On Symmetric Encryption and Point Obfuscation

Ran Canetti[1] *, Yael Tauman Kalai[2], Mayank Varia[3] **, and Daniel Wichs[4] ***

[1] School of Computer Science, Tel Aviv University. (`canetti@post.tau.ac.il`)
[2] Microsoft Research New England (`yael@microsoft.com`)
[3] Massachusetts Institute of Technology (`varia@csail.mit.edu`)
[4] New York University (`wichs@cs.nyu.edu`)

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.

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 naive) 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, 14, 21, 2, 25] 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, 25]). This models both direct attacks where the adversary gains access to the internal storage of the parties, such as the freezing attack of [18], 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 of 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 the *computational* setting.[5] 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.[6]

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 [19, 5, 9, 4, 20, 17, 8, 3], and several positive results emerged, most notably the results of [8, 3] 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, [17] 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,

---

[5] Many other models of leakage-resilience, such as the "only computation leaks information" model [23, 15], place further restrictions on the type of information that may be leaked, and are not considered in this work.

[6] In 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, $Dec'_k(r, c) = Dec_{k+r}(c))$.

25], [14, 3], 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 gives no information (or, rather, no computational ability) that cannot be determined given only oracle access to p. A central point here is that $\mathcal{O}$ should work correctly and securely for *any* 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, 26] under strong assumptions (and in [22] in the random oracle model). Obfuscators for MBPF are only known based on very strong and specific assumptions (specifically, 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, 14]. 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 $O(I_{(k,m)})$ 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 (i.e., it cannot be guessed in polynomial time.) Also, [14] 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 above primitives. Specifically, we show that 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, a generalized version of KDM security, which also withstands the case where the key is taken from a weakly random distribution, is also a special case of MBPF obfuscation. 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.

The remainder of the introduction overviews our results. We first present the general connections between obfuscation and symmetric encryption; next we sketch some conclusions and corollaries.

As a preliminary step towards drawing general 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 here the simulator knows the class of distributions, but not the distribution itself. This relaxation allows us to relate the different classes of strong encryption to MBPF obfuscators for different classes of distributions. Specifically:

Obfuscation vs. Weak-Key and Leakage Resilient Encryption: We say that an MBPF obfuscator is α-entropic with independent messages if it is an *MBPF obfuscator* for product distributions on k, m, where the distribution of k has min-entropy at least α, and m is drawn *independently* of k, but need not have any entropy. We say that the obfuscator is a fully-entropic IM MBPF if it has α-entropic security for all super-logarithmic α. We show:

From IM MBPF obfuscators to encryption. Any α-entropic IM MBPF obfuscator with independent messages allows us to construct semantically secure encryption scheme with security for α-weak keys, via the transformation $\mathsf{Enc}_k(m) = \mathcal{O}(I_{(k,m)})$.

From encryption to IM MBPF obfuscators. Conversely, any encryption scheme with semantic security for α-weak keys allows us to construct α-entropic IM MBPF obfuscators. The transformation is simple: To obfuscate a pair k, m, simply encrypt m with key k to obtain a ciphertext c; then, the obfuscated program simply has a hard-coded ciphertext c, and on input x, runs the decryption algorithm on c with the key x. 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 for α-weak keys, then the resulting IM MBPF obfuscator $\mathcal{O}$ is *self-composable*, in the sense that security is preserved even if $\mathcal{O}$ is run multiple times on MBPFs with the *same input* k and (possibly) different outputs m_i. As was shown by [11], this property is not, in general, implied by obfuscation alone. Conversely, if we start with a self-composable IM MBPF obfuscator then we derive an encryption scheme which is CPA secure for α-weak keys.

Fully-entropic obfuscation and fully-weak key security. If we start with an IM MBPF obfuscator that has full-entropic security (i.e., it works for any distribution where k is independent from m and has some super-logarithmic min-entropy) then we obtain an encryption scheme with semantic-security for fully-weak keys. (i.e. security for any key-distribution with super-logarithmic entropy).

Computational leakage vs. auxiliary information. If we start from a computational leakage resilient encryption then the resulting MBPF obfuscator is secure with respect to dependent auxiliary input, as defined in [16]. Similarly, if we start from a MBPF obfuscator that's secure with dependent auxiliary input then the resulting encryption scheme is computationally leakage resilient.

KDM security: All of the above equivalence results in the preceding paragraph 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 a α-*entropic (dependent) MBPF obfuscator* if it withstands any *joint* distribution on k, m where the projection distribution on k has min-entropy at least α (and m may depend on k). We say that $\mathcal{O}$ is a fully-entropic (dependent) MBPF obfuscator if the above holds for all super-logarithmic α .

