

Instantiability of RSA-OAEP under Chosen-Plaintext Attack

EIKE KILTZ*

ADAM O'NEILL†

ADAM SMITH‡

Abstract

We show that the widely deployed RSA-OAEP encryption scheme of Bellare and Rogaway (Eurocrypt 1994), which combines RSA with two rounds of an underlying Feistel network whose hash (*i.e.*, round) functions are modeled as random oracles, meets indistinguishability under chosen-plaintext attack (IND-CPA) in the *standard model* based on simple, non-interactive, and non-interdependent assumptions on RSA and the hash functions. To prove this, we first give a result on a more general notion called “padding-based” encryption, saying that such a scheme is IND-CPA if (1) its underlying padding transform satisfies a “fooling” condition against small-range distinguishers on a class of high-entropy input distributions, and (2) its trapdoor permutation is sufficiently *lossy* as defined by Peikert and Waters (STOC 2008). We then show that the first round of OAEP satisfies condition (1) if its hash function is t -wise independent for appropriate t and that RSA satisfies condition (2) under the Φ -Hiding Assumption of Cachin *et al.* (Eurocrypt 1999).

This appears to be the first non-trivial *positive* result about the instantiability of RSA-OAEP. In particular, it increases our confidence that chosen-plaintext attacks are unlikely to be found against the scheme. In contrast, RSA-OAEP’s predecessor in PKCS #1 v1.5 was shown to be vulnerable to such attacks by Coron *et al.* (Eurocrypt 2000).

Keywords: RSA, OAEP, padding-based encryption, lossy trapdoor functions, leftover hash lemma, standard model.

*Ruhr-Universität Bochum, Germany. Email: eike.kiltz@ruhr-uni-bochum.de. URL: <http://www.kiltz.net>.

†Boston University, USA. Email: amoneill@bu.edu. URL: <http://cs-people.bu.edu/amoneill/>.

‡Pennsylvania State University, USA. Email: asmith@cse.psu.edu. URL: <http://www.cse.psu.edu/~asmith>.

Contents

1	Introduction	3
1.1	Our Contributions	3
1.2	Related Work	6
2	Preliminaries	7
3	Padding-Based Encryption from Lossy TDP + Fooling Extractor	8
3.1	Background and Tools	8
3.2	The Result	10
4	OAEP as a Fooling Extractor	11
4.1	OAEP	11
4.2	Analysis	12
5	Lossiness of RSA	15
5.1	Background on RSA and Notation	15
5.2	RSA Lossy TDP from Φ -Hiding	16
5.3	RSA Lossy TDP from Multi-Prime Φ -Hiding	18
5.4	Small-Exponent RSA LTDP from 2-vs- m Primes	19
6	Instantiating RSA-OAEP	20
A	Proof of Lemma 4.5	25
B	Security of OAEP Under Key-Independent Chosen-Plaintext Attack	27

1 Introduction

The RSA-OAEP encryption scheme was designed by Bellare and Rogaway [7] as a drop-in replacement for RSA PKCS #1 v1.5 [1] with provable security guarantees. In particular, it follows the same paradigm as RSA PKCS #1 v1.5 in that it encrypts a message of less than k bits to a k -bit ciphertext (where k is the modulus size) by first applying a fast, randomized, and invertible “padding transform” to the message before applying RSA. In the case of RSA-OAEP, the underlying padding transform (which is itself called ‘OAEP’¹) embeds a message m and random coins r as $s\|(H(s) \oplus r)$ where ‘ $\|$ ’ denotes concatenation, $s = (m\|0^{k_1}) \oplus G(r)$ for some parameter k_1 , and G and H are hash functions (see Figure 2 on p. 11). In contrast, PKCS #1 v1.5 essentially just concatenates m with r .

RSA-OAEP was designed using the random oracle (RO) methodology [6]. This means that, for the security analysis, its hash functions are modeled as truly random functions, available to all parties only via oracle access. When the scheme is implemented in practice, these oracles are heuristically “instantiated” in certain ways using a cryptographic hash function like SHA1. A cryptographic hash function (or a function built from one) is certainly not random nor computable only via an oracle (it has a short, public description), but schemes designed using this methodology are hoped to be secure. Unfortunately, a series of works, starting with the seminal paper of Canetti *et al.* [20] showed that there are schemes secure in the RO model that are insecure under *every* instantiation of its oracles; such RO model schemes are called *uninstantiable*. Thus, to gain confidence in an RO model scheme, we should show that it is not uninstantiable, *i.e.*, that its oracles admit a secure instantiation by efficiently computable functions under well-defined assumptions. Then, when we instantiate the scheme, we know that our goal is at least plausible. We feel this is especially important for a scheme such as RSA-OAEP, which is by now widely standardized and deployed.

Yet, while RO model schemes continue to be proposed, relatively few have been shown to be instantiable. In particular, we are not aware of *any* result showing instantiability of RSA-OAEP, even under a relatively modest security model. In fact, the scheme has come under criticism lately due to several works (discussed in Section 1.2) showing the impossibility of certain types of instantiations under chosen-ciphertext attack (IND-CCA). Fortunately, we bring some good news: We give reasonable assumptions under which RSA-OAEP is secure against *chosen-plaintext attack* (IND-CPA). We believe this is an important step towards a better understanding of the scheme’s security.

1.1 Our Contributions

Our result on the instantiability of RSA-OAEP is obtained via three steps or other results. (These other results may also be of independent interest.) First, we show a general result on the instantiability of “padding-based encryption,” of which f -OAEP is a special case, under the assumption that the underlying padding transform is what we call a *fooling extractor* and the trapdoor permutation is *lossy* [45]. We then show that OAEP and RSA satisfy the respective conditions under suitable assumptions.

PADDING-BASED ENCRYPTION WITHOUT ROS. Our first result is a general theorem about *padding-based encryption* (PBE), a notion formalized recently by Kiltz and Pietrzak [37].² PBE generalizes the design methodology of PKCS #1 and RSA-OAEP we already mentioned. Namely, we start with a k -bit to k -bit trapdoor permutation (TDP) that satisfies a weak security notion like one-wayness.

¹We often use the same terminology for ‘ f -OAEP,’ which refers to OAEP using an abstract TDP f , with the meaning hopefully clear from context.

²Such schemes were called “simple embedding schemes” by Bellare and Rogaway [7], who discussed them only on an intuitive level.

To “upgrade” the TDP to an encryption scheme satisfying a strong security notion like IND-CPA, we design an invertible “padding transform” which embeds a plaintext and random coins into a k -bit string, to which we then apply the TDP. This methodology is quite natural and has long been prevalent in practice, motivating the design of OAEP and later schemes such as SAEP [14] and PSS-E [24]. The latter were all designed and analyzed in the RO model.

We show that the RO model is *unnecessary* in the design and analysis of IND-CPA secure PBE. To do so, we formulate a connection between PBE and a new notion we call “fooling extractor for small-range distinguishers” or just “fooling extractor.” Intuitively, a fooling extractor is a kind of randomness extractor that transforms a high-entropy source into something that looks random to any function (or distinguisher) with a *small range*.³ Our result says that if the padding transform of a PBE scheme is an “adaptive” fooling extractor for sources of the form (m, R) — where m is a plaintext and R is the random coins (which we call “encryption sources”) — and its TDP is sufficiently *lossy* (the logarithm of its range size should be slightly less than the length of R) as defined by Peikert and Waters [45], then the PBE scheme is IND-CPA. Here “adaptive” means that m may *depend* on the choice of the extractor seed. We call such padding transforms “encryption-compatible.”

OAEP FOOLS SMALL-RANGE DISTINGUISHERS. Our second result says that the OAEP padding transform is encryption-compatible if the hash function G is t -wise independent for appropriate t (roughly, proportional to the allowed message length). Note that no restriction is put on hash function H ; in particular, neither hash function is modeled as an RO. The inspiration for our proof comes from the “Crooked” Leftover Hash Lemma (LHL) of Dodis and Smith [27], especially its application to deterministic encryption by Boldyreva *et al* [11] (who also gave a simpler proof). Qualitatively, the Crooked LHL says that $(K, f(\Pi(K, X)))$ looks like $(K, f(U))$ for any small-range function f , pairwise-independent function Π keyed by K , and high-entropy source X ; in our terminology, this says that a pairwise-independent function is a fooling extractor for such X . In our application, we might naïvely view Π as the OAEP. There are two problems with this. First, OAEP is *not* pairwise independent, even in the RO model. Second, showing that OAEP is encryption-compatible entails showing adaptivity (as defined above), whereas in the lemma K is independent of X .

To solve the first problem, we show that the Crooked LHL can be strengthened to say that $K, f(X, \Pi(K, X))$ looks like $K, f(X, U)$; *i.e.*, that $\Pi(K, X)$ looks random to f *even given* X . The proof is an extension of the proof of the Crooked LHL in [11]. Then, by viewing X as the random coins in OAEP and Π as the hash function G , we can conclude that OAEP is a fooling extractor for any fixed encryption source (m, R) , where m is *independent* of K (note that our analysis does not use any properties of H —the only fact we use about the second Feistel round is that it is invertible).

To solve the second problem, we extend an idea of Trevisan and Vadhan [50] to our setting and show that if G is t -wise independent for large enough t , the error probability for a particular encryption source is so small that we can take a union bound over all possible m and conclude that OAEP is in fact adaptive, meaning it is indeed encryption-compatible. Interestingly, we obtain better parameters in the case that f is *regular*, meaning every preimage set has the same size. However, our analysis still goes through assuming that every preimage set is sufficiently large, which we show can always be assumed with some loss in parameters.

LOSSINESS OF RSA. To instantiate RSA-OAEP, it remains to show lossiness of RSA. Our final result is that RSA is indeed lossy under reasonable assumptions. Intuitively, lossiness [45] means that there is an alternative, “lossy” key generation algorithm that outputs a public key indistinguishable from a normal one, but which induces a small-range (uninvertible) function. We first show lossiness of RSA

³In the formal definition we actually consider an “external” distinguisher who gets the extractor seed; see Section 3 for details.

under the Φ -Hiding Assumption (Φ A) of Cachin, Micali, and Stadler [17]. Φ A has been used as the basis for a number of efficient protocols, *e.g.*, [17, 16, 29, 31]. Φ A states roughly that given an RSA modulus $N = pq$, it is hard to distinguish primes that divide $\phi(N) = (p-1)(q-1)$ from those that do not. Normal RSA parameters (N, e) are such that $\gcd(e, \phi(N)) = 1$. Under Φ A, we may alternatively choose (N', e) such that e divides $p-1$. The range of the RSA function is then reduced by a factor $1/e$. To resist known attacks, we can take the bit-length of e up to almost $1/4$ that of N , giving RSA lossiness of almost $k/4$ bits, where k is the modulus length.⁴

In practice, however, e is usually chosen to be small for efficiency reasons. We observe that in this case more lossiness can be achieved by considering *multi-prime* RSA where $N = p_1 \cdots p_m$ for $m \geq 2$ (for a fixed modulus length), and in the lossy case choosing (N', e) such that e divides p_i for all $1 \leq i \leq m-1$; the range of the RSA function is then reduced by a factor $1/e^{m-1}$. The maximum bit-length of e in this case to avoid our best attack is roughly $k(1/m - 2/m^2)$ where k is the modulus length — this was recently improved to $k(2/3m^{2/3})$ by Herrmann [11] — so for a fixed modulus size we gain in lossiness only for small e . If we assume such multi-prime RSA moduli are indistinguishable from two-prime ones, we can achieve such lossiness in the case of standard (two-prime) RSA as well.

IMPLICATIONS FOR RSA-OAEP. Combining the above results gives that RSA-OAEP is IND-CPA in the standard model under (rather surprisingly, at least to us) simple, non-interactive, and non-interdependent assumptions on RSA and the hash functions. The parameters for RSA-OAEP supported by our proofs are discussed in Section 6. While they are considerably worse than what is expected in practice, we view the upshot of our results not as the concrete parameters they support, but rather that they increase the theoretical backing for the scheme’s security at a more qualitative level, showing it can be instantiated at least for larger parameters. In particular, our results give us greater confidence that chosen-plaintext attacks are unlikely to be found against the scheme; such attacks are known against the predecessor of RSA-OAEP in PKCS #1 v1.5 [23]. That said, we strongly encourage further research to try to improve the concrete parameters.

