Resisting Randomness Subversion: Fast Deterministic and Hedged Public-Key Encryption in the Standard Model

Mihir Bellare[1(✉)] and Viet Tung Hoang[2,3]

[1] Department of Computer Science and Engineering, University of California San Diego, San Diego, USA
mihir@eng.ucsd.edu
[2] Department of Computer Science, Georgetown University, Washington, DC, USA
[3] Department of Computer Science, University of Maryland, College Park, USA

Abstract. This paper provides the first *efficient*, *standard-model*, *fully-secure* schemes for some related and challenging forms of public-key encryption (PKE), namely deterministic and hedged PKE. These forms of PKE defend against subversion of random number generators, an end given new urgency by recent revelations on the nature and extent of such subversion.We resolve the (recognized) technical challenges in reaching these goals via a new paradigm that combines UCEs (universal computational extractors) with LTDFs (lossy trapdoor functions). Crucially, we rely only on a weak form of UCE, namely security for statistically (rather than computationally) unpredictable sources. We then define and achieve unique-ciphertext PKE as a way to defend against implementation subversion via algorithm-substitution attacks.

1 Introduction

Recent revelations about the prevalence of mass-surveillance and subversion raise new challenges for cryptography. This paper is concerned with subversion of public-key encryption (PKE). We first consider randomness-subversion attacks, namely ones that undermine randomness-generation processes. Forms of PKE resisting these have in fact already been defined, namely deterministic public-key encryption (D-PKE) [3] and hedged public-key encryption (H-PKE) [4]. However, good schemes —we mean efficient ones providing full security in the standard model— are not only lacking but a recognized challenge [53]. With the new impetus and urgency arising from the subversion perspective, we revisit these goals to provide such schemes. We achieve our ends via a new PKE paradigm in which universal computational extractors (UCEs) [8] —of the weaker ilk requiring only statistical rather than computational unpredictability— are combined with lossy trapdoor functions (LTDFs) [48].

We then turn to defending against subversion of encryption implementations via algorithm-substitution attacks [12,56]. Here we follow [12] to define the

E. Oswald and M. Fischlin (Eds.): EUROCRYPT 2015, Part II, LNCS 9057, pp. 627–656, 2015.
DOI: 10.1007/978-3-662-46803-6_21

new goal of unique ciphertext public-key encryption (U-PKE) and then reach it generically and efficiently from D-PKE.

Deterministic PKE. Technically, conceptually and historically, D-PKE is the core goal in this domain, and we begin there. The encryption algorithm of a D-PKE scheme takes public encryption key ek and message m to deterministically return a ciphertext c. We use the IND formalization of [6] which they show equivalent to the PRIV formalization of [3]. These formalizations capture the best possible privacy, namely semantic security for unpredictable messages that do not depend on the public key.

The core IND requirement asks for privacy when messages are individually unpredictable but may be arbitrarily correlated. We call this *full IND security* for emphasis. Full security is important in practice. For example, I might upload an encrypted file, then make a small edit to the file, re-encrypt and re-upload, so that the messages underlying the successive ciphertexts are very similar. It is thus the desired goal.

The EwH —encrypt with hash— D-PKE scheme of [3] encrypts message m under a (any) randomized IND-CPA scheme RE with the coins set to a hash of m. When the hash function is a random oracle, they showed EwH achieves full IND security. Achieving full IND security in the standard model however seemed out of reach. Many standard-model D-PKE schemes, using sophisticated techniques [6,11,17,19,30,33,49], have been proposed, but the security they achieve is not full. They only achieve security for *block sources*, where each message is assumed unpredictable *even given prior ones*, which is not realistic in practice.

The elusiveness of full security in the standard model was explained by Wichs [53], who showed that it could not be achieved under any single-stage assumption. To achieve full security one thus needs a multi-stage assumption. However most assumptions are single stage and it was not immediately clear what would even be a candidate for a suitable multi-stage assumption.

Such a candidate emerged with the UCE class of assumptions of security for hash functions of BHK1 [8]. The latter showed that the RO in EwH could be securely instantiated with a function family H that is $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$ —UCE-secure for *computationally* unpredictable sources— to yield a standard model, fully IND secure D-PKE scheme. Unfortunately, soon after, Brzuska, Farshim and Mittelbach (BFM) [21] showed that $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$-security is not achievable if indistinguishability obfuscation (iO) [2,34,35] is possible. BFM [21] and BHK1 [8] independently proposed to instead use $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$— UCE-security for *statistically* unpredictable sources. BFM [21] give some evidence that their attacks will not extend to $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ and that this assumption is weaker.

This raises several questions. Can one show that the scheme EwH is secure under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$? If not, can one provide a new, different D-PKE scheme that achieves full IND-security under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$?

Results for D-PKE. Our first result is negative. We show that if iO is possible then the RO in EwH is not universally instantiable. In more detail, given *any* family of functions H —in particular a $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ one— we build a (pathological and H-dependent) randomized PKE scheme RE such that (1) RE is IND-CPA

secure, but (2) An attack shows that the D-PKE scheme $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ given by the EwH transform is not IND-secure. The starting point is ideas of BFM [21], but several new ideas are needed, including several applications of a variable-output-length PRF to allocate randomness for the iO and a base PKE scheme in such a way that both (1) and (2) are possible. We note that the same negative result was obtained independently and concurrently by [22]. A general framework to obtain RO un-instantiability results via iO is given in [38] but it applies to single-stage games and thus doesn't yield a result for D-PKE.

Let H be a $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ function family. Then our negative result rules out showing an analogue of BHK1 [8], namely that $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ is fully IND secure for *any* IND-CPA RE. But there is a loophole, namely that the negative result does not preclude showing this for a *particular* choice of RE. We exploit this loophole to arrive at the desired goal of a fully IND secure D-PKE scheme under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$, as follows. We take the ROM BR93 PKE scheme [13], instantiate its trapdoor function with a *lossy* trapdoor function (LTDF) [32,48], and instantiate its RO with H, to get a standard-model PKE scheme RE. Next, we take the D-PKE scheme $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$, which has two uses of H, under two independent keys. Our D-PKE scheme DE1 is obtained by implementing these two uses of H with a single key. We prove that DE1 is fully IND secure assuming the LTDF is secure and H is $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$. We remark that using a single H key is important to prove security under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$, not just an efficiency optimization.

The connection of LTDFs to D-PKE was first made by Boldyreva, Fehr and O'Neill (BFO) [17]. Their LTDF-based D-PKE schemes however only achieve security for block sources, not full IND security. The block source restriction seems quite inherent in their methods, and indeed due to Wichs [53] we do not expect to achieve fully IND secure D-PKE using LTDFs alone. Our approach combines LTDFs with $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ to surmount this obstacle.

DE1 is the first D-PKE scheme that is fully IND secure in the standard model. Beyond that, however, it has the following important practical attributes: it is competitive on short messages, very fast on long messages, and supports variable-length messages directly. These practical attributes are a first for standard-model D-PKE schemes.

LTDFs and $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ are a productive and (in retrospect) natural match. Intuitively, LTDFs allow us to move to a game with information-theoretic guarantees, at which point it becomes possible to exploit UCE under statistical unpredictability. We view DE1 as a relatively simple illustration of the power of the UCE+LTDF method. H-PKE brings new challenges, which we surmount via non-trivial extensions of the basic method. We believe the UCE+LTDF method will have applications beyond this as well.

Hedged PKE. The encryption algorithm of a H-PKE scheme takes public encryption key ek, message m and randomness r to deterministically return a ciphertext c. The H-IND requirement of BBNRSS [4] has two parts: (1) standard IND-CPA security if r is good, meaning uniform and independent across encryptions, and (2) semantic security of m if the pair (m, r) is unpredictable

and does not depend on the public key. This second requirement is formalized as indistinguishability under chosen-distribution attack (IND-CDA) [4].

H-IND-secure PKE aims to provide the best possible privacy in the face of untrusted randomness. If the randomness is good, it does as well as standard IND-CPA encryption. But, whereas schemes providing only IND-CPA can fail spectacularly under poor randomness [4,20,46], H-IND PKE will not. It will compensate for poor randomness by also exploiting any available entropy in the message, protecting the latter as long as the message and randomness *together* are unpredictable. This is as good as it can get, since if the message-randomness pair is predictable, trial re-encryption on candidate pairs will recover the message underlying a target ciphertext. IND-CDA is an extension of IND that coincides with the latter if the randomness has no entropy at all.

In practice the most desirable form of IND-CDA is, again, full, meaning privacy when message-randomness pairs, although individually unpredictable, may be arbitrarily correlated. By full H-IND, we mean IND-CPA plus full IND-CDA. In the ROM, fully H-IND PKE is achieved by an extension of EwH called REwH that encrypts m under an IND-CPA scheme with the coins set to the hash of $m \,\|\, r$ [4]. In the standard model, things are more difficult. Providing a fully IND-CDA PKE scheme is harder than providing a fully IND D-PKE scheme because the unpredictability pertains to (m, r) not just m and also, more importantly, because IND-CDA is formalized in [4] as an *adaptive* requirement. Additionally, while IND-CPA is easy in isolation, it is *not* in combination with IND-CDA. The reason is subtle, namely that IND-CDA breaks when m depends on the public key, but IND-CPA must remain secure in this case. This butting of heads of the IND-CPA and IND-CDA conditions doubles the challenge of achieving fully H-IND PKE compared to fully IND D-PKE.

These technical difficulties are reflected in the landscape of standard-model schemes, where fully H-IND PKE has not been achieved under *any* assumption. BBNRSS [4] build standard-model H-IND PKE schemes by composition of standard-model D-PKE and IND-CPA schemes, and also directly via anonymous LTDFs, but these schemes achieve IND-CDA only for block sources. (The latter now means that message-randomness pairs are assumed to be unpredictable even given prior ones.) It is instructive that full H-IND PKE has not even been achieved under $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$. To elaborate, recall that BHK1 [8] showed that $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$-instantiating the RO in EwH results in a fully IND secure standard-model D-PKE scheme. We can correspondingly $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$-instantiate the RO in REwH. But, even if the resulting scheme can be shown fully IND-CDA, there seems no reason it is IND-CPA. The reason is the difficulty alluded to above. Namely, a UCE hash function may not provide security on messages that are a function of the hashing key, but the latter is part of the public key and IND-CPA requires security for messages depending on the public key.

But the bar for us is even higher: due to the BFM attacks [21] on $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$, we want to use the weaker $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ assumption, just as we did for $\mathsf{DE1}$. We thus face at least two difficulties. The first is to achieve full IND-CDA under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$. Here the main challenge is handling adaptivity. But beyond that

the fundamental above-mentioned difficulty of achieving IND-CPA in the same scheme remains, because no form of UCE guarantees security for messages that depend on the hashing key.

Results for H-PKE. We surmount the technical difficulties discussed above to provide the first standard-model, fully H-IND PKE schemes. We specify three schemes, HE1, HE2 and HE3. All efficiently achieve our security goals, the second and third handle variable-length messages, and the third further adds better concrete security.

Recall that we obtained DE1 as EwH[H, BR93[LT, H]], where H is UCE[$\mathcal{S}^{\text{sup}}$] and LT is a LTDF. A natural idea is to similarly get H-PKE as REwH[H, BR93[LT, H]]. (In both cases we use one hash key rather than two.) We are able to show this achieves full IND-CDA. This is significant since handling adaptivity required anonymous LTDFs in [4] which we do not need. But we then hit the problem above, namely UCE[$\mathcal{S}^{\text{sup}}$] security of H may not be enough to provide IND-CPA. We resolve this by building a *particular*, suitable UCE[$\mathcal{S}^{\text{sup}}$] family H. We first build a particular family U of AU (almost universal) hash functions and then obtain H by applying the AU-then-Hash transform of BHK2 [9] to a fixed-input-length UCE[$\mathcal{S}^{\text{sup}}$] family $\overline{\mathsf{H}}$ and our U. We refer to the resulting PKE scheme as HE1. We are able to show that it is full IND-CDA as well as IND-CPA assuming UCE[$\mathcal{S}^{\text{sup}}$] security of $\overline{\mathsf{H}}$ and security of the LTDF.

This achieves, for the first time, the security goal of fully H-IND PKE in the standard model, which we consider already significant. But in terms of practicality, HE1 is not ideal because it can only handle fixed-length messages. HE2 efficiently encrypts variable and arbitrary length messages while retaining full H-IND security. It uses a variable-output-length PRF in addition to the primitives used by HE1. Finally, HE3 exploits some combinatorial techniques to obtain better security bounds, as a result of which it offers security for lower values of the message min-entropy than the other schemes.

Speed. Our D-PKE and H-PKE schemes are the first to achieve full security in the standard model, which we believe is a significant theoretical contribution. However, beyond that, they have important practical attributes, expanded on below and in Section 5.

It is well known that asymmetric primitives are orders of magnitude less efficient than symmetric ones. Central to making standard IND-CPA encryption efficient is hybrid encryption as represented by the KEM-DEM paradigm [25]. Encryption generates a random asymmetrically-protected per-message symmetric key and then symmetrically encrypts the message under the latter, leading to cheap encryption of long messages. But for standard model D-PKE and H-PKE the hybrid encryption paradigm breaks down, because, with the constraint of being deterministic or not trusting the randomness, it is not clear how to even pick the per-message key. This difficulty is recognized and seems quite fundamental and hard to bypass. As a result, prior standard-model D-PKE and H-PKE schemes fix the message length and rely only on asymmetric operations. Their cost in asymmetric operations becomes exorbitant on long messages and they also cannot encrypt variable-length messages.

Our methods break these efficiency bottlenecks to recover hybrid-encryption like performance. Our DE1, HE2 and HE3 schemes handle messages of variable and arbitrary length, and the asymmetric cost is fixed independent of the message length, so that we pay only in hashing as the message length grows. Placing us in a particularly good position to exploit this is the speed of $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ functions. Direct constructions based on HMAC-SHA-256 [8,45] are already efficient, but in fact still more efficient and even parallelizable constructions are given in BHK2 [9], along with software implementations and cost comparisons. Meanwhile LTDFs can be efficiently instantiated in a variety of ways [32,40,43,48], making the asymmetric component competitive. This leads overall to performance comparable to existing IND-CPA schemes while providing protection against randomness subversion.

In practice concrete security is important to know how to set parameters. Good bounds are important so that one may use smaller parameters. (The cost of the asymmetric operations is usually cubic in the key length so cutting the latter by one-half yields a factor eight speedup.) For this reason we not only state in our theorems the concrete security bounds of the reductions but also try to obtain good ones.

