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ (with hard-coded ciphertext c) defined by $C_c(x) = \text{Dec}_x(c)$. Then, \mathcal{O} is a **self-composable MBPF obfuscator** with $\alpha(n)$ -entropic security for dependent messages.*

Proof. First, the correctness property of the obfuscator and the polynomial slowdown property follow the same way as in the proof of Lemma 3.1.

We must show that \mathcal{O} is self-composable with $\alpha(n)$ -entropic security for dependent messages. For any PPT adversary \mathcal{A} and for any $t \leq \text{poly}(n)$, we define the simulator \mathcal{S} which, on input 1^n , finds the message size ℓ by querying the $I_{(k,m)}$ oracle on \perp , chooses t random ciphertexts c_1, \dots, c_t from the distribution $D(n, \ell)$ as defined by the multi-KDM security definition, and runs \mathcal{A} on a circuits $(C_{c_1}, \dots, C_{c_t})$ constructed using the ciphertexts c_1, \dots, c_t . Then, for any distribution ensemble $\{X_n\}_{n \in \mathbb{N}}$ where X_n is distributed over $\{0, 1\}^n$ with $H_\infty(X_n) \geq \alpha(n)$, and t functions f_1, \dots, f_t , we have

$$\begin{aligned} & \left| \Pr [\mathcal{A}(\mathcal{O}(I_{k,f_1(k)}), \dots, \mathcal{O}(I_{k,f_t(k)})) = 1 \mid k \leftarrow X_n] - \Pr [\mathcal{S}^{I_{k,f_1(k)}, \dots, I_{k,f_t(k)}}(1^n, 1^\ell) = 1] \right| \\ & \leq \left| \Pr \left[\mathcal{A}(C_{c_1}, \dots, C_{c_t}) = 1 \mid \begin{array}{c} k \leftarrow X_n \\ c_i \leftarrow \text{Enc}_k(f_i(k)) \end{array} \right] - \Pr [\mathcal{A}(C_{c_1}, \dots, C_{c_t}) = 1 \mid c_i \leftarrow D(n, \ell)] \right| \\ & \leq \text{negl}(n) \end{aligned}$$

where the last inequality follows from the fact that the encryption scheme (Enc, Dec) is multi-KDM secure. \square

6 Implications

We now show how to use the above equivalence results between encryption with weak keys and obfuscation of multi-bit point functions to derive new results in both directions.

6.1 Encryption with Fully Weak Keys

Encryption with $\alpha(n)$ -weak keys vs. fully-weak keys. Prior work on leakage-resilient encryption and encryption with weak-keys has given results of the following form:

1. Fix any constant $\varepsilon > 0$ and let $\alpha(n) = n^\varepsilon$.
2. Construct an encryption scheme, which depends on ε , and achieves security for $\alpha(n)$ -weak keys.

We note that there are several issues with the above two-step approach. Firstly, we may not know the exact level of key-entropy, or correspondingly the value of ε , at design time. Therefore, in practice, it may be difficult to decide on what ε to use when choosing the encryption scheme. A scheme which is designed for some specific ε does not provide any security guarantees for key-distributions whose entropy is strictly less than n^ε , and so we may be tempted to be conservative with the choice of ε at design time. On the other hand, when taking an excessively small value of ε in the above constructions, we are forced to reduce the exact-security of the system (e.g. working in a group of description-length n^ε) or reduce the efficiency of the system proportionally with $n^{1/\varepsilon}$, leading to poorer security or performance even if the system is later only used with uniformly random keys! Secondly, none of the prior results generalize to allow for specific super-logarithmic entropy thresholds such as $\alpha(n) = \log^{1+\varepsilon}(n)$, even if ε is specified a-priori.

In contrast, an encryption scheme with security for fully-weak keys provides the corresponding advantages. More specifically, the order of quantifiers now requires that there is a *single encryption scheme*, parameterized only by the security parameter n (but *not* by ε), which simultaneously achieves security for all $\alpha(n) \in \omega(\log(n))$. The exact-security of the scheme may depend on $\alpha(n)$ (since there is always a way to break the scheme in time $2^{\alpha(n)}$), but this relationship is now more fluid, with the exact-security gracefully degrading for smaller $\alpha(n)$. In particular, the security guarantees are meaningful even for $\alpha(n) = \log^{1+\varepsilon}(n)$, and there is no single threshold above which the scheme is secure and below which it is insecure. This is a significant advantage, as it does not require one to decide at design time on the tradeoff between allowed entropy levels and achieved security/efficiency.

New construction of encryption with fully-weak keys. We now describe the point-function obfuscation scheme of Canetti [10], and notice that it yields a self-composable MBPF obfuscator with *fully-entropic security* for independent messages. It is based on a strengthened version of the DDH assumption, which we describe shortly. Using this simple observation and our connection between obfuscation and encryption (Lemma 3.3), we get the first symmetric-key encryption scheme with CPA security for *fully-weak* keys (albeit under a strong assumption). We begin by defining the *strengthened DDH* assumption for a prime-ordered group \mathbb{G} .

