Secrecy without Perfect Randomness: Cryptography with (Bounded) Weak Sources

Michael Backes, Aniket Kate, Sebastian Meiser, and Tim Ruffing

CISPA, Saarland University, Germany
{backes@cs, aniket@mmci, meiser@cs, tim.ruffing@mmci}.uni-saarland.de

Abstract. Cryptographic protocols are commonly designed and their security proven under the assumption that the protocol parties have access to perfect (uniform) randomness. Physical randomness sources deployed in practical implementations of these protocols often fall short in meeting this assumption, but instead provide only a steady stream of bits with certain high entropy. Trying to ground cryptographic protocols on such imperfect, weaker sources of randomness has thus far mostly given rise to a multitude of impossibility results, including the impossibility to construct provably secure encryption, commitments, secret sharing, and zero-knowledge proofs based solely on a weak source. More generally, indistinguishability-based properties break down for such weak sources. In this paper, we show that the loss of security induced by using a weak source can be meaningfully quantified if the source is *bounded*, e.g., for the well-studied Santha-Vazirani (SV) sources. The quantification relies on a novel relaxation of indistinguishability by a quantitative parameter. We call the resulting notion *differential indistinguishability* in order to reflect its structural similarity to differential privacy. More concretely, we prove that indistinguishability with uniform randomness implies differential indistinguishability with weak randomness. We show that if the amount of weak randomness is limited (e.g., by using it only to seed a PRG), all cryptographic primitives and protocols still achieve differential indistinguishability.

Keywords: indistinguishability, randomness, weak sources, differential privacy, pseudorandom generators, Santha-Vazirani sources

1 Introduction

Cryptographic protocols are commonly designed and their security proven under the assumption that the protocol parties have access to perfect, i.e., uniform, randomness. Actual physical randomness sources that cryptographic implementations rely on, however, rarely meet this assumption: instead of providing uniform randomness, they provide only a stream of bits with a certain high amount of

This is the full version of a paper appearing at the 13th International Conference on Applied Cryptography and Network Security (ACNS 2015).

entropy. Moreover, these so-called *weak sources*, such as the Santha-Vazirani (SV) sources [33], are often non-extractable [15, 33], i.e., it is computationally infeasible to extract more than a super-logarithmic amount of (almost) uniform randomness from them.

There have been several attempts to bridge this gap, i.e., to ground the security guarantees of cryptographic systems on such weak sources. As soon as indistinguishability-based secrecy properties are being desired, however, this line of research has mostly given rise to a multitude of impossibility results [7, 15, 30], only complemented by a few constructive results if additional assumptions are being imposed. For instance, encryption can be realized using weak sources, if one imposes strong assumptions on the entropy of encrypted messages [5], or if the weak source is restricted to the key generation algorithm and a perfect source is available for the actual encryption algorithm [18]. The plurality of impossibility results in this area, as well as the absence of comprehensive constructive results, indicates that traditional indistinguishability-based secrecy notions fall short in capturing the impact of weak randomness on cryptography. This constitutes an unsatisfactory situation, with several open questions looking for an answer:

- Is it possible to quantify the secrecy loss of cryptographic operations and primitives, if a weak source (such as an SV source) is being used?
- Imagine that today a cryptographic protocol (e.g., an e-voting system) is executed and tomorrow it turns out that the employed randomness was weak. Given that there are strong impossibility results [7, 15, 30] for indistinguishability, is all lost or can we still give quantitative guarantees about the secrecy of the system?
- Given that these quantitative guarantees will necessarily be weaker than traditional cryptographic guarantees, under which assumptions do they still provide reasonable practical security guarantees?

In this paper we address all of these questions.

1.1 Our Contributions

Relaxing Indistinguishability to Quantify the Secrecy Loss. We derive quantitative guarantees for all indistinguishability-based cryptographic constructions that are used with arbitrary weak sources that are additionally *bounded* in the following sense: in addition to imposing an upper bound on the probability of each individual bitstring (i.e., requiring a sufficiently high min-entropy), one additionally imposes a lower bound on these probabilities. These *bounded weak sources* include SV sources [33] and resemble balanced sources [24].

To quantify the secrecy loss that weak randomness imposes on cryptography, we define *differential indistinguishability*, a quantitative relaxation of cryptographic indistinguishability in the spirit of differential privacy [19, 31] and pseudodensity [32]. The necessity of a new, relaxed notion arises from the impossibility result of Dodis et al. [15] who showed that whenever only weak sources of randomness are available, traditional indistinguishability is provably impossible for cryptographic primitives that have a secrecy requirement, e.g., encryption, commitments, and zero-knowledge proofs. More concretely, one cannot ensure that the advantage in distinguishing two challenger machines X_0 and X_1 is negligible for every probabilistic polynomial-time adversary. However, it might still be the case that no adversary has a non-negligible advantage in performing a practical attack that breaks the security *entirely*, e.g., by reaching a state in which it is *certain* whether it interacts with X_0 or X_1. The notion of differential indistinguishability consequently aims at quantifying the resulting loss of secrecy without overestimating the adversary's power to break the scheme entirely: Two games, i.e., 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 x are related by

$$\Pr\left[\langle \mathsf{A} | \mathsf{X}_0 \rangle = x\right] \leq 2^{\varepsilon} \cdot \Pr\left[\langle \mathsf{A} | \mathsf{X}_1 \rangle = x\right] + \delta,$$

where x is a possible output of A.[1] Here $\varepsilon \geq 0$ is a reasonably small constant or a decreasing function such as $1/p(\cdot)$ for a polynomial p. We allow only a negligible function for δ, which corresponds to a negligible probability to break the security of the scheme entirely. Differential indistinguishability thus offers quantitative parameters to reason about the loss of secrecy incurred by the use of imperfect randomness.

Guarantees for Cryptographic Primitives Using Weak Sources. As our main contribution we show that traditional indistinguishability (given a uniform randomness source) suffices to guarantee differential indistinguishability if the uniform source is replaced by an arbitrary bounded weak source. This result immediately entails meaningful quantitative lower security bounds in cases where indistinguishability-based definitions are provably impossible to achieve [15].

In particular, our methodology can be applied in hindsight and produces meaningful quantitative guarantees for all cryptographic primitives and protocols, provided that the amount of used imperfect randomness is bounded; there is no need for new cryptographic constructions for any of the existing primitives whose security is defined and proven by means of indistinguishability, including simulator-based notions. Moreover, we show that if the bounded weak randomness is used only to seed a secure PRG, differential indistinguishability suffers only a negligible quantitative (additional) security loss under composition – just as traditional indistinguishability.

Intuitively, is not surprising that the provided secrecy does not degrade substantially if the quality of the randomness degrades within certain small bounds, because otherwise virtually all practical implementations of cryptography would be insecure due to the inherent imperfection of physical sources. Our work confirms this intuition and provides a framework to analyze the resulting loss of secrecy quantitatively.

[1] In contrast to differential privacy and pseudodensity, we use 2 instead of e as a base for the exponential function, because the base 2 fits standard definitions of entropy better.

Technically, Theorem 1 states that the interactions with two machines X_0 and X_1 are differentially indistinguishable for bounded weak distributions if they are indistinguishable for the uniform distribution. These machines X_0 and X_1 can then be instantiated by arbitrary challenger machines to immediately derive results for cryptographic notions. Theorem 1 comprises arbitrary classes of adversaries and thus covers information-theoretical and computational indistinguishability. To derive quantitative guarantees, the theorem only imposes the requirement that the entropy of the bounded weak randomness used by the primitive or protocol is bounded in terms of the security parameter. Thus all existing primitives that use a bounded amount of randomness can immediately be analyzed and their secrecy loss quantified by an additional multiplicative factor that only depends on the quality of the random source.

Connection to Differential Privacy. We analyze the relation between differential indistinguishability and the well-studied notion of differential privacy [19,31], especially in terms of composition. Similar to the privacy loss in differential privacy when the privacy of several users is analyzed, differential indistinguishability suffers from a commensurate loss of entropy, which consequently leads to a secrecy loss in cases where several users use weak, potentially even dependent randomness. This relation is of particular interest in scenarios in which the users are not aware of using imperfect randomness and thus fail to deploy existing methods [12,25,27] to improve their randomness using multiple sources.

