

Replacing Probability Distributions in Security Games via Hellinger Distance

Kenji Yasunaga
Osaka University
yasunaga@ist.osaka-u.ac.jp

Abstract

Security of cryptographic primitives is usually proved by assuming “ideal” probability distributions. We need to replace them with approximated “real” distributions in the real-world systems without losing the security level. We demonstrate that the Hellinger distance is useful for this problem, while the statistical distance is mainly used in the cryptographic literature. First, we show that for preserving λ -bit security of a given security game, the closeness of $2^{-\lambda/2}$ to the ideal distribution is sufficient for the Hellinger distance, whereas $2^{-\lambda}$ is generally required for the statistical distance. The result can be applied to both search and decision primitives through the bit security framework of Micciancio and Walter (Eurocrypt 2018). We also show that the Hellinger distance gives a tighter evaluation of closeness than the max-log distance when the distance is small. Finally, we show that the leftover hash lemma can be strengthened to the Hellinger distance. Namely, a universal family of hash functions gives a strong randomness extractor with optimal entropy loss for the Hellinger distance. Based on the results, a λ -bit entropy loss in randomness extractors is sufficient for preserving λ -bit security. The current understanding based on the statistical distance is that a 2λ -bit entropy loss is necessary.

1 Introduction

Security of cryptographic primitives relies on the use of randomness sources. Secret keys and random bits are usually assumed to be sampled from uniform distributions. Various probability distributions other than uniform ones appear in cryptography. In lattice-based cryptography, discrete Gaussian distributions are used for the hardness of the Learning with Errors (LWE) problem [34, 33, 23, 29] and the tight reductions for the Short Integer Solution (SIS) problem [26, 25]. Adding noise from Laplace distributions enables data privacy of statistical databases in differential privacy [15, 14, 16].

To ensure the security of primitives, we usually define a security game played by an adversary and show that the adversary’s success probability is sufficiently close to some value. In the proof, we assume we can use “ideal” probability distributions. We need to replace them with approximated “real” distributions in real-world systems. For example, in a security game of an encryption scheme, the adversary receives a ciphertext and tries to guess which of the two plaintexts were encrypted. The scheme is secure if the success probability is sufficiently close to $1/2$. A secret key and random coins for encryption are assumed to be sampled from uniform distributions. One may employ the output of a randomness extractor [36, 11] as a randomness source since the output distribution is sufficiently close to the uniform distribution. However, the distance to the ideal distribution may affect the security level of primitives. A question is which closeness measure of distributions should be used when replacing distributions in security games.

In cryptographic literature, we mainly employ the *statistical distance* (a.k.a. the *total variation distance*) to measure distribution closeness. The main reason is that it enables a straightforward analysis of the resulting security levels. The statistical distance is defined as the maximum difference of probabilities of events between two distributions. By employing a distribution P that is close to ideal Q within ϵ in the statistical distance, we can guarantee that the adversary’s success probability only increases by at most ϵ . However, there may not be any other reason for using the statistical distance.

Also, achieving security by the statistical distance has some limitations. Radhakrishnan and Ta-Shma [32] showed a lower bound on the entropy loss of randomness extractors. Roughly, the result implies that to extract a uniformly random string from an entropy source, we need to lose $2 \log(1/\epsilon)$ of entropy, where ϵ is the distance to the uniform distribution. Based on this result, if we extract a random string from a source of 120-bit entropy by ensuring 50-bit security, the output bit should be of length at most $120 - 2 \cdot 50 = 20$. This loss of entropy is crucial when using biometric data as entropy sources [11, 8], where a limited amount of entropy can be used. Randomness extraction (or key derivation) from weak sources arises in many situations of cryptography, including Diffie-Hellman key exchange [17, 21] and random number generators from physical sources [5, 4], to name a few.

Our Contribution. In this work, we propose to use the *Hellinger distance* for replacing distributions in security games. Roughly speaking, we show that the closeness of $2^{-\lambda/2}$ in the Hellinger distance is sufficient to preserve λ -bit security. When using the statistical distance, the closeness of $2^{-\lambda}$ is, in general, necessary to achieve the same security level.

To discuss the *bit security*, we use the framework of Micciancio and Walter [28]. Their framework can smoothly connect the bit security between *search* and *decision* primitives. Their definition is the same as the standard one for search primitives, where the secret is chosen from a sufficiently large space. For decision primitives, in which the attacker tries to guess a secret bit, the definition of the advantage is different from the standard one. See Section 3 for the details. We show that the distance closeness of $2^{-\lambda/2}$ in the Hellinger distance is sufficient for preserving the bit security for both search and decision primitives.

Next, we show that the Hellinger distance gives a tighter evaluation of closeness than the *max-log distance*, the probability metric introduced in [27, 28]. The work showed that the closeness of $2^{-\lambda/2}$ in the max-log distance is sufficient for preserving λ -bit security. We proved that the Hellinger distance is bounded above by the max-log distance as long as the max-log distance is at most $\sqrt{2} - 1$. Also, we present a concrete example of a distribution pair such that their Hellinger distance is exponentially small, while their max-log distance is a constant.

Finally, we demonstrate the usefulness of using the Hellinger distance in the problem of randomness extraction (or information-theoretic key derivation). We show that the leftover hash lemma [6, 20] can be strengthened to the Hellinger distance without losing the security level. Namely, a *universal* family of hash functions gives a strong randomness extractor with optimal entropy loss even when measuring in the Hellinger distance. We can conclude that the entropy loss of λ -bit is sufficient for preserving λ -bit security. In general, the entropy loss of 2λ -bit is necessary to preserve bit security when using the statistical distance.

Techniques. We describe a technical overview of our results. Although the actual proofs seem different from the below, it reflects the difference between the statistical distance and the Hellinger distance. Let $P = (P_1, P_2, \dots)$ and $Q = (Q_1, Q_2, \dots)$ be a pair of probability distribution ensembles such that each P_i is close to Q_i . Let ϵ_A^Q be the probability that an adversary A succeeds in the security game in which samples from Q is used. We want to bound the probability ϵ_A^P , which is the success probability when using P instead of Q .

