Differential Indistinguishability for Cryptographic Primitives with Imperfect Randomness

Michael Backes[*] Aniket Kate[†] Sebastian Meiser[‡] Tim Ruffing[§]

Abstract

Indistinguishability-based definitions of cryptographic primitives such as encryption, commitments, and zero-knowledge proofs are proven to be impossible to realize in scenarios where parties only have access to *non-extractable* sources of randomness (Dodis et al., FOCS 2004). In this work we demonstrate that it is, nevertheless, possible to quantify this secrecy loss for non-extractable sources such as the (well-studied) Santha–Vazirani (SV) sources. In particular, to establish meaningful security guarantees in scenarios where such imperfect randomness sources are used, we define and study *differential indistinguishability*, a generalization of indistinguishability inspired by the notion of differential privacy.

We analyze strengths and weaknesses of differential indistinguishability both individually as well as under composition, and we interpret the resulting differential security guarantees for encryption, commitments, and zero-knowledge proofs.

Surprisingly, indistinguishability with uniform randomness carries over to differential indistinguishability with SV randomness: We show that all primitives that are secure under a traditional indistinguishibility-based definition are *differentially* secure when they use (a bounded amount of) SV randomness instead of uniform randomness.

[*] Department of Computer Science, Saarland University, and MPI-SWS, `backes@cs.uni-saarland.de`.

[†] MMCI, Saarland University, `aniket@mmci.uni-saarland.de`.

[‡] Department of Computer Science, Saarland University, `meiser@cs.uni-saarland.de`.

[§] Department of Computer Science, Saarland University, `tim.ruffing@stud.uni-saarland.de`.

Contents

1 Introduction

Most cryptographic protocols are designed and proven secure under the assumption that the protocol parties have access to perfect (uniform) randomness. Unfortunately, physical randomness sources that are employed in cryptographic implementations of those protocols are rarely the unbiased and independent perfect sources that are assumed in the security proofs of the protocols. In many cases such sources are not perfect, but only random in a weaker sense, e.g., they provide a certain amount of entropy. Moreover, many sources are found to be not extractable [13, 28], i.e., it is impossible to deterministically extract a super-logarithmic amount of nearly uniform randomness from them.

Santha–Vazirani (SV) sources [28] (and their generalizations [9, 13, 31]) are a prominent example of such non-extractable sources of randomness: a SV source produces an infinite bit sequence where each bit has almost one bit of entropy, but can have a small bias that might depend on all prior bits.

Although SV sources are non-extractable, their (reasonably high) entropy was found to be sufficient for simulating probabilistic algorithms [1, 9, 28, 31] and for implementing some cryptographic primitives [11, 13, 23] such as message authentication codes and signatures. Intuitively, the security of these primitives resides in the impossibility of predicting a whole bitstring, and consequently a high entropy is sufficient to render this task computationally hard.

However, many cryptographic definitions rely on the well-established notion of indistinguishability, i.e., an attacker only needs to distinguish two scenarios to break the cryptographic definition. Weak sources such as SV sources have been found to be insufficient [5, 13, 24] to achieve these definitions. More precisely, McInnes and Pinkas [24] show that unconditionally secure symmetric encryption of even a single bit cannot be based on a weak source. Dodis et al. [13] prove that block sources (a generalization of SV sources) are not sufficient for realizing any cryptographic primitive that requires a secrecy guarantee that is based on indistinguishability. In particular, traditional secrecy guarantees cannot be given for primitives such as encryption, commitments, zero-knowledge proofs, and secret sharing, even against a computationally bounded adversary.

This line of work indicates a gap between cryptographically secure primitives and their usage with imperfect randomness. The following points are currently unclear:

- Is it possible to quantify the secrecy loss of cryptographic primitives, e.g., encryption schemes, if they use imperfect randomness sources such as SV sources? How can we give meaningful quantitative guarantees in such cases?
- Do these quantitative guarantees require new constructions, or do they apply to existing schemes?
- Given that these quantitative guarantees are necessarily weaker than traditional cryptographic guarantees, what practical risks and drawbacks do they impose on the security of cryptographic systems?

1.1 Our Contributions

New Notion: Differential Indistinguishability. To address the first question, we define *differential indistinguishability*, a generalization of cryptographic indistinguishability in the spirit of (computational) differential privacy [25] and pseudodensity [27]. Differential indistinguishability quantifies the secrecy loss of cryptographic primitives in scenarios where uniform randomness is not available. To describe those scenarios we make use of a generalized form of the Santha-Vazirani (SV) randomness sources [28] that allows for reasoning about blocks of randomness at once (instead of working on a bit-by-bit basis). We show that traditional indistinguishability (under the assumption of uniform randomness) suffices to guarantee differential indistinguishability if a generalized SV source is used instead of the uniform source. We instantiate differential indistinguishability by several (indistinguishability-based) security definitions for cryptographic primitives such as the *hiding* property for commitment schemes, *indistinguishability against adaptive chosen ciphertext attacks* (IND-CCA-II) for encryption, or *zero-knowledge* for interactive proof systems.

Guarantees for Cryptographic Primitives. We present quantitative guarantees for all cryptographic constructions that are used with a generalized SV source, given that the construction has been proven secure on uniform randomness. This enables us to present a non-trivial lower bound for the security in cases where a general impossibility result by Dodis et al. [13] has ruled out traditional indistinguishability guarantees. Interestingly, new constructions for these primitives are *not* necessary in order to obtain such guarantees. We present theorems that reason immediately about all existing primitives whose security definitions are based on (information-theoretical or computational) indistinguishability, as long as the amount of non-uniform randomness can be bounded.

To demonstrate the general applicability of our results, we first show that commitment schemes that are information-theoretically hiding with uniform randomness are differentially indistinguishable against unbounded adversaries if they rely on a generalized SV source instead. This result is useful for security protocols such as electronic voting that promise everlasting privacy [26]. Second, we show that public-key encryption schemes that are indistinguishable under adaptive chosen ciphertext attack also are differentially indistinguishable under adaptive chosen-ciphertext attack if a generalized SV source is used instead of uniform randomness. This result applies both for the case when a generalized SV source is used for generating the keys as well as the case when it is used for encrypting a message. Third, we demonstrate that differential indistinguishability is applicable to zero-knowledge proofs as well. We show that proof systems that are zero-knowledge in the traditional sense are differentially zero-knowledge if the prover uses a generalized SV source instead of uniform randomness. For non-interactive zero-knowledge proofs, we also consider an imperfectly sampled common reference string.

The applicability to commitments, encryption and zero-knowledge proofs demonstrates how traditional indistinguishability-based definitions can be extended to achieve guarantees whenever a generalized SV source is employed in practice. We note that both our notions and our proofs can directly be applied to other indistinguishability-based security such as semantic security for encryption and computational hiding for commitments, and to other primitives such as secret sharing, pseudorandom generators, and oblivious transfer. Our analysis can also be used in scenarios in which one cannot influence the randomness source (anymore), e.g., for analyzing the security of a voting system after the election is over.

Risks and Drawbacks. We evaluate the relation between differential indistinguishability and the well-studied notion of differential privacy [15, 25]. Notably, we show that differential indistinguishability is composable in a similar way as differential privacy, i.e., from a guarantee for a single execution one can derive a guarantee for several (serial) executions. This composition comes with a loss of secrecy that is linearly bounded in the number of executions, similar as when composing several database queries in differential privacy.

1.2 An Overview of Our Results

Our differential indistinguishability notion is structurally related to differential privacy [15], which originally quantifies the privacy loss of a probabilistic statistical algorithm f by the probability of guessing on which of two similar input databases (D, D') the algorithm is executed. Recall that a statistical algorithm $f : \mathcal{D} \to \mathcal{R}$ is *ε-differentially private* [15] if for all neighboring databases D, D' and every set $S \subseteq \mathcal{R}$ of possible results of f, $\Pr[f(D) \in S] \leq e^{\varepsilon} \Pr[f(D') \in S]$, where $\varepsilon \geq 0$ is a "reasonably small" constant that quantifies the privacy loss of the statistical algorithm. In contrast to traditional information-theoretical or computational indistinguishability for differential privacy, the probability of guessing right can be *non-negligibly* higher than the probability of guessing wrong. However, as long as the disclosure is within a small multiplicative factor e^{ε}, the algorithm is considered to be differentially private. Such a weaker guarantee for statistical algorithms is necessary as traditional indistinguishability is not possible if the algorithms are to output (useful) statistics.

Differential Indistinguishability. For cryptographic primitives that are used with imperfect randomness a similar impossibility result applies. Dodis et al. [13] show that in the case of imperfect randomness, traditional indistinguishability is provably impossible for cryptographic primitives that have a secrecy requirement, e.g., encryption, commitments, and zero-knowledge proofs. As differential privacy avoids the impossibility of traditional indistinguishability for (useful) statistics about databases, we aim for a similar relaxation of the security guaranteed by the aforementioned cryptographic primitives in order to avoid the impossibility in the case of imperfect randomness.