Definition 6.1 (Strengthened DDH Assumption [10]). *Let \mathbb{G} be a group of prime order $p = 2^{\text{poly}(n)}$ and let g be a random generator of \mathbb{G} . The strengthened DDH assumption states that, for any distribution $\{X_n\}$ over \mathbb{Z}_p with entropy $H_\infty(X_n) \geq \omega(\log(n))$, we have $\langle g^a, g^b, g^{ab} \rangle \approx \langle g^a, g^b, g^c \rangle$ where $a \leftarrow_R X_n$, and $b, c \leftarrow_R \mathbb{Z}_p$.*

We now define the function $F : \mathbb{Z}_p \rightarrow \mathbb{G} \times \mathbb{G}$ by $F(k) = \langle r, r^k \rangle$ where $r \leftarrow_R \mathbb{G}$. In [10], this was shown to be a secure *point-function* obfuscator (with fully-entropic security) under the strengthened DDH assumption. In addition, this point-function obfuscator is self-composable since, given a (random) obfuscation $\langle g_1, g_2 \rangle$ of some point x , it is easy to generate freshly random (and independent) new obfuscation of x by taking $\langle g_1^u, g_2^u \rangle$ for a random $u \in \mathbb{Z}_p$. We use the construction of Canetti and Dakdouk [11] to turn a point-function obfuscator into a multi-bit point-function obfuscator. Define the function:

$$\mathcal{O}(I_{(k,m)}) = \begin{cases} \text{Sample } r_0, r_1, \dots, r_\ell \leftarrow_R \mathbb{G} \text{ for } \ell = |m|. \\ \text{Set } g_0 = r_0^k \\ \text{For each } i \in \{1, \dots, \ell\} : \text{ if } m_i = 1 \text{ set } g_i = r_i^k \text{ else } g_i \leftarrow_R \mathbb{G}. \\ \text{Output: } c = (\langle r_0, g_0 \rangle, \dots, \langle r_\ell, g_\ell \rangle). \end{cases}$$

Using the techniques of [11], it is easy to show that \mathcal{O} is a *self-composable* obfuscator with fully-entropic security for *independent messages* under the strengthened DDH assumption. Combining this with Lemma 3.3, we get the following theorem.

Theorem 6.1. *Under the strengthened DDH assumption, there exists a CPA-secure symmetric encryption scheme with security against fully-weak keys. In particular, this means that there is a single scheme, parameterized only by the security parameter n , such that security of the scheme is maintained when the key is chosen from any distribution of entropy $\alpha(n) \in \omega(\log(n))$.*

The strengthened DDH assumption is indeed a strong one. A potentially weaker formulation would be to limit the min-entropy of X_n to be at least some specific super-logarithmic function $\alpha(n)$. This way, we would obtain a parameterized version of Theorem 6.1 that relates the strength of the security guarantee to the strength of the assumption. It is important to note that the construction itself is independent of the parameter α . That is, we obtain a single encryption scheme that provides a range of security guarantees, depending on the strength of the assumption.

6.2 Obfuscation

Entropically Secure Obfuscation for Independent Messages: It is fairly simple to construct $\alpha(n)$ -entropically secure obfuscation for independent messages, when $\alpha(n) = n^\varepsilon$ for some constant $\varepsilon \geq 0$. First we construct a semantically secure encryption scheme with $\alpha(n)$ -weak keys. This can be done by simply extracting a sufficient amount of uniform randomness from the key k , using a strong randomness extractor Ext , and then using the result as a one time pad to encrypt the message. For variable-length messages, we also need to expand the extracted randomness to an appropriate size, using a pseudo-random generator PRG. In particular, we define

$$\text{Enc}_k(m) = \langle r, \text{PRG}(\text{Ext}(k; r)) \oplus m \rangle$$

where r is a uniformly random seed for the extractor. The output length of Ext and the input length of PRG are set to some value v which is sufficiently small that the outputs of the extractor is (statistically) close to uniform, and sufficiently large that the output of the PRG is pseudo-random.⁶

⁶For example, if we choose $v = n^\varepsilon/2$, then an extractor based on universal-hash functions will produce an output which is $2^{-v/2} = \text{negl}(n)$ -close to uniform, and the output of the PRG is $\text{negl}(n^\varepsilon/2) = \text{negl}(n)$ -pseudorandom. However, this does not generalize to smaller values of α such as, $\alpha(n) = \log^2(n)$.

One can use this encryption scheme to construct one which also has the wrong-key detection property using Lemma 2.1. Such a scheme yields an multi-bit point function obfuscator with $\alpha(n)$ -entropic security for independent messages, by Lemma 3.2.

Self-Composable Entropically Secure Obfuscation for Independent Messages: One problem with the above construction of semantically-secure encryption using extractors, is that it does not generalize to CPA security. In fact, achieving CPA secure encryption with weak keys seems to be a much harder problem, which has received much attention in recent works [1, 13, 24]. We now show how to use these results to achieve self-composable entropically secure obfuscation for independent messages. On a high level, we would simply like to just apply our result connecting such encryption and obfuscation (Lemma 3.4) “out of the box”. However, there are several issues that we must deal with first.