Unique-Ciphertext PKE. In an algorithm-substitution attack (ASA) [12,56], the prescribed encryption algorithm is replaced with a malicious one that may attempt to leak information about the message to "big brother" based on a shared key. BPR [12] formalize the attacker goal in an ASA as compromising privacy without detection. BPR [12] and ACMPS [1] indicate that randomized encryption will be subject to successful attack. In the symmetric setting, BPR [12] show that ASAs can be protected against by a form of deterministic encryption they call unique-ciphertext symmetric encryption.

We analogously define unique-ciphertext PKE. U-PKE requires that for every key pair (ek, dk) and message m, there is at most one ciphertext c that decrypts to m under dk. A U-PKE scheme is thus deterministic, but not every D-PKE scheme is U-PKE. For example, appending to a D-PKE ciphertext a zero bit ignored by decryption leaves D-PKE intact but violates U-PKE. In Section 6 we show however how to achieve U-PKE in a simple and generic way from D-PKE. Combining this with our efficient D-PKE scheme above yields efficient U-PKE, allowing us to better defend against ASAs.

Discussion and Related Work. In a world of subversion, there are no panaceas. As with BPR [12], our goals are deliberately restricted in scope. We aim to provide better (not perfect) security in the face of some (not all) subversion threats. Thus, we restrict attention to randomness-subversion attacks and algorithm-substitution attacks. We assume that key-generation, being one-time, can leverage good randomness.

We might view IND-CPA as the optimistic view (the randomness is excellent, use it), D-PKE as the pessimistic view (the randomness may be bad so, to be safe, ignore it) and H-IND PKE as the pragmatic view (I don't know how good the randomness is but I will just get the best out of it that I can). We would expect the extent and nature of randomness subversion to vary rather than be

ubiquitous and total, in part because subversion will aim to evade detection. In this light H-IND PKE emerges as the best defense in the face of randomness subversion.

Failures of randomness-generation processes [24,28,29,39,41,44] have in the past been attributed to error. Now we know better, namely that some should be attributed to subversion. This makes practical defenses more urgent and increases the motivation for work like ours that delivers such defenses.

At SXSW 2014, Snowden said "... we know that the encryption algorithms we are using today work ... it is the random number generators that are attacked as opposed to the encryption algorithms themselves ... ". We aim, in some sense, to turn this on its head. We suggest that the encryption algorithms *don't* work because they are not robust in the face of poor randomness. We pursue practical hedged encryption as a counter-measure.

We do not expect or aim to maintain, under subversion, the high level of security we can achieve in its absence. Security will unavoidably degrade. Our goal with H-IND PKE is for it to degrade as little as possible rather than disappear. This philosophy sets us apart from most of the related work on randomness subversion we will discuss in the next paragraph, which either aims to understand under what limitations on the class of attacks one can achieve the same security one would under perfect randomness, or shows that such security is not possible.

Yilek [55] studies randomness-reset attacks, where the randomness is uniform but the adversary can force its re-use across different encryptions. Paterson, Schuldt and Sibborn [47] introduce related-randomness attacks, where encryption is under adversary-specified functions of some initial uniform randomness, providing negative results, as well as positive results for some classes of attacks. Birrell, Chung, Pass and Telang [15] and Hemenway and Ostrovsky [40] study the encryption of randomness-dependent messages. Austrin, Chung, Mahmoody, Pass and Seth [1] show that encryption is insecure under even quite weak adversarial tampering of randomness. Authenticated key-exchange with bad randomness is studied in [31,54]. Negative results for cryptography with imperfect randomness are provided by [18,26,27]. Kamara and Katz [42] study symmetric encryption providing semantic security under good coins in the face of chosen-plaintext attacks involving bad coins.

Ristenpart and Yilek [50] study the use of H-IND PKE in real systems. Brakerski and Segev [19] study D-PKE security in the presence of auxiliary information about messages. Raghunathan, Segev and Vadhan [49] study security of D-PKE when the message may depend on the public key. Vergnaud and Xiao [52] study IND-CDA when the message and randomness may depend on the public key. In the symmetric setting, Rogaway and Shrimpton's misuse-resistant authenticated encryption [51] represents a form of hedging.

2 Preliminaries

We review basic notation and definitions including games, function families, VOL PRFs, LTDFs and UCE.

By $\lambda \in \mathbb{N}$ we denote the security parameter and by 1^λ its unary representation. We denote the number of coordinates of a vector $\mathbf{x}$ by $|\mathbf{x}|$, and the length of a string $x \in \{0,1\}^*$ by $|x|$. Algorithms are randomized unless otherwise indicated. Running time is worst case. "PT" stands for "polynomial-time," whether for randomized algorithms or deterministic ones. If A is an algorithm, we let $y \leftarrow A(x_1, \ldots; r)$ denote running A with randomness r on inputs $x_1, \ldots$ and assigning the output to y. We let $y \leftarrow\!\!\$\ A(x_1, \ldots)$ be the resulting of picking r at random and letting $y \leftarrow A(x_1, \ldots; r)$. We let $[A(x_1, \ldots)]$ denote the set of all possible outputs of A when invoked with inputs $x_1, \ldots$.

We use the code based game playing framework of [14]. (See Fig. 1 for an example.) By $\mathrm{G}^A(\lambda)$ we denote the event that the execution of game G with adversary A and security parameter λ results in output true, the game output being what is returned by Game.

For concrete security assessments, we adopt the notation of [10]. Let the *number of queries* of A to an oracle Proc be the function $\mathbf{Q}_A^{\mathrm{Proc}}$ that on input λ returns the maximum number of queries that A makes to Proc when executed with security parameter λ, the maximum over all coins and all possible replies to queries to all oracles of A. Time assessments are simplified by the convention that running time is that of the game rather than merely the adversary, and we let $\mathbf{T}(\mathrm{G}^{A_1, A_2, \cdots})$ denote the function of λ that returns the maximum execution time of game G with adversaries $A_1, A_2, \ldots$ and security parameter λ, the maximum over all coins, and the time being all inclusive, meaning the time taken by game procedures to compute replies is included.

Function Families. Our syntax for function families follows [8], in particular allowing variable output lengths. This is important in our applications to encrypt messages of variable length, which in turn is important in practice. A family of functions H specifies the following. On input the unary representation 1^λ of the security parameter $\lambda \in \mathbb{N}$, key generation algorithm $\mathsf{H.Kg}$ returns a key $hk \in \{0,1\}^{\mathsf{H.kl}(\lambda)}$, where $\mathsf{H.kl}\colon \mathbb{N} \to \mathbb{N}$ is the key length function associated to H. The deterministic, PT evaluation algorithm $\mathsf{H.Ev}$ takes 1^λ, key hk an input $x \in \{0,1\}^*$ with $|x| \in \mathsf{H.IL}(\lambda)$, and a unary encoding 1^ℓ of an output length $\ell \in \mathsf{H.OL}(\lambda)$ to return $\mathsf{H.Ev}(1^\lambda, hk, x, 1^\ell) \in \{0,1\}^\ell$. Here $\mathsf{H.IL}$ is the input-length function associated to H, so that $\mathsf{H.IL}(\lambda) \subseteq \mathbb{N}$ is the set of allowed input lengths, and similarly $\mathsf{H.OL}$ is the output-length function associated to H, so that $\mathsf{H.OL}(\lambda) \subseteq \mathbb{N}$ is the set of allowed output lengths. The latter allows us to cover functions of variable output length. If H has fixed input length then let $\mathsf{H.il}$ denote the function such that $\mathsf{H.IL}(\lambda) = \{\mathsf{H.il}(\lambda)\}$ for every $\lambda \in \mathbb{N}$. If H has fixed output length, define $\mathsf{H.ol}$ likewise.

Variable Output Length PRFs. A variable output length (VOL) PRF is a function family F such that $\mathsf{F.Kg}$ returns a uniformly distributed key in $\{0,1\}^{\mathsf{F.kl}}$ and $\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{F},A}(\lambda) = 2\Pr[\mathrm{PRF}^A_{\mathsf{F}}(\lambda)] - 1$ is negligible for every PT adversary A, where

$\underline{\text{GAME CPA}^A_{\mathsf{PKE}}(\lambda)}$	$\underline{\text{GAME PRF}^A_{\mathsf{F}}(\lambda)}$	$\underline{\text{GAME Lossy}^A_{\mathsf{LT}}(\lambda)}$
$(ek, dk) \leftarrow\!\!\$\ \mathsf{PKE.Kg}(1^\lambda)$	$b \leftarrow\!\!\$\ \{0,1\}$; $fk \leftarrow\!\!\$\ \{0,1\}^{\mathsf{F.kl}(\lambda)}$	$(ek, dk) \leftarrow\!\!\$\ \mathsf{LT.EKg}(1^\lambda)$
$b \leftarrow\!\!\$\ \{0,1\}$	$b' \leftarrow\!\!\$\ A^{\text{RR}}(1^\lambda)$	$lk \leftarrow\!\!\$\ \mathsf{LT.LKg}(1^\lambda)$
$(m_0, m_1, \_) \leftarrow\!\!\$\ A(1^\lambda, ek)$	Return $(b = b')$	$b \leftarrow\!\!\$\ \{0,1\}$
$c \leftarrow\!\!\$\ \mathsf{PKE.Enc}(ek, m_b)$		If $b = 1$ then $K \leftarrow ek$
$b' \leftarrow\!\!\$\ A(1^\lambda, t, c)$	$\underline{\text{RR}(x, 1^\ell)}$	Else $K \leftarrow lk$
Return $(b = b')$	If $b = 1$ then	$b' \leftarrow A(1^\lambda, K)$
	$y \leftarrow \mathsf{F.Ev}(1^\lambda, fk, x, 1^\ell)$	Return $(b' = b)$
	Else $y \leftarrow\!\!\$\ \{0,1\}^\ell$	
	Return y	

Fig. 1. Left: Game CPA defining IND-CPA security of a PKE scheme PKE. **Middle:** Game PRF defining the PRF security of a variable-output-length function family F. **Right:** Game Lossy defining the security of a lossy trapdoor function LT.

game $\text{PRF}^A_{\mathsf{F}}$ is defined in the middle panel of Fig. 1. In this game the adversary is given an oracle RR that either implements a random oracle or $\mathsf{F.Ev}(1^\lambda, fk, \cdot, \cdot)$, where $fk \leftarrow\!\!\$\ \{0,1\}^{\mathsf{F.kl}(\lambda)}$ is a random key. We assume that A doesn't repeat a prior RR query, and any RR query $(x, 1^\ell)$ must satisfy $x \in \mathsf{F.IL}(\lambda)$ and $\ell \in \mathsf{F.OL}(\lambda)$. This extends [36] to VOL families. A practical construction of a VOL PRF from a blockcipher is given in [16].

Public-Key Encryption. A PKE scheme PKE defines PT algorithms PKE.Kg, PKE.Enc, PKE.Dec, the last deterministic. Algorithm PKE.Kg takes as input 1^λ and outputs a public encryption key $ek \in \{0,1\}^{\mathsf{PKE.ekl}(\lambda)}$ and a secret decryption key dk, where $\mathsf{PKE.ekl}\colon \mathbb{N} \to \mathbb{N}$ is the public-key length of PKE. Algorithm PKE.Enc takes as input $1^\lambda, ek$ and a message m with $|m| \in \mathsf{PKE.IL}(\lambda)$ to return a ciphertext c, where PKE.IL is the input-length function associated to PKE, so that $\mathsf{PKE.IL}(\lambda) \subseteq \mathbb{N}$ is the set of allowed input (message) lengths. Algorithm PKE.Dec takes $1^\lambda, dk, c$ and outputs $m \in \{0,1\}^* \cup \{\perp\}$. Correctness requires that $\mathsf{PKE.Dec}(1^\lambda, dk, c) = m$ for all $\lambda \in \mathbb{N}$, all $(ek, dk) \in [\mathsf{PKE.Kg}(1^\lambda)]$ all m with $|m| \in \mathsf{PKE.IL}(\lambda)$ and all $c \in [\mathsf{PKE.Enc}(1^\lambda, ek, m)]$. Scheme PKE is IND-CPA secure [37] if $\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{PKE},A}(\lambda) = 2[\text{CPA}^A_{\mathsf{PKE}}(\lambda)] - 1$ is negligible for every PT adversary A, where game CPA is defined in the left panel of Fig. 1. We require that the messages m_0, m_1 output by A have the same length $|m_0| = |m_1| \in \mathsf{PKE.IL}(\lambda)$. Let $\mathsf{PKE.rl}\colon \mathbb{N} \to \mathbb{N}$ denote the randomness-length function of PKE, meaning $\mathsf{PKE.Enc}(1^\lambda, \cdot, \cdot)$ draws its coins at random from $\{0,1\}^{\mathsf{PKE.rl}(\lambda)}$. We say that PKE has input length $\mathsf{PKE.il}\colon \mathbb{N} \to \mathbb{N}$ if $\mathsf{PKE.IL}(\lambda) = \{\mathsf{PKE.il}(\lambda)\}$ for all $\lambda \in \mathbb{N}$, and refer to this as a PKE scheme that only allows fixed length messages. Our goal will be to allow variable and arbitrary-length messages, ideally $\mathsf{PKE.IL}(\cdot) = \mathbb{N}$, but at least some large subset thereof.