To this end, we introduce *differential indistinguishability*, a generalization of cryptographic indistinguishability in the spirit of computational differential privacy [25] and pseudodensity [27]. Two games, i.e., the interactions with two machines X_0 and X_1, are (ε, δ)-differentially indistinguishable if for all interactive distinguisher machines A, the output probabilities for all outputs are related by

$$\Pr[\langle \mathsf{A} | \mathsf{X}_0 \rangle = x] \leq e^{\varepsilon} \Pr[\langle \mathsf{A} | \mathsf{X}_1 \rangle = x] + \delta,$$

where x is a possible output of A. Intuitively, δ remains a negligible function, whereas $\varepsilon \geq 0$ is a reasonably small constant.

As our main contribution we show that traditional indistinguishability (for uniform randomness) implies differential indistinguishability for imperfect randomness. Consequently, all primitives that are secure for uniform randomness are also differentially secure for imperfect randomness. To give this strong result we first show that SV distributions and the uniform distribution are $(\varepsilon, 0)$-differentially indistinguishable with ε depending on the maximal bias of the distribution, even for a computationally unbounded distinguisher A (Lemma 1). We then extend this lemma to show our main technical result (Lemma 2), which states that two games/machines X_0 and X_1 are (ε, δ)-differentially indistinguishable for SV distributions with a negligible δ, if they are indistinguishable for the uniform distribution.

This result holds for arbitrary classes of adversaries and thus allows for deriving results for both information-theoretically as well as computationally indistinguishable schemes. One can instantiate the games from Lemma 2 with the indistinguishability games that are to be analyzed and directly compute the results.

An Illustrative Example. Moran and Naor [26] introduce the notion of *everlasting privacy* for electronic voting (e-voting) using information-theoretically hiding commitments. However, if the randomness employed for commitments is not perfect, traditional cryptographic indistinguishability notions cannot make a statement. Using our notion of differential indistinguishability, Lemma 2 gives meaningful bounds for the privacy loss of the vote for all generalized SV sources. Moreover, for SV sources the scheme still satisfies a weaker (non coercive) form of deniability in the following sense:

For any given vote a certain candidate, e.g., Alice, might appear more likely, but it could very well be a vote for any other candidate.

Interpretation. The non-negligible multiplicative factor in differential indistinguishability may weaken the security guarantees, and similar to differential privacy it needs to be interpreted carefully in the respective context. To derive quantitative guarantees, we require that the amount of (imperfect) randomness used by the challenger in the indistinguishability game for a scheme has to be polynomially bounded in terms of the security parameter. This requirement is natural, given that our goal is to reason about the influence of the randomness. Surprisingly, this is the only requirement if we consider SV sources. This leads to the observation that all existing primitives that use a bounded amount of randomness can immediately be analyzed and their secrecy loss can be quantified by an additional multiplicative factor that depends on the quality of the random source, without any need for new constructions, although in some cases better constructions could reduce the impact of imperfect randomness significantly.

As the parameter ε is a multiplicative parameter, the guarantee is not trivial, even for small constants ε. It is easy to see that a multiplicative bound can always be transformed into an additive bound, but the converse is not true.

Compostion. Recall that the privacy loss in differential privacy linearly accumulates under composition, which leads to the notion of a "privacy budget". Intuitively, such a budget describes the number of queries (executions of the algorithm) that are allowed on the sensitive database. The same argumentation applies to differential indistinguishability guarantees when we compose cryptographic primitives. We can easily give guarantees for the composition of two or more primitives, but the amount of used randomness and thus the "secrecy loss" increases linear in the number of compositions. This is critical, because in contrast to the traditional setting this secrecy loss is not negligible; thus, a "secrecy budget" (the counterpart of the "privacy budget") is necessary.

Future Directions. Our work presents a novel view on the relation between imperfect randomness and indistinguishability. This naturally leads to many more interesting questions. Which properties of indistinguishability propagate to differential indistinguishability? For example, many indistinguishability-based definitions such as *indistinguishability under chosen plaintext attack* and *semantic security* for public-key encryption have been shown to be equivalent. Which of these equivalences hold for their differential counterparts? What happens if a (programmable) random oracle provides imperfect randomness? For which sources other than SV sources can one give differential guarantees? Is differential indistinguishability helpful towards improving results on leakage-resilient cryptographic schemes [6, 17]? For example, can one give differential guarantees in cases where the adversary learns more than is considered by the existing leakage-resilient schemes?

1.3 Related Work

The effect of imperfect randomness on traditional cryptography is well studied. On the negative side, several papers demonstrate the inherent limitations of indistinguishability-based cryptographic guarantees with imperfect randomness [2, 5, 13, 14]. Most closely related to our work, Dodis et al. [13] show that traditional indistinguishability required for encryption, commitment, secret sharing, and zero-knowledge cannot be realized if a SV source is used. More precisely, they prove that no protocol for any of these primitives can be secure against certain block sources, which include SV sources. These sources sample blocks (i.e., several bits at once) that are $1/\mathsf{poly}(k)$ close to the uniform distribution [9, 13, 28] for an arbitrary polynomial, where k is the security parameter.

This impossibility result has been refined and generalized over the last few years. Dodis, Pietrzak, and Przydatek [14] proved that using imperfect randomness, secure encryption implies secure $(2, 2)$-secret sharing; however, the converse does not hold. Bosley and Dodis [5] showed that information-theoretically secure encryption of more than $\log(n)$ bits is possible only if more than $\log(n)$ almost-uniform bits can be extracted from the source in the first place. Recently, Austrin et. al. [2] refined the impossibility result [13] to show that it holds even when "efficient SV adversaries" are considered. They also consider a min-entropy-based multiplicative factor in their tamper-resilient signature analysis, which is an instantiation of our unbreakability analysis.

On the positive side, one line of research examines the extraction of (almost) perfect randomness from several kinds of imperfect randomness sources [4,9,10,21,29,30]. However, extraction generally requires the source to have a certain degree of independence, whereas the only main requirement for SV sources and their variants is to provide some entropy. Aiming at particular applications, some works have shown that a few primitives can be securely instantiated even if only imperfect randomness is available [12,13,19,20]. Dodis et al. [13] prove that signatures can be successfully instantiated using a block source instead of uniform randomness. Goldwasser, Sudan, and Vaikuntanathan [19] show that Byzantine agreement is possible for some suitable imperfect randomness. Dodis et al. [12] prove that differential privacy of statistical queries can be preserved even when the noise is generated using an imperfect random source. In particular, they ask the interesting question of whether or not differential privacy is possible if no uniform randomness is available, and give a positive answer for SV sources by presenting a γ-differentially private algorithm that works on these sources. Relevant to our observations, they note that traditional indistinguishability-based privacy is a stronger notion as compared to, e.g., unforgeability.

Kamara and Katz [20] propose a notion of security for symmetric-key encryption that is able to cope with imperfect randomness. However, their notion applies only if the challenge messages are encrypted using uniform randomness. We consider their work orthogonal to ours and refer to Appendix A.1 for more thoughts about a combination of the works. In the universal composability (UC) setting, Canetti, Pass, and Shelat [8] showed that even for (sampleable) sources for which a deterministic extractor exists, UC-secure commitments are not possible. Instead, they present a UC-secure commitment scheme assuming a collision-resistant hash function, dense cryptosystem and one-way function (with sub-exponential hardness), and requiring access to $O(1)$ instances of a gray-box source with sufficient entropy. Finally, Brandao et al. [7] show that in the quantum setting, single sources of SV randomness can be improved, where the running time is polynomial in the inversed distance to the uniform distribution. This result indicates that SV sources that are $1/\mathsf{poly}(k)$ close to the uniform distribution for some polynomial and the security parameter k might be a reasonable assumption for cryptography in general.

2 Preliminaries

We denote sampling an element r from a distribution D by $r \leftarrow D$. The probability of the probabilistic event $F(r)$, where r is sampled from the distribution D, is denoted by $\Pr[F(r) \mid r \leftarrow D]$ or more compactly by $\Pr[F(D)]$. To keep the notation simple, we may write f_k for the value of a function $f(\cdot)$ applied to k. We denote by $\{D_{\ell_k}\}_{k\in\mathbb{N}}$ a family of distributions such that for each $k \in \mathbb{N}$ the distribution D_{ℓ_k} samples elements from $\{0,1\}^{\ell_k}$. In particular, $\{U_{\ell_k}\}_{k\in\mathbb{N}}$ is the family of uniform distributions, where U_{ℓ_k} is the uniform distribution over $\{0,1\}^{\ell_k}$.

Throughout the paper we consider interactive Turing machines that have access to a uniform random tape, even if they additionally get an input drawn from some distribution. Unless we mention that they run in probabilistic polynomial time (ppt), those machines are *not bounded*. We write $x \leftarrow \langle \mathsf{X}|\mathsf{Y}\rangle$ for denoting the interaction of the two machines X and Y. The output x is always the output of the first machine X. We also shorten this to $\langle \mathsf{X}|\mathsf{Y}\rangle = x_0$ for denoting the event $x = x_0$, where $x \leftarrow \langle \mathsf{X}|\mathsf{Y}\rangle$.

Randomness Sources. We focus on randomness sources that are a generalization of Santha–Vazirani (SV) sources [28] to block sources [9,13]. Block sources are well-suited to describe both physical random sources as well as certain random sources that have been "tampered with" by an adversary [2].