For $\ell \in \mathbb{N}$, we define the probability μ_ℓ^Q that A succeeds in at least one out of ℓ independent plays of G_A^Q . As long as ℓ is small compared to $1/\epsilon_A^P$, it holds that $\mu_\ell^P \approx \ell \cdot \epsilon_A^P$. Since the number of sample queries in each game is bounded above by the running time T_A of A , $\mu_\ell^P \leq \mu_\ell^Q + \text{SD}(P^\ell, Q^\ell) \leq \mu_\ell^Q + \ell T_A \cdot \max_i \text{SD}(P_i, Q_i)$, where $\text{SD}(P_i, Q_i)$ is the statistical distance between P_i and Q_i , and P^ℓ is the ℓ -fold product of P . Note that we use the relation $\text{SD}(P^\ell, Q^\ell) \leq \ell T_A \cdot \max_i \text{SD}(P_i, Q_i)$. Now, it holds that $\epsilon_A^P \approx \ell^{-1} \cdot \mu_\ell^P \leq \ell^{-1} \cdot (\mu_\ell^Q + \ell T_A \cdot \max_i \text{SD}(P_i, Q_i)) \approx \epsilon_A^Q + T_A \cdot \max_i \text{SD}(P_i, Q_i)$. Thus, if the primitive has λ -bit security, i.e., $\epsilon_A^Q/T_A \leq 2^{-\lambda}$, then $\epsilon_A^P/T_A \leq 2^{-\lambda} + \max_i \text{SD}(P_i, Q_i)$. It implies that $\max_i \text{SD}(P_i, Q_i) \leq 2^{-\lambda}$ is required for preserving bit security. For the Hellinger distance $\text{HD}(P_i, Q_i)$, we provide a technical lemma (Lemma 1) showing that $\text{SD}(P^\ell, Q^\ell) \leq \sqrt{2\ell T_A} \cdot \max_i \text{HD}(P_i, Q_i)$. Therefore, we have $\epsilon_A^P \leq \epsilon_A^Q + \sqrt{2\ell^{-1} T_A} \cdot \max_i \text{HD}(P_i, Q_i)$. Hence, if the primitive has λ -bit security, $\epsilon_A^P/T_A \leq 2^{-\lambda} + \sqrt{2(\ell T_A)^{-1}} \cdot \max_i \text{HD}(P_i, Q_i)$, implying that, by choosing $\ell = 1/\epsilon_A^P$, it suffices to satisfy $\max_i \text{HD}(P_i, Q_i) \leq 2^{-\lambda/2}$ for preserving bit security.

The leftover hash lemma essentially gives an upper bound on the *collision probability* of the hash functions chosen from a universal family. If the collision probability is bounded, it is close to uniform in the Hellinger distance. This relation was provided by Chung and Vadhan [9] using Hölder’s inequality. Based on the relation, we show that a universal family of hash functions gives a strong randomness extractor for the Hellinger distance. Notably, we can achieve the same parameters as in the case of the statistical distance. Thus, the optimal entropy loss is achieved by universal hash functions.

Related Work. Barak et al. [3] initiated the study on improving the leftover hash lemma for a limited class of primitives. The work of [3, 13] showed that the bound of [32] could be improved for the search primitives and the square-friendly decision primitives, including stateless encryption schemes and weak pseudorandom functions. Specifically, the entropy loss of λ is sufficient for square-friendly primitives. For search primitives, Dodis, Pietrzak, and Wichs [12] achieved the entropy loss of $O(\log \lambda)$ in randomness extraction with $O(\lambda)$ -wise independent hash functions. Matsuda et al. [24] generalized the results of [13] by using the Rényi divergence for capturing the case that the ideal distribution is not uniform. Skorski [37] showed that being square-friendly is necessary to reduce entropy loss. Compared with the above work, our results for reducing entropy loss do not build on a specific class of primitives but need to rely on the bit security framework of [28], especially for the decision primitives.

In lattice-based cryptography, several probability metrics other than the statistical distance have been employed for improving the analysis of security proofs [30, 2, 27, 31, 39]. The metrics used in these work include the Kullback-Leibler divergence, the Rényi divergence, the max-log distance, and the relative error.

Micciancio and Walter [28] introduced a new framework of bit security that can smoothly connect the search primitives and the decision primitives quantitatively. A feature is that it allows the adversary to declare an attack failure. With their framework, we can say that a λ -bit secure pseudorandom generator (a decision primitive) is also a λ -bit secure one-way function (a search primitive). In the conventional definition, a $\lambda/2$ -bit secure pseudorandom generator strangely yields a λ -bit secure one-way function. While they showed that the max-log distance is beneficial in their framework, we show that the Hellinger distance has the same effect and gives a tighter evaluation of closeness.

Distances/divergences between distributions other than the statistical distance have appeared in other cryptographic literature. Chung and Vadhan [9] gave a tight analysis of hashing block sources using the Hellinger distance as a key tool. Agrawal [1] introduced the notion of randomness extractors for the Kullback-Leibler divergence and gave explicit/non-explicit constructions with almost the same parameters as standard extractors. Steinberger [38] used the Hellinger distance for the improved analysis of key-alternating ciphers. Dai, Hoang, and Tessaro [10] used the chi-square divergence to analyze the information-theoretic indistinguishability proofs. Berman et al. [7] studied the polarization lemma for various distance notions such as the triangular discrimination and the Jensen-Shannon divergence to extend the region of polarization.

2 Preliminaries

We define the distances for distributions used in this work. The basic properties and general relationships of various distances/divergences can be found in [18]. We also present a useful lemma for the Hellinger distance, which will be used later.

Let P and Q be probability distributions over a finite set Ω . For a distribution P over Ω and $A \subseteq \Omega$, we denote by $P(A)$ the probability of event A , which is equal to $\sum_{x \in A} P(x)$. The *statistical distance* (a.k.a. *total variation distance*) between P and Q is

$$\text{SD}(P, Q) = \max_{A \subseteq \Omega} |P(A) - Q(A)|.$$

The *data processing inequality* guarantees that for any function $f: \Omega \rightarrow \{0, 1\}^*$, we have

$$\text{SD}(f(P), f(Q)) \leq \text{SD}(P, Q). \tag{1}$$

The *Hellinger distance* between P and Q is

$$\text{HD}(P, Q) = \sqrt{\frac{1}{2} \sum_{x \in \Omega} \left(\sqrt{P(x)} - \sqrt{Q(x)} \right)^2} = \sqrt{1 - \sum_{x \in \Omega} \sqrt{P(x) \cdot Q(x)},}$$

which takes values in $[0, 1]$. It holds that

$$\text{HD}(P, Q)^2 \leq \text{SD}(P, Q) \leq \sqrt{2} \cdot \text{HD}(P, Q). \quad (2)$$

The *Hellinger affinity* is defined as

$$\text{HA}(P, Q) = 1 - \text{HD}(P, Q)^2 = \sum_{x \in \Omega} \sqrt{P(x) \cdot Q(x)},$$

which is also known as the Bhattacharyya coefficient or fidelity.

The Hellinger distance has the following useful property, which is weaker than the *Pythagorean probability preservation* defined in [27, 28].