Moreover, our analysis brings to light to some simple modifications that may increase the scheme’s security. The first is to key the hash function G . Although our results have some interpretation in the case that G is a fixed function (see below), it may be preferable for G to have an explicit, randomly selected key. It is an interesting open question whether our proof can be extended to function families that use shorter keys. The second possible modification is to increase the length of the randomness versus that of the redundancy in the message when encrypting short messages under RSA-OAEP. Of course, we suggest these modifications only in cases where they do not impact efficiency too severely.

USING UNKEYED HASH FUNCTIONS. Formally, our results assume G is randomly chosen from a large family (*i.e.*, it is a keyed hash function). However, our analysis actually shows that *almost every* function (*i.e.*, all but a negligible fraction) from the family yields a secure instantiation; we just do not know an explicit member that works. In other words, it is not strictly necessary that G be randomly chosen. When G is instantiated in practice using a cryptographic hash function, it is plausible that the resulting instantiation is secure.

ON CHOSEN-CIPHERTEXT SECURITY. Any extension of our results to security under chosen-ciphertext attack (IND-CCA) must get around the recent negative results of Kiltz and Pietrzak [37] (which we discuss in more detail in Section 1.2). We suggest two possible approaches.

The first is based on the fact that, by the results of Bellare and Palacio [5], the notion of plaintext awareness (PA) + IND-CPA implies IND-CCA. Thus, in order to show IND-CCA security of RSA-

⁴We remark that the recent attacks on Φ A [48] are for moduli of a special form that does not include RSA.

OAEP in the standard model it suffices, by our results, to show PA (which is an orthogonal property to privacy). To show the latter one could try to use non-black-box assumptions on H along the lines of [19]. We leave a detailed investigation to future work.

The second is to make stronger assumptions on the TDP. For example, although it does not hold for RSA, we can show IND-CCA security of OAEP by assuming the TDP satisfies a notion of “adaptive lossiness,” a strengthening to the notion of adaptivity for TDFs studied in [35]. Informally, we say $\mathcal{F}$ is *adaptive lossy* (cf. the definition of lossiness in Section 2) if the function $g_c(x) := f(x)$ where $(f, f^{-1}) \leftarrow \mathcal{F}(x; \text{Coins})$ is lossy for any Coins (here $\mathcal{F}$ takes an auxiliary input x), and moreover (f, x) where $(f, f^{-1}) \xleftarrow{\$} \mathcal{F}(x^*)$ is indistinguishable for any x^* *even oracle access* to f^{-1} . In the case of TDFs, we can show that the construction of adaptive TDF in [35] from lossy+ABO TDFs satisfies this strengthening (but we do not know a construction of adaptive lossy TDP).

1.2 Related Work

SECURITY OF OAEP IN THE RO MODEL. In their original paper [7], Bellare and Rogaway showed that OAEP is IND-CPA assuming the TDP is one-way. They further showed it achieves a notion they called “plaintext awareness.” Subsequently, Shoup [49] observed that the latter notion is too weak to imply security against chosen-ciphertext attacks, and in fact there is no black-box proof of IND-CCA security of OAEP based on one-wayness of the TDP. Fortunately, Fujisaki *et al.* [28] proved that OAEP is nevertheless IND-CCA assuming so-called “partial-domain” one-wayness, and that partial-domain one-wayness and (standard) one-wayness of RSA are equivalent.

SECURITY OF OAEP WITHOUT ROS. Results on instantiability of OAEP have so far mainly been negative. Boldyreva and Fischlin [12] showed that (contrary to a conjecture of Canetti [18]) one cannot securely instantiate even *one* of the two hash functions (while still modeling the other as an RO) of OAEP under IND-CCA by a “perfectly one-way” hash function [18, 21] if one assumes only that f is partial-domain one-way. Brown [15] and Paillier and Villar [43] later showed that there are no “key-preserving” black-box proofs of IND-CCA security of RSA-OAEP based on one-wayness of RSA. Recently, Kiltz and Pietrzak [37] (building on the earlier work of Dodis *et al.* [25] in the signature context) generalized these results and showed that there is no black-box proof of IND-CCA (or even NM-CPA) security of OAEP based on any property of the TDP satisfied by an *ideal* (truly random) permutation.⁵ In fact, their result can be extended to rule out a black-box proof of NM-CPA security of OAEP assuming the TDP is lossy [34], so our results are in some sense *optimal* given our assumptions.

INSTANTIATIONS OF RELATED SCHEMES. A positive instantiation result about a variant of OAEP called OAEP++ [38] (where part of the transform is output in the clear) was obtained by Boldyreva and Fischlin in [13]. They showed an instantiation that achieves (some weak form of) non-malleability under chosen-plaintext attacks (NM-CPA) for random messages, assuming the existence of non-malleable pseudorandom generators (NM-PRGs).⁶ We note that the approach of trying to obtain positive results for instantiations under security notions weaker than IND-CCA originates from their work, and the authors explicitly ask whether OAEP can be shown IND-CPA in the standard model based on reasonable assumptions on the TDP and hash functions.

⁵Note, however, that their result does not rule out such a proof based on other properties of the TDP, non-black-box assumptions on the hash functions, or in the case of a specific TDP like RSA.

⁶In particular, their security notion does *not* imply IND-CPA since they consider random messages. We also point out that it remains an open question whether NM-PRGs can be constructed.

Another line of work has looked at instantiating other RO model schemes related at least in spirit to OAEP. Canetti [18] showed that the IND-CPA scheme in [6] can be instantiated using (a strong form of) perfectly-one way probabilistic hash functions. More recently, the works of Canetti and Dakdouk [19], Pandey *et al.* [44], and Boldyreva *et al.* [10] obtained (partial) instantiations of the earlier IND-CCA scheme of [6]. Hofheinz and Kiltz [33] recently showed an IND-CCA secure instantiation of a variant of the DHIES scheme of [2].

2 Preliminaries

NOTATION AND CONVENTIONS. For a probabilistic algorithm A , by $y \stackrel{\$}{\leftarrow} A(x)$ we mean that A is executed on input x and the output is assigned to y , whereas if S is a finite set then by $s \stackrel{\$}{\leftarrow} S$ we mean that s is assigned a uniformly random element of S . We sometimes use $y \leftarrow A(x; \text{Coins})$ to make A 's random coins explicit. We denote by $\Pr[A(x) \Rightarrow y : \dots]$ the probability that A outputs y on input x when x is sampled according to the elided experiment. Unless otherwise specified, an algorithm may be probabilistic and its running-time includes that of any overlying experiment. We denote by 1^k the unary encoding of the security parameter k . We sometimes suppress dependence on k for readability. For $i \in \mathbb{N}$ we denote by $\{0, 1\}^i$ the set of all binary strings of length i . If s is a string then $|s|$ denotes its length in bits, whereas if S is a set then $|S|$ denotes its cardinality. By ‘||’ we denote string concatenation. All logarithms are base 2.

BASIC DEFINITIONS. Writing $P_X(x)$ for the probability that a random variable X puts on x , the *statistical distance* between random variables X and Y with the same range is given by $\Delta(X, Y) = \frac{1}{2} \sum_x |P_X(x) - P_Y(x)|$. If $\Delta(X, Y)$ is at most ε then we say X, Y are ε -close and write $X \approx_\varepsilon Y$. The *min-entropy* of X is $H_\infty(X) = -\log(\max_x P_X(x))$. A random variable X over $\{0, 1\}^n$ is called a (n, ℓ) -source if $H_\infty(X) \geq \ell$. Let $f : A \rightarrow B$ be a function. We denote by $R(f)$ the *range* of f , i.e., $\{b \in B \mid \exists a \in A, f(a) = b\}$. We call $|R(f)|$ the *range-size* of f . We call f *regular* if each pre-image set is the same size, i.e., $|\{x \in D \mid f(x) = y\}|$ is the same for all $y \in R$.

PUBLIC-KEY ENCRYPTION AND ITS SECURITY. A *public-key encryption scheme* with message-space MsgSp is a triple of algorithms $\mathcal{AE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. The key-generation algorithm $\mathcal{K}$ returns a public key pk and matching secret key sk . The encryption algorithm $\mathcal{E}$ takes pk and a plaintext m to return a ciphertext. The deterministic decryption algorithm $\mathcal{D}$ takes sk and a ciphertext c to return a plaintext. We require that for all messages $m \in \text{MsgSp}$

$$\Pr \left[\mathcal{D}(sk, \mathcal{E}(pk, m)) \neq m : (pk, sk) \stackrel{\$}{\leftarrow} \mathcal{K} \right]$$

is negligible.

To an encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ and an adversary $A = (A_1, A_2)$ we associate a chosen-plaintext attack experiment,

Experiment $\text{Exp}_{\Pi, A}^{\text{ind-cpa}}(k)$
 $b \stackrel{\$}{\leftarrow} \{0, 1\} ; (pk, sk) \stackrel{\$}{\leftarrow} \mathcal{K}(1^k)$
 $(m_0, m_1, \text{state}) \stackrel{\$}{\leftarrow} A_1(pk)$
 $c \stackrel{\$}{\leftarrow} \mathcal{E}(pk, m_b)$
 $d \stackrel{\$}{\leftarrow} A_2(pk, c, \text{state})$
 If $d = b$ then return 1 else return 0

where we require A 's output to satisfy $|m_0| = |m_1|$. Define the *ind-cpa advantage* of A against Π as

$$\mathbf{Adv}_{\Pi,A}^{\text{ind-cpa}}(k) = 2 \cdot \Pr \left[\mathbf{Exp}_{\Pi,A}^{\text{ind-cpa}}(k) \Rightarrow 1 \right] - 1 .$$

LOSSY TRAPDOOR PERMUTATIONS. A *lossy trapdoor permutation (LTDP) generator* [45]⁷ is a pair $\text{LTDP} = (\mathcal{F}, \mathcal{F}')$ of algorithms. Algorithm $\mathcal{F}$ is a usual trapdoor permutation (TDP) generator, namely it outputs a pair (f, f^{-1}) where f is a (description of a) permutation on $\{0, 1\}^k$ and f^{-1} its inverse. Algorithm $\mathcal{F}'$ outputs a (description of a) function f' on $\{0, 1\}^k$. We call $\mathcal{F}$ the “injective mode” and $\mathcal{F}'$ the “lossy mode” of LTDP respectively, and we call $\mathcal{F}$ “lossy” if it is the first component of some lossy TDP. For a distinguisher D , define its *ltdp-advantage* against LTDP as

$$\mathbf{Adv}_{\text{LTDP},D}^{\text{ltdp}}(k) = \Pr \left[D(f) \Rightarrow 1 : (f, f^{-1}) \xleftarrow{\$} \mathcal{F} \right] - \Pr \left[D(f') \Rightarrow 1 : f' \xleftarrow{\$} \mathcal{F}' \right] .$$

We say LTDP has *residual leakage* s if for all f' output by $\mathcal{F}'$ we have $|R(f')| \leq 2^s$. The *lossiness* of LTDP is $\ell = k - s$.

t -WISE INDEPENDENT HASHING. Let $H: \mathcal{K} \times D \rightarrow R$ be a hash function. We say that H is *t -wise independent* if for all distinct $x_1, \dots, x_t \in D$ and all $y_1, \dots, y_t \in R$

$$\Pr \left[H(K, x_1) = y_1 \wedge \dots \wedge H(K, x_t) = y_t : K \xleftarrow{\$} \mathcal{K} \right] = \frac{1}{|R|^t} .$$

In other words, $H(K, x_1), \dots, H(K, x_t)$ are all uniformly and independently random. The standard construction of a t -wise independent hash function uses polynomial evaluation over a finite field and has key length $t \log |D|$.

3 Padding-Based Encryption from Lossy TDP + Fooling Extractor

In this section, we show a general result on how to build IND-CPA secure padding-based encryption (PBE) without using random oracles, by combining a lossy TDP with a “fooling extractor” for small-range distinguishers.

3.1 Background and Tools

We first provide the definitions relevant to our result.

PADDING-BASED ENCRYPTION. The idea behind padding-based encryption (PBE) is as follows: We start with a k -bit to k -bit trapdoor permutation (*e.g.*, RSA) and wish to build a secure encryption scheme. As in [7], we are interested in encrypting messages of less than k bits to ciphertexts of length k . It is well-known that we cannot simply encrypt messages under the TDP directly to achieve strong security. So, in a PBE scheme we “upgrade” the TDP by first applying a randomized and invertible “padding transform” to a message prior to encryption.