Definition 1 (Block source) *A tuple of distributions $D = (D^1, \ldots, D^t)$, each over the set $\{0,1\}^n$ of bitstrings of length n, is* (n, γ)-Santha–Vazirani (SV) *(for $0 < \gamma < 1$) if for all $0 \leq i \leq t$ and for all $x_1, \ldots, x_i \in \{0,1\}^n$,*

$$(1-\gamma) \cdot 2^{-n} \leq \Pr\left[D^i = x_i | x_1 \leftarrow D^1, \ldots, x_{i-1} \leftarrow D^{i-1}\right] \leq (1+\gamma) \cdot 2^{-n}.$$

The original SV sources are a special case of Definition 1 that arises if we set $n = 1$. Note that if we consider the probability for yielding a combined string $x \in \{0,1\}^{nt}$ from all distributions $(D^1, \ldots, D^t)$, the probability for x is bounded by $((1-\gamma) \cdot 2^{-n})^t \leq \Pr[D = x] \leq ((1+\gamma) \cdot 2^{-n})^t$.

SV Sources. We call a family of distributions $\{D_{t_k n_k}\}_{k \in \mathbb{N}}$ an (n, γ)-SV source, if every element $D_{n_k t_k}$ can be described by a (n_k, γ)-SV block distribution consisting of t_k distributions, each over $\{0,1\}^{n_k}$.

3 Main Results

In this section we present our main lemmas that can be applied to a variety of cryptographic notions. We begin by describing and proving an important fact about SV distributions, namely that they are differentially hard to distinguish from the uniform distribution.

Although an (n, γ)-SV distribution is not negligibly close to a uniform distribution, the parameter γ gives a bound on the discrepancy between the uniform distribution and the γ-SV distribution if we restrict the number of queries ("samples") that a distinguisher is allowed to make by a positive polynomial t. The following lemma shows that there is no adversary that is able to distinguish the distributions better than with a multiplicative factor of e^{ε}, depending on the number of bits and the parameter γ of the distribution.

Lemma 1. *Let $\{D_{n_k t_k}\}_{k \in \mathbb{N}}$ be a family of (n, γ)-SV sources over $\{0,1\}^{n_k t_k}$ with $\gamma < 1/2$, and let $\{U_{n_k t_k}\}_{k \in \mathbb{N}}$ be a family of uniform sources over $\{0,1\}^{n_k t_k}$. For all probabilistic machines A and for all possible outputs x of A,*

$$\Pr\left[\mathsf{A}(1^k, D_{n_k t_k}) = x\right] \leq e^{t_k \gamma_k} \Pr\left[\mathsf{A}(1^k, U_{n_k t_k}) = x\right] \tag{a}$$

$$\text{and} \qquad \Pr\left[\mathsf{A}(1^k, U_{n_k t_k}) = x\right] \leq e^{2 t_k \gamma_k} \Pr\left[x \leftarrow \mathsf{A}(1^k, D_{n_k t_k}) = x\right]. \tag{b}$$

Proof. Let a (n, γ)-SV distribution $\{D_{n_k t_k}\}_{k \in \mathbb{N}}$ with $\gamma < 1/2$ over a set $\{0,1\}^{n_k t_k}$ be given. Let A be any probabilistic machine. Essentially, we have to prove the following: No matter which test A performs on a value $r \in \{0,1\}^{n_k t_k}$, the probability that r has been generated by $D_{n_k t_k}$ is very close to the probability that r has been generated by $U_{n_k t_k}$.

We start by proving (a): Note that for all values $r_0 \in \{0,1\}^{n_k t_k}$ the probability $\Pr[D_{n_k t_k} = r_0]$ is strictly larger than zero. For all values $r_0 \in \{0,1\}^{n_k t_k}$,

$$\log\left(\frac{\Pr[D_{n_k t_k} = r_0]}{\Pr[U_{n_k t_k} = r_0]}\right) \leq \log\left(\frac{\left(2^{-n_k}(1+\gamma_k)\right)^{t_k}}{2^{-n_k t_k}}\right) = t_k \cdot \log(1+\gamma_k) \leq t_k \gamma_k.$$

Let r_{A} be the arbitrary but fixed random choices of A from the randomness tape (i.e., not the input). By $\mathsf{A}(r, r_{\mathsf{A}})$ we denote the machine that simulated A on input r such that A takes the random choices r_{A}. Using this equation we can show (i) as follows. For all possible outputs x of A,

$$\begin{aligned} \Pr\left[\mathsf{A}(1^k, D_{n_k t_k}) = x\right] &= \sum_{r_0 \in \{0,1\}^{n_k t_k}} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \Pr\left[D_{n_k t_k} = r_0\right] \\ &\leq \sum_{r_0 \in \{0,1\}^{n_k t_k}} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] e^{t_k \gamma_k} \Pr\left[U_{n_k t_k} = r_0\right] \\ &= e^{t_k \gamma_k} \Pr\left[\mathsf{A}(1^k, U_{n_k t_k}) = x\right]. \end{aligned}$$

This shows (a). For (b), note that for all values $r_0 \in \{0,1\}^{n_k t_k}$ the probability $\Pr\left[D_{n_k t_k} = r_0\right]$ is strictly larger than zero. For all values $r_0 \in \{0,1\}^{n_k t_k}$,

$$\begin{aligned} \log\left(\frac{\Pr\left[U_{n_k t_k} = r_0\right]}{\Pr\left[D_{n_k t_k} = r_0\right]}\right) &\leq \log\left(\frac{\left(\frac{1}{2^{n_k}}\right)^{t_k}}{\left(\frac{1}{2^{n_k}}(1-\gamma_k)\right)^{t_k}}\right) \\ &= -t_k \cdot \log\left(1-\gamma_k\right) \\ &\leq t_k \cdot \left(\frac{\gamma_k}{1-\gamma_k}\right) \\ &\leq 2 t_k \gamma_k, \end{aligned}$$

where the last two bounds hold because $\gamma_k < 1/2$. Thus, using this inequality, we obtain for all possible outputs x of A and $\gamma_k < 1/2$ that

$$\begin{aligned} \Pr\left[\mathsf{A}(1^k, U_{n_k t_k}) = x\right] &= \sum_{r_0} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \cdot \Pr\left[U_{n_k t_k} = r_0\right] \\ &\leq \sum_{r_0} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \cdot e^{2 t_k \gamma_k} \Pr\left[D_{n_k t_k} = r_0\right] \\ &= e^{2 t_k \gamma_k} \Pr\left[\mathsf{A}(1^k, D_{n_k t_k}) = x\right]. \end{aligned}$$

This completes the proof. □

3.1 Differential Indistinguishability

Traditional cryptography defines two machines X_0 and X_1 to be *indistinguishable* for a certain class of distinguishers $\mathcal{A}$, if no distinguisher $\mathsf{A} \in \mathcal{A}$ in this class is able to notice a difference between an interaction with X_0 and an interaction with X_1. Formally, the concept of "noticing a difference" is captured by saying that any possible view of a distinguisher is (almost) equally likely for both X_0 and X_1, i.e., the difference of the probability that A outputs any given view in the interaction with X_0 from the probability of outputting the same view in the interaction with X_1 is negligible. We consider a variant of indistinguishability that allows these probabilities to be related also by a multiplicative factor $e^{\varepsilon} > 1$, similar to the concept of mutual pseudodensity [27] employed, e.g., for computational differential privacy [25].

Definition 2 (Differential Indistinguishability) *Two probabilistic machines X_0 and X_1 are (ε,δ)-differentially indistinguishable over a distribution $\{D_{\ell_k}\}_{k\in\mathbb{N}}$ over $\{0,1\}^{\ell_k}$ for a positive polynomial ℓ and a class $\mathcal{A}$ of adversaries (probabilistic machines), if for all $\mathsf{A} \in \mathcal{A}$, for all sufficiently large k, for all possible outputs x of A, and for all $b \in \{0,1\}$,*

$$\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_b(1^k, D_{n_k t_k})\right\rangle = x\right] \leq e^{\varepsilon} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_{1-b}(1^k, D_{n_k t_k})\right\rangle = x\right] + \delta_k.$$

This definition is constructed to be very general, which allows it to describe many of the traditional cryptographic indistinguishability notions [18, 22]. For the traditional case of $\varepsilon = 0$ we speak of δ-indistinguishability. The definition covers interactive and non-interactive notions, as well as simulation-based notions. For information-theoretic secrecy, the class $\mathcal{A}$ of adversaries is the class of all probabilistic (possibly unbounded) machines and we have $\delta_k = 0$. Statistical secrecy can be expressed by the same class of adversaries for $\delta_k > 0$. Cryptographic (computational) secrecy can be achieved with the class $\mathcal{A}_{\mathsf{poly}}$ of probabilistic polynomial-time machines with δ_k being a negligible function. This differential definition additionally allows for a quantitative factor of $\varepsilon > 0$ that can (and has to) be interpreted carefully.

For SV distributions, differential indistinguishability is directly implied by traditional indistinguishability (under uniform randomness). The following lemma is the main result we base our analysis on. It allows us to easily give guarantees for cryptographic primitives, whenever their security notions can be expressed in terms of Definition 2.