Lossy Trapdoor Functions. A lossy trapdoor function [48] LT specifies PT algorithms LT.EKg, LT.LKg, LT.Ev, LT.Inv, the last two deterministic, as well as an input length $\mathsf{LT.il}\colon \mathbb{N} \to \mathbb{N}$ and an output length $\mathsf{LT.ol}\colon \mathbb{N} \to \mathbb{N}$. Key-generation algorithm LT.EKg takes 1^λ and returns an "injective" key ek and a

$\underline{\text{GAME UCE}^{S,D}_{\mathsf{H}}(\lambda)}$	$\underline{\text{HASH}(x, 1^\ell)}$

$b \leftarrow\!\$\ \{0,1\}$; $hk \leftarrow\!\$\ \mathsf{H.Kg}(1^\lambda)$
$L \leftarrow\!\$\ S^{\text{HASH}}(1^\lambda)$; $b' \leftarrow\!\$\ D(1^\lambda, hk, L)$
Return $(b' = b)$
	If $T[x,\ell] = \perp$ then
	  If $b = 0$ then $T[x,\ell] \leftarrow\!\$\ \{0,1\}^\ell$
	  Else $T[x,\ell] \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^\ell)$
	Return $T[x,\ell]$

$\underline{\text{GAME Pred}^{P}_{S}(\lambda)}$	$\underline{\text{HASH}(x, 1^\ell)}$
$Q \leftarrow \emptyset$; $L \leftarrow\!\$\ S^{\text{HASH}}(1^\lambda)$; $Q' \leftarrow\!\$\ P(1^\lambda, L)$ Return $(Q' \cap Q \neq \emptyset)$	If $T[x,\ell] = \perp$ then $T[x,\ell] \leftarrow\!\$\ \{0,1\}^\ell$ $Q \leftarrow Q \cup \{x\}$; Return $T[x,\ell]$
$\underline{\text{GAME Reset}^{R}_{S}(\lambda)}$	$\underline{\text{HASH}(x, 1^\ell)}$

$\mathsf{Dom} \leftarrow \emptyset$; $L \leftarrow\!\$\ S^{\text{HASH}}(1^\lambda)$; $b \leftarrow\!\$\ \{0,1\}$
If $b = 0$ then // reset the array T
  For all $(x,\ell) \in \mathsf{Dom}$ do
  $T[x,\ell] \leftarrow\!\$\ \{0,1\}^\ell$
$b' \leftarrow R^{\text{HASH}}(1^\lambda, L)$; Return $(b' = b)$
	If $T[x,\ell] = \perp$ then $T[x,\ell] \leftarrow\!\$\ \{0,1\}^\ell$
	$\mathsf{Dom} \leftarrow \mathsf{Dom} \cup \{(x,\ell)\}$; Return $T[x,\ell]$

Fig. 2. Games UCE **(top),** Pred **(middle), and** Reset **(bottom) to define UCE security**

decryption key dk. Evaluation algorithm $\mathsf{LT.Ev}$ takes $1^\lambda, ek$ and $x \in \{0,1\}^{\mathsf{LT.il}(\lambda)}$ to return an $\mathsf{LT.ol}(\lambda)$-bit string. Inversion algorithm $\mathsf{LT.Inv}$ takes $1^\lambda, dk$ and $y \in \{0,1\}^{\mathsf{LT.ol}(\lambda)}$ to return a $\mathsf{LT.il}(\lambda)$-bit string. The *correctness requirement* demands that $\mathsf{LT.Inv}(1^\lambda, dk, \mathsf{LT.Ev}(1^\lambda, ek, x)) = x$ for every $\lambda \in \mathbb{N}$, every $(ek, dk) \in [\mathsf{LT.EKg}(1^\lambda)]$ and every $x \in \{0,1\}^{\mathsf{LT.il}(\lambda)}$. Algorithm $\mathsf{LT.LKg}$, given 1^λ, returns a "lossy" key lk. Let $\tau : \mathbb{N} \to \mathbb{N}$ be a function such that $2^{-\tau(\cdot)}$ is negligible. We say that LT is τ-lossy if the size of the set $\{\mathsf{LT.Ev}(1^\lambda, lk, x) \mid x \in \{0,1\}^{\mathsf{LT.il}(\lambda)}\}$ is at most $2^{\mathsf{LT.il}(\lambda)-\tau(\lambda)}$ for every $\lambda \in \mathbb{N}$ and every $lk \in [\mathsf{LT.LKg}(1^\lambda)]$. Security of an LTDF demands two things. First, lossy and injective keys are indistinguishable. Formally, $\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},A}(\lambda) = 2\Pr[\text{Lossy}^{A}_{\mathsf{LT}}(\cdot)] - 1$ must be negligible for every PT adversary A, where game Lossy is defined in the right panel of Fig. 1. Second, LTDF is τ-lossy for some τ such that $2^{-\tau(\cdot)}$ is negligible. To simplify concrete security analyses, we assume that $\mathsf{LT.LKg}$'s worst-case running time is at most that of $\mathsf{LT.EKg}$.

There are by now many constructions of LTDFs known [32,40,43,48]. As an example, RSA is shown to be lossy [43] under the Φ-hiding assumption of [23]. For a 2048-bit modulus, one may choose $\tau = 430$ for 80-bit security.

UCE. We recall the Universal Computational Extractor (UCE) framework of BHK1 [8]. Let H be a family of functions as defined above. Let S be an adversary called the *source* and D an adversary called the *distinguisher*. We associate to them and H the game $\text{UCE}^{S,D}_{\mathsf{H}}(\lambda)$ at the left panel of Fig. 2. The source has access to an oracle HASH and we require that any query $(x, 1^\ell)$ made to this oracle satisfy $|x| \in \mathsf{H.IL}(\lambda)$ and $\ell \in \mathsf{H.OL}(\lambda)$. When the challenge bit b is 1 (the "real" case) the oracle responds via $\mathsf{H.Ev}$ under a key hk that is chosen by the game and *not* given to the source. When $b = 0$ (the "random" case) it responds as a random oracle. The source then leaks a string L to its accomplice distinguisher.

The latter *does* get the key hk as input and must now return its guess $b' \in \{0,1\}$ for b. The game returns true iff $b' = b$, and the uce-advantage of (S, D) is defined for $\lambda \in \mathbb{N}$ via $\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\lambda) = 2\Pr[\mathrm{UCE}^{S,D}_{\mathsf{H}}(\lambda)] - 1$. If $\mathcal{S}$ is a class (set) of sources, we say that H is $\mathsf{UCE}[\mathcal{S}]$-secure if $\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\cdot)$ is negligible for all sources $S \in \mathcal{S}$ and all PT distinguishers D. Trivial attacks from [8] show that $\mathsf{UCE}[\mathcal{S}]$-security is not achievable if $\mathcal{S}$ is the class of all PT sources. To obtain meaningful notions of security, BHK1 [8] impose restrictions on the source. There are many ways to do this; below we'll focus on what they call statistically unpredictable and reset-secure sources.

A source is unpredictable if it is hard to guess the source's Hash queries even given the leakage, in the *random case* of UCE game. Formally, let S be a source and P an adversary called a predictor. Consider game $\mathrm{Pred}^{P}_{S}(\lambda)$ in the middle panel of Fig. 2. Given the leakage, P outputs a set Q'; we require that $|Q'|$ is polynomially bounded. The predictor wins if this set contains a Hash-query of the source. For $\lambda \in \mathbb{N}$ we let $\mathsf{Adv}^{\mathsf{pred}}_{S,P}(\lambda) = \Pr[\mathrm{Pred}^{P}_{S}(\lambda)]$. We say that S is statistically unpredictable if $\mathsf{Adv}^{\mathsf{pred}}_{S,P}(\cdot)$ is negligible for all (even computationally unbounded) predictors P. We say that H is $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$-secure if $\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\cdot)$ is negligible for all statistically unpredictable PT sources and all PT distinguishers.

The second restriction on sources from [8] is reset security. Let S be a source and R an adversary called a reset adversary. The source again is executed with its Hash being a random oracle. The reset adversary is either given access to the same random oracle or to an *independent* one. The requirement is that it should not be able to tell which. Consider game $\mathrm{Reset}^{R}_{S}(\lambda)$ at the right panel of Fig. 2; we require that R make only polynomial number of queries to Hash. For $\lambda \in \mathbb{N}$ we let $\mathsf{Adv}^{\mathsf{reset}}_{S,R}(\lambda) = 2\Pr[\mathrm{Reset}^{R}_{S}(\lambda)] - 1$. We say S is statistically reset-secure if $\mathsf{Adv}^{\mathsf{reset}}_{S,R}(\cdot)$ is negligible for all reset adversaries R. We say that H is $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$-secure if $\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\cdot)$ is negligible for all statistically reset-secure PT sources and all PT distinguishers.

BHK1 [8] show that $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$-security of H implies $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$-security of H. BFM [21] show that if indistinguishability obfuscation for all circuits is possible then $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$ —UCE for *computationally* unpredictable sources— is not achievable in the standard model. However $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ and $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ are not subject to their attack and emerge as weaker and plausible assumptions. Moving to the statistical versions was independently suggested by BHK1 [8] and BFM [21]. These statistical assumptions will be the basis of our constructs.

While $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ and $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ may seem like strong assumptions, we know that multi-stage assumptions are necessary to reach our goals [53]. There are very few candidate multi-stage assumptions and amongst them the ones we use are the more plausible.

$\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ and $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ families may be efficiently instantiated via HMAC-SHA-256 [8,45] or super-efficiently via [9], which we will exploit for efficient schemes.

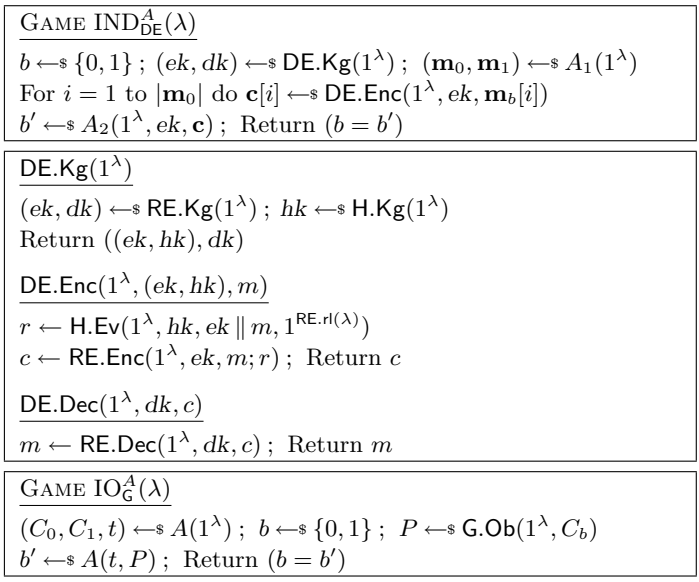

Fig. 3. **Top:** Game defining IND security of D-PKE scheme DE. **Middle:** D-PKE scheme $\mathsf{DE} = \mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$. **Bottom:** Game defining iO security of an indistinguishability obfuscator G.

3 Efficient, Fully IND Secure D-PKE

This section begins with a negative result —that assuming iO the random oracle (RO) in EwH is not *universally* instantiable— and then provides a complementary positive result —that there is a *particular* instantiation of the RO and IND-CPA scheme in EwH that results in a fully IND secure D-PKE scheme. The latter, which is the main result of this section, showcases our UCE+LTDF method and brings a new D-PKE scheme with two attributes: (1) On the theoretical front, it is the first D-PKE scheme shown *fully* IND secure in the standard model, and (2) On the practical front, it encrypts variable-input length messages and achieves hybrid-encryption like efficiency on long messages.

D-PKE and EwH. We say that a PKE scheme DE is a deterministic public-key encryption (D-PKE) [3] if the encryption algorithm $\mathsf{DE.Enc}$ is deterministic. We use the IND formalization of security of BFOR [6], which they show equivalent to the PRIV formalization of [3]. Game IND defining the IND notion is shown in the left panel of Fig. 3. An IND adversary $A = (A_1, A_2)$ is a pair of PT algorithms, where A_1 on input 1^λ returns a pair of message vectors $(\mathbf{m}_0, \mathbf{m}_1)$. We require that (i) there be a polynomial v such that $|\mathbf{m}_0| = |\mathbf{m}_1| \leq v(\lambda)$ and $|\mathbf{m}_0[i]| = |\mathbf{m}_1[i]| \in \mathsf{DE.IL}(\lambda)$, for every $i \leq |\mathbf{m}_0|$, and (ii) messages $\mathbf{m}_0[1], \ldots, \mathbf{m}_0[|\mathbf{m}_0|]$ are distinct and also messages $\mathbf{m}_1[1], \ldots, \mathbf{m}_1[|\mathbf{m}_1|]$ are distinct. The guessing probability $\mathrm{Guess}_A(\cdot)$ of A is the function that on input $\lambda \in \mathbb{N}$ returns the maximum, over all b, m, i, of $\Pr[\mathbf{m}_b[i] = m]$, the probability over $(\mathbf{m}_0, \mathbf{m}_1) \leftarrow\$ A_1(1^\lambda)$. We say

that A has *high min-entropy* if $\mathrm{Guess}_A(\cdot)$ is negligible. We let $\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{DE},A}(\lambda) = 2\Pr[\mathrm{IND}^A_{\mathsf{DE}}(\lambda)] - 1$ and say that DE is IND-secure if $\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{DE},A}(\cdot)$ is negligible for all PT A of high min-entropy.

We stress that this definition captures *full* security because the messages in the message vectors may be arbitrarily correlated. This is what is needed in practice. In contrast, security for block sources [17] requires that each message in each vector has high min entropy even given prior ones. This is often not true in practice and security only for block sources is quite weak, yet prior standard-model schemes have only been able to achieve this.

EwH [3] is a simple and natural transform that takes a family of functions H and a randomized PKE scheme RE to return the D-PKE scheme $\mathsf{DE} = \mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ whose algorithms are shown in the middle panel of Fig. 3. We let $\mathsf{DE.IL} = \mathsf{RE.IL}$. We require that $\mathsf{RE.rl}(\lambda) \in \mathsf{H.OL}(\lambda)$ and $\mathsf{RE.ekl}(\lambda) + \ell \in \mathsf{H.IL}(\lambda)$ for all $\lambda \in \mathbb{N}$ and all $\ell \in \mathsf{RE.IL}(\lambda)$.

Indistinguishability Obfuscation. We recall the definition of [34], which extends that of [2] to allow auxiliary information. We say that circuits C_0 and C_1 are *functionally equivalent*, denoted $C_0 \equiv C_1$, if they have the same size, the same number n of inputs, and $C_0(x) = C_1(x)$ for every input $x \in \{0,1\}^n$. An indistinguishability obfuscator (iO) G defines PT algorithms $\mathsf{G.Ob}, \mathsf{G.Ev}$ and a randomness length function $\mathsf{G.rl}\colon \mathbb{N} \to \mathbb{N}$. Algorithm $\mathsf{G.Ob}$ takes as input 1^λ and a circuit C, and outputs a string P using randomness of length $\mathsf{G.rl}(\lambda)$. Deterministic algorithm $\mathsf{G.Ev}$ takes as input strings P, x and returns $y \in \{0,1\}^* \cup \{\perp\}$. We require that for any circuit C, any input x for C any $\lambda \in \mathbb{N}$, and any $P \in [\mathsf{G.Ob}(1^\lambda, C)]$, it holds that $\mathsf{G.Ev}(P, x) = C(x)$. An adversary A is *well-formed* if $\Pr[C_0 \not\equiv C_1 \colon (C_0, C_1, t) \leftarrow\!\!\$\ A(1^\lambda)]$ is negligible. We say that G is iO-secure if $\mathsf{Adv}^{\mathsf{io}}_{\mathsf{G},A}(\lambda) = 2\Pr[\mathrm{IO}^A_{\mathsf{G}}(\lambda)] - 1$ is negligible for every PT well-formed adversary A, where game IO is defined at the right panel of Fig. 3.

