

Nonce-Based Cryptography: Retaining Security when Randomness Fails

Mihir Bellare¹ and Björn Tackmann¹

Department of Computer Science and Engineering
University of California San Diego, USA

<http://cseweb.ucsd.edu/~mihir/>, <http://cseweb.ucsd.edu/~btackmann/>

Abstract. We take nonce-based cryptography beyond symmetric encryption, developing it as a broad and practical way to mitigate damage caused by failures in randomness, whether inadvertent (bugs) or malicious (subversion). We focus on definitions and constructions for nonce-based public-key encryption and briefly treat nonce-based signatures. We introduce and construct hedged extractors as a general tool in this domain. Our nonce-based PKE scheme guarantees that if the adversary wants to violate IND-CCA security then it must do *both* of the following: (1) fully compromise the RNG (2) penetrate the sender system to exfiltrate a seed used by the sender.

1 Introduction

An old security adage says there is no point putting strong locks on the door if you leave the window open. The lock here is modern public-key encryption, proven to meet the strong IND-CCA goal. The window is the assumption made in these proofs that, at every encryption, the encryptor has access to perfect, fresh (independent from prior) randomness. To allow encryption to fulfill in practice the promise it makes in theory, we must close the window. This paper develops nonce-based public-key encryption as a practical way to do this. It goes on to develop nonce-based digital signatures.

RANDOMNESS. That randomness failures occur and lead to cryptographic failures is by now very well known and does not need to be belabored. The news of interest is perhaps that it is getting worse. Let us explain. There are two sources of randomness failures. The first, which has been with us a while and is not going away, is *bugs*. A good example is the Debian Linux vulnerability present from September 2006 to May 2008 where a programmer removed some lines of code from the OpenSSL source, resulting in there being only 15 bits of entropy in the seed for the PRNG [1]. HDWH [19] finds cryptographic vulnerabilities in a significant fraction of TLS and SSH servers due to what they call “malfunctioning RNGs.” And the list goes on. The second source of randomness failures, which may have been with us for a while but of which we have learned only recently due to the mass-surveillance revelations, is *subversion*, the deliberate and targeted attempt to weaken randomness. At the South by South West (SXSW) 2014

conference, 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.” The prime example here is Dual EC, a RNG the NSA designed to have a backdoor and then pushed into standards and adoption. The ability to compromise security in practice via the backdoor has been demonstrated in [14].

PRIOR WORK. The basic definitions of security for public-key encryption (PKE), namely IND-CPA [18] and IND-CCA [5, 15], provide no guarantees if the randomness is bad. There is a long line of work giving new definitions of security that do provide such guarantees, and building PKE schemes to meet these definitions.

The simplest way to avoid vulnerabilities due to poor randomness is to not use randomness at all. Deterministic PKE [3, 12, 6, 17] however only provides security when the messages have high min-entropy. This limits utility (for example, we may want to encrypt votes, which have low min-entropy) and, although in a different context, CGPR [13] show that in practice the entropy of “real” data is often quite low. Hedged PKE [4, 7, 27] extends Deterministic PKE to provide privacy as long as the message and randomness together have sufficient min-entropy. This is a significant benefit and we recommend that one use Hedged PKE whenever possible. But the limitations remain. Since messages regularly do not in fact have entropy [13], and the “randomness” can be entirely predictable (this happened both with the Debian Linux bug and the Dual EC subversion), the message and randomness together still may not have enough entropy for hedged encryption to provide security. A further limitation of both Deterministic and Hedged PKE is that security is only provided for messages that do not depend on the public key. (This second limitation can be partially addressed but at some cost [26].) Yilek [32] defines and achieves security against randomness-reset attacks, where the randomness is perfect but the adversary can force its re-use across different encryptions. This is useful in the context of virtual machine resets but not more broadly. PSS [24] introduce related-randomness attacks, where encryption is under adversary-specified functions of some initial uniform randomness. However, negative results they provide show that for many functions one cannot achieve security.

In summary, all these notions have some limitations and the practical benefit they provide is not clear. Most importantly, these were all designed in the older mindset of RNG failures due to bugs and can break down severely when the RNG is subverted. The latter is the new reality against which we need to defend.

NONCES. Rogaway [28, 29] introduced nonce-based symmetric-key encryption, where the encryption algorithm is deterministic, taking the shared key, message and a quantity called a nonce. Security is provided as long as the nonce does not repeat. The notion was strengthened by Rogaway and Shrimpton (RS) [30]. Rogaway suggests that packet sequence numbers may play the role of the nonce. Motivations he provides include reducing implementor error and achieving stronger notions of security. We suggest that nonces can be used much more broadly and

are a good defense against poor randomness. The goal we pursue in depth is nonce-based PKE.

THE OBVIOUS EXTENSION FAILS. Towards this we begin by noting that a direct extension of nonce-based symmetric-key encryption as defined in [28–30] to the public-key setting does not work. Such an extension would have the encryption algorithm \mathcal{E} be deterministic, taking input the encryption key ek , message m and a nonce n to return the ciphertext $C = \mathcal{E}(ek, m, n)$. The privacy game would give the adversary an oracle that takes messages m_0, m_1 and a nonce n (in the definitions of [28–30], nonces are adversarially chosen subject to not repeating) to return $C = \mathcal{E}(ek, m_b, n)$, where b is a challenge bit chosen at random by the game. Security would require that the adversary has little advantage in guessing b . But such security would not be achieved because in the public-key setting the adversary has ek and can itself encrypt. Thus it could query its oracle with any m_0, m_1, n of its choice to get back ciphertext $C = \mathcal{E}(ek, m_b, n)$, and itself compute $C_0 = \mathcal{E}(ek, m_0, n)$. If $C = C_0$ it knows that the challenge b is 0, else it is 1.

NONCE-BASED PKE. We define a nonce-based PKE scheme NPE as follows. The receiver runs key-generation algorithm NPE.Kg as usual to get an encryption key ek and decryption key dk . Not as usual, the sender begins by locally running a *seed generation* algorithm NPE.sKg to get a *seed* xk . The encryption algorithm NPE.Enc is *deterministic*, taking in addition to the usual ek and message m , two new inputs, a nonce n and the *seed* xk , and returning ciphertext $C = \text{NPE.Enc}(ek, xk, m, n)$. Decryption is unchanged, taking dk and a ciphertext C to return a message. The receiver does *not* need to know the seed, and the keys and seed are entirely independent. The sender can either re-use the same seed across multiple encryptions, or generate a fresh one at every encryption, or anything in between, and the receiver will be oblivious to all of this.

Security is captured via two games and corresponding requirements. Nonce-based privacy One (NBP1) asks that IND-CCA privacy be maintained as long as message-nonce pairs do not repeat. That is, the only way security fails is if, for the same message, a nonce is re-used. This is a very strong guarantee. However there is one caveat, namely this holds when the seed is kept private from the adversary. Nonce-based privacy Two (NBP2) addresses the possibility that the adversary compromises the sender’s system and obtains the seed. Even in this case, it guarantees IND-CCA privacy as long as nonces are unpredictable to the adversary. The formalizations are in terms of a stateful *nonce generator* NG that takes an adversary specified input η to return the next nonce, so that nonces are (indirectly) under adversary control.

In practice we would expect a combination of a variety of things to be used as the nonce, for example the current time (this does not repeat) and prior ciphertexts, but, also, randomness from the system RNG, since, for NBP2, nonce unpredictability is required. (This is a departure from the symmetric setting.) However guarantees in the face of poor randomness are much better than before, as we now explain.

WHAT THIS BUYS US. Intuitively our definitions are saying the following. Consider two cases. The expected and good case is that the sender seed stays private. In this case we get IND-CCA privacy *regardless of the quality of the randomness*, the only requirement being that message-nonce pairs do not repeat. The latter is a mild condition, such repetition being unlikely with reasonable nonces, even simply using the date and time as the nonce. The other case is that the sender’s system is compromised and the seed is exposed. In this case, we are effectively in the setting of standard PKE and we cannot deterministically provide IND-CCA. We guarantee that we do no worse than standard PKE, meaning we provide IND-CCA as long as the randomness (here part of the nonce) is good. But in fact we do better, since the requirement on nonces is only unpredictability. This means that we are safe even if the outputs of the RNG are correlated and structured, as long as they remain unpredictable.

Put another way, if a subvertor wanted to compromise privacy, it would not suffice to compromise the RNG. They would have to also break in to the sender’s system, find the seed, and exfiltrate it. Frequent rotation of seeds (which has effectively no cost) makes this even harder. This ups the ante. Now, it is true that with the NSA’s capabilities in malware, we should not under-estimate their ability to penetrate a target sender. But this would have to be done on per-sender basis, making mass surveillance harder.

HEDGED EXTRACTORS AND OUR SCHEME. It is easy to achieve either of NBP1 or NBP2 in isolation. We can get an NBP1 nonce-based PKE scheme by encrypting under a conventional (randomized) IND-CCA PKE scheme with the coins set to the result of a PRF keyed by the sender seed and applied to the message and nonce, but there is no reason this scheme would also be NBP2 secure. We can get an NBP2 nonce-based PKE scheme by encrypting under the conventional IND-CCA PKE scheme with the coins set to the result of an extractor keyed by the seed and applied to the (message and) nonce, but there is no reason this scheme would also be NBP1 secure. To simultaneously get both properties, we introduce and use *hedged extractors*.

A hedged extractor HE takes a seed (also called a key) xk , a message m and a nonce n to deterministically return a string $r = \text{HE}(xk, (m, n))$. It has two properties: (1) It is a PRF, meaning if xk is random and hidden then the outputs look random even to an adversary that picks m, n , and (2) it is an extractor, meaning if xk is random but known, then r looks random if (m, n) is unpredictable (meaning, has enough min-entropy). Again, achieving either goal in isolation is trivial. The task is to achieve them simultaneously, in the same construction. We give two solutions, one in the ROM, the other in the standard model. The first is easy but practical, and likely to be what we would use, namely to simply apply the RO to $xk, (m, n)$. The second combines a PRF with a strong randomness extractor via XOR. The ROM solution delivers optimal security, the standard-model one a bit less due to the inherent limitations of strong randomness extractors, namely that they are only guaranteed to work for seed-independent inputs that retain min-entropy even conditioned on prior inputs.

Our nonce-based PKE scheme is then simply defined via the same paradigm as for the in-isolation cases, namely we encrypt under a conventional (randomized) IND-CCA PKE scheme with the coins set to the result of a hedged extractor keyed by the sender seed and applied to the message and nonce. Both NBP1 and NBP2 security are proven for this scheme assuming IND-CCA of the conventional scheme and security of the hedged extractor.

DISCUSSION AND PRAGMATICS. We can view nonce-based cryptography as moving the traditional abstraction boundary between cryptography and system RNGs closer to the cryptography. The view is that, in the presence of bad RNGs, a safer and better architecture is that the cryptography take on as large a share of the burden of providing security as possible, in other words, rely on its environment as little as possible. Our suggestion here is that the environment is relied on only to produce nonces with relatively weak requirements. This view is in some ways the opposite of that represented by work that aims to strengthen RNGs against failure or subversion [16]. In practice the two can co-exist and their combined presence will increase security.

Our nonce-based encryption scheme is simple and modular, a way to transform any given conventional IND-CCA scheme into an NBP1+NBP2 secure nonce-based scheme. With a practical choice of hedged extractor such as our ROM one, we retain the efficiency attributes of the initial PKE scheme. In our scheme, decryption is unchanged. The decryptor does not need to change its software or even know that nonces are being used. These attributes make it easier to deploy nonce-based PKE as a practical defense against poor randomness.

In the above-discussed prior work aimed at increasing resistance of PKE to randomness failures, the *model* was unchanged in the sense that the object whose security was being considered continued in syntax to be a classical public-key encryption scheme as per [18]. Nonce-based encryption is a new model (because the sender has a seed) and a new syntax (there is a seed generation algorithm and the encryption algorithm is different). It is these changes, and in particular not just the nonce, but the combination of nonce and seed, that are a game changer and result in significantly better guarantees against poor randomness compared to prior work.