Lemma 2. *If two probabilistic machines X_0 and X_1 are δ-indistinguishable for a class of probabilistic machines $\mathcal{A}$ and the family of uniform sources $\{U_{n_k t_k}\}_{k\in\mathbb{N}}$ over $\{0,1\}^{n_k t_k}$, then X_0 and X_1 are also $(\varepsilon, e^{\varepsilon}\delta_k)$-differentially indistinguishable for $\mathcal{A}$ and any family of (n,γ)-SV sources $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$, where $\gamma_k \le \min\{1/2,\ \varepsilon/3t_k\}$.*

Proof. Let $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$ be a (n,γ)-SV source, and $\{U_{n_k t_k}\}_{k\in\mathbb{N}}$ be the uniform source, both over $\{0,1\}^{n_k t_k}$. Furthermore, let X_0, X_1, be probabilistic (not necessarily polynomially bounded) machines, and let $\mathsf{A} \in \mathcal{A}$ be an adversary machine such that for a function δ

$$\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_0(1^k, U_{n_k t_k}) \right\rangle = x\right] \le \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_1(1^k, U_{n_k t_k}) \right\rangle = x\right] + \delta_k.$$

Using Lemma 1 we show that A behaves similarly on D, as otherwise a machine that simulates $\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_0(1^k, r) \right\rangle$ (or $\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_1(1^k, r) \right\rangle$) could distinguish $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$ and $\{U_{n_k t_k}\}_{k\in\mathbb{N}}$.

$$\begin{aligned}
\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_0(1^k, D_{n_k t_k}) \right\rangle = x\right] &\le e^{t_k\gamma_k} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_0(1^k, U_{n_k t_k}) \right\rangle = x\right] && (1)\\
&\le e^{t_k\gamma_k} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_1(1^k, U_{n_k t_k}) \right\rangle = x\right] + e^{t_k\gamma_k}\delta_k && (2)\\
&\le e^{3t_k\gamma_k} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_1(1^k, D_{n_k t_k}) \right\rangle = x\right] + e^{t_k\gamma_k}\delta_k && (3)
\end{aligned}$$

Here, inequalities (1) and (3) follow from inequalities (a) and (b) in Lemma 1, respectively. The remaining inequality (2) holds by assumption. □

Intuition on the counter direction of Lemma 2. Differentially private mechanisms are an intuitive example that shows why the counter direction of Lemma 2 does not hold. If a mechanism is differentially private, it is not necessarily computationally or even information-theoretically indistinguishable, as that would conflict with a reasonably high utility. Such a mechanism might reach a given value for ε when using imperfect randomness (as in [12]). However, this does not imply that with access to uniform randomness the mechanism is δ-indistinguishable for neighboring databases for a negligible function δ.

3.2 Unbreakability

We further analyze how imperfect randomness influences the probability for guessing a whole bitstring, e.g., breaking the binding property of a commitment. The corresponding security definitions typically require that no adversary has more than a negligible chance to reach a certain bad event. We generalize the intuition of *breaking a scheme* by dividing a game Z into two parts. The "normal game" Z_0 and a judge Z_1 that decides whether or not a given string constitutes a break of the scheme.

Definition 3 (Unbreakability) *Let $\mathsf{Z} = (\mathsf{Z}_0, \mathsf{Z}_1)$ be a probabilistic machine that may keep state. We say that Z is δ-*unbreakable *for a class $\mathcal{A}$ of adversaries and for a distribution $\{D_{\ell_k}\}_{k\in\mathbb{N}}$ over $\{0,1\}^{\ell_k}$, if for all $\mathsf{A} \in \mathcal{A}$ and for sufficiently large k,*

$$\Pr\left[\mathsf{Z}_1(a) = 1 \;\middle|\; a \leftarrow \left\langle \mathsf{A}(1^k) \middle| \mathsf{Z}_0(1^k, D_{\ell_k}) \right\rangle\right] \le \delta_k.$$

We can show that for all games that can be described as a unbreakability game and for which the probability to win is negligible under uniform randomness, the probability is still negligible if a SV source with matching parameters is used.

Lemma 3. *If a probabilistic machine $\mathsf{Z} = (\mathsf{Z}_0, \mathsf{Z}_1)$ that may keep state is δ-unbreakable for a class of probabilistic machines $\mathcal{A}$ and uses a uniform randomness source $\{U_{n_k t_k}\}_{k\in\mathbb{N}}$ over $\{0,1\}^{n_k t_k}$, then Z is $(e^{\varepsilon}\delta)$-unbreakable for $\mathcal{A}$ if it uses a (n, γ)-SV source $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$ instead, where $\gamma \le \min\{1/2, \varepsilon/3t_k\}$.*

Proof. We reduce this lemma to Lemma 1 as follows: Let $\mathsf{Z} = (\mathsf{Z}_0, \mathsf{Z}_1)$ be a probabilistic (not necessarily polynomially bounded) machine that may keep state. Given any adversary $\mathsf{A} \in \mathcal{A}$, we construct a probabilistic machine B on input $r \in \{0,1\}^{n_k t_k}$ as follows. B simulates the interaction between A and $\mathsf{Z}_0(1^k, r)$, yields an output a and simulates Z_1 on a. If Z_0 keeps state for Z_1, B also simulates this behavior. It holds that

$$\begin{aligned}
\Pr&\left[\mathsf{Z}_1(a) = 1 \;\middle|\; a \leftarrow \left\langle \mathsf{A}(1^k) \middle| \mathsf{Z}_0(1^k, D_{n_k t_k}) \right\rangle\right] \\
&= \Pr\left[\mathsf{B}(1^k, D_{n_k t_k}) = 1\right] \\
&\le e^{t_k \gamma_k} \Pr\left[\mathsf{B}(1^k, U_{n_k t_k}) = 1\right] && (4)\\
&= e^{t_k \gamma_k} \Pr\left[\mathsf{Z}_1(a) = 1 \;\middle|\; a \leftarrow \left\langle \mathsf{A}(1^k) \middle| \mathsf{Z}_0(1^k, U_{n_k t_k}) \right\rangle\right] \\
&\le e^{t_k \gamma_k}\, \delta_k, && (5)
\end{aligned}$$

where inequality (4) holds by Lemma 1 and inequality (5) follows by assumption. □

In the remainder of this paper we first give some intuition about how differential indistinguishability can be interpreted, how it relates to differential privacy and how and why there is a secrecy loss under composition. Finally we exemplify the instantiability of differential indistinguishability in Section 5 by presenting differential guarantees for commitments, encryption and zero-knowledge proofs.

4 Interpretation and Analysis

Impact of Multiplicative Factor. Similar to differential privacy, differential indistinguishability adds a multiplicative factor to the inequality used in the traditional indistinguishability notion. We observe that a multiplicative bound may express properties that are inexpressible by an additive bound. While every multiplicative bound of the form $\Pr[A] \le e^{\varepsilon} \Pr[B] + \delta_k$ implies a purely additive bound $\Pr[A] \le \Pr[B] + \delta_k + e^{\varepsilon} - 1 \approx \Pr[B] + \delta_k + \varepsilon$, the converse does not hold in general. No matter which additive bound can be shown between two probabilistic events, there does not necessarily exist a multiplicative bound. In particular, there are machines that are δ-indistinguishable for some δ but not (ε, δ')-indistinguishable for any ε such that $\delta' < \delta$. We refer to Appendix B.1 for a formal counterexample.

For secrecy properties, traditional indistinguishability intuitively states that no adversary can learn any information about the secret, except with negligible probability. Our multiplicative factor generalizes indistinguishability to additionally allow the adversary to learn information about the secret with more than a negligible probability, as long as the loss of secrecy is bounded; e.g., if ε is a small constant then differential indistinguishability ensures that the owner of the secret retains deniability by introducing doubt for the adversary.

As an illustrative example, consider the e-voting protocol based on a commitment scheme that is $(\varepsilon, 0)$-differentially hiding such that ε is a small constant or decreases with the security parameter. Assume that, among two candidates (say) Alice and Bob, a voter voted for Alice. In the traditional indistinguishability case, a non-negligible additive difference would have resulted in a non-negligible probability for leaking the vote. The multiplicative factor in differential indistinguishability, however, does not lead to such events. A multiplicative factor means that two cases (i.e., the case where the vote was casted for Alice and the case where it was casted for Bob) can be differentiated with a (possibly) non-negligible advantage, but both cases are still almost equally probable. No distinguisher can be sure that it has observed a vote for Alice or a vote for Bob. In other words, there cannot be a distinguisher that is certain of its answer. Consider a distinguisher that only outputs (say) 1 if it is certain that the vote was casted for Alice, and that outputs 0 in all other cases. Such a distinguisher is affected by the multiplicative bound as the output 1 is almost equally probable in all cases. Moreover, if the probability of outputting 1 is zero when the vote was casted for Bob, then differential indistinguishability implies that the probability of outputting 1 is zero when the vote was casted for Alice.

Notice that the same analysis applies if a negligible additive value $\delta \neq 0$ is present. In this case, there might be a negligible chance for the adversary to be certain about the vote (e.g., if one bit of the vote is leaked), but in all other cases deniability is preserved.

Relation to Differential Privacy and Sensitivity. The close relation of differential indistinguishability to differential privacy (DP) is helpful in interpreting the guarantees and in understanding the drawbacks. Differential privacy is influenced by the *sensitivity* of a statistical function, i.e., the amount of influence individual database records can have on the function. The sensitivity is directly proportional to the amount of randomness to be added to the output of the function in order to guarantee a certain ε-level of privacy. Consequently, the sensitivity directly influences to what degree the utility of the noisy output is decreased (in comparison to the original function).