Implausibility of Universal Instantiation of EwH. BBO [3] showed that if H is implemented via a RO then $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ is IND-secure for *any* IND-CPA RE. A basic theoretical and practical question is whether the RO in this result can be securely instantiated. The most desirable instantiation is *universal*, by which we mean there is a function family H such that $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ is IND-secure for any IND-CPA RE. Here we show that if iO exists then there is no such universal instantiation. Given any function family H we build an IND-CPA PKE scheme RE such that $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ is not IND-secure. We stress that this does not preclude providing specific H, RE such that $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ is IND-secure, and indeed it is in this way that we will later obtain our positive result.

Our findings strengthen, and are consistent with, prior work. BHK1 [8] showed that a $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$ family will provide a universal instantiation of EwH, but $\mathsf{UCE}[\mathcal{S}^{\mathrm{cup}}]$ is ruled out under iO by BFM [21], so there is no contradiction. However, following BFM, it remained possible that some other class of function families might be able to universally instantiate EwH. Under iO, we rule this out.

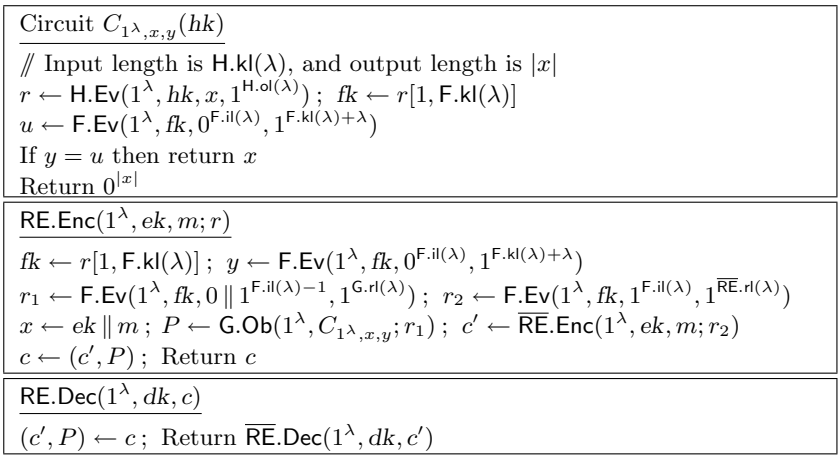

Fig. 4. Middle, Bottom: Encryption and decryption algorithm of the counter-example PKE scheme RE for Proposition 1. **Top:** Circuit constructed and obfuscated in RE.Enc.

We let H be a function family with input length H.il and output length H.ol. We will build the counter-example PKE scheme RE from H and the following auxiliary primitives: an arbitrary, base IND-CPA scheme $\overline{\mathsf{RE}}$, a VOL PRF F and an iO scheme G. The result is as follows.

Proposition 1. *Let* H *be a function family with input length* H.il *and output length* H.ol. *Let* F *be a VOL PRF with* $\mathsf{F.IL} = \mathsf{F.OL} = \mathbb{N}$. *Assume* $\mathsf{F.kl} \leq \mathsf{H.ol}$. *Let* $\overline{\mathsf{RE}}$ *be an IND-CPA PKE scheme with fixed input length* $\overline{\mathsf{RE}}.\mathsf{il}$ *and public key length* $\overline{\mathsf{RE}}.\mathsf{pkl}$ *satisfying* $\overline{\mathsf{RE}}.\mathsf{il} + \overline{\mathsf{RE}}.\mathsf{pkl} = \mathsf{H.il}$. *Let* G *be an iO-secure iO scheme. Define PKE scheme* RE *as follows. Let* $\mathsf{RE.il} = \overline{\mathsf{RE}}.\mathsf{il}$. *Let* $\mathsf{RE.Kg} = \overline{\mathsf{RE}}.\mathsf{Kg}$. *Let the encryption and decryption algorithms of* RE *be as shown in Fig. 4. Then (1)* $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ *is not* IND*-secure, but (2)* RE *is IND-CPA secure.* ■

The proof of Proposition 1 is in [7]. Here we will sketch the ideas. An encryption $c = (c', P)$ of a message m under RE with public key ek will have two parts. The first, c', is an encryption of m under $\overline{\mathsf{RE}}$ with ek. The second, P, is an obfuscated circuit that will (1) help attack $\mathsf{DE} = \mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ yet (2) not compromise IND-CPA security of RE. The question is how to construct RE to ensure *both* properties. (Ensuring either alone is trivial.)

The starting idea, inspired by BFM [21], is to have RE.Enc, given $1^\lambda, ek, m$ and coins r, create the following circuit:

$C_{1^\lambda,ek,m,r}(hk)$: If $\mathsf{H}(1^\lambda, hk, ek \| m, 1^{\mathsf{RE.rl}(\lambda)}) = r$ then return m else return $0^{|m|}$.

The input to the circuit is a key hk for H, and the hardwired values $1^\lambda, ek, m, r$ are the inputs to the algorithm RE.Enc that creates the circuit. Now RE.Enc lets P be an obfuscation of this circuit. Pretend for now that the obfuscation process

$\underline{\mathsf{DE1.Kg}(1^\lambda)}$
$(ek, dk) \leftarrow\!\!\$\ \mathsf{LT.EKg}(1^\lambda)$; $hk \leftarrow\!\!\$\ \mathsf{H.Kg}(1^\lambda)$; Return $((ek, hk), (dk, hk))$

$\underline{\mathsf{DE1.Enc}(1^\lambda, (ek, hk), m)}$
$r \leftarrow \mathsf{H.Ev}(1^\lambda, hk, m, 1^{\mathsf{LT.il}(\lambda)})$; $trap \leftarrow \mathsf{LT.Ev}(1^\lambda, ek, r)$
$c \leftarrow m \oplus \mathsf{H.Ev}(1^\lambda, hk, r, 1^{|m|})$; Return $(trap, c)$

$\underline{\mathsf{DE1.Dec}(1^\lambda, (dk, hk), (trap, c))}$
$r \leftarrow \mathsf{LT.Inv}(1^\lambda, dk, trap)$; Return $c \oplus \mathsf{H.Ev}(1^\lambda, hk, r, 1^{|c|})$

Fig. 5. The algorithms of our DE1 D-PKE scheme

is deterministic, which of course is not true, and also that no coins are used to create c', which is also not true. Under these assumptions, if an attacker has an EwH ciphertext $(c', P) = \mathsf{DE.Enc}(1^\lambda, (ek, hk), m)$, and also has the public key (ek, hk) of DE, then it can run P on hk which, due to the structure of EwH and the construction of $C_{1^\lambda, ek, m, r}$, returns m, breaking the IND-security of DE. But there are a number of difficulties. One is that there seems no reason that this RE retains IND-CPA security assuming only iO security of the obfuscation. Another is that the obfuscation and $\overline{\mathsf{RE}}$ are randomized, and RE has to provide coins for both from r yet be able to create P to allow the attack when r is produced via the hash in EwH.

We will use the VOL PRF F to allocate pseudorandom coins for the obfuscation process and $\overline{\mathsf{RE}}$. The key for F will be a prefix $fk \leftarrow r[1, \mathsf{F.kl}(\lambda)]$ of the coins r provided to RE.Enc. Recall that in our definition of a VOL PRF, the key generation always samples $fk \leftarrow\!\!\$\ \{0,1\}^{\mathsf{F.kl}(\lambda)}$, so if r is truly random then we give F a correctly generated key. Instead of hardwiring r to the circuit, we hardwire $y \leftarrow \mathsf{F.Ev}(1^\lambda, fk, 0^{\mathsf{F.il}(\lambda)}, 1^\ell)$ for an appropriate ℓ. We also hardwire $x = ek \| m$ rather than ek, m separately. Our circuit $C_{1^\lambda, x, y}$ is shown in the left panel of Fig. 4. We need (1) an attack on $\mathsf{DE} = \mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ and (2) a proof that RE is IND-CPA. For (1) our claim is that if $C_{1^\lambda, ek\|m, y}$ is produced by RE.Enc within DE then $C_{1^\lambda, ek\|m, y}(hk)$ will return $ek\|m$, and thus running an obfuscation P of $C_{1^\lambda, ek\|m, y}$ on hk will return the same. For (2), r is truly random so $C_{1^\lambda, ek\|m, y}$ as produced during encryption is indistinguishable from $C_{1^\lambda, ek\|m, u}$ with u a random ℓ-bit string, by PRF security of F. To use iO security, we want that when u is random the probability that there exists a $\mathsf{H.kl}(\lambda)$-bit z such that $C_{1^\lambda, ek\|m, u}(z) \neq 0^{|ek\,\|\,m|}$ is negligible. This is established via a counting argument which relies on having set ℓ to be large enough. See [7] for details.

The DE1 Scheme. We now provide our positive result on D-PKE, namely an efficient, fully IND standard model scheme under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$. Let H be a $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ function family with $\mathsf{H.IL}(\cdot) = \mathsf{H.OL}(\cdot) = \mathbb{N}$. From the above we know that $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$ will not be IND for all IND-CPA RE. We consider instead a particular choice of IND-CPA RE. Recall that BR93 [13] present a simple TDF-based PKE scheme proven IND-CPA in the ROM. We instantiate their TDF with a LTDF and then instantiate the RO with H to get a standard-model

PKE scheme we denote $\mathsf{RE} = \mathsf{BR93}[\mathsf{LT}, \mathsf{H}]$. We now consider the standard-model D-PKE scheme $\mathsf{EwH}[\mathsf{H}, \mathsf{RE}]$. In this scheme, H is used twice, with two independent keys. Our final $\mathsf{DE1}$ D-PKE scheme is obtained by using the same key for both invocations of H. The algorithms of this scheme are shown in Fig. 5. Importantly, $\mathsf{DE1.IL}(\cdot) = \mathsf{H.OL}(\cdot) = \mathbb{N}$, meaning we can encrypt messages of arbitrary and varying length. We note that using a single H key is not only an optimization in key size but also avoids using multi-key variants of UCE [8] and is important to prove security under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$. The following says that $\mathsf{DE1}$ is IND-secure.

Theorem 2. *Let* LT *be a lossy trapdoor function and* H *a* $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ *function family with* $\mathsf{H.IL}(\cdot) = \mathsf{H.OL}(\cdot) = \mathbb{N}$. *Let* $\mathsf{DE1}$ *be constructed as in Fig. 5. Then*

Asymptotic result: $\mathsf{DE1}$ *is IND-secure.*

Concrete result: Let A be an adversary and P a predictor. We can construct an adversary B, a source S, and a distinguisher D such that

$$\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{DE1},A}(\cdot) \leq 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\cdot) + 2\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\cdot) + \frac{3v^2}{2^{\mathsf{LT.il}}} \tag{1}$$

$$\mathsf{Adv}^{\mathsf{pred}}_{S,P}(\cdot) \leq \frac{1.5v^2}{2^{\mathsf{LT.il}}} + qv \cdot \mathrm{Guess}_A(\cdot) + \frac{qv}{2^{\tau}} \tag{2}$$

where q is the maximum of the size of P's output in the execution of Pred^P_S*, v is the maximum of the size of A's message vector in the execution of* $\mathrm{IND}^A_{\mathsf{DE}}$*, and τ is the lossiness of* LT*. Furthermore,* $\mathbf{T}(\mathrm{UCE}^{S,D}_{\mathsf{H}}) \leq \mathbf{T}(\mathrm{IND}^A_{\mathsf{DE1}})$*;* $\mathbf{Q}^{\mathrm{HASH}}_S \leq v$*; and* $\mathbf{T}(\mathrm{Lossy}^B_{\mathsf{LT}}) \leq \mathbf{T}(\mathrm{IND}^A_{\mathsf{DE1}})$*.* ■

The proof is in [7]. Here we discuss some of the ideas. To construct a source S and a distinguisher D, a naive method is to let them run A to simulate game $\mathrm{IND}^A_{\mathsf{DE1}}$. However this won't produce a statistically unpredictable source. The key idea is to let our source generate a *lossy* key *lk*. instead of an injective key *ek* as in game $\mathrm{IND}^A_{\mathsf{DE1}}$. The statistical unpredictability of S then follows from the lossiness of LT, as represented by (2). On the other hand, game $\mathrm{UCE}^{S,D}_{\mathsf{H}}$ for challenge bit $b = 1$ no longer coincides with game $\mathrm{IND}^A_{\mathsf{DE1}}$. Still, this gap can be bounded by constructing B attacking LT, so that (1) holds.

In Section 5 we discuss how, under appropriate instantiations of the $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family, $\mathsf{DE1}$ is extremely efficient compared to prior standard-model D-PKE schemes.

BFOR [6] originally defined an IND adversary as a triple (A_0, A_1, A_2), where A_0 specifies state information that is passed on to A_1, A_2. Results from [5] indicate this is important to ensure that security in the standard model implies security in the ROM. For notational simplicity, here we omit A_0. Our construction and proof work for the original IND definition with the following modification. One first needs to redefine Guess_A as the conditional min-entropy of the messages, given the state, and then include the state as a part of the leakage of S.

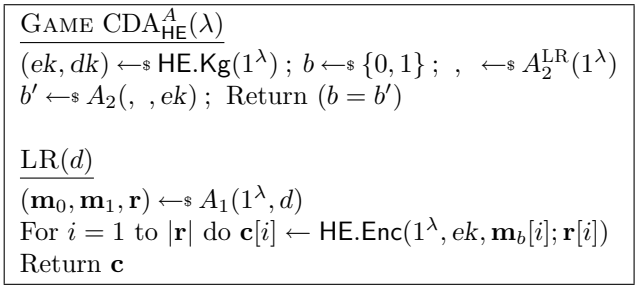

Fig. 6. Game defining IND-CDA security of PKE scheme HE

4 Fully Secure Hedged PKE

In this section we provide the first fully H-IND PKE schemes in the standard model. Additionally our schemes are efficient. HE1 is our base scheme encrypting fixed-length messages; HE2 encrypts variable-length messages; HE3 has a tighter security analysis. Our schemes provide pragmatic and effective defense against subversion of encryption randomness.

