

Improved security proofs in lattice-based cryptography: using the Rényi divergence rather than the statistical distance

Shi Bai¹, Adeline Langlois², Tancrede Lepoint³, Damien Stehlé¹,
Ron Steinfeld⁴

¹ ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France.
shi.bai@ens-lyon.fr, damien.stehle@ens-lyon.fr –
<http://perso.ens-lyon.fr/shi.bai>, <http://perso.ens-lyon.fr/damien.stehle>

² EPFL, Lausanne, Switzerland.

adeline.langlois@epfl.ch – <http://lasec.epfl.ch/~alangloi/>

³ CryptoExperts, France.

tancrede.lepoint@cryptoexperts.com –

<https://www.cryptoexperts.com/tlepoint/>

⁴ Clayton School of Information Technology, Monash University, Clayton, Australia.
ron.steinfeld@monash.edu – <http://users.monash.edu.au/~rste/>

Abstract. The Rényi divergence is a measure of closeness of two probability distributions. We show that it can often be used as an alternative to the statistical distance in security proofs for lattice-based cryptography. Using the Rényi divergence is particularly suited for security proofs of primitives in which the attacker is required to solve a search problem (e.g., forging a signature). We show that it may also be used in the case of distinguishing problems (e.g., semantic security of encryption schemes), when they enjoy a public sampleability property. The techniques lead to security proofs for schemes with smaller parameters, and sometimes to simpler security proofs than the existing ones.

1 Introduction

Let D_1 and D_2 be two non-vanishing probability distributions over a common measurable support X . Let $a \in (1, +\infty)$. The *Rényi divergence* [Rén61,EH12] (RD for short) $R_a(D_1||D_2)$ of order a between D_1 and D_2 is defined as the expected value of $(D_1(x)/D_2(x))^{a-1}$ over the randomness of x sampled from D_1 . For notational convenience, our definition of the RD is the exponential of the classical definition [EH12]. The RD is an alternative to the statistical distance (SD for short) $\Delta(D_1, D_2) = \frac{1}{2} \sum_{x \in X} |D_1(x) - D_2(x)|$ as measure of distribution closeness, where we replace the difference in SD, by the ratio in RD. RD enjoys several properties that are analogous of those enjoyed by SD, where addition in the property of SD is replaced by multiplication in the analogous property of RD (see Subsection 2.3).

SD is ubiquitous in cryptographic security proofs. One of its most useful properties is the so-called *probability preservation property*: For any measurable

event $E \subseteq X$, we have $D_2(E) \geq D_1(E) - \Delta(D_1, D_2)$. RD enjoys the analogous property $D_2(E) \geq D_1(E)^{\frac{a}{a-1}} / R_a(D_1 \| D_2)$. If the event E occurs with significant probability under D_1 , and if the SD (resp. RD) is small, then the event E also occurs with significant probability under D_2 . These properties are particularly handy when the success of an attacker against a given scheme can be described as an event whose probability should be negligible, e.g., the attacker outputs a new valid message-signature pair for a signature scheme. If in the attacker succeeds with good probability in the real scheme based on distribution D_1 , then it also succeeds with good probability in the simulated scheme (of the security proof) based on distribution D_2 .

To make the SD probability preservation property useful, it must be ensured that the SD $\Delta(D_1, D_2)$ is smaller than any $D_1(E)$ that the security proof must handle. Typically, the quantity $D_1(E)$ is assumed to be greater than some success probability lower bound ε , which is of the order of $1/\text{poly}(\lambda)$ where λ refers to the security parameter, or even $2^{-o(\lambda)}$ if the proof handles attackers whose success probabilities can be sub-exponentially small (which we believe better reflects practical objectives). As a result, the SD $\Delta(D_1, D_2)$ must be $< \varepsilon$ for the SD probability preservation property to be relevant. Similarly, the RD probability preservation property is non-vacuous when the RD $R_a(D_1 \| D_2)$ is $\leq \text{poly}(1/\varepsilon)$. In many cases, the latter seems less demanding than the former: in all our applications of RD, the RD between D_1 and D_2 is small while their SD is too large for the SD probability preservation to be applicable. In fact, as we will see in Subsection 2.3, the RD becomes sufficiently small to be useful before the SD when $\sup_x D_1(x)/D_2(x)$ tends to 1. This explains the superiority of the RD in several of our applications.