Although there are neither databases nor the concept of utility (in the same sense as DP) in our setting, we can regard distinguishing individual blocks of an (n, γ)-SV distribution from uniform blocks as a DP game. We can then consider the sensitivity of a scheme corresponding to a security definition as the number of randomness blocks that are necessary for honest parties in the definition; e.g., for an encryption scheme that uses ℓ_k random bits from a $(1, \gamma)$-SV source, the sensitivity would be ℓ_k.

A practical implication of this relation is that the sensitivity is directly connected to γ of the source and ε of the guarantee. The higher the sensitivity (i.e., the more randomness is drawn by honest parties) the smaller γ must be to allow for guaranteeing ε-differential indistinguishability. Clearly, the bias in a $(1, \gamma)$-SV source can be arbitrarily increased by drawing more random bits; e.g., by drawing three bits and taking the majority vote. Although this technique does not make a difference for uniform randomness, it may increase the bias of the bits for imperfect randomness. Therefore, the amount of randomness (and the sensitivity) is a necessary parameter in the definition of differential indistinguishability.

Composability. Traditional indistinguishability with a negligible function δ and $\varepsilon = 0$ allows for polynomially many compositions, i.e., seeing multiple samples does not help the adversary substantially, because a polynomial factor does not influence the negligibility. For differential indistinguishability, this argument does not apply as the (non negligible) multiplicative factors can be accumulated as well. On the positive side, if a scheme is indistinguishable and composable, and there is a bound on the number of executions, then Lemma 2 allows for directly deriving a guarantee for an (n, γ)-SV source. On the negative side, the composition increases the sensitivity linearly in the number of executions, like for DP. As a result, to guarantee $(\varepsilon, e^{\varepsilon}\delta)$-differential indistinguishability for q compositions (using Lemma 2) for a (n, γ)-SV source, γ must decrease faster in comparison to the case without composition, i.e., $\gamma_k \leq \min\{1/2,\ \varepsilon/(3q \cdot t_k)\}$.

In real-life scenarios, where an adversary might have auxiliary information, this concern is important; e.g., consider our illustrative example of e-voting with everlasting privacy [26] with a SV source. If an adversary observes, e.g., several everlastingly private votes of members from the same family, and knows that all family members have voted for the same party, then the adversary can improve its success probability (ε is increased at most by a factor of 5).

On the other hand, similar to DP, as long as an individual security game (e.g., a challenge commitment) is independent from other observations, the (ε, δ)-guarantee holds. Thus, we can speak of "neighboring scenarios" (in the DP terms) in which tuples of (independent) games $(\mathsf{X}_0^1, \ldots, \mathsf{X}_0^q)$ and $(\mathsf{X}_1^1, \ldots, \mathsf{X}_1^q)$ that might only differ in one element $\mathsf{X}_0^i \neq \mathsf{X}_1^i$, and are the same otherwise.

Reduction Proofs on Differential Indistinguishability. Reduction proofs are a typical way for proving a scheme secure, if this scheme is based on secure primitives. In such a reduction proof we usually assume an adversary against the scheme and construct an adversary against the primitive. While this technique is still possible if a differentially indistinguishable scheme is used in a black-box manner, the quantitative guarantee naturally degenerates if the primitive is used more than once. The sensitivity of the scheme can be derived depending on the number of times the primitive is used within the reduction. The reduction itself is usually assumed to draw uniform randomness, even if the underlying primitive relies on imperfect randomness. If the reduction itself makes use of imperfect randomness, these guarantees can further decrease.

5 Cryptographic Primitives

In this section we instantiate differential indistinguishability with several popular secrecy definitions, namely hiding (for commitment schemes), indistinguishabilily under chosen ciphertext attacks (for encryption schemes) and zero-knowledge (for interactive proofs). We refer the reader to Appendix A.3 for results for pseudorandom generators. These definitions serve as examples for how to instantiate the notion and how to apply our main results to quantify the secrecy loss.

5.1 Commitments

An (interactive) *commitment scheme* $\mathcal{C} = (\mathsf{S}, \mathsf{R})$ is a pair of ppt machines S and R that interact with each other over two phases, called Commit and Open. In the Commit phase, the sender S commits to a message m in the message space $\mathcal{M}$. We write $\mathsf{Commit}\left\langle \mathsf{R}(1^k) \middle| \mathsf{S}(1^k, m) \right\rangle$ to denote the random variable that describes the output of R during the commit phase (in interaction with the sender that commits to m). In the Open phase, the sender S sends m to the recipient R and proves that m indeed is the content of the commitment.

A commitment is *hiding* if the Commit phase does not leak information about the message m. We formalize this intuition with a general hiding notion that allows for a multiplicative secrecy loss.

Definition 4 (ε-differentially hiding commitment) *A commitment scheme $\mathcal{C} = (\mathsf{S}, \mathsf{R})$ over a message space $\mathcal{M}$ is* ε-differentially hiding *for a random source $\{D_{\ell_k}\}_{k \in \mathbb{N}}$, if for all messages $m_0, m_1 \in \mathcal{M}$, for every probabilistic receiver machine R^* interacting with the sender machine S, and for all possible outputs x of R^* during the* Commit *phase,*

$$\Pr\left[\mathsf{Commit}\left\langle \mathsf{R}^*(1^k) \middle| \mathsf{S}(1^k, m_0, D_{\ell_k}) \right\rangle = x\right] \leq e^{\varepsilon} \Pr\left[\mathsf{Commit}\left\langle \mathsf{R}^*(1^k) \middle| \mathsf{S}(1^k, m_1, D_{\ell_k}) \right\rangle = x\right].$$

For $\varepsilon = 0$ and a uniform random source this is a standard definition for *information-theoretically hiding* commitments. With the notion of differential hiding at hand, we can calculate a quantitative guarantee on the security of a information-theoretically hiding commitment scheme if it is used with an SV source.

Theorem 4. *Let $\mathcal{C} = (\mathsf{S}, \mathsf{R})$ be any commitment scheme that is information-theoretically hiding (i.e., 0-hiding) and requires the sender S to use $n_k t_k$ bits of uniform randomness. $\mathcal{C}$ is* ε-differentially hiding *if S uses an (n, γ)-SV randomness source $\{D_{t_k n_k}\}_{k \in \mathbb{N}}$ instead of uniform randomness, where $\gamma_k \leq \min\{1/2,\ \varepsilon/3t_k\}$.*

We refer to Appendix B.2 for a proof. Note that Theorem 4 covers non-interactive commitments. The same analysis holds for *statistically* and *computationally* hiding commitments with the difference that in the former case, we introduce an additive negligible value δ on the right hand side of Definition 4, in the latter case we further only consider ppt adversaries.

Discussion. If a commitment scheme is ε-differentially hiding, the adversary may learn that the probability that a commitment contains a message m_0 is e^{ε} times higher than the probability that it contains another message m_1. However, if ε is reasonably small, e.g., $\varepsilon = 0.001$ (and thus $e^{\varepsilon} \approx 1.001$), both m_0 and m_1 are a plausible content of the commitment. In particular, the adversary cannot reasonably believe or even convince a third party that m_0 is the value that has been committed to. On the other hand, the sender S, no matter if she has committed to m_0 or m_1, retains (a weak form of) deniability: She could indeed have committed to any certain message.

Although Theorem 4 presents guarantees for single commitments, it can be extended for the composition of multiple commitments. We refer to Section 4 for a discussion about composability.

Binding Property with Imperfect Randomness. Whenever Theorem 4 is used to show a scheme to be ε-differentially hiding for an (n, γ)-SV source, the binding property is preserved (with a constant factor of e^{ε}). The reason is that binding is an "unbreakability property" (Definition 3) and thus Lemma 3 can be applied (and the source fulfills it).

5.2 Public-Key Encryption

Differential indistinguishability makes it possible to relax standard security definitions for public-key encryption, e.g., *indistinguishability under adaptive chosen ciphertext attack* (IND-CCA-II) [18].

Definition 5 ((ε, δ)-DIF-IND-CCA-II) *A public-key encryption scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has* (ε, δ)-differentially indistinguishable encryptions under adaptive chosen ciphertext attack *if for every pair of ppt oracle machines $\mathsf{A} = (\mathsf{A}_0, \mathsf{A}_1)$ and for all sufficiently large k and bitstrings z of polynomial length in k, it holds that $\Pr\left[\mathsf{P}_{k,z}^{(0)} = 1\right] < e^{\varepsilon} \Pr\left[\mathsf{P}_{k,z}^{(1)} = 1\right] + \delta_k$, where $P_{k,z}^{(i)}$ is the following* *probabilistic machine:*

$$\begin{aligned}
&\mathsf{P}^{(i)}_{k,z} := (e,d) \leftarrow \mathsf{Gen}(1^k)\\
&((x_0,x_1),s) \leftarrow \mathsf{A}_0^{\mathsf{Dec}(d,\cdot)}(1^k,e,z)\\
&c \leftarrow \mathsf{Enc}(e,x_i)\\
&\textit{return } \mathsf{A}_1^{\mathsf{Dec}_c(d,\cdot)}(1^k,s,c)
\end{aligned}$$

Here, s denotes an internal state of the adversary, and $\mathsf{Dec}_c(d,\cdot)$ denotes an encryption oracle that answers on all ciphertexts except for c, where it returns an error symbol $\bot$.