- *Efficiently-Sampleable Distributions:* The works of [1, 13, 24] are concerned with “key leakage”, where the adversary gets to learn some (short) function of the secret key, whose output length is bounded by λ bits. Conditioned on such leakage, the key can be thought of as being derived from a (special type) of weak source with entropy $\alpha(n) \approx n - \lambda$. It turns out that the constructions are also secure when the key is chosen from an *arbitrary, but efficiently-sampleable* weak source of entropy $\alpha(n)$ [24]. Therefore, our results for obfuscation will only translate to the case where the distribution obfuscated program is efficiently sampleable.
- *Public Keys/Parameters:* Only the scheme of [13] is explicitly designed for the symmetric key setting. The schemes of [1, 24] are public-key encryption schemes. As noted, such schemes are secure when the key-generation procedure uses randomness that comes from a weak source. Therefore such schemes naturally translate to the symmetric key setting, where the randomness of the key-generation algorithm is the shared secret key. Unfortunately, these schemes also rely on *public parameters* which are chosen uniformly at random, and are available to the key-generation algorithm. Therefore, we will only get an obfuscator in the presence of public parameters. Note that in the context of standard obfuscation, public parameters are never needed since the obfuscator \mathcal{O} could always sample fresh parameters each time it runs. However, when considering *composable obfuscation*, this equivalence does not hold since future uses of the obfuscator might compromise security of prior uses. Therefore, having randomness in the form of public parameters, which are re-used for all invocations of the obfuscator, can be useful in this context.
- *Uniform Ciphertexts:* Recall that our definition of CPA security is slightly different than the standard (we require that the ciphertexts of any message are indistinguishable from some universally specified distribution) and has not been explicitly analyzed by these schemes. However, in all of these schemes explicitly show in their proofs that the ciphertexts are indistinguishable from uniform, which satisfies our definition.
- *Wrong-Key Detection:* The wrong-key detection property is explicitly analyzed in [13]. For the schemes of [1, 24], we get the property that, given the public parameters it is computationally difficult to find k, k' such that $\text{Dec}_{k'}(\text{Enc}_k(m)) \neq \perp$. This translates to a *computational-correctness* property for the obfuscator where, given the public parameters, it is computationally difficult to find k, m, x such that $\mathcal{O}(I_{(k,m)})(x) \neq I_{(k,m)}(x)$.

Using our connection between CPA-secure symmetric key encryption and self-composable obfuscation with independent messages, we get the following new constructions of obfuscators as a corollary of Lemma 3.4, using the schemes of [1, 13, 24].

Theorem 6.2. *For any constant $\varepsilon > 0$, there exists a self-composable MBPF obfuscator with independent messages under any of the following assumptions:*

1. *Decisional Diffie-Hellman (DDH) with n^ε -entropic security, based on [24]. ^(*,†)*
2. *Learning With Errors (LWE) with n^ε -entropic security, based on [1]. ^(*,†)*
3. *Learning Subspaces with Noise (LSN) with εn -entropic security, based on [13]. ^(*)*

where the restrictions are:

* Only works for efficiently sampleable key-distributions.

† Requires public parameters and only achieves computational-correctness.

Difficulty of Achieving Obfuscation with Dependent Messages. The connection between encryption and obfuscation also yields new negative results for the more standard notion of obfuscation that allows for *dependent* messages, and in particular for the standard VBB notion. We rely on a recent result of Haitner and Holenstein [16], which shows that there can be *no* black-box reduction from a semantically secure encryption scheme with security against key-dependent messages to, essentially, *any standard cryptographic assumption*. The notion of “cryptographic assumption” is formalized in [16] as (essentially) any game between an attacker and a challenger in which we assume that all PPT attackers have a negligible success probability. In particular, this includes all standard assumptions such as existence of Trapdoor One-Way Permutations or Claw-Free Permutations, as well as specific algebraic assumptions like the hardness of factoring, DDH, Learning with Errors and many others.⁷ Since, by Theorem 5.1, we have a reduction from a semantically secure encryption schemes with security against key-dependent messages to obfuscation of multi-bit point functions with n -entropic security (i.e. even uniformly random keys), we see that this latter notion of obfuscation cannot be realized from essentially any cryptographic assumption under black-box reductions.

Theorem 6.3. *No construction of an MBPF obfuscator with $\alpha(n)$ -entropic security for dependent messages can be proven secure via a black-box reduction to any “standard cryptographic assumption”, even for $\alpha(n) = n$ (i.e. even uniformly random keys).*

We note that Canetti and Dakdouk [11] showed that *composable obfuscation of point functions (with no output)* (i.e. functions $I_k(x)$ which output 1 when $x = k$ and \perp otherwise) implies multi-bit point function obfuscators *with dependent messages*. Thus we get the following as a corollary.

Corollary 6.1. *No construction of a composable obfuscator for single-value point functions with $\alpha(n)$ -entropic security can be proven secure via a black-box reduction to any “standard cryptographic assumption”, for any $\alpha()$ (even for $\alpha(n) = n$, namely uniformly random keys).*

We note that the impossibility result of [16] only considers semantically secure encryption with *variable length messages* and does not rule out KDM security when the message size is shorter than the key. Correspondingly, the work of [11] constructs MBPF obfuscators with $\alpha(n)$ -entropic security (for some $\alpha(n) \ll n$) and *for dependent messages* in this special case, where the message size is (significantly) smaller than the key size (i.e. functions $I_{(k,m)}$ where $|m| < |k|$). These constructions only relied on *standard cryptographic assumptions* such as collision-resistant hash functions. The above theorem shows that such constructions do not generalize to variable-length messages, where the message size can exceed the key size. Alternatively, in this work we show how to leverage prior results on leakage-resilient cryptography to construct self-composable MBPF obfuscators with $\alpha(n)$ -entropic security (for some $\alpha(n) \ll n$), under standard assumptions, in the special case of (variable-length) *independent messages*. It seems that there is little hope in generalizing this approach to the standard notion of obfuscation, which also allows key-dependent messages.