Picking a seed, like picking a key, does require (good) randomness. The viewpoint here, as in all the prior work discussed above, is that there is a difference between static and dynamic randomness usage. We assume good randomness for key generation because effort can be invested in it. Current key-generation software often has the user generate coins by waving their mouse around. Good seed generation would require similar effort, but one would expect to use a seed for some time so this effort is not frequent. This is flexible. If a seed is lost due to a system reboot, or compromised, the user can elect to make the effort to pick a new one. If you are encrypting from multiple systems (your desktop, laptop and phone) each can have its own, independently chosen seed.

NONCE-BASED SIGNATURES AND BEYOND. We define nonce-based signatures, where the signing algorithm is deterministic, and takes not only the signing key and message, but also a seed and nonce. We require that (1) if the seed remains

hidden then we have regular security (unforgeability) regardless of how nonces are generated, and (2) if the seed is exposed, then we have security as long as the nonces were unpredictable. Section 5 formalizes this and shows how to convert any signature scheme into a nonce-based one with these security properties using a hedged extractor.

Due to their speed and short signature sizes, the most attractive signature schemes for practice are elliptic-curve versions of DSA, El Gamal and Schnorr [31]. However, they are randomized, and fail spectacularly when the randomness is bad. Discussions on the `cfrg` forum show overwhelming support for making these schemes deterministic. This is easily done, by deriving the coins either as a PRF, keyed by a seed that is part of the secret key and applied to the message, or as a RO applied to the secret key and message [23, 21, 9, 25], and the popular Ed25519 signature scheme of [11] already embodies this. Making a scheme nonce-based complements this traditional de-randomization, retaining the benefits of deterministic signing while adding further ones. See Section 5 for more extensive background and discussion.

Nonces in combination with seeds can similarly be used in many other areas of cryptography to provide resilience in the face of poor randomness or even provide other gains. Our work aims to be illustrative rather than exhaustive.

RELATED WORK. BKS [8] introduce stateful PKE. Here also the sender can maintain a seed. They show that this leads to significant efficiency gains. Their schemes are however randomized, and there are no nonces. An interesting direction for future work is to combine their methods with ours to get similar efficiency gains for nonce-based PKE.

Rogaway [29] discusses nonces as “surfacing the IV.” As motivation, he says that when IVs are implicit, implementors and even books get things wrong. He says that often nonces are readily available, for example packet sequence numbers. He does not seem to explicitly mention robustness in the face of randomness failure as a goal in the symmetric case. Intriguingly, in the final section of the paper, he goes on to say: “... it makes just as much sense to consider nonce-based public-key encryption schemes as it does to consider nonce-based symmetric encryption schemes. This provides an approach to effectively weakening the requirement for randomness on the sender.” Our work has pursued this suggestion. It is surprising that this waited 12 years.

2 Notation and standard definitions

NOTATION. We let ε denote the empty string. If X is a finite set, we let $x \leftarrow_s X$ denote picking an element of X uniformly at random and assigning it to x . Algorithms may be randomized unless otherwise indicated. Running time is worst case. If A is an algorithm, we let $y \leftarrow A(x_1, \dots; r)$ denote running A with random coins r on inputs x_1, \dots and assigning the output to y . We let $y \leftarrow_s A(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We let $[A(x_1, \dots)]$ denote the set of all possible outputs of A when invoked with inputs x_1, \dots . We use the code based game playing framework of [10]. (See Fig. 1

Game $\mathbf{G}_F^{\text{prf}}(\mathcal{A})$	Game $\mathbf{G}_{\text{PE}}^{\text{ind}}(\mathcal{A})$	Game $\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{A})$
$fk \leftarrow_s \mathbf{F}.\text{Keys}$	$(ek, dk) \leftarrow_s \text{PE}.\text{Kg}$	$(sk, vk) \leftarrow_s \text{DS}.\text{Kg}$
$c \leftarrow_s \{0, 1\}$	$b \leftarrow_s \{0, 1\} ; C \leftarrow \emptyset$	$M \leftarrow \emptyset$
$c' \leftarrow_s \mathcal{A}^{\text{FN,RO}}$	$b' \leftarrow_s \mathcal{A}^{\text{ENC,DEC}}(ek)$	$(m, s) \leftarrow_s \mathcal{A}^{\text{SIG}}(vk)$
Return $(c = c')$	Return $(b = b')$	$v \leftarrow (\text{DS}.\text{Ver}(vk, m, s) = 1)$
<u>FN(x)</u>	<u>ENC(m_0, m_1)</u>	Return $(v \wedge (m \notin M))$
If $(c = 1)$ then	If $(m_0 \neq m_1)$ then return \perp	<u>SIG(m)</u>
$S[x] \leftarrow \mathbf{F}^{\text{RO}}(fk, x)$	$c \leftarrow_s \text{PE}.\text{Enc}(ek, m_b)$	$s \leftarrow \text{DS}.\text{Sig}(sk, m)$
Else	$C \leftarrow C \cup \{c\}$	$M \leftarrow M \cup \{m\}$
If $S[x] = \perp$ then	Return c	Return s
$S[x] \leftarrow_s \mathbf{F}.\text{Rng}$	<u>DEC(c)</u>	
Return $S[x]$	If $(c \in C)$ then return \perp	
<u>RO(x, l)</u>	$m \leftarrow \text{PE}.\text{Dec}(ek, dk, c)$	
If $T[x, l] = \perp$ then	Return m	
$T[x, l] \leftarrow_s \{0, 1\}^l$		
Return $T[x, l]$		

Fig. 1. Games for defining PRF security of a function family F , standard IND-CCA security of a standard PKE scheme PE and EUF-CMA security of a signature scheme DS.

for an example.) By $\text{Pr}[G]$ we denote the event that the execution of game G results in the game returning true. Random oracles are variable output length, represented by a game procedure RO that takes x, l and returns a random string of length l . The min-entropy of a random variable X over \mathcal{X} is defined as $H_\infty(X) = -\log(\max_{x \in \mathcal{X}}(\text{Pr}[X = x]))$.

FUNCTION FAMILIES. A family of functions $F: \mathbf{F}.\text{Keys} \times \mathbf{F}.\text{Dom} \rightarrow \mathbf{F}.\text{Rng}$ is a two-argument function that takes a key K in the key space $\mathbf{F}.\text{Keys}$, an input x in the domain $\mathbf{F}.\text{Dom}$ and returns an output $F(K, x)$ in the range $\mathbf{F}.\text{Rng}$. In the ROM, F takes an oracle RO.

PSEUDO-RANDOM FUNCTIONS. The security of F as a PRF is defined via game $\mathbf{G}_F^{\text{prf}}(\mathcal{A})$ that is associated to adversary \mathcal{A} and shown in Fig. 1. Here F could have access to a RO and thus the game is in the ROM. Tables S, T are assumed initially \perp everywhere. The advantage of \mathcal{A} is defined as $\mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) = 2 \text{Pr}[\mathbf{G}_F^{\text{prf}}(\mathcal{A})] - 1$.

PUBLIC-KEY ENCRYPTION. A *public-key encryption scheme* PE specifies the following. Receiver key-generation algorithm $\text{PE}.\text{Kg}$ returns an encryption key ek and associated decryption key dk . Encryption algorithm $\text{PE}.\text{Enc}$ takes ek and message $m \in \{0, 1\}^*$ to return a ciphertext c . Deterministic decryption algorithm $\text{PE}.\text{Dec}$ takes ek, dk and ciphertext c to return a value in $\{0, 1\}^* \cup \{\perp\}$,

and we require standard decryption correctness. The advantage of an adversary \mathcal{A} in breaking the IND-CCA security of PE is defined as $\mathbf{Adv}_{\text{PE}}^{\text{ind}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_{\text{PE}}^{\text{ind}}(\mathcal{A})] - 1$ for the game $\mathbf{G}_{\text{PE}}^{\text{ind}}(\mathcal{A})$ described in Fig. 1. This represents a conventional (not nonce-based scheme), and thus PE.Enc is randomized. We will use such schemes as base schemes and we will need to surface their coins, writing $c \leftarrow \text{PE.Enc}(ek, m; r)$ to mean that PE.Enc is run with coins r to deterministically return c . The length of the coins is denoted PE.rl .

DIGITAL SIGNATURES. A digital signature scheme DS specifies the following. Signer key-generation algorithm DS.Kg returns a signature key sk and a verification key vk . Signing algorithm DS.Sig takes sk and message $m \in \{0, 1\}^*$ to return a signature $s \in \{0, 1\}^{\text{DS.ol}}$. Verification algorithm DS.Ver takes vk , message $m \in \{0, 1\}^*$, and signature $s \in \{0, 1\}^{\text{DS.ol}}$, to return a bit $b \in \{0, 1\}$. The advantage of an adversary \mathcal{A} in breaking the EUF-CMA security of DS is defined as $\mathbf{Adv}_{\text{DS}}^{\text{uf}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{A})]$ for the game $\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{A})$ described in Fig. 1. Again we may need to surface the coins $r \in \{0, 1\}^{\text{DS.rl}}$ of DS.Sig , writing $s \leftarrow \text{DS.Sig}(sk, m; r)$.

3 Hedged Extractors

Our nonce-based schemes work simply by supplying coins to a base scheme via a *hedged extractor* keyed by the sender seed. This primitive, that we introduce and build here, is a function family that has two security properties. The first is that it is a PRF. The second, which we define and call *ror* (real or random) security, formalizes randomness of outputs when the key (seed) is known. Clearly this can only be achieved with some restrictions, and the type of *ror* security achieved will vary across constructions, from the “best possible” achieved by our ROM construction to a weaker, but we think still meaningful, version for our standard model construction. To make the goals precise we first introduce the notion of a nonce generator.

NONCE GENERATORS. A nonce generator is an algorithm NG that, on input a *nonce selector* η and a current state St , returns a nonce n , belonging to the range set $\text{NG.Rng} \subseteq \{0, 1\}^*$ of NG , together with an updated state, written $(n, St) \leftarrow \text{NG}(\eta, St)$. We say the generator has nonce length $\text{NG.nl} \in \mathbb{N}$ if $\text{NG.Rng} = \{0, 1\}^{\text{NG.nl}}$. Let \mathcal{P} be an adversary called a predictor and consider game $\mathbf{G}_{\text{NG}}^{\text{pred}}(\mathcal{P})$ of Fig. 2. Let

$$\mathbf{Adv}_{\text{NG}}^{\text{pred}}(\mathcal{P}) = \Pr[\mathbf{G}_{\text{NG}}^{\text{pred}}(\mathcal{P})] \quad \text{and} \quad \mathbf{Adv}_{\text{NG}}^{\text{pred}}(q_1, q_2) = \max_{\mathcal{P}} \mathbf{Adv}_{\text{NG}}^{\text{pred}}(\mathcal{P}),$$

where the maximum is over all \mathcal{P} making at most $q_1 \in \mathbb{N}$ queries to GEN and $q_2 \in \{0, 1\}$ queries to EXPOSE . Now let us explain. A call to GEN generates the next nonce in the sequence, returning nothing to the adversary. The adversary can influence the choice of nonces through its choice of the selector η . The EXPOSE oracle allows to additionally get access to the state of the nonce generator. To win, the adversary needs to guess some generated nonce or create a collision between generated nonces.

