

Hedged Public-Key Encryption: How to Protect Against Bad Randomness

Mihir Bellare¹, Zvika Brakerski², Moni Naor², Thomas Ristenpart¹,
Gil Segev², Hovav Shacham¹, and Scott Yilek¹

¹ Dept. of Computer Science & Engineering, University of California at San Diego,
9500 Gilman Drive, La Jolla, CA 92093, USA.

{mihir,tristenp,hovav,syilek}@cs.ucsd.edu

² Dept. of Computer Science and Applied Mathematics,
Weizmann Institute of Science, Rehovot 76100, Israel

{zvika.brakerski,moni.naor,gil.segev}@weizmann.ac.il

Abstract. Public-key encryption schemes rely for their IND-CPA security on per-message fresh randomness. In practice, randomness may be of poor quality for a variety of reasons, leading to failure of the schemes. Expecting the systems to improve is unrealistic. What we show in this paper is that we can, instead, improve the cryptography to offset the lack of possible randomness. We provide public-key encryption schemes that achieve IND-CPA security when the randomness they use is of high quality, but, when the latter is not the case, rather than breaking completely, they achieve a weaker but still useful notion of security that we call IND-CDA. This hedged public-key encryption provides the best possible security guarantees in the face of bad randomness. We provide simple RO-based ways to make in-practice IND-CPA schemes hedge secure with minimal software changes. We also provide non-RO model schemes relying on lossy trapdoor functions (LTDFs) and techniques from deterministic encryption. They achieve adaptive security by establishing and exploiting the anonymity of LTDFs which we believe is of independent interest.

1 Introduction

Cryptography ubiquitously assumes that parties have access to sufficiently good randomness. In practice this assumption is often violated. This can happen because of faulty implementations, side-channel attacks, system resets or for a variety of other reasons. The resulting cryptographic failures can be spectacular [22, 24, 29, 2, 15]. What can we do about this? One answer is that system designers should build “better” systems, but this is clearly easier said than done. The reality is that random number generation is a complex and difficult task, and it is unrealistic to think that failures will never occur. We propose a different approach: designing schemes in such a way that poor randomness will have as little as possible impact on the security of the scheme in the following sense. With good randomness the scheme achieves whatever (strong) security notion

one is targeting, but when the same scheme is fed bad (even adversarially chosen) randomness, rather than breaking completely, it achieves some weaker but still useful notion of security that is the best possible under the circumstances. We call this “hedged” cryptography.

Previous work by Rogaway [32], Rogaway and Shrimpton [33], and Kamara and Katz [27] considers various forms of hedging for the symmetric encryption setting. In this paper, we initiate a study of hedged public-key encryption. We address two central foundational questions, namely to find appropriate definitions and to efficiently achieve them. Let us now look at all this in more detail.

THE PROBLEM. Achieving the standard IND-CPA notion of privacy [23] *requires* the encryption algorithm to be randomized. In addition to the public key and message, it takes as input a random string that needs to be freshly and independently created for *each and every* encryption.

Weak (meaning, low-entropy) randomness does not merely imply a loss of theoretical security. It can lead to catastrophic attacks. For example, weak-randomness based encryption is easily seen to allow recovery of the plaintext from the ciphertext for the quadratic residuosity scheme of [23] as well as the El Gamal encryption scheme [21]. Brown [15] presents such an attack on RSA-OAEP [10] with encryption exponent 3. Ouafi and Vaudenay [30] present such an attack on Rabin-SAEP [13]. We present an alternative attack in [7].

The above would be of little concern if we could guarantee good randomness. Unfortunately, this fails to be true in practice. Here, an “entropy-gathering” process is used to get a seed which is then stretched to get “random” bits for the application. The theory of cryptographically strong pseudorandom number generators [11] implies that the stretching can in principle be sound, and extractors further allow us to reduce the requirement on the seed from being uniformly distributed to having high min-entropy, but we still need a sufficiently good seed. (No amount of cryptography can create randomness out of nothing!) In practice, entropy might be gathered from timing-related operating system events or user keystrokes. As evidence that this process is error-prone, consider the recent randomness failure in Debian Linux, where a bug in the OpenSSL package led to insufficient entropy gathering and thence to practical attacks on the SSH [29] and SSL [2, 36] protocols. Other exploits include [25, 19].

THE NEW NOTION. The idea is to provide two tiers of security. First, when the “randomness” is really random, the scheme should meet the standard IND-CPA notion of security. Otherwise, rather than failing completely, it should gracefully achieve some weaker but as-good-as-possible notion of security. The first important question we then face is to pick and formally define this fallback notion.

Towards this, we begin by suggesting that the *message* being encrypted may also have entropy or uncertainty from the point of view of the adversary. (If not, what privacy is there to be preserved by encryption?) We propose to harvest this. In this regard, the first requirement that might come to mind is that encryption with weak (even adversarially-known) randomness should be as secure as deterministic encryption, meaning achieve an analog of the PRIV notion of [6]. But

achieving this would require that the message by itself have high min-entropy. We can do better. Our new target notion of security, that we call Indistinguishability under a Chosen Distribution Attack (IND-CDA), asks that security is guaranteed as long as the *joint* distribution of the message and randomness has sufficiently high min-entropy. In this way, we can exploit for security whatever entropy might be present in the randomness *or* the message, and in particular achieve security even if neither taken alone is random enough.

Notice that if the message and randomness together have low min-entropy, then we cannot hope to achieve security, because an adversary can recover the message with high probability by trial encryption with all message-randomness pairs that occur with a noticeable probability. In a nutshell, our new notion asks that this necessary condition is also sufficient, and in this way is requiring security that is as good as possible.

We denote by H-IND our notion of *hedged* security that is satisfied by encryption schemes that are secure both in the sense of IND-CPA and in the sense of IND-CDA.

ADAPTIVITY. Our IND-CDA definition generalizes the indistinguishability-style formalizations of PRIV-secure deterministic encryption [8, 12], which in turn extended entropic security [18]. But we consider a new dimension, namely, adaptivity. Our adversary is allowed to specify joint message-randomness distributions on to-be-encrypted challenges. The adversary is said to be adaptive if these queries depend on the replies to previous ones. Non-adaptive H-IND means IND-CPA plus non-adaptive IND-CDA and adaptive H-IND means IND-CPA plus adaptive IND-CDA.

Non-adaptive IND-CDA is a notion of security for randomized schemes that becomes identical to PRIV in the special case that the scheme is deterministic. Adaptive IND-CDA, when restricted to deterministic schemes, is an adaptive strengthening of PRIV that we think is interesting in its own right. As a consequence of the results discussed below, we get the first deterministic encryption schemes that achieve this stronger notion.

SCHEMES WITH RANDOM ORACLES. Our random oracle (RO) model schemes and their attributes are summarized in the first two rows of the table of Figure 1. Both REwH1 and REwH2 efficiently transform an arbitrary (randomized) IND-CPA scheme into a H-IND scheme with the aid of the RO. They are simple ways to make in-practice encryption schemes H-IND secure with minimal software changes. REwH1 has the advantage of not changing the public key and thus not requiring new certificates. It always provides non-adaptive H-IND security. It provides adaptive H-IND security if the starting scheme has the extra property of being anonymous in the sense of [4]. Anonymity is possessed by some deployed schemes like DHIES [1], making REwH1 attractive in this case. But some in-practice schemes, notably RSA ones, are not anonymous. If one wants adaptive H-IND security in this case we suggest REwH2, which provides it assuming only that the starting scheme is IND-CPA. It does this by adding a randomizer to

	Non-adaptive H-IND	Adaptive H-IND
REwH1	IND-CPA	IND-CPA + ANON-CPA
REwH2	IND-CPA	IND-CPA
RtD	IND-CPA, PRIV	IND-CPA, (u-)LTDF
PtD	(u-)LTDF	(u-)LTDF

Fig. 1. Table entries for the first two rows indicate the assumptions made on the (randomized) encryption scheme that underlies the RO-model hedged schemes in question. The entries for standard model scheme RtD are the assumptions on the underlying randomized and deterministic encryption schemes, respectively, and for PtD, on the underlying deterministic encryption scheme, which is the only primitive it uses.

the public key, so it does require new certificates. The schemes are extensions of the EwH deterministic encryption scheme of [6] and similar to [20].