Organization. The rest of the paper is organized as follows: We recall important concepts and introduce our notation in Section 2. We define differential indistinguishability and present our main results in Section 3. We then demonstrate the utility of differential indistinguishability to public-key encryption and study composability of differentially indistinguishable primitives in Section 4. We interpret and analyze differential indistinguishability in Section 5, including a comparison between differential indistinguishability with differential privacy. Finally, we discuss related work in Section 6 and possible future directions in Section 7. To improve readability, we have shifted several proofs to Appendix A.

2 Preliminaries and Notation

We denote sampling an element r from a distribution D by $r \leftarrow D$. The probability of the 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 write f_k for the value of a function $f(\cdot)$ applied to k, where k is typically the security parameter. We drop the explicit dependence of parameters and security bounds $(\alpha, \beta, \varepsilon, \gamma)$ on k whenever it is clear from the context. We denote by $\{D_k\}_{k\in\mathbb{N}}$ a family of distributions such that for each $k \in \mathbb{N}$ the distribution D_k samples elements from $\{0,1\}^k$. In particular, $\{U_k\}_{k\in\mathbb{N}}$ is the family of uniform distributions, where U_k is the uniform distribution over $\{0,1\}^k$.

Throughout the paper we consider (possibly interactive) Turing machines X that always have implicitly access to a *random tape with an infinite sequence of* *uniformly distributed random bits*, even if the machines get an additional input drawn from some random source. Unless we mention that they run in probabilistic polynomial time (ppt) in the length of their first input, those machines are *not bounded*. The distribution on the outputs of X when run on input x is denoted by $\mathsf{X}(x)$. Similarly, we write $\langle \mathsf{X}(x) | \mathsf{Y}(y) \rangle$ to denote the distribution on the output of the machine X on input x in an interaction with the machine Y on input y. We write $\log := \log_2$ for the logarithm to base 2.

Randomness Sources. In addition to the commonly used min-entropy, we make use of a symmetrically defined counterpart, coined *max-entropy* by Haitner et. al. [24]: whereas min-entropy bounds the maximum likelihood event, max-entropy bounds the minimum likelihood event (and consequently requires probability distributions with full support).[2]

Definition 1. *Let D be a distribution over the set S. The* min-entropy *of D is $H_{min}(D) := \min_{x \in S}(-\log \Pr[D = x])$; the* max-entropy *of D is $H_{max}(D) := \max_{x \in S}(-\log \Pr[D = x])$.*

These entropy measures allow us to define *bounded weak* sources, which must additionally provide a certain amount of max-entropy in comparison to weak sources.

Definition 2. *A family of distributions $\{D_n\}_{n \in \mathbb{N}}$, each over the set $\{0,1\}^n$ of bitstrings of length n, is a (α, β)-*bounded weak source*, if every D_n satisfies the following entropy requirements:*

(i) D_n has min-entropy at least $n - \alpha$, and
(ii) D_n has max-entropy at most $n + \beta$.

*If a family of distributions $\{D_n\}_{n \in \mathbb{N}}$ satisfies only requirement (i), but not requirement (ii), we call it an α-*weak source (or a min-entropy source) *instead.*

The following generalization of Santha-Vazirani (SV) sources [33] to block sources [11,15] is a special case of (α, β)-bounded weak sources. 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 [1].

Definition 3 (SV 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 $x_1, \ldots, x_i \in \{0,1\}^n$,*

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

The original SV sources are a special case of Definition 3 that arises for $n = 1$. Every (n, γ)-SV block source over $\{0,1\}^{tn}$ is an (α, β)-bounded weak source where $\alpha = t \cdot \log(1 + \gamma)$ and $\beta = -t \cdot \log(1 - \gamma)$.

Remark 1. Our complete analysis is also possible for sources that are only *statistically close* to (α, β)-bounded weak sources such as sources in [24] that have a limited number of outliers. We refer to Appendix D for both definitions and results for such sources.

[2] This notion of *max-entropy* is not to be confused with *Hartley entropy*, which is also sometimes called max-entropy.

3 Differential Indistinguishability

In this section we present our main results, which can be applied to a variety of cryptographic notions. 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 requiring that any possible view of a distinguisher is (almost) equally likely for both X_0 and X_1, i.e., the difference between the probability that A outputs any given view in the interaction with X_0 and the probability that A outputs the same view in the interaction with X_1 is negligible. We consider a variant of indistinguishability that allows these probabilities to be also related by a multiplicative factor $2^\varepsilon > 1$, similar to the concept of mutual pseudodensity [32] and differential privacy [19, 31].

Definition 4 (Differential Indistinguishability). *Two probabilistic machines X_0 and X_1 are (ε,δ)-differentially indistinguishable for a distribution $\{D_\ell\}_{\ell\in\mathbb{N}}$ over $\{0,1\}^\ell$ 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 $b \in \{0,1\}$,*

$$\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_b(1^k, D_\ell)\right\rangle = x\right] \le 2^\varepsilon \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_{1-b}(1^k, D_\ell)\right\rangle = x\right] + \delta_k.$$

This definition allows to express many of the traditional cryptographic indistinguishability notions [22, 28]. We discuss the impact of the multiplicative factor, that can (and must) be interpreted carefully, in Section 5. 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 perfect (information-theoretic) indistinguishability, the class of adversaries is the class $\mathcal{A}_\infty$ of all probabilistic (possibly unbounded) machines and we have $\delta = 0$.[3] Statistical indistinguishability can be expressed with the same class of adversaries for $\delta > 0$. Cryptographic (computational) indistinguishability can be achieved with the class $\mathcal{A}_{ppt}$ of ppt machines with δ being a negligible function.[4]

3.1 Main Result

Traditional indistinguishability for uniform randomness directly implies differential indistinguishability for (α, β)-bounded weak sources. This is captured by the following theorem. It allows us to easily give guarantees for cryptographic primitives whenever their security notions can be expressed in terms of Definition 4.

Theorem 1. *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 $\{U_n\}_{n\in\mathbb{N}}$ over $\{0,1\}^n$, then X_0 and X_1 are also $(\alpha+\beta, 2^\alpha \cdot \delta)$-differentially indistinguishable for $\mathcal{A}$ and any (α, β)-bounded weak source over $\{0,1\}^n$.*

[3] We additionally drop the formulation "for sufficiently large k" in the case of information-theoretic security.

[4] Note that this is equivalent to requiring a negligible function for every adversary [4].

Proof. We show the theorem by first proving a technical lemma about bounded weak distributions: Even though an (α, β)-bounded weak distribution is not negligibly close to a uniform distribution, the parameters α and β give a bound on the discrepancy between the uniform distribution and the bounded weak distribution.

Lemma 1. *Let $\{D_n\}_{n\in\mathbb{N}}$ be an (α, β)-bounded weak source over $\{0,1\}^n$ and let $\{U_n\}_{n\in\mathbb{N}}$ be a family of uniform sources over $\{0,1\}^n$. For all probabilistic machines A, for all $k \in \mathbb{N}$ and for all possible outputs x of A,*

$$\Pr\left[\mathsf{A}(1^k, D_n) = x\right] \le 2^{\alpha} \Pr\left[\mathsf{A}(1^k, U_n) = x\right] \tag{a}$$

$$\text{and} \qquad \Pr\left[\mathsf{A}(1^k, U_n) = x\right] \le 2^{\beta} \Pr\left[\mathsf{A}(1^k, D_n) = x\right]. \tag{b}$$

Proof. Let $\{D_n\}_{n\in\mathbb{N}}$ be an (α, β)-bounded weak distribution over $\{0,1\}^n$. By Definition 2, D_n has min-entropy at least $n - \alpha$ and max-entropy at most $n + \beta$. We start with (a). For all values $r_0 \in \{0,1\}^n$,

$$\begin{aligned}\log\left(\frac{\Pr[D_n = r_0]}{\Pr[U_n = r_0]}\right) &= \log\left(\Pr[D_n = r_0]\right) - \log\left(2^{-n}\right)\\ &\le -\min_{y\in\{0,1\}^n}\left(-\log\left(\Pr[D_n = y]\right)\right) - \log\left(2^{-n}\right)\\ &\le -(n-\alpha) + n = \alpha.\end{aligned}$$