References

- [1] A. Akavia, S. Goldwasser, and V. Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *TCC*, pages 474–495, 2009.

⁷On the other hand, the impossibility result does not exclude proofs of security in the Random Oracle model, reductions to non-standard assumptions (which cannot be formulated as a game between an adversary and a challenger) such as “Knowledge of Exponent”, or non-black-box reductions.

- [2] J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In *CRYPTO*, pages 36–54, 2009.
- [3] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618, 2009.
- [4] M. Backes, M. Dürmuth, and D. Unruh. Oaep is secure under key-dependent messages. In *ASIACRYPT*, pages 506–523, 2008.
- [5] M. Backes, B. Pfitzmann, and A. Scedrov. Key-dependent message security under active attacks - brsim/uc-soundness of symbolic encryption with key cycles. In *CSF*, pages 112–124, 2007.
- [6] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2001.
- [7] J. Black, P. Rogaway, and T. Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In *Selected Areas in Cryptography*, pages 62–75, 2002.
- [8] D. Boneh, S. Halevi, M. Hamburg, and R. Ostrovsky. Circular-secure encryption from decision diffie-hellman. In *CRYPTO*, pages 108–125, 2008.
- [9] J. Camenisch, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *EUROCRYPT*, pages 351–368, 2009.
- [10] R. Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In *CRYPTO*, volume 1294 of *Lecture Notes in Computer Science*, pages 455–469. Springer, 1997.
- [11] R. Canetti and R. R. Dakdouk. Obfuscating point functions with multibit output. In *EUROCRYPT*, volume 4965 of *Lecture Notes in Computer Science*, pages 489–508. Springer, 2008.
- [12] R. Canetti, D. Micciancio, and O. Reingold. Perfectly one-way probabilistic hash functions. In *Proceedings of the 30th ACM Symposium on Theory of Computing*, pages 131–140, 1998.
- [13] Y. Dodis, Y. T. Kalai, and S. Lovett. On cryptography with auxiliary input. In *STOC*, pages 621–630, 2009.
- [14] S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302, 2008.
- [15] S. Goldwasser and Y. T. Kalai. On the impossibility of obfuscation with auxiliary input. In *FOCS*, pages 553–562, 2005.
- [16] I. Haitner and T. Holenstein. On the (im)possibility of key dependent encryption. In *TCC*, pages 202–219, 2009.
- [17] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52(5):91–98, 2009.
- [18] S. Halevi and H. Krawczyk. Security under key-dependent inputs. In *ACM Conference on Computer and Communications Security*, pages 466–475, 2007.
- [19] D. Hofheinz and D. Unruh. Towards key-dependent message security in the standard model. In *EUROCRYPT*, pages 108–126, 2008.
- [20] J. Katz and V. Vaikuntanathan. Signature schemes with bounded leakage resilience, 2009. To Appear in Asiacrypt '09. <http://www.mit.edu/~vinodv/papers/asiacrypt09/KV-Sigs.pdf>.
- [21] B. Lynn, M. Prabhakaran, and A. Sahai. Positive results and techniques for obfuscation. In *EUROCRYPT*, volume 3027 of *Lecture Notes in Computer Science*, pages 20–39. Springer, 2004.
- [22] S. Micali and L. Reyzin. Physically observable cryptography (extended abstract). In *TCC*, pages 278–296, 2004.
- [23] D. Micciancio and B. Warinschi. Completeness theorems for the abadi-rogaway language of encrypted expressions. *Journal of Computer Security*, 12(1):99–130, 2004.
- [24] M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 18–35. Springer, 2009.
- [25] H. Wee. On obfuscating point functions. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 523–532, 2005.

A Fully-entropic security and the virtual black-box property

In this section, we prove Theorem 2.1, which states that an obfuscator \mathcal{O} with fully-entropic security satisfies the virtual black-box property in Definition 2.2. The virtual black-box definition appears to be stronger than fully-entropic security because it does not impose any constraint on the min-entropy of the distribution X_n . However, we show in this section that this is not the case.

We do not work with the virtual black-box definition directly, but rather use two intermediate definitions.

