

A Unified Approach to Deterministic Encryption: New Constructions and a Connection to Computational Entropy

Benjamin Fuller¹, Adam O’Neill², and Leonid Reyzin²

¹ Boston University and MIT Lincoln Laboratory

² Boston University

Abstract. We propose a general construction of deterministic encryption schemes that unifies prior work and gives novel schemes. Specifically, its instantiations provide:

- A construction from any trapdoor function that has sufficiently many hardcore bits.
- A construction that provides “bounded” multi-message security from lossy trapdoor functions.

The security proofs for these schemes are enabled by three tools that are of broader interest:

- A weaker and more precise sufficient condition for semantic security on a high-entropy message distribution. Namely, we show that to establish semantic security on a distribution M of messages, it suffices to establish indistinguishability for all conditional distribution $M|E$, where E is an event of probability at least $1/4$. (Prior work required indistinguishability on *all* distributions of a given entropy.)
- A result about computational entropy of conditional distributions. Namely, we show that conditioning on an event E of probability p reduces the quality of computational entropy by a factor of p and its quantity by $\log_2 1/p$.
- A generalization of leftover hash lemma to correlated distributions.

We also extend our result about computational entropy to the average case, which is useful in reasoning about leakage-resilient cryptography: leaking λ bits of information reduces the quality of computational entropy by a factor of 2^λ and its quantity by λ .

1 Introduction

Public-key cryptosystems require randomness: indeed, if the encryption operation is deterministic, the adversary can simply use the public key to verify that the ciphertext c corresponds to its guess of the plaintext m by encrypting m . However, such an attack requires the adversary to have a reasonably likely guess for m in the first place. Recent results on deterministic public-key encryption (DE) (building on work in the information-theoretic symmetric-key setting [38,17,14]) have studied how to achieve security when the randomness

comes only from m itself [3,5,7,27,8,40]. DE has a number of practical applications, such as efficient search on encrypted data and securing legacy protocols (cf. [3]). It is also interesting from a foundational standpoint; indeed, its study has proven useful in other contexts: Bellare et al. [4] showed how it extends to a notion of “hedged” public-key encryption that reduces dependence on external randomness for probabilistic encryption more generally, and Dent et al. [13] adapted its notion of privacy to a notion of confidentiality for digital signatures.

However, our current understanding of DE is somewhat lacking. The constructions of [3,5,7,27], as well as their analysis techniques, are rather disparate, and some natural questions arise from them. Namely, does the scheme of [5] inherently require using the Goldreich-Levin hardcore bit? Can it be made to work with trapdoor functions rather than permutations? Is the single-message security achieved by [5,7,27] an inherent limitation of standard model (i.e., non-random-oracle) schemes? In this work our main goal is to provide a *unified framework* for the construction of DE and to shed light on these questions.

1.1 Our Results

A SCHEME BASED ON TRAPDOOR FUNCTIONS. We propose a general *Encrypt-with-Hardcore* (EwHCORE) construction of DE from trapdoor functions (TDFs), which generalizes the basic idea behind the schemes of [3,5] and leads to a unified framework for the construction of DE. Let f be a TDF with a hardcore function hc , and let \mathcal{E} be any probabilistic public-key encryption algorithm. Our scheme encrypts an input message x by computing $y = f(x)$ and then encrypting y using \mathcal{E} with $hc(x)$ as the coins; that is, the encryption of x is $\mathcal{E}(f(x); hc(x))$.

Intuitively, this scheme requires that the output of hc be sufficiently long to provide enough random coins for \mathcal{E} (in fact, it need only be sufficiently long to be used as a seed for a pseudorandom generator), and that it not reveal any partial information about x (because \mathcal{E} does not necessarily protect the privacy of its random coins). There are two nontrivial technical steps needed to make intuition precise. First, we define a condition required of hc (which we call “robustness”) and show that it is sufficient for security of the resulting DE. Second, through a computational entropy argument, we show how to make *any* sufficiently long hc robust by applying a randomness extractor.

This general scheme admits a number of instantiations depending of f and hc . For example, when f is any trapdoor function and hc is a random oracle (RO), we obtain the construction of [3]¹. When f is an iterated trapdoor permutation (TDP) and hc is a collection Goldreich-Levin (GL) [23] bits extracted at each iteration, we obtain the construction of [5]. When f is a lossy trapdoor function (LTDF) [35] and hc is a pairwise-independent hash, we get a variant of the construction of [7] (which is less efficient but has a more straightforward analysis). We also obtain a variant of the construction of Hemenway et al. [27]

¹ Technically, this construction does not even need a TDF because of the random oracle model; however, it may be prudent to use a TDF because then it seems more likely that the instantiation of the random oracle will be secure as it may be hardcore for the TDF.

under the same assumption as they use (see Section 5.2 for details). Note that in all but the last of these cases, the hardcore function is *already* robust (without requiring an extractor), which shows that in prior work this notion played an implicit role.

Moreover, this general scheme not only explains past constructions, but also gives us new ones. Specifically, if f is a trapdoor function with enough hardcore bits, we obtain:

- DE that works on the uniform distribution of messages;
- DE that works on any distribution of messages whose min-entropy is at most logarithmically smaller than maximum possible;
- assuming sufficient hardness distinguishing the output of hc from uniform (so in particular of inverting f), DE that works on even-lower entropy message distributions.

Prior results require more specific assumptions on the trapdoor function (such as assuming that it is a permutation or that it is lossy—both of which imply enough hardcore bits) in order to get constructions that work even just on the uniform distribution of messages. Furthermore, our results yield more efficient schemes (though sometimes under stronger assumptions) even in the permutation case, by avoiding iteration.

Notably, we obtain the *first* DE scheme without random oracles based on the hardness of syndrome decoding using the Niederreiter trapdoor function [32], which was shown to have linearly many hardcore bits by Freeman et al. [19] (and moreover to be “correlated input” secure) but is not known to be lossy. (A scheme in the random oracle model follows from [3].) Additionally, the RSA [37] and Paillier [34] trapdoor permutations have linearly many hardcore bits under certain computational assumptions (the “Small Solutions RSA” [39] and “Bounded Computational Composite Residuosity” [9] assumptions respectively). Therefore, we can use these TDPs to instantiate our scheme efficiently under the same computational assumptions. Before our work, DE schemes from RSA and Paillier either required many iterations [5] or decisional assumptions that imply lossiness of these TDPs [30,19,7].

SECURITY FOR MULTIPLE MESSAGES: DEFINITION AND CONSTRUCTION. An important caveat is that, as in [5,7], we can prove the above standard-model DE schemes secure only for the encryption of a *single* high-entropy plaintext, or, what was shown equivalent in [7], an unbounded number of messages drawn from a *block source* [10], where each subsequent message brings “fresh” entropy. On the other hand, the strongest and most practical security model for DE introduced by [3] considers the encryption of an unbounded number of plaintexts that have individual high entropy but may not have any conditional entropy. In order for EwHCORE to achieve this, the hardcore function hc must also be robust on *correlated inputs*. (A general study of correlated-input security for the case of hash functions rather than hardcore functions was concurrently initiated in [25].) In particular, it follows from the techniques of [3] that a RO hash satisfies such a notion. This leads to a multi-message secure scheme in the RO model

(as obtained in [3]). We thus have a large gap between what is (known to be) achievable with random oracles versus in the standard model.