SCHEMES WITHOUT RANDOM ORACLES. It is easy to see that even the existence of a non-adaptively secure IND-CDA encryption scheme implies the existence of a PRIV-secure deterministic encryption (DE) scheme. Achieving PRIV without ROs is already hard. Indeed, fully PRIV-secure DE without ROs has not yet been built. Prior work, however, does show how to construct PRIV-secure DE without ROs for block sources [12]. (Messages being encrypted have high min-entropy even conditioned on previous messages.) But H-IND introduces three additional challenges: (1) the min-entropy guarantee is on the joint message-randomness distribution rather than merely on the message; (2) we want a single scheme that is not only IND-CDA secure but also IND-CPA-secure; and (3) the adversary’s queries may be adaptive.

We are able to overcome these challenges to the best extent possible. We provide schemes that are H-IND-secure in the same setting as the best known PRIV ones, namely, for block sources, where we suitably extend the latter notion to consider both randomness and messages. Furthermore, we achieve these results under the same assumptions as previous work.

Our standard model schemes and their attributes are summarized in the last two rows of the table of Figure 1. RtD is formed by the generic composition of a deterministic scheme and a randomized scheme and achieves non-adaptive H-IND security as long as the base schemes meet their regular conditions. (That is, the former is PRIV-secure for block sources and the latter is IND-CPA.) Adaptive security requires that the deterministic scheme be a u-LTDF. (A lossy trapdoor function whose lossy branch is a universal hash function [31, 12].) PtD is simpler, merely concatenating the message to the randomness and then applying deterministic encryption. It achieves both non-adaptive and adaptive H-IND under the assumption that the deterministic scheme is a u-LTDF. For both schemes, the universality assumption on the LTDF can be dropped by modifying the scheme and using the crooked leftover hash lemma as per [12]. (This is why the “u” is parenthesized in the table of Figure 1.)

ANONYMOUS LTDFs. Also of independent interest, we show that any u-LTDF is anonymous. Here we refer to a new notion of anonymity for trapdoor functions that we introduce, one that strengthens the notion of [4]. This step exploits an adaptive variant of the leftover hash lemma of [26].

Why anonymity? It is exploited in our proofs of adaptive security. Our new notion of anonymity for trapdoor functions is matched by a corresponding one for encryption schemes. We show that any encryption scheme that is both anonymous and non-adaptive H-IND secure is also adaptively H-IND secure. Anonymity of the u-LTDF, in our encryption schemes based on the latter primitive, allows us to show that these schemes are anonymous and thereby lift their non-adaptive security to adaptive.

RELATED WORK. In the symmetric setting, several works have recognized and addressed the problem of security in the face of bad randomness. Concern over the quality of available randomness is one of Rogaway’s motivations for introducing nonce-based symmetric encryption [32], where security relies on the nonce never repeating rather than being random. Rogaway and Shrimpton [33] provide a symmetric authenticated encryption scheme that defaults to a PRF when the randomness is known.

Kamara and Katz [27] provide symmetric encryption schemes secure against chosen-randomness attack (CRA). Here the adversary can obtain encryption under randomness of its choice but privacy is only required for messages encrypted with perfect, hidden randomness. Entropy in the messages is not considered or used. We in contrast seek privacy even when the randomness is bad as long as there is compensating entropy in the message. Also we deal with the public key setting.

Many works consider achieving strong cryptography given only a “weak random source” [28, 16, 14]. This is a source that does have high min-entropy but may not produce truly random bits. They show that many cryptographic tasks including symmetric encryption [28], commitment, secret-sharing, and zero knowledge [16] are impossible in this setting. We are not in this setting. We do assume a small amount of initial good randomness to produce keys. (This makes sense because it is one-time and because otherwise we can’t hope to achieve anything anyway.) On the other hand our assumption on the randomness available for encryption is even weaker than in the works mentioned. (We do not even assume it has high min-entropy.) Our key idea is to exploit the entropy in the message, which is not done in [28, 16, 14]. This allows us to circumvent their negative results.

Waters independently proposed hedge security as well as the PtD construction as a way to achieve it [35].

2 Preliminaries

NOTATION. Vectors are written in boldface, e.g. \mathbf{x} . If \mathbf{x} is a vector then $|\mathbf{x}|$ denotes its length and $\mathbf{x}[i]$ denotes its i^{th} component for $1 \leq i \leq |\mathbf{x}|$. We say that \mathbf{x} is a vector over D if $\mathbf{x}[i] \in D$ for all $1 \leq i \leq |\mathbf{x}|$. Throughout, $k \in \mathbb{N}$ denotes the

security parameter and 1^k its unary encoding. Unless otherwise indicated, an algorithm is randomized. The set of possible outputs of algorithm A on inputs x_1, x_2, \dots is denoted $[A(x_1, x_2, \dots)]$. “PT” stands for polynomial-time.

GAMES. Our security definitions and proofs use code-based games [9], and so we recall some background from [9]. A game (look at Figure 2 for examples) has an **Initialize** procedure, procedures to respond to adversary oracle queries, and a **Finalize** procedure. A game G is executed with an adversary A as follows. First, **Initialize** executes, and its outputs are the inputs to A . Then A executes, its oracle queries being answered by the corresponding procedures of G . When A terminates, its output becomes the input to the **Finalize** procedure. The output of the latter is called the output of the game, and we let $G^A \Rightarrow y$ denote the event that this game output takes value y . Our convention is that the running time of an adversary is the time to execute the adversary with the game that defines security, so that the running time of all game procedures is included.

PUBLIC-KEY ENCRYPTION. A public-key encryption (PKE) scheme is a tuple of PT algorithms $\mathcal{AE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ with associated message length parameter $n(\cdot)$ and randomness length parameter $\rho(\cdot)$. The parameter generation algorithm \mathcal{P} takes as input 1^k and outputs a parameter string par . The key generation algorithm \mathcal{K} takes input par and outputs a key pair (pk, sk) . The encryption algorithm \mathcal{E} takes inputs pk , message $m \in \{0, 1\}^{n(k)}$ and coins $r \in \{0, 1\}^{\rho(k)}$ and returns the ciphertext denoted $\mathcal{E}(pk, m; r)$. The deterministic decryption algorithm \mathcal{D} takes input sk and ciphertext c and outputs either \perp or a message in $\{0, 1\}^{n(k)}$. For vectors \mathbf{m}, \mathbf{r} with $|\mathbf{m}| = |\mathbf{r}| = v$ we denote by $\mathcal{E}(pk, \mathbf{m}; \mathbf{r})$ the vector $(\mathcal{E}(pk, \mathbf{m}[1]; \mathbf{r}[1]), \dots, \mathcal{E}(pk, \mathbf{m}[v]; \mathbf{r}[v]))$. We say that \mathcal{AE} is deterministic if \mathcal{E} is deterministic. (That is, $\rho(\cdot) = 0$.)