Definition A.1 (Distributional indistinguishability [10]). *For any PPT adversary \mathcal{A} with 1 bit output and for any distribution $\{X_n, Y_n\}_{n \in \mathbb{N}}$ with $H_\infty(X_n) \in \omega(\log(n))$, the distributions $\langle k, m, \mathcal{A}(\mathcal{O}(k, m)) \rangle$ and $\langle k, m, \mathcal{A}(\mathcal{O}(k', m')) \rangle$ are computationally indistinguishable, where $(k, m), (k', m') \leftarrow \{X_n, Y_n\}$ independently.*

Definition A.2 (Oracle indistinguishability [10]). *For any PPT adversary \mathcal{A} with 1 bit output and any polynomial p , there exists a polynomial-sized family of sets of keys $\{L_n\}_{n \in \mathbb{N}}$ such that for all large enough n , all $k, k' \notin L_n$ and all m, m' ,*

$$|\Pr[A(\mathcal{O}(k, m)) = 1] - \Pr[A(\mathcal{O}(k', m')) = 1]| < \frac{1}{p(n)}$$

We prove Theorem 2.1 by showing that fully-entropic security implies the distributional indistinguishability property, which in turn implies oracle indistinguishability, which finally implies the virtual black-box property. The main ideas in these proofs are adapted from [10] to the multi-bit setting.

Lemma A.1. *If an obfuscator \mathcal{O} satisfies fully-entropic security, then it satisfies distributional indistinguishability.*

Proof. Suppose for the sake of contradiction that \mathcal{O} does not satisfy distributional indistinguishability, so there exists an adversary \mathcal{A} , distribution $\{X_n, Y_n\}$ with $H_\infty(X_n) \in \omega(\log(n))$, distinguisher D , and polynomial p such that for infinitely many values of n ,

$$|\Pr[D(k, m, \mathcal{A}(\mathcal{O}(k, m))) = 1] - \Pr[D(k, m, \mathcal{A}(\mathcal{O}(k', m')) = 1)]| \geq \frac{1}{10p(n)}, \quad (19)$$

where (k, m) and (k', m') are independently sampled from $\{X_n, Y_n\}$. Let $\beta(n) = H_\infty(X_n)$. We show that the same adversary \mathcal{A} breaks the fully-entropic security of \mathcal{O} because it breaks α -entropic security for $\alpha(n) = \beta(n) - \log(10p(n))$. Note that $\alpha \in \omega(\log(n))$ as desired.

Define $P_{k,m} = \Pr[\mathcal{A}(\mathcal{O}(k, m)) = 1]$, where the probability is taken over the randomness of \mathcal{A} and \mathcal{O} . It follows from (19) that there exist two sets $Z_1, Z_0 \subset \{0, 1\}^n$ such that:

- For any $(k, m) \in Z_1, (k', m') \in Z_0$ we have $P_{k,m} - P_{k',m'} > \frac{1}{10p(n)}$.
- The sets are large: for $(k, m) \leftarrow \{X_n, Y_n\}$, $\Pr[(k, m) \in Z_1] = \Pr[(k, m) \in Z_0] = \frac{1}{10p(n)}$.

Let $\{X_n^1, Y_n^1\}$ and $\{X_n^0, Y_n^0\}$ be the distributions formed by taking $\{X_n, Y_n\}$ and conditioning on the event that a key-message pair is chosen from Z_1 or Z_0 , respectively. We claim that the two distributions $\{X_n^b, Y_n^b\}$ for $b \in \{0, 1\}$ each have min-entropy α . This holds because for any k, m ,

$$\Pr[\{X_n^b, Y_n^b\} = (k, m)] \leq \Pr[\{X_n, Y_n\} = (k, m)] \cdot 10p(n)$$

since equality holds if $(k, m) \in Z_b$ or the left side probability equals 0 if $(k, m) \notin Z_b$. Therefore, by the union bound,

$$\Pr[X_n^b = k] \leq \Pr[X_n = k] \cdot 10p(n) \leq 2^{-\beta} \cdot 10p(n) = 2^{-\alpha}.$$

As a result, given any PPT simulator $\mathcal{S}^{I(k,m)}$ where (k, m) is chosen from either $\{X_n^1, Y_n^1\}$ or $\{X_n^0, Y_n^0\}$, the simulator only queries the correct key k with negligible probability. Hence,

$$\Pr[\mathcal{S}^{I(k,m)}(1^{|k|}) = 1 \mid (k, m) \leftarrow \{X_n^1, Y_n^1\}] - \Pr[\mathcal{S}^{I(k,m)}(1^{|k|}) = 1 \mid (k, m) \leftarrow \{X_n^0, Y_n^0\}]$$

is negligible. On the other hand, we know that

$$\Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1 \mid (k, m) \leftarrow \{X_n^1, Y_n^1\}] - \Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1 \mid (k, m) \leftarrow \{X_n^0, Y_n^0\}] > \frac{1}{10p(n)}.$$

As a result, it follows by the triangle inequality that for any simulator \mathcal{S} , either

$$\Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1 \mid (k, m) \leftarrow \{X_n^1, Y_n^1\}] - \Pr[\mathcal{S}^{I_{(k,m)}}(1^{|k|}) = 1 \mid (k, m) \leftarrow \{X_n^1, Y_n^1\}] > \frac{1}{20p(n)},$$

or the corresponding inequality holds for $\{X_n^0, Y_n^0\}$. Hence, one of these distributions breaks the α -entropic security of \mathcal{O} and thus the fully-entropic security of \mathcal{O} , as desired. \square

Lemma A.2. *If an obfuscator \mathcal{O} satisfies distributional indistinguishability, then it satisfies oracle indistinguishability.*