To help bridge this gap, we propose a notion of “ q -bounded” security for DE, where up to q high-entropy but arbitrarily correlated messages may be encrypted under the same public key (whose size may depend polynomially on q). We feel that if one is limited to the standard model, this notion is useful. Indeed, it seems that the requirement of previous results in the standard model—that messages come from a block source—may be difficult to guarantee: all that’s needed to violate it is a single message that has low conditional entropy. Following [7], we also extend our security definition to unbounded multi-message security where messages are drawn from what we call a “ q -block source” (essentially, a block source where each “block” consists of q messages which may be arbitrarily correlated but have individual high entropy); Theorem 4.2 of [7] extends to show that q -bounded multi-message security and unbounded multi-message security for q -block sources are equivalent for a given min-entropy.

Using our EwHCore construction and a generalization of the leftover hash lemma discussed below, we show q -bounded DE schemes (for long enough messages), for any polynomial q , based on LTDFs losing an $1 - O(1/q)$ fraction of the input. It is known how to build such LTDFs from the decisional Diffie-Hellman [35], d -linear [19], and decisional composite residuosity [7,19] assumptions.

1.2 Our Tools

Our results are enabled by three tools that may be of more general applicability.

A MORE PRECISE CONDITION FOR SECURITY OF DE. We revisit the definitional equivalences for DE proven by [5] and [7]. At a high level, they showed that the semantic security style definition for DE (called PRIV) introduced in the initial work of [3], which asks that a scheme hides all public-key independent² functions of messages drawn from some distribution is in some sense equivalent to an indistinguishability based notion for DE, which asks that it is hard to distinguish ciphertexts of messages drawn from one of two possible distributions. Notice that while PRIV can be meaningfully said to hold for a given message distribution, IND inherently talks of *pairs* of distributions. The works of [5,7] compensated for this by giving an equivalences in terms of *min-entropy*. That is, they showed that PRIV for all message distributions of min-entropy μ is implied by indistinguishability with respect to all pairs of plaintext distributions of min-entropy slightly less than μ .

We demonstrate a more precise equivalence that, for a *fixed* distribution \mathbf{M} , identifies a class of pairs of distributions such that if IND holds on those pairs, then PRIV holds on \mathbf{M} . By re-examining the equivalence proof of [5], we show that PRIV on \mathbf{M} is implied by IND on all pairs of “slightly induced” distributions of $\mathbf{M} \mid \mathbf{E}$, where \mathbf{E} is an arbitrary event of probability at least $1/4$.

² As shown in [3], the restriction to public-key independent functions is inherent here.

This first tool is needed to argue that “robustness” of hc is sufficient for security EwHCORE (essentially, a robust hardcore function is one that remains hardcore on a slightly induced distribution³).

CONDITIONAL COMPUTATIONAL ENTROPY. We investigate how conditioning reduces computational entropy of a random variable X . Suppose you have a distribution that has *computational* entropy (such as the pair $f(r), \text{hc}(r)$ for a random r). Suppose you condition that distribution on an event E of probability p . How much computational entropy is left?

To make this question more precise, we should note that computational entropy is parameterized by quality (how distinguishable is X from a variable Z that has true entropy) and quantity (how much true entropy is there in Z).

We prove an intuitively natural result: conditioning on an event of probability p reduces the quality of metric entropy by a factor of p and the quantity of metric entropy by $\log_2 1/p$ (note that this means that the reduction in quantity and quality is the same, because the quantity of entropy is measured on log scale). Naturally, the answer becomes so simple only once the correct notion of entropy is in place. Our result holds for Metric^* entropy (defined in [2,18]). This entropy is convertible (with some loss) to HILL entropy [26,2], which can then be used with randomness extractors to get pseudorandom bits.

Our result improves the bounds of Dziembowski and Pietrzak [18, Lemma 3], where the loss in the *quantity* of entropy was related to its original *quality*. The use of metric entropy simplifies the analogous result of Reingold et al. [36, Theorem 1.3] for HILL entropy. (See [20] for information on other related work [22, Lemma 3.1] and [11, Lemma 16].)

We use this result to show that randomness extractors can be used to convert a hardcore function into a robust one, through a computational entropy argument for slightly induced distributions. The result is also applicable to leakage-resilient cryptography, as demonstrated by [18]. To make the result useful in more contexts, we also provide an average-case entropy formulation, which can be helpful in situations in which not all leakage is equally informative. For the information-theoretic case, it is known that leakage of λ bits reduces the average entropy by at most λ ([15, Lemma 2.2]). We show essentially the same⁴ for the computational case: if λ bits of information are leaked, then the amount of computational Metric^* entropy decreases by at most λ and its quality decreases by at most 2^λ (again, this entropy can be converted to HILL entropy and be used in randomness extractors [15,28]).

(CROOKED) LEFTOVER HASH LEMMA FOR CORRELATED DISTRIBUTIONS. We show that the leftover hash lemma (LHL) [26, Lemma 4.8], as well its generalized form [15, Lemma 2.4] and the “crooked” LHL [16]) extend in a natural way to

³ One could alternatively define robustness as one that remains hardcore on inputs of slightly lower entropy; however, in our proofs of robustness we would then need to go through an additional argument that distributions of lower entropy are induced by distributions of higher entropy.

⁴ In case of randomized leakage, the information-theoretic result of [15, Lemma 2.2(b)] gives better bounds.

“correlated” distributions. That is, suppose we have t random variables (sources) X_1, \dots, X_t , where each X_i individually has high min-entropy but may be fully determined by the outcome of some other X_j (though we assume $X_i \neq X_j$ for all $i \neq j$). We would like to apply a hash function H such that $H(X_1), \dots, H(X_t)$ is indistinguishable from t independent copies of the uniform distribution on the range of H (also over the choice of the key for H , which is made public). We show that this is the case assuming H is $2t$ -wise independent. (The standard LHL is thus $t = 1$; previously, Kiltz et al. [31] showed this for $t = 2$.) Naturally, this requires the output size of H to be about a $1/t$ fraction of its input size, so there is enough entropy to extract.

2 Preliminaries

We omit standard cryptographic definitions (see the full version for precise definitions [20]). The security parameter is denoted by k , and 1^k denotes the string of k ones. Vectors are denoted in boldface, for example \mathbf{x} . For convenience, we extend algorithmic notation to operate on each vector of inputs component-wise. For example, if A is an algorithm and \mathbf{x}, \mathbf{y} are vectors then $\mathbf{z} \stackrel{\$}{\leftarrow} A(\mathbf{x}, \mathbf{y})$ denotes that $\mathbf{z}[i] \stackrel{\$}{\leftarrow} A(\mathbf{x}[i], \mathbf{y}[i])$ for all $1 \leq i \leq |\mathbf{x}|$. We write P_X for the distribution of random variable X and $P_X(x)$ for the probability that X puts on value $x \in \mathcal{X}$, i.e., $P_X(x) = \Pr[X = x]$. Denote by $|X|$ the size of the support of X , i.e., $|X| = |\{x \text{ s.t. } P_X(x) > 0\}|$. We often identify X with P_X when there is no danger of confusion. For a function $f : \mathcal{X} \rightarrow \mathbb{R}$, we denote the expectation of f over X by $\mathbb{E} f(X) \stackrel{\text{def}}{=} \mathbb{E}_{x \in X} f(x) \stackrel{\text{def}}{=} \sum_{x \in \mathcal{X}} P_X(x) f(x)$.