Although RD seems more amenable than SD for search problems, it seems less so for distinguishing problems. A typical cryptographic example is semantic security of an encryption scheme. Semantic security requires an adversary \mathcal{A} to distinguish between the encryption distributions of two plaintext messages of its choosing: the distinguishing advantage $\text{Adv}_{\mathcal{A}}(D_1, D_2)$, defined as the difference of probabilities that \mathcal{A} outputs 1 using D_1 or D_2 , should be large. In security proofs, algorithm \mathcal{A} is often called on distributions D'_1 and D'_2 that are close to D_1 and D_2 (respectively). If the SDs between D_1 and D'_1 and D_2 and D'_2 are both bounded from above by ε , then, by the SD probability preservation property (used twice), we have $\text{Adv}_{\mathcal{A}}(D'_1, D'_2) \geq \text{Adv}_{\mathcal{A}}(D_1, D_2) - 2\varepsilon$. As a result, SD can be used to distinguishing problems in a similar fashion as for search problems. The multiplicativity of the RD probability preservation property seems to prevent RD from being applicable to distinguishing problems.

We replace the statistical distance by the Rényi divergence in several security proofs for lattice-based cryptographic primitives. *Lattice-based cryptography* is a relatively recent cryptographic paradigm in which cryptographic primitives are shown at least as secure as it is hard to solve standard problems over lattices (see the survey [MR09]). Security proofs in lattice-based cryptography involve different types of distributions, often over infinite sets, such as continuous Gaussian distributions and Gaussian distributions with lattice supports. The RD seems

particularly well suited to quantify the closeness of Gaussian distributions. Consider for example two continuous distributions over the reals, both with standard deviation 1, but one with center 0 and the other one with center c . Their SD is linear in c , so that c must remain extremely small for the SD probability preservation property to be useful. On the other hand, their RD of order $a = 2$ is bounded as $\exp(O(c^2))$ so that the RD preservation property remains useful even for slightly growing c .

RD was first used in lattice-based cryptography by [LPR13], in the decision to search reduction for the Ring Learning With Errors problem (which serves as a security foundation for many asymptotically fast primitives). It was then exploited in [LSS14] to decrease the parameters of the Garg et al. (approximation to) cryptographic multilinear maps [GGH13]. In the present work, we present a more extensive study of the power of RD in lattice-based cryptography, by showing several independent applications of RD. In some cases, it leads to security proofs allowing to take smaller parameters in the cryptographic schemes, hence leading to efficiency improvements. In other cases, this leads to alternative security proofs that are conceptually simpler.

Our applications of RD also include distinguishing problems. To circumvent the aforementioned a priori limitation of the RD probability preservation property for distinguishing problems, we propose an alternative approach that handles a class of distinguishing problems, enjoying a special property that we call public sampleability. This public sampleability allows to estimate success probabilities via Hoeffding’s bound.

The applications we show in lattice-based cryptography are as follows:

- Smaller signatures for the Hash-and-Sign GPV signature scheme [GPV08].
- Smaller storage requirement for the Fiat-Shamir BLISS signature scheme [DDLL13,PDG14,Duc14].
- Alternative proof that the Learning With Errors (LWE) problem with noise chosen uniformly in an interval is no easier than the Learning With Errors problem with Gaussian noise [DMQ13]. Our reduction does not require the latter problem to be hard, and it is hence marginally more general as it also applies to distributions with smaller noises. Further, our reduction preserves the LWE dimension n , and is hence much tighter than the one from [DMQ13] (the latter degrades the LWE dimension by a $\log n$ factor).
- Smaller parameters in the dual-Regev encryption scheme from [GPV08].

We think RD is likely to have further applications in lattice-based cryptography, for search and for distinguishing problems.

Related works. The framework for using RD in distinguishing problems was used in [LPSS14], in the context of the k -LWE problem (a variant of LWE in which the attacker is given extra information). In [PDG14], Pöppelmann, Ducas and Güneysu used the Kullback-Leibler divergence (which is the RD of order $a = 1$) to lower the storage requirement of [DDLL13]. Asymptotically, using the Kullback-Leibler divergence rather than SD only leads to a constant factor improvement. Our approach allows bigger savings in the case where the number of signature queries is limited, as explained in Section 3.