Using this inequality we can show (a) as follows. For all possible outputs x of A,

$$\begin{aligned}\Pr\left[\mathsf{A}(1^k, D_n) = x\right] &= \sum_{r_0\in\{0,1\}^n} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \Pr[D_n = r_0]\\ &\le \sum_{r_0\in\{0,1\}^n} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \cdot 2^{\alpha} \cdot \Pr[U_n = r_0]\\ &\le 2^{\alpha} \Pr\left[\mathsf{A}(1^k, U_n) = x\right].\end{aligned}$$

This shows (a). For (b), note that for all values $r_0 \in \{0,1\}^n$, the probability $\Pr[D_n = r_0]$ is strictly larger than zero because $\beta < \infty$. For all values $r_0 \in \{0,1\}^n$,

$$\begin{aligned}\log\left(\frac{\Pr[U_n = r_0]}{\Pr[D_n = r_0]}\right) &= \log\left(2^{-n}\right) - \log\left(\Pr[D_n = r_0]\right)\\ &\le \log\left(2^{-n}\right) + \max_{y\in\{0,1\}^n}\left(-\log\left(\Pr[D_n = y]\right)\right)\\ &\le -n + (n+\beta) = \beta.\end{aligned}$$

Using this equation we can show (b) as follows. For all possible outputs x of A,

$$\begin{aligned}\Pr\left[\mathsf{A}(1^k, U_n) = x\right] &= \sum_{r_0\in\{0,1\}^n} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \Pr[U_n = r_0]\\ &\le \sum_{r_0\in\{0,1\}^n} \Pr\left[\mathsf{A}(1^k, r_0) = x\right] \cdot 2^{\beta} \cdot \Pr[D_n = r_0]\\ &\le 2^{\beta} \Pr\left[\mathsf{A}(1^k, D_n) = x\right].\end{aligned}$$

This completes the proof of Lemma 1. □

Now we use the lemma to prove our main theorem. Let $\{D_n\}_{n\in\mathbb{N}}$ be an (α,β)-bounded weak source, and $\{U_n\}_{n\in\mathbb{N}}$ be the uniform source, both over $\{0,1\}^n$. 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,

$$\Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_0(1^k,U_n)\right\rangle = x\right] \leq \Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_1(1^k,U_n)\right\rangle = x\right] + \delta.$$

Using Lemma 1, we show that A behaves similarly on D_n, 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 $\{D_n\}_{n\in\mathbb{N}}$ and $\{U_n\}_{n\in\mathbb{N}}$.

$$\begin{aligned}\Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_0(1^k,D_n)\right\rangle = x\right] &\leq 2^{\alpha}\Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_0(1^k,U_n)\right\rangle = x\right] && (1)\\ &\leq 2^{\alpha}\Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_1(1^k,U_n)\right\rangle = x\right] + 2^{\alpha}\cdot\delta && (2)\\ &\leq 2^{\alpha+\beta}\Pr\left[\left\langle \mathsf{A}(1^k)\middle|\mathsf{X}_1(1^k,D_n)\right\rangle = x\right] + 2^{\alpha}\cdot\delta && (3)\end{aligned}$$

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

Recall that every (n,γ)-SV block source over $\{0,1\}^{tn}$ (Definition 3) is an (α,β)-bounded weak source where $\alpha = t\cdot\log(1+\gamma)$ and $\beta = -t\cdot\log(1-\gamma)$. With $\gamma < 1/2$, it holds that $\beta \leq 2t\gamma$ and $\alpha \leq 2t\gamma$. Thus, we can instantiate Theorem 1 for SV block sources as follows:

Corollary 1. *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 $\{U_{nt}\}_{nt\in\mathbb{N}}$ over $\{0,1\}^{nt}$, then X_0 and X_1 are also $(\varepsilon, 2^{\varepsilon}\delta)$-differentially indistinguishable for $\mathcal{A}$ and any family of (n,γ)-SV block sources $\{D_{nt}\}_{nt\in\mathbb{N}}$ over $\{0,1\}^{tn}$ with $\gamma\leq\frac{1}{2}$, where $\varepsilon=\gamma\cdot 4t$.*

Remark 2. Lemma 1 can also be interesting for sources with unbounded max-entropy. In this case, β is infinitely large and consequently, inequality (b) does not yield interesting guarantees anymore. However, for restricting undesirable events that are not based on indistinguishability, inequality (a) suffices, which is in line with the results of Dodis and Yu [18]. We refer to Appendix B for a discussion.

3.2 Computational Differential Indistinguishability Guarantees

In the computational setting where adversaries are ppt machines, we can achieve a stronger result: If we rely on a pseudorandom generator (PRG), we can expand a short seed from a randomness source to polynomially many bits of pseudorandomness. This well-known property is especially interesting here, as it allows us to apply Theorem 1 in a much broader form: Virtually every classically secure protocol is differentially secure when only a short random seed has been drawn from a bounded weak source and then expanded via a PRG, as this puts a limit on the entropy loss imposed by the actual bounded weak source. We formalize this observation in the following corollary, which is central to our work.

Corollary 2. *If two probabilistic machines X_0 and X_1 are computationally indistinguishable for a class of ppt machines $\mathcal{A}$ and uniform randomness, then X_0 and X_1 are also $(\alpha+\beta, 2^\alpha \cdot \delta)$-differentially indistinguishable for $\mathcal{A}$ and for a negligible function δ, if they draw their randomness from a PRG that is seeded with a (α, β)-bounded weak source.*

The corollary also gives guarantees for protocols and security proofs in which the amount of necessary randomness can be influenced by the adversary, e.g., by sending requests to the machine.

4 Application to Cryptography

We apply differential indistinguishability to a common secrecy definition, namely indistinguishability under chosen ciphertext attacks for public-key encryption. This definition serves as example for how to instantiate the notion and how to apply our main results to quantify the secrecy loss under imperfect randomness.

Moreover, we analyze differential indistinguishability under composition. We obtain a general composability result for differential indistinguishability that comes, similar to the composability of differential privacy, with a loss of secrecy. We refer to Appendix C for a discussion about additional application examples (commitment schemes and zero-knowledge proofs).

4.1 Public-Key Encryption

For PKE, standard security definitions, e.g., *indistinguishability under adaptive chosen ciphertext attack* ($\mathsf{IND\text{-}CCA}$) [22] can naturally be relaxed to use differential indistinguishability instead of traditional indistinguishability.

Definition 5 **((ε,δ)-$\mathsf{DIF\text{-}IND\text{-}CCA}$).** *A pair $\mathsf{A} = (\mathsf{A}_0, \mathsf{A}_1)$ of ppt oracle machines is an $\mathsf{IND\text{-}CCA}$ adversary if A_0 outputs two messages x_0, x_1 of the same length together with a state s, A_1 outputs a bit, and both A_0 and A_1 have access to decryption oracles as defined below. A PKE scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has* (ε,δ)-differentially indistinguishable encryptions under adaptive chosen ciphertext attack for a randomness source $\{D_n\}_{n\in\mathbb{N}}$ *if for all $\mathsf{IND\text{-}CCA}$ adversaries and for all sufficiently large k and bitstrings z of polynomial length in k, it holds that* $\Pr\left[\mathsf{P}_{k,z}^{(0)} = 1\right] \leq 2^\varepsilon \Pr\left[\mathsf{P}_{k,z}^{(1)} = 1\right] + \delta$*, where $P_{k,z}^{(i)}$ is defined as:*

$$\mathsf{P}_{k,z}^{(i)} := (e,d) \leftarrow \mathsf{Gen}(1^k);\ ((x_0,x_1),s) \leftarrow \mathsf{A}_0^{\mathsf{Dec}(d,\cdot)}(1^k, e, z)$$
$$c \leftarrow \mathsf{Enc}(e, x_i; D_n);\ \text{output } \mathsf{A}_1^{\mathsf{Dec}_c(d,\cdot)}(1^k, s, c)$$

Here, $\mathsf{Dec}_c(d,\cdot)$ denotes a decryption oracle that answers on all ciphertexts except for c, where it returns an error symbol $\perp$. The randomness used by the encryption algorithm Enc is drawn from D_n.