We will use the notions of min-entropy and average min-entropy (defined in [15]). For vector-valued \mathbf{X} the min-entropy is the minimum of the components (see [3,5]). We use the standard notions of *collision probability* of X denoted $\text{Col}(X)$ and *statistical distance* of X and Y denoted $\Delta(X, Y)$. We denote the *computational distance* between two random variables X, Y with respect to a distinguisher D as $\delta^D(X, Y)$.

Dodis et al. [15, Lemma 2.2] characterized the effect of auxiliary information on average min-entropy, namely, $\tilde{H}_\infty(A|(B, C)) \geq \tilde{H}_\infty((A, B)|C) - |B| \geq \tilde{H}_\infty(A|C) - |B|$.

We will use extractors (defined in [33]) and average-case extractors (defined in [15, Section 2.5]) and denote both by **ext**.

For a (*probabilistic*) *public-key encryption scheme*, which is a triple of algorithms $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ defined in the usual way, we will use the standard notion of *IND-CPA security* as defined in [24].

We use the standard definition of a *lossy trapdoor function (LTDF) generator* (defined in [35]) which we denote as a pair $\text{LTDF} = (\mathcal{F}, \mathcal{F}')$ of algorithms.

COMPUTATIONAL ENTROPY. We use the standard notion of HILL entropy as defined in [26]. Additionally, we use a notion known as “metric-star” entropy (this notion was used in [18,21]):

Definition 1. A distribution X has **Metric*** entropy at least k , denoted $H_{\epsilon,s}^{\text{Metric}^*}(X) \geq k$ if for all deterministic distinguishers D of size at most s , with outputs in $[0, 1]$, there exists a distribution Y with $H_\infty(Y) \geq k$ and $\delta^D(X, Y) \leq \epsilon$.

Equivalence (with a loss in quality) between **Metric*** and HILL entropy was shown in [2, Theorem 5.2]. Extractors can be applied to distributions with computational entropy to obtain pseudorandom outputs. This is well-known for HILL entropy, but the only known way to extract from **Metric*** entropy is first to convert **Metric*** to HILL entropy by using [2, Theorem 5.2]. Conditional entropy has been extended to the computational case (for both HILL [28] and Metric entropy [21]). Conditional **Metric*** can be defined similarly, by making the distinguisher deterministic with outputs in $[0, 1]$. The **Metric*** to HILL conversion can be extended to the computational case as shown in [11, Lemma 18], [21, Theorem 2.7]. Average-case extractors can be used on distributions with conditional **Metric*** entropy by first using applying [21, Theorem 2.7].

2.1 Deterministic Encryption

An encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is *deterministic* if \mathcal{E} is deterministic.

SEMANTIC SECURITY OF DE. We recall the semantic-security style PRIV notion for DE from [3]. (More specifically, it is a “comparison-based” semantic-security style notion; this was shown equivalent to a “simulation-based” formulation in [5].) To encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$, an adversary $A = (A_0, A_1, A_2)$, and $k \in \mathbb{N}$ we associate the left-most and middle experiments in Figure 1. We require that there are functions $v = v(k), \ell = \ell(k)$ such that (1) $|\mathbf{x}| = v$, (2) $|\mathbf{x}[i]| = \ell$ for all $1 \leq i \leq v$, and (3) the $\mathbf{x}[i]$ are all distinct with probability 1 over $(\mathbf{x}, t) \stackrel{\$}{\leftarrow} A_1(\text{state})$ for any state output by A_0 . (Since in this work we only consider the definition relative to deterministic Π requirement (3) is without loss of generality.) In particular we say A *outputs vectors of size v* for v as above. Define the *PRIV advantage* of A against Π as

$$\text{Adv}_{\Pi,A}^{\text{priv}}(k) = \Pr \left[\text{Exp}_{\Pi,A}^{\text{priv}-1}(k) \Rightarrow 1 \right] - \Pr \left[\text{Exp}_{\Pi,A}^{\text{priv}-0}(k) \Rightarrow 1 \right].$$

Let \mathbb{M} be a class of distributions on message vectors. Define $\mathbb{A}_{\mathbb{M}}$ to be the class of adversaries $\{A = (A_0, A_1, A_2)\}$ such that for each $A \in \mathbb{A}_{\mathbb{M}}$ there is a $\mathbf{M} \in \mathbb{M}$ for which \mathbf{x} has distribution \mathbf{M} over $(\mathbf{x}, t) \stackrel{\$}{\leftarrow} A_1(\text{state})$ for any state output by A_0 . We say that Π is *PRIV secure for \mathbb{M}* if $\text{Adv}_{\Pi,A}^{\text{priv}}(\cdot)$ is negligible for any PPT $A \in \mathbb{A}_{\mathbb{M}}$. Note that (allowing non-uniform adversaries as usual) we can without loss of generality consider only those A with “empty” A_0 , since A_1 can always be hardwired with the “best” state. However, following [5] we explicitly allow state because it greatly facilitates some proofs.

INDISTINGUISHABILITY OF DE. Next we recall the indistinguishability-based formulation of security for DE [5,7]. To an encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$, an adversary $D = (D_1, D_2)$, and $k \in \mathbb{N}$ we associate the right-most experiment in Figure 1. We make the analogous requirements on D_1 as on A_1 in the PRIV definition. Define the *IND advantage* of D against Π as $\text{Adv}_{\Pi,D}^{\text{ind}}(k) =$

2 · Pr [$\mathbf{Exp}_{\Pi, D}^{\text{ind}}(k) \Rightarrow 1$] - 1. Let \mathbb{M}^* be a class of *pairs* of distributions on message vectors. Define $\mathbb{D}_{\mathbb{M}^*}$ to be the class of adversaries $\{D = (D_1, D_2)\}$ such that for each $D \in \mathbb{D}_{\mathbb{M}^*}$, there is a pair of distributions $(\mathbf{M}_0, \mathbf{M}_1) \in \mathbb{M}^*$ such that for each $b \in \{0, 1\}$ the distribution of $\mathbf{x} \stackrel{\$}{\leftarrow} D_1(b)$ is \mathbf{M}_b . We say that Π is *IND secure for \mathbb{M}^** if $\mathbf{Adv}_{\Pi, D}^{\text{ind}}(\cdot)$ is negligible for any PPT $D \in \mathbb{D}_{\mathbb{M}^*}$.