Our definition of PBE largely follows the recent formalization in [37]. Let k, μ, ρ be three integers such that $\mu + \rho \leq k$. A *padding transform* $(\pi, \hat{\pi})$ consists of two mappings $\pi: \{0, 1\}^{\mu+\rho} \rightarrow \{0, 1\}^k$ and $\hat{\pi}: \{0, 1\}^k \rightarrow \{0, 1\}^\mu \cup \{\perp\}$ such that π is injective and the following consistency requirement is fulfilled:

$$\forall m \in \{0, 1\}^\mu, r \in \{0, 1\}^\rho : \hat{\pi}(\pi(m \parallel r)) = m .$$

A *padding transform generator* is an algorithm Π that on input 1^k outputs a (description of a) padding transform $(\pi, \hat{\pi})$. Let $\mathcal{F}$ be a k -bit trapdoor permutation generator and Π be a padding transform

⁷We note that [45] actually defines lossy trapdoor *functions*, but the extension to permutations is straightforward.

generator. Define the associated *padding-based encryption scheme* $\mathcal{AE}_\Pi[\mathcal{F}] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ with message-space $\{0, 1\}^\mu$ by

$$\begin{array}{l|l|l} \text{Alg } \mathcal{K}(1^k) & \text{Alg } \mathcal{E}((\pi, f), m) & \text{Alg } \mathcal{D}((\pi, f^{-1}), y) \\ (\pi, \hat{\pi}) \xleftarrow{\$} \Pi(1^k) & r \xleftarrow{\$} \{0, 1\}^\rho ; x \leftarrow \pi(m \| r) & x \leftarrow f^{-1}(y) \\ \pi \leftarrow (\pi, \hat{\pi}) & y \leftarrow f(x) & m \leftarrow \hat{\pi}(x) \\ (f, f^{-1}) \xleftarrow{\$} \mathcal{F}(1^k) & \text{Return } y & \text{Return } m \\ \text{Return } ((\pi, f), (\pi, f^{-1})) & & \end{array}$$

Padding-based encryption schemes have long been prevalent in practice, for example PKCS #1 [1]. While OAEP [7] is the best-known, the notion also captures later schemes such as SAEP [14] and PSS-E [24].

FOOLING EXTRACTORS. We define a new notion that we call “fooling extractor for small-range distinguishers” or just “fooling extractor.” Intuitively, fooling extractors are a type of randomness extractor that “fools” distinguishers with small-range output. We give some more intuition after the formal definition.

Definition 3.1 Let $\text{FExt}: \{0, 1\}^c \times \{0, 1\}^n \rightarrow \{0, 1\}^k$ be a function and let $\mathcal{X} = \{X_1, \dots, X_q\}$ be a class of n -bit sources. We say that FExt *fools range- 2^s distinguishers on $\mathcal{X}$ with probability $1 - \varepsilon$* (or is an (s, ε) -fooling extractor for $\mathcal{X}$) if for all functions f on $\{0, 1\}^k$ with range-size at most 2^s and all $1 \leq i \leq q$:

$$(K, f(\text{FExt}(K, X_i))) \approx_\varepsilon (K, f(U)) ,$$

where K is uniform on $\{0, 1\}^c$ and U is uniform and independent on $\{0, 1\}^n$. We call K the *key* or *seed* of FExt . Note that K is *independent* of i above.

We say that FExt *adaptively* fools range- 2^s distinguishers on $\mathcal{X}$ with probability $1 - \varepsilon$ (or is an adaptive (s, ε) -fooling extractor for $\mathcal{X}$) if for all functions f on $\{0, 1\}^k$ with range-size at most 2^s :

$$\mathbf{E}_{k \xleftarrow{\$} \{0, 1\}^c} \left[\max_{1 \leq i \leq q} \Delta(f(\text{FExt}(k, X_i)), f(U)) \right] \leq \varepsilon .$$

Since $\Delta(K, f(\text{FExt}(K, X_i)), (K, f(U))) = \mathbf{E}_k \Delta(k, f(\text{FExt}(k, X_i)), (k, f(U)))$, the above implies that $(K, f(\text{FExt}(K, X_i))) \approx_\varepsilon (K, f(U))$ for i depending on K (or, put differently, $(K, f(\text{FExt}(K, X_i))) \approx_\varepsilon (K, f(U))$ holds for every i over the *same* choice of K).

As a useful special case, we say that FExt fools range- 2^s regular distinguishers on $\mathcal{X}$ with probability $1 - \varepsilon$ (or is a regular (s, ε) -fooling extractor for $\mathcal{X}$) if we quantify only over *regular* f in the definition. An *adaptive regular* (s, ε) -fooling extractor for $\mathcal{X}$ is defined analogously.

We note that while the intuition given prior to the definition describes fooling the function f , it actually requires fooling an “implicit” or “external” distinguisher that sees both the output $f(\text{FExt}(K, X_i))$ of f and the extractor seed K . This crucial for the definition to be meaningful. Indeed, just asking that $f(\text{FExt}(K, X_i))$ be indistinguishable from $f(U)$ for all small-range functions f is equivalent to asking only that $\text{FExt}(K, X_i)$ be indistinguishable from U . This latter requirement is trivial to achieve (if one is not concerned with key length)—for example, by using K as a one-time pad.

We also note that the concept of fooling extractors was implicit in the work of Dodis and Smith [27] on error-correction without leaking partial information, whose “Crooked” Leftover Hash Lemma establishes in our language that a pairwise-independent function is a (s, ε) -fooling extractor for every singleton (n, ℓ) -source X where $s \leq \ell - 2 \log(1/\varepsilon) + 2$. This lemma was later applied in the context of deterministic public-key encryption by Boldyreva et al. [11] (and indeed this application inspired our work), who also gave a simpler proof.

3.2 The Result

To state our result, we first formalize the concept of *encryption-compatible* padding transforms.

Definition 3.2 Let Π be a padding transform generator whose coins are drawn from Coins . Define the associated function $h_\Pi : \text{Coins} \times \{0, 1\}^{\mu+\rho} \rightarrow \{0, 1\}^k$ by $h(c, m \| r) = \pi(m \| r)$ for all $c \in \text{Coins}, m \in \{0, 1\}^\mu, r \in \{0, 1\}^\rho$, where $(\pi, \hat{\pi}) \leftarrow \Pi(1^k; \text{Coins})$. Define the class $\mathcal{X}_\Pi$ of *encryption sources* associated to Π as containing all sources of the form (m, R) , where $m \in \{0, 1\}^\mu$ is fixed and $R \in \{0, 1\}^\rho$ is uniformly random. (Note that the class $\mathcal{X}_\Pi$ therefore contains 2^μ distinct $(\mu + \rho)$ -bit sources.) We say that Π is (s, ε) -*encryption-compatible* if h_Π as above is an adaptive (s, ε) -fooling extractor for $\mathcal{X}_\Pi$. A *regular* (s, ε) -encryption-compatible padding transform generator is defined analogously.

Theorem 3.3 Let $\text{LTDP} = (\mathcal{F}, \mathcal{F}')$ be an LTDP with residual leakage s , and let Π be an (s, ε) -encryption-compatible padding transform generator. Then for any IND-CPA adversary A against $\mathcal{AE}_\Pi[\mathcal{F}]$ there is a adversary D against LTDP such that for all $k \in \mathbb{N}$

$$\mathbf{Adv}_{\mathcal{AE}, A}^{\text{ind-cpa}}(k) \leq \mathbf{Adv}_{\text{LTDP}, D}^{\text{ltdp}}(k) + \varepsilon.$$

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

Proof: Given $A = (A_1, A_2)$, we define three games, called G_0, G_1, G_2 , in Figure 1. Note that game G_0 is the experiment $\mathbf{Exp}_{\Pi, A}^{\text{ind-cpa}}(k)$ defining IND-CPA security. We claim that for a distinguisher D against LTDP that is simple to construct, we have

$$\frac{1}{2} + \mathbf{Adv}_{\mathcal{AE}_\Pi[\mathcal{F}], A}^{\text{ind-cpa}}(k) = \Pr[G_0 \Rightarrow 1] \tag{1}$$

$$\leq \Pr[G_1 \Rightarrow 1] + \mathbf{Adv}_{\text{LTDP}, D}^{\text{ltdp}}(k) \tag{2}$$

$$\leq \Pr[G_2 \Rightarrow 1] + \mathbf{Adv}_{\text{LTDP}, D}^{\text{ltdp}}(k) + \varepsilon \tag{3}$$

$$= \frac{1}{2} + \mathbf{Adv}_{\text{LTDP}, D}^{\text{ltdp}}(k) + \varepsilon, \tag{4}$$

from which the theorem follows by re-arranging terms. So let us justify the above.

Equation (1) is true by the definition of IND-CPA security.

For (2) we can construct a distinguisher D as required since G_0, G_1 do not use f^{-1} in any way.

Equation (3) is true by the definition of encryption compatibility. Namely, since h_Π in the definition is an adaptive (s, ε) -fooling extractor for $\mathcal{X}_\Pi$, we know the expectation over $\hat{\pi}$ of $\Delta(f(\hat{\pi}(m, R)), f(U))$ is at most ε for m depending on $\hat{\pi}$, so in particular it holds for $m = m_b$ in game G_1 .

Finally, (4) uses the fact that in G_2 no information about b is given to A . Note that the final two steps in the proof are information-theoretic, meaning they do not use any assumption about A 's running-time. ■

Remark 3.4 The analogous result to the above holds for regular LTDPs and regular encryption-compatible padding transforms. That is, if the LTDP is *regular* (meaning $\mathcal{F}'$ is) then it suffices to use a regular encryption-compatible padding transform to obtain the same conclusion. The latter may be easier to design or more efficient than in the general case; indeed, we get better parameters for OAEP in the regular case in Section 4. Furthermore, known examples of LTDPs (including RSA, as shown in Section 5) are regular, although a technical issue about the domain of RSA versus the output range

Game G_0: $b \xleftarrow{\$} \{0, 1\}; (f, f^{-1}) \xleftarrow{\$} \mathcal{F}$ $(\pi, \hat{\pi}) \xleftarrow{\$} \Pi; \mathbf{\Pi} \leftarrow (\pi, \hat{\pi})$ $(m_0, m_1, \text{state}) \xleftarrow{\$} A_1(f, \mathbf{\Pi})$ $r \leftarrow \{0, 1\}^\rho; x \leftarrow \hat{\pi}(m_b \ r)$ $d \xleftarrow{\$} A_2((f, \mathbf{\Pi}), f(x), \text{state})$ If $d = b$ then Return 1 Else Return 0	Game G_1: $b \xleftarrow{\$} \{0, 1\}; f \xleftarrow{\$} \mathcal{F}'$ $(\pi, \hat{\pi}) \xleftarrow{\$} \Pi; \mathbf{\Pi} \leftarrow (\pi, \hat{\pi})$ $(m_0, m_1, \text{state}) \xleftarrow{\$} A_1(f, \mathbf{\Pi})$ $r \xleftarrow{\$} \{0, 1\}^\rho; x \leftarrow \hat{\pi}(m_b \ r)$ $d \xleftarrow{\$} A_2((f, \mathbf{\Pi}), f(x), \text{state})$ If $d = b$ then Return 1 Else Return 0	Game G_2: $b \xleftarrow{\$} \{0, 1\}; f \xleftarrow{\$} \mathcal{F}'$ $(\pi, \hat{\pi}) \xleftarrow{\$} \Pi; \mathbf{\Pi} \leftarrow (\pi, \hat{\pi})$ $(m_0, m_1, \text{state}) \xleftarrow{\$} A_1(f, \mathbf{\Pi})$ $x \xleftarrow{\$} \{0, 1\}^k$ $d \xleftarrow{\$} A_2((f, \mathbf{\Pi}), f(x), \text{state})$ If $d = b$ then Return 1 Else Return 0
--	--	---

Figure 1: Games for the proof of Theorem 3.3. Shaded areas indicate the differences between games.

of OAEP makes it difficult to exploit this for RSA-OAEP and raises an interesting open problem; see Section 6.

4 OAEP as a Fooling Extractor

In this section, we show that the OAEP padding transform of Bellare and Rogaway [7] is encryption-compatible as defined in Section 3 if its initial hash function is t -wise independent for t depending on the message length and lossiness of the TDP.