Lemma 1. Let $Q = (Q_1, \dots, Q_\ell)$ and $P = (P_1, \dots, P_\ell)$ be probability distribution ensembles over a finite support $\prod_i \Omega_i$. Then,

$$\text{SD}(P, Q) \leq \sqrt{2\ell} \cdot \max_{a_i \in \prod_{j < i} \Omega_j} \text{HD}(P_i|a_i, Q_i|a_i).$$

Proof. Let $\epsilon = \max_{a_i \in \prod_{j < i} \Omega_j} \text{HD}(P_i|a_i, Q_i|a_i)$. Then, $\text{HA}(P_i|a_i, Q_i|a_i) = 1 - \text{HD}(P_i|a_i, Q_i|a_i)^2 \geq 1 - \epsilon^2$ for any i and $a_i \in \prod_{j < i} \Omega_j$. It holds that

$$\begin{aligned} & \text{HA}(P, Q) \\ &= \sum_{b_1, \dots, b_\ell \in \prod_i \Omega_i} \sqrt{P(b_1, \dots, b_\ell) \cdot Q(b_1, \dots, b_\ell)} \\ &= \sum_{b_1 \in \Omega_1} \sqrt{P_1(b_1) \cdot Q_1(b_1)} \cdot \left(\sum_{b_2 \in \Omega_2} \sqrt{P_2(b_2|P_1 = b_1) \cdot Q_2(b_2|Q_1 = b_1)} \cdot \left(\dots \right. \right. \\ & \quad \left. \left. \cdot \left(\sum_{b_\ell \in \Omega_\ell} \sqrt{P_\ell(b_\ell|P_{1,\ell-1} = (b_1, \dots, b_{\ell-1})) \cdot Q_\ell(b_\ell|Q_{1,\ell-1} = (b_1, \dots, b_{\ell-1}))} \right) \dots \right) \right) \\ &\geq (1 - \epsilon^2)^\ell \geq 1 - \ell\epsilon^2, \end{aligned}$$

where $P_{1,\ell-1} = (P_1, \dots, P_{\ell-1})$ and $Q_{1,\ell-1} = (Q_1, \dots, Q_{\ell-1})$. Thus, $\text{HD}(P, Q) = \sqrt{1 - \text{HA}(P, Q)} \leq \sqrt{\ell}\epsilon$. The statement follows from (2). \square

3 Replacing Distributions in Security Games

We consider replacing probability distributions in security games. Let $Q = (Q_\theta)_\theta$ be an ideal distribution ensemble in a security game. We want to replace Q with an approximated distribution ensemble $P = (P_\theta)_\theta$ without compromising security. We define a general security game by following the definitions of [28, 27].

An *n-bit security game* G_A is a game played by an adversary A interacting with a challenger C . At the beginning of the game, the challenger chooses a uniformly random secret $x \in \{0, 1\}^n$, represented by the random variable X . At the end of the game, A outputs some value v , represented by the random variable V . The goal of the adversary is to output v such that $R(x, v) = 1$, where R is a Boolean function. The adversary may output a special symbol \perp such that $R(x, \perp) = 0$ for any x . During the game, A or C may obtain a sample from a distribution Q_θ by querying θ . The success probability of A is $\epsilon_A^Q = \Pr[R(X, V) = 1]$,

where the probability is taken over the randomness of the entire security game, including the randomness of A . We may denote the game by G_A^Q since we intend to replace Q with another distribution ensemble.

Micciancio and Walter [28] defined the bit security based on an *advantage* that is different from most of the literature for the case $n = 1$. We use their framework for evaluating the security loss by replacing distributions in security games.

Definition 1 (Bit Security of [28]). Let Π be a primitive for which an n -bit security game G_A^Q is defined. Let X and V be random variables representing a random secret $x \in \{0, 1\}^n$ and an output value v of A in G_A^Q , respectively. We define the *output probability* $\alpha_A = \Pr[V \neq \perp]$ and the *conditional success probability* $\beta_A = \Pr[R(X, V) = 1 \mid V \neq \perp]$. The *advantage* of A is defined to be

$$\text{adv}_A = \begin{cases} \alpha_A \beta_A & n > 1 \\ \alpha_A (2\beta_A - 1)^2 & n = 1 \end{cases}.$$

The *bit security* of Π is defined to be

$$\min_A \log_2 \frac{T_A}{\text{adv}_A},$$

where T_A is the running time of A . We say the primitive is λ -bit secure if its bit security is at least λ .

We say Π is a *search* primitive if its n -bit security game G is defined for $n > 1$, and a *decision* primitive if G is a 1-bit security game.

For search primitives, it is not difficult to see that $Q = (Q_\theta)_\theta$ can be replaced with $P = (P_\theta)_\theta$ if their statistical distance between P_θ and Q_θ is sufficiently small and the number of queries is not so much. Specifically, if a search primitive Π^Q is λ -bit secure and $\text{SD}(P_\theta, Q_\theta) \leq 2^{-\lambda}$, then Π^P is $(\lambda - \log q)$ -bit secure, where we denote by Π^Q a primitive for which a security game G_A^Q is defined and q is the number of queries. This fact implies that it is sufficient to choose P that is close to Q within $2^{-\lambda}$ in the statistical distance for preserving the bit security.

Micciancio and Walter [27, 28] demonstrated that if distributions are close in the *max-log distance*, the closeness requirement may be relaxed. The max-log distance between distributions P and Q over Ω with the same support $S \subseteq \Omega$ is

$$\text{ML}(P, Q) = \max_{x \in S} |\ln P(x) - \ln Q(x)|.$$

They showed that the closeness of $2^{-\lambda/2}$ is sufficient to preserve the bit security for search primitives in [27] and decision primitives in [28].

Lemma 2 ([27, 28]). Let $Q = (Q_i)_i$ and $P = (P_i)_i$ be distribution ensembles over the support $\prod_i \Omega_i$ satisfying $\text{ML}(P_i|a_i, Q_i|a_i) \leq 2^{-\lambda/2} \leq 1/4$ for any i and $a_i \in \prod_{j < i} \Omega_j$. If a search primitive Π^Q is λ -bit secure, then Π^P is $(\lambda - 3)$ -bit secure. If a decision primitive Π^Q is λ -bit secure, then Π^P is $(\lambda - 8)$ -bit secure.

They showed the above results for a more general class of λ -efficient divergences [27, 28]. We demonstrate that similar effects can be obtained by using the Hellinger distance.

3.1 Security for Search Primitives

Let $Q = (Q_i)_i$ and $P = (P_i)_i$ be distribution ensembles over the same support $\prod_i \Omega_i$. We consider P and Q satisfying $\text{HD}(P_i|a_i, Q_i|a_i) \leq 2^{-\lambda/2}$ for any i and $a_i \in \prod_{j < i} \Omega_j$. We call such a pair (P, Q) a $2^{-\lambda/2}$ -Hellinger close pair. We show that this closeness is sufficient for preserving bit security.