Road-map. In Section 2, we provide necessary background on lattice-based cryptography, and on the Rényi divergence. In Section 3, we use RD to improve signature scheme parameters. Section 4 contains the description of the framework in which we can use RD for distinguishing problems, which we apply to the dual-Reggev encryption scheme. In Section 5, we describe an alternative hardness proof for LWE with noise uniformly chosen in an interval.

Notations. If x is a real number, we let $\lfloor x \rfloor$ denote a closest integer to x . The notation \ln refers to the natural logarithm and the notation \log refers to the base 2 logarithm. We define $\mathbb{T} = ([0, 1], +)$. For an integer q , we let \mathbb{Z}_q denote the ring of integers modulo q . We let \mathbb{T}_q denote the group $\mathbb{T}_q = \{i/q \bmod 1 : i \in \mathbb{Z}\} \subseteq \mathbb{T}$. Vectors are denoted in bold. If \mathbf{b} is a vector in \mathbb{R}^d , we let $\|\mathbf{b}\|$ denote its Euclidean norm. By default, all our vectors are column vectors.

If D is a probability distribution, we let $\text{Supp}(D) = \{x : D(x) \neq 0\}$ denote its support. For a set X of finite weight, we let $U(X)$ denote the uniform distribution on X . The statistical distance between two distributions D_1 and D_2 over a countable support X is $\Delta(D_1, D_2) = \frac{1}{2} \sum_{x \in X} |D_1(x) - D_2(x)|$. This definition is extended in the natural way to continuous distributions. If $f : X \rightarrow \mathbb{R}$ takes non-negative values, then for all countable $Y \subseteq X$, we define $f(Y) = \sum_{y \in Y} f(y) \in [0, +\infty]$. For any vector $\mathbf{c} \in \mathbb{R}^n$ and any real $s > 0$, the (spherical) Gaussian function with standard deviation parameter s and center \mathbf{c} is defined as follows: $\forall \mathbf{x} \in \mathbb{R}^n, \rho_{s, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / s^2)$. The Gaussian distribution is $D_{s, \mathbf{c}} = \rho_{s, \mathbf{c}} / s^n$. When $\mathbf{c} = \mathbf{0}$, we may omit the subscript \mathbf{c} .

We use the usual Landau notations. A function $f(\lambda)$ is said negligible if it is $\lambda^{-\omega(1)}$. A probability $p(\lambda)$ is said overwhelming if it is $1 - \lambda^{-\omega(1)}$.

The distinguishing advantage of an algorithm \mathcal{A} between two distributions D_0 and D_1 is defined as $\text{Adv}_{\mathcal{A}}(D_0, D_1) = |\Pr_{x \leftarrow D_0}[\mathcal{A}(x) = 1] - \Pr_{x \leftarrow D_1}[\mathcal{A}(x) = 1]|$, where the probabilities are taken over the randomness of the input x and the internal randomness of \mathcal{A} . Algorithm \mathcal{A} is said to be an (ε, T) -distinguisher if it runs in time $\leq T$ and if $\text{Adv}_{\mathcal{A}}(D_0, D_1) \geq \varepsilon$.

2 Preliminaries

We assume the reader is familiar with standard cryptographic notions, as well as with lattices and lattice-based cryptography. We refer to [Reg09a, MR09] for introductions on the latter topic.

2.1 Lattices

A (full-rank) n -dimensional *Euclidean lattice* $\Lambda \subseteq \mathbb{R}^n$ is the set of all integer linear combinations $\sum_{i=1}^n x_i \mathbf{b}_i$ of some \mathbb{R} -basis $(\mathbf{b}_i)_{1 \leq i \leq n}$ of \mathbb{R}^n . In this setup, the tuple $(\mathbf{b}_i)_i$ is said to form a \mathbb{Z} -basis of Λ . For a lattice Λ and any $i \leq n$, the i th successive minimum $\lambda_i(\Lambda)$ is the smallest radius r such that Λ contains i linearly independent vectors of norm at most r . The dual Λ^* of a lattice Λ is defined as $\Lambda^* = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{y}^t \Lambda \subseteq \mathbb{Z}^n\}$.