4.1 OAEP

We recall the OAEP padding transform of Bellare and Rogaway [7], lifted to the “instantiated” setting where hash functions may be keyed. Let $G: \mathcal{K}_G \times \{0, 1\}^\rho \rightarrow \{0, 1\}^\mu$ and $H: \mathcal{K}_H \times \{0, 1\}^\mu \rightarrow \{0, 1\}^\rho$ be hash functions. The associated padding transform generator $\text{OAEP}[G, H]$ on input 1^k returns $(\pi_{K_G, K_H}, \hat{\pi}_{K_G, K_H})$, where $K_G \xleftarrow{\$} \mathcal{K}_G(1^k)$ and $K_H \xleftarrow{\$} \mathcal{K}_H(1^k)$, defined via

Algorithm $\pi_{K_G, K_H}(m \ r)$ $s \leftarrow m \oplus G(K_G, r)$ $t \leftarrow r \oplus H(K_H, s)$ $x \leftarrow s \ t$ Return x	Algorithm $\hat{\pi}_{K_G, K_H}(x)$ $s \ t \leftarrow x$ $r \leftarrow t \oplus H(K_H, s)$ $m \leftarrow s \oplus G(K_G, r)$ Return m
---	--

See Figure 2 for a graphical illustration.




Figure 2: Algorithms $\pi_{K_G, K_H}(m, r)$ and $\hat{\pi}_{K_G, K_H}(s, t)$ for $\text{OAEP}[G, H]$.

Remark 4.1 Since we mainly study IND-CPA security, for simplicity we define above the “no-redundancy” version of the OAEP, *i.e.*, corresponding to the “basic scheme” in [7]. However, all our results also holds for the redundant version. Additionally, as is typical in the literature we have defined OAEP to apply the G -function to the least-significant bits of the input; in standards and implementations it is typically the most significant bits (where the order of m and r are switched). Again, we stress that our results hold in either case.

4.2 Analysis

The following establishes that OAEP is encryption-compatible if the hash function G is t -wise independent for appropriate t . No restriction is put on the other hash function H . Indeed, our result also applies to SAEP [14] (although the latter is neither standardized nor known to provide CCA security in the RO model, except in certain cases).

Theorem 4.2 Let $G: \mathcal{K}_G \times \{0, 1\}^\rho \rightarrow \{0, 1\}^\mu$ and $H: \mathcal{K}_H \times \{0, 1\}^\mu \rightarrow \{0, 1\}^\rho$ be hash functions, and suppose G is t -wise independent. Let $\text{OAEP} = \text{OAEP}[G, H]$. Then

- (1) OAEP is (s, ε) -encryption-compatible where $\varepsilon = 2^{-u}$ for $u = \frac{t}{3t+2}(\rho - s - \log t + 2) - \frac{2(\mu+s)}{3t+2} - 1$.
- (2) OAEP is *regular* (s, ε) -encryption-compatible where $\varepsilon = 2^{-u}$ for $u = \frac{t}{2t+2}(\rho - s - \log t + 2) - \frac{\mu+s+2}{t+1} - 1$.
- (3) When $t = 2$, OAEP is (s, ε) -encryption-compatible where $\varepsilon = 2^{-u}$ for $u = (\rho - s - 2\mu)/4 - 1$.

Note that parts (2) and (3) capture special cases of (1) in which we get better bounds. The techniques used in the proof were first developed in the context of the classical LHL by Trevisan and Vadhan [50] and Dodis, Sahai and Smith [26], though the style of presentation of our theorem statement and proof are inspired by Barak *et al.* [3, Lemma 1]. We mention that due to our use of (variants of) the Crooked LHL rather than the classical one and the structure of OAEP, some of the technical details differ in our case and require new ideas.

Corollary 4.3 Let $G: \mathcal{K}_G \times \{0, 1\}^\rho \rightarrow \{0, 1\}^\mu$ and $H: \mathcal{K}_H \times \{0, 1\}^\mu \rightarrow \{0, 1\}^\rho$ be hash functions and suppose that G is t -wise independent for $t \geq 3\frac{\mu+s}{\rho-s}$. Then $\text{OAEP}[G, H]$ is (s, ε) -encryption-compatible where $\varepsilon = \exp(-c(\rho - s - \log t))$ for a constant $c > 0$.

In particular, $c \approx 1/2$ for regular functions. For such a function, if $\rho - s$ is at least 180 then ε is roughly 2^{-80} for $t = 10$ and message lengths $\mu \leq 2^{15}$ (which for practical purposes does not restrict the message-space). Applying Theorem 3.3, we see that if G is 10-wise independent and the number of random bits used in OAEP is at least 180 bits larger than the residual lossiness of the TDP, then the security of OAEP is tightly related to that of the lossy TDP.

Remark 4.4 To show security of OAEP against what we call *key-independent* chosen-plaintext attack, it suffices to argue that $\text{OAEP}[G, H]$ is a fooling extractor for any *fixed* encryption source $X = (m, R)$ where $m \in \{0, 1\}^\mu$. The latter holds for any $\varepsilon > 0$ and $s \leq \rho - 2\log(1/\varepsilon) + 2$ assuming G is only pairwise-independent (*i.e.*, $t = 2$). See Appendix B for details.

Proof: (of Theorem 4.2) We now prove the above theorem.

OVERVIEW. We write OAEP for $\text{OAEP}[G, H]$. The high-level idea for all three parts of the theorem is the same. Fix a lossy function f with range size at most 2^s . We first show that for every *fixed*

message $m \in \{0, 1\}^\mu$, with high probability (say $1 - \delta$) over the choice of K_G , the statistical distance between $f(\text{OAEP}(m, R))$ and $f(U)$ is small (say $\hat{\varepsilon}$). This aspect of the proof changes from part to part. We then take a union bound to show that the above holds for *all* messages over the same choice of K_G with probability at least $1 - 2^\mu \delta$. This means that the statistical distance between the pair $(K_G, f(\text{OAEP}(m, R)))$ and $(K_G, f(\text{OAEP}(U)))$ is at most $\varepsilon = \hat{\varepsilon} + 2^\mu \delta$ for all messages over the same choice of K_G . Finally, we express δ as a function of $\hat{\varepsilon}$, and select $\hat{\varepsilon}$ to minimize this sum. Note that the entire argument works for any choice of H .

We first prove part (3) of the theorem, then part (2), and finally part (1).

PROOF OF PART (3). To prove part (3) of the theorem, we strengthen the Crooked LHL of [27] to give the distinguisher access to the *input* to the fooling function as well its output.

Lemma 4.5 (Augmented Crooked LHL.) Let $h: \mathcal{K} \times A \rightarrow B$ be a pairwise-independent function and let $g: A \times B \rightarrow S$ be a function. Let X be a random variable on A such that $H_\infty(X) \geq \lg |S| + 2 \lg(1/\hat{\varepsilon}) - 2$ for some $\hat{\varepsilon} > 0$. Then

$$\Delta((K, g(X, h(K, X))), (K, g(X, U))) \leq \hat{\varepsilon},$$

where $K \xleftarrow{\$} \mathcal{K}$ and U is uniform and independent on B .

The proof, which extends the proof of the Crooked LHL given in [11], is in Appendix A.

We let G play the role of h in Lemma 4.5 and let $\{0, 1\}^\rho$ and $\{0, 1\}^\mu$ play the roles of A and B , respectively. Let g in the lemma be defined by $g(a, b) = f(m \oplus a \| b \oplus H(K_H, m \oplus a))$ for arbitrary but *fixed* $m \in \{0, 1\}^\mu$, $K_H \in \mathcal{K}_H$. It follows that OAEP is a $(s, \hat{\varepsilon})$ -fooling extractor for every *fixed* encryption source X of the form (m, R) . Part (3) of the theorem now follows by applying Markov's inequality and taking a union bound over *all* such sources; we omit the details.

PROOF OF PART (2). Instead of Markov's inequality, the proof of part (2) of the theorem uses a stronger tail inequality for t -wise independent random variables, due to Bellare and Rompel [8] (our application was inspired by the use of t -wise independence by Trevisan and Vadhan [50] and Dodis, Sahai and Smith [26]).

Let f be any function on $\{0, 1\}^k$ to a set $\mathcal{Y}$ of size at most 2^s . For this part of the theorem, assume that f is regular, that is, that each preimage set has size exactly 2^{k-s} . Let $X = (m, R)$ be any $(\mu + \rho, \rho)$ -source, where $m \in \{0, 1\}^\mu$ is fixed and R is uniform over $\{0, 1\}^\rho$. For each $r \in \{0, 1\}^\rho$ and $y \in \mathcal{Y}$, define the random variable

$$Z_{r,y} = \begin{cases} 2^{-\rho} & \text{if } f(\pi_{K_G, K_H}(m \| r)) = y, \\ 0 & \text{otherwise,} \end{cases}$$

where, here and in what follows, the probability is over the random choices of K_G and K_H (although the distribution on K_H does not matter – we use only the fact that it is independent of m, R, K_G). Let $Z_y = \sum_r Z_{r,y}$. We claim that $\mathbf{E}[Z_y] = 2^{-s}$. To see this, note that

$$\mathbf{E}[Z_y] = \sum_r 2^{-\rho} \cdot \Pr[f(U \| r) = y] = \Pr[f(U \| R) = y] = 2^{-s}$$

where we use the fact that R is uniform and f is regular.

To bound the deviation of Z_y from its mean, note that for a fixed y , the variables $\{Z_{r,y}\}_{r \in \{0,1\}^\rho}$ are t -wise independent (by the t -wise independence of G) and take values in $[0, 2^{-\rho}]$. We can apply the following tail bound (modified from the original to apply to random variables in $[0, M]$ rather than $[0, 1]$).

Lemma 4.6 (Bellare and Rompel [8]) Let $A_1, \dots, A_n$ be t -wise independent random variables taking values in $[0, M]$. Let $A = \sum_i A_i$ and $\delta \leq 1$. Then

$$\Pr[|A - \mathbf{E}[A]| \geq \delta \cdot \mathbf{E}[A]] \leq c_t \left(\frac{t \cdot M}{\delta^2 \cdot \mathbf{E}[A]} \right)^{t/2}$$

where $c_t < 3$ and $c_t < 1$ when $t \geq 8$.

Setting $\delta = 2\hat{\varepsilon}$, we get that for every $y \in \mathcal{Y}$,

$$\Pr[|Z_y - 2^{-s}| \geq 2\hat{\varepsilon} \cdot 2^{-s}] \leq c_t \left(\frac{t}{4\hat{\varepsilon}^2 \cdot 2^{-s+\rho}} \right)^{t/2}. \quad (5)$$

By a union bound, the probability that there exists a $y \in \mathcal{Y}$ such that $|Z_y - 2^{-s}| \geq 2\hat{\varepsilon} \cdot 2^{-s}$ is at most

$$2^s c_t \left(\frac{t}{4\hat{\varepsilon}^2 \cdot 2^{-s}} \right)^{t/2}.$$

Observe that if (4.2) holds for *all* $y \in \mathcal{Y}$ then, letting Y denote the random variable $f(\pi_{K_G, K_H}(m, R))$, we have

$$\Delta((K_G, K_H, Y), (K_G, K_H, f(U))) \leq \frac{1}{2} \sum_{y \in \mathcal{Y}} |Z_y - 2^{-s}| = \sum_{y \in \mathcal{Y}} \hat{\varepsilon} \cdot 2^{-s} = \hat{\varepsilon}.$$

By another union bound, the probability that the above holds simultaneously for all 2^μ possible $(\mu + \rho, \rho)$ -sources $X = (m, R)$ is at least $1 - \delta_{\hat{\varepsilon}}$, where

$$\delta_{\hat{\varepsilon}} = 2^{\mu+s} c_t \left(\frac{t}{4\hat{\varepsilon}^2 \cdot 2^{-s+\rho}} \right)^{t/2}. \quad (6)$$

Thus, OAEP is (s, ε) -encryption-compatible with $\varepsilon = \hat{\varepsilon} + \delta_{\hat{\varepsilon}}$. Note that $\delta_{\hat{\varepsilon}}$ can be written in the form $\gamma \cdot \hat{\varepsilon}^{-t}$ (where γ depends on t, ρ, s, μ but not $\hat{\varepsilon}$). Setting $\hat{\varepsilon} = \gamma^{1/(t+1)}$ yields $\varepsilon \leq 2\gamma^{1/(t+1)}$ and part (2) of the Theorem follows by observing that

$$\begin{aligned} u = -\log \varepsilon &\geq -\frac{1}{t+1} \cdot \log \gamma - 1 \\ &= -\frac{1}{t+1} \cdot \left(\frac{t}{2}(\rho - s - \log t + 2) + \mu + s + \log c_t \right) - 1 \\ &\geq \frac{t}{2t+2} \cdot (\rho - s - \log t + 2) - \frac{\mu + s + 2}{t+1} - 1. \end{aligned}$$

PROOF OF PART (1). We now turn to proving the lemma for general (not necessarily balanced) functions f . We first give a proof for approximately balanced functions, in which no pre-image set is too small; we then show that this implies a bound for arbitrary functions.