Theorem 1. Let Π^Q be a primitive for which an n -bit security game G_A^Q is defined for $n > 1$. For any $2^{-\lambda/2}$ -Hellinger close pair (P, Q) , if Π^Q is λ -bit secure, then Π^P is $(\lambda - 2.973)$ -bit secure.

Proof. Let ϵ_A^Q be the success probability of an adversary A in G_A^Q , and T_A the running time of A . Since Π is λ -bit secure, it holds that $\epsilon_A^Q/T_A \leq 2^{-\lambda}$ for any A . It is sufficient to show that $\epsilon_A^P/T_A \leq 2^{-(\lambda - 2.973)}$, where ϵ_A^P is the success probability of A in G_A^P .

We consider ℓ independent plays of G_A^Q and define μ_ℓ^Q to be the probability that A succeeds in at least one out of ℓ plays of G_A^Q . Namely, $\mu_\ell^Q = 1 - (1 - \epsilon^Q)^\ell$. We define μ_A^P analogously. Since the number of queries to the distribution ensemble is at most T_A in each play, it holds that

$$\left| \mu_\ell^P - \mu_\ell^Q \right| \leq \text{SD}(P^\ell, Q^\ell) \leq \sqrt{2\ell T_A} \cdot 2^{-\lambda/2},$$

where P^ℓ is the ℓ -fold product of P , the first inequality is by the data processing inequality, and the second inequality follows from Lemma 1. Thus,

$$(1 - \epsilon_A^Q)^\ell \leq \sqrt{2\ell T_A} \cdot 2^{-\lambda/2} + (1 - \epsilon_A^P)^\ell.$$

By the fact that $(1 - x)^\ell \geq 1 - \ell x$ for $x \in [0, 1]$ and setting $\ell = 1/\epsilon_A^P$, it holds that

$$1 - \frac{\epsilon_A^Q}{\epsilon_A^P} \leq \sqrt{\frac{2T_A \cdot 2^{-\lambda}}{\epsilon_A^P}} + (1 - \epsilon_A^P)^{1/\epsilon_A^P} < \sqrt{\frac{2T_A \cdot 2^{-\lambda}}{\epsilon_A^P}} + e^{-1},$$

where we use the relation that $(1 - 1/x)^x < e^{-1}$ for $x > 0$. By rewriting the inequality,

$$\left(\sqrt{\epsilon_A^P} - \frac{\sqrt{T_A \cdot 2^{-\lambda}}}{\sqrt{2}(1 - e^{-1})} \right)^2 < \frac{\epsilon_A^Q}{1 - e^{-1}} + \frac{T_A \cdot 2^{-\lambda}}{2(1 - e^{-1})^2}.$$

It holds that

$$\sqrt{\epsilon_A^P} < \sqrt{\frac{\epsilon_A^Q}{1 - e^{-1}} + \frac{T_A \cdot 2^{-\lambda}}{2(1 - e^{-1})^2}} + \frac{\sqrt{T_A \cdot 2^{-\lambda}}}{\sqrt{2}(1 - e^{-1})}.$$

Squaring both sides gives that

$$\frac{\epsilon_A^P}{T_A} < \frac{\epsilon_A^Q}{(1 - e^{-1})T_A} + \frac{2^{-\lambda}}{(1 - e^{-1})^2} + \frac{\sqrt{2} \cdot 2^{-\lambda/2}}{1 - e^{-1}} \sqrt{\frac{\epsilon_A^Q}{(1 - e^{-1})T_A} + \frac{2^{-\lambda}}{2(1 - e^{-1})^2}}.$$

Since $\epsilon_A^Q/T_A \leq 2^{-\lambda}$, we have $\epsilon_A^P/T_A < 7.851 \cdot 2^{-\lambda} < 2^{2.973} \cdot 2^{-\lambda}$. Therefore, the statement follows. \square

3.2 Security for Decision Primitives

Next, we show that the closeness of $2^{-\lambda/2}$ in the Hellinger distance is sufficient for preserving λ -bit security even for decision primitives.

Theorem 2. Let Π^Q be a primitive for which a 1-bit security game G_A^Q is defined. For any $2^{-\lambda/2}$ -Hellinger close pair (P, Q) , if Π^Q is λ -bit secure, then Π^P is $(\lambda - 6.847)$ -bit secure.

Proof. Suppose for contradiction that Π^P is not $(\lambda - 6.847)$ -bit secure. Namely, there exists an adversary A for Π^P with running time T_A such that $\alpha_A^P(2\beta_A^P - 1)^2 > T_A/2^{\lambda - 6.847}$, where α_A^Q and β_A^Q are the output probability and the conditional success probability of A . Since Π^Q is λ -bit secure, we have $\alpha_A^Q(2\beta_A^Q - 1)^2 \leq T_A/2^\lambda$, where α_A^Q and β_A^Q are the corresponding probabilities for Π^Q . Let $\alpha = \min\{\alpha_A^Q, \alpha_A^P\}$.

We define the games \tilde{G}_A^Q and \tilde{G}_A^P such that they are the same as G_A^Q and G_A^P with the difference that the adversary can restart the game with fresh randomness at any time. Consider the adversary B that runs A repeatedly until either the output value is different from \perp or B runs A in total $1/\alpha$ times, and outputs the same value as A does in the former and \perp in the latter. Let α_B^Q and β_B^Q be the output probability and the conditional success probability, respectively when playing \tilde{G}_B^Q . We also define α_B^P and β_B^P analogously. Then, it holds that $\beta_B^Q = \beta_A^Q$ and $\beta_B^P = \beta_A^P$. The running time of B satisfies $T_B \leq T_A/\alpha$. It follows from the data processing inequality and Lemma 1 that

$$\beta_B^P - \beta_B^Q \leq \sqrt{2T_B} \cdot 2^{-\lambda/2}.$$

Hence, we have

$$2\beta_B^P - 1 \leq 2\beta_B^Q - 1 + \sqrt{\frac{8T_B}{2^\lambda}}.$$

Since $\beta_B^Q = \beta_A^Q$, it holds that

$$2\beta_B^Q - 1 = 2\beta_A^Q - 1 \leq \sqrt{\frac{T_A}{\alpha 2^\lambda}}.$$

It follows from the above inequalities that

$$2\beta_A^P - 1 = 2\beta_B^P - 1 \leq \sqrt{\frac{8T_B}{2^\lambda}} + \sqrt{\frac{T_A}{\alpha 2^\lambda}} \leq (\sqrt{8} + 1) \sqrt{\frac{T_A}{\alpha 2^\lambda}}.$$

Then, we have

$$\frac{T_A}{\alpha (2\beta_A^P - 1)^2} \geq \frac{2^\lambda}{(\sqrt{8} + 1)^2} > 2^{\lambda - 3.874}.$$

If $\alpha = \alpha_A^P$, the above inequality implies that Π^P is $(\lambda - 3.874)$ -bit secure. Otherwise, we have

$$\alpha_A^Q = \alpha < \frac{T_A}{2^{\lambda - 3.874} (2\beta_A^P - 1)^2} < \frac{2^{\lambda - 6.847} \cdot \alpha_A^P}{2^{\lambda - 3.874}} = 2^{-2.973} \cdot \alpha_A^P,$$

where the last inequality follows from the assumption. In the proof of Theorem 1, if we consider the event that A outputs values other than \perp instead of the event that A succeeds, it implies that $\alpha_A^P < 2^{2.973} \cdot \alpha_A^Q$ for $2^{-\lambda/2}$ -Hellinger close pair (P, Q) . This contradicts the above inequality. Therefore, the statement follows. \square

Limitation of the Statistical Distance

We show that a similar result to Theorem 2 does not hold for the statistical distance. Namely, the closeness of $2^{-\lambda/2}$ in the statistical distance is not sufficient for preserving security.