$\mathbf{Expr} \mathbf{Exp}_{\Pi, A}^{\text{priv-1}}(k):$ $(pk, sk) \stackrel{\$}{\leftarrow} \mathcal{K}(1^k)$ $state \stackrel{\$}{\leftarrow} A_0(1^k)$ $(\mathbf{x}_1, t_1) \stackrel{\$}{\leftarrow} A_1(state)$ $\mathbf{c} \stackrel{\$}{\leftarrow} \mathcal{E}(pk, \mathbf{x}_1)$ $g \stackrel{\$}{\leftarrow} A_2(pk, \mathbf{c}, state)$ If $g = t_1$ ret 1 else ret 0	$\mathbf{Expr} \mathbf{Exp}_{\Pi, A}^{\text{priv-0}}(k):$ $(pk, sk) \stackrel{\$}{\leftarrow} \mathcal{K}(1^k)$ $state \stackrel{\$}{\leftarrow} A_0(1^k)$ $(\mathbf{x}_1, t_1), (\mathbf{x}_0, t_0) \stackrel{\$}{\leftarrow} A_1(state)$ $\mathbf{c} \stackrel{\$}{\leftarrow} \mathcal{E}(pk, \mathbf{x}_0)$ $g \stackrel{\$}{\leftarrow} A_2(pk, \mathbf{c}, state)$ If $g = t_1$ ret 1 else ret 0	$\mathbf{Expr} \mathbf{Exp}_{\Pi, A}^{\text{ind}}(k):$ $(pk, sk) \stackrel{\$}{\leftarrow} \mathcal{K}(1^k)$ $b \stackrel{\$}{\leftarrow} \{0, 1\}; (\mathbf{x}, t) \stackrel{\$}{\leftarrow} D_1(b)$ $\mathbf{c} \stackrel{\$}{\leftarrow} \mathcal{E}(pk, \mathbf{x})$ $d \stackrel{\$}{\leftarrow} D_2(pk, \mathbf{c})$ If $b = d$ ret 1 else ret 0
---	--	---

Fig. 1. Security experiments for deterministic encryption

3 Our Tools

3.1 A Precise Definitional Equivalence for DE

While the PRIV definition is meaningful with respect a single message distribution \mathbf{M} , the IND definition must inherently talk of *pairs* of different message distributions. Thus, in proving an equivalence between the two notions, the best we can hope to show is that PRIV security for a message distribution \mathbf{M} is equivalent to IND security for some *class of pairs* of message distributions (depending on \mathbf{M}). However, prior works [5,7] did not provide such a statement. Instead, they showed that PRIV security on *all* distributions of a given entropy μ is equivalent to IND security on all pairs of distributions of slightly less entropy.

INDUCED DISTRIBUTIONS. To state our result we first give some definitions relating to a notion of “induced distributions.” Let X, X' be distributions (or random variables) on the same domain. For $\alpha \in \mathbb{N}$, we say that X' is an α -*induced distribution of X* if X' is a conditional distribution $X' = X \mid \mathbf{E}$ for an event \mathbf{E} such that $\Pr[\mathbf{E}] \geq 2^{-\alpha}$. We call \mathbf{E} the *corresponding event* to X' . We require that the pair (X, \mathbf{E}) is efficiently samplable. Define $X[\alpha]$ to be the class of all α -induced distributions of X . Furthermore, let X_0, X_1 be two α -induced distributions of X with corresponding events $\mathbf{E}_0, \mathbf{E}_1$ respectively. Define $X^*[\alpha] = \{(X_0, X_1)\}$ to be the class of all pairs (X_0, X_1) for which there is a pair (X'_0, X'_1) of α -induced distributions of X such that X_0 (resp. X_1) is statistically close to X'_0 (resp. X'_1).⁵

THE EQUIVALENCE. We are now ready to state our result. The following theorem captures the “useful” direction that IND implies PRIV:

⁵ We need to allow a negligible statistical distance for technical reasons. Since we will be interested in indistinguishability of functions of these distributions this will not make any appreciable difference, and hence we mostly ignore this issue in the remainder of the paper.

Theorem 1. *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. For any distribution \mathbf{M} on message vectors, PRIV security of Π with respect to \mathbf{M} is implied by IND security of Π with respect to $\mathbf{M}^*[2]$. In particular, let $A \in \mathbb{A}_{\mathbf{M}}$ be a PRIV adversary against Π . Then there is a IND adversary $D \in \mathbb{D}_{\mathbf{M}^*[2]}$ such that for all $k \in \mathbb{N}$*

$$\text{Adv}_{\Pi, A}^{\text{priv}}(k) \leq 162 \cdot \text{Adv}_{\Pi, D}^{\text{ind}}(k) + \left(\frac{3}{4}\right)^k .$$

Furthermore, the running-time of D is the time for at most that for k executions of A (but 4 in expectation).

The theorem essentially follows from the techniques of [5]; details are given in [20]. Thus, our contribution here is not in providing any new technical tools used in proving this result but rather in extracting it from the techniques of [5]. In particular, our more precise statement allows us to use results about entropy of conditional distributions, which we explain next. Looking ahead, it also simplifies proofs for schemes based on one-wayness, because it is easy to argue that one-wayness is preserved on slightly induced distributions (the alternative would require an argument that distributions of lower entropy are induced by distributions of higher entropy).

To establish a definitional *equivalence*; that is, also show that PRIV implies IND, we need to further restrict the latter to pairs (that are statistically close to pairs) of *complementary* 2-induced distributions of \mathbf{M} (which we did not do above for conceptual simplicity), where we call X_0, X_1 *complementary* if $E_1 = \overline{E_0}$. We stress that this further restriction is not needed for the “useful” implication above and for our security proofs.

3.2 Measuring Computational Entropy of Induced Distributions

We study how conditioning a distribution reduces its computational entropy. This result is used later in the work to show that randomness extractors can convert a hardcore function into a robust one; it also applicable to leakage-resilient cryptography. This result is simplest to understand when stated in terms of **Metric*** computational entropy (defined in [18]) It is easy to see that conditioning on an event E with probability P_E reduces (information-theoretic) min-entropy by at most $\log P_E$. We show that the same holds for the computational notion of **Metric*** entropy if one considers reduction in both quantity and quality:

Lemma 1. *Let X, Y be discrete random variables. Then*

$$H_{\epsilon/P_Y(y), s'}^{\text{Metric}^*}(X|Y = y) \geq H_{\epsilon, s}^{\text{Metric}^*}(X) - \log 1/P_Y(y) \text{ where } s' \approx s.$$

The use of **Metric*** entropy and an improved proof allow for a simpler and tighter formulation than results of [18, Lemma 3] and [36, Theorem 1.3] (see the full version for a comparison [20]). The proof is similar to [36] and can be found in the full version [20].

If we now consider averaging over all values of Y , we obtain the following simple formulation that expresses how much average entropy is left in X from the point of view of someone who knows Y . (This scenario naturally occurs in leakage-resilient cryptography, as exemplified in [18]).

Theorem 2. *Let X, Y be discrete random variables. Then*

$$H_{\epsilon|Y|,s'}^{\text{Metric}^*}(X|Y) \geq H_{\epsilon,s}^{\text{Metric}^*}(X) - \log |Y|, \text{ where } s' \approx s.$$

This statement is similar to the statement in the information-theoretic case (where the reduction is only in quantity) from [15, Lemma 2.2]. In the full version [20], we compare the theorem to [11, Lemma 16] and [22, Lemma 3.1].

To apply a randomness extractor, we must convert conditional **Metric*** to conditional **HILL** entropy using [21, Theorem 2.7], this conversion loses some quality. Thus, the conversion should be applied only when necessary (for instance, repeated conditioning is best measured in **Metric*** entropy, and then converted to **HILL** entropy once at the end). Here we provide a “**HILL**-to-**HILL**” formulation of Lemma 1.