Assume for now that $\min_{y \in \mathcal{Y}} |f^{-1}(y)| \geq \lambda \cdot 2^{k-s}$ for some real number $0 < \lambda \leq 1$ (note that regularity corresponds to $\lambda = 1$). We sketch how to modify the proof of part (2) under this assumption; essentially, we end up with an extra factor of λ in the denominator of Equation 6. We use the same definition of Z_y as in part (2). Instead of $\mathbf{E}[Z_y] = 2^{-s}$, we now have $\mathbf{E}[Z_y] = \Pr[f(U||R) = y] = |f^{-1}(y)|/2^k$. Thus, instead of Equation (5), we have

$$\Pr \left[|Z_y - |f^{-1}(y)|/2^k| \geq 2\hat{\varepsilon} \cdot |f^{-1}(y)|/2^k \right] \leq c_t \left(\frac{t}{4\hat{\varepsilon}^2 \cdot |f^{-1}(y)|/2^k \cdot 2^\rho} \right)^{t/2}.$$

Using $\min_{y \in \mathcal{Y}} |f^{-1}(y)| \geq \lambda \cdot 2^{k-s}$ and taking a union bound, we get that the probability that there exists $y \in \mathcal{Y}$ such that

$$|Z_y - |f^{-1}(y)|/2^k| \geq 2\hat{\varepsilon} \cdot |f^{-1}(y)|/2^k \quad (7)$$

is at most

$$2^s c_t \left(\frac{t}{4\hat{\varepsilon}^2 \cdot \lambda \cdot 2^{-s} \cdot 2^\rho} \right)^{t/2}. \quad (8)$$

We can obtain a bound for arbitrary functions f by noting that every function f is “close” to a function with no small pre-images. Specifically:

Claim 4.7 Let $f: \{0,1\}^k \rightarrow \mathcal{Y}$ where $|\mathcal{Y}| \leq 2^s$ be a function. For any real number $\lambda > 0$, there exists a function $g: \{0,1\}^k \rightarrow \mathcal{Y}$ such that (i) $\min_{y \in \mathcal{Y}} |g^{-1}(y)| \geq \lambda \cdot 2^{k-s}$; and (ii) the function g agrees with f on a $1 - \lambda$ fraction of its domain. In particular, $\Delta(f(U), g(U)) \leq \lambda$.

We can now prove part (3) of the theorem from Equation (8) by choosing $\lambda = \hat{\varepsilon}$ in the claim and then completing the analysis as in part (2). It remains to prove the claim.

Proof (of Claim 4.7): The idea is that we will take all the small pre-image sets of f and merge them together with some larger preimage set (e.g., if 0 had a large pre-image set, then for all elements x such that $f^{-1}(f(x))$ is small, we set $f(x) = 0$). How many elements can belong to small pre-image sets? There are at most 2^s pre-image sets, each of which contains at most $\lambda \cdot 2^{k-s}$ elements. So there are at most $\lambda \cdot 2^k$ elements of the domain on which f has to be changed. ■■

5 Lossiness of RSA

In this section, we show that the RSA trapdoor permutation is lossy under reasonable assumptions. In particular, we show that, for large enough encryption exponent e , RSA is considerably lossy under the Φ -Hiding Assumption of [17]. We then show that by generalizing this assumption to multi-prime RSA we can get even more lossiness. Finally, we propose a “Two-Or- m -Primes” Assumption that, when combined with the former, amplifies the lossiness of standard (two-prime) RSA for small e .

5.1 Background on RSA and Notation

We denote by $\mathcal{RSA}_k$ the set of all tuples (N, p, q) such that $N = pq$ is the product of two distinct $k/2$ -bit primes. Such an N is called an *RSA modulus*. By $(N, p, q) \xleftarrow{\$} \mathcal{RSA}_k$ we mean that (N, p, q)

is sampled according to the uniform distribution on $\mathcal{RSA}_k$. An *RSA TDP generator* [47] is an algorithm $\mathcal{F}$ that returns $(N, e), (N, d)$, where N is an RSA modulus and $ed \equiv 1 \pmod{\phi(N)}$. (Here $\phi(\cdot)$ denotes Euler's totient function, so in particular $\phi(N) = (p-1)(q-1)$.) The tuple (N, e) defines the permutation on $\mathbb{Z}_N^*$ given by $f(x) = x^e \bmod N$, and similarly (N, d) defines its inverse. We say that a lossy TDP generator $\text{LTDP} = (\mathcal{F}, \mathcal{F}')$ is an RSA LTDP if $\mathcal{F}$ is an RSA TDP generator.

To define the Φ -Hiding Assumption and later some extensions of it, the following notation is also useful. For $i \in \mathbb{N}$ we denote by $\mathcal{P}_i$ the set of all i -bit primes. Let R be a relation on p and q . By $\mathcal{RSA}_k[R]$ we denote the subset of $\mathcal{RSA}_k$ for that the relation R holds on p and q . For example, let e be a prime. Then $\mathcal{RSA}_k[p = 1 \bmod e]$ is the set of all (N, p, q) , where $N = pq$ is the product of two distinct $k/2$ -bit primes p, q and $p = 1 \bmod e$. That is, the relation $R(p, q)$ is true if $p = 1 \bmod e$ and q is arbitrary. By $(N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[R]$ we mean that (N, p, q) is sampled according to the uniform distribution on $\mathcal{RSA}_k[R]$.

5.2 RSA Lossy TDP from Φ -Hiding

Φ -HIDING ASSUMPTION (ΦA). We recall the Φ -Hiding Assumption of [17]. For an RSA modulus N , we say that N *ϕ -hides* a prime e if $e \mid \phi(N)$. Intuitively, the assumption is that, given RSA modulus N , it is hard to distinguish primes which are ϕ -hidden by N from those that are not. Formally, let $0 < c < 1/2$ be a (public) constant determined later. Consider the following two distributions:

$$\begin{aligned} \mathcal{R}_1 &= \{(e, N) : e, e' \xleftarrow{\$} \mathcal{P}_{ck} ; (N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[p = 1 \bmod e']\} \\ \mathcal{L}_1 &= \{(e, N) : e \xleftarrow{\$} \mathcal{P}_{ck} ; (N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[p = 1 \bmod e]\} . \end{aligned}$$

To a distinguisher D we associate its ΦA *advantage* defined as

$$\mathbf{Adv}_{c,D}^{\Phi A}(k) = \Pr [D(\mathcal{R}_1) \Rightarrow 1] - \Pr [D(\mathcal{L}_1) \Rightarrow 1] .$$

As shown in [17], distributions $\mathcal{R}_1, \mathcal{L}_1$ can be sampled efficiently assuming the widely-accepted Extended Riemann Hypothesis.⁸

RSA LTDP FROM ΦA . We construct an RSA LTDP based on ΦA . In injective mode the public key is (N, e) where e is not ϕ -hidden by N , whereas in lossy mode it is. Namely, define $\text{LTDP}_1 = (\mathcal{F}_1, \mathcal{F}'_1)$ as follows:

Algorithm $\mathcal{F}_1$

$e, e' \xleftarrow{\$} \mathcal{P}_{ck}$
 $(N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[p = 1 \bmod e']$
 If $\gcd(e, \phi(N)) \neq 1$ then return $\perp$
 $d \leftarrow e^{-1} \bmod \phi(N)$
 Return $((N, e), (N, d))$

Algorithm $\mathcal{F}'_1$

$e \xleftarrow{\$} \mathcal{P}_{ck}$
 $(N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[p = 1 \bmod e]$
 Return (N, e)

The fact that algorithm $\mathcal{F}_1$ has only a negligible probability of failure (returning $\perp$) follows from the fact that $\phi(N)$ can have only a constant number of prime factors of length ck and Bertrand's Postulate.

Proposition 5.1 Suppose there is a distinguisher D against LTDP_1 . Then there is a distinguisher D' such that for all $k \in \mathbb{N}$

$$\mathbf{Adv}_{\text{LTDP}_1, D}^{\text{ldp}}(k) \leq 2 \cdot \mathbf{Adv}_{c, D'}^{\Phi A}(k) .$$

⁸This is done by choosing a uniform $(1/2 - c)k$ -bit number x until $p = xe + 1$ is a prime.

Furthermore, the running-time of D' is that of D . LTDP_1 has lossiness ck .

From a practical perspective, a drawback of LTDP_1 is that $\mathcal{F}_1$ chooses $N = pq$ in a non-standard way, so that it hides a prime of the same length as e . Moreover, for small values of e it returns $\perp$ with high probability. This is done for consistency with how [17] formulated ΦA . But, to address this, we also propose what we call the *Enhanced ΦA* ($\text{E}\Phi\text{A}$), which says that N generated in the non-standard way (*i.e.*, by $\mathcal{F}_1$) is indistinguishable from one chosen at random subject to $\gcd(e, \phi(N)) = 1$.⁹ We conjecture that $\text{E}\Phi\text{A}$ holds for all values of c that ΦA does. Details follow.

ENHANCED Φ -HIDING ASSUMPTION. We say that the *Enhanced Φ -Hiding Assumption* ($\text{E}\Phi\text{A}$) holds for c if the following two distributions $\mathcal{R}_{1*}$ and $\mathcal{L}_{1*}$ are computationally indistinguishable:

$$\begin{aligned}\mathcal{R}_{1*} &= \{(e, N) : e \xleftarrow{\$} \mathcal{P}_{ck} ; (N, p, q) \xleftarrow{\$} \mathcal{RSA}_k\} \\ \mathcal{L}_{1*} &= \{(e, N) : e \xleftarrow{\$} \mathcal{P}_{ck} ; (N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[p = 1 \bmod e]\} .\end{aligned}$$

To a distinguisher D we associate its *E ΦA advantage* defined as

$$\mathbf{Adv}_{c,D}^{\text{E}\Phi\text{A}}(k) = \Pr[D(\mathcal{R}_{1*}) \Rightarrow 1] - \Pr[D(\mathcal{L}_{1*}) \Rightarrow 1] .$$

As before, distributions $\mathcal{R}_{1*}, \mathcal{L}_{1*}$ can be sampled efficiently assuming the widely-accepted Extended Riemann Hypothesis. We conjecture that $\text{E}\Phi\text{A}$ holds for all values of $\mathcal{K}_\phi, c$ that ΦA does.

RSA LTDP FROM $\text{E}\Phi\text{A}$. Now define $\text{LTDP}_{1*} = (\mathcal{F}_{1*}, \mathcal{F}'_{1*})$ where

Algorithm $\mathcal{F}_{1*}$

$e \xleftarrow{\$} \mathcal{P}_{ck}$
 $(N, p, q) \xleftarrow{\$} \mathcal{RSA}_k$
 If $\gcd(e, \phi(N)) \neq 1$ then Return $\perp$
 Else Return $(N, e), (N, d)$

and $\mathcal{F}'_{1*} = \mathcal{F}'_1$ in Section 5.2. Again we have the probability that $\mathcal{F}_{1*}$ returns $\perp$ is negligible. We stress that $\mathcal{F}_{1*}$, unlike $\mathcal{F}_1$, chooses p, q at random as is typical in practice. We have the following proposition.

Proposition 5.2 If the Enhanced Φ -Hiding Assumption holds for c then $\text{LTDP}_{1*} = (\mathcal{F}_{1*}, \mathcal{F}'_{1*})$ is an RSA LTDP with lossiness ck . In particular, suppose there is a distinguisher D against LTDP_{1*} . Then there is a distinguisher D' such that

$$\mathbf{Adv}_{\text{LTDP}_{1*}, D}^{\text{ltdp}}(k) \leq 2 \cdot \mathbf{Adv}_{c, D'}^{\text{E}\Phi\text{A}}(k) .$$

Furthermore, the running-time of D' is that of D .