The (spherical) *discrete Gaussian distribution* over a lattice $\Lambda \subseteq \mathbb{R}^n$, with standard deviation parameter $s > 0$ and center \mathbf{c} is defined as:

$$\forall \mathbf{x} \in \Lambda, D_{\Lambda, s, \mathbf{c}} = \frac{\rho_{s, \mathbf{c}}(\mathbf{x})}{\rho_{s, \mathbf{c}}(\Lambda)}.$$

When the center is $\mathbf{0}$, we omit the subscript \mathbf{c} .

The *smoothing parameter* [MR07] of an n -dimensional lattice Λ with respect to $\varepsilon > 0$, denoted by $\eta_\varepsilon(\Lambda)$, is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. We use the following properties.

Lemma 2.1 ([MR07, Le. 3.3]). *Let Λ be an n -dimensional lattice and $\varepsilon > 0$. Then*

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1 + 1/\varepsilon))}{\pi}} \cdot \lambda_n(\Lambda).$$

Lemma 2.2 (Adapted from [GPV08, Le. 5.3]). *Let $m, n \geq 1$ and q a prime integer, with $m \geq 2n \ln q$. For $A \in \mathbb{Z}_q^{n \times m}$ we define A^\perp as the lattice $\{\mathbf{x} \in \mathbb{Z}^m : A\mathbf{x} = \mathbf{0} \pmod{q}\}$. Then,*

$$\forall \varepsilon < 1/2 : \Pr_{A \leftarrow U(\mathbb{Z}_q^{n \times m})} \left[\eta_\varepsilon(A^\perp) \geq 4\sqrt{\frac{\ln(4m/\varepsilon)}{\pi}} \right] \leq q^{-n}.$$

Lemma 2.3 (Adapted from [GPV08, Cor. 2.8]). *Let Λ, Λ' be n -dimensional lattices with $\Lambda' \subseteq \Lambda$ and $\varepsilon \in (0, 1/2)$. Then for any $\mathbf{c} \in \mathbb{R}^n$ and $s \geq \eta_\varepsilon(\Lambda')$ and any $x \in \Lambda/\Lambda'$ we have*

$$(D_{\Lambda, s, \mathbf{c}} \pmod{\Lambda'})(x) \in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, \frac{1 + \varepsilon}{1 - \varepsilon} \right] \cdot \frac{\det(\Lambda)}{\det(\Lambda')}.$$

2.2 The SIS and LWE problems

The Small Integer Solution (SIS) problem was introduced by Ajtai in [Ajt96]. It serves as a security foundation for numerous cryptographic primitives, including, among many others, hash functions [Ajt96] and signatures [GPV08, DDLL13].

Definition 2.4. *Let $m \geq n \geq 1$ and $q \geq 2$ be integers, and β a positive real. The $\text{SIS}_{n, m, q, \beta}$ problem is as follows: given $A \leftarrow U(\mathbb{Z}_q^{n \times m})$, the goal is to find $\mathbf{x} \in \mathbb{Z}^m$ such that $A\mathbf{x} = \mathbf{0} \pmod{q}$ and $0 < \|\mathbf{x}\| \leq \beta$.*

The SIS problem was proven by Ajtai [Ajt96] to be at least as hard as some standard worst-case problems over Euclidean lattices, under specific parameter constraints. We refer to [GPV08] for an improved (and simplified) reduction.

The Learning With Errors (LWE) problem was introduced in 2005 by Regev [Reg05, Reg09b]. LWE is also extensively used as a security foundation, for encryption schemes [Reg09b, GPV08], fully homomorphic encryption schemes [BV11], and pseudo-random functions [BPR12, AKPW13], among many others. Its definition involves the following distribution. Let χ be a distribution over \mathbb{T} , $q \geq 2$, $n \geq 1$ and $\mathbf{s} \in \mathbb{Z}_q^n$. A sample from $A_{\mathbf{s}, \chi}$ is of the form $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{T}$, with $\mathbf{a} \leftarrow U(\mathbb{Z}_q^n)$, $b = \frac{1}{q} \langle \mathbf{a}, \mathbf{s} \rangle + e$ and $e \leftarrow \chi$.

Definition 2.5. Let χ be a distribution over \mathbb{T} , $q \geq 2$, and $m \geq n \geq 1$. The search variant $\text{sLWE}_{n,q,\chi,m}$ of the LWE problem is as follows: given m samples from $A_{\mathbf{s},\chi}$ for some $\mathbf{s} \in \mathbb{Z}_q^n$, the goal is to find \mathbf{s} . The decision variant $\text{LWE}_{n,q,\chi,m}$ consists in distinguishing between the distributions $(A_{\mathbf{s},\chi})^m$ and $U(\mathbb{Z}_q^n \times \mathbb{T})^m$, where $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$.