Corollary 1. *Let X be a discrete random variable over χ and let Y be a discrete random variable. Then,*

$$H_{\epsilon',s'}^{\text{HILL}}(X|Y = y) \geq H_{\epsilon,s}^{\text{HILL}}(X) - \log 1/P_Y(y)$$

where $\epsilon' = \epsilon/P_Y(y) + \sqrt[3]{\frac{\log |\chi|}{s}}$, and $s' = \Omega(\sqrt[3]{s/\log |\chi|})$.

The Corollary follows by combining Lemma 1, [2, Theorem 5.2], and setting $\epsilon_{\text{HILL}} = \sqrt[3]{\log |\chi|/s}$ (see the full version for justification of parameters [20]).

3.3 A (Crooked) Leftover Hash Lemma for Correlated Distributions

The following generalization of the (Crooked) LHL to correlated input distributions will be very useful to us when considering bounded multi-message security in Section 6. Since our generalization of the classical LHL is a special case of our generalization of the Crooked LHL, we just state the latter here.

Lemma 2. (CLHL for Correlated Sources) *Let $\mathcal{H}: \mathcal{K} \times D \rightarrow R$ be a $2t$ -wise δ -dependent function for $t > 0$ with range R , and let $f: R \rightarrow S$ be a function. Let $\mathbf{X} = (X_1, \dots, X_t)$ where the X_i are random variables over D such that $H_\infty(X_i) \geq \mu$ for all $1 \leq i \leq n$ and moreover $\Pr[X_i = X_j] = 0$ for all $1 \leq i \neq j \leq t$. Then*

$$\Delta((K, f(\mathcal{H}(K, \mathbf{X}))), (K, f(\mathbf{U}))) \leq \frac{1}{2} \sqrt{|S|^t (t^2 2^{-\mu} + 3\delta)}$$

where $K \stackrel{\$}{\leftarrow} \mathcal{K}$ and $\mathbf{U} = (U_1, \dots, U_t)$ where the U_i are all uniform and independent over R (recall that functions operate on vectors \mathbf{X} and \mathbf{U} component-wise).

One can further extend Lemma 2 to the case of average conditional min-entropy using the techniques of [15]. Note that the lemma implies the corresponding generalization of the classical LHL by taking \mathcal{H} to have range S and f to be the identity function. The proof of the lemma, which extends the proof of the Crooked LHL in [7], is given in the full version [20].

4 Encrypt-with-Hardcore Scheme from Robust HCFs

We define a new notion of *robust* HCFs. Intuitively, robust HCFs are those that remain hardcore when the input is conditioned on any event that occurs with good probability.

Definition 2. *Let \mathcal{F} be a TDF generator and let hc be a HCF such that hc is hardcore for \mathcal{F} with respect to a distribution \mathbf{X} on input vectors. For $\alpha = \alpha(k)$, we say hc is α -robust for \mathcal{F} on \mathbf{X} if hc is also hardcore for \mathcal{F} with respect to the class $\mathbf{X}[\alpha]$ of α -induced distributions of \mathbf{X} .*

DISCUSSION. Robustness is interesting even for the classical definition of hardcore bits, where hc is boolean and a single uniform input x is generated in the security experiment. Here robustness means that hc remains hardcore even when x is conditioned on an event that occurs with good probability. It is clear that not every hardcore bit in the classical sense is robust — note, for example, that while every bit of the input to RSA is well-known to be hardcore assuming RSA is one-way [1], they are not even 1-robust since we may condition on a particular bit of the input being a fixed value.

THE SCHEME. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a probabilistic encryption scheme, \mathcal{F} be a TDF generator, and hc_f be a HCF. Assume that hc outputs binary strings of the same length as the random string r needed by \mathcal{E} . Define the associated “Encrypt-with-Hardcore” deterministic encryption scheme $\text{EwHCore}[\Pi, \mathcal{F}, \text{hc}] = (\mathcal{DK}, \mathcal{DE}, \mathcal{DD})$ with plaintext-space $\text{PtSp} = \{0, 1\}^k$ via

$$\begin{array}{l}
 \mathbf{Alg} \mathcal{DK}(1^k): \\
 (pk, sk) \xleftarrow{\$} \mathcal{K}(1^k) \\
 (f, f^{-1}) \xleftarrow{\$} \mathcal{F}(1^k) \\
 \text{Return } ((pk, f), (sk, f^{-1}))
 \end{array}
 \left|
 \begin{array}{l}
 \mathbf{Alg} \mathcal{DE}((pk, f), x): \\
 r \leftarrow \text{hc}_f(x) \\
 c \leftarrow \mathcal{E}(pk, f(x); r) \\
 \text{Return } c
 \end{array}
 \right|
 \begin{array}{l}
 \mathbf{Alg} \mathcal{DD}((sk, f^{-1}), c): \\
 y \leftarrow \mathcal{D}(sk, c) \\
 x \leftarrow f^{-1}(y) \\
 \text{Return } x
 \end{array}$$

SECURITY ANALYSIS. To gain some intuition, suppose hc is hardcore for \mathcal{F} on some distribution \mathbf{X} on input vectors. One might think that PRIV security of $\text{EwHCore} = \text{EwHCore}[\Pi, \mathcal{F}, \text{hc}]$ on \mathbf{X} then follows by IND-CPA security of Π . However, this is not true. For example, hc may be a “natural” hardcore function (i.e., that outputs some bits of the input), and \mathcal{E} may output some of its coins in the clear. This is how our notion of robustness comes into play, giving us the following theorem (for a proof and further discussion, see [20]):

Theorem 3. *Suppose Π is IND-CPA secure, hc is 2-robust for \mathcal{F} on a distribution \mathbf{M} on input vectors. Then $\text{EwHCore}[\Pi, \mathcal{F}, \text{hc}]$ is PRIV-secure on \mathbf{M} .*

5 Single-Message Instantiations of EwHCore

5.1 Getting Robust Hardcore Functions

AUGMENTED TRAPDOOR FUNCTIONS. In order to describe the conversion procedure, it is useful to introduce the notion of an “augmented” version of a TDF, which augments the description of the TDF with keying material for a HCF. More formally, let \mathcal{F} be a trapdoor function generator and let H be a keyed function with keyspace \mathcal{K} . Define the H -augmented version of \mathcal{F} , denoted $\mathcal{F}[H]$, that on input 1^k returns $(f, K), (f^{-1}, K)$ where $(f, f^{-1}) \xleftarrow{\$} \mathcal{F}(1^k)$ and $K \xleftarrow{\$} \mathcal{K}$; evaluation is defined for $x \in \{0, 1\}^k$ as $f(x)$ (i.e., evaluation just ignores K) and inversion is defined analogously.

MAKING ANY LARGE HARDCORE FUNCTION ROBUST. We show that by applying a randomness extractor in a natural way, one can convert *any* large hardcore function in the standard sense to one that is robust (with some loss in parameters). However, while the conversion procedure is natural, proving that it works turns out to be non-trivial.