As a concrete example, we consider a modified one-time pad encryption scheme Π^Q . The probabilistic encryption function for messages over $\{0, 1\}^\lambda$ is defined to be

$$\text{Enc}_k(m) = \begin{cases} (1, m) & \text{with probability } 2^{-\lambda} \\ (0, m \oplus k) & \text{with probability } 1 - 2^{-\lambda} \end{cases},$$

where $k \in \{0, 1\}^\lambda$ is a key sampled according to a distribution Q . Here we assume that Q is the uniform distribution over $\{0, 1\}^\lambda$. Consider a distinguishing game in which, for a random secret $b \in \{0, 1\}$, an attacker tries to predict b given m_0, m_1 , and $\text{Enc}_k(m_b)$. The attacker can easily find the corresponding message if the first bit of the ciphertext is 1. Otherwise, the scheme is perfectly secure, and thus the attacker has no advantage in the distinguishing game. Let A be an attacker such that given $m_0, m_1, \text{Enc}_k(m_b) = (c_1, c_2)$, where $c_1 \in \{0, 1\}, c_2 \in \{0, 1\}^\lambda$, A outputs b such that $c_2 = m_b$ if $c_1 = 1$, and \perp otherwise. Then,

$$\frac{T_A}{\alpha_A^Q (2\beta_A^Q - 1)^2} \geq \frac{T_A}{2^{-\lambda}} \geq 2^\lambda.$$

Since other adversaries cannot achieve a higher advantage than $2^{-\lambda}$, Π^Q has λ -bit security.

Let P be a distribution over $\{0, 1\}^\lambda$ such that

$$P(x) = \begin{cases} 2^{-\lambda} + 2^{-\lambda/2} & x = 0^\lambda \\ 0 & x \in S \\ 2^{-\lambda} & \text{otherwise} \end{cases},$$

where $|S| = 2^{\lambda/2}$. One may consider S a set of strings starting with $1^{\lambda/2}$. It holds that $\text{SD}(P, Q) = 2^{-\lambda/2}$. Consider an adversary A' such that when $c_1 = 1$, A' outputs b satisfying $c_2 = m_b$. When $c_1 = 0$, A' outputs b such that $c_2 = m_b$ if $c_2 \in \{m_0, m_1\}$, and \perp if $c_2 \notin \{m_0, m_1\}$. For this adversary A' , it is not difficult to see that $\alpha_{A'}^P = 2^{-\lambda} + (1 - 2^{-\lambda})(2^{-\lambda} + 2^{-\lambda/2}) \geq 2^{-\lambda/2}$ and $\beta_{A'}^P = 1$. Thus, the bit security of Π^P is at most $\lambda/2$. This indicates that the closeness of $2^{-\lambda/2}$ in the statistical distance may reduce the bit security by half.

4 Relations between Max-Log Distance and Hellinger Distance

We show that the Hellinger distance is bounded above by the max-log distance when the max-log distance is less than $\sqrt{2} - 1$. Namely, the Hellinger distance gives a tighter evaluation of closeness when the distance is small.

Proposition 1. Let P and Q be distributions over Ω with the same support $S \subseteq \Omega$. Then, $\text{HD}(P, Q) \leq \text{ML}(P, Q)$ as long as $\text{ML}(P, Q) \leq \sqrt{2} - 1$.

Proof. It follows from the relation between the Hellinger distance and the chi-square divergence (cf. [18]) that

$$\text{HD}(P, Q) \leq \sqrt{\frac{1}{2} \sum_{x \in S} \frac{(P(x) - Q(x))^2}{Q(x)}},$$

where $S \subseteq \Omega$ is the support of P and Q . Then,

$$\begin{aligned} \text{HD}(P, Q) &\leq \sqrt{\frac{1}{2} \sum_{x \in S} Q(x) \left(\frac{P(x)}{Q(x)} - 1 \right)^2} \\ &\leq \sqrt{\frac{1}{2} \sum_{x \in S} Q(x) \cdot \max_{x \in S} \left| \frac{P(x)}{Q(x)} - 1 \right|^2} \\ &= \frac{1}{\sqrt{2}} \max_{x \in S} \left| \frac{P(x)}{Q(x)} - 1 \right|. \end{aligned}$$

Let $\text{ML}(P, Q) = \epsilon$. By definition, for any $x \in S$,

$$e^{-\epsilon} \leq \frac{P(x)}{Q(x)} \leq e^{\epsilon}.$$

Since we have the relations $e^y - 1 \leq y + y^2$ and $1 - e^{-y} \leq y + y^2$ for $y \in [0, \sqrt{2}]$, it holds that

$$\text{HD}(P, Q) \leq \frac{1}{\sqrt{2}}(\epsilon + \epsilon^2) \leq \epsilon,$$

where the last inequality holds for $0 \leq \epsilon \leq \sqrt{2} - 1$. □

Next, we give a concrete example of distributions for which an exponential gap exists. We show that, for a uniform distribution Q over $\{0, 1\}^n$, there is a distribution P such that $\text{ML}(P, Q) = 0.6$ and $\text{HD}(P, Q) \leq 0.6 \cdot 2^{-n/2}$.

Proposition 2. Let Q be the uniform distribution over Ω with $|\Omega| \geq 4$. There is a distribution P over Ω such that

$$\text{ML}(P, Q) = \epsilon \quad \text{and} \quad \text{HD}(P, Q) \leq \sqrt{\frac{3(\epsilon + \epsilon^2)}{8|\Omega|}}$$

for any $\epsilon \in [0, 0.618]$.