Game $\mathbf{G}_{\text{NG}}^{\text{pred}}(\mathcal{P})$	Game $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})$
$St \leftarrow \varepsilon ; s \leftarrow 0 ; N \leftarrow \emptyset$	$St \leftarrow \varepsilon ; xk \leftarrow \text{HE.Keys} ; c \leftarrow \{0, 1\} ; s \leftarrow 0$
$n \leftarrow \mathcal{P}^{\text{GEN,EXPOSE}}$	$c' \leftarrow \mathcal{G}^{\text{RoR,EXPOSE,RO}}(xk)$
Return $((n \in N) \text{ OR } \text{coll})$	Return $(c = c')$
<u>GEN(η)</u>	<u>RoR(m, η)</u>
If $s = 1$ then return \perp	If $s = 1$ then return \perp
$(n, St) \leftarrow \text{NG}(\eta, St)$	$(n, St) \leftarrow \text{NG}(\eta, St)$
If $(n \in N)$ then $\text{coll} \leftarrow \text{true}$	If $(c = 1)$ then $r \leftarrow \text{HE}^{\text{RO}}(xk, (m, n))$
$N \leftarrow N \cup \{n\}$	Else $r \leftarrow \text{HE.Rng}$
<u>EXPOSE</u>	Return r
$s \leftarrow 1$	<u>EXPOSE</u>
Return St	$s \leftarrow 1 ; \text{Return } St$
	<u>RO(x, l)</u>
	If $T[x, l] = \perp$ then $T[x, l] \leftarrow \{0, 1\}^l$
	Return $T[x, l]$

Fig. 2. Games for defining predictability of the nonce generator NG and real-or-random security of function family HE.

Nonce generators represent another departure from nonce-based symmetric encryption. In the latter the adversary picks the nonce, but we saw in Section 1 that this does not work in the public-key setting. Instead, we model the process of a sender picking a nonce via a nonce generator.

In discussions, we refer to NG as weakly unpredictable if it is unpredictable for adversaries making no EXPOSE query, meaning $\text{Adv}_{\text{NG}}^{\text{pred}}(q_1, 0)$ is “small” for “practical” values of q_1 , and strongly unpredictable if it is unpredictable even for adversaries making an EXPOSE query, meaning $\text{Adv}_{\text{NG}}^{\text{pred}}(q_1, 1)$ is “small” for “practical” values of q_1 . If NG is strongly unpredictable it is also weakly unpredictable, but not necessarily vice versa. That is, the class of weakly unpredictable nonce generators is larger than the class of strongly unpredictable nonce generators.

REAL OR RANDOM SECURITY. Let $\text{HE}: \text{HE.Keys} \times \text{HE.Dom} \rightarrow \text{HE.Rng}$ be an oracle family of functions (this means it may have access to a random oracle). The first input is referred to as the “key” or the “seed.” The domain has the form $\text{HE.Dom} = \{0, 1\}^* \times \text{HE.NS}$, so that an input is a pair of strings, the first referred to as the “message” and the second as the “nonce,” the latter drawn from a nonce space HE.NS associated to HE. Consider game $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})$ of Fig. 2 associated to HE, nonce generator NG and an adversary \mathcal{G} . The number of queries to EXPOSE is either 0 or 1, and the number to other oracles is arbitrary. Let

$$\text{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) = 2 \text{Pr}[\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})] - 1 .$$

Note that here the adversary is given the key (seed) xk as input, unlike in the PRF notion, modeling exposure of the sender seed. Security asks that outputs of $\text{HE}(xk, (\cdot, \cdot))$, for adversary-chosen messages and nonces from the nonce generator, are indistinguishable from random. Clearly, this will be possible only with certain restrictions, which will emerge when we discuss our constructions below.

In game $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})$, we say that adversary \mathcal{G} is *agnostic* if its ROR queries do not depend on the seed xk . More formally, there exists a pair $(\mathcal{G}_1, \mathcal{G}_2)$ of algorithms such that $\mathcal{G}^{\text{ror,EXPOSE,RO}}(xk)$ does the following:

$$St \leftarrow_s \mathcal{G}_1^{\text{ror,EXPOSE,RO}}(\varepsilon); c' \leftarrow \mathcal{G}_2^{\text{EXPOSE,RO}}(xk, St); \text{Return } c'.$$

This represents one of the restrictions we will impose to achieve ror security in the standard model.

HEDGED EXTRACTORS. A *hedged extractor* $\text{HE}: \text{HE.Keys} \times \text{HE.Dom} \rightarrow \text{HE.Rng}$ is an oracle family of functions as above where the goal is that (1) HE is a PRF, meaning $\text{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{A})$ is low for any adversary \mathcal{A} of practical resources, and also (2) $\text{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})$ is small for some class of nonce generators NG and some class of ror adversaries \mathcal{G} , both specified via results for individual hedged extractors. We give a ROM construction and standard model one. Both achieve PRF security, but differ in the type of ror security achieved. The ROM construction achieves ror security for unpredictable nonce generators (both weak and strong) and for all ror adversaries. This is “best possible” because the unpredictability assumption is easily seen to be necessary. The standard model construction achieves ror security for strongly unpredictable generators and agnostic ror adversaries. These restrictions reflect limitations of the randomness extractors that are our underlying tool. The restriction to agnostic adversaries reflects that randomness extractors only work on seed-independent distributions, and the strong unpredictability requirement on the generator reflects that when extracting from a sequence of inputs, one needs not only that each has some min-entropy, but that it does even given the others.

ROM HEDGED EXTRACTOR. We start by giving a simple and efficient construction $\mathbf{HE1}$ of a hedged extractor in the ROM. Let ℓ be a desired number of output bits for the extractor, and k a desired seed (key) length. Associated to ℓ, k is the hedged extractor $\text{HE} = \mathbf{HE1}[\ell, k]: \{0, 1\}^k \times (\{0, 1\}^* \times \{0, 1\}^*) \rightarrow \{0, 1\}^\ell$ defined by

$$\text{HE}^{\text{RO}}(xk, x) = \text{RO}((xk, x), \ell).$$

Here $\text{HE.Keys} = \{0, 1\}^k$, $\text{HE.Dom} = \{0, 1\}^* \times \text{HE.NS}$ with $\text{HE.NS} = \{0, 1\}^*$, and $\text{HE.Rng} = \{0, 1\}^\ell$.

The following lemma states that this construction achieves PRF security and also achieves real or random security assuming only that the nonce generator is unpredictable. Note the latter only requires each nonce to individually be unpredictable, but nonces may be arbitrarily correlated, and it could be that given n_1 one can easily predict n_2 . But the extractor works nonetheless.

Lemma 1. *Let $\ell, k \geq 1$ be integers and let $\text{HE} = \mathbf{HE1}[\ell, k]$ be the ROM function family associated to ℓ and k as above.*

1. If \mathcal{A} is an adversary making q_2 queries to its RO oracle, then

$$\mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{A}) \leq \frac{q_2}{2^k}. \quad (1)$$

2. Let NG be a nonce generator. If \mathcal{G} is an adversary making q_1 queries to its RoR oracle, q_2 queries to its RO oracle, and $q_3 \in \{0, 1\}$ queries to its EXPOSE oracle, then

$$\mathbf{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) \leq q_2 \cdot \mathbf{Adv}_{\text{NG}}^{\text{pred}}(q_1, q_3). \quad (2)$$

Note that in the 2nd case, the reduction preserves the number of EXPOSE queries, meaning the number reflected by q_3 is the number made by \mathcal{G} . This is the best one could hope for.

Proof (Lemma 1). For the proof of Equation (1), consider the games G_0, G_1 of Fig. 3, where G_1 contains the boxed code and G_0 does not. Letting c denote the challenge bit in game $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{A})$, the following, justified below, establishes Equation (1):

$$\mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{A}) \mid c = 1] - \left(1 - \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{A}) \mid c = 0]\right) \quad (3)$$

$$= \Pr[G_1] - \Pr[G_0] \quad (4)$$

$$\leq \Pr[G_0 \text{ sets bad}] \quad (5)$$

$$\leq \frac{q_2}{2^k}. \quad (6)$$

Equation (3) is a standard re-formulation of the definition of the advantage. In game G_0 , replies to queries to the FN and RO oracles are independently distributed, so that it is equivalent to the $c = 0$ case of game $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{A})$, up to the flipping of the outcomes from **true** to **false**. In game G_1 , the replies to FN queries are given by HE^{RO} , making it equivalent to the $c = 1$ case of game $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{A})$. This justifies Equation (4). Games G_0, G_1 are identical until **bad** (differ only in statements following the setting of **bad** to **true**), so the Fundamental Lemma of Game Playing [10] justifies Equation (5). In game G_0 , replies to all oracle queries are random and independent of xk so the probability that the latter is queried as part of an RO query is at most the quantity of Equation (6).

For Equation (2), consider the games G_2, G_3 of Fig. 3, where G_3 contains the boxed code and G_2 does not. Let predictor adversary \mathcal{P} be as specified in Fig. 3. Letting c denote the challenge bit in game $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})$, the following, justified below, establishes Equation (2):

$$\mathbf{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) = \Pr[\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) \mid c = 1] - \left(1 - \Pr[\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) \mid c = 0]\right) \quad (7)$$

$$= \Pr[G_3] - \Pr[G_2] \quad (8)$$

$$\leq \Pr[G_2 \text{ sets bad}] \quad (9)$$

$$\leq q_2 \cdot \mathbf{Adv}_{\text{NG}}^{\text{pred}}(\mathcal{G}) \quad (10)$$

$$\leq q_2 \cdot \mathbf{Adv}_{\text{NG}}^{\text{pred}}(q_1, q_3). \quad (11)$$

<p><u>Game $G_0 / \overline{G_1}$</u> $xk \leftarrow_s \{0, 1\}^k$; $c' \leftarrow_s \mathcal{A}^{\text{FN}, \text{RO}}$; Return ($c' = 1$)</p> <p><u>FN($w$)</u> If $S[w] = \perp$ then $S[w] \leftarrow_s \{0, 1\}^\ell$; Return $S[w]$</p> <p><u>RO(w, l)</u> If $T[w, l] = \perp$ then $T[w, l] \leftarrow_s \{0, 1\}^l$; $(u, x) \leftarrow w$ If $((u = xk) \text{ and } (l = \ell))$ then If $(S[w] = \perp)$ then $S[w] \leftarrow_s \{0, 1\}^\ell$; bad \leftarrow true; $\boxed{T[w, l] \leftarrow S[w]}$ Return $T[w, l]$</p>	<p><u>Game $G_2 / \overline{G_3}$</u> $St \leftarrow \varepsilon$; $xk \leftarrow_s \{0, 1\}^k$ $S \leftarrow \emptyset$ $c' \leftarrow_s \mathcal{G}^{\text{ROR}, \text{EXPOSE}, \text{RO}}(xk)$; Return ($c' = 1$)</p> <p><u>ROR($m, \eta$)</u> If $s = 1$ then return \perp $(n, St) \leftarrow_s \text{NG}(\eta, St)$ $r \leftarrow R[m, n]$; $R[m, n] \leftarrow_s \{0, 1\}^\ell$ If $((m, n) \in S)$ then bad \leftarrow true; $\boxed{R[m, n] \leftarrow r}$ If $(T[(xk, (m, n)), \ell] \neq \perp)$ then bad \leftarrow true; $\boxed{R[m, n] \leftarrow T[(xk, (m, n)), \ell]}$ $S \leftarrow S \cup \{(m, n)\}$ Return $R[m, n]$</p> <p><u>EXPOSE</u> $s \leftarrow 1$ Return St</p> <p><u>RO(w, l)</u> If $(T[w, l] = \perp)$ then $T[w, l] \leftarrow_s \{0, 1\}^l$; $(u, x) \leftarrow w$; $(m, n) \leftarrow x$ If $((u = xk) \text{ and } (l = \ell) \text{ and } ((m, n) \in S))$ then bad \leftarrow true; $\boxed{T[w, l] \leftarrow R[m, n]}$ Return $T[w, l]$</p>
	<p><u>Adversary \mathcal{P}^{GEN}</u> $xk \leftarrow_s \{0, 1\}^k$; $i \leftarrow 0$ $g \leftarrow_s \{1, \dots, q_2\}$ $c' \leftarrow_s \mathcal{G}^{\text{ROR}, \text{EXPOSE}, \text{RO}}(xk)$ Return n_g</p> <p><u>ROR(m, η)</u> If $s = 1$ then return \perp GEN(η) $r \leftarrow_s \{0, 1\}^\ell$ Return r</p> <p><u>EXPOSE</u> $s \leftarrow 1$ $St \leftarrow \text{EXPOSE}$ Return St</p> <p><u>RO(w, l)</u> If $(T[w, l] = \perp)$ then $T[w, l] \leftarrow_s \{0, 1\}^l$ $(u, x) \leftarrow w$; $(m, n) \leftarrow x$ $i \leftarrow i + 1$; $n_i \leftarrow n$ Return $T[w, l]$</p>

Fig. 3. Games and adversary for proof of Lemma 1.