Hedged PKE. To achieve standard IND-CPA security, PKE schemes demand truly random coins. Many well-known PKE schemes fail spectacularly, allowing message recovery from the ciphertext, if the latter is created with even somewhat weak coins [4,20,46]. BBNRSS [4] introduce security under chosen-distribution attack (IND-CDA) to provide meaningful security when bad randomness is used. A secure hedged PKE scheme must provide IND-CPA security when the coins are truly random, and fall back to IND-CDA security when bad coins are provided. Formally, for a PKE scheme HE, we say that HE is H-IND secure if (1) HE is IND-CPA secure, and (2) HE is IND-CDA secure. Game CDA defining the IND-CDA notion is given in Fig. 6. An IND-CDA adversary $A = (A_1, A_2)$ is a pair of algorithms. In the first part of the attack, A_2 can adaptively query oracle LR, each query taking a distribution-specifier string d and returning a challenge ciphertext vector $\mathbf{c}$. In this phase A_2 does not get ek. Once this stage ends, it gets ek and must then render its decision. Algorithm A_1 defines a distribution over triples $(\mathbf{m}_0, \mathbf{m}_1, \mathbf{r})$ that is a function of d. We require that (i) there be a polynomial v such that $|\mathbf{m}_0| = |\mathbf{m}_1| = |\mathbf{r}| \leq v(\lambda)$, (ii) $|\mathbf{m}_0[i]| = |\mathbf{m}_1[i]| \in \mathsf{HE.IL}(\lambda)$ and $|\mathbf{r}[i]| = \mathsf{HE.rl}(\lambda)$ for every $i \leq |\mathbf{r}|$, and (iii) for each $b \in \{0,1\}$ the $|\mathbf{r}|$ pairs $(\mathbf{m}_b[i], \mathbf{r}[i])$ are distinct, where $1 \leq i \leq |\mathbf{r}|$. Let $\mathrm{Guess}_A(\cdot)$ be the function that on input $\lambda \in \mathbb{N}$ returns the maximum, over all b, i, m, r, d, of $\Pr[(\mathbf{m}_b[i], \mathbf{r}[i]) = (m, r)]$, the probability over $(\mathbf{m}_0, \mathbf{m}_1, \mathbf{r}) \leftarrow\!\$\ A_1(1^\lambda, d)$. We say that A has high min-entropy if $\mathrm{Guess}_A(\cdot)$ is negligible. We say that HE is IND-CDA-secure if $\mathsf{Adv}^{\mathsf{cda}}_{\mathsf{HE},A}(\cdot) = 2\Pr[\mathrm{CDA}^A_{\mathsf{HE}}(\cdot)] - 1$ is negligible for every PT adversary A of high min-entropy. We stress that this captures full IND-CDA since the messages in the message vectors may be arbitrarily correlated.

The HE1 Scheme. Recall we obtained our D-PKE scheme DE1 via a BR93-based instantiation of EwH. In analogy it is natural to try to obtain an H-

$\underline{\mathsf{Hedge}[\mathsf{H},\mathsf{LT}].\mathsf{Kg}(1^\lambda)}$	$\underline{\mathsf{Hedge}[\mathsf{H},\mathsf{LT}].\mathsf{Enc}(1^\lambda,(ek,hk),m;r)}$
$hk \leftarrow\!\!\$\ \mathsf{H.Kg}(1^\lambda)$	$x \leftarrow \mathsf{H.Ev}(1^\lambda, hk, r \,\|\, m, 1^{\mathsf{LT.il}(\lambda)})$
$(ek,dk) \leftarrow\!\!\$\ \mathsf{LT.EKg}(1^\lambda)$	$trap \leftarrow \mathsf{LT.Ev}(1^\lambda, ek, x)$
Return $((ek,hk),(dk,hk))$	$c \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^{\lvert m\rvert}) \oplus m$
	Return $(trap, c)$
	$\underline{\mathsf{Hedge}[\mathsf{H},\mathsf{LT}].\mathsf{Dec}(1^\lambda,(dk,hk),(trap,c))}$
	$x \leftarrow \mathsf{LT.Inv}(1^\lambda, dk, trap)$
	$m \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^{\lvert c\rvert}) \oplus c$; Return m

$\underline{\mathsf{H.Kg}(1^\lambda)}$	$\underline{\mathsf{H.Ev}(1^\lambda, hk, x, 1^\ell)}$
$uk \leftarrow\!\!\$\ \mathsf{U.Kg}(1^\lambda)$; $\overline{hk} \leftarrow\!\!\$\ \overline{\mathsf{H}}.\mathsf{Kg}(\lambda)$	$(\overline{hk}, uk) \leftarrow hk$; $u \leftarrow \mathsf{U.Ev}(1^\lambda, uk, x)$
$hk \leftarrow (\overline{hk}, uk)$; Return hk	$y \leftarrow \overline{\mathsf{H}}.\mathsf{Ev}(1^\lambda, \overline{hk}, u, 1^\ell)$; Return y

$\underline{\mathsf{U.Kg}(1^\lambda)}$	$\underline{\mathsf{U.Ev}(1^\lambda, (\overline{uk}, rk), x)}$
$\overline{uk} \leftarrow\!\!\$\ \overline{\mathsf{U}}.\mathsf{Kg}(1^\lambda)$	If $\lvert x\rvert < \mathsf{U.ol}(\lambda)$ then
$mk \leftarrow\!\!\$\ \{0,1\}^{\mathsf{U.ol}(\lambda)}$	Return $mk \oplus (x \,\|\, 10^{\mathsf{U.ol}(\lambda) - \lvert x\rvert})$
$rk \leftarrow\!\!\$\ \mathrm{GF}(2^{\mathsf{U.ol}(\lambda)}) \setminus \{0^{\mathsf{U.ol}(\lambda)}\}$	$x_1 \leftarrow x[1, \mathsf{U.ol}(\lambda)]$; $x_2 \leftarrow x[\mathsf{U.ol}(\lambda)+1, \lvert x\rvert]$
Return $(\overline{uk}, rk, mk)$	$y \leftarrow \overline{\mathsf{U}}.\mathsf{Ev}(1^\lambda, \overline{uk}, x_2) \oplus (x_1 \times rk)$
	Return y

Fig. 7. **Top:** The PKE scheme $\mathsf{Hedge}[\mathsf{H},\mathsf{LT}]$ associated to function family H and LTDF LT. **Middle:** The $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U},\overline{\mathsf{H}}]$ VIL $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family built from an AU hash U and a FIL $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family $\overline{\mathsf{H}}$. **Bottom:** The $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$ AU family built from an AU family $\overline{\mathsf{U}}$. The operator $\times$ is multiplication in the finite field $\mathrm{GF}(2^{\mathsf{U.ol}(\lambda)})$ and the string $0^{\mathsf{U.ol}(\lambda)}$ encodes the zero element of $\mathrm{GF}(2^{\mathsf{U.ol}(\lambda)})$. **HE1:** Our HE1 PKE scheme is obtained from an LTDF LT, a FIL $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family $\overline{\mathsf{H}}$ and an AU family $\overline{\mathsf{U}}$ as $\mathsf{HE1} = \mathsf{Hedge}[\mathsf{H},\mathsf{LT}]$ with $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U},\overline{\mathsf{H}}]$ and $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$.

IND scheme via a similar BR93-based instantiation of the REwH transform of BBNRSS [4]. This results in the candidate scheme $\mathsf{Hedge}[\mathsf{H},\mathsf{LT}]$, associated to a function family H and LTDF LT, whose algorithms are shown in the left panel of Fig. 7. Here $\mathsf{Hedge}[\mathsf{H},\mathsf{LT}].\mathsf{IL}(\cdot) = \mathsf{H.OL}(\cdot)$, meaning we can encrypt messages of length matching the allowed output lengths of H.

We first ask if one can show IND-CDA security of $\mathsf{Hedge}[\mathsf{H},\mathsf{LT}]$ assuming $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ security of H. This involves two new difficulties relative to Theorem 2. The first, more minor, is the presence of the randomness. The second is more major, namely that the IND-CDA notion is adaptive. To address this, BBNRSS [4] needed quite involved techniques including anonymous LTDFs and an adaptive LHL, and yet only achieved security for block sources, not the full IND-CDA security that we target. However we are able to show that $\mathsf{Hedge}[\mathsf{H},\mathsf{LT}]$ does achieve (full) IND-CDA assuming only that LT is a (standard) LTDF and H is $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$.

But recall that H-IND requires also that $\mathsf{Hedge}[\mathsf{H},\mathsf{LT}]$ is IND-CPA. But it is quite unclear why this would be true under $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ security of H. The reason is that UCE guarantees nothing for inputs depending on hk but messages in

IND-CPA can depend on the public key, which contains hk. This difficulty is quite fundamental and at first seemed impossible to circumvent within the UCE framework. We resolve it by using a *particular* $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family H. Let $\overline{\mathsf{H}}$ be a fixed input length $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family. Recall that the AU-then-Hash transform of BHK2 [10] takes an AU (almost universal) family U and $\overline{\mathsf{H}}$ to return a variable input length family $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U}, \overline{\mathsf{H}}]$ that they show is itself $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$. We will take an (arbitrary) AU family $\overline{\mathsf{U}}$ and construct another, special AU family $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$ via a transform called Hash-then-Mask that we introduce. Then our $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family is $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U}, \overline{\mathsf{H}}]$. With this choice we will be able to show that $\mathsf{HE1} = \mathsf{Hedge}[\mathsf{H}, \mathsf{LT}]$ —this is our scheme— is IND-CPA. In conjunction with our prior claim, HE1 is then H-IND as desired.

We now detail this. We recall some definitions from BHK2 [9]. Let V be a fixed output length (FOL) function family. Let $\lambda, m \in \mathbb{N}$. Let

$$\begin{aligned}
\mathbf{Coll1}_{\mathsf{V}}(\lambda, m) &= \max\{\, \Pr[y = \mathsf{V.Ev}(1^{\lambda}, vk, x)] : |y| = \mathsf{V.ol}(\lambda) \text{ and } |x| \le m \,\} \\
\mathbf{Coll2}_{\mathsf{V}}(\lambda, m_0, m_1) &= \max\{\, \Pr[\mathsf{V.Ev}(1^{\lambda}, vk, x_0) = \mathsf{V.Ev}(1^{\lambda}, vk, x_1)] \ : \\
&\qquad |x_0| \le m_0,\ |x_1| \le m_1 \text{ and } x_0 \ne x_1 \,\} \\
\mathbf{Coll}_{\mathsf{V}}(\lambda, m_0, m_1) &= \max\{\, \mathbf{Coll2}_{\mathsf{V}}(\lambda, m_0, m_1),\ \mathbf{Coll1}_{\mathsf{V}}(\lambda, \min\{m_0, m_1\}) \,\} \ .
\end{aligned}$$

In the first and second equations, the probability is over $vk \leftarrow\!\!\$\ \mathsf{V.Kg}(1^{\lambda})$. A FOL family V is *almost universal* (AU) if for all polynomials $M_0, M_1\colon \mathbb{N} \to \mathbb{N}$ the function f_{M_0,M_1} is negligible, where for $\lambda \in \mathbb{N}$ we let $f_{M_0,M_1}(\lambda) = \mathbf{Coll}_{\mathsf{V}}(\lambda, M_0(\lambda), M_1(\lambda))$.

Now let $\overline{\mathsf{U}}$ be a (FOL) AU family having $\overline{\mathsf{U}}.\mathsf{IL} = \mathbb{N}$. We introduce a transform called Hash-then-Mask that given $\overline{\mathsf{U}}$ returns the family $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$ defined in the right panel of Fig. 7. It has $\mathsf{U.ol} = \overline{\mathsf{U}}.\mathsf{ol}$ and $\mathsf{U.IL} = \mathbb{N}$. Lemma 3 below shows that U is itself an AU family.

Lemma 3. *Let $\overline{\mathsf{U}}$ be a (FOL) AU hash of $\overline{\mathsf{U}}.\mathsf{IL} = \mathbb{N}$. Let $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$. Then for any $\lambda, m, m' \in \mathbb{N}$ we have (a) $\mathbf{Coll1}_{\mathsf{U}}(\lambda, m) \le \mathbf{Coll1}_{\overline{\mathsf{U}}}(\lambda, m) + 2^{-\overline{\mathsf{U}}.\mathsf{ol}(\lambda)}$ and (b) $\mathbf{Coll2}_{\mathsf{U}}(\lambda, m, m') \le \mathbf{Coll2}_{\overline{\mathsf{U}}}(\lambda, m, m') + 2/2^{\overline{\mathsf{U}}.\mathsf{ol}(\lambda)}$.* ■

The proof of Lemma 3 is in [7]. Note that BHK2 [9] provide an extremely fast construction of an AU family $\overline{\mathsf{U}}$, running at 0.4 cycles per byte. Our Hash-then-Mask does not degrade speed much, and thus the family $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$ used in our scheme is also fast.

Now let $\overline{\mathsf{H}}$ be a function family with FIL $\overline{\mathsf{H}}.\mathsf{il}$ and with $\overline{\mathsf{H}}.\mathsf{OL} = \mathbb{N}$. Let U be a FOL AU function family with $\mathsf{U.ol} = \overline{\mathsf{H}}.\mathsf{il}$ and with $\mathsf{U.IL} = \overline{\mathsf{H}}.\mathsf{OL} = \mathbb{N}$. The AU-then-Hash transform of BHK2 [9] takes $\mathsf{U}, \overline{\mathsf{H}}$ and returns the family $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U}, \overline{\mathsf{H}}]$ shown in the middle panel of Fig. 7. It has $\mathsf{H.OL} = \mathsf{H.IL} = \mathbb{N}$. BHK2 [9] show that if $\overline{\mathsf{H}}$ is $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ then so is H.

We are finally ready to define our HE1 scheme. Let $\overline{\mathsf{H}}$ be a function family with FIL $\overline{\mathsf{H}}.\mathsf{il}$ and with $\overline{\mathsf{H}}.\mathsf{OL} = \mathbb{N}$. Let $\overline{\mathsf{U}}$ be a (FOL) AU family having $\overline{\mathsf{U}}.\mathsf{IL} = \mathbb{N}$. Let LT be an LTDF. Let $\ell\colon\ \mathbb{N} \to \mathbb{N}$ be a polynomial. Then let $\mathsf{HE1} = \mathsf{Hedge}[\mathsf{H}, \mathsf{LT}]$ with $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U}, \overline{\mathsf{H}}]$ and $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$. A subtle point is that we set $\mathsf{HE1.il} = \ell$, meaning HE1 is restricted to encrypt messages of length ℓ.

Why this is needed is not evident from the scheme description but will be needed in the proof of security. We also set $\mathsf{HE1.rl} = \overline{\mathsf{U}}.\mathsf{ol}$. Theorem 4 below shows that $\mathsf{HE1}$ is H-IND secure. The concrete security statements refer to

$$\mathsf{Adv}^{\mathsf{coll}}_{\mathsf{U}}(\lambda, p, \sigma) = \max\left\{ \sum_{i=1}^{k}\sum_{j=1}^{k'} \mathbf{Coll}_{\mathsf{U}}(\lambda, m_i, m'_j) \ : \ k \right\} \le p, k' \le p, \sum_{i=1}^{k} m_i \le \sigma, \sum_{i=1}^{k'} m'_i \le \sigma \ .$$

Theorem 4. *Let $\overline{\mathsf{H}}$ be a $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ function family with FIL $\overline{\mathsf{H}}.\mathsf{il}$ and with $\overline{\mathsf{H}}.\mathsf{OL} = \mathbb{N}$. Let $\overline{\mathsf{U}}$ be a (FOL) AU family having $\overline{\mathsf{U}}.\mathsf{IL} = \mathbb{N}$. Let LT be an LTDF. Let $\ell\colon \mathbb{N} \to \mathbb{N}$ be a polynomial. Let $\mathsf{HE1}$ be defined from $\overline{\mathsf{H}}, \overline{\mathsf{U}}, \mathsf{LT}, \ell$ as above.*

Asymptotic result: $\mathsf{HE1}$ is H-IND secure.