We consider the standard IND-CPA notion of security, captured by the game $\text{IND}_{\mathcal{AE}}$ where $\mathcal{AE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ is an encryption scheme. In the game, **Initialize** chooses a random bit b , generates parameters $par \leftarrow \mathcal{P}(1^k)$ and generates a key pair $(pk, sk) \leftarrow \mathcal{K}(par)$ before returning pk to the adversary. Procedure **LR**, on input messages m_0 and m_1 , returns $c \leftarrow \mathcal{E}(pk, m_b)$. Lastly, procedure **Finalize** takes as input a guess bit b' and outputs **true** if $b = b'$ and **false** otherwise. An IND-CPA adversary makes a single query (m_0, m_1) to **LR** with $|m_0| = |m_1|$. For IND-CPA adversary A we let $\text{Adv}_{\mathcal{AE}, A}^{\text{ind-cpa}}(k) = 2 \cdot \Pr [\text{IND}_{\mathcal{AE}, k}^A \Rightarrow \text{true}] - 1$. We say \mathcal{AE} is IND-CPA secure if $\text{Adv}_{\mathcal{AE}, A}^{\text{ind-cpa}}(\cdot)$ is negligible for all PT IND-CPA adversaries A .

SOURCES. We generalize the notion of a source to consider a joint distribution on the messages and the randomness with which they will be encrypted. A t -source ($t \geq 1$) with message length $n(\cdot)$ and randomness length $\rho(\cdot)$ is a probabilistic algorithm \mathcal{M} that on input 1^k returns a $(t+1)$ -tuple $(\mathbf{m}_0, \dots, \mathbf{m}_{t-1}, \mathbf{r})$ of equal-length vectors, where $\mathbf{m}_0, \dots, \mathbf{m}_{t-1}$ are over $\{0, 1\}^{n(k)}$ and \mathbf{r} is over $\{0, 1\}^{\rho(k)}$. We say that \mathcal{M} has min-entropy $\mu(\cdot)$ if

$$\Pr [(\mathbf{m}_b[i], \mathbf{r}[i]) = (m, r)] \leq 2^{-\mu(k)}$$

for all $k \in \mathbb{N}$, all $b \in \{0 \dots, t-1\}$, all i and all $(m, r) \in \{0, 1\}^{n(k)} \times \{0, 1\}^{\rho(k)}$. We say it has conditional min-entropy $\mu(\cdot)$ if

$$\Pr [(\mathbf{m}_b[i], \mathbf{r}[i]) = (m, r) \mid \forall j < i (\mathbf{m}_b[j], \mathbf{r}[j]) = (\mathbf{m}'[j], \mathbf{r}'[j])] \leq 2^{-\mu(k)}$$

for all $k \in \mathbb{N}$, all $b \in \{0 \dots, t-1\}$, all i , all (m, r) , and all vectors \mathbf{m}', \mathbf{r}' . A t -source with message length $n(\cdot)$, randomness length $\rho(\cdot)$, and min-entropy $\mu(\cdot)$ is referred to as a (μ, n, ρ) -**mr**-source when $t = 1$ and $\rho(\cdot) > 0$; a (μ, n) -**m**-source when $t = 1$ and $\rho(\cdot) = 0$; a (μ, n, ρ) -**mmr**-source when $t = 2$ and $\rho(\cdot) > 0$; and (μ, n) -**mm**-source when $t = 2$ and $\rho(\cdot) = 0$. Each “**m**” indicates the source outputting one message vector and an “**r**” indicates a randomness vector. When the source has *conditional* min-entropy $\mu(\cdot)$ we write block-source instead of source for each of the above. A $v(\cdot)$ -vector source outputs vectors of size $v(k)$ for all k .

UNIVERSAL HASH FUNCTIONS. A family of functions is a tuple $\mathcal{H} = (\mathcal{P}, \mathcal{K}, F)$ with associated message length $n(\cdot)$. It is required that the domain of $F(K, \cdot)$ is $\{0, 1\}^n$ for every k , every $par \in [\mathcal{P}(1^k)]$, and every $K \in [\mathcal{K}(par)]$. We say that \mathcal{H} is universal if for every k , all $par \in [\mathcal{P}(1^k)]$, and all distinct $x_1, x_2 \in \{0, 1\}^{n(k)}$, the probability that $F(K, x_1) = F(K, x_2)$ is at most $1/|R(par)|$ where $R(par) = \{F(K, x) : K \in [\mathcal{K}(par)] \text{ and } x \in \{0, 1\}^n\}$ and the probability is over $K \leftarrow_s \mathcal{K}(par)$.

LOSSY TRAPDOOR FUNCTIONS (LTDFs). To a deterministic PKE scheme (recall that a family of injective trapdoor functions and a deterministic encryption scheme are, syntactically, the same object) $\mathcal{AE} = (\mathcal{P}_d, \mathcal{K}_d, \mathcal{E}_d, \mathcal{D}_d)$ with message length $n_d(\cdot)$ we can associate an (n_d, ℓ) -lossy key generator \mathcal{K}_l . This is a PT algorithm that, on input par , outputs a value pk for which the map $\mathcal{E}_d(pk, \cdot)$ has image size at most $2^{n_d(k) - \ell(k)}$. The parameter ℓ is called the lossiness of the lossy key generator. We associate to \mathcal{AE} , lossy key generator \mathcal{K}_l , and a LOS adversary A the function $\mathbf{Adv}_{\mathcal{AE}, \mathcal{K}_l, A}^{\text{los}}(k) = 2 \cdot \Pr [\text{LOS}_{\mathcal{AE}, \mathcal{K}_l, k}^A \Rightarrow \text{true}] - 1$, where game $\text{LOS}_{\mathcal{AE}, \mathcal{K}_l}$ works as follows. **Initialize** chooses a random bit b and generates parameters $par \leftarrow_s \mathcal{P}_d(1^k)$, if $b = 0$ runs $(pk, sk) \leftarrow_s \mathcal{K}_d(par)$ and if $b = 1$ runs $pk \leftarrow_s \mathcal{K}_l(par)$. It then returns pk (to the adversary A). When A finishes, outputting guess b' , **Finalize** returns **true** if $b = b'$. We say \mathcal{K}_l is *universal-inducing* if $\mathcal{H} = (\mathcal{P}_d, \mathcal{K}_l, \mathcal{E}_d)$ is a family of universal hash functions with message length n_d .

A deterministic encryption scheme \mathcal{AE} is a (n_d, ℓ) -lossy trapdoor function (LTDF) if there exists a (n_d, ℓ) -lossy key generator such that $\mathbf{Adv}_{\mathcal{AE}, \mathcal{K}_l, A}^{\text{los}}(\cdot)$ is negligible for all PT A . We say it is a universal (n_d, ℓ) -lossy trapdoor function (u-LTDF) if in addition \mathcal{K}_l is universal-inducing.

Lossy trapdoor functions were introduced by Peikert and Waters [31], and can be based on a variety of number-theoretic assumptions, including the hardness of the decisional Diffie-Hellman problem, the worst-case hardness of lattice problems, and the hardness of Paillier’s composite residuosity problem [31, 12, 34]. Boldyreva et al. [12] observed that the DDH-based construction is universal.