Equation (7) is a standard re-formulation of the definition of the advantage. In game G_2 , replies to queries to the ROR and RO oracles are independently distributed, so that it is equivalent to the $c = 0$ case of game $\mathbf{G}_{\text{HE}, \text{NG}}^{\text{TOR}}(\mathcal{G})$, up

to the flipping of the outcomes from true to false. In game G_3 , the replies to RO-R queries are given by HE^{RO} , making it equivalent to the $c = 1$ case of game $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G})$. This justifies Equation (8). Games G_2, G_3 are identical until bad, so the Fundamental Lemma of Game Playing [10] justifies Equation (9). In game G_2 , replies to all oracle queries are random and independent, so adversary \mathcal{P} can simulate the oracles of adversary \mathcal{G} directly. Its output is a random one of the nonces in a RO-query of \mathcal{G} , whence Equation (10). Equation (11) is because \mathcal{P} makes q_1 calls to its GEN oracle and q_3 queries to its EXPOSE oracle. \square

STANDARD-MODEL HEDGED EXTRACTOR. Next we give a standard-model hedged extractor based on a PRF F and an almost-XOR-universal hash function H . We use the latter essentially as a strong extractor. The construction is simple: the PRF is evaluated on the message and nonce, and the hash function is evaluated only on the nonce. The results are combined via a simple XOR operation. The intuition behind this is that as long as at least one of the outputs generated by the two schemes is random, then the result is also random. PRF security of the hedged extractor is proved assuming only on the assumed PRF security of F . Real-or-random security of the hedged extractor is shown for a restricted class of nonce generators NG and adversaries \mathcal{G} . Namely NG must retain unpredictability even in the presence of an EXPOSE query revealing the state, and \mathcal{G} 's RO-R queries must not depend on the seed. These restrictions reflect inherent limitations of strong extractors.

We start with some definitions. For $\epsilon \in [0, 1]$, function family H is ϵ -almost XOR-universal [22] if $H.\text{Rng} = \{0, 1\}^\ell$ for some $\ell \in \mathbb{N}$ and, for all distinct $x, y \in H.\text{Dom}$ and all $s \in \{0, 1\}^\ell$, we have

$$\Pr[H(hk, x) \oplus H(hk, y) = s : hk \leftarrow_s H.\text{Keys}] \leq \epsilon.$$

Our standard model construction is as follows. Let ℓ be a desired number of output bits for the extractor. Let $F: F.\text{Keys} \times (\{0, 1\}^* \times \{0, 1\}^*) \rightarrow \{0, 1\}^\ell$ be a function family assumed to be a PRF, and let $H: H.\text{Keys} \times H.\text{Dom} \rightarrow \{0, 1\}^\ell$ be an almost-XOR-universal hash function with $H.\text{Dom} \subseteq \{0, 1\}^*$. We associate to ℓ, F, H the standard-model hedged extractor $\text{HE} = \mathbf{HE2}[F, H]: (F.\text{Keys} \times H.\text{Keys}) \times (\{0, 1\}^* \times H.\text{Dom}) \rightarrow \{0, 1\}^\ell$ defined by

Algorithm $\text{HE}(xk, (x, y))$

$(hk, fk) \leftarrow xk ; z_1 \leftarrow H(hk, y) ; z_2 \leftarrow F(fk, (x, y)) ; \text{Return } z_1 \oplus z_2$

Here $\text{HE}.\text{Keys} = F.\text{Keys} \times H.\text{Keys}$, $\text{HE}.\text{Dom} = \{0, 1\}^* \times \text{HE}.\text{NS}$ with $\text{HE}.\text{NS} = H.\text{Dom}$, and $\text{HE}.\text{Rng} = \{0, 1\}^\ell$. The following says this hedged extractor achieves PRF security and restricted real-or-random security.

Lemma 2. *Let $\ell \geq 1$ be an integer. Let $F: F.\text{Keys} \times (\{0, 1\}^* \times \{0, 1\}^*) \rightarrow \{0, 1\}^\ell$ be a function family. Let $H: H.\text{Keys} \times H.\text{Dom} \rightarrow \{0, 1\}^\ell$ be a $(1 + \gamma) \cdot 2^{-\ell}$ -almost-XOR-universal hash function. Let $\text{HE} = \mathbf{HE2}[F, H]$ be the function family associated to ℓ, F and H as above.*

1. *If \mathcal{A} is an adversary making q queries to its FN oracle then there is an adversary \mathcal{B} (described in the proof) such that*

$$\text{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_{\text{F}}^{\text{prf}}(\mathcal{B}). \quad (12)$$

Adversary \mathcal{B} also makes q queries to its FN oracle and has running time that of \mathcal{A} plus the time for q computations of H .

2. Let NG be a nonce generator that produces outputs in the set $H.\text{Dom}$. If \mathcal{G} is an agnostic adversary making q queries to its ROR oracle and $\text{Adv}_{\text{NG}}^{\text{pred}}(q, 1) \leq 2^{-m}$, then

$$\text{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) \leq \frac{q}{2} \sqrt{\gamma + \frac{2^\ell}{2^m}}. \quad (13)$$

To prove this we first need some more definitions. Recall that for $\epsilon \in [0, 1]$, a function family H is ϵ -almost universal if $\Pr[H(hk, x) = H(hk, y) : hk \leftarrow_{\$} H.\text{Keys}] \leq \epsilon$ for all distinct $x, y \in H.\text{Dom}$. For a function family H with $H.\text{Dom} = \{0, 1\}^* \times \{0, 1\}^*$, we say that H is ϵ -almost universal in the second input component if for all $x_1, y_1 \in \{0, 1\}^*$ and all $x_2, y_2 \in \{0, 1\}^*$ with $x_2 \neq y_2$, $\Pr[H(hk, (x_1, x_2)) = H(hk, (y_1, y_2)) : hk \leftarrow_{\$} H.\text{Keys}] \leq \epsilon$. For $k \in \mathbb{N}$ and $\epsilon \in [0, 1]$, a (k, ϵ) -strong extractor SE is a (standard model) family of functions such that $\text{Adv}_{\text{E,NG}}^{\text{ror}}(1, 0, 1) \leq \epsilon$ for all NG with $\text{Adv}_{\text{NG}}^{\text{pred}}(1, 1) \leq 2^{-k}$. This is a re-formulation of the standard requirement that is clearly equivalent to it. The Leftover Hash Lemma, a celebrated result by Impagliazzo, Levin, and Luby [20], states that an almost universal hash function is a strong extractor:

Lemma 3 (Leftover Hash Lemma). *Let $\gamma, k > 0$. Let H be an $(1+\gamma)/|H.\text{Rng}|$ -almost universal hash function family. Then H is an (k, ϵ) -strong extractor where*

$$\epsilon = \frac{1}{2} \sqrt{\gamma + \frac{|H.\text{Rng}|}{2^k}}.$$

The original result is stated in a different formalism and a slightly more restricted form, but it generalizes to the form stated here. (The formulation of the bounds comes from [2].)

Since we use the nonce that is used in the hash function H also outside of it, namely in the evaluation of the PRF F , we cannot immediately apply the leftover-hash lemma. However, we show that if the function H is almost XOR-universal, then the hedged extractor HE obtained by XORing the output of F to the output of H is almost universal and therefore serves as a strong extractor.

Lemma 4. *Let $H : H.\text{Keys} \times H.\text{Dom} \rightarrow \{0, 1\}^\ell$ be an ϵ -almost-XOR-universal hash function. Let $F : F.\text{Keys} \times (\{0, 1\}^* \times \{0, 1\}^*) \rightarrow \{0, 1\}^\ell$ a function family and $H.\text{Dom} \subseteq \{0, 1\}^*$. Then HE as defined above is ϵ -almost-universal in the second input component.*

Proof. For all $m, m' \in \{0, 1\}^*$ and $n, n' \in H.\text{Dom}$ with $n \neq n'$:

$$\begin{aligned} & \Pr [\text{HE}(xk, (m, n)) = \text{HE}(xk, (m', n')) : xk \leftarrow_{\$} \text{HE.Keys}] \\ &= \Pr \left[H(hk, n) \oplus F(fk, (m, n)) = H(hk, n') \oplus F(fk, (m', n')) \right. \\ & \quad \left. : (hk, fk) \leftarrow_{\$} H.\text{Keys} \times F.\text{Keys} \right] \end{aligned}$$

<p> <u>Game $G(\mathcal{G})$</u> $St \leftarrow \varepsilon ; xk \leftarrow_{\\$} \mathbf{E.Keys} ; c \leftarrow_{\\$} \{0, 1\} ; s \leftarrow 0$ $St_{\mathcal{G}} \leftarrow_{\\$} \mathcal{G}_{\text{prep}}^{\text{GEN}}$ $c' \leftarrow_{\\$} \mathcal{G}_{\text{chal}}^{\text{ROR, EXPOSE}}(St_{\mathcal{G}}, xk)$ Return $(c = c')$ </p> <p> <u>ROR(m, η)</u> If $s = 1$ then return \perp $(n, St) \leftarrow_{\\$} \mathbf{NG}(\eta, St)$ If $(c = 1)$ then $r \leftarrow \mathbf{E}^{\text{RO}}(xk, m, n)$ Else $r \leftarrow_{\\$} \mathbf{E.Rng}$ Return r </p> <p> <u>GEN(η)</u> $(n, St) \leftarrow_{\\$} \mathbf{NG}(\eta, St)$ </p> <p> <u>EXPOSE</u> $s \leftarrow 1$ Return St </p>

Fig. 4. Hybrid game for E, as needed in the proof of Lemma 2.

$$\leq \max_{fk \in \mathbf{F.Keys}} \Pr \left[\mathbf{H}(hk, n) \oplus \mathbf{H}(hk, n') = \mathbf{F}(fk, (m, n)) \oplus \mathbf{F}(fk, (m', n')) \right. \\ \left. : hk \leftarrow_{\$} \mathbf{H.Keys} \right] \leq \epsilon,$$

since for each $fk \in \mathbf{F.Keys}$ the term $\mathbf{F}(fk, (m, n)) \oplus \mathbf{F}(fk, (m', n'))$ describes a fixed value $s \in \{0, 1\}^\ell$ and \mathbf{H} is ϵ -almost XOR-universal. \square

A strong extractor will only guarantee the outputs to be random and independent of the seed if the inputs to the extractor, that is the nonces, do not depend on the seed. Once the seed is exposed to the adversary, no guarantee on further outputs can be given. Therefore, for the game $\mathbf{G}_{\text{HE,NG}}^{\text{TOR}}(\mathcal{G})$ we restrict our attention to agnostic adversaries \mathcal{G} .

Proof (Lemma 2). We first prove Equation (12). Adversary \mathcal{B} starts by choosing a seed $hk \leftarrow_{\$} \mathbf{H.Keys}$ uniformly at random. It then uses the assumed adversary \mathcal{A} , and whenever \mathcal{A} makes a query $\text{FN}(x, y)$, then \mathcal{B} first computes $z_1 \leftarrow \mathbf{H}(hk, y)$ and then makes a query $z_2 \leftarrow \text{FN}(x, y)$, and returns $z_1 \oplus z_2$. Finally, \mathcal{B} provides the same output as \mathcal{A} . The view of \mathcal{A} has the same distribution in $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{A})$ and in $\mathbf{G}_{\text{F}}^{\text{prf}}(\mathcal{B})$, and therefore the distribution of the output is also the same. This implies Equation (12).