Let \mathcal{F} be a TDF generator, and let $\text{hc} : \{0, 1\}^k \rightarrow \{0, 1\}^\ell$ be an HCF for \mathcal{F} on an input distribution X such that $H_\infty(X) \geq \mu$. Let $\text{ext} : \{0, 1\}^\ell \times \{0, 1\}^d \rightarrow \{0, 1\}^m \times \{0, 1\}^d$ be a strong average-case $(\ell - \alpha, \epsilon_{\text{ext}})$ -extractor for $\alpha \in \mathbb{N}$. (Here we view ext as a keyed function with the *second* argument as the key.) Define a new “extractor-augmented” HCF $\text{hc}[\text{ext}]$ for $\mathcal{F}[\text{ext}]$ such that $\text{hc}[\text{ext}]_s(x) = \text{ext}(\text{hc}(x), s)$ for all $x \in \{0, 1\}^k$ and $s \in \{0, 1\}^d$. The following characterizes the α -robustness of $\text{hc}[\text{ext}]$.

Lemma 3. *Fix $X' \in X[\alpha]$, and suppose there is a distinguisher D' against $\text{hc}[\text{ext}]$ on X' . Then there is a distinguisher D against hc on X such that for all $k \in \mathbb{N}$*

$$\text{Adv}_{\mathcal{F}, X', \text{hc}[\text{ext}], D'}^{\text{hcf}}(k) \leq O\left(\sqrt[3]{\text{Adv}_{\mathcal{F}, X, \text{hc}, D}^{\text{hcf}}(k)} + 2^\alpha \cdot \text{Adv}_{\mathcal{F}, X, \text{hc}, D}^{\text{hcf}}(k)\right) + \epsilon_{\text{ext}}.$$

Furthermore, the running-time of D is $O((t_{D'}(k+\ell))^3)$, where $t_{D'}$ is the running-time of D' .

Note that when $\alpha = \log(k)$ the security loss in the reduction is polynomial (in our application we just need $\alpha = 2$). The proof, which appears in the full version [20], relies crucially on Corollary 1.

The above conversion procedure notwithstanding, we give specific examples of hardcore functions that are already robust.

ROBUST GOLDREICH-LEVIN BITS FOR ANY TDF. In [20] we show that the Goldreich-Levin [23] (GL) hardcore function is robust. Specifically, if the function that extracts i -many independent GL bits is hardcore for \mathcal{F} , then it is also $O(\log k)$ -robust for \mathcal{F} .

ROBUST BITS FOR ANY LTDF. Peikert and Waters [35] showed that LTDFs admit a simple, large hardcore function, namely a pairwise-independent hash

function (the same argument applies also to universal hash functions or, more generally, randomness extractors). By using average conditional min-entropy, in [20] we show that this hardcore function is $O(\log k)$ robust.

5.2 Putting It Together

Equipped with the above results, we describe instantiations of the Encrypt-with-Hardcore scheme that both explain prior constructions and produce novel ones.

USING AN ITERATED TRAPDOOR PERMUTATION. The prior trapdoor permutation based DE scheme of Bellare et al. [5] readily provides an instantiation of EwHCore by using an iterated trapdoor permutation as the TDF. Let \mathcal{F} be a TDP and hc be a hardcore bit for \mathcal{F} . For $i \in \mathbb{N}$ denote by \mathcal{F}^i the TDP that iterates \mathcal{F} i -many times. Define the Blum-Micali-Yao (BM) [6,41] hardcore function for \mathcal{F}^i via $\text{BM}^i[\text{hc}](f, x) = \text{hc}(x) \parallel \text{hc}(f(x)) \parallel \dots \parallel \text{hc}(f^{i-1}(x))$. Bellare et al. [5] used the specific choice of $\text{hc} = \mathcal{GL}$ (the Goldreich-Levin bit) in their scheme, which is explained by the fact that the GL bit is robust, and one can show that BM iteration expands one robust hardcore bit to many (on a non-uniform distribution, the bit should be hardcore on all “permutation distributions” of the former).

However, due to our augmentation procedure to make any large hardcore function robust, we are no longer bound to any specific choice of hc . For example, we may choose hc to be a natural hardcore bit. In fact, it may often be the case that \mathcal{F} has many simultaneously hardcore natural bits, and therefore our construction will require fewer iterations of the TDP than the construction of [5].

USING A LOSSY TDF. Using the fact that extractors are robust hardcore functions for LTDFs, we get an instantiation of the Encrypt-with-Hardcore scheme from LTDFs that is an alternative to the prior scheme of Boldyreva et al. [7] and the concurrent work of Wee [40]. Our scheme requires an LTDF with residual leakage $s \leq H_\infty(X) - 2 \log(1/\epsilon) - r$, where r is the number of random bits needed in \mathcal{E} (or the length of a seed to a pseudorandom generator that can be used to obtain those bits).

USING 2-CORRELATED PRODUCT TDFs. Hemenway et al. [27] show a construction of DE from a *decisional 2-correlated product TDF*, namely where \mathcal{F} has the property that $f_1(x), f_2(x)$ is indistinguishable from $f_1(x_1), f_2(x_2)$ where x_1, x_2 are sampled independently (in both cases for two independent public instances f_1, f_2 of \mathcal{F}). They show such a trapdoor function is a secure DE scheme for uniform messages. To obtain an instantiation of EwHCore under the same assumption, we can use \mathcal{F} as the TDF, and an independent instance of the TDF as hc . When a randomness extractor is applied to the latter, robustness follows from Lemma 3, taking into account that the lemma holds even if the output of the hardcore function is not uniform, as long as it has high HILL entropy.

USING ANY TDF WITH A LARGE HCF. Our most novel instantiations in the single-message case come from considering TDFs that have a sufficiently large HCF but are not necessarily lossy or an iterated TDP. Let us first consider instantiations on the uniform message distribution Freeman et al. [19] shown that

the Niederreiter TDF [32] has linearly many (simultaneous) hardcore bits under the “Syndrome Decoding Assumption (SDA)” and “Indistinguishability Assumption (IA)” (as defined in [19, Section 7.2]). Furthermore, the RSA [37] and Paillier [34] TDPs have linearly many hardcore bits under certain computational assumptions, namely the “Small Solutions RSA (SS-RSA) Assumption” [39] and the “Bounded Computational Composite Residuosity (BCCR) Assumption” [9] respectively. Because these hardcore functions are sufficiently long, they can be made robust via Lemma 3 and give us a linear number of *robust* hardcore bits—enough to use as randomness for \mathcal{E} (expanded by a pseudorandom generator if necessary). Thus, by Theorem 3, we obtain:

Corollary 2. *Under SDA+IA for the Niederreiter TDF, DE for the uniform message distribution exists. Similarly, under SS-RSA the RSA TDP or BCCR for the Paillier TDP respectively, DE for the uniform message distribution exists.*

In particular, the first statement provides the first DE scheme without random oracles based on the hardness of syndrome decoding. (A scheme in the random oracle model follows from [3].) Moreover, the schemes provided by the second statement are nearly as efficient as the ones obtained from lossy TDFs (since they do not use iteration), and the latter typically requires decisional assumptions (in contrast to the computational assumptions used here).

If we do not wish to rely on specific assumptions, we can also get DE from strong but general assumptions, such as sub-exponential hardness. We can also obtain DE for nonuniform message distributions (the strength of the assumption needed will depend on how far the entropy of the message space is from the maximum). See [20] for details.