Proof. Assume for the sake of contradiction that there exists a PPT adversary \mathcal{A} and a polynomial p that break oracle indistinguishability. We define the following constants:

$$P_{k,m} = \Pr[\mathcal{A}(\mathcal{O}(k, m)) = 1], \quad m_k^1 = \arg \max_m \{P_{k,m}\}, \quad m_k^0 = \arg \min_m \{P_{k,m}\}, \quad P_k^1 = P_{k,m_k^1}, \quad P_k^0 = P_{k,m_k^0}.$$

Because \mathcal{A} and p break oracle indistinguishability, for any polynomial-sized family of sets $\{L_n\}_{n \in \mathbb{N}}$ and for infinitely many values of n , there exist $k, k' \notin L_n$ and m, m' such that $P_{k,m} - P_{k',m'} \geq \frac{1}{p(n)}$. Without loss of generality, we can assume $m = m_k^1$ and $m' = m_{k'}^0$, since $P_{k,m_k^1} \geq P_{k,m}$ and $P_{k',m_{k'}^0} < P_{k',m'}$.

Hence, for any polynomial-sized family of sets $\{L_n\}_{n \in \mathbb{N}}$ and for infinitely many values of n , there exist $k, k' \notin L_n$ such that

$$P_k^1 - P_{k'}^0 \geq \frac{1}{p(n)}. \quad (20)$$

For a given constant $c \in \mathbb{N}$, construct the family of sets $\{L_n^c\}_{n \in \mathbb{N}}$ in the following manner. The set $L_n^c = S_n^c \cup T_n^c$, where S_n^c is the set of the n^c keys k with the maximal values of P_k^1 , and T_n^c is the set of n^c keys k' with the minimal $P_{k'}^0$. Clearly, $|L_n^c| \leq |S_n^c| + |T_n^c| = 2n^c$ so the family $\{L_n^c\}$ is polynomially-bounded in size. Hence, for any $c \in \mathbb{N}$, and for all n such that (20) holds for the family $\{L_n^c\}$, we have that any keys $k \in S_n^c$ and $k' \in T_n^c$ satisfy

$$P_k^1 - P_{k'}^0 \geq \frac{1}{p(n)}.$$

Next, we form the families $\{\tilde{S}_n\}_{n \in \mathbb{N}}$ and $\{\tilde{T}_n\}_{n \in \mathbb{N}}$ as follows. Given n , let c_n be the largest value such that (20) is satisfied with respect to n and $L_{c_n}^{c_n}$. Then, \tilde{S}_n is defined recursively, as follows.

1. The base case is $\tilde{S}_0 = S_0^{c_0}$ and $\tilde{T}_0 = T_0^{c_0}$.
2. For $n > 0$, let n' be such that $\tilde{S}_{n-1} = S_{n-1}^{c_{n-1}}$. Then, \tilde{S}_n equals the largest set out of $S_n^{c_n}$ and $S_{n'}^{c_{n'}}$.

We define \tilde{T}_n analogously. Finally, we form the distribution $\{X_n, Y_n\}$ that is uniform over the key-message pairs (k, m_k^1) for all $k \in \tilde{S}_n$ and the key-message pairs $(k', m_{k'}^0)$ for all $k' \in \tilde{T}_n$. This distribution is well-spread, because given any polynomial n^d , there exists a value n_0 such that $|\tilde{S}_n| = |\tilde{T}_n| > n^d$ for all $n > n_0$.

We show that there exists a distinguisher D such that for infinitely many values of n ,

$$\Pr[D(k, m, \mathcal{A}(\mathcal{O}(k, m))) = 1] - \Pr[D(k, m, \mathcal{A}(\mathcal{O}(k', m')))) = 1] \geq \frac{1}{3p(n)}, \quad (21)$$

where (k, m) and (k', m') are independently drawn from $\{X_n, Y_n\}$.

We construct the distinguisher D as follows. Let \tilde{P} be a constant such that $P_k^1 - \tilde{P} \geq \frac{1}{2p(n)}$ for all $k \in \tilde{S}_n$ and $\tilde{P} - P_k^0 \geq \frac{1}{2p(n)}$ for all $k \in \tilde{T}_n$. This is known to D by non-uniformity. The distinguisher receives as input a

key k , a message m , and a bit b . It estimates $P_{k,m}$ by sampling $\mathcal{A}(\mathcal{O}(k, m))$ for many independent choices of the randomness for \mathcal{A} and \mathcal{O} . If its estimate of $P_{k,m}$ is at least \tilde{P} , then D outputs b . Otherwise, it outputs $1 - b$.

We demonstrate that the distinguisher D satisfies (21) for all n such that $\tilde{S}_n = S_n^c$ for some c . There are infinitely many such n 's. Our distribution has the property that for $(k, m) \leftarrow \{X_n, Y_n\}$, the value $P_{k,m}$ is bigger than \tilde{P}_n with probability $\frac{1}{2}$ and is smaller with probability $\frac{1}{2}$. Also, the distinguisher will make the correct determination on whether $P_{k,m}$ is bigger or smaller than \tilde{P}_n with overwhelming probability.