Concrete IND-CPA result: Let A be an adversary and $\overline{P}$ be a predictor. We can construct a source $\overline{S}$, a distinguisher $\overline{D}$ and an adversary B such that

$$\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{HE1},A}(\cdot) \le 2\mathsf{Adv}^{\mathsf{uce}}_{\overline{\mathsf{H}},\overline{S},\overline{D}}(\cdot) + 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\cdot) + 2^{1-\overline{\mathsf{U}}.\mathsf{ol}}$$

$$\mathsf{Adv}^{\mathsf{pred}}_{\overline{S},\overline{P}}(\cdot) \le \frac{\sqrt{q}}{2^{\tau/2}} + \sqrt{q \cdot \mathbf{Coll2}_{\overline{\mathsf{U}}}(\cdot, \mathsf{LT}.\mathsf{il})} + \frac{2\sqrt{q}}{2^{\overline{\mathsf{U}}.\mathsf{ol}/2}}$$

where q is the maximum of the size of $\overline{P}$'s output in the execution of $\mathrm{Pred}^{\overline{P}}_{\overline{S}}$ and τ is the lossiness of LT. Furthermore, $\mathbf{T}(\mathrm{Lossy}^{B}_{\mathsf{LT}}), \mathbf{T}(\mathrm{UCE}^{\overline{S},\overline{D}}_{\overline{\mathsf{H}}})$; and $\mathbf{Q}^{\mathrm{HASH}}_{S} = 2$.

Concrete IND-CDA result: Let A be an adversary and $\overline{P}$ be a predictor. We can construct a source $\overline{S}$, a distinguisher $\overline{D}$ and an adversary B such that

$$\mathsf{Adv}^{\mathsf{cda}}_{\mathsf{HE1},A}(\cdot) \le 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\cdot) + 2\mathsf{Adv}^{\mathsf{uce}}_{\overline{\mathsf{H}},\overline{S},\overline{D}}(\cdot) + 2\mathsf{Adv}^{\mathsf{coll}}_{\overline{\mathsf{U}}}(\cdot, 2p, s) + 3p^2 \cdot \mathrm{Guess}_A(\cdot) + \frac{19p^2}{2^{\min\{\overline{\mathsf{U}}.\mathsf{ol}, \mathsf{LT}.\mathsf{il}\}}}$$

$$\mathsf{Adv}^{\mathsf{pred}}_{\overline{S},\overline{P}}(\cdot) \le \sqrt{2q \cdot \mathsf{Adv}^{\mathsf{coll}}_{\overline{\mathsf{U}}}(\cdot, 2p, s)} + 2p\sqrt{q \cdot \mathrm{Guess}_A(\cdot)} + \frac{6p\sqrt{q}}{2^{\min\{\overline{\mathsf{U}}.\mathsf{ol}, \tau\}/2}}$$

where p is the maximum of the total number of messages that A produces in the execution of $\mathrm{CDA}^{A}_{\mathsf{HE1}}$, $s = p \cdot (\overline{\mathsf{U}}.\mathsf{ol} + \mathsf{LT}.\mathsf{il} + \ell)$, q is the maximum of the size of $\overline{P}$'s output in the execution of $\mathrm{Pred}^{\overline{P}}_{\overline{S}}$, and τ is the lossiness of LT. Moreover, $\mathbf{T}(\mathrm{Lossy}^{B}_{\mathsf{LT}}), \mathbf{T}(\mathrm{UCE}^{\overline{S},\overline{D}}_{\overline{\mathsf{H}}}) \le \mathbf{T}(\mathrm{CDA}^{A}_{\mathsf{HE1}})$; and $\mathbf{Q}^{\mathrm{HASH}}_{\overline{S}} \le 2p$. ■

The proof of Theorem 4 is in [7]. Here we discuss some of the ideas. For IND-CPA security, recall that the adversary A makes only a single LR query. The transform $\mathsf{Hash\text{-}then\text{-}Mask}$ ensures that, for any string m, if r is a random $\overline{\mathsf{U}}.\mathsf{ol}(\lambda)$-bit string and $uk \xleftarrow{\$} \mathsf{U}.\mathsf{Kg}(1^\lambda)$ then $u \leftarrow \mathsf{U}(1^\lambda, uk, r \,\|\, m)$ is also uniformly random, independent of m. Therefore, one doesn't need to know m to sample $r \xleftarrow{\$} \{0,1\}^{\overline{\mathsf{U}}.\mathsf{ol}(\lambda)}$

$\underline{\mathsf{HE2.Enc}(1^\lambda, (ek, hk), m; r)}$	$\underline{\mathsf{HE2.Dec}(1^\lambda, (dk, hk), (trap, c))}$
$x \leftarrow \mathsf{H.Ev}(1^\lambda, hk, r \,\|\, m, 1^{\mathsf{LT.il}(\lambda)})$	$x \leftarrow \mathsf{LT.Inv}(1^\lambda, dk, trap)$
$trap \leftarrow \mathsf{LT.Ev}(1^\lambda, ek, x)$	$seed \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^{\mathsf{F.kl}(\lambda)+\overline{\mathsf{U}}.\mathsf{ol}(\lambda)})$
$seed \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^{\mathsf{F.kl}(\lambda)+\overline{\mathsf{U}}.\mathsf{ol}(\lambda)})$	$y \leftarrow seed[1, \overline{\mathsf{U}}.\mathsf{ol}(\lambda)]$
$y \leftarrow seed[1, \overline{\mathsf{U}}.\mathsf{ol}(\lambda)]$	$fk \leftarrow seed[\overline{\mathsf{U}}.\mathsf{ol}(\lambda) + 1, \lvert seed\rvert]$
$fk \leftarrow seed[\overline{\mathsf{U}}.\mathsf{ol}(\lambda) + 1, \lvert seed\rvert]$	$mask \leftarrow \mathsf{F.Ev}(1^\lambda, fk, 0^{\mathsf{F.il}(\lambda)}, 1^{\lvert c\rvert})$
$mask \leftarrow \mathsf{F.Ev}(1^\lambda, fk, 0^{\mathsf{F.il}(\lambda)}, 1^{\lvert m\rvert})$	$m \leftarrow \mathsf{H.Ev}(1^\lambda, hk, y, 1^{\lvert c\rvert}) \oplus mask \oplus c$
$c \leftarrow \mathsf{H.Ev}(1^\lambda, hk, y, 1^{\lvert m\rvert}) \oplus mask \oplus m$	Return m
Return $(trap, c)$	

Fig. 8. Encryption and decryption algorithms of HE2, where $\overline{\mathsf{U}}$ is an AU family, $\overline{\mathsf{H}}$ is a FIL $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family, F is a VOL PRF, LT is a LTDF. Here $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$ and $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U}, \overline{\mathsf{H}}]$.

and compute $x \leftarrow \mathsf{H.Ev}(1^\lambda, hk, r \,\|\, m, 1^{\mathsf{LT.il}(\lambda)})$, because one can instead sample $u \leftarrow\!\!\$\ \{0,1\}^{\overline{\mathsf{U}}.\mathsf{ol}(\lambda)}$ and compute $x \leftarrow \overline{\mathsf{H}}.\mathsf{Ev}(1^\lambda, \overline{hk}, u, 1^{\mathsf{LT.il}(\lambda)})$. The source will leak $\mathsf{H.Ev}(1^\lambda, hk, x, 1^{|m|})$ so that the distinguisher can run A to get m and xor the two strings to complete the ciphertext. Still, computing $\mathsf{H.Ev}(1^\lambda, hk, x, 1^{|m|})$ requires knowing $|m|$; it's why HE1 can only handle fixed-length messages. For IND-CDA security, we can actually prove that $\mathsf{Hedge}[\mathsf{H}, \mathsf{LT}]$ is IND-CDA secure for *any* $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ H. The source will run A_1 and the first phase of A_2 to create the ciphertexts via the HASH oracle. Note that during the first phase, A_2 only receives what the source sees, and therefore doesn't get to learn the hash key hk. UCE then allows us to switch to a game in which the adversary has to fight an RO-based scheme, and thus its adaptivity is futile. Moreover, it can only specify *distributions*, and thus despite the adaptivity, the chance that the source repeats a HASH query is about $p^2 \cdot \mathrm{Guess}_A$. We again exploit the lossiness of LT to allow statistical unpredictability.

The HE2 Scheme. With HE1 we reach our goal of the first fully H-IND secure PKE scheme in the standard model. Additionally it is more efficient than prior standard-model schemes that only achieved non-full security. However, like prior standard-model schemes, it is FIL, meaning only encrypts messages of a fixed length. We now provide the HE2 scheme that retains the security properties of HE1 but additionally can encrypt messages of variable and arbitrary length. Furthermore it can do this with hybrid-encryption like performance, meaning the asymmetric cost is fixed as message length grows.

The additional tool that we need is a VOL PRF F —this means $\mathsf{F.OL}(\cdot) = \mathbb{N}$— such that $\lambda \in \mathsf{F.IL}(\lambda)$ for every $\lambda \in \mathbb{N}$. As before let $\overline{\mathsf{H}}$ be a function family with FIL $\overline{\mathsf{H}}.\mathsf{il}$ and with $\overline{\mathsf{H}}.\mathsf{OL}(\cdot) = \mathbb{N}$. Let $\overline{\mathsf{U}}$ be a (FOL) AU family having $\overline{\mathsf{U}}.\mathsf{IL}(\cdot) = \mathbb{N}$. Let LT be an LTDF. Let $\mathsf{U} = \mathsf{Hash\text{-}then\text{-}Mask}[\overline{\mathsf{U}}]$ and $\mathsf{H} = \mathsf{AU\text{-}then\text{-}Hash}[\mathsf{U}, \overline{\mathsf{H}}]$. The encryption and decryption algorithms of HE2 are specified in Fig. 8. The key-generation algorithm HE2.Kg is the same as HE1.Kg. We let $\mathsf{HE2.rl} = \overline{\mathsf{U}}.\mathsf{ol}$. But this time $\mathsf{HE2.IL}(\cdot) = \mathbb{N}$, meaning we can encrypt messages of any length. Theorem 5 below formally confirms that HE2 is H-IND secure.

Theorem 5. *Let* F *be a PRF with* $\mathsf{F.OL}(\cdot) = \mathbb{N}$ *and* $\lambda \in \mathsf{F.IL}(\lambda)$ *for every* $\lambda \in \mathbb{N}$. *Let* $\overline{\mathsf{H}}$ *be a* $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ *function family with FIL* $\overline{\mathsf{H}}.\mathsf{il}$ *and with* $\overline{\mathsf{H}}.\mathsf{OL}(\cdot) = \mathbb{N}$. *Let* $\overline{\mathsf{U}}$ *be a (FOL) AU family having* $\overline{\mathsf{U}}.\mathsf{IL}(\cdot) = \mathbb{N}$. *Let* LT *be an LTDF. Let* $\mathsf{HE2}$ *be defined from* $\mathsf{F}, \overline{\mathsf{H}}, \overline{\mathsf{U}}, \mathsf{LT}$ *as above.*

Asymptotic result: $\mathsf{HE2}$ *is H-IND secure.*

Concrete IND-CPA result: Let A *be an adversary and* $\overline{P}$ *be a predictor. We can construct a source* $\overline{S}$*, a distinguisher* $\overline{D}$*, adversaries* B *and* C *such that*

$$\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{HE2},A}(\cdot) \le 2\mathsf{Adv}^{\mathsf{uce}}_{\overline{\mathsf{H}},\overline{S},\overline{D}}(\cdot) + 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\cdot) + 2\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{F},C}(\cdot) + 2^{1-\overline{\mathsf{U}}.\mathsf{ol}}$$

$$\mathsf{Adv}^{\mathsf{pred}}_{\overline{S},\overline{P}}(\cdot) \le \frac{\sqrt{q}}{2^{\tau/2}} + \sqrt{q \cdot \mathbf{Coll2}_{\overline{\mathsf{U}}}(\cdot, \mathsf{LT.il})} + \frac{2\sqrt{q}}{2^{\overline{\mathsf{U}}.\mathsf{ol}/2}}$$

where q *is the maximum of the size of* $\overline{P}$*'s output in the execution of* $\mathrm{Pred}^{\overline{P}}_{\overline{S}}$ *and* τ *is the lossiness of* LT*. Furthermore,* $\mathbf{T}(\mathrm{Lossy}^{B}_{\mathsf{LT}}), \mathbf{T}(\mathrm{UCE}^{\overline{S},\overline{D}}_{\overline{\mathsf{H}}}), \mathbf{T}(\mathrm{PRF}^{C}_{\mathsf{F}}) \le \mathbf{T}(\mathrm{CPA}^{A}_{\mathsf{HE3}})$*;* $\mathbf{Q}^{\mathrm{RR}}_{C} = 1$*; and* $\mathbf{Q}^{\mathrm{HASH}}_{S} = 2$.