Note that $(0,\delta)$-DIF-IND-CCA-II security is just IND-CCA-II security if δ is a negligible function.

Encryption with Imperfect Randomness. Both the encryption algorithm and the key generation algorithm require randomness. First, we consider the case that the encryption key is generated with uniform randomness, but the messages are encrypted with an SV source.

Theorem 5. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be any IND-CCA-II secure public-key encryption scheme that requires the encryption algorithm Enc to use at most $n_k t_k$ bits of uniform randomness. $\mathcal{E}$ is $(\varepsilon, e^{\varepsilon}\delta)$-DIF-IND-CCA-II secure if Enc uses an (n,γ)-SV source $D_{n_k t_k}$ instead of uniform randomness, where $\gamma_k \le \min\{1/2,\ \varepsilon/3t_k\}$.*

Although Theorem 5 presents a guarantee for a single encryption, it can be extended for the composition of multiple encryptions. We refer to Section 4 for a discussion about composability. This composability comes with a loss of secrecy (ε increases linearly in the number of compositions).

Key Generation with Imperfect Randomness. We can also give a positive result for the case that the key is generated with an SV source but the encryption algorithm uses proper randomness.

Theorem 6. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be any IND-CCA-II secure public-key encryption scheme that requires the key generation algorithm Gen to use at most $n_k t_k$ uniform random bit. $\mathcal{E}$ is also $(\varepsilon, e^{\varepsilon}\,\delta)$-DIF-IND-CCA-II secure if Gen uses an (n,γ)-SV source instead of uniform randomness, where $\gamma_k \le \min\{1/2,\ \varepsilon/3t_k\}$.*

We refer to Appendix B.3 for the proofs of Theorems 5 and 6.

Theorem 6 implies that imperfect randomness for key generation can be less problematic than imperfect randomness used in the encryption algorithm. The number of messages that are encrypted with a single imperfect key does not increase the randomness that is drawn for generating this key and thus will only influence δ (in the same way it is influenced in the traditional setting), whereas for Theorem 5 our analysis on composability from Section 4 applies. Theorems 5 and 6 can also be combined to yield a result for a scenario where both the key was generated using imperfect randomness and the encryption uses imperfect randomness (naturally the linear loss for composability that applies to Theorem 5 carries over).

Discussion. We have seen that it is still possible to give some guarantees if a IND-CCA-II secure encryption scheme uses imperfect randomness. Encryption schemes suffer from the same composability problems as commitment schemes: the guarantees become much worse if an adversary is allowed to send multiple challenge queries. This renders encryption schemes which use imperfect randomness for encrypting large quantities of messagescompletely insecure if the encrypted messages are not sufficiently independent of each other.

Other Security Definitions and Private-Key Encryption. Although we focus on IND-CCA-II for public-key encryption schemes in this section, our results are analogously applicable to, e.g., information-theoretically secure schemes and of course schemes that are secure under chosen plaintext attack (IND-CPA) as well as schemes that are secure under a priori chosen ciphertext attack (IND-CCA), both in the public-key setting and (under some constraints) in the private-key setting, which we discuss in Appendix A.1 in detail.

5.3 Zero-Knowledge Proofs

We can also relax traditional definitions for zero-knowledge (ZK) proofs. The relaxation applies to the indistinguishability of real views and simulated views: An (ε, δ)-ZK proof system is secure in the sense that the output of a simulator is *almost* indistinguishable, i.e., mutually (ε, δ)-differentially indistinguishable, from the output of a verifier interacting with the real prover. In other words, a distinguisher that has access to the output of the verifier can have only a small chance (quantified by ε) to claim to a third party that an interaction with the real prover has been taken place, i.e., that new knowledge could have been learned at all. For sufficiently small values of ε, such a claim is not convincing at all. For a malicious verifier, that means that everything that has been learned about the witness could have been learned from the simulator with almost the same probability.

An interactive proof system $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ for an $\mathcal{NP}$-language L is a pair of ppt machines P and V that both run on the same input $x \in L$. The prover P gets a witness w from the set $W(x)$ of witnesses for x as additional input, whereas the verifier V gets an auxiliary string z of polynomial length, capturing previous knowledge.

Definition 6 ((ε, δ)-Differential Zero-Knowledge[1]) *A proof system* $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ *is* (ε, δ)-differentially zero-knowledge *for a randomness source* $\{D_{\ell_k}\}_{k\in\mathbb{N}}$ *if for every polynomial* p *and every ppt verifier machine* V^**, there is a ppt machine* S *(the simulator), such that the following distribution ensembles are* (ε, δ)*-differentially indistinguishable in* $k = |x|$ *for all ppt adversaries:*

1. $\{\langle \mathsf{P}\,(x, w, D_{\ell_k}) | \mathsf{V}^*(x, z)\rangle\}_{x\in L, z\in\{0,1\}^{p_k}}$*, i.e., the output of* V^* *(for arbitrary* $w \in W(x)$*)*
2. $\{\mathsf{S}(x, z)\}_{x\in L, z\in\{0,1\}^{p_k}}$

For $\varepsilon = 0$ and a negligible function δ, this is the definition for *computational ZK* [18].

Theorem 7. *Let* $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ *be any proof system that is computationally ZK (i.e.,* (ε, δ)*-ZK for negligible* δ*) and requires the prover to use* $n_k t_k$ *bits of uniform randomness.* $\mathcal{P}$ *is* (ε, δ')*-differentially ZK if the prover* P *uses an* (n, γ)*-SV randomness source instead of a uniform randomness source, where* $\delta'_k = e^{t_k\gamma_k}\delta_k$ *and* $\gamma_k \le \min\{1/2,\ \varepsilon/3t_k\}$*.*

The proof can be found in Appendix B.4. Note that the theorem includes ZK proofs of knowledge because they do not differ from proofs of existence in the ZK property (but only in the existence of an extractor).

Soundness. The *soundness* property is preserved if the proof system uses SV randomness instead of uniform randomness, similar to the binding property of commitments.

Randomness Source of the Simulator. Definition 6 assumes that the simulator has access to uniform randomness. The intuition behind the definition of ZK is that everything that is generated from an interaction with the prover could have been generated without any interaction, using the simulator. Given that uniform randomness is available in general, but the prover does not use it, the same intuition applies if we allow the simulator to access uniform randomness. In contrast, a variant of the definition where the simulator has only access to imperfect randomness, is at least

[1] Note that our definition is distinct from the related concept of ε-knowledge [16], which allows the probabilities of the output bits of a distinguisher to be related by a non-negligible *additive* constant instead of a multiplicative factor.

as strong as Definition 6 if the imperfect randomness source is efficiently sampleable. However, Lemmas 1 and 2 are not applicable to constant-round ZK proof systems that are secure in the traditional sense, i.e., with $\varepsilon = 0$ using a uniform source of randomness. The reason is that for those systems, the bound on the running time of the simulator is not strict, i.e., the simulator only runs in expected polynomial time.[2] Consequently, we cannot bound the amount of randomness drawn by these simulators.

Non-interactive ZK proofs. Our results are also applicable to the case of non-interactive proof systems, in particular to the case, where the common reference string is drawn from a SV source. We refer to Appendix A.2 for a discussion.

References

[1] A. E. Andreev, A. E. F. Clementi, J. D. P. Rolim, and L. Trevisan. Weak random sources, hitting sets, and bpp simulations. *SIAM J. Comput.*, 28(6):2103–2116, 1999.

[2] P. Austrin, K.-M. Chung, M. Mahmoody, R. Pass, and K. Seth. On the (im)possibility of tamper-resilient cryptography: Using fourier analysis in computer viruses. IACR Cryptology ePrint Archive, Report 2013/194, 2013.

[3] B. Barak and Y. Lindell. Strict polynomial-time in simulation and extraction. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC'02)*, pages 484–493. ACM, 2002.

[4] C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy Amplification by Public Discussion. *SIAM J. Comput.*, 17(2):210–229, 1988.

[5] C. Bosley and Y. Dodis. Does privacy require true randomness? In *Proceedings of the 4th Theory of Cryptography Conference (TCC'07)*, pages 1–20. Springer, 2007.

[6] Z. Brakerski, Y. T. Kalai, J. Katz, and V. Vaikuntanathan. Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In *Proceedings of the 51st IEEE Symposium on Foundations of Computer Science (FOCS'10)*, pages 501–510. IEEE, 2010.

[7] F. G. S. L. Brandão, R. Ramanathan, A. Grudka, K. Horodecki, M. Horodecki, and P. Horodecki. Robust Device-Independent Randomness Amplification with Few Devices. arXiv.org, Report arXiv:1310.4544v1 [quant-ph], 2013.

[8] R. Canetti, R. Pass, and A. Shelat. Cryptography from sunspots: How to use an imperfect reference string. In *Proceedings of the 48th IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 249–259. IEEE, 2007.

[9] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. In *Proceedings of the 26th IEEE Symposium on Foundations of Computer Science (FOCS'85)*, pages 429–442. IEEE, 1985.