PARAMETERS FOR LTDP_1 . When e is too large, ΦA can be broken by using Coppersmith's method for finding small roots of a univariate modulo an unknown divisor of N [22, 40]. (No other attack on ΦA here is known.) Namely, consider the polynomial $r(x) = ex + 1 \bmod p$. Coppersmith's method allows us to find all roots of r smaller than $N^{1/4}$, and thus factor N , in lossy mode in polynomial time if $c \geq 1/4$. (This is essentially the “factoring with high bits known” attack.) More specifically, applying [40, Theorem 1], N can be factored in time $\text{poly}(\log N)$ and $O(N^\varepsilon)$ if $c = 1/4 - \varepsilon$ (*i.e.*, $\log e \geq \log N(1/4 - \varepsilon)$). For example, with modulus size $k = 2048$ we can set $\varepsilon = .04$ for 80-bit security (to enforce $k\varepsilon \geq 80$) and obtain $2048 \cdot (1/4 - 0.04) = 430$ bits of lossiness.

⁹Additionally, in practice the encryption exponent e is usually fixed. This can be addressed by parameterizing $\text{E}\Phi\text{A}$ by a fixed e instead of choosing it at random. Note that for $e = 3$ one should make both $e \mid p - 1$ and $e \mid q - 1$ in the lossy case (otherwise the assumption is false; cf. [17, Remark 2, p. 6]).

5.3 RSA Lossy TDP from Multi-Prime Φ -Hiding

Multi-prime RSA (according to [39] the earliest reference is [46]) is a generalization of RSA to moduli $N = p_1 \cdots p_m$ of length k with $m \geq 2$ prime factors of equal bit-length. Multi-prime RSA is of interest to practitioners since it allows to speed up decryption and is included in RSA PKCS #1 v2.1. We are interested in it here because for it we can show greater lossiness and even with smaller encryption exponent e .

NOTATION AND TERMINOLOGY. Let $m \geq 2$ be fixed. We denote by $\mathcal{MRSA}_k$ the set of all tuples $(N, p_1, \dots, p_m)$, where $N = p_1 \cdots p_m$ is the product of distinct k/m -bit primes. Such an N is called an *m-prime RSA modulus*. By $(N, p_1, \dots, p_m) \xleftarrow{\$} \mathcal{MRSA}_k$ we mean that $(N, p_1, \dots, p_m)$ is sampled according to the uniform distribution on $\mathcal{MRSA}_k$. The rest of the notation and terminology of Section 5 is extended to the multi-prime setting in the obvious way.

MULTI Φ -HIDING ASSUMPTION. For an m -prime RSA modulus N , let us say that N *m ϕ -hides* a prime e if $e \mid p_i - 1$ for all $1 \leq i \leq m - 1$. Intuitively, the assumption is that, given such N , it is hard to distinguish primes which are *m ϕ -hidden* by N from those that do not divide $p_i - 1$ for any $1 \leq i \leq m$. Formally, let $m = m(k) \geq 2$ be a polynomial and let $c = c(k)$ be an inverse polynomial determined later. Consider the following two distributions:

$$\begin{aligned} \mathcal{R}_2 &= \{(e, N) : e, e' \xleftarrow{\$} \mathcal{P}_{ck} ; (N, p_1, \dots, p_t) \xleftarrow{\$} \mathcal{MRSA}_k [p_{i \leq m-1} = 1 \bmod e']\} \\ \mathcal{L}_2 &= \{(e, N) : e \xleftarrow{\$} \mathcal{P}_{ck} ; (N, p_1, \dots, p_t) \xleftarrow{\$} \mathcal{MRSA}_k [p_{i \leq m-1} = 1 \bmod e]\} . \end{aligned}$$

Above and in what follows, by $p_{i \leq m-1} = 1 \bmod e$ we mean that $p_i = 1 \bmod e$ for all $1 \leq i \leq m - 1$. To a distinguisher D we associate its *M Φ A advantage* defined as

$$\text{Adv}_{m,c,D}^{\text{M}\Phi\text{A}}(k) = \Pr [D(\mathcal{R}_2) \Rightarrow 1] - \Pr [D(\mathcal{L}_2) \Rightarrow 1] .$$

As before, distributions $\mathcal{R}_2, \mathcal{L}_2$ can be sampled efficiently assuming the widely-accepted Extended Riemann Hypothesis.

Note that if we had required that in the lossy case $N = p_1 \cdots p_m$ is such that $e \mid p_i$ for all $1 \leq i \leq m$, then in this case we would always have $N = 1 \bmod e$. But in the injective case $N \bmod e$ is random, which would lead to a trivial distinguishing algorithm. This explains why we do not impose $e \mid p_m$ in the lossy case above.

MULTI-PRIME RSA LTDP FROM M Φ A. We construct a multi-prime RSA LTDP based on M Φ A having lossiness $(m - 1) \log e$, where in lossy mode N *m ϕ -hides* e . Namely, define $\text{LTDP}_2 = (\mathcal{F}_2, \mathcal{F}'_2)$ as follows:

Algorithm $\mathcal{F}_2$

```

 $e, e' \xleftarrow{\$} \mathcal{P}_{ck}$ 
 $(N, p_1, \dots, p_m)$ 
 $\xleftarrow{\$} \mathcal{MRSA}_k [p_{i \leq m-1} = 1 \bmod e']$ 
If  $\gcd(e, \phi(N)) \neq 1$  then Return  $\perp$ 
 $d \leftarrow e^{-1} \bmod \phi(N)$ 
Else return  $(N, e), (N, d)$ 
```

Algorithm $\mathcal{F}'_2$

```

 $e \xleftarrow{\$} \mathcal{P}_{ck}$ 
 $(N, p_1, \dots, p_m)$ 
 $\xleftarrow{\$} \mathcal{MRSA}_k [p_{i \leq m-1} = 1 \bmod e]$ 
Return  $(N, e)$ 
```

Proposition 5.3 Suppose there is a distinguisher D against LTDP_2 . Then there is a distinguisher D' such that for all $k \in \mathbb{N}$

$$\text{Adv}_{\text{LTDP}_2, D}^{\text{ltdp}}(k) \leq 2 \cdot \text{Adv}_{m,c,D'}^{\text{M}\Phi\text{A}}(k) .$$

Furthermore, the running-time of D' is that of D . LTDP_2 has lossiness $(m-1)ck$.

PARAMETERS FOR LTDP_2 . We use the recent cryptanalysis of the $\text{M}\Phi\text{A}$ for $m \geq 3$ due to Herrmann [32].¹⁰ Using [32, Section 3] we can break the $\text{M}\Phi\text{A}$ in time $\text{poly}(\log N)$ and $O(N^\varepsilon)$ if

$$c \geq 1/m - \frac{2}{3\sqrt{m^3}} - \varepsilon.$$

(For $m \geq 3$ this improves the bound with $c \geq 1/m - 1/m^2 - \varepsilon$ obtained from “factoring with high bits known”; for $m \geq 4$ this improves the bound with $c \geq 1/m - 2 \frac{(1/m)^{(1/(m-1))} - (1/m)^{m/(m-1)}}{m(m-1)} - \varepsilon$ from the preliminary version [36].)

For example, with modulus size $k = 2048$ and $m = 3$ ($m = 4, 5$) we set $\varepsilon = .04$ (for about 80-bit security) and obtain 676 (778, 822) bits of lossiness for LTDP_2 , according to Proposition 5.3.

We note that this may not be the best attack possible based on Coppersmith’s method (in particular the coefficients of the polynomial we use are highly correlated). We also remark that for a *fixed* modulus length, m cannot be too large since the Elliptic Curve Method for factoring can compute a factor p_i of N faster than the Number Field Sieve one if p_i is significantly smaller than $N^{1/2}$ [39].

5.4 Small-Exponent RSA LTDP from 2-vs- m Primes

For efficiency reasons, the public RSA exponent e is typically not chosen to be too large in practice. (For example, researchers at UC San Diego [51] found that 99.5% of the certificates in the campus’s TLS corpus had $e = 2^{16} + 1$.) Therefore, we investigate the possibility of using an additional assumption to amplify the lossiness of RSA for small e .

The high-level idea is to assume that it is hard to distinguish $N = pq$ where p, q are primes of length $k/2$ from $N = p_1 \cdots p_m$ for $m > 2$, where $p_1, \dots, p_m$ are primes of length k/m (which we call the “2-vs- m Primes” Assumption). This assumption is a generalization of the “2-vs-3 Primes” Assumptions introduced in [9] and also used contemporaneously to our work to construct a “slightly lossy” TDF based on modular squaring [42]. Combined with the $\text{M}\Phi\text{A}$ Assumption of Section 5.3, we obtain $(m-1) \log e$ bits of lossiness from *standard* (two-prime) RSA. Let us state our assumption and construction formally.

2-VS- m PRIMES ASSUMPTION. We say that the *2-vs- m primes assumption* holds for m if the following two distributions $\mathcal{N}_2$ and $\mathcal{N}_m$ are computationally indistinguishable:

$$\begin{aligned} \mathcal{N}_2 &= \{N : e \stackrel{\$}{\leftarrow} \mathcal{P}_{ck} ; (N, p, q) \stackrel{\$}{\leftarrow} \mathcal{RSA}_k[p \equiv 1 \pmod{e}]\} \\ \mathcal{N}_m &= \{N : e \stackrel{\$}{\leftarrow} \mathcal{P}_{ck} ; (N, p, q) \stackrel{\$}{\leftarrow} \mathcal{MRSA}_k[p_{i \leq m-1} \equiv 1 \pmod{e}]\}. \end{aligned}$$

To a distinguisher D we associate its *HFA-advantage* defined as

$$\mathbf{Adv}_m^{2\text{vmp}}(D) = \Pr[D(\mathcal{N}_2) \Rightarrow 1] - \Pr[D(\mathcal{N}_m) \Rightarrow 1].$$

RSA LTDP FROM 2-VS- m PRIMES + $\text{M}\Phi\text{A}$. Define $\text{LTDP}_3 = (\mathcal{F}_3, \mathcal{F}'_3)$ as follows:

¹⁰In the preliminary version of this paper we gave a weaker cryptanalysis which was subsequently improved in [32] for the case $m \geq 4$.

Algorithm $\mathcal{F}_3$

$e, e' \xleftarrow{\$} \mathcal{P}_{ck}$
 $(N, p, q) \xleftarrow{\$} \mathcal{RSA}_k[p = 1 \bmod e']$
 If $\gcd(e, \phi(N)) \neq 1$ then Return $\perp$
 Else Return $(N, e), (N, d)$

Algorithm $\mathcal{F}'_3$

$e \xleftarrow{\$} \mathcal{P}_{ck}$
 $(N, p_1, \dots, p_m) \xleftarrow{\$} \mathcal{MRSA}_k[p_{i \leq m-1} = 1 \bmod e]$
 Return (N, e)

Proposition 5.4 If the 2-vs- m Primes Assumption holds for m and the Multi-Prime Φ -Hiding Assumption holds for m, e , then $\text{LTDP}_3 = (\mathcal{F}_3, \mathcal{F}'_3)$ is an RSA LTDP with lossiness $(m-1)ck$. In particular, suppose there is a distinguisher D against LTDP_3 . Then there is a distinguisher D_1, D_2 such that

$$\text{Adv}_{\text{LTDP}_3}^{\text{ldp}}(D) \leq 2 \cdot (\text{Adv}_m^{2\text{vmp}}(D_1) + \text{Adv}_{m,c}^{\text{M}\Phi\text{A}}(D_2)) .$$

Furthermore, the running-time of D_1, D_2 is that of D .

The proof is a standard hybrid argument.

PARAMETERS FOR LTDP_3 . We note that m in the construction cannot be too large, otherwise a small factor of N in the lossy case can be recovered by the elliptic curve factoring method due to Lenstra [39], whose running-time is proportional to the smallest factor of N . The largest factor recovered by the method so far was 223-bits in length [52]. Thus, for example using 2048-bit RSA with $e = 2^{16} - 1$, if we assume it is hard to recover factors larger than that we can get about $8 \cdot 16 = 128$ bits of provable lossiness under the HFA plus $\text{M}\Phi\text{A}$ where $m = 9$.