<p>proc. Initialize(1^k):</p> $par \leftarrow \mathcal{P}(1^k)$ $(pk, sk) \leftarrow \mathcal{K}(par)$ $b \leftarrow \{0, 1\}$ Ret par	<p>proc. LR(\mathcal{M}):</p> If $pkout = \text{true}$ then Ret \perp $(\mathbf{m}_0, \mathbf{m}_1, \mathbf{r}) \leftarrow \mathcal{M}(1^k)$ Ret $\mathcal{E}(pk, \mathbf{m}_b; \mathbf{r})$	<p>proc. RevealPK(\cdot):</p> $pkout \leftarrow \text{true}$ Ret pk
		<p>proc. Finalize(b'):</p> Ret $(b = b')$

Fig. 2. Game $\text{CDA}_{\mathcal{AE}, k}$

3 Security against Chosen Distribution Attack

Let $\mathcal{AE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. A CDA adversary is one whose **LR** queries are all **mmr**-sources. Game $\text{CDA}_{\mathcal{AE}}$ of Figure 2 provides the adversary with two oracles. The advantage of CDA adversary A is

$$\text{Adv}_{\mathcal{AE}, A}^{\text{cda}}(k) = 2 \cdot \Pr [\text{CDA}_{\mathcal{AE}, k}^A \Rightarrow \text{true}] - 1.$$

In the random oracle model we allow all algorithms in Game CDA to access the random oracle; importantly, this includes the **mmr**-sources.

DISCUSSION. Adversary A can query **LR** with an **mmr**-source of its choice, an output $(\mathbf{m}_0, \mathbf{m}_1, \mathbf{r})$ of which represents choices of message vectors to encrypt and randomness with which to encrypt them. (An alternative formulation might have CDA adversaries query two **mr**-sources, and distinguish between the encryption of samples taken from one of these. But this would mandate that schemes ensure privacy of messages *and* randomness.) This allows A to dictate a joint distribution on the messages and randomness. In this way it conservatively models even adversarially-subverted random number generators. Multiple **LR** queries are allowed. In the most general case these queries may be adaptive, meaning depend on answers to previous queries.

Given that multiple **LR** queries are allowed, one may ask why an **mmr**-source needs to produce message and randomness vectors rather than simply a single pair of messages and a single choice of randomness. The reason is that the coordinates in a vector all depend on the same coins underlying an execution of \mathcal{M} , but the coins underlying the execution of the sources in different queries are independent.

Note that **Initialize** does not return the public key pk to A . A can get it at any time by calling **RevealPK** but once it does this, **LR** will return \perp . The reason is that we inherit from deterministic encryption the unavoidable limitation that encryption cannot hide public-key related information about the plaintexts [6]. (When the randomness has low entropy, the ciphertext itself is such information.)

As we saw in the previous section, no encryption scheme is secure when both messages and randomness are predictable. Formally, this means chosen-distribution attacks are trivial when adversaries can query **mmr**-sources of low

min-entropy. Our notions (below) will therefore require security only for sources that have high min-entropy or high conditional min-entropy.

EQUALITY PATTERNS. Suppose A makes a query \mathcal{M} which returns $(\mathbf{m}_0, \mathbf{m}_1, \mathbf{r}) = ((a, a), (a, a'), (r, r))$ for some $a \neq a'$ and random r . Then it can win trivially because the (two) components of the returned vector \mathbf{c} are equal if $b = 0$ and unequal otherwise. This limitation, again inherited from deterministic encryption [6], is inherent. To capture it we associate to an mmr -source \mathcal{M} an equality-pattern probability

$$\zeta(k) = \Pr [\text{eq}((\mathbf{m}_0, \mathbf{r}), (\mathbf{m}_1, \mathbf{r})) = 0 : (\mathbf{m}_0, \mathbf{m}_1, \mathbf{r}) \leftarrow_s \mathcal{M}(1^k)]$$

where $\text{eq}((\mathbf{x}_1, \mathbf{x}_2), (\mathbf{y}_1, \mathbf{y}_2))$ is 1 if for all i, j

$$(\mathbf{x}_1[i], \mathbf{x}_2[i]) = (\mathbf{x}_1[j], \mathbf{x}_2[j]) \text{ iff } (\mathbf{y}_1[i], \mathbf{y}_2[i]) = (\mathbf{y}_1[j], \mathbf{y}_2[j]),$$

and 0 otherwise. We point out that **LR** queries that are mmr -block-sources (and not, just, mmr -sources) with high conditional min-entropy have negligible equality-pattern probability.

NOTIONS. We can assume (without loss of generality) that a CDA adversary makes a single **RevealPK** query and then no further **LR** queries. We say A is a (μ, n, ρ) -adversary if all of its **LR** queries are (μ, n, ρ) - mmr -sources. We say that a PKE scheme \mathcal{AE} with message length $n(\cdot)$ and randomness length $\rho(\cdot)$ is IND-CDA secure for (μ, n, ρ) - mmr -sources if for all PT (μ, n, ρ) adversaries A the function $\text{Adv}_{\mathcal{AE}, A}^{\text{cda}}(\cdot)$ is negligible. Scheme \mathcal{AE} is H-IND secure for (μ, n, ρ) - mmr -sources if it is IND-CPA secure and IND-CDA secure for (μ, n, ρ) - mmr -sources. We can extend these notions to mmr -block-sources by restricting to adversaries that query mmr -block-sources.

ON ADAPTIVITY. We can consider non-adaptive IND-CDA security by restricting attention in the notions above to adversaries that only make a single **LR** query. Why do we not focus solely on this (simpler) security goal? The standard IND-CPA setting (implicitly) provides security against multiple, adaptive **LR** queries. This is true because in that setting a straightforward hybrid argument shows that security against multiple adaptive **LR** queries is implied by security against a single **LR** query [5, 3]. We wish to maintain the same standard of adaptive security in the IND-CDA setting. Unfortunately, in the IND-CDA setting, unlike the IND-CPA setting, adaptive security is not implied by non-adaptive security. In short this is because a CDA adversary necessarily cannot learn the public key before (or while) making **LR** queries. To see the separation, consider a PKE scheme that appends to every ciphertext the public key used. This will not affect the security of the scheme when an adversary can only make a single query. However, an adaptive CDA adversary can query an mmr -source, learn the public key, and craft a second source that uses the public key to ensure ciphertexts which leak the challenge bit.

Given this, our primary goal is the stronger notion of adaptive security. That said, non-adaptive hedge security is also relevant because in practice adaptive adversaries might be rare and (as we will see in Section 5) one can find

non-adaptively-secure schemes that are more efficient and/or have proofs under weaker assumptions.

ADAPTIVE PRIV. A special case of our framework occurs when the PKE scheme \mathcal{AE} being considered has randomness length $\rho(k) = 0$ for all k (meaning also that adversaries query **mm**-sources, instead of **mmr**-sources). In this case we are considering deterministic encryption, and the IND-CDA definition and notions give a strengthening (by way of adaptivity) of the PRIV security notion from [6, 8, 12]. (For non-adaptive adversaries the definitions are equivalent.) For clarity we will use PRIV to refer to this special case, and let $\mathbf{Adv}_{\mathcal{AE},A}^{\text{priv}}(k) = \mathbf{Adv}_{\mathcal{AE},A}^{\text{cda}}(k)$.

RESOURCE USAGE. Recall that by our convention, the running time of a CDA adversary is the time for the execution of the adversary with game $\text{CDA}_{\mathcal{AE},k}$. Thus, A being PT implies that the **mmr**-sources that comprise A 's **LR** queries are also PT. This is a distinction from [12] which will be important in our results. Note that in practice we do not expect to see sources that are not PT, so our definition is not restrictive. Non-PT sources were needed in [12] for showing that single-message security implied (non-adaptive) multi-message security for deterministic encryption of block sources.