Note that $(0, \delta)$-$\mathsf{DIF\text{-}IND\text{-}CCA}$ security is equivalent to traditional δ-$\mathsf{IND\text{-}CCA}$ security.

Encryption with Imperfect Randomness. Both the encryption algorithm and the key generation algorithm require randomness. Dodis and Yu [18] show that even if weak sources are used for the key generation of $\mathsf{IND\text{-}CCA}$ secure encryption schemes, the security is preserved. However, this result does not apply when imperfect randomness is used by the *encryption algorithm*. The next theorem, an application of Theorem 1, quantifies the secrecy loss whenever the encryption algorithm has only access to an (α, β)-bounded weak source.

Theorem 2. *Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be any PKE scheme that is δ-$\mathsf{IND\text{-}CCA}$ secure under the assumption that Enc consumes at most n bits of uniform randomness. Then $\mathcal{E}$ is $(\alpha + \beta, 2^{\alpha}\delta)$-$\mathsf{DIF\text{-}IND\text{-}CCA}$ secure if Enc uses an (α, β)-bounded weak source $\{D_n\}_{n\in\mathbb{N}}$ instead of a uniform source.*

We refer to Appendix A.1 for a proof.

Discussion. Theorem 2 enables us to provide meaningful guarantees if an $\mathsf{IND\text{-}CCA}$ secure encryption scheme relies on imperfect randomness, as long as the randomness used to encrypt the ciphertext in question is drawn from a bounded weak source. If an encryption scheme is (ε,δ)-$\mathsf{DIF\text{-}IND\text{-}CCA}$ secure, the adversary may learn that the probability that a ciphertext contains a particular message m_0 is 2^{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 $2^{\varepsilon} \approx 1.001$), both m_0 and m_1 are a plausible content of the ciphertext. In particular, the adversary cannot reasonably believe or even convince a third party that m_0 is the value that has been encrypted. Moreover, the encryptor retains (a weak form of) deniability: She could indeed have encrypted any message.

Imperfect Randomness in Both Key Generation and Encryption. Our results also enable us to give a differential indistinguishability guarantee in the case when both the key generation algorithm Gen and the encryption algorithm Enc make use of a bounded weak source. If a PRG was used, seeded by a bounded weak random source, then we can immediately apply Corollary 2 to derive a differential indistinguishability guarantee. In contrast to the result of Dodis and Yu that requires the encryption scheme to be simulatable as defined by [18], which excludes, e.g., stateful schemes, we do not require any such structural property of the scheme.[5] If, for some reason, no PRG was used, one can still apply Theorem 1, but this will naturally yield weaker guarantees, as the combined randomness of Gen and Enc needs to be taken into account (and moreover the security loss under composition is significant, as discussed below).

Multiple Encryptions. Theorem 2 states a guarantee only for a single encryption (namely the encryption of one challenge message). However, it can be extended to the encryption of a message vector. In particular, if a PRG is used

[5] We discuss simulatability as well as the relation between our result and the result by Dodis and Yu [18] in Section 5.2.

(and thus the amount of bounded weak randomness is limited to the seed of the PRG), Corollary 2 yields immediately a differential indistinguishability guarantee with ε being independent of the number of encrypted messages. If however, the encryption algorithm Enc is run several times with (fresh) imperfect randomness, the entropy loss of the randomness can increase linearly in the number of messages in the vector for SV block sources, and consequently, ε increases significantly.

Other Security Definitions. Although we focus on $\mathsf{IND\text{-}CCA}$ security for PKE in this section, the broad applicability of Theorem 1 allows to handle other security definitions such as *indistinguishability under chosen plaintext attack* ($\mathsf{IND\text{-}CPA}$) similarly.

4.2 Composability

Traditional indistinguishability with a negligible function δ and $\varepsilon = 0$ allows for polynomially many compositions, because a polynomial factor for the advantage of an adversary, which might come from from seeing multiple samples, does not help the adversary substantially (the advantage remains negligible). This is not true for differential indistinguishability in general, because the (non-negligible) multiplicative factors can, under certain conditions, be accumulated as well.

For individual users we have shown that sequential composition of one or more primitives is possible without an (additional) loss of secrecy if a PRG is used (Corollary 2). If, however, several users within a protocol use imperfect randomness, the secrecy can degrade. Interestingly, we can give a bound on the loss of secrecy that is similar to the composition that occurs for differential privacy. We formulate a general composition lemma that we can instantiate to cope with several situations.

Lemma 2. *Let $\mathcal{A}$ be a class of adversaries. If X_0 and X_1 are (ε, δ)-differentially indistinguishable for $\mathcal{A}$, and X_1 and X_2 are (ε', δ')-differentially indistinguishable for $\mathcal{A}$, then X_0 and X_2 are $(\varepsilon'', \delta'')$-differentially indistinguishable for $\mathcal{A}$ where $\varepsilon'' = \varepsilon + \varepsilon'$ and $\delta'' = 2^{\varepsilon'}\delta + 2^{\varepsilon}\delta'$.*

We refer to Appendix A.2 for a proof.

A direct application of the lemma is the above described scenario in which multiple users (sequentially or concurrently) contribute to a protocol and use bad randomness. In this case, the machine X_1 can express an intermediate scenario that is used in a straightforward hybrid argument, where for two users X_1 is the only hybrid. Moreover, the lemma is applicable to scenarios where an individual user draws from a random source several times (for several primitives or protocols) instead of using a PRG, and also to compositions of differential indistinguishability guarantees in information-theoretical settings, where a PRG cannot be employed in the first place.

5 Interpretation and Analysis

In this section, we analyze and interpret the security guarantees provided by differential indistinguishability. In particular, we study the impact of a multiplication factor, and the influence of min- and max-entropy on differential indistinguishability. Furthermore, we discuss the relation between differential indistinguishability and differential privacy.

5.1 Impact of a 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 $\Pr[A] \leq 2^{\varepsilon}\Pr[B] + \delta$ implies a purely additive bound $\Pr[A] \leq \Pr[B] + \delta + 2^{\varepsilon} - 1 \approx \Pr[B] + \delta + \varepsilon$, 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 A.3 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. The 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.

Besides differential privacy, a multiplicative factor has also been used to achieve a specialized relaxation of semantic security in the presence of efficient adversaries that may tamper with an SV source [1, App B.4], and additionally for a security analysis of anonymous communication protocols [2].

Example. Let us assume that Alice participates in an e-voting protocol based on, e.g., a commitment scheme. If the random source that she uses to seed her PRG turns out to be an (α, β)-bounded weak source, the commitments are still *ε-differentially hiding* (see Appendix C.1 for a formal definition), where $\varepsilon = \alpha + \beta$ is a small constant. Assume that Alice can vote for one of two popular candidates, say, Bob and Charlie, and she chooses to vote for Bob. In the traditional indistinguishability case, a non-negligible additive difference in the guarantee could result from a non-negligible probability of leaking the vote, which is highly unsatisfactory. The multiplicative factor 2^{ε}, however, allows us to guarantee that both cases will still maintain non-zero probability and no distinguisher can be sure whether Alice voted for Bob or for Charlie. Consider a distinguisher that only outputs, say '1' if it is certain that the vote was cast for Bob, and '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 cast for Charlie, then differential indistinguishability implies that the probability of outputting '1' is zero when the vote was cast for Bob.

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, but in all other cases, deniability is preserved.

5.2 Influence of Min- and Max-Entropy

The literature on imperfect randomness has focused on "weak (entropy) sources" (called α-weak sources in this paper), because a non-trivial amount of min-entropy suffices for many applications. It is known to be sufficient to achieve unpredictability-based definitions, i.e., security notions in which the adversary has to guess a whole bitstring, e.g., the *binding* property of commitments and *unforgeability* of signatures and message authentication codes [13,15,29] (see also Appendix B).