We also define α-KDM encryption schemes which provide security *even* when the key is taken from any distribution of entropy α, *and* the message can be an arbitrary function of the secret key. We show:

Obfuscation vs. encryption. Any α-entropic (dependent) MBPF obfuscator provides, via the same transformation as before, an α-KDM semantically secure encryption scheme.

Multi message resilience vs. self composability. If the encryption scheme we start with is *multi-message* α-KDM secure, in the sense that it withstands the case where the adversary obtains encryptions of any polynomial number of functions of the secret key, then the resulting (dependent) MBPF obfuscator has α-entropic security and is *self composable*. The converse implication holds as well.

To connect our new α-entropic definition to previous works, we show that any MBPF obfuscator that is α-entropic for any super-logarithmic α 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 VBB case no such bound exists.)

1.2 Implications

We show some implications of the above correspondence results:

Secure encryption w.r.t. (fully) weak keys. Known constructions of encryption schemes that are secure w.r.t. weak keys are parameterized by the min-entropy α tolerated. That is, a bound α must be chosen in advance, and then a 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 is secure w.r.t. α-weak keys, for *any* super-logarithmic function α. The main advantage is that the min-entropy α does not need to be chosen in advance. More specifically, we obtain a single encryption scheme, parameterized only by the security parameter n (and *not* by α), which simultaneously achieves security for all $\alpha(n) \in \omega(\log n)$.

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. See Section 6.1 for further details.

Impossibility for MBPF Obfuscators and fully composable point function obfuscators. Using our transformations, the negative result due to Haitner and Holenstein [17] implies that there are no constructions of MBPF obfuscators that can 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. See Section 6.2 for further details.

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. See Section 6.2 for further details.

Organization. Section 2 contains some basic definitions for obfuscation and encryption. Section 3 draws connections between obfuscation and weak key and leakage resilient encryption. Section 4 draws connections between obfuscation and encryption resilient to key dependent messages. Section 6 states the corollaries that we draw from the general connections. Many proofs are left out and appear only in the full version [12].

2 Definitions

2.1 Obfuscation of Point Functions with Multi-bit Output

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

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

which outputs the *message* m given the *key* k, and $\bot$ 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 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:*

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

*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 (Virtual black-box property [10, 6, 26]). *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:*

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

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

Note the difference 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 1. *If $\mathcal{O}$ is a MBPF obfuscator that satisfies* fully-entropic security *(as in Definition 1) then $\mathcal{O}$ also satisfies* virtual black-box obfuscation *(as in Definition 2).*

The proof of this theorem appears in the full version of this paper [12]. 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 there 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 use this to prove that fully-entropic security provides 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 [13] 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 the full version [12], 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 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 $I_{(k,m_1)}, \ldots, I_{(k,m_t)}$ with the *same key k* but *potentially different messages* $m_1, \ldots, m_t$. (For point functions, self composition boils down to the case of many obfuscated versions of the same function.)

Definition 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 $X_n = X_n^{(1)}, \ldots, X_n^{(t)}$, $Y_n = Y_n^{(1)}, \ldots, Y_n^{(t)}$, 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:*

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

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

If the above holds only for the distributions X_n where $\Pr[k_1 = k_2 \ldots = k_t] = 1$, 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 $(\mathsf{Enc}, \mathsf{Dec})$.[7] 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 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 $\mathrm{poly}(n, \ell)$, 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:*

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

---

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

where the games $\mathrm{CPA}_b^{X,D}(\mathcal{A}, n)$ *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 $\mathrm{CPA}_0^{X,D}$, *the challenger sets* $c \leftarrow \mathsf{Enc}_k(m)$.
   In game $\mathrm{CPA}_1^{X,D}$, *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* $\mathrm{SEM}_0^{X,D}, \mathrm{SEM}_1^{X,D}$, *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 $\mathsf{Enc}_k(0^\ell)$, 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 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 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* $m \in \{0,1\}^{\mathrm{poly}(n)}$, $\Pr[\mathsf{Dec}_{k'}(\mathsf{Enc}_k(m)) \neq \bot] \leq \mathrm{negl}(n)$.

We note that a similar, but weaker, property called confusion freeness, was defined in [24]. 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 1 (see the full version [12] for proof) 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 1. *Let* $(\mathsf{Enc}, \mathsf{Dec})$ *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* $(\mathsf{Enc}', \mathsf{Dec}')$ *by:*

$$\mathsf{Enc}'_k(m) \triangleq \left\{ \begin{array}{c} \textit{Choose: } h \leftarrow \mathcal{H}, r \leftarrow U_n \\ \textit{Output: } \langle r, h, c = \mathsf{Enc}_{h(k)}(r\|m) \rangle \end{array} \right.$$

$$\mathsf{Dec}'_k(\langle r, h, c \rangle) \triangleq \left\{ \begin{array}{c} \textit{Compute: } (r'\|m') = \mathsf{Dec}_{h(k)}(c) \\ \textit{Output: } m' \textit{ if } r' = r \textit{ and } \perp \textit{ otherwise} \end{array} \right.$$

Then $(\mathsf{Enc}', \mathsf{Dec}')$ *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".*

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 2. *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 Lemma 2 and Lemma 3, respectively.