We prove Equation (13) using a hybrid argument similarly to Zuckerman [33, Lemma 6]. The hybrid argument involves a game G as specified in Fig. 4 and adversaries $\mathcal{G}_1, \dots, \mathcal{G}_q$ that each query the oracle ROR only once. This is achieved by having \mathcal{G}_i answer the i th query of \mathcal{G} by using the ROR oracle, by using for all

previous queries uniformly random values and computing all subsequent values by evaluating HE. In more detail, for all queries $j = 1, \dots, i-1$, the adversary \mathcal{G}_i will call its $\text{GEN}(\eta)$ oracle and then sample a value from $\text{HE.Rng} = \{0, 1\}^\ell$ uniformly at random. For the i th query, the adversary \mathcal{G}_i will call $\text{ROR}(m, \eta)$ and obtain the output r , as well as the nonce generator's state St through a subsequent EXPOSE query. The answer to \mathcal{G} will be r in this case. For all queries $j = i+1, \dots, q$, adversary \mathcal{G}_i will compute the nonce via $(n_j, St) \leftarrow \text{NG}(\eta, St)$ and the output as $\text{HE}(xk, m, n_j)$. This is possible because it uses the seed xk only after the challenge query. We can therefore specify the adversary as a pair $\mathcal{G}_i = (\mathcal{G}_{i,\text{prep}}, \mathcal{G}_{i,\text{chal}})$ as described in Fig. 4.

To conclude the hybrid argument, we first observe that the view of \mathcal{G} in $\mathbf{G}_{\text{HE,NG}}^{\text{TOR}}(\mathcal{G})$ with $c = 1$ is the same as its view in $\mathbf{G}(\mathcal{G}_1)$ with $c = 1$. Subsequently, for each $i = 1, \dots, q-1$, the view of \mathcal{G} in $\mathbf{G}(\mathcal{G}_i)$ with $c = 0$ is the same as in $\mathbf{G}(\mathcal{G}_{i+1})$ with $c = 1$. Finally, the view of \mathcal{G} in $\mathbf{G}(\mathcal{G}_q)$ with $c = 0$ is the same as in $\mathbf{G}_{\text{HE,NG}}^{\text{TOR}}(\mathcal{G})$ with $c = 0$.

We still have to bound the advantage of adversary \mathcal{G}_i . By Lemma 4, we know that HE is $(1+\gamma)/2^\ell$ -almost universal in its second input and therefore by a slight generalization of Lemma 3 an (m, ϵ) -strong extractor for $\epsilon = 1/2\sqrt{\gamma + 2^{\ell-m}}$. Fix i , then $\mathcal{G}_{i,\text{prep}}$ makes $i-1$ queries to its GEN oracle. We can therefore define a nonce generator $\text{RG}(\mathcal{G}_{i,\text{guess}}, \text{NG})$ as follows: Set $St \leftarrow \varepsilon$ and execute $(n, St_{\mathcal{G}}) \leftarrow_s \mathcal{G}_{i,\text{guess}}^{\text{GEN}}$ with the oracle as defined in the game. Compute a random x via $(x, St) \leftarrow_s \text{NG}(\eta, St)$ and output a pair $(x, (St, St_{\mathcal{G}}))$.

This implies that $\text{Adv}_{\text{RG}(\mathcal{G}_{i,\text{prep}}, \text{NG})}^{\text{pred}}(1, 1) \leq \text{Adv}_{\text{NG}}^{\text{pred}}(i, 1)$, since otherwise $\mathcal{G}_{i,\text{prep}}$ could be used in the statement against the assumption on NG. But now we can view $\mathcal{G}_{i,\text{chal}}$ as an adversary against HE in the original \mathbf{G}^{TOR} game, and $2 \Pr[\text{HYB}_{\text{HE,NG}}(\mathcal{G}_i)] - 1 \leq \text{Adv}_{\text{SE, RG}(\mathcal{G}_{i,\text{prep}}, \text{NG})}^{\text{TOR}}(\mathcal{G}_{i,\text{chal}}) \leq \epsilon$. We obtain a factor q through the hybrid argument; this completes the proof. \square

4 Nonce-based public-key encryption

In this section we define nonce-based public-key encryption, giving first a syntax and then two security goals, NBP1 and NPB2. We give a construction, simultaneously meeting both goals, based on a hedged extractor. Instantiating the latter via our constructions of Section 3 yields concrete nonce-based PKE schemes, one in the ROM and the other in the standard model.

SYNTAX. A *nonce-based public-key encryption scheme* NPE specifies the following. Receiver key-generation algorithm NPE.Kg returns an encryption key ek and associated decryption key dk . Sender seed-generation algorithm NPE.sKg returns a seed xk . Encryption algorithm NPE.Enc takes ek, xk , message $m \in \{0, 1\}^*$ and nonce n from nonce set NPE.NS to return a ciphertext c . This algorithm is *deterministic*. Decryption algorithm NPE.Dec (also deterministic) takes ek, dk and ciphertext c to return a value in $\{0, 1\}^* \cup \{\perp\}$. The scheme NPE is correct if for all $(ek, dk) \in [\text{NPE.Kg}]$, all $xk \in [\text{NPE.sKg}]$, all $m \in \{0, 1\}^*$ and all $n \in \text{NPE.NS}$ we have $\text{NPE.Dec}(ek, dk, \text{NPE.Enc}(ek, xk, m, n)) = m$.

<p>Game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})$</p> <p>$(ek, dk) \leftarrow \text{NPE.Kg} ; xk \leftarrow \text{NPE.sKg}$ $b \leftarrow \{0, 1\} ; St \leftarrow \varepsilon ; C, S_0, S_1 \leftarrow \emptyset$ $b' \leftarrow \mathcal{A}^{\text{ENC,DEC,RO}}(ek)$ Return $(b = b')$</p> <p><u>ENC</u>(m_0, m_1, η)</p> <p>If $(m_0 \neq m_1)$ then return \perp $(n, St) \leftarrow \text{NG}(\eta, St)$ If $((m_0, n) \in S_0)$ then return \perp If $((m_1, n) \in S_1)$ then return \perp $S_0 \leftarrow S_0 \cup \{(m_0, n)\}$ $S_1 \leftarrow S_1 \cup \{(m_1, n)\}$ $c \leftarrow \text{NPE.Enc}^{\text{RO}}(ek, xk, m_b, n)$ $C \leftarrow C \cup \{c\}$ Return c</p> <p><u>DEC</u>(c)</p> <p>If $(c \in C)$ then $m \leftarrow \perp$ Else $m \leftarrow \text{NPE.Dec}^{\text{RO}}(ek, dk, c)$ Return m</p> <p><u>RO</u>(x, l)</p> <p>If $T[x, l] = \perp$ then $T[x, l] \leftarrow \{0, 1\}^l$ Return $T[x, l]$</p>	<p>Game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$</p> <p>$(ek, dk) \leftarrow \text{NPE.Kg}$ $xk \leftarrow \text{NPE.sKg}$ $b \leftarrow \{0, 1\} ; St \leftarrow \varepsilon ; C \leftarrow \emptyset$ $b' \leftarrow \mathcal{A}^{\text{ENC,DEC,RO}}(ek, xk)$ Return $(b = b')$</p> <p><u>ENC</u>(m_0, m_1, η)</p> <p>If $(m_0 \neq m_1)$ then return \perp $(n, St) \leftarrow \text{NG}(\eta, St)$ $c \leftarrow \text{NPE.Enc}^{\text{RO}}(ek, xk, m_b, n)$ $C \leftarrow C \cup \{c\}$ Return c</p> <p><u>DEC</u>(c)</p> <p>If $(c \in C)$ then return \perp $m \leftarrow \text{NPE.Dec}^{\text{RO}}(ek, dk, c)$ Return m</p> <p><u>RO</u>(x, l)</p> <p>If $T[x, l] = \perp$ then $T[x, l] \leftarrow \{0, 1\}^l$ Return $T[x, l]$</p>
---	--

Fig. 5. Games for security goals for nonce-based asymmetric encryption scheme NPE. Both are relative to nonce generator NG.

The receiver key-generation algorithm has the same role as in a randomized PKE scheme. The sender seed-generation algorithm is a new element of nonce-based PKE schemes. Encryption is changed to take a nonce rather than randomness and, importantly, is now deterministic. Decryption is as in a standard PKE scheme. Unlike symmetric nonce-based encryption, the decryption algorithm is not given the nonce.

Nonce-based encryption is a sender-side hardening and can be added to an existing encryption scheme in such a way that the receiver is oblivious to its presence and the receiver implementation needs no changes.

SECURITY DEFINITIONS. Let NPE be a nonce-based PKE scheme. Let NG be a nonce generator returning nonces in NPE.NS. We associate to NPE, NG and adversary \mathcal{A} the games of Fig. 5. We define the advantages of \mathcal{A} in these games

via

$$\mathbf{Adv}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})] - 1$$

and

$$\mathbf{Adv}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})] - 1,$$

respectively. The games are described in the ROM; standard-model definitions are derived by considering only schemes and adversaries that do not query the RO oracle. The notation “nbp” stands for “nonce-based privacy.” We proceed to discuss the definitions.

Game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})$ formalizes security in the case where the sender’s seed is not exposed, captured in the formalism by the fact that the adversary is not given the seed as input, while game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$ formalizes security in the case where the sender’s seed is exposed, formalized by its being given to the adversary as input. Both ask for indistinguishability-style security under a chosen-ciphertext attack. In $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})$, the natural restriction that one would consider is to ask that nonces not repeat. The restriction we make is weaker, resulting in a stronger security condition, namely that message-nonce pairs may not repeat. (Thus, security is provided even if a nonce repeats, as long as the message is different.) We will achieve this notion for *any* nonce generator, meaning we get very good privacy with minimal restrictions on nonces. In $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$, no restriction is made, so a priori the notion is stronger, but we will achieve it only if the hedged extractor is ror-secure, which will be further reduced to unpredictability of the nonce generator. Thus, in this case, security requires unpredictable nonces.

In game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$, we say that adversary \mathcal{A} is *agnostic* if its ENC, DEC queries do not depend on the seed xk . More formally, there exists a pair $(\mathcal{A}_1, \mathcal{A}_2)$ of algorithms such that $\mathcal{A}^{\text{ENC,DEC,RO}}(ek, xk)$ does the following:

$$St \leftarrow_s \mathcal{A}_1^{\text{ENC,DEC,RO}}(ek); c' \leftarrow \mathcal{A}_2^{\text{RO}}(xk, St); \text{Return } c'.$$

SCHEME. We specify a scheme NPE that achieves both the NPB1 and NPB2 security notions simultaneously, that is, it guarantees security if either the sender’s state remains secret and as long as the message-nonce pairs are unique, or even if the sender’s state is leaked to the adversary as long as the nonces are sufficiently unpredictable. The construction is actually a transform **R2NPE** that takes a base, randomized PKE scheme PE and a hedged extractor HE to return a nonce-based PKE scheme $\text{NPE} = \mathbf{R2NPE}[\text{PE}, \text{HE}]$ whose algorithms are described in Fig. 6. The nonce space is that of the hedged extractor, i.e. $\text{NPE.NS} = \text{HE.NS}$. The construction requires that the randomness provided by HE is sufficient, i.e., $\text{PE.rl} = \text{HE.ol}$.

We first prove that the scheme achieves NBP1-security, that is, it is secure as a nonce-based scheme as long as the sender’s seed remains secret. The theorem bounds the adversaries advantage by advantages of other, related adversaries against the underlying probabilistic public-key scheme and the PRF-property of the hedged extractor HE. For the constructions described in Section 3, the latter advantage is then bounded by Lemmas 1 (for the ROM construction) and 2 (for the standard-model construction), respectively.