Recently, Dodis and Yu [18] have extended this result significantly by showing that if such an unpredictability game can be considered a part of an indistinguishability game (e.g., for an encryption scheme with a weakly generated key) and if a *simulatability* condition proposed by the authors holds, then min-entropy also suffices for the indistinguishability game. In particular, they consider a primitive that can be divided into a setup phase (generating setup elements such as a key pair) and a simulatable (i.e., stateless and repeatable) indistinguishability game phase. They show that indistinguishability for such a primitive that can be preserved despite the setup phase (but not the game phase!) employing an α-weak source instead of uniform randomness. Here, the security notion under consideration is indeed divided. The setup phase has some, usually not explicitly specified, unpredictability notion (e.g., no adversary must be able to guess a correct key), and a corresponding game. Nevertheless, due to the impossibility result by Dodis et al. [15], whenever only min-entropy is ensured, a secrecy guarantee cannot be achieved in general, but only for certain schemes and under certain conditions. We discuss this in detail for public-key encryption in Section 4.1.

If, however, the randomness source has additionally a bounded max-entropy (and thus, among other properties, a full support), generic results are possible. In particular, a differential secrecy guarantee is still possible for a secrecy notion that is not simulatable (as defined by Dodis and Yu [18]), when an (α, β)-bounded weak source is used for generating the key. More importantly, such a differential guarantee is achievable when bounded weak randomness is used by the encryption algorithm itself.

Interestingly, max-entropy on its own is not sufficient for giving meaningful guarantees. If only the max-entropy of a source is bounded, the source could still output one individual element with a very high probability such that the probability over the other elements is evenly distributed. Therefore, we require both min-entropy and max-entropy measures for giving reasonable quantitative guarantees in all cases for which none of the specialized (e.g., unpredictability-based) solutions is applicable.

5.3 Relation to Differential Privacy and Sensitivity

Differential privacy [19] quantifies the privacy provided by database query mechanisms: Intuitively, differential privacy requires that the output of a query mechanism should not allow to distinguish similar databases better than with a small multiplicative factor. Both in terms of the definition and in terms of the small but usually non-negligible multiplicative factor, differential privacy and differential indistinguishability are closely related. We find this relation to be helpful for interpreting the guarantees and for understanding the drawbacks of differential indistinguishability. Differential privacy is influenced by the *sensitivity* of a statistical query, i.e., the amount of influence individual database records can have on the output of the query. Typical differential private mechanisms sanitize their output by adding random noise to guarantee a certain ε-level of privacy; the amount of added noise directly depends on the sensitivity.

Although there are neither databases nor the concept of utility (in the same sense as in differential privacy) in our setting, the fact that a bounded weak source is differentially indistinguishable from a uniform source is analogous to the differential privacy of a query mechanism. From this point of view, the missing entropy of the weak source corresponds to the sensitivity in differential privacy.

This relation between sensitivity and entropy is interesting for sources that can be analyzed in a block-by-block manner, e.g., (n, γ)-SV sources. For such a source the entropy loss and thus the “sensitivity” is directly associated with the parameter γ and the amount of blocks that are drawn from this source. 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 for a given value of ε. Clearly, the bias and thus the entropy loss in a $(1, \gamma)$-SV source can be arbitrarily increased, e.g., by drawing more random bits and taking the majority vote over them. Although this amplification does not make a difference for uniform randomness, it may increase the bias of the bits for SV sources. Therefore, for SV sources, the amount of randomness is a necessary parameter that influences the security.

6 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 [1,7,15,16]. Remarkably, Dodis et al. [15] show that traditional indistinguishability required for encryption, commitments, secret sharing, and zero-knowledge cannot be realized if a bounded weak source is used, which constitutes the main motivation for our work. More precisely, they prove that no protocol for any of these primitives can be secure against certain block sources, which include bounded weak sources. These sources sample blocks (i.e., several bits at once) that are $1/\mathit{poly}(k)$ close to the uniform distribution [11, 15, 33] for an arbitrary polynomial, where k is the security parameter.

This impossibility result has been refined and generalized over the last few years. Bosley and Dodis [7] show 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. In the universal composability (UC) setting [9], Canetti, Pass, and Shelat [10] show that even for (sampleable) sources for which a deterministic extractor exists, UC-secure commitments are not possible. Austrin et. al. [1] refined the impossibility result by Dodis et. al. [15] to show that it holds even when the adversary that tampers with the SV source is required to be efficient. Recently, Dodis and Yao [17] proposed a novel classification of random sources that groups them into "separable" and "expressive" sources. They apply their notions to rule out even one-bit encryption, commitment, and zero-knowledge proofs for many weak sources.

On the positive side, one line of research examines the extraction of (almost) perfect randomness from several kinds of imperfect randomness sources [6, 11, 12, 27, 34, 35]. However, extraction generally requires the source to have a certain degree of independence, whereas the only main requirement for bounded weak sources is to provide some entropy.

Aiming at particular applications, it has been shown that a few primitives can be securely instantiated even if only imperfect randomness is available [1, 14, 26], e.g., signatures [15] and Byzantine agreement [23].

Dodis et al. [14] 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 whether 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.

A multiplicative factor as in this work has also been used to achieve a specialized relaxation of semantic security in the presence of efficient adversaries that may tamper with an SV source [1, App. B.4]. Moreover, such a factor has proven useful for a security analysis of anonymous communication protocols [2, 3].

Most closely related to our work, Dodis and Yu [18] show that for all unpredictability-based primitives as well as for a class of restricted indistinguishability-based primitives, randomness sources with high min-entropy suffice to guarantee security whenever a uniform random source already guarantees security. While this is related to our result for unpredictability-based primitives (Corollary 3), Dodis and Yu establish a traditional indistinguishability guarantee (i.e., $\varepsilon = 0$) for a restricted class of indistinguishability-based primitives under weaker assumptions on the randomness source, clearly surpassing our results in these cases. However, the imposed gray-box requirements on indistinguishability games rule out many common and interesting cases. In particular, their analysis applies only to scenarios in which imperfect randomness is used at the beginning of a game, i.e., typically as input to a key generation algorithm. This leads to the observation that, e.g., for encryption, their result is restricted to imperfectly generated keys, and does not take care of the case where the *encryption algorithm* has access only to imperfect randomness.[6] In contrast, while our method provides only a differential guarantee, it is capable of obliviously analyzing essentially all indistinguishability games that make use of imperfect randomness, without imposing restrictions on the usage of this imperfect randomness. We refer to Section 5.2 for a more thorough analysis of our requirements on randomness and the possible results.

Kamara and Katz [26] 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. While we consider their approach orthogonal to ours, it turns out that a combination with our approach is possible. In the public-key setting, Bellare et al. [5] define and realize the notion of hedged public-key encryption, which provides secrecy guarantees even in the case of randomness failures, as long as the encrypted *message* has enough entropy.

7 Future Directions

Our work presents a novel view on the relation between weak randomness and indistinguishability, and it naturally leads to many more interesting questions.

From a theoretical point of view, we can ask whether it can be used in more scenarios such as for leakage-resilient cryptography [8, 21]. In particular, is it possible to give differential guarantees in cases where the adversary learns more than allowed by existing leakage-resilient schemes?

On the practical side, a natural next step is to apply our results to real applications and to random sources that are used in practice: Can we use entropy measurements of real randomness generators (both hardware generators and software generators) together with differential indistinguishability to give cryptographic guarantees?

Acknowledgments. We would like to thank Jalaj Upadhyay for insightful discussions about randomness sources and differential privacy. This work was supported by the German Ministry for Education and Research (BMBF) through funding for the Center for IT-Security, Privacy and Accountability (CISPA) and the German Universities Excellence Initiative.

A Postponed Proofs

A.1 Proof of Theorem 2 (Public-Key Encryption)

Proof. Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a public-key encryption scheme, let $\mathcal{A}_{ppt}$ be the class of ppt machines, and let $\{D_n\}_{n\in\mathbb{N}}$ be an (α, β)-bounded weak source.

[6] We note that this restriction cannot be circumvented by storing enough imperfect randomness at the beginning of the game in order to use it later during encryption. This approach would require the challenger to remember what parts of the stored randomness have already been used, which is implicitly excluded in [18]. We refer to Section 4.1 for a discussion.