The distinguisher receives as input k, m , and a bit $b = \mathcal{A}(\mathcal{O}(k', m'))$, where (k', m') either equals (k, m) or is an independent sample from $\{X_n, Y_n\}$. Some basic probability calculations show that

$$\Pr[D = 1] \begin{cases} \geq \frac{1}{2} + \frac{1}{2p(n)} - \text{negl}(n) & \text{when } P_{k,m} \text{ and } P_{k',m'} \text{ are either both larger or both smaller than } \tilde{P}_n, \\ \leq \frac{1}{2} - \frac{1}{2p(n)} + \text{negl}(n) & \text{when } P_{k,m} \text{ and } P_{k',m'} \text{ are separated by } \tilde{P}_n, \end{cases}$$

where this probability is taken over $(k, m) \leftarrow \{X_n, Y_n\}$ and the randomness of D, \mathcal{A} , and \mathcal{O} . When $(k', m') = (k, m)$, the first case always holds, and when (k', m') is an independent sample, the two cases each hold with probability $\frac{1}{2}$. Therefore,

$$\Pr[D = 1 | (k', m') = (k, m)] - \Pr[D = 1 | (k', m') \text{ is indep sample from } \{X_n, Y_n\}] \geq \frac{1}{2p(n)} - \text{negl}(n),$$

as desired. \square

Lemma A.3. *If an obfuscator \mathcal{O} satisfies oracle indistinguishability, then it satisfies the virtual black-box property.*

Proof. Assume that oracle indistinguishability holds. Let \mathcal{A} be a PPT adversary that outputs 1 bit and let $\{L_n\}$ be the polynomial-sized family of sets associated to \mathcal{A} . We form a simulator $\mathcal{S}^{I(k,m)}$ that queries its oracle on all of the keys in L_n . If $k \in L_n$, then the simulator learns k and m , and it emulates an execution of $\mathcal{A}(\mathcal{O}(k, m))$. In this case, its simulation is perfect.

Otherwise, the simulator can run $\mathcal{A}(\mathcal{O}(k', m'))$ for any $k' \notin L_n$ and any m' . By α -oracle indistinguishability,

$$\Pr[\mathcal{S}^{I(k,m)}(1^{|k|}) = 1] = \Pr[\mathcal{A}(\mathcal{O}(I_{(k',m')})) = 1] \approx \Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1]$$

where the \approx denotes a negligible difference in probability. Finally, the simulator's runtime is bounded by the size of L_n and the runtime of \mathcal{A} , so \mathcal{S} runs in polynomial time as desired. \square

B Comparison of α -obfuscation definitions

In this section, we discuss the definition of α -entropic security. The goal of the definition is to weaken the virtual black-box property in Definition 2.2, which is very strong and thus far can only be satisfied under non-standard assumptions such as a strong variant of the DDH assumption [10], an exponentially hard to invert one-way function [25], or the random oracle model [21].

We believe there are two natural ways to weaken the virtual black-box property. First, we can increase the min-entropy requirement, as α -entropic security does. Second, we can give the simulator a super-polynomial runtime so it can make more queries to its oracle.

In the paper, we chose to do the former. This section justifies that decision by showing that α -entropic security implies a virtual black-box style definition in which the simulator receives a boost to its running time.

Definition B.1 (α -runtime security). *For any PPT adversary \mathcal{A} with 1 bit output, there exists a negligible function $\varepsilon(n)$ and a simulator \mathcal{S} running in time $\varepsilon(n) \cdot 2^{\alpha(n)}$ such that for all distributions $\{X_n, Y_n\}_{n \in \mathbb{N}}$ with X_n taking values in $\{0, 1\}^n$ and Y_n taking values in $\{0, 1\}^{\text{poly}(n)}$, we have:*

$$\left| \Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr[\mathcal{S}^{\mathcal{I}(k,m)}(1^{|k|}) = 1] \right| \leq \text{negl}(n)$$

where the probability is taken over the randomness of $(k, m) \leftarrow (X_n, Y_n)$, \mathcal{A} , \mathcal{S} , and \mathcal{O} .

In this definition, we choose to give the simulator $\text{negl}(n) \cdot 2^\alpha$ running time so that it does not quite have enough time to query everything in the support of a distribution with min-entropy α , but other than this restriction \mathcal{S} has the “largest” runtime possible.

Theorem B.1. *If an obfuscator satisfies α -entropic security, then it satisfies α -runtime security.*

The rest of this section is devoted to a proof of this theorem. We do not prove the theorem directly, but rather go through an intermediate definition from [10].

Definition B.2 (α -oracle indistinguishability [10]). *For any PPT adversary \mathcal{A} with 1 bit output, there exists a negligible function $\varepsilon(n)$ and a family of sets $\{L_n\}_{n \in \mathbb{N}}$ such that $|L_n| \leq \varepsilon(n) \cdot 2^{\alpha(n)}$ and for all $k, k' \notin L_n$ and all m, m' ,*

$$|\Pr[A(\mathcal{O}(k, m)) = 1] - \Pr[A(\mathcal{O}(k', m')) = 1]| < \text{negl}(n)$$

Lemma B.1. *If an obfuscator satisfies α -oracle indistinguishability, then it satisfies α -runtime security.*

Proof. Let \mathcal{A} be a PPT adversary that outputs 1 bit. We form a simulator $\mathcal{S}^{I(k, m)}$ that queries its oracle on all of the keys in L_n . If $k \in L_n$, then the simulator learns k and m , and it emulates an execution of $\mathcal{A}(\mathcal{O}(k, m))$. In this case, its simulation is perfect.