6 Bounded Multi-message Security and its Instantiations

6.1 The New Notion and Variations

THE NEW NOTION. Our notion of q -bounded multi-message security (or just q -bounded security) for DE is quite natural, and can be viewed as analogous to other forms of “bounded” security (see e.g. [12]). In a nutshell, it asks for security on up to q arbitrarily correlated but high-entropy messages (where we allow the public-key size to depend on q). Fix an encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. For $q = q(k)$ and $\mu = \mu(k)$, let $\mathbb{M}^{q,\mu}$ be the class of distributions on message vectors $M^{\mu,q} = (M_1^{\mu,q}, \dots, M_q^{\mu,q})$ where $H_\infty(M_i^{\mu,q}) \geq \mu$ and for all $1 \leq i \leq q$ and $M_{1,q}^\mu, \dots, M_{q,q}^\mu$ are distinct with probability 1. We say that Π is q -bounded multi-message PRIV (resp. IND) secure for μ -sources if it is PRIV (resp. IND) secure for $\mathbb{M}^{q,\mu}$. By Theorem 1, PRIV on $\mathbb{M}^{q,\mu}$ is equivalent to IND on $\mathbb{M}^{q,\mu-2}$.

UNBOUNDED MULTI-MESSAGE SECURITY FOR q -BLOCK SOURCES. We also consider unbounded multi-message security for what we call a q -block source, a generalization of a block-source [10] where every q -th message introduces some “fresh” entropy. Fix an encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. For $q = q(k)$, $n =$

$n(k)$, and $\mu = \mu(k)$, let $\mathbb{M}^{q,n,\mu}$ be the class of distributions on message vectors $M^{q,n,\mu} = (M_1^{q,n,\mu}, \dots, M_{qn}^{q,n,\mu})$ such that $H_\infty(X_{qi+j} \mid X_1 = x_1, \dots, X_{qi-1} = x_{qi-1}) \geq \mu$ for all $1 \leq i \leq n$, all $0 \leq j \leq q-1$, and all outcomes x_1, \dots, x_{qi-1} of X_1, \dots, X_{qi-1} . We say that Π is q -bounded multi-message PRIV (resp. IND) secure for (μ, n) -block-sources if Π is PRIV (resp. IND) secure on $\mathbb{M}^{q,n,\mu}$. Using a similar argument to [7, Theorem 4.2], one can show equivalence of PRIV on $\mathbb{M}^{q,n,\mu}$ to IND on $\mathbb{M}^{q,n,\mu}$.

6.2 Our Basic Scheme

We cannot trivially achieve q -bounded security by running, say, q copies of a scheme secure for one message in parallel (and encrypting the i -th message under the i -th public key), since this approach would lead to a stateful scheme. The main technical tool we use to achieve the notion is Lemma 2. Combined with [15, Lemma 2.2], this tells us that a $2q$ -wise independent hash function is robust on correlated input distributions of sufficient min-entropy:

Proposition 1. *For any q , let $\text{LTDF} = (\mathcal{F}, \mathcal{F}')$ be an LTDF generator with input length n and residual leakage s , and let $\mathcal{H}: \mathcal{K} \times D \rightarrow R$ where $r = \log |R|$ be a $2q$ -wise independent hash function. Then \mathcal{H} is a 2-robust hardcore function for \mathcal{F} on any input distribution $X = (X_1, \dots, X_q)$ such that $H_\infty(X) \geq q(s + r) + 2 \log q + 2 \log(1/\epsilon) - 2$ for negligible ϵ .*

By Theorem 3, we obtain a q -bounded DE scheme based on lossy trapdoor functions that lose a $1 - O(1/q)$ fraction of its input. Specifically, we can use the DDH-based construction of Peikert and Waters [35], the Paillier-based one of [7,19], or the one from d -linear of [19] for any polynomial q .

6.3 Our Optimized Scheme

We show that by extending some ideas of [7], we obtain a more efficient DE scheme meeting q -bounded security that achieves better parameters.

INTUITION AND PRELIMINARIES. Intuitively, for the optimized scheme we modify the scheme of [7] to first pre-process an input message using a $2q$ -wise independent permutation (instead of pairwise as in [7]). However, there are two issues to deal with here. First, for $q > 1$ such a permutation is not known to exist (in an explicit and efficiently computable sense). Second, Lemma 2 applies to t -wise independent functions rather than permutations.

To solve the first problem, we turn to $2q$ -wise “ δ -dependent” permutations (as constructed in e.g. [29]). Namely, say that a permutation $H: \mathcal{K} \times D \rightarrow D$ is t -wise δ -dependent if for all distinct $x_1, \dots, x_t \in D$

$$\Delta((H(K, x_1), \dots, H(K, x_t)), (P_1, \dots, P_t)) \leq \delta,$$

where $K \stackrel{\$}{\leftarrow} \mathcal{K}$ and P_1, \dots, P_t are defined iteratively by taking P_1 to be uniform on D and, for all $2 \leq i \leq t$, taking P_i to be uniform on $D \setminus \{p_1, \dots, p_{i-1}\}$ where p_1, \dots, p_{i-1} are the outcomes of P_1, \dots, P_{i-1} respectively.

To solve the second problem, we show that a t -wise δ -dependent permutation is a t -wise δ' -dependent function where δ' is a bit bigger than δ (see [20] for details, where we also restate Lemma 2 in terms of δ -dependent permutations).

THE CONSTRUCTION. We now detail our construction. Let $\text{LTDF} = (\mathcal{F}, \mathcal{F}')$ be an LTDF and let $\mathcal{P}: \mathcal{K} \times \{0, 1\}^k \rightarrow \{0, 1\}^k$ be an efficiently invertible family of permutations on k bits. Define the associated deterministic encryption scheme $\Pi[\text{LTDF}, \mathcal{P}] = (\mathcal{DK}, \mathcal{DE}, \mathcal{DD})$ with plaintext-space $\text{PtSp} = \{0, 1\}^k$ via

$$\begin{array}{l} \mathbf{Alg} \mathcal{DK}(1^k): \\ (f, f^{-1}) \xleftarrow{\$} \mathcal{F}(1^k); K \xleftarrow{\$} \mathcal{K} \\ \text{Return } ((f, K), (f^{-1}, K)) \end{array} \quad \left| \begin{array}{l} \mathbf{Alg} \mathcal{DE}((f, K), x): \\ c \leftarrow f(\mathcal{P}(K, x)) \\ \text{Return } c \end{array} \right. \quad \left| \begin{array}{l} \mathbf{Alg} \mathcal{DD}((sk, f^{-1}), c): \\ x \leftarrow f^{-1}(\mathcal{P}^{-1}(K, c)) \\ \text{Return } x \end{array} \right.$$

We have the following result:

Theorem 4. *Suppose LTDF is a lossy trapdoor function on $\{0, 1\}^n$ with residual leakage s , and let $q, \epsilon > 0$. Suppose \mathcal{P} is a $2q$ -wise δ -dependent permutation on $\{0, 1\}^n$ for $\delta = t^2/2^n$. Then for any q -message IND adversary $B \in \mathbb{D}_{\mathbb{M}^{q,\mu}}$ with min-entropy $\mu \geq qs + 2 \log q + \log(1/\epsilon) + 5$, there is an LTDF distinguisher D such that for all $k \in \mathbb{N}$*

$$\text{Adv}_{\Pi[\text{LTDF}, \mathcal{P}], B}^{\text{ind}}(k) \leq \text{Adv}_{\text{LTDF}, D}^{\text{ltdf}}(k) + \epsilon.$$

Furthermore, the running-time of D is the time to run B .