4 Constructions

Here we present several constructions for hedged encryption. The first scheme uses a random oracle and an IND-CPA secure probabilistic encryption scheme. The next two schemes derive from composing a randomized encryption scheme with a deterministic one (there are two ways of ordering composition). Interestingly, only one ordering will end up providing security. The final scheme converts a deterministic encryption scheme to a hedged one by padding the message with random bits. For the following, let $\mathcal{AE}_r = (\mathcal{P}_r, \mathcal{K}_r, \mathcal{E}_r, \mathcal{D}_r)$ be a (randomized) PKE scheme with message length $n_r(\cdot)$ and randomness length $\rho(\cdot)$. Let $\mathcal{AE}_d = (\mathcal{P}_d, \mathcal{K}_d, \mathcal{E}_d, \mathcal{D}_d)$ be a (deterministic) PKE scheme with message length $n_d(\cdot)$ and randomness length always 0. Associate to \mathcal{AE}_c for $c \in \{d, r\}$ the function $\text{maxcLen}_c(k)$ mapping any k to the maximum length (over all possible public keys, messages, and if applicable, randomness) of a ciphertext output by \mathcal{E}_c .

RANDOMIZED-ENCRYPT-WITH-HASH. Let $\mathcal{R} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be a random oracle. Let $\text{REwH}[\mathcal{AE}_r] = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be the scheme parameterized by randomizer length κ that works as follows. Parameter generation, and decryption are the same as in \mathcal{AE}_r . Key generation runs $\mathcal{K}_r(\text{par}_r)$ to get (pk_r, sk_r) , chooses $K \leftarrow_{\$} \{0, 1\}^{\kappa(k)}$, and lets $pk = (pk_r \parallel K)$ and $sk = sk_r$. Algorithm $\mathcal{E}^{\mathcal{R}}$, on input (pk, m) where $pk = (pk_r \parallel K)$, chooses $r \leftarrow_{\$} \{0, 1\}^{\rho(k)}$ and computes $r' \leftarrow \mathcal{R}(pk_r \parallel K \parallel r \parallel m)$ (where here we take \mathcal{R} 's output to be of length $\rho(k)$) and outputs $\mathcal{E}_r(pk_r, m; r')$. Intuitively, the random oracle provides perfect and (as long as m and r are hard to predict) private randomness. When the key length $\kappa(k) = 0$ for all k , we refer to the scheme as **REwH1**, while when $\kappa(k) > 0$ for all k we refer to the scheme as **REwH2**. The scheme extends the Encrypt-with-Hash deterministic encryption scheme from [6], which is a special case of **REwH1** when

r has length 0, and is also reminiscent of constructions in the symmetric setting that utilize a PRF to ensure good randomness [27, 33], as well as schemes using the Fujisaki-Okamoto transform [20].

DETERMINISTIC-THEN-RANDOMIZED. Our first standard model attempt is to perform hedged encryption via first applying deterministic encryption and then randomized. More formally let $\text{DtR}[\mathcal{AE}_r, \mathcal{AE}_d] = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be the scheme that works as follows. The parameter generation algorithm \mathcal{P} runs $par_r \leftarrow \mathcal{P}_r(1^k)$ and $par_d \leftarrow \mathcal{P}_d(1^k)$ and outputs $par = (par_r, par_d)$. Key generation \mathcal{K} just runs $(pk_r, sk_r) \leftarrow \mathcal{K}_r(par_r)$ and $(pk_d, sk_d) \leftarrow \mathcal{K}_d(par_d)$ and outputs $pk = (pk_r, pk_d)$ and $sk = (sk_r, sk_d)$. We define encryption by

$$\mathcal{E}((pk_r, pk_d), m; r) = \mathcal{E}_r(pk_r, c \parallel 10^\ell; r),$$

where $c = \mathcal{E}_d(pk_d, m)$ and $\ell = n_r - |c| - 1$. Here we need that $n_r(k) > \max_{\text{clen}_d}(k)$ for all k . Decryption is defined in the natural way. The scheme will clearly inherit IND-CPA security from the application of \mathcal{E}_r . If the deterministic encryption scheme is PRIV secure for min-entropy μ , then the composition will also be secure if the *message* has min-entropy at least μ . However, our strong notion of IND-CDA security requires that schemes be secure if the *joint* distribution on the message and randomness has high min-entropy. If the entropy is unfortuitously split between both the randomness and the message, then there is no guarantee that the composition will be secure. In fact, many choices for instantiating \mathcal{AE}_r and \mathcal{AE}_d lead to a composition for which attacks can be exhibited (even when the schemes are, separately, secure).

RANDOMIZED-THEN-DETERMINISTIC. We can instead apply randomized encryption first, and then apply deterministic encryption. Define $\text{RtD}[\mathcal{AE}_r, \mathcal{AE}_d] = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ to work as follows. The parameter and key generation algorithms are as for scheme DtR. Encryption is defined by

$$\mathcal{E}((pk_r, pk_d), m; r) = \mathcal{E}_d(pk_d, c \parallel 10^\ell).$$

where $c = \mathcal{E}_r(pk_r, m; r)$ and $\ell = n_d - |c| - 1$. Here we need that $n_d(k) > \max_{\text{clen}_r}(k)$ for all k . The decryption algorithm \mathcal{D} works in the natural way. As we will see, this construction avoids the security issues of the previous, as long as the randomized encryption scheme preserves the min-entropy of its inputs. (For example, if for all k , all $par_r \in [\mathcal{P}_r(1^k)]$, and all $(pk_r, sk_r) \in [\mathcal{K}_r(par_r)]$, $\mathcal{E}_r(pk_r, \cdot)$ is injective in (m, r) .) Many encryption schemes have this property; El Gamal [21] is one example.

PAD-THEN-DETERMINISTIC. Our final construction dispenses entirely with the need for a dedicated randomized encryption scheme, instead using simple padding to directly construct a (randomized) encryption scheme from a deterministic one. Let $\text{PtD}[\mathcal{AE}_d] = (\mathcal{P}_d, \mathcal{K}_d, \mathcal{E}, \mathcal{D})$ work as follows. Parameter and key generation are inherited from the underlying (deterministic) encryption scheme. Encryption is defined by

$$\mathcal{E}(pk_d, m; r) = \mathcal{E}_d(pk_d, r \parallel m).$$

Decryption proceeds by applying \mathcal{D}_d , to retrieve $r \parallel m$, and then returning m .

5 Non-adaptive Hedge Security

In this section we investigate the non-adaptive hedge security of REwH, RtD and PtD, leaving adaptive security to future sections.

RANDOMIZED-ENCRYPT-WITH-HASH. Intuitively, the security of $\text{REwH}[\mathcal{AE}_r]$ follows from the IND-CPA security of \mathcal{AE}_r and the random oracle providing “perfect” randomness. Following [6], for any k let $\text{maxpk}_{\mathcal{AE}}(k)$ be the maximum of $\Pr[pk = w : (pk, sk) \leftarrow_s \mathcal{K}(\text{par})]$, where the maximum is taken over all $w \in \{0, 1\}^*$ and all $\text{par} \in [\mathcal{P}(1^k)]$.