To simplify the notation we write $P_{k,z}^{(b,r)}$ for simulating $P_{k,z}^{(b)}$ and using $r \in \{0,1\}^n$ as the randomness for Enc. Let $\mathsf{X}_0(1^k, r) := P_{k,z}^{(0,r)}$ and $\mathsf{X}_1 := P_{k,z}^{(1,r)}$ with the modification that X_0 and X_1 additionally provide a decryption oracle (as defined in Definition 5) to the adversary. Observe that by our definition of X_0 and X_1, the following two statements hold:

(i) $\mathsf{X}_0(1^k, U_n)$ and $\mathsf{X}_1(1^k, U_n)$ are indistinguishable for the class $\mathcal{A}_{ppt}$ of adversaries if and only if $\mathcal{E}$ is $\mathsf{IND\text{-}CCA}$.

(ii) $\mathsf{X}_0(1^k, D_n)$ and $\mathsf{X}_1(1^k, D_n)$ are (ε, δ)-differential indistinguishability for the class $\mathcal{A}_{ppt}$ of adversaries if and only if $\mathcal{E}$ is (ε, δ)-$\mathsf{DIF\text{-}IND\text{-}CCA}$ for $\{D_n\}_{n\in\mathbb{N}}$.

Thus, the claim follows immediately from Theorem 1. □

A.2 Proof of Lemma 2 (General Composition)

Proof. Given any adversary $\mathsf{A} \in \mathcal{A}$, for sufficiently large k and every possible output x of A, applying the definition of differential indistinguishability for X_0 and X_1 as well as X_1 and X_2 yields

$$\begin{aligned}
\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_0(1^k)\right\rangle = x\right] &\le 2^{\varepsilon} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_1(1^k)\right\rangle = x\right] + \delta \\
&\le 2^{\varepsilon}(2^{\varepsilon'} \Pr\left[\left\langle \mathsf{A}(1^k \middle| \mathsf{X}_2(1^k)\right\rangle = x\right] + \delta') + \delta \\
&\le 2^{\varepsilon+\varepsilon'} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_2(1^k)\right\rangle = x\right] + 2^{\varepsilon'}\delta + 2^{\varepsilon}\delta'.
\end{aligned}$$

Symmetrically, we obtain the opposite bound

$$\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_2(1^k)\right\rangle = x\right] \le 2^{\varepsilon'+\varepsilon} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{X}_0(1^k)\right\rangle = x\right] + 2^{\varepsilon}\delta' + 2^{\varepsilon'}\delta. \qquad \square$$

A.3 On Additive and Multiplicative Bounds (Section 5.1)

Given any arbitrary function δ with $1 \ge \delta_k > 0$, we construct 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. No matter which additive bound can be shown between two probabilistic events, there does not necessarily exist a non-trivial multiplicative bound, i.e., a multiplicative bound that could be used to improve on the additive bound.

Proof. 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 $\varepsilon \ge 0$ and δ be functions with $\delta'_k < \delta_k$ for sufficiently large k. For such k,

$$\begin{aligned}
\Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{S}(1^k, m_0)\right\rangle = 0\right] &= \delta > \delta' = 2^{\varepsilon}\, 0 + \delta' \\
&= 2^{\varepsilon} \Pr\left[\left\langle \mathsf{A}(1^k) \middle| \mathsf{S}(1^k, m_1)\right\rangle = 0\right] + \delta'.
\end{aligned}$$

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

B Unpredictability

So far we only considered the effect of (bounded) weak randomness on cryptographic indistinguishability notions. The security games for notions such as the *binding* property of commitments, *unforgeability* of signatures and message authentication codes, or guessing the key of an encryption scheme do not require indistinguishability. Instead, the adversary typically has to predict a particular bitstring, which should only be possible with negligible probability. It is well-known that such unpredictability (or unbreakability) notions are achievable even if an α-weak source is employed [13, 15, 18, 29].

We further analyze how imperfect randomness influences the probability for guessing a whole bitstring, e.g., for 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. Technically, the output of an adversary A in interaction with Z_0 is fed into Z_1, which finally outputs a bit $b \in \{0,1\}$ indicating whether the adversary has won.

Definition 6 (Unpredictability). *Let $\mathsf{Z} = (\mathsf{Z}_0, \mathsf{Z}_1)$ be a probabilistic machine that may keep state. We say that Z is δ-*unpredictable *for a class $\mathcal{A}$ of adversaries and for a distribution $\{D_n\}_{n\in\mathbb{N}}$, if for all $\mathsf{A} \in \mathcal{A}$ and for sufficiently large k,*

$$\Pr\left[\mathsf{Z}_1\left(\left\langle \mathsf{A}(1^k) \middle| \mathsf{Z}_0(1^k, D_n)\right\rangle\right) = 1\right] \leq \delta.$$

We show that for all games that can be described as a unpredictability game and for which the probability to win is negligible under uniform randomness, the probability is still negligible if an α-weak source is used. Similar to our comments in Remark 2, we notice that min-entropy suffices for this result.

Corollary 3. *If a probabilistic machine $\mathsf{Z} = (\mathsf{Z}_0, \mathsf{Z}_1)$ that may keep state is δ-unpredictable for a class of probabilistic machines $\mathcal{A}$ and consumes at most n bits of uniform randomness, then Z is $(2^\alpha\delta)$-unpredictable for $\mathcal{A}$ for any α-weak source $\{D_n\}_{n\in\mathbb{N}}$.*

Proof. We reduce this corollary 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 $r \in \{0,1\}^n$ 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

$$\begin{aligned} &\Pr\left[\mathsf{Z}_1(a) = 1 \,\middle|\, a \leftarrow \left\langle \mathsf{A}(1^k) \middle| \mathsf{Z}_0(1^k, D_n)\right\rangle\right] && (4)\\ &= \Pr\left[\mathsf{B}(1^k, D_n) = 1\right] \leq 2^\alpha \Pr\left[\mathsf{B}(1^k, U_n) = 1\right] && (5)\\ &= 2^\alpha \Pr\left[\mathsf{Z}_1(a) = 1 \,\middle|\, a \leftarrow \left\langle \mathsf{A}(1^k) \middle| \mathsf{Z}_0(1^k, U_n)\right\rangle\right] \leq 2^\alpha\,\delta. && (6) \end{aligned}$$

Inequality (5) follows from Lemma 1 and inequality (6) holds by assumption. □

C Commitments and Zero-Knowledge Proofs

In addition to public-key encryption (see Section 4.1), we apply differential indistinguishability to the hiding property of commitment schemes and the zero-knowledge property of proof systems.

C.1 Commitments

A non-interactive *commitment scheme* $\mathcal{C}$ consists of a three algorithms Setup, Commit, and Open. The Setup algorithm is run by a recipient and outputs public parameters *pp*. A sender runs the Commit algorithm, which takes as input the public parameters *pp* and a message m in the message space $\mathcal{M}$. It outputs a commitment *com* as well as opening information *op*. Correspondingly, the recipient runs the Open algorithm that takes as input the public parameters *pp*, the commitment *com* and the opening information *op*. It outputs the message m that has been committed to, or $\perp$ if *op* is not valid opening information for the commitment *com* under the public parameters *pp*.

A commitment scheme is *information-theoretically hiding* if the recipient, given only the public parameters *pp* and the commitment *com*, cannot determine any information about the message m. We relax this notion to a general hiding notion that allows for a multiplicative secrecy loss.

Definition 7 (ε-Differentially Hiding Commitment Scheme). *A non-interactive commitment scheme* $\mathcal{C} = (\mathsf{Setup}, \mathsf{Commit}, \mathsf{Open})$ *over a message space* $\mathcal{M}$ *is* information-theoretically ε-differentially hiding *for a random source* $\{D_n\}_{n\in\mathbb{N}}$*, if for all adversaries* $\mathsf{A} \in \mathcal{A}$*, for all pairs of messages* $m_0, m_1 \in \mathcal{M}$ *of the same length, and for all bitstrings pp (that represent public parameters),*

$$\Pr\left[\mathsf{A}(\mathsf{Commit}(pp, m_0)) = 1\right] \leq 2^{\varepsilon} \Pr\left[\mathsf{A}(\mathsf{Commit}(pp, m_1)) = 1\right].$$

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 provide a quantitative guarantee on the security of a information-theoretically hiding commitment scheme if it is used with an (α, β)-bounded weak source.