Proof. Let $M = |\Omega|$. We define P such that

$$P(x) = \begin{cases} M^{-1} \cdot e^{\epsilon} & x = x_0 \\ M^{-1} \cdot e^{-\epsilon} & x = x_1 \\ (M - 2)^{-1} \cdot (1 - M^{-1}(e^{\epsilon} + e^{-\epsilon})) & x \notin \{x_0, x_1\} \end{cases}.$$

First we show that $\text{ML}(P, Q) = \epsilon$. It is clear from the definition that $|\ln P(x) - \ln Q(x)| = \epsilon$ for $x \in \{x_0, x_1\}$. For $x \notin \{x_0, x_1\}$, we need to show that

$$M^{-1} \cdot e^{-\epsilon} \leq (M-2)^{-1} \cdot (1 - M^{-1}(e^\epsilon + e^{-\epsilon})) \leq M^{-1} \cdot e^\epsilon,$$

which can be rewritten as

$$M \geq \max \left\{ \frac{e^\epsilon - e^{-\epsilon}}{1 - e^{-\epsilon}}, \frac{e^\epsilon - e^{-\epsilon}}{e^\epsilon - 1} \right\}.$$

Since the right-hand side is at most 4 for $\epsilon \geq 0$, we have $\text{ML}(P, Q) = \epsilon$.

Next, we give an upper bound on $\text{HD}(P, Q)$. Recall that $\text{HD}(P, Q) = \sqrt{1 - \text{HA}(P, Q)}$ and $\text{HA}(P, Q) = \sum_{x \in \Omega} \sqrt{P(x) \cdot Q(x)}$. For $x \notin \{x_0, x_1\}$,

$$P(x) = \frac{1}{M-2} \left(1 - \frac{1}{M}(e^\epsilon + e^{-\epsilon}) \right) \geq \frac{1}{M-2} \left(1 - \frac{1}{M}(2 + \epsilon + \epsilon^2) \right) = \frac{1}{M} \left(1 - \frac{\epsilon + \epsilon^2}{M-2} \right),$$

where the inequality follows from the fact that $e^x + e^{-x} \leq 2 + x + x^2$ for $0 \leq x \leq 1$. By using the relation that $e^x + e^{-x} \geq 2$ for $0 \leq x \leq 1$, we have

$$\begin{aligned} \text{HA}(P, Q) &= \sqrt{P(x_0)Q(x_0)} + \sqrt{P(x_1)Q(x_1)} + \sum_{x \in \Omega \setminus \{x_0, x_1\}} \sqrt{P(x)Q(x)} \\ &\geq \frac{2}{M} + \frac{M-2}{M} \cdot \sqrt{1 - \frac{\epsilon + \epsilon^2}{M-2}}. \end{aligned}$$

Thus,

$$\begin{aligned} \text{HD}(P, Q)^2 &\leq 1 - \frac{2}{M} - \frac{M-2}{M} \cdot \sqrt{1 - \frac{\epsilon + \epsilon^2}{M-2}} \\ &\leq 1 - \frac{2}{M} - \frac{M-2}{M} \left(1 - \frac{\epsilon + \epsilon^2}{2(M-2)} - \frac{1}{2} \left(\frac{\epsilon + \epsilon^2}{M-2} \right)^2 \right) \\ &= \frac{\epsilon + \epsilon^2}{2M} \left(1 + \frac{\epsilon + \epsilon^2}{2(M-2)} \right) \leq \frac{3(\epsilon + \epsilon^2)}{8M}, \end{aligned}$$

where the second and the last inequalities follow from $\sqrt{1-x} \geq 1 - x/2 - x^2/2$ for $0 \leq x \leq 1$ and $(\epsilon + \epsilon^2)/(2(y-2)) \leq 1/4$ for $\epsilon \in [0, 0.618]$ and $y \geq 4$, respectively. Hence, the statement follows. \square

5 Randomness Extraction for Hellinger Distance

We focus on the problem of randomness extraction from entropy sources. The *min-entropy* of random variable X over $\{0, 1\}^n$ is $H_{\min}(X) = \min_{x \in \{0, 1\}^n} \log_2(1/\Pr[X = x])$. *Randomness extractors* are usually defined as a seeded function that maps any entropy source to a distribution that is close to the uniform distribution in the statistical distance. For $n \in \mathbb{N}$, we denote by U_n the uniform distribution over $\{0, 1\}^n$.

Definition 2 (Randomness Extractor). A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is said to be a (k, ϵ) -*(strong) extractor* if for every distribution X over $\{0, 1\}^n$ of $H_{\min}(X) \geq k$, it holds that $\text{SD}(\text{Ext}(X, U_d), U_{m+d}) \leq \epsilon$, where X and U_d are independent.

For (strong) extractors, the input entropy is $k + d$, and the output length is $m + d$. The difference $(k + d) - (m + d) = k - m$ is called the *entropy loss* of extractors. The entropy loss is unavoidable. Radhakrishnan and Ta-Shma [32] showed that it must be at least $2 \log(1/\epsilon) - O(1)$.

It is known that a *universal family of hash functions* gives an extractor with optimal entropy loss. A random hash function $H: \{0, 1\}^n \rightarrow \{0, 1\}^m$ from a family \mathcal{H} of hash functions is called *universal* if for any

distinct $x, x' \in \{0, 1\}^n$, $\Pr[H(x) = H(x')] \leq 2^{-m}$. Specifically, let $|\mathcal{H}| = 2^d$ and $m = k - 2 \log(1/\epsilon)$. Then, extractor Ext defined by $\text{Ext}(x, H) = H(x)$ is a $(k, \epsilon/2)$ -strong extractor. This result is known as the *leftover hash lemma* [6, 20]. The main technical lemma is a bound on the *collision probability*. For a random variable X , the collision probability of X is

$$\text{cp}(X) = \Pr[X = X'] = \sum_x \Pr[X = x]^2,$$

where X' is an independent copy of X .

Lemma 3 (The Leftover Hash Lemma [6, 20]). Let X be a random variable over $\{0, 1\}^n$ with $H_{\min}(X) \geq k$. Let $H: \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a random hash function from a universal family \mathcal{H} . Then, $\text{cp}(H(X), H) \leq 2^{-d} \cdot (2^{-m} + 2^{-k})$.

We define a notion of extractors for which the output distribution is close to uniform in the Hellinger distance.

Definition 3 (Hellinger extractor). A function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is said to be a (k, ϵ) -*(strong) Hellinger extractor* if for every distribution X over $\{0, 1\}^n$ of $H_{\min}(X) \geq k$, it holds that $\text{HD}((\text{Ext}(X, U_d), U_d), U_{m+d}) \leq \epsilon$, where X and U_d are independent.

It follows from (2) that if Ext is a (k, ϵ) -Hellinger extractor, then it is also a $(k, \sqrt{2}\epsilon)$ -extractor.

We use the following useful lemma of Chung and Vadhan [9] for proving a leftover hash lemma for the Hellinger distance.

Lemma 4 ([9, Lemma 3.12]). Let X be a random variable over $\{0, 1\}^n$. If $\text{cp}(X) \leq \alpha/2^n$, then $\text{HA}(X, U_n) \geq \sqrt{1/\alpha}$.