Concrete IND-CDA result: Let A *be an adversary and* $\overline{P}$ *be a predictor. We can construct a source* $\overline{S}$*, a distinguisher* $\overline{D}$*, adversary* B *such that*

$$\mathsf{Adv}^{\mathsf{cda}}_{\mathsf{HE},A}(\cdot) \le 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\cdot) + 2\mathsf{Adv}^{\mathsf{uce}}_{\overline{\mathsf{H}},\overline{S},\overline{D}}(\cdot) + 2\mathsf{Adv}^{\mathsf{coll}}_{\overline{\mathsf{U}}}(\cdot, 3p, s) + 5p^2 \cdot \mathrm{Guess}_A(\cdot) + \frac{44p^2}{2^{\min\{\overline{\mathsf{U}}.\mathsf{ol},\mathsf{LT.il}\}}}$$

$$\mathsf{Adv}^{\mathsf{pred}}_{\overline{S},\overline{P}}(\cdot) \le \sqrt{2q\mathsf{Adv}^{\mathsf{coll}}_{\overline{\mathsf{U}}}(\cdot, 3p, s)} + 2.5p\sqrt{q \cdot \mathrm{Guess}_A(\cdot)} + \frac{9.5p\sqrt{q}}{2^{\min\{\overline{\mathsf{U}}.\mathsf{ol},\tau\}/2}}$$

where p *is the maximum of the total number of messages that* A *produces in the execution of* $\mathrm{CDA}^{A}_{\mathsf{HE2}}$*,* s *is* $3p \cdot \max\{\overline{\mathsf{U}}.\mathsf{ol}, \mathsf{LT.il}\}$ *plus the maximum of the total length of messages that* A *produces in the execution of* $\mathrm{CDA}^{A}_{\mathsf{HE2}}$*,* q *is the maximum of the size of* $\overline{P}$*'s output in the execution of* $\mathrm{Pred}^{\overline{P}}_{\overline{S}}$*, and* τ *is the lossiness of* LT*. In addition,* $\mathbf{T}(\mathrm{Lossy}^{B}_{\mathsf{LT}}), \mathbf{T}(\mathrm{UCE}^{\overline{S},\overline{D}}_{\overline{\mathsf{H}}}) \le \mathbf{T}(\mathrm{CDA}^{A}_{\mathsf{HE2}})$*; and* $\mathbf{Q}^{\mathrm{HASH}}_{\overline{S}} \le 3p$. ■

The proof of Theorem 5 is in [7]. Here we give some intuition about why $\mathsf{HE2}$ can securely handle variable-length messages. We'll only discuss the IND-CPA case, in which the message length may depend on the public key. The source will be responsible for producing a PRF key *fk*, whose length is independent of the public key, and will leak it along with some other information. The UCE security is only used to ensure that *fk* looks random to the distinguisher. The task of generating the two pads $\mathsf{F.Ev}(1^\lambda, fk, 0^{\mathsf{F.il}(\lambda)}, 1^{|m|})$ and $\mathsf{H.Ev}(1^\lambda, hk, y, 1^{|m|})$ is left to the distinguisher who runs A to get m. Note that the distinguisher always creates $\mathsf{H.Ev}(1^\lambda, hk, y, 1^{|m|})$ regardless of the challenge bit of game UCE. We then use the PRF security of F to ensure that the first pad looks random to A. Consequently, in the string $(trap, c)$ that A receives, the first component is independent of the message, and the second component is indistinguishable from a random string.

The HE3 Scheme. Consider the $p\sqrt{q \cdot \text{Guess}_A(\cdot)}$ term in the concrete bound for IND-CDA security in Theorem 5. This is worse than the "optimal" bound $p(q+p) \cdot \text{Guess}_A(\cdot)$ if one uses a random oracle. Why does this gap matter? Asymptotically, we know that $\text{Guess}_A(\cdot)$ is negligible, and hence this entire term is negligible too, under either of the two bounds. But concretely, the first bound means that we must have more min-entropy in the messages to get security. This is not desirable in practice. For example if we encrypt passwords, their min-entropy may be borderline. Thus it would be desirable to have a better bound. Moreover, it would also be desirable to give a simple construction based on a *generic* UCE-secure hash. We achieve both goals with our HE3 scheme.

The only ingredients we need this time are a PRF F (with fixed input length F.il and $\mathsf{F.OL}(\cdot) = \mathbb{N}$), a $\mathsf{UCE}[\mathcal{S}^{\text{srs}}]$ family H (with $\mathsf{H.IL}(\cdot) = \mathsf{H.OL}(\cdot) = \mathbb{N}$) and a LTDF LT. We let ρ: $\mathbb{N} \to \mathbb{N}$ be a polynomial that is a parameter of the scheme. The encryption and decryption algorithms of HE3 are shown in Fig. 9 and the key-generation algorithm HE3.Kg is the same as HE1.Kg. We let $\mathsf{HE3.rl} = \rho$. We also let $\mathsf{HE3.IL}(\cdot) = \mathbb{N}$, meaning the scheme encrypts variable and arbitrary length messages. While the scheme is quite simple it's challenging to find an analysis to match the desired bound $p(q+p) \cdot \text{Guess}_A(\cdot)$ for the reset-advantage in the IND-CDA setting. A naive analysis will end up in an inferior bound $q^2 p \cdot \text{Guess}_A(\cdot)$. Let $(m_1, r_1), \ldots, (m_p, r_p)$ be the message-coin pairs specified by A's IND-CDA queries. The reset adversary R is given a random oracle RO that on input (x, ℓ), returns a random string of length ℓ. Let Bad be the event that R queries $y \leftarrow \mathsf{RO}(m_k, \rho(\lambda))$ and then queries $\mathsf{RO}(y \oplus r_k, \mathsf{F.kl}(\lambda) + \lambda)$ for some $k \le p$. For HE3 to be IND-CDA secure, Bad must not occur. Suppose the reset adversary R queries $\mathsf{RO}(x_1, \rho(\lambda)), \ldots, \mathsf{RO}(x_{\lfloor q/2 \rfloor}, \rho(\lambda))$, and then queries $\mathsf{RO}(z_1, \mathsf{F.kl}(\lambda) + \lambda), \ldots, \mathsf{RO}(z_{\lfloor q/2 \rfloor}, \mathsf{F.kl}(\lambda) + \lambda)$. If there are $i, j \le \lfloor q/2 \rfloor$ and $k \le p$ such that $x_i = m_k$ and $\mathsf{RO}(x_i, \rho(\lambda)) \oplus z_j = r_k$ then Bad occurs. This seems to happen with probability $\frac{q^2 p}{4} \text{Guess}_A(\cdot)$, because R can *adaptively* choose z_j after seeing $\mathsf{RO}(x_1, \rho(\lambda)), \ldots, \mathsf{RO}(x_{\lfloor q/2 \rfloor}, \rho(\lambda))$.

To tackle this problem, we exploit a combinatorial technique on the coin length ρ—a parameter that we fully control. From Lemma 6 below, the chance that Bad occurs is at most $qp \cdot \text{Guess}_A(\cdot) + q^2 p \cdot 2^{-\rho(\lambda)/3}$. If ρ is large enough, say $\rho(\lambda) \ge 4.5\lambda$ for every $\lambda \in \mathbb{N}$, then this matches the optimal bound. The proof of Lemma 6 is in [7].

Lemma 6. *Let U, V be random variables over $\{0,1\}^*$ and $\{0,1\}^\ell$, respectively. Assume that the maximum, over all u, v, of $\Pr[(U,V) = (u,v)]$, is at most ϵ. Let* RO *be a random oracle and let $W = \mathsf{RO}(U, \ell) \oplus V$. For any adversary A that makes at most q queries to* RO*, the probability that the first component of one of A's* RO *queries is W is at most $q\epsilon + q^2 \cdot 2^{-\ell/3}$.* ■

Theorem 7 below confirms that HE3 is H-IND secure with very good concrete security bounds. While $\mathsf{UCE}[\mathcal{S}^{\text{sup}}]$ is enough for IND-CPA security, IND-CDA requires the stronger $\mathsf{UCE}[\mathcal{S}^{\text{srs}}]$ assumption. The proof is in [7].

Theorem 7. *Let* F *be a PRF with* $\mathsf{F.OL}(\cdot) = \mathbb{N}$ *and fixed input length* F.il*. Let* H *be a* $\mathsf{UCE}[\mathcal{S}^{\text{srs}}]$ *function family with* $\mathsf{H.IL}(\cdot) = \mathsf{H.OL}(\cdot) = \mathbb{N}$*. Let* LT *be an LTDF*

$\underline{\mathsf{HE3.Enc}(1^\lambda,(ek,hk),m;r)}$	$\underline{\mathsf{HE3.Dec}(1^\lambda,(dk,hk),(trap,c))}$
$w \leftarrow \mathsf{H.Ev}(1^\lambda, hk, m, 1^{\lvert r\rvert}) \oplus r$	$x \leftarrow \mathsf{LT.Inv}(1^\lambda, dk, trap)$
$x \leftarrow \mathsf{H.Ev}(1^\lambda, hk, w, 1^{\mathsf{LT.il}(\lambda)})$	$seed \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^{\mathsf{F.kl}(\lambda)+\lambda})$
$trap \leftarrow \mathsf{LT.Ev}(1^\lambda, ek, x)$	$y \leftarrow seed[1,\lambda]$; $fk \leftarrow seed[\lambda+1, \lvert seed\rvert]$
$seed \leftarrow \mathsf{H.Ev}(1^\lambda, hk, x, 1^{\mathsf{F.kl}(\lambda)+\lambda})$	$mask \leftarrow \mathsf{F.Ev}(1^\lambda, fk, 0^{\mathsf{F.il}(\lambda)}, 1^{\lvert c\rvert})$
$y \leftarrow seed[1,\lambda]$; $fk \leftarrow seed[\lambda+1, \lvert seed\rvert]$	$m \leftarrow \mathsf{H.Ev}(1^\lambda, hk, y, 1^{\lvert c\rvert}) \oplus mask \oplus c$
$mask \leftarrow \mathsf{F.Ev}(1^\lambda, fk, 0^{\mathsf{F.il}(\lambda)}, 1^{\lvert m\rvert})$	Return m
$c \leftarrow \mathsf{H.Ev}(1^\lambda, hk, y, 1^{\lvert m\rvert}) \oplus mask \oplus m$	
Return $(trap, c)$	

Fig. 9. Encryption and decryption algorithms of HE3, where H is a $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ family, F is a VOL PRF and LT is a LTDF

such that $\mathsf{LT.il}(\lambda) \geq \lambda$ *for all* $\lambda \in \mathbb{N}$. *Let* $\rho\colon \mathbb{N} \to \mathbb{N}$ *be a polynomial such that* $\rho(\lambda) \geq \lambda$ *for all* $\lambda \in \mathbb{N}$. *Let* HE3 *be defined from* $\mathsf{F}, \mathsf{H}, \mathsf{LT}, \rho$ *as above.*

Asymptotic result: HE3 *is H-IND secure.*

Concrete IND-CPA result: Let A be an adversary and P be a predictor. We can construct a source S, a distinguisher D, adversaries B and C such that

$$\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{HE3},A}(\cdot) \leq 2\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\cdot) + 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\cdot) + 2\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{F},C}(\cdot) + 2^{1-\rho}$$

$$\mathsf{Adv}^{\mathsf{pred}}_{S,P}(\cdot) \leq \frac{2q}{2^\rho} + \frac{q}{2^\tau}$$

where q is the maximum of the size of P's output in the execution of Pred^P_S *and* τ *is the lossiness of* LT. *Furthermore,* $\mathbf{T}(\mathrm{Lossy}^B_{\mathsf{LT}}), \mathbf{T}(\mathrm{UCE}^{S,D}_{\mathsf{H}}), \mathbf{T}(\mathrm{PRF}^C_{\mathsf{F}}) \leq \mathbf{T}(\mathrm{CPA}^A_{\mathsf{HE3}})$; $\mathbf{Q}^{\mathrm{RR}}_C = 1$; *and* $\mathbf{Q}^{\mathrm{Hash}}_S = 2$.

Concrete IND-CDA result: Let A be an adversary and R be a predictor. We can construct a source S, a distinguisher D, adversary B such that

$$\mathsf{Adv}^{\mathsf{cda}}_{\mathsf{HE},A}(\lambda) \leq 2\mathsf{Adv}^{\mathsf{ltdf}}_{\mathsf{LT},B}(\lambda) + 2\mathsf{Adv}^{\mathsf{uce}}_{\mathsf{H},S,D}(\lambda) + p^2 \cdot \mathrm{Guess}_A(\lambda) + \frac{8p^2}{2^\lambda} + \frac{12p^2}{2^{\min\{\tau(\lambda),\rho(\lambda)\}}}$$

$$\mathsf{Adv}^{\mathsf{reset}}_{S,R}(\lambda) \leq p(p+q) \cdot \mathrm{Guess}_A(\lambda) + \frac{5p^2}{2^\lambda} + \frac{6.5p^2}{2^{\min\{\tau(\lambda),\rho(\lambda)\}}} + \frac{pq^2}{2^{\rho(\lambda)/3}}$$

where p is the maximum of the total number of messages that A produces in the execution of $\mathrm{CDA}^A_{\mathsf{HE3}}$, $q = \mathbf{Q}^{\mathrm{Hash}}_R$, *and* τ *is the lossiness of* LT. *Furthermore,* $\mathbf{T}(\mathrm{Lossy}^B_{\mathsf{LT}}), \mathbf{T}(\mathrm{UCE}^{S,D}_{\mathsf{H}}) \leq \mathbf{T}(\mathrm{CDA}^A_{\mathsf{HE3}})$; *and* $\mathbf{Q}^{\mathrm{Hash}}_S \leq 3p$. ■

5 Efficiency and Comparisons with Prior Schemes

Our schemes improve on prior work on both the theoretical and practical fronts. On the theoretical front, DE1 is the first standard-model D-PKE scheme that is

fully IND secure and HE1, HE2, HE3 are the first standard-model PKE schemes achieving full H-IND, meaning IND-CPA plus *full* IND-CDA. Prior standard-model D-PKE (resp. PKE) schemes only achieved IND (resp. IND-CDA) for block sources, which assumes messages (resp. message-randomness pairs) are unpredictable even given prior ones, which is unlikely to be true in applications.

On the practical front, prior standard-model schemes fix a message length, create keys depending on it, and use only asymmetric operations, making them inflexible and inefficient. Our schemes handle variable input length messages with hybrid-encryption like efficiency, meaning the asymmetric cost is fixed and one pays only in hashing as message length grows. Exploiting fast instantiations of $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ and $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ functions [9,45], this yields high performance.

To elaborate, recall that asymmetric primitives are orders of magnitude more expensive than symmetric ones. Crucial to making IND-CPA PKE efficient is the hybrid encryption paradigm as represented by the KEM-DEM framework [25]. Here, $\mathsf{PKE.Enc}(1^\lambda, ek, m)$ uses its coins to generate a random symmetric key K along with an encapsulation c_a of K under ek, and returns ciphertext (c_a, c_s) where c_s is a symmetric encryption of m under K. The asymmetric cost is thus fixed regardless of message length and is amortized out for long messages. Ideally, we would like a similar generic hybrid encryption paradigm for D-PKE and H-PKE. But, despite interest and search, this has not been found. The reason in part is the apparently crucial use of randomness in the choice of K. As a result, prior standard-model D-PKE and H-PKE schemes have used only asymmetric operations. This has resulted not only in fixed message lengths but in costs that are exorbitant for long messages.

Our methods and schemes change this. Although we do not provide a generic hybrid encryption paradigm for these domains, our DE1, HE2 and HE3 schemes achieve hybrid-encryption like performance, meaning the asymmetric cost is fixed regardless of message length, and one pays only in symmetric operations — in our case this means hashing via the $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ or $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ functions— as the message length grows.