Lemma 2. *Let* $\alpha(n) \in \omega(\log(n))$ *and let* $\mathcal{O}$ *be a MBPF obfuscator with* $\alpha(n)$*-entropic security for independent messages. Let* $\mathsf{Enc}_k(m) \triangleq \mathcal{O}(I_{(k,m)})$, $\mathsf{Dec}_k(C) \triangleq C(k)$ *where the ciphertext* C *is interpreted as a circuit. Then the encryption scheme* $(\mathsf{Enc}, \mathsf{Dec})$ *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

$$\left| \Pr[\mathrm{SEM}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\mathrm{SEM}_1^{X,D}(\mathcal{A}, n) = 1] \right|$$

$$= \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)}[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (U_n, U_\ell)}[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] \right|$$

$$\leq \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)}[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] - \Pr_{(k,m) \leftarrow (X_n, Y_n)}\left[\mathcal{S}^{I_{(k,m)}(\cdot)}(1^n) = 1\right] \right| \quad (1)$$

$$+ \left| \Pr_{(k,m) \leftarrow (X_n, Y_n)}\left[\mathcal{S}^{I_{(k,m)}}(1^n) = 1\right] - \Pr_{(k,m) \leftarrow (U_n, U_\ell)}\left[\mathcal{S}^{I_{(k,m)}}(1^n) = 1\right] \right| \quad (2)$$

$$+ \left| \Pr_{(k,m) \leftarrow (U_n, U_\ell)}\left[\mathcal{S}^{I_{(k,m)}}(1^n) = 1\right] - \Pr_{(k,m) \leftarrow (U_n, U_\ell)}[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1] \right| \quad (3)$$

$$\leq \mathrm{negl}(n)$$

where (1),(3) follow by the definition of entropic security of obfuscation, and (2) 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 $\alpha(n)$. □

Lemma 3. *Let* $(\mathsf{Enc}, \mathsf{Dec})$ *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* $c = \mathsf{Enc}_k(m)$ *and outputs the circuit* $C_c(\cdot)$ *defined by* $C_c(x) \triangleq \mathsf{Dec}_x(c)$. *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 $m \in \{0,1\}^{\mathrm{poly}(n)}$. If $k = x$ then

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

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

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

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

$$\left| \Pr_{(k,m)\leftarrow(X_n, Y_n)} \left[\mathcal{A}(\mathcal{O}(I_{(k,m)})) = 1 \right] - \Pr_{(k,m)\leftarrow(X_n, Y_n)} \left[\mathcal{S}^{I_{(k,m)}}(1^n, 1^\ell) = 1 \right] \right| \quad (4)$$

$$= \left| \Pr \left[\mathcal{A}(C_c) = 1 \ \middle| \ \begin{array}{c} (k, m) \leftarrow (X_n, Y_n) \\ c \leftarrow \mathsf{Enc}_k(m) \end{array} \right] - \Pr \left[\mathcal{A}(C_c) = 1 \mid c \leftarrow D(n, \ell) \right] \right|$$

$$\leq \mathrm{negl}(n) \quad (5)$$

Where (5) follows by semantic-security. □

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. *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 4) and then we show that encryption implies obfuscation (Lemma 5).

In the next lemma, going from obfuscation to encryption, it would be natural to define $\mathsf{Enc}_k(m) = \mathcal{O}(I_{(k,m)})$. However, we instead define $\mathsf{Enc}_k(m) = (\mathcal{O}(I_{(k,r)}), m \oplus r)$ 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 4. *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 $(\mathsf{Enc}, \mathsf{Dec})$ by*

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

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

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

Lemma 5. *Let $(\mathsf{Enc}, \mathsf{Dec})$ 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 $c = \mathsf{Enc}_k(m)$ and outputs the circuit $C_c(\cdot)$ defined by $C_c(x) = \mathsf{Dec}_x(c)$. Then, $\mathcal{O}$ is a **self-composable** MBPF obfuscator with $\alpha(n)$-entropic security for independent messages.*

4   KDM Encryption and MBPF Obfuscation

First, we define the notion of semantically-secure encryption with *key dependent messages* (KDM) and $\alpha(n)$-weak keys.

Definition 7 (Semantic KDM Encryption with Weak Keys).