ENHANCED HFA. As in the previous cases, to address the fact that in practice $N = pq$ is chosen at random and not subject to p hiding a prime of the same bit-length as e , we may define an *enhanced* version of HFA. Then under the enhanced HFA + enhanced $\text{M}\Phi\text{A}$ assumptions we obtain the same amount of lossiness for standard 2-prime RSA.

6 Instantiating RSA-OAEP

By combining the results of Section 3, Section 4, and Section 5, we obtain standard model instantiations of RSA-OAEP under chosen-plaintext attack.

REGULARITY. In particular, we would like to apply part (2) of Theorem 4.2 in this case, as it is not hard to see that under all of the assumptions discussed in Section 5, RSA is a *regular* lossy TDP on the domain $\mathbb{Z}_N^*$. Unfortunately, this domain is different from $\{0, 1\}^{\rho+\mu}$ (identified as integers), the range of OAEP. In RSA PKCS #1 v2.1, the mismatch is handled by selecting $\rho + \mu = \lfloor \log N \rfloor - 16$, and viewing OAEP's output as an integer less than $2^{\rho+\mu} < N/2^{16}$ (i.e., the most significant two bytes of the output are zeroed out). The problem is that in the lossy case RSA may not be regular on the subdomain $\{0, \dots, 2^{\rho+\mu} - 16\}$.

We can prove, in some cases, that in the lossy case RSA is *approximately* regular on this subdomain, and in those cases it follows from the proof of part (1) of Theorem 4.2 that we obtain essentially the better parameters given by part (2). However, here we just use the weaker parameters given by part (1) of Theorem 4.2. We leave a detailed discussion of approximate regularity to future work. In particular, understanding the regularity of RSA on subintervals of the domain is a first step towards improving the concrete parameters we obtain.

CONCRETE PARAMETERS. Since the results in Section 5 have several cases and the parameter settings are rather involved, we avoid stating an explicit theorem about RSA-OAEP. From part (1)

of Theorem 4.2 one can see that for $u = 80$ bits security, messages of roughly $\mu \approx k - s - 3 \cdot 80$ bits can be encrypted (for sufficiently large t). For concreteness, we give two example parameter settings. Using the Multi Φ -Hiding Assumption with $N = 1024$ bits and 3 primes, we obtain $\ell = k - s = 291$ bits of lossiness and hence can encrypt messages of length $\mu = 40$ bits (for $t \approx 400$); using the Φ -Hiding Assumption with $N = 2048$, we obtain $\ell = k - s = 430$ bits of lossiness and hence can encrypt messages of length $\mu = 160$ bits (for $t \approx 150$). We stress that while we view our results as providing important theoretical backing for the scheme at a more qualitative level, we strongly encourage further research to try to improve the concrete parameters.

Acknowledgements

We thank Mihir Bellare, Alexandra Boldyreva, Dan Brown, Yevgeniy Dodis, Jason Hinek, Arjen Lenstra, Alex May, Phil Rogaway, and the anonymous reviewers of Crypto 2010 for helpful comments. In particular, we thank Dan for reminding us of [17, Remark 2, p. 6], Alex for pointing out the improved attack in Section 5.3, Phil for encouraging us to consider the case of small e more closely and Yevgeniy for suggesting the statement of Lemma 4.5 (our original lemma was specific to OAEP). This work was primarily done while A.O. was a Ph.D. student at Georgia Institute of Technology, supported in part by Alexandra Boldyreva’s NSF CAREER award #0545659 and NSF Cyber Trust award #0831184. A.S. was supported in part by NSF awards #0747294, 0729171.

References

- [1] Rsa public-key cryptography standards (pkcs). <http://www.rsa.com/rsalabs/node.asp?id=2124>. (Cited on page 3, 9.)
- [2] Michel Abdalla, Mihir Bellare, and Phillip Rogaway. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In David Naccache, editor, *Topics in Cryptology – CT-RSA 2001*, volume 2020 of *Lecture Notes in Computer Science*, pages 143–158, San Francisco, CA, USA, April 8–12, 2001. Springer, Berlin, Germany. (Cited on page 7.)
- [3] Boaz Barak, Ronen Shaltiel, and Eran Tromer. True random number generators secure in a changing environment. In Colin D. Walter, Çetin Kaya Koç, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems – CHES 2003*, volume 2779 of *Lecture Notes in Computer Science*, pages 166–180, Cologne, Germany, September 8–10, 2003. Springer, Berlin, Germany. (Cited on page 12.)
- [4] Mihir Bellare, Alexandra Boldyreva, and Adam O’Neill. Deterministic and efficiently searchable encryption. In Alfred Menezes, editor, *Advances in Cryptology – CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 535–552, Santa Barbara, CA, USA, August 19–23, 2007. Springer, Berlin, Germany. (Cited on page 27.)
- [5] Mihir Bellare and Adriana Palacio. Towards plaintext-aware public-key encryption without random oracles. In Pil Joong Lee, editor, *Advances in Cryptology – ASIACRYPT 2004*, volume 3329 of *Lecture Notes in Computer Science*, pages 48–62, Jeju Island, Korea, December 5–9, 2004. Springer, Berlin, Germany. (Cited on page 5.)
- [6] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press. (Cited on page 3, 6, 7.)

- [7] Mihir Bellare and Phillip Rogaway. Optimal asymmetric encryption. In Alfredo De Santis, editor, *Advances in Cryptology – EUROCRYPT’94*, volume 950 of *Lecture Notes in Computer Science*, pages 92–111, Perugia, Italy, May 9–12, 1994. Springer, Berlin, Germany. (Cited on page 3, 6, 8, 9, 11, 12.)
- [8] Mihir Bellare and John Rompel. Randomness-efficient oblivious sampling. In *35th Annual Symposium on Foundations of Computer Science*, pages 276–287, Santa Fe, New Mexico, November 20–22, 1994. IEEE Computer Society Press. (Cited on page 13, 14.)
- [9] Manuel Blum, Paul Feldman, and Silvio Micali. Proving security against chosen cyphertext attacks. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO’88*, volume 403 of *Lecture Notes in Computer Science*, pages 256–268, Santa Barbara, CA, USA, August 21–25, 1990. Springer, Berlin, Germany. (Cited on page 19.)
- [10] Alexandra Boldyreva, David Cash, Marc Fischlin, and Bogdan Warinschi. Foundations of non-malleable hash and one-way functions. In Mitsuru Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 524–541, Tokyo, Japan, December 6–10, 2009. Springer, Berlin, Germany. (Cited on page 7.)
- [11] Alexandra Boldyreva, Serge Fehr, and Adam O’Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 335–359, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Berlin, Germany. (Cited on page 4, 9, 13.)
- [12] Alexandra Boldyreva and Marc Fischlin. Analysis of random oracle instantiation scenarios for OAEP and other practical schemes. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 412–429, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Berlin, Germany. (Cited on page 6.)
- [13] Alexandra Boldyreva and Marc Fischlin. On the security of OAEP. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology – ASIACRYPT 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 210–225, Shanghai, China, December 3–7, 2006. Springer, Berlin, Germany. (Cited on page 6.)
- [14] Dan Boneh. Simplified OAEP for the RSA and Rabin functions. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 275–291, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Berlin, Germany. (Cited on page 4, 9, 12.)
- [15] Daniel R. L. Brown. What hashes make rsa-oaep secure? Cryptology ePrint Archive, Report 2006/223, 2006. <http://eprint.iacr.org/>. (Cited on page 6.)
- [16] Christian Cachin. Efficient private bidding and auctions with an oblivious third party. In *ACM CCS 99: 6th Conference on Computer and Communications Security*, pages 120–127, Kent Ridge Digital Labs, Singapore, November 1–4, 1999. ACM Press. (Cited on page 5.)
- [17] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In Jacques Stern, editor, *Advances in Cryptology – EUROCRYPT’99*, volume 1592 of *Lecture Notes in Computer Science*, pages 402–414, Prague, Czech Republic, May 2–6, 1999. Springer, Berlin, Germany. (Cited on page 5, 15, 16, 17, 21.)
- [18] Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In Burton S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO’97*, volume 1294 of *Lecture Notes in Computer Science*, pages 455–469, Santa Barbara, CA, USA, August 17–21, 1997. Springer, Berlin, Germany. (Cited on page 6.)
- [19] Ran Canetti and Ronny Ramzi Dakdouk. Extractable perfectly one-way functions. In Luca

- Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008: 35th International Colloquium on Automata, Languages and Programming, Part II*, volume 5126 of *Lecture Notes in Computer Science*, pages 449–460, Reykjavik, Iceland, July 7–11, 2008. Springer, Berlin, Germany. (Cited on page 6, 7.)
- [20] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *Journal of the ACM*, 51(4):557–594, 2004. (Cited on page 3.)
 - [21] Ran Canetti, Daniele Micciancio, and Omer Reingold. Perfectly one-way probabilistic hash functions (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 131–140, Dallas, Texas, USA, May 23–26, 1998. ACM Press. (Cited on page 6.)
 - [22] Don Coppersmith. Small solutions to polynomial equations, and low exponent RSA vulnerabilities. *Journal of Cryptology*, 10(4):233–260, 1997. (Cited on page 17.)
 - [23] Jean-Sébastien Coron, Marc Joye, David Naccache, and Pascal Paillier. New attacks on PKCS#1 v1.5 encryption. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 369–381, Bruges, Belgium, May 14–18, 2000. Springer, Berlin, Germany. (Cited on page 5.)
 - [24] Jean-Sébastien Coron, Marc Joye, David Naccache, and Pascal Paillier. Universal padding schemes for RSA. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 226–241, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Berlin, Germany. (Cited on page 4, 9.)
 - [25] Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak. On the generic insecurity of the full domain hash. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 449–466, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Berlin, Germany. (Cited on page 6.)
 - [26] Yevgeniy Dodis, Amit Sahai, and Adam Smith. On perfect and adaptive security in exposure-resilient cryptography. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 301–324, Innsbruck, Austria, May 6–10, 2001. Springer, Berlin, Germany. (Cited on page 12, 13.)
 - [27] Yevgeniy Dodis and Adam Smith. Correcting errors without leaking partial information. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 654–663, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press. (Cited on page 4, 9, 13.)
 - [28] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP is secure under the RSA assumption. *Journal of Cryptology*, 17(2):81–104, March 2004. (Cited on page 6.)
 - [29] Craig Gentry, Philip D. Mackenzie, and Zulfikar Ramzan. Password authenticated key exchange using hidden smooth subgroups. In Vijayalakshmi Atluri, Catherine Meadows, and Ari Juels, editors, *ACM CCS 05: 12th Conference on Computer and Communications Security*, pages 299–309, Alexandria, Virginia, USA, November 7–11, 2005. ACM Press. (Cited on page 5.)
 - [30] Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004. (Cited on page 27.)
 - [31] Brett Hemenway and Rafail Ostrovsky. Public-key locally-decodable codes. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 126–143, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Berlin, Germany. (Cited on page 5.)
 - [32] Mathias Herrmann. Improved cryptanalysis of the multi-prime ϕ -hiding assumption. In Abderahmane Nitaj and David Pointcheval, editors, *AFRICACRYPT 11: 4th International Conference on Cryptology in Africa*, volume 6737 of *Lecture Notes in Computer Science*, pages 92–99, Dakar,

- Senegal, July 5–7, 2011. Springer, Berlin, Germany. (Cited on page 19.)
- [33] Dennis Hofheinz and Eike Kiltz. The group of signed quadratic residues and applications. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 637–653, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Berlin, Germany. (Cited on page 7.)
 - [34] E. Kiltz and K. Pietrzak. Personal communication, 2009. (Cited on page 6.)
 - [35] Eike Kiltz, Payman Mohassel, and Adam O’Neill. Adaptive trapdoor functions and chosen-ciphertext security. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 673–692, French Riviera, May 30 – June 3, 2010. Springer, Berlin, Germany. (Cited on page 6.)
 - [36] Eike Kiltz, Adam O’Neill, and Adam Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 295–313, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Berlin, Germany. (Cited on page 19.)
 - [37] Eike Kiltz and Krzysztof Pietrzak. On the security of padding-based encryption schemes - or - why we cannot prove OAEP secure in the standard model. In Antoine Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 389–406, Cologne, Germany, April 26–30, 2009. Springer, Berlin, Germany. (Cited on page 3, 5, 6, 8.)
 - [38] Kazukuni Kobara and Hideki Imai. Oaep++ : A very simple way to apply oaep to deterministic ow-cpa primitives. Cryptology ePrint Archive, Report 2002/130, 2002. <http://eprint.iacr.org/>. (Cited on page 6.)
 - [39] Arjen K. Lenstra. Unbelievable security. matching AES security using public key systems (invited talk). In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 67–86, Gold Coast, Australia, December 9–13, 2001. Springer, Berlin, Germany. (Cited on page 18, 19, 20.)
 - [40] Alexander May. Using lll-reduction for solving rsa and factorization problems: A survey. In *LLL+25 Conference in honour of the 25th birthday of the LLL algorithm*, 2007. (Cited on page 17.)
 - [41] Silvio Micali, Charles Rackoff, and Bob Sloan. The notion of security for probabilistic cryptosystems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, pages 381–392, Santa Barbara, CA, USA, August 1987. Springer, Berlin, Germany. (Cited on page 27.)
 - [42] Petros Mol and Scott Yilek. Chosen-ciphertext security from slightly lossy trapdoor functions. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010: 13th International Conference on Theory and Practice of Public Key Cryptography*, volume 6056 of *Lecture Notes in Computer Science*, pages 296–311, Paris, France, May 26–28, 2010. Springer, Berlin, Germany. (Cited on page 19.)
 - [43] Pascal Paillier and Jorge L. Villar. Trading one-wayness against chosen-ciphertext security in factoring-based encryption. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology – ASIACRYPT 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 252–266, Shanghai, China, December 3–7, 2006. Springer, Berlin, Germany. (Cited on page 6.)
 - [44] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive one-way functions and applications. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 57–74, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Berlin, Germany. (Cited on page 7.)
 - [45] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E.

- Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 187–196, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press. (Cited on page 3, 4, 8.)
- [46] Ronald L. Rivest, Adi Shamir, and Len Adelman. U.S. patent 4,405,829: Cryptographic communications system and method. (Cited on page 18.)
 - [47] Ronald L. Rivest, Adi Shamir, and Len Adelman. A method for obtaining public-key cryptosystems and digital signatures. Technical Memo MIT/LCS/TM-82, Massachusetts Institute of Technology, Laboratory for Computer Science, 1977. (Cited on page 16.)
 - [48] Christian Schridde and Bernd Freisleben. On the validity of the phi-hiding assumption in cryptographic protocols. In Josef Pieprzyk, editor, *Advances in Cryptology – ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 344–354, Melbourne, Australia, December 7–11, 2008. Springer, Berlin, Germany. (Cited on page 5.)
 - [49] Victor Shoup. OAEP reconsidered. *Journal of Cryptology*, 15(4):223–249, 2002. (Cited on page 6.)
 - [50] Luca Trevisan and Salil P. Vadhan. Extracting randomness from samplable distributions. In *41st Annual Symposium on Foundations of Computer Science*, pages 32–42, Redondo Beach, California, USA, November 12–14, 2000. IEEE Computer Society Press. (Cited on page 4, 12, 13.)
 - [51] Scott Yilek, Eric Rescorla, Hovav Shacham, Brandon Enright, and Stefan Savage. When private keys are public: Results from the 2008 debian openssl vulnerability. In *Internet Measurement Conference*. (Cited on page 19.)
 - [52] Paul Zimmerman. Integer factoring records. <http://www.loria.fr/~zimmerma/records/factor.html>. (Cited on page 20.)

A Proof of Lemma 4.5

We introduce the following notation for the proof. For a random variable V with range $\mathcal{V}$, we define the *collision probability* of V as $\text{Col}(V) = \Pr[V = V'] = \sum_{v \in \mathcal{V}} P_V(v)^2$ where V' is an independent copy of V , and for an event $\mathcal{E}$ we define the *conditional collision probability* $\text{Col}_{\mathcal{E}}(V) = \Pr[V = V' \mid \mathcal{E}]$. For random variables V, W , we define the *square of the 2-distance* as $D(V, W) = \sum_v (P_V(v) - P_W(v))^2$.

Writing $\mathbf{E}_k$ for expectation over the choice of random k from $\mathcal{K}$, we have

$$\Delta((K, g(X, h(K, X))), (K, g(X, U))) = \mathbf{E}_k[\Delta(g(X, h(k, X)), g(X, U))] \quad (9)$$

$$\leq \frac{1}{2} \mathbf{E}_k \left[\sqrt{|S| D(g(X, h(k, X)), g(X, U))} \right] \leq \frac{1}{2} \sqrt{|S| \mathbf{E}_k[D(g(X, h(k, X)), g(X, U))]} \quad (10)$$

where the first inequality is by Cauchy-Swartz and the second is by Jensen's inequality. We now show

$$\mathbf{E}_k[D(g(X, h(k, X)), g(X, U))] \leq \text{Col}(X)$$

from which the theorem follows. Write $(X, Y_k) = (X, h(k, X))$ for an arbitrary but fixed k . Then

$$\begin{aligned} D(g(X, Y_k), g(X, U)) &= \sum_s (P_{g(X, Y_k)}(s) - P_{g(X, U)}(s))^2 \\ &= \sum_s P_{g(X, Y_k)}(s)^2 - 2 \sum_s P_{g(X, Y_k)}(s) P_{g(X, U)}(s) + \sum_s P_{g(X, U)}(s)^2. \end{aligned}$$

Using the Kronecker delta $\delta_{s, s'}$ which equals 1 if $s = s'$ and else 0 for all $s, s' \in S$, we can write

$P_{g(X,Y_k)}(s) = \sum_x P_X(x) \delta_{g(x,h(k,x)),s}$, and thus

$$\begin{aligned} \sum_s P_{g(X,Y_k)}(s)^2 &= \sum_s \left(\sum_x P_X(x) \delta_{g(x,h(k,x)),s} \right) \left(\sum_{x'} P_X(x') \delta_{g(x',h(k,x')),s} \right) \\ &= \sum_{x,x'} P_X(x) P_X(x') \delta_{g(h(k,x)),g(h(k,x'))} . \end{aligned}$$

We use the pairwise independence of h to rewrite this in terms of collision probabilities:

$$\begin{aligned} \mathbf{E}_k \left[\sum_s P_{g(X,Y_k)}(s)^2 \right] &= \sum_{x,x'} P_X(x) P_X(x') \mathbf{E}_k [\delta_{g(x,h(k,x)),g(x',h(k,x'))}] \\ &= \text{Col}(X) + \text{Col}_{\mathcal{E}}(g(X,U))(1 - \text{Col}(X)), \quad (11) \end{aligned}$$

where the subscript $\mathcal{E}$ denotes (conditioning on) the event that $X \neq X'$. That is,

$$\text{Col}_{\mathcal{E}}(g(X,U)) = \Pr [g(X,U) = g(X',U') \mid X \neq X'] .$$

Similarly,

$$\begin{aligned} \sum_s P_{g(X,Y_k)}(s) P_{g(X,U)}(s) &= \sum_s \left(\sum_x P_X(x) \delta_{g(x,h(k,x)),s} \right) \left(\sum_{x',u} P_X(x') P_U(u) \delta_{g(x',u),s} \right) \\ &= \sum_x \sum_{x'} \sum_u P_X(x) P_X(x') P_U(u) \delta_{g(x,h(k,x)),g(x',u)} \end{aligned}$$

so that

$$\begin{aligned} \mathbf{E}_k \left[\sum_s P_{g(X,Y_k)}(s) P_{g(X,U)}(s) \right] &= \sum_x \sum_{x'} \sum_u P_X(x) P_X(x') P_U(u) \mathbf{E}_k [\delta_{g(x,h(k,x)),g(x',u)}] \\ &= \text{Col}(g(X,U)) = \text{Col}_{\bar{\mathcal{E}}}(g(X,U)) \text{Col}(X) + \text{Col}_{\mathcal{E}}(g(X,U))(1 - \text{Col}(X)) . \end{aligned}$$

where $\mathcal{E}$ is defined as above. Note that the only difference between the expression above and that in (11) is that even when $X = X'$, a collision is not guaranteed.

Finally,

$$\sum_s P_{g(X,U)}(s)^2 = \text{Col}(g(X,U)) = \text{Col}_{\bar{\mathcal{E}}}(g(X,U)) \text{Col}(X) + \text{Col}_{\mathcal{E}}(g(X,U))(1 - \text{Col}(X))$$

as well. By combining the above, we have

$$\begin{aligned} \mathbf{E}_k [D(g(X,Y_k), f(X,U))] &= \text{Col}(X) + \text{Col}_{\mathcal{E}}(g(X,U))(1 - \text{Col}(X)) \\ &\quad - 2(\text{Col}_{\bar{\mathcal{E}}}(g(X,U)) \text{Col}(X) + \text{Col}_{\mathcal{E}}(g(X,U))(1 - \text{Col}(X))) \\ &\quad + \text{Col}_{\bar{\mathcal{E}}}(g(X,U)) \text{Col}(X) + \text{Col}_{\mathcal{E}}(g(X,U))(1 - \text{Col}(X)) \\ &= (1 - \text{Col}_{\bar{\mathcal{E}}}(g(X,U))) \text{Col}(X) \\ &\leq \text{Col}(X) . \end{aligned}$$

To complete the proof, we can plug the bound above into (10):

$$\Delta((K, g(X, h(K, X))), (K, g(X, U))) \leq \frac{1}{2} \sqrt{|S| \mathbf{E}_k [D(g(X, h(k, X)), g(X, U))]} \leq \frac{1}{2} \sqrt{|S| \text{Col}(X)}.$$

By the assumption on the min-entropy of X , the collision probability $\text{Col}(X)$ is at most $4\hat{\varepsilon}^2/|S|$. So the statistical distance $\Delta((K, g(X, h(K, X))), (K, g(X, U)))$ is at most $\hat{\varepsilon}$, as desired. ■

B Security of OAEP Under Key-Independent Chosen-Plaintext Attack

The commonly-accepted notions of security for encryption ask for privacy with respect to messages that may depend on the public-key. We define here a notion of privacy for messages *not* depending on the public key. We mention that such a definition appears for example in the work of Micali et al. [41] (under the name “three-pass,” versus “one-pass,” cryptosystem), in the text of Goldreich [30], and in the context of the recent work on deterministic encryption [4].

THE DEFINITION. To an encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ and an adversary $B = (B_1, B_2)$ we associate

Experiment $\text{Exp}_{\Pi, B}^{\text{indki-cpa}}(k)$
 $b \xleftarrow{\$} \{0, 1\} ; (m_0, m_1, s) \xleftarrow{\$} B_1$
 $(pk, sk) \xleftarrow{\$} \mathcal{K} ; c \xleftarrow{\$} \mathcal{E}(pk, m_b)$
 $d \xleftarrow{\$} B_2(pk, c, s)$
 If $d = b$ then Return 1 Else Return 0

We require $|m_0| = |m_1|$ above. Define the *indki-cpa advantage* of B against Π as

$$\text{Adv}_{\Pi, B}^{\text{indki-cpa}}(k) = 2 \cdot \Pr \left[\text{Exp}_{\Pi, B}^{\text{indki-cpa}}(k) \Rightarrow 1 \right] - 1.$$

REMARKS. While non-standard, KI security seems adequate for some applications. For example, in [30] Goldreich points out that high-level applications that use encryption as a tool do so in a key-oblivious manner, and Bellare et al. [4] argue that in real life public keys are abstractions hidden in our software, so messages are unlikely to depend on them. KI security also suffices for hybrid encryption.

THE RESULT. We can show a standard model instantiation under KI security directly from Lemma 4.5, where G is any *pairwise-independent* function. This is captured by the theorem below.

Theorem B.1 Let $\text{LTDP} = (\mathcal{F}, \mathcal{F}')$ be an LTDP with residual leakage ℓ , and let OAEP be the encryption scheme associated to $\mathcal{F}$, hash functions G, H , and a parameter $k_0 < k$. Suppose G is pairwise-independent. Let $\varepsilon > 0$. Then for any $k_0 \geq \ell + 2\log(1/\varepsilon) - 2$ and any INDKI-CPA adversary B against OAEP, there is a distinguisher D against LTDP such that

$$\text{Adv}_{\text{OAEP}, B}^{\text{indki-cpa}}(k) \leq \text{Adv}_{\text{LTDP}, D}^{\text{ltdp}}(k) + \varepsilon.$$

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

As we mentioned, the proof is a simple hybrid argument concluding by Lemma 4.5.