To capitalize on this for performance, good and careful instantiation of the UCE hash functions is needed. We need UCE functions H that are both VIL —variable input length, $\mathsf{H.IL}(\cdot) = \mathbb{N}$— and VOL —variable output length, $\mathsf{H.OL}(\cdot) = \mathbb{N}$. We now discuss how best to obtain these.

A simple instantiation of a UCE family is based on HMAC-SHA-256, as suggested in [8] and justified in [45]. While this yields a VIL family, it is FOL (fixed output length). A method to turn FOL UCE families into VOL ones is given in [8], but is slow. A better and faster transform is provided in [7]. With this we get $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ and $\mathsf{UCE}[\mathcal{S}^{\mathrm{srs}}]$ families with very good performance. These suffice for DE1, HE1 and HE3.

But one can do even better. BHK2 [9] provide a fast FIL, VOL $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ function $\overline{\mathsf{H}}$ based on AES. They also provide a fast AU family $\overline{\mathsf{U}}$. Applying their AU-then-Hash transform will return a VIL, VOL $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family H that is significantly faster than the HMAC-SHA-256 based instantiation. This suffices for DE1 and HE1.

$\underline{\mathsf{UE.Kg}(1^\lambda)}$	$\underline{\mathsf{UE.Enc}(1^\lambda, ek, m)}$	$\underline{\mathsf{UE.Dec}(1^\lambda, (ek, dk), c)}$
$(ek, dk) \leftarrow\!\!\$\ \mathsf{DE.Kg}(1^\lambda)$	$c \leftarrow \mathsf{DE.Enc}(ek, m)$	$m \leftarrow \mathsf{DE.Dec}(dk, c)$
Return $(ek, (ek, dk))$	Return c	If $m \neq \bot$ then
		$\quad c' \leftarrow \mathsf{DE.Enc}(ek, m)$
		$\quad$ If $c' \neq c$ then return $\bot$
		Return m

Fig. 10. U-PKE scheme $\mathsf{UE} = \mathsf{UniqueCtx}[\mathsf{DE}]$ constructed from D-PKE scheme DE

Recall $\mathsf{HE2}$ needs a $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family H of a special form, but it is based on $\mathsf{AU\text{-}then\text{-}Hash}$ and thus amenable to an efficient instantiation. Start again from $\overline{\mathsf{H}}, \overline{\mathsf{U}}$ from BHK2 as above. This time turn $\overline{\mathsf{U}}$ into U via our $\mathsf{Hash\text{-}then\text{-}Mask}$ transform —this preserves performance— and apply $\mathsf{AU\text{-}then\text{-}Hash}$ to this to get H. This $\mathsf{UCE}[\mathcal{S}^{\mathrm{sup}}]$ family is again exceptionally fast and of the special form required for $\mathsf{HE2}$.

6 Unique-Ciphertext PKE

In an algorithm-substitution attack (ASA) [12], the prescribed encryption algorithm is replaced with a subverted one that may attempt to leak information about the message to "big brother." The latter and the subverted algorithm may even share a key based on which they communicate. BPR [12] formalize the attacker goal in an ASA as compromising privacy while evading detection, the latter meaning that subverted ciphertexts are indistinguishable from real ones even given the decryption key. They focus on the symmetric setting. They give attacks showing that randomized, stateless schemes will succumb to attack. They show however that security against ASAs may be achieved by what they call unique ciphertext symmetric encryption schemes.

BPR [12] initiate the study of ASAs for PKE. Continuing that theme, we define unique ciphertext PKE. We say that a PKE scheme PKE has unique ciphertexts, or is a U-PKE scheme, if for every $\lambda \in \mathbb{N}$, every $(ek, dk) \in [\mathsf{PKE.Kg}(1^\lambda)]$, and every message m, there is at most one ciphertext $c \in \{0,1\}^*$ such that $\mathsf{PKE.Dec}(1^\lambda, dk, c) \neq \bot$. Coupled with correctness, this means that for every $\lambda \in \mathbb{N}$, every $(ek, dk) \in [\mathsf{PKE.Kg}(1^\lambda)]$ and every $m \in \{0,1\}^*$ with $|m| \in \mathsf{PKE.IL}(\lambda)$ the set $[\mathsf{PKE.Enc}(1^\lambda, ek, m)]$ has size exactly one. The latter means that a unique ciphertext scheme is deterministic, meaning a D-PKE scheme.

We now ask how to design a U-PKE scheme. The natural thought is that any D-PKE scheme is a U-PKE scheme. This is not true. As an example, take any IND D-PKE scheme, and modify it so that encryption pre-pends a bit to the ciphertext that is ignored by decryption. This is still an IND D-PKE scheme, but it does not have unique ciphertexts, because if c is the encryption of m under $1^\lambda, ek$ in the starting D-PKE scheme then both $0 \,\|\, c$ and $1 \,\|\, c$ are valid ciphertexts in the new D-PKE scheme.

However, we show that one can transform any given D-PKE scheme DE into a U-PKE scheme UE. The U-PKE public key is the same as the D-PKE one,

but the secret key is the pair (ek, dk) consisting of the D-PKE public key and matching secret key. Encryption is as in D-PKE. U-PKE decryption of ciphertext c first recovers the candidate message m via D-PKE decryption of c under dk and then checks that re-encrypting m under ek yields c, rejecting otherwise. $\mathsf{UE} = \mathsf{UniqueCtx}[\mathsf{DE}]$ is is formally specified in Fig. 10.

The security requirement for U-PKE contains to be IND, meaning a U-PKE scheme is treated just as a D-PKE scheme in the context of security. Applying our UniqueCtx to DE1 thus yields a very efficient IND U-PKE scheme.

In the symmetric setting, unique-ciphertext encryption could be stateful and thus attain IND-CPA type security [12]. Here, a synchronized state is shared between sender and receiver. In the PKE setting, however, it is does not seem practical to assume that the sender and receiver share a synchronized state. Indeed, this would go against the spirit of public-key cryptography. As a consequence, for the benefit of unique ciphertexts, security must drop compared to IND-CPA, meaning we pay in security to protect against ASAs.

Acknowledgments. Bellare is supported in part by NSF grants CNS-1116800 and CNS-1228890. Hoang is supported in part by NSF grant 1223623. Part of the work was done when Hoang was working at UCSD, supported in part by NSF grants CNS-1116800 and CNS-1228890.

References

1. Austrin, P., Chung, K.-M., Mahmoody, M., Pass, R., Seth, K.: On the impossibility of cryptography with tamperable randomness. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 462–479. Springer, Heidelberg (2014)
2. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer, Heidelberg (2001)
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., Dowsley, R., Keelveedhi, S.: How secure is deterministic encryption? In: Public-Key Cryptography-PKC 2015. Springer (2015)
6. 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)
7. Bellare, M., Hoang, V.T.: Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. Cryptology ePrint Archive, Report 2014/786 (2014)
8. Bellare, M., Hoang, V.T., Keelveedhi, S.: Instantiating random oracles via UCEs. Cryptology ePrint Archive, Report 2013/424 (2013). Preliminary version in CRYPTO 2013

9. Bellare, M., Hoang, V.T., Keelveedhi, S.: Cryptography from compression functions: The UCE bridge to the ROM. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 169–187. Springer, Heidelberg (2014)
10. Bellare, M., Hoang, V.T., Keelveedhi, S.: Cryptography from compression functions: The UCE bridge to the ROM. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 169–187. Springer, Heidelberg (2014)
11. Bellare, M., Kiltz, E., Peikert, C., Waters, B.: Identity-based (lossy) trapdoor functions and applications. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 228–245. Springer, Heidelberg (2012)
12. Bellare, M., Paterson, K.G., Rogaway, P.: Security of symmetric encryption against mass surveillance. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 1–19. Springer, Heidelberg (2014)
13. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, pp. 62–73. ACM Press, November 1993
14. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006)
15. Birrell, E., Chung, K.-M., Pass, R., Telang, S.: Randomness-dependent message security. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 700–720. Springer, Heidelberg (2013)
16. Boldyreva, A., Chenette, N., Lee, Y., O'Neill, A.: Order-preserving symmetric encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 224–241. Springer, Heidelberg (2009)
17. 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)
18. Bosley, C., Dodis, Y.: Does privacy require true randomness? In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 1–20. Springer, Heidelberg (2007)
19. 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)
20. Brown, D.R.L.: A weak-randomizer attack on RSA-OAEP with e = 3. Cryptology ePrint Archive, Report 2005/189 (2005). http://eprint.iacr.org/2005/189
21. Brzuska, C., Farshim, P., Mittelbach, A.: Indistinguishability obfuscation and UCEs: The case of computationally unpredictable sources. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 188–205. Springer, Heidelberg (2014)
22. Brzuska, C., Farshim, P., Mittelbach, A.: Random oracle uninstantiability from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2014/867 (2014). http://eprint.iacr.org/2014/867
23. Cachin, C., Micali, S., Stadler, M.A.: Computationally private information retrieval with polylogarithmic communication. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 402–414. Springer, Heidelberg (1999)
24. Checkoway, S., Niederhagen, R., Everspaugh, A., Green, M., Lange, T., Ristenpart, T., Bernstein, D.J., Maskiewicz, J., Shacham, H., Fredrikson, M.: On the practical exploitability of dual EC in TLS implementations. In: Proceedings of the 23rd USENIX Security Symposium, pp. 319–335, August 2014
25. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM Journal on Computing **33**(1), 167–226 (2003)

26. Dodis, Y., López-Alt, A., Mironov, I., Vadhan, S.: Differential privacy with imperfect randomness. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 497–516. Springer, Heidelberg (2012)
27. Dodis, Y., Ong, S.J., Prabhakaran, M., Sahai, A.: On the (im)possibility of cryptography with imperfect randomness. In: 45th FOCS, pp. 196–205. IEEE Computer Society Press, October 2004
28. Dodis, Y., Pointcheval, D., Ruhault, S., Vergnaud, D., Wichs, D.: Security analysis of pseudo-random number generators with input: /dev/random is not robust. Cryptology ePrint Archive, Report 2013/338 (2013). http://eprint.iacr.org/2013/338
29. Dorrendorf, L., Gutterman, Z., Pinkas, B.: Cryptanalysis of the windows random number generator. In: Ning, P., di Vimercati, S.D.C., Syverson, P.F. (eds.) ACM CCS 2007, pp. 476–485. ACM Press, October 2007
30. Escala, A., Herranz, J., Libert, B., Ràfols, C.: Identity-based lossy trapdoor functions: new definitions, hierarchical extensions, and implications. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 239–256. Springer, Heidelberg (2014)
31. Feltz, M., Cremers, C.: On the limits of authenticated key exchange security with an application to bad randomness. Cryptology ePrint Archive, Report 2014/369 (2014). http://eprint.iacr.org/2014/369
32. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More constructions of lossy and correlation-secure trapdoor functions. Journal of Cryptology **26**(1), 39–74 (2013)
33. Fuller, B., O'Neill, A., Reyzin, L.: A unified approach to deterministic encryption: new constructions and a connection to computational entropy. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 582–599. Springer, Heidelberg (2012)
34. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS, pp. 40–49. IEEE Computer Society Press, October 2013
35. Gentry, C., Lewko, A., Sahai, A., Waters, B.: Indistinguishability obfuscation from the multilinear subgroup elimination assumption. Cryptology ePrint Archive, Report 2014/309 (2014). http://eprint.iacr.org/2014/309
36. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. Journal of the ACM **33**(4), 792–807 (1986)
37. Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences **28**(2), 270–299 (1984)
38. Green, M.D., Katz, J., Malozemoff, A.J., Zhou, H.-S.: A unified approach to idealized model separations via indistinguishability obfuscation. Cryptology ePrint Archive, Report 2014/863 (2014). http://eprint.iacr.org/2014/863
39. Gutterman, Z., Pinkas, B., Reinman, T.: Analysis of the linux random number generator. In: 2006 IEEE Symposium on Security and Privacy, pp. 371–385. IEEE Computer Society Press, May 2006
40. Hemenway, B., Ostrovsky, R.: Building lossy trapdoor functions from lossy encryption. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 241–260. Springer, Heidelberg (2013)
41. Heninger, N., Durumeric, Z., Wustrow, E., Halderman, J.A.: Mining your Ps and Qs: Detection of widespread weak keys in network devices. In: Proceedings of the 21st USENIX Security Symposium, pp. 205–220, August 2012
42. Kamara, S., Katz, J.: How to encrypt with a malicious random number generator. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 303–315. Springer, Heidelberg (2008)

43. Kiltz, E., O'Neill, A., Smith, A.: Instantiability of RSA-OAEP under chosen-plaintext attack. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 295–313. Springer, Heidelberg (2010)
44. Lenstra, A.K., Hughes, J.P., Augier, M., Bos, J.W., Kleinjung, T., Wachter, C.: Public keys. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 626–642. Springer, Heidelberg (2012)
45. Mittelbach, A.: Salvaging indifferentiability in a multi-stage setting. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 603–621. Springer, Heidelberg (2014)
46. Ouafi, K., Vaudenay, S.: Smashing SQUASH-0. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 300–312. Springer, Heidelberg (2009)
47. Paterson, K.G., Schuldt, J.C.N., Sibborn, D.L.: Related randomness attacks for public key encryption. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 465–482. Springer, Heidelberg (2014)
48. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C., (eds.) 40th ACM STOC, pp. 187–196. ACM Press, May 2008
49. Raghunathan, A., Segev, G., Vadhan, S.: Deterministic public-key encryption for adaptively chosen plaintext distributions. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 93–110. Springer, Heidelberg (2013)
50. Ristenpart, T., Yilek, S.: When good randomness goes bad: Virtual machine reset vulnerabilities and hedging deployed cryptography. In: NDSS 2010. The Internet Society, February / March 2010
51. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006)
52. Vergnaud, D., Xiao, D.: Public-key encryption with weak randomness: Security against strong chosen distribution attacks. Cryptology ePrint Archive, Report 2013/681 (2013). http://eprint.iacr.org/2013/681
53. Wichs, D.: Barriers in cryptography with weak, correlated and leaky sources. In: Kleinberg, R.D. (ed.). In: ITCS 2013, pp. 111–126. ACM, January 2013
54. Yang, G., Duan, S., Wong, D.S., Tan, C.H., Wang, H.: Authenticated key exchange under bad randomness. Cryptology ePrint Archive, Report 2011/688 (2011). http://eprint.iacr.org/2011/688
55. Yilek, S.: Resettable public-key encryption: How to encrypt on a virtual machine. In: Pieprzyk, J. (ed.) CT-RSA 2010. LNCS, vol. 5985, pp. 41–56. Springer, Heidelberg (2010)
56. Young, A., Yung, M.: Kleptography: Using cryptography against cryptography. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 62–74. Springer, Heidelberg (1997)