Algorithm NPE.Kg $(ek, dk) \leftarrow_s \text{PE.Kg}$ Return (ek, dk)	Algorithm NPE.Enc ^{RO} (ek, xk, m, n) $r \leftarrow \text{HE}^{\text{RO}}(xk, (m, n))$ $c \leftarrow \text{PE.Enc}(ek, m; r)$ Return c
Algorithm NPE.sKg $xk \leftarrow_s \text{HE.Keys}$ Return xk	Algorithm NPE.Dec (ek, dk, c) $m \leftarrow \text{PE.Dec}(ek, dk, c)$ Return m

Fig. 6. The nonce-based public-key encryption scheme NPE, based on probabilistic public-key encryption scheme PE and hedged extractor HE. The nonce space is the same as for the hedged extractor.

Theorem 5. Let PE be a (standard, randomized) public-key encryption scheme. Let HE be a hedged extractor. Let nonce-based public-key encryption scheme $\text{NPE} = \mathbf{R2NPE}[\text{PE}, \text{HE}]$ be associated to them as above. Let NG be a nonce generator. Let \mathcal{A} be an adversary making at most q_1 queries to its ENC oracle, q_2 queries to its DEC oracle, and q_3 queries to its RO oracle. Then the proof specifies adversaries \mathcal{B} and \mathcal{G} such that

$$\mathbf{Adv}_{\text{NPE,NG}}^{\text{nbpl}}(\mathcal{A}) \leq 2\mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{G}) + \mathbf{Adv}_{\text{PE}}^{\text{ind}}(\mathcal{B}), \quad (14)$$

where adversary \mathcal{B} makes at most q_1 queries to its ENC oracle and q_2 queries to its DEC oracle; in terms of computation it evaluates NG for q_1 times and manages an array of generated ciphertexts. Adversary \mathcal{G} makes at most q_1 queries to its FN oracle and q_3 queries to its RO oracle; it generates keys for PE and compute q_1 encryptions and q_2 decryptions in addition to the computation of \mathcal{A} .

Proof. Adversary \mathcal{B} obtains ek and starts by setting $St \leftarrow \varepsilon$, $S_0 \leftarrow \emptyset$, and $S_1 \leftarrow \emptyset$. Adversary \mathcal{B} then runs the adversary \mathcal{A} internally on input ek . For queries $\text{ENC}(m_0, m_1, \eta)$, adversary \mathcal{B} computes $(n, St) \leftarrow_s \text{NG}(\eta, St)$, checks whether $(m_0, n) \in S_0$ or $(m_1, n) \in S_1$, and returns \perp in that case. Otherwise, it queries its own oracle $\text{ENC}(m_0, m_1)$ and returns the result to \mathcal{A} . Queries $\text{DEC}(c)$ are answered using the respective oracles in the game, and potential queries $\text{RO}(x, l)$ are answered by emulating a random oracle via lazy sampling as also described in the game.

For each $d \in \{0, 1\}$, adversary \mathcal{G}_d against the PRF initializes $St \leftarrow \varepsilon$, $S_0 \leftarrow \emptyset$, and $S_1 \leftarrow \emptyset$ as \mathcal{B} and generates a key pair $(ek, dk) \leftarrow_s \text{PE.Kg}$. Adversary \mathcal{G}_d then runs adversary \mathcal{A} on input ek . Upon a query $\text{ENC}(m_0, m_1, \eta)$ from \mathcal{A} , adversary \mathcal{G}_d computes $(n, St) \leftarrow_s \text{NG}(\eta, St)$, checks whether $(m_0, n) \in S_0$ or $(m_1, nonce) \in S_1$, and returns \perp in that case. \mathcal{G}_d gets r from its oracle $\text{FN}(m_d, n)$, computes $c \leftarrow \text{PE.Enc}(ek, m_d; r)$ and returns c . Queries $\text{DEC}(c)$ are answered by computing $m \leftarrow \text{PE.Dec}(ek, dk, c)$ and returning m . Potential queries $\text{RO}(x, l)$ are referred to the random oracle provided in the game. When \mathcal{A} provides its output b' , then \mathcal{G}_d outputs $b' \oplus d$.

We define a hybrid game G in which all computations are performed as before, except that in the encryption with a new pair of message m_b and nonce n a fresh random string $r \leftarrow \{0, 1\}^\ell$ is used. The difference with the game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})$ can be highlighted as described in game G in Fig. 7

Game G	Game H
\vdots \vdots $S_1 \leftarrow S_1 \cup \{(m_1, n)\}$ $c \leftarrow \text{NPE.Enc}^{\text{RO}}(ek, xk, m_b, n)$	\vdots \vdots $(n, St) \leftarrow_s \text{NG}(\eta, St)$ $c \leftarrow \text{NPE.Enc}^{\text{RO}}(ek, xk, m_b, n)$
<div style="border: 1px solid black; padding: 2px; display: inline-block;"> $r \leftarrow_s \{0, 1\}^\ell$ $c \leftarrow \text{PE.Enc}(ek, m_b; r)$ </div>	<div style="border: 1px solid black; padding: 2px; display: inline-block;"> $r \leftarrow_s \{0, 1\}^\ell$ $c \leftarrow \text{PE.Enc}(ek, m_b; r)$ </div>
$C \leftarrow C \cup \{c\}$ Return c	$C \leftarrow C \cup \{c\}$ Return c
\vdots \vdots	\vdots \vdots

Fig. 7. Intermediate games used in the proofs of Theorem 5 (left) and Theorem 6 (right).

The view of \mathcal{A} in G is exactly the same as in $\mathbf{G}_{\text{PE}}^{\text{ind}}(\mathcal{B})$. Furthermore, we observe that

$$2 \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_0)] = \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})|b = 0] - \Pr[G|b = 0] + 1,$$

the reason is that if the bit $b = 0$ is chosen in $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_0)$, then the view of \mathcal{A} is exactly as in G with $b = 0$; all ciphertexts are encryptions of m_0 with fresh randomness. Analogously, if $b = 1$ is chosen in $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_0)$, then the view of \mathcal{A} is exactly as in $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})$ with $b = 0$; all ciphertexts are encryptions of m_0 with randomness computed via NG and HE from the message m_0 and the input η , but in this case \mathcal{G}_0 outputs the “wrong” bit. In the same sense,

$$2 \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_1)] = \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})|b = 1] - \Pr[G|b = 1] + 1,$$

since if the bit b is chosen as $b = 0$ in $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_0)$, then the view of \mathcal{A} is exactly as in G with $b = 1$; all ciphertexts are encryptions of m_1 with fresh randomness. This is the reason for \mathcal{G}_1 to invert the output of \mathcal{A} . Overall, we obtain

$$\begin{aligned} \mathbf{Adv}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A}) &= \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})|b = 0] + \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})|b = 1] - 1 \\ &= \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})|b = 0] - \Pr[G|b = 0] \\ &\quad + \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})|b = 1] - \Pr[G|b = 1] + 2 \Pr[G] - 1 \\ &\leq 2 \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_0)] - 1 + 2 \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G}_1)] - 1 + \mathbf{Adv}_{\text{PE}}^{\text{nbp1}}(\mathcal{B}) \end{aligned}$$

$$= \mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{G}_0) + \mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{G}_1) + \mathbf{Adv}_{\text{PE}}^{\text{nbp1}}(\mathcal{B}).$$

The proof concludes by defining \mathcal{G} as choosing \mathcal{G}_0 or \mathcal{G}_2 with probability $1/2$ each. \square

The property of a hedged extractor to serve simultaneously as extractor and as a PRF implies that the scheme described above is also secure even if the sender's seed leaks, as long as the nonces are sufficiently unpredictable. The reduction in the theorem below preserves agnosticity, meaning if \mathcal{A} is agnostic, so is \mathcal{G} . This allows us to draw conclusions based on Lemma 2 in the case of the standard-model hedged extractor.

Theorem 6. *Let PE be a (standard, randomized) public-key encryption scheme. Let HE be a hedged extractor. Let nonce-based public-key encryption scheme $\text{NPE} = \mathbf{R2NPE}[\text{PE}, \text{HE}]$ be associated to them as above. Let NG be a nonce generator. Let \mathcal{A} be an adversary making at most q_1 queries to its ENC oracle, q_2 queries to its DEC oracle, and q_3 queries to its RO oracle. Then the proof specifies adversaries \mathcal{B} and \mathcal{G} such that*

$$\mathbf{Adv}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A}) \leq 2\mathbf{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) + \mathbf{Adv}_{\text{PE}}^{\text{ind}}(\mathcal{B}), \quad (15)$$

where adversary \mathcal{B} makes at most q_1 queries to its ENC oracle and q_2 queries to its DEC oracle, and emulates a random oracle for q_3 queries. Adversary \mathcal{G} makes at most q_1 queries to its ROR oracle and q_3 queries to its RO oracle; it generates keys for PE and computes q_1 encryptions and q_2 decryptions in addition to the computation of \mathcal{A} . Furthermore, if \mathcal{A} is agnostic, then \mathcal{G} is also agnostic.

Proof. The proof follows the same ideas as the one for Theorem 5. Adversary \mathcal{B} against the underlying public-key encryption scheme PE behaves as follows. It obtains an encryption key ek as an input and generates a seed $xk \leftarrow_{\$} \text{HE.Keys}$, and then runs adversary \mathcal{A} on input (ek, xk) . Upon a query $\text{ENC}(m_0, m_1, \eta)$ by \mathcal{A} , adversary \mathcal{B} queries its oracle $\text{ENC}(m_0, m_1)$ and returns the result to \mathcal{A} . Queries $\text{DEC}(c)$ are answered by making the same query to its own oracle; the queries $\text{RO}(x, l)$ are answered by emulating a random oracle via lazy sampling. Adversary \mathcal{B} outputs the same output bit as \mathcal{A} .

For each $d \in \{0, 1\}$, adversary \mathcal{G}_d against the extractor behaves as follows. It obtains an extractor seed $xk \in \text{HE.Keys}$, generates a key pair $(ek, dk) \leftarrow_{\$} \text{PE.Kg}$, and then runs adversary \mathcal{A} on input (ek, xk) . Upon a query $\text{ENC}(m_0, m_1, \eta)$ from \mathcal{A} , adversary \mathcal{G}_d calls $\text{ROR}(m_d, \eta)$ to obtain a random string r . It then computes $c \leftarrow \text{PE.Enc}(ek, m_d; r)$ and returns c . Queries $\text{DEC}(c)$ are answered by computing $m \leftarrow \text{PE.Dec}(ek, dk, c)$ and returning m . Potential queries $\text{RO}(x, l)$ are referred to the random oracle provided in the game. When \mathcal{A} provides its output b' , then \mathcal{G}_d outputs $b' \oplus d$. Note that \mathcal{G}_d uses the seed only through \mathcal{A} , and if \mathcal{A} uses it only after all encryption queries, then \mathcal{G}_d uses it only after all ROR queries. Thus, if \mathcal{A} is agnostic, then \mathcal{G}_d is also agnostic.

We define a hybrid game H in which all computations are performed as before, except that in the encryption a fresh random string $r \leftarrow \{0, 1\}^\ell$ is used;

the difference with the game $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$ is highlighted by the boxed code in Fig. 7.