We show that a universal family of hash functions gives a Hellinger extractor with optimal entropy loss.

Theorem 3 (Leftover Hash Lemma for Hellinger). Let $H: \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a random hash function from a universal family \mathcal{H} with $|\mathcal{H}| = 2^d$, $m = k + 1 - 2 \log(1/\epsilon)$. Then, function $\text{Ext}: \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ defined by $\text{Ext}(x, H) = H(x)$ is a (k, ϵ) -Hellinger extractor.

Proof. Let X be a random variable over $\{0, 1\}^n$ with $H_{\min}(X) \geq k$. It follows from Lemma 3 that

$$\text{cp}(H(X), H) \leq 2^{-d} \cdot (2^{-m} + 2^{-k}) = \frac{\alpha}{2^{m+d}},$$

where $\alpha = 1 + 2^{m-k}$. By Lemma 4, we have that $\text{HA}((H(X), H), U_{m+d}) \geq \alpha^{-1/2}$. Then, it holds that

$$\begin{aligned} \text{HD}((H(X), H), U_{m+d}) &= \sqrt{1 - \text{HA}((H(X), H), U_{m+d})} \\ &\leq \sqrt{1 - \alpha^{-1/2}} = \sqrt{1 - (1 + 2^{m-k})^{-1/2}} \\ &\leq \sqrt{1 - (1 - 2^{m-k-1})} = \sqrt{2^{m-k-1}} = \epsilon, \end{aligned}$$

where the last inequality follows from the fact that $(1+x)^{-1/2} \geq 1 - x/2$ for $x \geq 0$. Hence, the statement follows. \square

Since we have the relation that $\text{SD}(P, Q) \leq \sqrt{2} \cdot \text{HD}(P, Q)$, the lower bound of [32] implies that the entropy loss of Theorem 3 is also optimal.

Entropy Loss of Randomness Extractors in Security Games

We consider the situations in which a uniform distribution is employed in security games, and we would like to replace it with an output of randomness extractors. Let Π be a primitive with an n -bit security game G_A^Q such that the uniform distribution $Q = U_m$ is employed. Suppose that Π has λ -bit security.

Theorems 1 and 2 imply that for preserving the bit security when replacing Q with P , it is enough to hold $\text{HD}(P, Q) \leq 2^{-\lambda/2}$. Regarding the statistical distance, the closeness of $2^{-\lambda}$ is sufficient for preserving security.

A universal family of hash functions can achieve the security of extractors for both distances. When using the statistical distance, the entropy loss for achieving $\text{SD}(P, Q) \leq 2^{-\lambda}$ is $k - m = 2(\lambda - 1)$. By Theorem 3, the entropy loss for $\text{HD}(P, Q) \leq 2^{-\lambda/2}$ is $k - m = \lambda - 1$. Thus, by analyzing security games via the Hellinger distance, the entropy loss for preserving λ -bit security can be reduced by half.

Acknowledgments

This work was supported in part by JSPS Grants-in-Aid for Scientific Research Numbers 16H01705, 17H01695, and 18K11159.

References

- [1] Agrawal, R.: Samplers and extractors for unbounded functions. In: Achlioptas, D., Végé, L.A. (eds.) Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, AP-PROX/RANDOM 2019, September 20-22, 2019, Massachusetts Institute of Technology, Cambridge, MA, USA. LIPIcs, vol. 145, pp. 59:1–59:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
- [2] Bai, S., Lepoint, T., Roux-Langlois, A., Sakzad, A., Stehlé, D., Steinfeld, R.: Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. *J. Cryptology* **31**(2), 610–640 (2018)
- [3] Barak, B., Dodis, Y., Krawczyk, H., Pereira, O., Pietrzak, K., Standaert, F., Yu, Y.: Leftover hash lemma, revisited. In: Rogaway, P. (ed.) Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6841, pp. 1–20. Springer (2011)
- [4] Barak, B., Halevi, S.: A model and architecture for pseudo-random generation with applications to /dev/random. In: Atluri, V., Meadows, C.A., Juels, A. (eds.) Proceedings of the 12th ACM Conference on Computer and Communications Security, CCS 2005, Alexandria, VA, USA, November 7-11, 2005. pp. 203–212. ACM (2005)
- [5] Barak, B., Shaltiel, R., Tromer, E.: True random number generators secure in a changing environment. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) Cryptographic Hardware and Embedded Systems - CHES 2003, 5th International Workshop, Cologne, Germany, September 8-10, 2003, Proceedings. Lecture Notes in Computer Science, vol. 2779, pp. 166–180. Springer (2003)
- [6] Bennett, C.H., Brassard, G., Robert, J.: How to reduce your enemy’s information (extended abstract). In: Williams, H.C. (ed.) Advances in Cryptology - CRYPTO ’85, Santa Barbara, California, USA, August 18-22, 1985, Proceedings. Lecture Notes in Computer Science, vol. 218, pp. 468–476. Springer (1985)
- [7] Berman, I., Degwekar, A., Rothblum, R.D., Vasudevan, P.N.: Statistical difference beyond the polarizing regime. *Electronic Colloquium on Computational Complexity (ECCC)* **26**, 38 (2019)