Theorem 1. [REwH is non-adaptive H-IND secure] *Let $\mathcal{AE}_r = (\mathcal{P}_r, \mathcal{K}_r, \mathcal{E}_r, \mathcal{D}_r)$ be a PKE scheme with message length $n(\cdot)$ and randomness length ρ and let $\mathcal{AE} = \text{REwH}[\mathcal{AE}_r] = (\mathcal{P}_r, \mathcal{K}_r, \mathcal{E}, \mathcal{D}_r)$ be the PKE scheme constructed from it.*

- (IND-CPA) *Let A be an IND-CPA adversary. Then there exists an IND-CPA adversary B such that for all k*

$$\text{Adv}_{\mathcal{AE}, A}^{\text{ind-cpa}}(k) = \text{Adv}_{\mathcal{AE}_r, B}^{\text{ind-cpa}}(k)$$

where B runs in time that of A and makes the same number of queries.

- (IND-CDA) *Let A be an adversary that makes a single **LR** query consisting of a $v(\cdot)$ -vector (μ, n, ρ) -mmr-source with equality-pattern probability $\zeta(\cdot)$ and making at most $h(\cdot)$ random oracle queries. Then there exists an IND-CPA adversary B such that for all k*

$$\text{Adv}_{\mathcal{AE}, A}^{\text{cda}}(k) \leq v(k) \left(\text{Adv}_{\mathcal{AE}_r, B}^{\text{ind-cpa}}(k) + \frac{2 \cdot h(k)}{2^{\mu(k)}} + 8 \cdot \text{maxpk}_{\mathcal{AE}_r}(k) \right) + \zeta(k)$$

Adversary B runs in time that of A and $\text{maxpk}_{\mathcal{AE}_r}$ is the maximum public key probability of \mathcal{AE}_r . \square

The first part of the theorem is straightforward to prove. The second follows from an adaptation of the proof of security for the similar Encrypt-with-Hash deterministic encryption scheme in [6]. Notice that the theorem holds for both REwH1 and REwH2; the only difference is that with the latter the $\text{maxpk}_{\mathcal{AE}}(k)$ term improves depending on the length κ .

RANDOMIZED-THEN-DETERMINISTIC. Intuitively, the non-adaptive hedged security of the RtD construction is inherited from the IND-CPA security of the underlying randomized scheme \mathcal{AE}_r and the (non-adaptive) PRIV security of the underlying deterministic scheme \mathcal{AE}_d . As alluded to before, we have one technical requirement on \mathcal{AE}_r for the IND-CDA proof to work. We say $\mathcal{AE}_r = (\mathcal{P}_r, \mathcal{K}_r, \mathcal{E}_r, \mathcal{D}_r)$ with message length $n_r(\cdot)$ and randomness length $\rho(\cdot)$ is *min-entropy preserving* if for any k , any $\text{par}_r \in [\mathcal{P}_r(1^k)]$, any $(pk_r, sk_r) \in [\mathcal{K}_r(\text{par}_r)]$, and for all $c \in \{0, 1\}^*$ it is the case for any (μ, n_r, ρ) -mr-source \mathcal{M} outputting vectors of size one that $\Pr[c = \mathcal{E}_r(pk_r, m; r) : (m, r) \leftarrow_s \mathcal{M}(1^k)] \leq 2^{-\mu}$. In words, encryption preserves the min-entropy of the input message and randomness. We have the following theorem.

Theorem 2. [RtD is non-adaptive H-IND secure] Let $\mathcal{AE}_r = (\mathcal{P}_r, \mathcal{K}_r, \mathcal{E}_r, \mathcal{D}_r)$ be a min-entropy preserving PKE scheme with message length $n_r(\cdot)$ and randomness length $\rho(\cdot)$. Let $\mathcal{AE}_d = (\mathcal{P}_d, \mathcal{K}_d, \mathcal{E}_d, \mathcal{D}_d)$ be a (deterministic) encryption scheme with message length $n_d(\cdot)$ so that $n_d(\cdot) \geq \max\text{clen}_r(\cdot)$. Let $\mathcal{AE} = \text{RtD}[\mathcal{AE}_r, \mathcal{AE}_d] = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be the PKE scheme defined in Section 4.

- (IND-CPA) Let A be an IND-CPA adversary. Then there exists an IND-CPA adversary B such that for any k

$$\text{Adv}_{\mathcal{AE}, A}^{\text{ind-cpa}}(k) = \text{Adv}_{\mathcal{AE}_r, B}^{\text{ind-cpa}}(k)$$

where B runs in time that of A plus the time to run \mathcal{E}_d once.

- (IND-CDA) Let A be a CDA adversary that makes one LR query consisting of a $v(\cdot)$ -vector (μ, n_r, ρ) -mmr-source (resp. block-source). Then there exists a PRIV adversary B such that for any k

$$\text{Adv}_{\mathcal{AE}, A}^{\text{cda}}(k) \leq \text{Adv}_{\mathcal{AE}_d, B}^{\text{priv}}(k)$$

where B runs in time that of A plus the time to run $v(k)$ executions of \mathcal{E}_r and makes one LR query consisting of a $v(\cdot)$ -vector $(\mu, \max\text{clen}_r)$ -mm-source (resp. block-source). \square

Note that the second part of the theorem states the result for either sources or just block-sources. We briefly sketch the proof. The first part of the theorem is immediate from the IND-CPA security of \mathcal{AE}_r . For the second part, any mmr-source \mathcal{M} queried by A is converted into an mm-source \mathcal{M}' to be queried by B . This is done by having \mathcal{M}' run \mathcal{M} to get $(\mathbf{m}_0, \mathbf{m}_1, \mathbf{r})$ and then outputting the pair of vectors $(\mathcal{E}_r(pk, \mathbf{m}_0; \mathbf{r}), \mathcal{E}_r(pk, \mathbf{m}_1; \mathbf{r}))$. (The ciphertexts are the “messages” for \mathcal{E}_d .) Because \mathcal{AE}_r is min-entropy preserving, \mathcal{M}' is a source of the appropriate type.

PAD-THEN-DETERMINISTIC. The security of the PtD scheme is more difficult to establish. The IND-CDA security is inherited immediately from the PRIV security of the \mathcal{AE}_d scheme. Here the challenge is, in fact, proving IND-CPA security. For this we will need a stronger assumption on the underlying deterministic encryption scheme — that it is a u-LTDF.

Theorem 3. [PtD is non-adaptive H-IND secure] Let $\mathcal{AE}_d = (\mathcal{P}_d, \mathcal{K}_d, \mathcal{E}_d, \mathcal{D}_d)$ be a deterministic encryption scheme with message length $n_d(\cdot)$. Let $\mathcal{AE} = \text{PtD}[\mathcal{AE}_d] = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be the PKE scheme defined in Section 4 with message length $n(\cdot)$ and randomness length $\rho(\cdot)$ such that $n(k) = n_d(k) - \rho(k)$ for all k .

- (IND-CPA) Let \mathcal{K}_l be a universal-inducing (n_d, ℓ) -lossy key generation algorithm for \mathcal{AE}_d . Let A be an IND-CPA adversary. Then there exists a LOS adversary B such that for all k

$$\text{Adv}_{\mathcal{AE}, A}^{\text{ind-cpa}}(k) \leq \text{Adv}_{\mathcal{AE}_d, \mathcal{K}_l, B}^{\text{los}}(k) + \sqrt{2^{3n(k) - \ell(k) + 2}}.$$

B runs in time that of A .