The view of \mathcal{A} in \mathbf{H} is exactly the same as in $\mathbf{G}_{\text{PE}}^{\text{nbp2}}(\mathcal{B})$. Furthermore, we observe that

$$2 \Pr[\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}_0)] = \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})|b = 0] - \Pr[\mathbf{H}|b = 0] + 1,$$

the reason is that if the bit $b = 0$ is chosen in $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}_0)$, then the view of \mathcal{A} is exactly as in \mathbf{H} with $b = 0$; all ciphertexts are encryptions of m_0 with fresh randomness. Analogously, if $b = 1$ is chosen in $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}_0)$, then the view of \mathcal{A} is exactly as in $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$ with $b = 0$; all ciphertexts are encryptions of m_0 with randomness computed via HE and NG from the message m_0 and the input η , but in this case \mathcal{G}_0 outputs the “wrong” bit. In the same sense,

$$2 \Pr[\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}_1)] = \Pr[\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})|b = 1] - \Pr[\mathbf{H}|b = 1] + 1,$$

since if the bit b is chosen as $b = 0$ in $\mathbf{G}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}_0)$, then the view of \mathcal{A} is exactly as in \mathbf{H} with $b = 1$; all ciphertexts are encryptions of m_1 with fresh randomness. This is the reason for \mathcal{G}_1 to invert the output of \mathcal{A} . The final computation follows exactly as in Theorem 5. \square

SETTINGS WITH MULTIPLE SENDERS AND MULTIPLE RECEIVERS. The security properties defined above take into account only a single sender and a single receiver. Realistic settings, however, involve multiple senders and multiple receivers. This means that, on the one hand, encryptions toward the same receivers will be made with respect to different sender seeds. On the other hand, senders will use the same seed to generate randomness for encryptions toward different receivers. To achieve security in these settings, we extend the games $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp2}}(\mathcal{A})$ and $\mathbf{G}_{\text{NPE,NG}}^{\text{nbp1}}(\mathcal{A})$ by the following oracle:

```

ENC2( $ek, m, \eta$ )
( $n, St$ )  $\leftarrow_s$  NG( $\eta, St$ )
 $c \leftarrow$  NPE.Enc( $ek, xk, m, n$ )
Return  $c$ 

```

The scheme discussed above can easily be shown to achieve the correspondingly modified games; the hedged extractor provides uniform and independent randomness for each encryption.

Extending this game to multiple senders and multiple receivers is done by generating multiple sets of keys and seeds in the game and extending the oracles with arguments to select the desired sender and/or receiver of the messages to be processed. The proof then follows by two hybrid arguments, one for reducing the number of senders, and one for reducing the number of receivers. The oracle ENC2 is required in the step to reduce the number of receivers to simulate encryptions toward the receivers not captured in the game.

5 Nonce-based signatures

In this section we define and construct nonce-based digital signature schemes.

BACKGROUND. Eliminating randomness in signing is not new and is easily done. A simple way to convert a given randomized EUF-CMA digital signature scheme DS into a deterministic one is as follows. Let F be a PRF. The key-generation algorithm lets $(sk, vk) \leftarrow_s \text{DS.Kg}$ and $fk \leftarrow_s F.\text{Keys}$, and stores the pair (fk, sk) as the secret signing key of the new scheme. A signature on a message $m \in \{0, 1\}^*$ is then computed by first evaluating $r \leftarrow F(fk, m)$ and then $s \leftarrow \text{DS.Sig}(sk, m; r)$. This method goes back to MNPV [23] and it is easy to show that it works, meaning the constructed, deterministic signature scheme retains the EUF-CMA security of the starting randomized one assuming F is a PRF.

The above solution, however, changes the secret key, which is not always desirable. For example it may be a problem to retrofit deployed schemes with the modification, or if the same signature key is used by multiple applications and the format cannot easily be changed. A folklore solution is to leave the keys unchanged and obtain the coins r via a random oracle applied to the existing secret key sk and the message. In the case that DS is ECDSA, this was proven to work by KM [21]. It was proven to work in general (meaning, for any base EUF-CMA scheme) by BPS [9]. Such de-randomization is used in the Ed25519 signature scheme [11] and is specified for DSA and ECDSA in an RFC by Pornin [25].

NONCE-BASED SIGNATURES. In our model, the signer has a secret key as well as a seed. Signing uses both these and a nonce, and is deterministic. If the seed (and secret key) are kept private, we get the usual EUF-CMA level of security, regardless of how nonces are generated. So far this is providing the same security as deterministic signature schemes. The added condition is that if the seed is exposed (but the secret key isn't) then we still retain security as long as the nonces are unpredictable.

The secret key and seed are held by the same entity, namely the signer. So one may ask how it could be that the seed is exposed but the secret key isn't. That is, either the system is secure, in which case both are secure, or not, in which case both are exposed. If so, indeed, nonce-based signatures do not provide anything over and above classical deterministic signatures. However, we can imagine settings where the level of security for the secret key and seed are different. For example the secret key may be already stored in hardware, and the seed not. The seed may be stored at a different place than the signature scheme's secret key, and it may be re-generated at any frequency that seems appropriate for the application (ranging from never, at fixed time intervals, at every system reboot, or at every signature operation). Being a signer-only modification of the signature generation, a user may also use different seeds on different machines, or use the modified scheme with the standard probabilistic one.

The scheme we propose is again based on hedged extractors.

DEFINITIONS. A *nonce-based signature scheme* NDS specifies the following. Signer key-generation algorithm NDS.Kg returns a signature key sk , a verification key

<p>Game $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A})$ / Game $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A})$</p> <p>$(sk, vk, xk) \leftarrow_{\\$} \text{NDS.Kg}$ $St \leftarrow \varepsilon ; M \leftarrow \emptyset$ $(m, s) \leftarrow_{\\$} \mathcal{A}^{\text{SIG,RO}}(vk, \boxed{xk})$ $v \leftarrow \text{NDS.vrf}(vk, m, s)$ Return $((v = 1) \wedge (m \notin M))$</p> <p>$\text{SIG}(m, \eta)$ $(n, St) \leftarrow_{\\$} \text{NG}(\eta, St)$ $s \leftarrow \text{NDS.sign}(sk, xk, m, n)$ $M \leftarrow M \cup \{m\}$ Return s</p> <p>$\text{RO}(x, l)$ If $T[x, l] = \perp$ then $T[x, l] \leftarrow_{\\$} \{0, 1\}^l$ Return $T[x, l]$</p>	<p>Algorithm NDS.Kg $xk \leftarrow_{\\$} \text{HE.Keys}$ $(sk, vk) \leftarrow_{\\$} \text{DS.Kg}$ Return (sk, vk, xk)</p> <p>Algorithm $\text{NDS.sign}^{\text{RO}}(sk, xk, m, n)$ $r \leftarrow \text{HE}^{\text{RO}}(xk, (m, n))$ $c \leftarrow \text{DS.Sig}(sk, m; r)$ Return c</p>
---	--

Fig. 8. Games for security goals for nonce-based digital scheme NDS relative to nonce-generator NG, and our nonce-based digital signature scheme.

vk , and a seed xk . *Deterministic* signature algorithm NDS.sign takes sk , xk , message $m \in \{0, 1\}^*$, and nonce n from nonce set NDS.NS , to return a signature $s \in \{0, 1\}^{\text{NDS.ol}}$. Deterministic verification algorithm NDS.vrf takes vk , message $m \in \{0, 1\}^*$, and candidate signature $s \in \{0, 1\}^{\text{NDS.ol}}$, to return a bit $b \in \{0, 1\}$. The scheme NDS is correct if for all $(sk, vk, xk) \in [\text{NDS.Kg}]$, all $m \in \{0, 1\}^*$ and $n \in \text{NDS.NS}$, the verification of true signatures succeeds: $\text{NDS.vrf}(vk, m, \text{NDS.sign}(sk, xk, m, n)) = 1$.

To formalize security, we consider the games $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A})$ and $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A})$ in Fig. 8 associated to nonce-based signature scheme NDS , nonce generator NG returning nonces in NDS.NS , and adversary \mathcal{A} , where the second game includes the boxed code and the first does not. We let

$$\text{Adv}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A})], \text{ and}$$

$$\text{Adv}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A})].$$

As usual the games are described in the ROM, with standard-model definitions are derived by considering only schemes and adversaries that do not query the RO oracle. The difference between the games is tiny, and in just one line of the code, namely that in the second game, the adversary gets the seed xk as an additional input. The first game captures the case that the seed is not exposed, and we will guarantee security for any nonce generator. The second game captures the case that the seed is exposed, in which case we will provide security for unpredictable nonce generators.

Regular signature schemes can be viewed as a special case of nonce-based ones where the seed xk is defined to be the empty string and security is measured relative only to the nonce generator that always returns a uniformly random string.

In game $\mathbf{Adv}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A})$, we say that adversary \mathcal{A} is *agnostic* if its SIG queries do not depend on the seed xk . More formally, there exists a pair $(\mathcal{A}_1, \mathcal{A}_2)$ of algorithms such that $\mathcal{A}^{\text{SIG,RO}}(vk, xk)$ does the following:

$$St \leftarrow_s \mathcal{A}_1^{\text{SIG,RO}}(vk); (m, s) \leftarrow \mathcal{A}_2^{\text{RO}}(xk, St); \text{Return } (m, s).$$

SCHEME. We specify a transform **R2NDS** that takes a (standard, randomized) signature scheme **DS** and a hedged extractor **HE** and returns the nonce-based signature scheme $\text{NDS} = \mathbf{R2NDS}[\text{DS}, \text{HE}]$ whose algorithms are described in Fig. 8. The nonce space is the same as for the hedged extractor, i.e. $\text{NDS.NS} = \text{HE.NS}$. The length of the signatures is preserved, $\text{NDS.ol} = \text{DS.ol}$. The construction requires that **HE** provides sufficient randomness, i.e., $\text{DS.rl} = \text{HE.ol}$.

We first show that the described scheme **NDS** is indeed secure according to the game $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A})$, that is, in case the seed is not exposed. In this case and if the hedged extractor is a good pseudo-random function, we achieve the same security guarantees as achieved by the original scheme if proper randomness is used.

Theorem 7. *Let **DS** be a (standard, randomized) digital-signature scheme. Let **HE** be a hedged extractor. Let nonce-based digital signature scheme $\text{NDS} = \mathbf{R2NDS}[\text{DS}, \text{HE}]$ be associated to them as above. Let **NG** be a nonce generator. Let \mathcal{A} be an adversary making at most q_1 queries to its SIG oracle and q_2 queries to its RO oracle. Then the proof specifies adversaries \mathcal{B} and \mathcal{G} such that*

$$\mathbf{Adv}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{G}) + \mathbf{Adv}_{\text{DS}}^{\text{uf}}(\mathcal{B}), \quad (16)$$

where adversary \mathcal{B} makes at most q_1 queries to its SIG oracle; besides emulating a Random Oracle it performs almost the same computation as \mathcal{A} . Adversary \mathcal{G} makes at most q_1 queries to its ROR oracle and q_2 queries to its RO oracle; in terms of computation it generates a key pair for **DS** and computes q_1 signatures.

Proof. Adversary \mathcal{B} against **DS** behaves as follows. When started with input vk , it executes $\mathcal{A}(vk)$. Upon a query $\text{SIG}(m, \eta)$ from \mathcal{A} , adversary \mathcal{B} queries $\text{SIG}(m)$, obtaining a signature s , and returns s to \mathcal{A} . Queries to the oracle **RO** are exactly as in the game, that is, by lazy sampling of a random function.

Adversary \mathcal{G} against the hedged extractor behaves as follows. It generates a key pair $(sk, vk) \leftarrow_s \text{DS.Kg}$, and then runs adversary \mathcal{A} on input vk . Upon a query $\text{SIG}(m, \eta)$ from \mathcal{A} , adversary \mathcal{G} calls $\text{FN}(m, \eta)$ to obtain a random string r . It then computes $s \leftarrow \text{DS.Sig}(sk, m; r)$ and returns s . Potential queries $\text{RO}(x, l)$ are referred to the random oracle provided in the game. When \mathcal{A} provides its output (m, s) , then \mathcal{G} outputs the result of $\text{DS.Ver}(vk, m, s) = 1$.

We define a hybrid game **G** that is defined almost identically with $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A})$, with the only difference that the randomness used in SIG queries is uniformly random instead of derived via the hedged extractor. The difference with the game $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf1}}(\mathcal{A})$ is highlighted in Fig. 9.