- [8] Boyen, X., Dodis, Y., Katz, J., Ostrovsky, R., Smith, A.D.: Secure remote authentication using biometric data. In: Cramer, R. (ed.) *Advances in Cryptology - EUROCRYPT 2005*, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings. *Lecture Notes in Computer Science*, vol. 3494, pp. 147–163. Springer (2005)
- [9] Chung, K., Vadhan, S.P.: Tight bounds for hashing block sources. In: Goel, A., Jansen, K., Rolim, J.D.P., Rubinfeld, R. (eds.) *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, 11th International Workshop, APPROX 2008, and 12th International Workshop, RANDOM 2008, Boston, MA, USA, August 25-27, 2008. Proceedings. *Lecture Notes in Computer Science*, vol. 5171, pp. 357–370. Springer (2008)
- [10] Dai, W., Hoang, V.T., Tessaro, S.: Information-theoretic indistinguishability via the chi-squared method. In: Katz, J., Shacham, H. (eds.) *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part III. *Lecture Notes in Computer Science*, vol. 10403, pp. 497–523. Springer (2017)
- [11] Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.D.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.* **38**(1), 97–139 (2008)
- [12] Dodis, Y., Pietrzak, K., Wichs, D.: Key derivation without entropy waste. In: Nguyen, P.Q., Oswald, E. (eds.) *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Copenhagen, Denmark, May 11-15, 2014. Proceedings. *Lecture Notes in Computer Science*, vol. 8441, pp. 93–110. Springer (2014)
- [13] Dodis, Y., Yu, Y.: Overcoming weak expectations. In: Sahai, A. (ed.) *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013*, Tokyo, Japan, March 3-6, 2013. Proceedings. *Lecture Notes in Computer Science*, vol. 7785, pp. 1–22. Springer (2013)
- [14] Dwork, C.: Differential privacy: A survey of results. In: Agrawal, M., Du, D., Duan, Z., Li, A. (eds.) *Theory and Applications of Models of Computation*, 5th International Conference, TAMC 2008, Xi’an, China, April 25-29, 2008. Proceedings. *Lecture Notes in Computer Science*, vol. 4978, pp. 1–19. Springer (2008)
- [15] Dwork, C., McSherry, F., Nissim, K., Smith, A.D.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006*, New York, NY, USA, March 4-7, 2006, Proceedings. *Lecture Notes in Computer Science*, vol. 3876, pp. 265–284. Springer (2006)
- [16] Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* **9**(3-4), 211–407 (2014)
- [17] Gennaro, R., Krawczyk, H., Rabin, T.: Secure hashed diffie-hellman over non-ddh groups. In: Cachin, C., Camenisch, J. (eds.) *Advances in Cryptology - EUROCRYPT 2004*, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings. *Lecture Notes in Computer Science*, vol. 3027, pp. 361–381. Springer (2004)
- [18] Gibbs, A.L., Su, F.E.: On choosing and bounding probability metrics. *INTERNAT. STATIST. REV.* pp. 419–435 (2002)
- [19] Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: Johnson, D.S. (ed.) *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, May 14-17, 1989, Seattle, Washington, USA. pp. 25–32. ACM (1989). <https://doi.org/10.1145/73007.73010>, <https://doi.org/10.1145/73007.73010>
- [20] Impagliazzo, R., Levin, L.A., Luby, M.: Pseudo-random generation from one-way functions (extended abstracts). In: Johnson, D.S. (ed.) *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, May 14-17, 1989, Seattle, Washington, USA. pp. 12–24. ACM (1989)

- [21] Krawczyk, H.: Cryptographic extraction and key derivation: The HKDF scheme. In: Rabin, T. (ed.) *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference*, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings. *Lecture Notes in Computer Science*, vol. 6223, pp. 631–648. Springer (2010)
- [22] Levin, L.A.: Randomness and non-determinism. *Journal of Symbolic Logic* **58**(3), 1102–1103 (1993). <https://doi.org/10.1137/S0895480197329508>
- [23] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings. *Lecture Notes in Computer Science*, vol. 6110, pp. 1–23. Springer (2010)
- [24] Matsuda, T., Takahashi, K., Murakami, T., Hanaoka, G.: Improved security evaluation techniques for imperfect randomness from arbitrary distributions. In: Lin, D., Sako, K. (eds.) *Public-Key Cryptography - PKC 2019 - 22nd IACR International Conference on Practice and Theory of Public-Key Cryptography*, Beijing, China, April 14-17, 2019, Proceedings, Part I. *Lecture Notes in Computer Science*, vol. 11442, pp. 549–580. Springer (2019)
- [25] Micciancio, D., Peikert, C.: Hardness of SIS and LWE with small parameters. In: Canetti, R., Garay, J.A. (eds.) *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference*, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I. *Lecture Notes in Computer Science*, vol. 8042, pp. 21–39. Springer (2013)
- [26] Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.* **37**(1), 267–302 (2007)
- [27] Micciancio, D., Walter, M.: Gaussian sampling over the integers: Efficient, generic, constant-time. In: Katz, J., Shacham, H. (eds.) *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference*, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II. *Lecture Notes in Computer Science*, vol. 10402, pp. 455–485. Springer (2017)
- [28] Micciancio, D., Walter, M.: On the bit security of cryptographic primitives. In: Nielsen, J.B., Rijmen, V. (eds.) *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I. *Lecture Notes in Computer Science*, vol. 10820, pp. 3–28. Springer (2018)
- [29] Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Mitzenmacher, M. (ed.) *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, Bethesda, MD, USA, May 31 - June 2, 2009. pp. 333–342. ACM (2009)
- [30] Pöppelmann, T., Ducas, L., Güneysu, T.: Enhanced lattice-based signatures on reconfigurable hardware. In: Batina, L., Robshaw, M. (eds.) *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop*, Busan, South Korea, September 23-26, 2014. Proceedings. *Lecture Notes in Computer Science*, vol. 8731, pp. 353–370. Springer (2014)
- [31] Prest, T.: Sharper bounds in lattice-based cryptography using the rényi divergence. In: Takagi, T., Peyrin, T. (eds.) *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security*, Hong Kong, China, December 3-7, 2017, Proceedings, Part I. *Lecture Notes in Computer Science*, vol. 10624, pp. 347–374. Springer (2017)
- [32] Radhakrishnan, J., Ta-Shma, A.: Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM J. Discrete Math.* **13**(1), 2–24 (2000)
- [33] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, Baltimore, MD, USA, May 22-24, 2005. pp. 84–93. ACM (2005)

- [34] Regev, O.: The learning with errors problem (invited survey). In: Proceedings of the 25th Annual IEEE Conference on Computational Complexity, CCC 2010, Cambridge, Massachusetts, USA, June 9-12, 2010. pp. 191–204. IEEE Computer Society (2010)
- [35] Sahai, A., Vadhan, S.P.: Manipulating statistical difference. In: Pardalos, P.M., Rajasekaran, S., Rolim, J. (eds.) Randomization Methods in Algorithm Design, Proceedings of a DIMACS Workshop, Princeton, New Jersey, USA, December 12-14, 1997. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 43, pp. 251–270. DIMACS/AMS (1997). <https://doi.org/10.1090/dimacs/043/14>, <https://doi.org/10.1090/dimacs/043/14>
- [36] Shaltiel, R.: An introduction to randomness extractors. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part II. Lecture Notes in Computer Science, vol. 6756, pp. 21–41. Springer (2011)
- [37] Skorski, M.: Lower bounds on key derivation for square-friendly applications. In: Vollmer, H., Vallée, B. (eds.) 34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany. LIPIcs, vol. 66, pp. 57:1–57:12. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017). <https://doi.org/10.4230/LIPIcs.STACS.2017.57>, <https://doi.org/10.4230/LIPIcs.STACS.2017.57>
- [38] Steinberger, J.P.: Improved security bounds for key-alternating ciphers via hellinger distance. IACR Cryptology ePrint Archive **2012**, 481 (2012)
- [39] Walter, M.: Sampling the integers with low relative error. In: Buchmann, J., Nitaj, A., Rachidi, T. (eds.) Progress in Cryptology - AFRICACRYPT 2019 - 11th International Conference on Cryptology in Africa, Rabat, Morocco, July 9-11, 2019, Proceedings. Lecture Notes in Computer Science, vol. 11627, pp. 157–180. Springer (2019)