[10] Y. Dodis, A. Elbaz, R. Oliveira, and R. Raz. Improved Randomness Extraction from Two Independent Sources. In *Proceedings of the 8th International Workshop on Randomization and Computation (RANDOM'04)*. Springer, 2004.

[2] Non-black-box techniques are required to avoid this restriction, as shown in [3].

[11] Y. Dodis, J. Katz, L. Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In *Proceedings of the 26th International Cryptology Conference (CRYPTO'06)*, pages 232–250. Springer, 2006.

[12] Y. Dodis, A. López-Alt, I. Mironov, and S. Vadhan. Differential privacy with imperfect randomness. In *Proceedings of the 32nd International Cryptology Conference (CRYPTO'12)*, pages 497–516. Springer, 2012.

[13] Y. Dodis, S. J. Ong, M. Prabhakaran, and A. Sahai. On the (im)possibility of cryptography with imperfect randomness. In *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science (FOCS'04)*, pages 196–205. IEEE, 2004.

[14] Y. Dodis, K. Pietrzak, and B. Przydatek. Separating sources for encryption and secret sharing. In *Proceedings of the 3rd Theory of Cryptography Conference (TCC'06)*, pages 601–616. Springer, 2006.

[15] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference (TCC'06)*, pages 265–284. Springer, 2006.

[16] C. Dwork, M. Naor, and A. Sahai. Concurrent zero-knowledge. *J. ACM*, 51(6):851–898, 2004.

[17] S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 293–302. IEEE, 2008.

[18] O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Foundations of Cryptography. Cambridge University Press, 2001.

[19] S. Goldwasser, M. Sudan, and V. Vaikuntanathan. Distributed computing with imperfect randomness. In *Proceedings of the 19th International Conference on Distributed Computing (DISC'05)*, pages 288–302. Springer, 2005.

[20] S. Kamara and J. Katz. How to encrypt with a malicious random number generator. In *Proceedings of the 15th International Workshop on Fast Software Encryption (FSE'08), Revised Selected Papers*, pages 303–315. Springer, 2008.

[21] J. Kamp and D. Zuckerman. Deterministic Extractors for Bit-Fixing Sources and Exposure-Resilient Cryptography. In *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science (FOCS'03)*, pages 92–101. IEEE, 2003.

[22] J. Katz and Y. Lindell. *Introduction to modern cryptography*. CRC Press, 2007.

[23] U. M. Maurer and S. Wolf. Privacy amplification secure against active adversaries. In *Proceedings of the 17th International Cryptology Conference (CRYPTO'97)*, pages 307–321. Springer, 1997.

[24] J. L. McInnes and B. Pinkas. On the impossibility of private key cryptography with weakly random keys. In *Proceedings of the 10th International Cryptology Conference (CRYPTO'90)*, pages 421–435. Springer, 1990.

[25] I. Mironov, O. Pandey, O. Reingold, and S. Vadhan. Computational Differential Privacy. In *Proceedings of the 29th International Cryptology Conference (CRYPTO'09)*. Springer, 2009.

[26] T. Moran and M. Naor. Receipt-free universally-verifiable voting with everlasting privacy. In *Proceedings of the 26th International Cryptology Conference (CRYPTO'06)*, pages 373–392. Springer, 2006.

[27] O. Reingold, L. Trevisan, M. Tulsiani, and S. Vadhan. Dense Subsets of Pseudorandom Sets. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science (FOCS'08)*, pages 76–85. IEEE, 2008.

[28] M. Santha and U. V. Vazirani. Generating quasi-random sequences from slightly-random sources. In *Proceedings of the 25th IEEE Symposium on Foundations of Computer Science (FOCS'84)*, pages 434–440. IEEE, 1984.

[29] L. Trevisan and S. Vadhan. Extracting randomness from samplable distributions. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 32–42. IEEE, 2000.

[30] J. Von Neumann. Various techniques used in connection with random digits. *National Bureau of Standards Applied Mathematics Series*, 12:36–38, 1951.

[31] D. Zuckerman. General weak random sources. In *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science (FOCS'90)*, pages 534–543. IEEE, 1990.

A Other Cryptographic Primitives

A.1 Private-Key Encryption

In the case of private-key encryption, the adversary usually has access to an encryption oracle. For analyzing the influence of imperfect randomness on such schemes, the following strategies are possible. One can assume that the encryption oracle has access to uniform randomness or, one can restrict the number of encryption queries, or one can use stronger notions.

Uniform encryption oracles. If the encryption oracle uses uniform randomness, the analysis from Section 5.2 can be applied without any need for modification. This guarantee corresponds to a (maybe unrealistic) scenario where an adversary might see several encryptions, but only the randomness used to encrypt a single message is not uniform.

Restricting the queries. An adversary that makes at most q_k queries to the encryption oracle for a given security parameter k can be covered by an analysis much similar to the one from 5.2. However, now the machines X and X_1 use $(q_k + 1) \cdot n_k t_k$ many bits from the randomness source to answers all q_k oracle queries and to compute the challenge ciphertext. Consequently, a private-key encryption scheme is ε-secure against adversaries that only make q_k oracle queries when using an (n, γ)-SV source, if $\gamma_k \leq \min\{1/2,\ \varepsilon/3t_k q_k\}$.

Stronger notions. Another possibility is to rely on security definitions that are explicitly introduced to handle improper randomness, e.g., *indistinguishability against chosen randomness attack* (CRA) [20]. In a private-key encryption that is secure under this definition, the adversary can choose the randomness for the encryption queries Enc on its own (and thus the encryption oracle does not need to sample randomness). Originally, the CRA challenger uses uniform randomness to encrypt the challenge message (which is the opposite to our first strategy). We can show now that any CRA secure scheme is also ε-CRA, if the challenge is encrypted using an (n, γ)-SV distribution, if $\gamma \leq \min\{1/2,\ \varepsilon/3t_k\}$.

A.2 Non-interactive Zero-Knowledge Proofs

Similar to interactive zero-knowledge proofs (Section 5.3), a differential relaxation is also possible for the security definition of non-interactive zero-knowledge proofs.

Definition 7 (Non-interactive (ε, δ)-differential-zero-knowledge) *A non-interactive proof system* $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ *is* single-theorem adaptive (ε, δ)-differential zero-knowledge *if there exists a polynomially bounded simulator machine* $\mathsf{S} = (\mathsf{S}_1, \mathsf{S}_2)$ *such that for every function* F *(which is supposed to get the CRS* σ *and select a statement* x *and a witness* $w \in W(x)$ *adaptively) the following two ensembles are* (ε, δ)*-differentially indistinguishable.*

1. $\{(\sigma, F(\sigma), \pi) \mid \pi \leftarrow \mathsf{P}(x, w), (x, w) \leftarrow F(\sigma), \sigma \leftarrow U_{t_k n_k}\}_{k \in \mathbb{N}}$
2. $\left\{(\sigma, F(\sigma), \pi) \;\middle|\; \pi \leftarrow \mathsf{S}_2(x, s), (x, w) \leftarrow F(\sigma), (\sigma, s) \leftarrow \mathsf{S}_1(1^k)\right\}_{k \in \mathbb{N}}$

Note that, for the sake of simplicity, we consider adaptive single-theorem definitions, i.e., the CRS can only be used once. Additionally, we do not consider auxiliary input that is available to the adversary. It is straight-forward to extend our results to a variant with auxiliary input as well as to the multi-theorem setting. In the latter, the security guarantees decrease similar as described in Section 4 if the prover (aside from the CRS) uses imperfect randomness.

Theorem 8 (δ-NIZK with $\{U_{t_k n_k}\}_{k \in \mathbb{N}} \implies (\varepsilon, \delta')$-NIZK with $\{D_{t_k n_k}\}_{k \in \mathbb{N}}$) *Let* P *be a prover machine that is* δ*-zero-knowledge if it uses the uniform source* $\{U_{t_k n_k}\}_{k \in \mathbb{N}}$ *on* $\{0,1\}^{t_k n_k}$ *and the CRS* *is generated based on input from a uniform source. If* P *uses a* (n,γ)*-SV-distribution* $\{D_{t_k n_k}\}_{k\in\mathbb{N}}$ *(with* $\gamma(k) < \frac{1}{2}$*) over* $\{0,1\}^{t_k \cdot n_k}$*, and the CRS has been generated based on distinct random bits from the same source,*[3] *then* P *is* (ε,δ')*-zero-knowledge with* $\delta'_k = e^{t_k\gamma_k}\delta_k$*.*

Proof. The proof is analogous to the proof of Theorem 7. □

A.3 Pseudorandom Generators

Differential indistinguishability can also be instantiated by the definition of pseudorandomness. For the definitions of pseudorandom generators we follow the notation of [22].