Game G	Game H
\vdots $\text{SIG}(m, \eta)$ $(n, St) \leftarrow_s \text{NG}(\eta, St)$ $s \leftarrow \text{NDS.sign}(sk, xk, m, n)$ <div style="border: 1px solid black; padding: 2px; display: inline-block; margin: 2px;"> $r \leftarrow_s \{0, 1\}^\ell$ $s \leftarrow \text{DS.Sig}(sk, m; r)$ </div> $M \leftarrow M \cup \{m\}$ Return s \vdots	\vdots $\text{SIG}(m, \eta)$ $(n, St) \leftarrow_s \text{NG}(\eta, St)$ $s \leftarrow \text{NDS.sign}(sk, xk, m, n)$ <div style="border: 1px solid black; padding: 2px; display: inline-block; margin: 2px;"> $r \leftarrow_s \{0, 1\}^\ell$ $s \leftarrow \text{DS.Sig}(sk, m; r)$ </div> $M \leftarrow M \cup \{m\}$ Return s \vdots

Fig. 9. Intermediate games used in the proofs of Theorem 7 (left) and Theorem 8 (right).

The view of \mathcal{A} in G is the same as in $\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{B})$; in both cases the signatures are computed with fresh randomness.

We observe that

$$\begin{aligned}
 2 \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})] &= \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})|b = 1] + \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})|b = 0] \\
 &= \Pr[\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})|b = 1] + 1 - \Pr[\neg \mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})|b = 0] \\
 &= \Pr[\mathbf{G}_{\text{NDS,NG}}^{\text{nbufl}}(\mathcal{A})] + 1 - \Pr[\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{B})],
 \end{aligned}$$

the reason is that if the bit b is chosen as $b = 0$ in $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})$, then the view of \mathcal{A} is exactly as in $\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{B})$; all signatures are generated using fresh randomness. Analogously, if b is chosen as $b = 1$ in $\mathbf{G}_{\text{HE}}^{\text{prf}}(\mathcal{G})$, then the view of \mathcal{A} is exactly as in $\mathbf{G}_{\text{NDS,NG}}^{\text{nbufl}}(\mathcal{A})$ with $b = 0$; all signatures are generated with randomness computed via HE and NG from the message m and the input η . For $b = 1$ the probability of \mathcal{G} guessing correctly is the same as the probability of \mathcal{A} forging the signature, but for $b = 0$ the probability of \mathcal{G} guessing correctly is the same as the probability of \mathcal{A} *not* forging a signature. The above equation implies

$$\mathbf{Adv}_{\text{HE}}^{\text{prf}}(\mathcal{G}) = \mathbf{Adv}_{\text{NDS,NG}}^{\text{nbufl}}(\mathcal{A}) - \mathbf{Adv}_{\text{DS}}^{\text{uf}}(\mathcal{B})$$

and therefore the inequality claimed in the theorem statement. \square

The second statement concerns the security in the sense of $\mathbf{G}_{\text{NDS,NG}}^{\text{nbufl2}}(\mathcal{A})$, that is, the signatures are indeed secure even if the seed is exposed, as long as the nonces contain a sufficient amount of min-entropy. For the scheme based on our standard-model hedged extractor, we again restrict the statement to agnostic adversaries \mathcal{A} . The scheme based on our ROM-based hedged extractor is again secure against all (i.e., not necessarily agnostic) adversaries.

Theorem 8. *Let DS be a (standard, randomized) digital-signature scheme. Let HE be a hedged extractor. Let nonce-based digital signature scheme NDS = R2NDS[DS, HE] be associated to them as above. Let NG be a nonce generator. Let \mathcal{A} be an adversary making at most q_1 queries to its SIG oracle and q_2 queries to its RO oracle. Then the proof specifies adversaries \mathcal{B} and \mathcal{G} such that*

$$\mathbf{Adv}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{HE,NG}}^{\text{ror}}(\mathcal{G}) + \mathbf{Adv}_{\text{DS}}^{\text{uf}}(\mathcal{B}), \quad (17)$$

where adversary \mathcal{B} makes at most q_1 queries to its SIG oracle; besides emulating a Random Oracle it performs almost the same computation as \mathcal{A} . Adversary \mathcal{G} makes at most q_1 queries to its ROR oracle and q_2 queries to its RO oracle; in terms of computation it generates a key pair for DS and computes q_1 signatures. Furthermore, if \mathcal{A} is agnostic, then \mathcal{G} is also agnostic.

Proof. Adversary \mathcal{B} against DS behaves as follows. When started with input vk , it samples $xk \leftarrow_s \text{HE.Keys}$ and executes $\mathcal{A}(vk, xk)$. Upon a query $\text{SIG}(m, \eta)$ from \mathcal{A} , adversary \mathcal{B} queries $\text{SIG}(m)$, obtaining a signature s , and returns s to \mathcal{A} . Queries to the oracle RO are exactly as in the game, that is, by lazy sampling of a random function.

Adversary \mathcal{G} against the hedged extractor behaves as follows. It obtains a seed $xk \in \text{HE.Keys}$ and generates a key pair $(sk, vk) \leftarrow_s \text{DS.Kg}$, and then runs adversary \mathcal{A} on input (vk, xk) . Upon a query $\text{SIG}(m, \eta)$ from \mathcal{A} , adversary \mathcal{G} calls $\text{ROR}(m, \eta)$ to obtain a random string r . It then computes $s \leftarrow \text{DS.Sig}(sk, m; r)$ and returns s . Potential queries $\text{RO}(x, l)$ are referred to the random oracle provided in the game. When \mathcal{A} provides its output (m, s) , then \mathcal{G} outputs the result of $\text{DS.Ver}(vk, m, s) = 1$.

We define a hybrid game H that is defined almost identically with $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A})$, with the only difference that the randomness used in SIG queries is uniformly random instead of derived via the hedged extractor. The difference with the game $\mathbf{G}_{\text{NDS,NG}}^{\text{nbuf2}}(\mathcal{A})$ is highlighted in Fig. 9.

The view of \mathcal{A} in H is the same as in $\mathbf{G}_{\text{DS}}^{\text{uf}}(\mathcal{B})$; in both cases the signatures are computed with fresh randomness. The remainder of the proof follows almost exactly as in Theorem 7. \square

Acknowledgments

Bellare was supported in part by NSF grants CNS-1526801 and CNS-1228890, ERC Project ERCC FP7/615074 and a gift from Microsoft. Tackmann was supported in part by the Swiss National Science Foundation (SNF) via Fellowship No. P2EZP2_155566 and by NSF grant CNS-1228890. We thank the anonymous reviewers for their helpful and detailed comments.

References

1. OpenSSL predictable random number generator. Debian security advisory, May 2008. Available at <http://www.debian.org/security/2008/dsa-1571>.

2. B. Barak, Y. Dodis, H. Krawczyk, O. Pereira, K. Pietrzak, F.-X. Standaert, and Y. Yu. Leftover hash lemma, revisited. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 1–20. Springer, Heidelberg, Aug. 2011.
3. M. Bellare, A. Boldyreva, and A. O’Neill. Deterministic and efficiently searchable encryption. In A. Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 535–552. Springer, Heidelberg, Aug. 2007.
4. M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged public-key encryption: How to protect against bad randomness. In M. Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 232–249. Springer, Heidelberg, Dec. 2009.
5. M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among notions of security for public-key encryption schemes. In H. Krawczyk, editor, *CRYPTO’98*, volume 1462 of *LNCS*, pages 26–45. Springer, Heidelberg, Aug. 1998.
6. M. Bellare, M. Fischlin, A. O’Neill, and T. Ristenpart. Deterministic encryption: Definitional equivalences and constructions without random oracles. In D. Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 360–378. Springer, Heidelberg, Aug. 2008.
7. M. Bellare and V. T. Hoang. Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 627–656. Springer, Heidelberg, Apr. 2015.
8. M. Bellare, T. Kohno, and V. Shoup. Stateful public-key cryptosystems: How to encrypt with one 160-bit exponentiation. In A. Juels, R. N. Wright, and S. Vimercati, editors, *ACM CCS 06*, pages 380–389. ACM Press, Oct. / Nov. 2006.
9. M. Bellare, B. Poettering, and D. Stebila. From identification to signatures, tightly: A framework and generic transforms. Cryptology ePrint Archive, Report 2015/1157, 2015. <http://eprint.iacr.org/2015/1157>.
10. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
11. D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. In B. Preneel and T. Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, Sept. / Oct. 2011.
12. A. Boldyreva, S. Fehr, and A. O’Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In D. Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 335–359. Springer, Heidelberg, Aug. 2008.
13. D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In I. Ray, N. Li, and C. Kruegel, editors, *ACM CCS 15*, pages 668–679. ACM Press, Oct. 2015.
14. S. Checkoway, M. Fredrikson, R. Niederhagen, A. Everspaugh, M. Green, T. Lange, T. Ristenpart, D. J. Bernstein, J. Maskiewicz, and H. Shacham. On the practical exploitability of dual ec in tls implementations. In *USENIX Security*, volume 1, 2014.
15. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In H. Krawczyk, editor, *CRYPTO’98*, volume 1462 of *LNCS*, pages 13–25. Springer, Heidelberg, Aug. 1998.
16. Y. Dodis, C. Ganesh, A. Golovnev, A. Juels, and T. Ristenpart. A formal treatment of backdoored pseudorandom generators. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 101–126. Springer, Heidelberg, Apr. 2015.

17. B. Fuller, A. O'Neill, and L. Reyzin. A unified approach to deterministic encryption: New constructions and a connection to computational entropy. In R. Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 582–599. Springer, Heidelberg, Mar. 2012.
18. S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
19. N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman. Mining your ps and qs: Detection of widespread weak keys in network devices. In *USENIX Security Symposium*, pages 205–220, 2012.
20. R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24. ACM Press, May 1989.
21. N. Kobitz and A. Menezes. The random oracle model: A twenty-year retrospective. Cryptology ePrint Archive, Report 2015/140, 2015. <http://eprint.iacr.org/2015/140>.
22. H. Krawczyk. LFSR-based hashing and authentication. In Y. Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 129–139. Springer, Heidelberg, Aug. 1994.
23. D. M'Raihi, D. Naccache, D. Pointcheval, and S. Vaudenay. Computational alternatives to random number generators. In S. E. Tavares and H. Meijer, editors, *SAC 1998*, volume 1556 of *LNCS*, pages 72–80. Springer, Heidelberg, Aug. 1999.
24. K. G. Paterson, J. C. N. Schuldt, and D. L. Sibborn. Related randomness attacks for public key encryption. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 465–482. Springer, Heidelberg, Mar. 2014.
25. T. Pornin. Deterministic usage of the digital signature algorithm (DSA) and elliptic curve digital signature algorithm (ECDSA). IETF RFC 6979, 2013.
26. A. Raghunathan, G. Segev, and S. P. Vadhan. Deterministic public-key encryption for adaptively chosen plaintext distributions. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 93–110. Springer, Heidelberg, May 2013.
27. T. Ristenpart and S. Yilek. When good randomness goes bad: Virtual machine reset vulnerabilities and hedging deployed cryptography. In *NDSS 2010*. The Internet Society, Feb. / Mar. 2010.
28. P. Rogaway. Authenticated-encryption with associated-data. In V. Atluri, editor, *ACM CCS 02*, pages 98–107. ACM Press, Nov. 2002.
29. P. Rogaway. Nonce-based symmetric encryption. In B. K. Roy and W. Meier, editors, *FSE 2004*, volume 3017 of *LNCS*, pages 348–359. Springer, Heidelberg, Feb. 2004.
30. P. Rogaway and T. Shrimpton. A provable-security treatment of the key-wrap problem. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 373–390. Springer, Heidelberg, May / June 2006.
31. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
32. S. Yilek. Resettable public-key encryption: How to encrypt on a virtual machine. In J. Pieprzyk, editor, *CT-RSA 2010*, volume 5985 of *LNCS*, pages 41–56. Springer, Heidelberg, Mar. 2010.
33. D. Zuckerman. Simulating BPP using a general weak random source. In *32nd FOCS*, pages 79–89. IEEE Computer Society Press, Oct. 1991.