In some cases, it is convenient to use an error distribution χ whose support is \mathbb{T}_q . In the case, the definition of LWE is adapted in that $U(\mathbb{Z}_q^n \times \mathbb{T})$ is replaced by $U(\mathbb{Z}_q^n \times \mathbb{T}_q)$. Note also that for a fixed number of samples m , we can represent the LWE samples using matrices. The \mathbf{a}_i 's form the rows of a matrix A uniform in $\mathbb{Z}_q^{m \times n}$, and the scalar product is represented by the product between A and \mathbf{s} .

Regev [Reg09b] gave a quantum reduction from standard worst-case problems over Euclidean lattices to sLWE and LWE , under specific parameter constraints. Classical (but weaker) reductions have later been obtained (see [Pei09, BLP⁺13]). We will use the following sample-preserving search to decision reduction for LWE .

Theorem 2.6 (Adapted from [MM11, Prop. 4.10]). If $q \leq \text{poly}(m, n)$ is prime and the error distribution χ has support in \mathbb{T}_q , then there exists a reduction from $\text{sLWE}_{n,q,\chi,m}$ to $\text{LWE}_{n,q,\chi,m}$ that is polynomial in n and m .

2.3 The Rényi divergence

For any two discrete probability distributions P and Q such that $\text{Supp}(P) \subseteq \text{Supp}(Q)$ and $a \in (1, +\infty)$, we define the Rényi divergence of order a by

$$R_a(P\|Q) = \left(\sum_{x \in \text{Supp}(P)} \frac{P(x)^a}{Q(x)^{a-1}} \right)^{\frac{1}{a-1}}.$$

We omit the a subscript when $a = 2$. We define the Rényi divergences of orders 1 and $+\infty$ by

$$R_1(P\|Q) = \exp \left(\sum_{x \in \text{Supp}(P)} P(x) \log \frac{P(x)}{Q(x)} \right) \quad \text{and} \quad R_\infty(P\|Q) = \max_{x \in \text{Supp}(P)} \frac{P(x)}{Q(x)}.$$

The definitions are extended in the natural way to continuous distributions. The divergence R_1 is the (exponential of) the Kullback-Leibler divergence.

For any fixed P, Q , the function $a \mapsto R_a(P\|Q) \in (0, +\infty]$ is non-decreasing, continuous over $(1, +\infty)$, tends to $R_\infty(P\|Q)$ when a grows to infinity, and if $R_a(P\|Q)$ is finite for some a , then $R_a(P\|Q)$ tends to $R_1(P\|Q)$ when a tends to 1 (we refer to [EH12] for proofs). A direct consequence is that if $P(x)/Q(x) \leq c$ for all $x \in \text{Supp}(P)$ and for some constant c , then $R_a(P\|Q) \leq R_\infty(P\|Q) \leq c$. In the same setup, we have $\Delta(P, Q) \leq c/2$.

The following properties can be considered the multiplicative analogues of those of the SD. We refer to [EH12, LSS14] for proofs.

Lemma 2.7. *Let $a \in [1, +\infty]$. Let P and Q denote distributions with $\text{Supp}(P) \subseteq \text{Supp}(Q)$. Then the following properties hold:*

- **Log. Positivity:** $R_a(P\|Q) \geq R_a(P\|P) = 1$.
- **Data Processing Inequality:** $R_a(P^f\|Q^f) \leq R_a(P\|Q)$ for any function f , where P^f (resp. Q^f) denotes the distribution of $f(y)$ induced by sampling $y \leftarrow P$ (resp. $y \leftarrow Q$).
- **Multiplicativity:** Assume P and Q are two distributions of a pair of random variables (Y_1, Y_2) . For $i \in \{1, 2\}$, let P_i (resp. Q_i) denote the marginal distribution of Y_i under P (resp. Q), and let $P_{2|1}(\cdot|y_1)$ (resp. $Q_{2|1}(\cdot|y_1)$) denote the conditional distribution of Y_2 given that $Y_1 = y_1$. Then we have:
 - $R_a(P\|Q) = R_a(P_1\|Q_1) \cdot R_a(P_2\|Q_2)$ if Y_1 and Y_2 are independent.
 - $R_a(P\|Q) \leq R_\infty(P_1\|Q_1) \cdot \max_{y_1 \in X} R_a(P_{2|1}(\cdot|y_1)\|Q_{2|1}(\cdot|y_1))$.
- **Probability Preservation:** Let $A \subseteq \text{Supp}(Q)$ be an arbitrary event. If $a \in (1, +\infty)$, then $Q(A) \geq P(A)^{\frac{a}{a-1}}/R_a(P\|Q)$. Further, we have