- (IND-CDA) Let A be a CDA adversary that makes one LR query consisting of a $v(\cdot)$ -vector (μ, n, ρ) -mmr-source (resp. block-source). Then there exists

a PRIV adversary B such that for all k

$$\mathbf{Adv}_{\mathcal{AE},A}^{\text{cda}}(k) \leq \mathbf{Adv}_{\mathcal{AE}_d,B}^{\text{priv}}(k)$$

where B runs in time that of A and makes one LR query consisting of a $v(\cdot)$ -vector (μ, n_d) -mm-source (resp. block-source). \square

One might think that concluding IND-CPA can be based just on PtD being IND-CDA secure, since the padded randomness provides high min-entropy. However, this approach does not work because an IND-CPA adversary expects knowledge of the public-key *before* making any **LR** queries, while a CDA adversary only learns the public-key *after* making its **LR** queries. This issue is discussed in more detail in [8]. We use a different approach (which may be of independent interest) to prove this part of Theorem 3; the details are given in the full version [7]. Our proof strategy, intuitively, corresponds to using the standard LHL $2^{n(k)}$ times, once for each possible message the IND-CPA adversary might query.

6 Anonymity for Chosen Distribution Attacks

In the previous section we proved non-adaptive security for the RtD and PtD constructions. But, as established in Section 3, we actually want to meet the stronger goal of adaptive security. In the adaptive setting, adversaries can make multiple **LR** queries, specifying sources that are generated as a function of previously-seen ciphertexts. Recall that one reason adaptivity is difficult to achieve is because ciphertexts might leak information about the public key. In turn, knowledge of the public key leads to trivial IND-CDA attacks. This suggests a natural relationship with key privacy, also called anonymity [4]. Anonymity requires (informally) that ciphertexts leak no information about the public key used to perform encryption. In this section we formalize a notion of anonymity for chosen-distribution attacks. In the next section we'll use this definition as a step towards adaptive IND-CDA security.

DEFINITIONS. Let $\mathcal{AE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. Game $\text{ANON}_{\mathcal{AE}}$ shown in Figure 3 provides the adversary with two oracles. An ANON adversary A is one whose queries are all **mr**-sources. The advantage of ANON adversary A is

$$\mathbf{Adv}_{\mathcal{AE},A}^{\text{anon}}(k) = 2 \cdot \Pr [\text{ANON}_{\mathcal{AE},k}^A \Rightarrow \text{true}] - 1 .$$

We say that a PKE scheme \mathcal{AE} with message length $n(\cdot)$ and randomness length $\rho(\cdot)$ is ANON secure for (μ, n, ρ) -**mr**-sources if for all PT adversaries A that only query (μ, n, ρ) -**mr**-sources the function $\mathbf{Adv}_{\mathcal{AE},A}^{\text{anon}}(\cdot)$ is negligible. We can extend this notion to **mr**-block-sources in the obvious way. In the special case that the randomness length of \mathcal{AE} is always zero, the ANON definition formalizes anonymity for deterministic encryption or, equivalently, trapdoor functions, generalizing a definition from [4].

DISCUSSION. Anonymity for PKE in the sense of key privacy was first formalized by Bellare et al. [4], but their notion (analogously to traditional semantic

<u>proc. Initialize(k):</u>	<u>proc. Enc(\mathcal{M}):</u>	<u>proc. LR(\mathcal{M}):</u>	<u>proc. Finalize(a'):</u>
$par \leftarrow \mathcal{P}(1^k)$	If $pkout = true$	$(\mathbf{m}, \mathbf{r}) \leftarrow \mathcal{M}(1^k)$	Ret ($a = a'$)
$(pk_0, sk_0) \leftarrow \mathcal{K}(par)$	Ret \perp	$\mathbf{c} \leftarrow \mathcal{E}(pk_a, \mathbf{m}; \mathbf{r})$	
$(pk_1, sk_1) \leftarrow \mathcal{K}(par)$	$(\mathbf{m}, \mathbf{r}) \leftarrow \mathcal{M}(1^k)$	$pkout \leftarrow true$	
$a \leftarrow \{0, 1\}$	Ret $\mathcal{E}(pk_0, \mathbf{m}; \mathbf{r})$	Ret (pk_0, pk_1, \mathbf{c})	
Ret par			

Fig. 3. Game $ANON_{\mathcal{AE}, k}$.

security) only works in the context of good randomness. The ANON notion, akin to IND-CDA, formalizes key privacy in the face of bad randomness. While we will use it mainly as a technical tool to simplify showing that schemes meet adaptive IND-CDA, it is also of independent interest as a new security target for PKE schemes when key privacy is important. (That is, one might want to hedge against bad randomness for anonymity as well as message privacy.)

7 Adaptive Hedge Security

The following theorem, whose proof appears in the full version [7], shows that achieving ANON security and non-adaptive IND-CDA security are sufficient for achieving adaptive IND-CDA security.

Theorem 4. *Let $\mathcal{AE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme with message length $n(\cdot)$ and randomness length $\rho(\cdot)$. Let A be a IND-CDA adversary making $q(\cdot)$ **LR** queries, each being a $v(\cdot)$ -vector (μ, n, ρ) -mmr-source (resp. block-source). Then there exist IND-CDA adversary B and ANON adversary C such that for all k*

$$\mathbf{Adv}_{\mathcal{AE}, A}^{\text{cda}}(k) \leq 2q(k) \cdot \mathbf{Adv}_{\mathcal{AE}, B}^{\text{cda}}(k) + 4q(k) \cdot \mathbf{Adv}_{\mathcal{AE}, C}^{\text{anon}}(k).$$

*B makes one **LR** query consisting of a $v(\cdot)$ -vector (μ, n, ρ) -mmr-source (resp. block-source). C makes at most $q(k) - 1$ **Enc** queries and one **LR** query, all these consisting of $v(\cdot)$ -vector (μ, n, ρ) -mr-sources (resp. block-sources). Both B and C run in the same time as A . \square*

Given a non-adaptively IND-CDA secure scheme, Theorem 4 reduces the task of showing it adaptively secure to that of showing it meets the ANON definition. Of course, ANON is still an adaptive notion. (Adversaries can formulate their **LR** query to be a source that's a function of previously seen ciphertexts.) Nevertheless, it formalizes a sufficient condition for adaptive CDA security of any PKE scheme and captures the relationship between adaptivity and anonymity. We believe this is an interesting (and novel) application of anonymity.

We can show that our random oracle scheme REwH is ANON secure when the underlying randomized scheme meets the traditional notions of anonymity for PKE [4]. We also want to show that the RtD and PtD schemes are ANON secure.

We first show something more general: that any u-LTDF is anonymous. Then, that RtD and PtD are anonymous follows when using deterministic schemes that are also u-LTDFs.

UNIVERSAL LTDFs ARE ANONYMOUS. Intuitively u-LTDFs are anonymous because the lossy mode admits a universal hash, implying that no information about the public key is leaked by outputs (generated from sources with high conditional min-entropy). One might expect that formalizing this intuition would follow from straightforward application of the Leftover Hash Lemma (LHL) [26]. However our anonymity definitions are adaptive, so one cannot apply the LHL (or even the generalized LHL [17]) directly. Rather, we first show an adaptive variant of the LHL is implied by the standard LHL via a hybrid argument. See the full version for details. Here we use it to prove the following theorem; details appear in the full version [7].