Otherwise, the simulator can run $\mathcal{A}(\mathcal{O}(k', m'))$ for any $k' \notin L_n$ and any m' . By α -oracle indistinguishability,

$$\Pr[\mathcal{S}^{I(k, m)}(1^{|k|}) = 1] = \Pr[\mathcal{A}(\mathcal{O}(I_{(k', m')})) = 1] \approx \Pr[\mathcal{A}(\mathcal{O}(I_{(k, m)})) = 1]$$

where the \approx denotes a negligible difference in probability, as desired. \square

Lemma B.2. *If an obfuscator satisfies α -entropic security, then it satisfies α -oracle indistinguishability.*

Proof. Let \mathcal{O} be an obfuscator satisfying α -entropic security, and let \mathcal{A} be an adversary. We wish to show the existence of a negligible function ε and family of sets $\{L_n\}_{n \in \mathbb{N}}$ that satisfy α -oracle indistinguishability.

Given any k and m , we define $P_{k, m} = \Pr[\mathcal{A}(\mathcal{O}(k, m)) = 1]$. Also, we define the following constants:

$$\mu_k = \text{average}_m \{P_{k, m}\}, \quad m_k^1 = \arg \max_m \{P_{k, m}\}, \quad m_k^0 = \arg \min_m \{P_{k, m}\}, \quad \sigma_k = P_{k, m_k^1} - P_{k, m_k^0}$$

(Note that if the $\arg \max$ or $\arg \min$ are simultaneously fulfilled by many messages, then it suffices to pick any one of them arbitrarily.) Also, let \mathcal{S} be the PPT simulator associated with \mathcal{A} by α -entropic security. Clearly, $\mathcal{S}^{I(k, m)}$ does not learn any information about m unless it queries the correct key k , which it can only do for polynomially many keys. By α -entropic security, it follows that there exists a negligible function ε' such that at most $\varepsilon' \cdot 2^\alpha$ keys have a non-negligible σ_k .

If this were not the case, then there exists some polynomials p, q such that there are at least $\frac{2^{\alpha(n)}}{p(n)}$ keys k with $\sigma_k > \frac{1}{q(n)}$. Let X_n be the distribution that is uniform over these $\frac{2^{\alpha(n)}}{p(n)}$ keys and $2^{\alpha(n)} - \frac{2^{\alpha(n)}}{p(n)}$ other keys chosen arbitrarily, and let Y_n^1 and Y_n^0 be two distributions on messages such that for a given key k , the distribution Y_n^b always chooses the message m_k^b . Both distributions $\{X_n, Y_n^b\}$ have min-entropy α , and we know that

$$\Pr[A(\mathcal{O}(k, m)) = 1 \mid (k, m) \leftarrow \{X_n, Y_n^1\}] - \Pr[A(\mathcal{O}(k, m)) = 1 \mid (k, m) \leftarrow \{X_n, Y_n^0\}] > \frac{1}{p(n)q(n)}.$$

On the other hand, the simulator cannot distinguish between these two distributions, so by the triangle inequality property, either

$$\Pr[A(\mathcal{O}(k, m)) = 1 \mid (k, m) \leftarrow \{X_n, Y_n^1\}] - \Pr[\mathcal{S}^{I(k, m)}(1^{|k|}) = 1 \mid (k, m) \leftarrow \{X_n, Y_n^1\}] > \frac{1}{2p(n)q(n)},$$

or the corresponding inequality holds for $\{X_n, Y_n^0\}$, which breaks the α -entropic security.

Hence, there are at most $\varepsilon' \cdot 2^\alpha$ keys k such that σ_k is non-negligible. Let L'_n be the set of these keys. Next, we look at μ_k , and claim that there exists a negligible function ε'' and a set L''_n of size at most $\varepsilon'' \cdot 2^\alpha$ such that for all $k, k' \notin L''_n$, $|\mu_k - \mu_{k'}|$ is negligible.

If this were not the case, then there exists some polynomial $r(n)$ such that at least $\frac{2^\alpha}{r}$ keys have μ_k that are noticeably separated from $\mu = \text{average}_k\{\mu_k\}$. Using a similar proof to the one for σ_k above, we can then form two well-spread distributions over these keys such that the adversary can distinguish them but the simulator cannot.

Finally, let $\varepsilon = \varepsilon' + \varepsilon''$ and let $\{L_n\}_{n \in \mathbb{N}}$ be a family of sets with $L_n = L'_n \cup L''_n$. For all keys $k, k' \notin L_n$, we know that σ_k , $|\mu_k - \mu_{k'}|$, and $\sigma_{k'}$ are negligible, which means that for all messages m and m' ,

$$\Pr[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] \approx \mu_k \approx \mu_{k'} \approx \Pr[\mathcal{A}(\mathcal{O}(I_{(k',m')})) = 1],$$

where \approx denotes a negligible difference in probability, as desired. \square

Finally, Theorem B.1 follows immediately from Lemmas B.1 and B.2.