An efficiently invertible $2q$ -wise δ -dependent permutation on $\{0, 1\}^n$ for $\delta = t^2/2^n$ can be obtained from [29] using key length $nt + \log(1/\delta) = n(t + 1) - 2t$. Comparing the above to Proposition 1, we see that we have dropped the r in the entropy bound (indeed, there is no hardcore function here).

Acknowledgements. The authors are grateful to Mihir Bellare, Alexandra Boldyreva, Kai-Min Chung, Sebastian Faust, Marc Fischlin, Serge Fehr, Péter Gács, Bhavana Kanukurthi, Fenghao Liu, Payman Mohassel, Krzysztof Pietrzak, Adam Smith, Ramarathnam Venkatesan, and Hoeteck Wee for helpful discussions, improvements to our analysis, and useful references. The work was supported, in part, by NSF awards 0546614, 0831281, 1012910, 1012798, CNS-0915361, CNS-0952692, NSF CAREER award 0545659, and NSF Cyber Trust award 0831184.

References

1. Alexi, W., Chor, B., Goldreich, O., Schnorr, C.P.: RSA and Rabin functions: Certain parts are as hard as the whole. *SIAM J. Comput.* 17(2) (1988)
2. Barak, B., Shaltiel, R., Wigderson, A.: Computational analogues of entropy. In: 11th International Conference on Random Structures and Algorithms, pp. 200–215 (2003)
3. Bellare, M., Boldyreva, A., O’Neill, A.: Deterministic and Efficiently Searchable Encryption. In: Menezes, A. (ed.) *CRYPTO 2007*. LNCS, vol. 4622, pp. 535–552. Springer, Heidelberg (2007)

4. Bellare, M., Brakerski, Z., Naor, M., Ristenpart, T., Segev, G., Shacham, H., Yilek, S.: Hedged Public-Key Encryption: How to Protect against Bad Randomness. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 232–249. Springer, Heidelberg (2009)
5. Bellare, M., Fischlin, M., O'Neill, A., Ristenpart, T.: Deterministic Encryption: Definitional Equivalences and Constructions without Random Oracles. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 360–378. Springer, Heidelberg (2008)
6. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.* 13(4), 850–864 (1984)
7. Boldyreva, A., Fehr, S., O'Neill, A.: On Notions of Security for Deterministic Encryption, and Efficient Constructions without Random Oracles. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 335–359. Springer, Heidelberg (2008)
8. Brakerski, Z., Segev, G.: Better Security for Deterministic Public-Key Encryption: The Auxiliary-Input Setting. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 543–560. Springer, Heidelberg (2011)
9. Catalano, D., Gennaro, R., Howgrave-Graham, N.: Paillier's trapdoor function hides up to $O(n)$ bits. *J. Cryptology* (2002)
10. Chor, B., Goldreich, O.: Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.* 17(2) (1988)
11. Chung, K.-M., Kalai, Y.T., Liu, F.-H., Raz, R.: Memory Delegation. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 151–168. Springer, Heidelberg (2011)
12. Cramer, R., Hanaoka, G., Hofheinz, D., Imai, H., Kiltz, E., Pass, R., Shelat, A., Vaikuntanathan, V.: Bounded CCA2-Secure Encryption. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 502–518. Springer, Heidelberg (2007)
13. Dent, A.W., Fischlin, M., Manulis, M., Stam, M., Schröder, D.: Confidential Signatures and Deterministic Signcryption. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 462–479. Springer, Heidelberg (2010)
14. Desrosiers, S.P.: Entropic security in quantum cryptography. *Quantum Information Processing* 8(4), 331–345 (2009)
15. Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.* 38(1), 97–139 (2008)
16. Dodis, Y., Smith, A.: Correcting errors without leaking partial information. In: STOC, pp. 654–663 (2005)
17. Dodis, Y., Smith, A.: Entropic Security and the Encryption of High Entropy Messages. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 556–577. Springer, Heidelberg (2005)
18. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: FOCS, pp. 293–302 (2008)
19. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More Constructions of Lossy and Correlation-Secure Trapdoor Functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 279–295. Springer, Heidelberg (2010)
20. Fuller, B., O'Neill, A., Reyzin, L.: A unified approach to deterministic encryption: New constructions and a connection to computational entropy. *Cryptology ePrint Archive* (2012)
21. Fuller, B., Reyzin, L.: Computational entropy and information leakage. Tech. rep., Boston University (2011), <http://cs-people.bu.edu/bfuller/metricEntropy.pdf>
22. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: STOC, pp. 99–108. ACM, New York (2011)

23. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: STOC, pp. 25–32 (1989)
24. Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* 28(2), 270–299 (1984)
25. Goyal, V., O’Neill, A., Rao, V.: Correlated-Input Secure Hash Functions. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 182–200. Springer, Heidelberg (2011)
26. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM J. Comput.* 28(4), 1364–1396 (1999)
27. Hemenway, B., Lu, S., Ostrovsky, R.: Correlated product security from any one-way function and the new notion of decisional correlated product security. *Cryptology ePrint Archive*, Report 2010/100 (2010), <http://eprint.iacr.org/>
28. Hsiao, C.-Y., Lu, C.-J., Reyzin, L.: Conditional Computational Entropy, or Toward Separating Pseudoentropy from Compressibility. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 169–186. Springer, Heidelberg (2007)
29. Kaplan, E., Naor, M., Reingold, O.: Derandomized constructions of k -wise (almost) independent permutations. *Algorithmica* 55(1), 113–133 (2009)
30. Kiltz, E., O’Neill, A., Smith, A.: Instantiability of RSA-OAEP under chosen-plaintext attack. *IACR Cryptology ePrint Archive* 2011, 559 (2011)
31. Kiltz, E., Pietrzak, K., Stam, M., Yung, M.: A New Randomness Extraction Paradigm for Hybrid Encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 590–609. Springer, Heidelberg (2009)
32. Niederreiter, H.: Knapsack-type cryptosystems and algebraic coding theory. *Problems of Control and Information Theory* 15, 367–391 (1986)
33. Nisan, N., Zuckerman, D.: Randomness is linear in space. *Journal of Computer and System Sciences*, 43–52 (1993)
34. Paillier, P.: Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999)
35. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: STOC, pp. 187–196 (2008)
36. Reingold, O., Trevisan, L., Tulsiani, M., Vadhan, S.: Dense subsets of pseudorandom sets. In: 2008 49th Annual IEEE Symposium on Foundations of Computer Science, pp. 76–85. IEEE (2008)
37. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21(2), 120–126 (1978)
38. Russell, A., Wang, H.: How to fool an unbounded adversary with a short key. *IEEE Transactions on Information Theory* 52(3), 1130–1140 (2006)
39. Steinfeld, R., Pieprzyk, J., Wang, H.: On the Provable Security of an Efficient RSA-Based Pseudorandom Generator. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 194–209. Springer, Heidelberg (2006)
40. Wee, H.: Dual projective hashing and its applications—lossy trapdoor functions and more. In: Eurocrypt (2012)
41. Yao, A.C.C.: Theory and applications of trapdoor functions (extended abstract). In: FOCS, pp. 80–91 (1982)