Theorem 3. *Let* $\mathcal{C} = (\mathsf{Setup}, \mathsf{Commit}, \mathsf{Open})$ *be a non-interactive commitment scheme over a message space* $\mathcal{M}$ *that is information-theoretically hiding (i.e., 0-hiding) as in Definition 7 in which* Commit *uses* n *bits of uniform randomness.* $\mathcal{C}$ *is* $(\alpha + \beta)$-differentially hiding *if* Commit *uses an* (α, β)*-bounded weak source* $\{D_n\}_{n\in\mathbb{N}}$ *instead of a uniform source.*

Proof. Let $\mathcal{C} = (\mathsf{Setup}, \mathsf{Commit}, \mathsf{Open})$ be an information-theoretically hiding non-interactive commitment scheme over the message space $\mathcal{M}$ such that Commit uses at most $\{0,1\}^n$ random bits. Let $\{D_n\}_{n\in\mathbb{N}}$ be an (α, β)-bounded weak source and $\{U_n\}_{n\in\mathbb{N}}$ be the uniform distribution over $\{0,1\}^n$. Furthermore, let $\mathcal{A}$ be the set of all probabilistic (not necessarily polynomially bounded) machines.

We define the machines X_0 and X_1 as: $\mathsf{X}_b(1^k, r) :=$ receive pp and (m_0, m_1) from A and output $\mathsf{Commit}(pp, m_b)$ using r as randomness. Observe that by our definition of X_0 and X_1, the following two statements hold:

(i) $\mathsf{X}_0(1^k, U_n)$ and $\mathsf{X}_1(1^k, U_n)$ are indistinguishable for the class $\mathcal{A}$ of adversaries if and only if $\mathcal{C}$ is information-theoretically hiding.[7]

(ii) $\mathsf{X}_0(1^k, D_n)$ and $\mathsf{X}_1(1^k, D_n)$ are $(\varepsilon, 0)$-differential indistinguishability for the class $\mathcal{A}$ of adversaries if and only if $\mathcal{C}$ is ε-differentially hiding for $\{D_n\}_{n\in\mathbb{N}}$.

Thus, the claim follows immediately from Theorem 1. □

Other Definitions and Multiple Commitments. Essentially the same analysis can be carried out for *statistically hiding* 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 7, in the latter case we further only consider ppt adversaries. Furthermore, a similar analysis is applicable to *interactive* commitment schemes. Theorem 3 presents guarantees for single commitments. However, if the adversary is ppt and if the bounded weak randomness was only used to seed a PRG, then by using Corollary 2 instead of Theorem 1 we can give a (computational) ε-differentially hiding guarantee for polynomially many commitments for $\varepsilon = \alpha + \beta$.

Binding Property with Imperfect Randomness. Whenever Theorem 3 is used to show a scheme to be ε-differentially hiding for an (α, β)-bounded weak source, the binding property is preserved (with a constant factor of 2^α), if *the receiver* uses an (α, β)-bounded weak source, which is in line with the results by Dodis and Yu [18]. The reason is that binding is an "unpredictability property" as discussed in Appendix B.

C.2 Zero-Knowledge Proofs

Our method also allows for relaxing traditional definitions based on the simulation paradigm, e.g., zero-knowledge (ZK) proofs. The relaxation applies to the indistinguishability of real views and simulated views: An (ε, δ)-ZK proof system is differentially secure in the sense that the output of a simulator is *almost* indistinguishable, i.e., (ε, δ)-differential indistinguishability, from the output of a verifier interacting with the real prover. In other words, a distinguisher with access to the output of the verifier can have only a small multiplicative advantage (quantified by ε) in guessing 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 guess 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

[7] Since Definition 7 required the scheme to be secure for all values pp, m_0, m_1, the values can also be chosen by the (unbounded) adversary A.

a witness w from the set $W(x)$ of witnesses for x as additional input, whereas the verifier V gets an auxiliary string z.

Definition 8 ((ε, δ)-Differential Zero-Knowledge[8]). *A proof system* $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ *is* (ε, δ)-differentially zero-knowledge *for a randomness source* $\{D_n\}_{n\in\mathbb{N}}$ *if for 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 $|x|$ for all ppt adversaries:*

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

For $\varepsilon = 0$ and a negligible function δ, this is the definition of *computational ZK* [22].

Theorem 4. *Let* $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ *be any proof system that is computationally ZK (i.e.,* $(0, \delta)$*-ZK for negligible* δ*) and requires the prover to use* n *bits of uniform randomness.* $\mathcal{P}$ *is* $(\alpha + \beta, 2^\alpha\delta)$*-differentially ZK if the prover* P *uses an* (α, β)*-bounded weak randomness source* $\{D_n\}_{n\in\mathbb{N}}$ *instead of a uniform randomness source.*

Proof. Let A be a machine in the class $\mathcal{A}_{ppt}$ of all probabilistic polynomial-time adversaries. Further, let F be an arbitrary 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.

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).[9] $\mathsf{X}_0(1^k, r)$ runs $\langle \mathsf{P}(x, w; r)|\mathsf{V}^*(x, z)\rangle$ 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)-differentially indistinguishable if and only if $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ is (ε, δ)-zero-knowledge. In particular, the goal of the polynomially bounded adversary A is to distinguish between the machines X_0 and X_1, which simulate $\{\langle \mathsf{V}^*(x, z)|\mathsf{P}\,(x, w; r)\rangle\}_{x\in L, z\in\{0,1\}^*}$ and $\{\mathsf{S}(x, z)\}_{x\in L, z\in\{0,1\}^*}$, 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, Theorem 1 implies the claim. □

Note that Theorem 4 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).

[8] Note that this definition is distinct from ε-knowledge [20], which allows the probabilities of the output bits of a distinguisher to be related by a non-negligible *additive* value.

[9] Note that F might not be computable. However, it can be verified that Theorem 1 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.

Soundness. The *soundness* property is preserved if the proof system uses a weak source instead of uniform randomness, similar to the binding property of commitments, which we discuss in Appendix C.1.

Randomness Source of the Simulator. Definition 8 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. Under the assumption 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.

Non-interactive Zero-Knowledge Proofs. Similarly, a differential relaxation is also possible for the security definition of non-interactive zero-knowledge proofs. In particular, we consider the case that not only the prover uses a bounded weak source but also the common random string (CRS) is generated by a bounded weak source.

Definition 9 (Non-interactive (ε, δ)-Differential Zero-Knowledge). *The non-interactive proof system $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ is* single-theorem adaptive (ε, δ)-differential zero-knowledge for a randomness source $\{D_n\}_{n\in\mathbb{N}}$ *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.*

(i) $\{(\sigma, F(\sigma), \pi) \mid \pi \leftarrow \mathsf{P}(x, w; r), (x, w) \leftarrow F(\sigma), (\sigma, r) \leftarrow D_n\}_{k\in\mathbb{N}}$
(ii) $\left\{(\sigma, F(\sigma), \pi) \mid \pi \leftarrow \mathsf{S}_2(x, s), (x, w) \leftarrow F(\sigma), (\sigma, s) \leftarrow \mathsf{S}_1(1^k)\right\}_{k\in\mathbb{N}}$

Note that both the CRS σ and the randomness r used by the prover are drawn from D_n. For the sake of simplicity, we consider an adaptive single-theorem definition, 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 4.2 if the prover (aside from the CRS) uses imperfect randomness.

Theorem 5. *Let $\mathcal{P} = (\mathsf{P}, \mathsf{V})$ be a single-theorem adaptive non-interactive proof system that is δ-zero-knowledge if the prover and the generation of the CRS together require at most n bits of uniform randomness. Then $\mathcal{P}$ is $(\alpha + \beta, 2^{\alpha}\delta)$-differential zero-knowledge, if an (α, β)-bounded weak source $\{D_n\}_{n\in\mathbb{N}}$ is used instead of a uniform source.*

The proof is analogous to the proof of Theorem 4.

Note that Theorem 5 also covers the case that the 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.

D Approximate (α, β)-Bounded Weak Sources

In this section we give a definition for sources that are only statistically close to bounded weak sources. To do so, we first introduce a relaxed variant of entropy, which is in line with [24].