*A symmetric encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ 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 $\mathrm{poly}(n, \ell)$, 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:*

$$|\Pr[\mathrm{KDM}_0^{X,D}(\mathcal{A}, n) = 1] - \Pr[\mathrm{KDM}_1^{X,D}(\mathcal{A}, n) = 1]| \leq \mathrm{negl}(n), \qquad (6)$$

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

> $k \leftarrow X_n$
> $c_0 \leftarrow \mathsf{Enc}_k(f(k)), c_1 \leftarrow D(n, \ell)$ *where ℓ is the output size of f.*
> *Output: $\mathcal{A}(c_b)$*

We now show that semantically secure encryption with KDM and security for weak keys is equivalent to MBPF obfuscation.

Theorem 4. *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 theorems follows from essentially the same arguments as in Lemma 2 and Lemma 3. We simply observe that allowing the adversary to get encryption of a value $f(k)$ in the proofs of those lemmas, corresponds to having a distribution Y_n that depends on X_n, that it $Y_n = f(X_n)$. Conversely, for any joint distribution $\{X_n, Y_n\}$, we can define some (probabilistic, and possibly inefficient) function f so that $Y_n = f(X_n)$.

In the full version of this paper [12], we also 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.

5 Encryption/Obfuscation with Auxiliary Input

In the full version of this work [12] we also define encryption with semantic/CPA security *with auxiliary input family* $\mathcal{F}$, where the adversary gets to learn $f(k)$ for any $f \in \mathcal{F}$. [8] 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 defines for $\alpha(n)$-weak keys as well as fully weak keys.

We show that all of the results of Section 3 extend naturally to the auxiliary input setting. That is:

– We extend Theorem 2, to show an equivalence between semantically secure encryption with auxiliary-input family $\mathcal{F}$ and wrong-key detection, and obfuscation of MBPF with auxiliary-input family $\mathcal{F}$ and independent messages. The equivalence holds for $\alpha(n)$-weak keys or "fully weak" keys. The constructions are the same as those of Lemma 2 and Lemma 3.

– We similarly extend Theorem 3 showing a similar equivalence for CPA secure encryption and self-composable onfuscation with auxiliary input. The constructions are the same as those of Lemma 4 and Lemma 5.

---

[8] This 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 . . .

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 4), 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 8 (Strengthened DDH Assumption [10]). *Let $\mathbb{G}$ be a group of prime order $p = 2^{\mathrm{poly}(n)}$ 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 $F : \mathbb{Z}_p \to \mathbb{G} \times \mathbb{G}$ 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:

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

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 4, we get the following theorem.

Theorem 5. *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 5 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

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

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.[9]

One can use this encryption scheme to construct one which also has the wrong-key detection property using Lemma 1. Such a scheme yields an multi-bit point function obfuscator with $\alpha(n)$-entropic security for independent messages, by Lemma 3.

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, 14, 25]. 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 5) "out of the box". However, there are several issues that we must deal with first.

– *Efficiently-Sampleable Distributions:* The works of [1, 14, 25] 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 $\alpha(n)$ [25]. 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 [14] is explicitly designed for the symmetric key setting. The schemes of [1, 25] 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 indistinguish-

---

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

able 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 [14]. For the schemes of [1, 25], we get the property that, given the public parameters it is computationally difficult to find k, k' such that $\mathsf{Dec}_{k'}(\mathsf{Enc}_k(m)) \neq \bot$. 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 5, using the schemes of [1, 14, 25].

Theorem 6. *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 [25].* [*,†].
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 [14].* [*].

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 [17], 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 [17] 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.[10] Since, by Theorem 4, 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.

---

[10] On 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.

Theorem 7. *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 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 [17] 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.
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, Y. T. Kalai, M. Varia, and D. Wichs. On symmetric encryption and point obfuscation. 2010. Full Version. Cryptology ePrint Archive.

13. 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.

14. Y. Dodis, Y. T. Kalai, and S. Lovett. On cryptography with auxiliary input. In *STOC*, pages 621–630, 2009.

15. S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302, 2008.

16. S. Goldwasser and Y. T. Kalai. On the impossibility of obfuscation with auxiliary input. In *FOCS*, pages 553–562, 2005.

17. I. Haitner and T. Holenstein. On the (im)possibility of key dependent encryption. In *TCC*, pages 202–219, 2009.

18. 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.

19. S. Halevi and H. Krawczyk. Security under key-dependent inputs. In *ACM Conference on Computer and Communications Security*, pages 466–475, 2007.

20. D. Hofheinz and D. Unruh. Towards key-dependent message security in the standard model. In *EUROCRYPT*, pages 108–126, 2008.

21. 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.

22. 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.

23. S. Micali and L. Reyzin. Physically observable cryptography (extended abstract). In *TCC*, pages 278–296, 2004.

24. D. Micciancio and B. Warinschi. Completeness theorems for the abadi-rogaway language of encrypted expressions. *Journal of Computer Security*, 12(1):99–130, 2004.

25. 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.

26. H. Wee. On obfuscating point functions. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 523–532, 2005.