Definition 8 ((ε,δ)-Pseudorandomness) *A source* $\{D_{\ell_k}\}_{k\in\mathbb{N}}$ *is* (ε,δ)*-pseudorandom, if it is* (ε,δ)*-indistinguishable from* $\{U_{\ell_k}\}_{k\in\mathbb{N}}$*. We call* $\{D_{\ell_k}\}_{k\in\mathbb{N}}$ ε-pseudorandom*, if it is* (ε,δ')*-pseudorandom for a negligible function* δ'*.*

Definition 9 (Pseudorandom generators) *Let* ℓ_{out} *be a positive polynomial and let* G *be a deterministic polynomial-time machine, where for all* x *it holds that* $|G(x)| = \ell_{out}(|x|)$*.* G *is a* (ε,δ)-pseudorandom generator *for a distribution* $\{D_{\ell_k}\}_{k\in\mathbb{N}}$*, if*

1. *for every* $k \in \mathbb{N}$*, it holds that* $\ell_{out}(k) > \ell(k)$*.*
2. $\{G_k(D_k)\}_{k\in\mathbb{N}}$ *is* (ε,δ)*-pseudorandom.*

Theorem 9. *If* G *is a pseudorandom generator for the uniform source* $\{U_{n_k t_k}\}_{k\in\mathbb{N}}$ *over* $\{0,1\}^{n_k t_k}$ *and* $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$ *is a* (n,γ)*-SV distribution, then* G *is a* ε*-pseudorandom generator on* $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$*, where* $\gamma_k \le \min\{1/2,\ \varepsilon/3t_k\}$*.*

Proof. Let G be a pseudorandom generator on the uniform distribution $\{U_{n_k t_k}\}_{k\in\mathbb{N}}$ and let $\{D_{n_k t_k}\}_{k\in\mathbb{N}}$ be a (n,γ)-SV distribution. Let $\mathsf{X}_0(1^k, r) := G(r)$ and let $\mathsf{X}_1(1^k, \cdot) := U_{\ell_k}$. Since G is pseudorandom, there is a negligible function δ such that G on uniform input is δ-indistinguishable from U_{ℓ_k}.

(i) G is pseudorandom if and only if X_0 and X_1 are indistinguishable.

(ii) G is (ε,δ)-pseudorandom if and only if X_0 and X_1 are (ε,δ)-differentially indistinguishable.

The theorem directly follows from Lemma 2 and from the fact that for all constant ε and all negligible functions δ, the function $\delta'_k := e^{\varepsilon}\delta_k$ is still negligible. □

B Postponed Proofs

B.1 On Additive and Multiplicative Bounds

No matter which additive bound can be shown between two probabilistic events, there does not necessarily exist a multiplicative bound. In particular, there are machines that are δ-indistinguishable for some δ but not (δ',ε)-indistinguishable for any ε and $\delta' < \delta$. That means, e.g, that for every for $0 \le \delta \le 1$ there is a commitment scheme that is δ-hiding but not (ε,δ')-differentially hiding for any value $\delta' < \delta$.

[3] Note that the theorem also holds if the SV sources of the prover and of the trusted party that generates the CRS are independent, because the combination of sources can be considered as one single source.

Proof. Given any arbitrary function δ with $1 \geq \delta_k > 0$, we construct a counter example, i.e., a commitment scheme $\mathcal{C}$ such that for every adversary there is an additive bound of δ ($\mathcal{C}$ is δ-hiding), but there is no pair (ε, δ') with $\delta'_k < \delta_k$ (for sufficiently large k) such that $\mathcal{C}$ is (ε, δ')-differentially hiding.

Let $\mathcal{C}_{IT}$ be an information theoretically hiding commitment scheme. We construct $\mathcal{C} = (\mathsf{S}, \mathsf{R})$ from $\mathcal{C}_{IT}$ as follows. For security parameter k, $\mathcal{C}$ behaves like $\mathcal{C}_{IT}$ but with probability δ_k, the algorithm S additionally leaks the message. Clearly the scheme is δ-hiding. Consider the distinguisher A that sends two messages m_0, m_1 to the challenger for the hiding game. Only if S leaks m_0, A outputs 0. In all other cases, A outputs 1. Let $\varepsilon \geq 0$ and δ be functions with $\delta'_k < \delta_k$ for sufficiently large k. For such k,

$$\begin{aligned}
\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{S}(1^k, m_0) \right\rangle = 0\right] &= \delta_k \\
&> \delta'_k \\
&= e^{\varepsilon_k}\, 0 + \delta'_k \\
&= e^{\varepsilon_k} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{S}(1^k, m_1) \right\rangle = 0\right] + \delta'_k.
\end{aligned}$$

Consequently, $\mathcal{C}$ is not (ε, δ')-differentially hiding. □

B.2 Proof of Theorem 4 (Commitments)

Proof. Let $\mathcal{C} = (\mathsf{S}, \mathsf{R})$ be an information-theoretically hiding interactive commitment scheme over the message set $\mathcal{M}$ such that the sender S uses at most $\{0,1\}^{n_k t_k}$ random bits. Let $\{D_{n_k t_k}\}_{k \in \mathbb{N}}$ be an (n, γ)-SV-randomness distribution over $\{0,1\}^{n_k t_k}$ and $U_{n_k t_k}$ be the uniform distribution over $\{0,1\}^{n_k t_k}$. Furthermore, let $m_0, m_1 \in \mathcal{M}$ be two arbitrary messages, and let $\mathcal{A}$ be the set of all probabilistic (not necessarily polynomially bounded) machines.

We define $\mathsf{X}_0(1^k, r) := \mathsf{S}(1^k, m_0, r)$ and $\mathsf{X}_1(1^k, r) := \mathsf{S}(1^k, m_1, r)$. Observe that by our definition of X and X_1, the following two statements hold:

(i) X_0 and X_1 are indistinguishable for the class $\mathcal{A}$ of adversaries and only if $\mathcal{C}$ is information-theoretically hiding.

(ii) X_0 and X_1 are $(\varepsilon, 0)$-differentially indistinguishable for the class $\mathcal{A}$ of adversaries if and only if $\mathcal{C}$ is ε-differentially hiding.

Thus, the theorem follows immediately from Lemma 2. □

B.3 Proof of Theorem 5 (Public-Key Encryption)

Proof (of Theorem 5). Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a public-key encryption scheme, let $\mathcal{A}_{\mathsf{poly}}$ be the class of pairs of probabilistic polynomial-time machines with decryption oracles (Definition 5), and let $\{D_{n_k t_k}\}_{k \in \mathbb{N}}$ be an (n, γ)-SV distribution family. To simplify the notation we write $P^{(b,r)}_{k,z}$ for simulating $P^{(b)}_{k,z}$ and using $r \in \{0,1\}^{n_k t_k}$ as the randomness for Enc. Let $\mathsf{X}_0(1^k, r) := P^{(0,r)}_{k,z}$ and $\mathsf{X}_1 := P^{(1,r)}_{k,z}$. The rest of the proof is analogous to the proof of Theorem 4. Observe that by our definition of X_0 and X_1, the following two statements hold:

(i) X_0 and X_1 are indistinguishable for the class $\mathcal{A}_{\mathsf{poly}}$ of adversaries if and only if $\mathcal{E}$ is $\mathsf{IND\text{-}CCA\text{-}II}$.

(ii) X_0 and X_1 are (ε, δ)-differentially indistinguishable for the class $\mathcal{A}_{\mathsf{poly}}$ of adversaries if and only if $\mathcal{E}$ is (ε, δ)-$\mathsf{DIF\text{-}IND\text{-}CCA\text{-}II}$.

Thus, the claim follows immediately from Lemma 2. □

The proof of Theorem 6 is analogous to the previous proof.

B.4 Proof of Theorem 7 (Zero-Knowledge Proofs)

Proof. Let A be a machine in the class $\mathcal{A}_{\mathsf{poly}}$ of all probabilistic polynomial-time adversaries. Further, let F be a function that maps each security parameter k to a triple (x, w, z) consisting of a statement $x \in L$ with $|x| = k$, a corresponding $w \in W(x)$, and a auxiliary string z of p_k for a polynomial p.

We define machines $\mathsf{X}_0(1^k, r)$ and $\mathsf{X}_1(1^k, r)$ as follows: Both X_0 and X_1 use $F(1^k)$ to generate a triple (x, w, z). $\mathsf{X}_0(1^k, r)$ runs $\mathsf{V}^*(x, z, r)$ in $\langle \mathsf{P}(x, w) | \mathsf{V}^*(x, z) \rangle$ and sends the output of $\mathsf{V}^*(x, z)$, whereas $\mathsf{X}_1(1^k, r)$ ignores r, runs $\mathsf{S}(x, z)$ and sends its output. (Recall that the simulator has access to uniform randomness.)

Observe that X_0 and X_1 are (ε, δ)-indistinguishable if and only if (P, V) are (ε, δ)-zero-knowledge. In particular, the goal of the polynomially bounded A is to distinguish between the machines X_0 and X_1, which simulate $\{\langle \mathsf{P}\left(x, w; D_{\ell_k}\right) | \mathsf{V}^*(x, z) \rangle\}_{x \in L, z \in \{0,1\}^{p_k}}$ and $\{\mathsf{S}(x, z)\}_{x \in L, z \in \{0,1\}^{p_k}}$, respectively. Note that A has access to the statement x and the auxiliary string z, because it can be contained in the output of $\mathsf{V}^*(x, z)$.

Thus, Lemma 2 implies the claim. Note that F might not be computable. However, it can be verified that Lemma 2 as well as the underlying Lemma 1 hold even in the case that the adversary has to distinguish between the outputs of general functions. We have chosen to present the current formulation to stay consistent with common notions. □