D.1 Approximate Definitions

Definition 10. *Given a distribution D over the set X and another set $Y \subseteq X$, we use the following approximate measures of entropy:*

- *The* Y-min-entropy *of D is $H_{min}(D) := \min_{y \in Y} - \log \Pr[D = y]);$*
- *The* Y-max-entropy *of D is $H_{max}(D) := \max_{y \in Y} - \log \Pr[D = y]).$*

Definition 11. *A family of distributions $\{D_n\}_{n \in \mathbb{N}}$, each over the set $\{0,1\}^n$ of bitstrings of length n, is a δ-approximate* (α, β)-bounded weak source*, if for every D_n, there is a set $Y \subseteq \{0,1\}^n$ such that the following entropy requirements are satisfied:*

(i) *D_n has Y-min-entropy at least $n - \alpha$.*
(ii) *D_n has Y-max-entropy at most $n + \beta$.*
(iii) $\max(\Pr[D_n \notin Y], \Pr[U_n \notin Y]) \leq \delta$.

Using these approximate entropy measures, we can now define approximate bounded weak sources, which are a slight generalization of balanced sources (with bias at most d) from [24], in which $\alpha = \beta = \log(1 + d)$ and where δ is a negligible function.

D.2 Main Result for Approximate Bounded Weak Sources

We continue by adopting our main result for δ-approximate (α, β)-bounded weak sources.

Theorem 6. *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 $\{U_n\}_{n \in \mathbb{N}}$ over $\{0,1\}^n$, then X_0 and X_1 are also $(\alpha + \beta, 2^{\alpha}(\Delta + \delta) + \delta)$-differentially indistinguishable for $\mathcal{A}$ and any δ-approximate (α, β)-bounded weak source.*

The proof is analogous to the proof of Theorem 1.

Note that this result is also applicable for non-negligible δ; however, the resulting guarantees are not applicable to most cryptographic secrecy notions, as they would lead to a non-negligible additive factor.

References

1. P. Austrin, K. Chung, M. Mahmoody, R. Pass, and K. Seth. On the impossibility of cryptography with tamperable randomness. In *Proc. of the 34th International Cryptology Conference (CRYPTO'14)*, pages 462–479. Springer, 2014.
2. M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi. AnoA: A framework for analyzing anonymous communication protocols. In *Proc. of the 26th Computer Security Foundations Symposium (CSF'13)*, pages 163–178. IEEE, 2013.
3. M. Backes, A. Kate, S. Meiser, and E. Mohammadi. (Nothing else) MATor(s): monitoring the anonymity of Tor's path selection. In *Proc. of the 21st Conference on Computer and Communications Security (CCS'14)*. ACM, 2014.
4. M. Bellare. A note on negligible functions. *J. Cryptology*, 15(4):271, 2002.
5. M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged public-key encryption: How to protect against bad randomness. In *Proc. of the 15th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'09)*, pages 232–249. Springer, 2009.
6. C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy amplification by public discussion. *SIAM J. Comput.*, 17(2):210–229, 1988.
7. C. Bosley and Y. Dodis. Does privacy require true randomness? In *Proc. of the 4th Theory of Cryptography Conference (TCC'07)*, pages 1–20. Springer, 2007.
8. 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 *Proc. of the 51st Symposium on Foundations of Computer Science (FOCS'10)*, pages 501–510. IEEE, 2010.
9. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proc. of the 42nd Symposium on Foundations of Computer Science (FOCS'01)*, pages 136–145. IEEE, 2001.
10. R. Canetti, R. Pass, and A. Shelat. Cryptography from sunspots: How to use an imperfect reference string. In *Proc. of the 48th Symposium on Foundations of Computer Science (FOCS'07)*, pages 249–259. IEEE, 2007.
11. B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. In *Proc. of the 26th Symposium on Foundations of Computer Science (FOCS'85)*, pages 429–442. IEEE, 1985.
12. Y. Dodis, A. Elbaz, R. Oliveira, and R. Raz. Improved randomness extraction from two independent sources. In *Proc. of the 8th International Workshop on Randomization and Computation (RANDOM'04)*, pages 334–344. Springer, 2004.
13. Y. Dodis, J. Katz, L. Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In *Proc. of the 26th International Cryptology Conference (CRYPTO'06)*, pages 232–250. Springer, 2006.
14. Y. Dodis, A. López-Alt, I. Mironov, and S. Vadhan. Differential privacy with imperfect randomness. In *Proc. of the 32nd International Cryptology Conference (CRYPTO'12)*, pages 497–516. Springer, 2012.
15. Y. Dodis, S. J. Ong, M. Prabhakaran, and A. Sahai. On the (im)possibility of cryptography with imperfect randomness. In *Proc. of the 45th Symposium on Foundations of Computer Science (FOCS'04)*, pages 196–205. IEEE, 2004.
16. Y. Dodis, K. Pietrzak, and B. Przydatek. Separating sources for encryption and secret sharing. In *Proc. of the 3rd Theory of Cryptography Conference (TCC'06)*, pages 601–616. Springer, 2006.

17. Y. Dodis and Y. Yao. Privacy and imperfect randomness. IACR Cryptology ePrint Archive, Report 2014/623, 2014.
18. Y. Dodis and Y. Yu. Overcoming weak expectations. In *Proc. of the 10th Theory of Cryptography Conference (TCC'13)*, pages 1–22. Springer, 2013.
19. C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proc. of the 3rd Theory of Cryptography Conference (TCC'06)*, pages 265–284. Springer, 2006.
20. C. Dwork, M. Naor, and A. Sahai. Concurrent zero-knowledge. *J. ACM*, 51(6):851–898, 2004.
21. S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *Proc. of the 48th Symposium on Foundations of Computer Science (FOCS'07)*, pages 293–302. IEEE, 2008.
22. O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Foundations of Cryptography. Cambridge University Press, 2001.
23. S. Goldwasser, M. Sudan, and V. Vaikuntanathan. Distributed computing with imperfect randomness. In *Proc. of the 19th International Conference on Distributed Computing (DISC'05)*, pages 288–302. Springer, 2005.
24. I. Haitner, O. Horvitz, J. Katz, C.-Y. Koo, R. Morselli, and R. Shaltiel. Reducing complexity assumptions for statistically-hiding commitment. In *Proc. of the 24th International Conference on Theory and Applications of Cryptographic Techniques (EUROCRYPT'05)*, pages 58–77. Springer, 2005.
25. Y. T. Kalai, X. Li, A. Rao, and D. Zuckerman. Network extractor protocols. In *Proc. of the 49th Symposium on Foundations of Computer Science (FOCS'08)*, pages 654–663. IEEE, 2008.
26. S. Kamara and J. Katz. How to encrypt with a malicious random number generator. In *Proc. of the 15th International Workshop on Fast Software Encryption (FSE'08)*, pages 303–315. Springer, 2008.
27. J. Kamp and D. Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. In *Proc. of the 44th Symposium on Foundations of Computer Science (FOCS'03)*, pages 92–101. IEEE, 2003.
28. J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. CRC Press, 2007.
29. U. M. Maurer and S. Wolf. Privacy amplification secure against active adversaries. In *Proc. of the 17th International Cryptology Conference (CRYPTO'97)*, pages 307–321. Springer, 1997.
30. J. L. McInnes and B. Pinkas. On the impossibility of private key cryptography with weakly random keys. In *Proc. of the 10th International Cryptology Conference (CRYPTO'90)*, pages 421–435. Springer, 1990.
31. I. Mironov, O. Pandey, O. Reingold, and S. Vadhan. Computational differential privacy. In *Proc. of the 29th International Cryptology Conference (CRYPTO'09)*. Springer, 2009.
32. O. Reingold, L. Trevisan, M. Tulsiani, and S. Vadhan. Dense subsets of pseudorandom sets. In *Proc. of the 49th Symposium on Foundations of Computer Science (FOCS'08)*, pages 76–85. IEEE, 2008.
33. M. Santha and U. V. Vazirani. Generating quasi-random sequences from slightly-random sources. In *Proc. of the 25th Symposium on Foundations of Computer Science (FOCS'84)*, pages 434–440. IEEE, 1984.
34. L. Trevisan and S. Vadhan. Extracting Randomness from Samplable Distributions. In *Proc. of the 41st Symposium on Foundations of Computer Science (FOCS'00)*, pages 32–42. IEEE, 2000.
35. J. Von Neumann. Various techniques used in connection with random digits. *National Bureau of Standards Applied Mathematics Series*, 12:36–38, 1951.