Theorem 5. *Let $\mathcal{AE}_d = (\mathcal{P}_d, \mathcal{K}_d, \mathcal{E}_d, \mathcal{D}_d)$ be a (deterministic) encryption scheme with message length $n(\cdot)$ and an associated universal-inducing (n, ℓ) -lossy key generator \mathcal{K}_l . Let A be an ANON adversary making $q(\cdot)$ **Enc** queries and a single **LR** query, all of these being $v(\cdot)$ -vector (μ, n) - m -block-sources. Then there exists LOS adversary B such that for all k*

$$\mathbf{Adv}_{\mathcal{AE}_d, A}^{\text{anon}}(k) \leq 2 \cdot \mathbf{Adv}_{\mathcal{AE}_d, B}^{\text{los}}(k) + 3 \cdot q(k) \cdot v(k) \cdot \sqrt{2^{n(k) - \ell(k) - \mu(k)}} .$$

B runs in time that of A . \square

Consider RtD and PtD when instantiated with a deterministic encryption scheme that is a u-LTDF. We can apply Theorem 5 to conclude ANON security for both schemes. Combining this with Theorems 2 and 4 yields proof of adaptive hedge security for RtD. Likewise, combining it with Theorems 3 and 4 yields proof of adaptive hedge security for PtD. Also Theorems 4 and 5 combine with [12, Th. 5.1] to give the first adaptively-secure deterministic encryption scheme (based on u-LTDFs).

REWH2 IS ADAPTIVELY SECURE. As we show above, we can get adaptive security from REWH when the underlying IND-CPA randomized scheme is anonymous in the sense of [4]. We observe that scheme REWH2 is adaptively secure when instantiated with *any* IND-CPA randomized scheme (not just anonymous ones). To show this, we give a direct proof in the full version [7]. Since popular encryption schemes such as RSA are not anonymous, we believe scheme REWH2 could be relevant in practice. That being said, we still think REWH1 is important since non-adaptive security is still a strong notion, and the scheme does not require any changes to the structure of the public key.

EXTENSIONS. In the full version [7] we discuss extensions and variants of RtD and PtD, where we improve the (adaptive) concrete security and show how to securely use LTDFs that are not necessarily universal.

Acknowledgements

We thank the Asiacrypt 2009 reviewers for detailed and thoughtful comments. Mihir Bellare and Thomas Ristenpart are supported by NSF grant CNS-0627779 and a gift from Intel Corporation. Moni Naor is Incumbent of the Judith Kleeman Professorial Chair. His research is supported in part by a grant from the Israel Science Foundation. Gil Segev is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. Hovav Shacham is supported in part by a MURI grant administered by the Air Force Office of Scientific Research. Scott Yilek is supported by NSF grant CNS-0831536.

References

1. M. Abdalla, M. Bellare, and P. Rogaway. The oracle diffie-hellman assumptions and an analysis of dhies. In *CT-RSA 2001*, volume 2020 of *LNCS*. Springer.
2. P. Abeni, L. Bello, and M. Bertacchini. Exploiting DSA-1571: How to break PFS in SSL with EDH, July 2008. http://www.lucianobello.com.ar/exploiting_DSA-1571/index.html.
3. O. Baudron, D. Pointcheval, and J. Stern. Extended notions of security for multi-cast public key cryptosystems. In *ICALP 2000*, volume 1853 of *LNCS*. Springer.
4. M. Bellare, A. Boldyreva, A. Desai, and D. Pointcheval. Key-privacy in public-key encryption. In *ASIACRYPT 2001*, volume 2248 of *LNCS*. Springer.
5. M. Bellare, A. Boldyreva, and S. Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In *EUROCRYPT 2000*, volume 1807 of *LNCS*. Springer.
6. M. Bellare, A. Boldyreva, and A. O'Neill. Deterministic and efficiently searchable encryption. In *CRYPTO 2007*, volume 4622 of *LNCS*. Springer.
7. M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged public-key encryption: How to protect against bad randomness. IACR ePrint Archive, 2009. Full Version of this paper.
8. M. Bellare, M. Fischlin, A. O'Neill, and T. Ristenpart. Deterministic encryption: Definitional equivalences and constructions without random oracles. In *CRYPTO 2008*, volume 5157 of *LNCS*. Springer.
9. M. Bellare and P. Rogaway. Code-based game-playing proofs and the security of triple encryption. In *EUROCRYPT 2006*, volume 4004 of *LNCS*. Springer.
10. M. Bellare and P. Rogaway. Optimal asymmetric encryption – how to encrypt with RSA. In *EUROCRYPT 1994*, volume 950 of *LNCS*. Springer.
11. M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudo random bits. In *FOCS 1982*. IEEE.
12. A. Boldyreva, S. Fehr, and A. O'Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In *CRYPTO 2008*, volume 5157 of *LNCS*. Springer.
13. D. Boneh. Simplified OAEP for the RSA and Rabin functions. In *CRYPTO 2001*, volume 2139 of *LNCS*. Springer.
14. C. Bosley and Y. Dodis. Does privacy require true randomness? In *TCC 2007*, volume 4392 of *LNCS*. Springer.
15. D. R. Brown. A weak randomizer attack on RSA-OAEP with $e=3$. IACR ePrint Archive, 2005.

16. Y. Dodis, S. J. Ong, M. Prabhakaran, and A. Sahai. On the (im)possibility of cryptography with imperfect randomness. In *FOCS 2004*. IEEE.
17. Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM Journal of Computing*, 38(1):97–139, 2008.
18. Y. Dodis and A. Smith. Entropic security and the encryption of high entropy messages. In *TCC 2005*, volume 3378 of *LNCS*. Springer.
19. L. Dorrendorf, Z. Gutterman, and B. Pinkas. Cryptanalysis of the windows random number generator. In *CCS 2007*. ACM.
20. E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. In *PKC 1999*, volume 1560 of *LNCS*. Springer.
21. T. E. Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *CRYPTO 1984*, volume 196 of *LNCS*. Springer.
22. I. Goldberg and D. Wagner. Randomness in the Netscape browser. *Dr. Dobbs's Journal*, January 1996.
23. S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
24. Z. Gutterman and D. Malkhi. Hold your sessions: An attack on Java session-id generation. In *CT-RSA 2005*, volume 3376 of *LNCS*. Springer.
25. Z. Gutterman, B. Pinkas, and T. Reinman. Analysis of the linux random number generator. In *IEEE Symposium on Security and Privacy*, pages 371–385, 2006.
26. R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *STOC 1989*. ACM.
27. S. Kamara and J. Katz. How to encrypt with a malicious random number generator. In *FSE 2008*, volume 5086 of *LNCS*. Springer.
28. J. L. McInnes and B. Pinkas. On the impossibility of private key cryptography with weakly random keys. In *CRYPTO 1990*, volume 537 of *LNCS*. Springer.
29. M. Mueller. Debian OpenSSL predictable PRNG bruteforce SSH exploit, May 2008. <http://milw0rm.com/exploits/5622>.
30. K. Ouafi and S. Vaudenay. Smashing SQUASH-0. In *EUROCRYPT 2009*, volume 5479 of *LNCS*. Springer.
31. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC 2008*. ACM.
32. P. Rogaway. Nonce-based symmetric encryption. In *FSE 2004*, volume 3017 of *LNCS*. Springer.
33. P. Rogaway and T. Shrimpton. Deterministic authenticated-encryption: A provable-security treatment of the key-wrap problem. In *EUROCRYPT 2006*, volume 4004 of *LNCS*. Springer.
34. A. Rosen and G. Segev. Efficient lossy trapdoor functions based on the composite residuosity assumption. *Cryptology ePrint Archive*, Report 2008/134, 2008.
35. B. Waters. Personal Communication to Hovav Shacham, December 2008.
36. S. Yilek, E. Rescorla, H. Shacham, B. Enright, and S. Savage. When private keys are public: Results from the 2008 Debian OpenSSL vulnerability. In *IMC 2009*. ACM. To appear.