$$Q(A) \geq P(A)/R_\infty(P\|Q).$$

Let P_1, P_2, P_3 be three distributions with $\text{Supp}(P_1) \subseteq \text{Supp}(P_2) \subseteq \text{Supp}(P_3)$. Then we have:

- **Weak Triangle Inequality:**

$$R_a(P_1\|P_3) \leq \begin{cases} R_a(P_1\|P_2) \cdot R_\infty(P_2\|P_3), \\ R_\infty(P_1\|P_2)^{\frac{a}{a-1}} \cdot R_a(P_2\|P_3) \end{cases} \quad \text{if } a \in (1, +\infty).$$

Getting back to the setup in which $P(x)/Q(x) \leq c$ for all $x \in \text{Supp}(P)$ and for some constant c , the RD probability preservation property above is relevant even for large c , whereas the analogous SD probability preservation property starts making sense only when $c < 2$.

Pinsker's inequality is the analogue of the probability preservation property for $\alpha = 1$: for an arbitrary event $A \subseteq \text{Supp}(Q)$, we have $Q(A) \geq P(A) - \sqrt{\ln R_1(P\|Q)}/2$ (see [PDG14, Le. 1] for a proof). Analogously to the statistical distance, this probability preservation property is useful for unlikely events A only if $\ln R_1(P\|Q)$ is very small. We refer to Subsection 3.1 for additional comments on this property.

2.4 RD bounds

We will use the following result, adapted from [LSS14].

Lemma 2.8. *For any n -dimensional lattice $A \subseteq \mathbb{R}^n$ and $s > 0$, let P be the distribution $D_{A,s,\mathbf{c}}$ and Q be the distribution $D_{A,s,\mathbf{c}'}$ for some fixed $\mathbf{c}, \mathbf{c}' \in \mathbb{R}^n$. If $\mathbf{c}, \mathbf{c}' \in A$, let $\varepsilon = 0$. Otherwise, fix $\varepsilon \in (0, 1)$ and assume that $s \geq \eta_\varepsilon(A)$. Then, for any $a \in (1, +\infty)$:*

$$R_a(P\|Q) \in \left[\left(\frac{1-\varepsilon}{1+\varepsilon} \right)^{\frac{2}{a-1}}, \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^{\frac{2}{a-1}} \right] \cdot \exp \left(a\pi \frac{\|\mathbf{c} - \mathbf{c}'\|^2}{s^2} \right).$$

It may be checked that also $R_1(P\|Q)$ is of the order of $\exp(\|\mathbf{c}' - \mathbf{c}\|^2/s^2)$, $R_\infty(P\|Q) = +\infty$ and $\Delta(P, Q)$ is of the order of $\|\mathbf{c}' - \mathbf{c}\|^2/s^2$. In that setup, the RD of order $a = \infty$ is useless, and the probability preservation properties of the SD and RD of order $a = 1$ lead to interesting bounds for events occurring only when $\|\mathbf{c}' - \mathbf{c}\|/s = o(\varepsilon)$. Oppositely, for any $a \in (1, +\infty)$, the probability preservation property for the RD of order $a \in (1, +\infty)$ may be used with $\|\mathbf{c}' - \mathbf{c}\|/s = O(\sqrt{\log(1/\varepsilon)})$ while still leading to probabilistic lower bounds of the order of $\varepsilon^{O(1)}$.

As we have already seen, if two distributions are close in a uniform sense, then their RD is small. We observe the following immediate consequence of Lemma 2.3, that allows replacing the SD with the RD in the context of smoothing arguments, in order to save on the required parameter s . In applications of Lemma 2.3, it is customary to use $s \geq \eta_\varepsilon(\Lambda')$ with $\varepsilon \leq 2^{-\lambda}$, in order to make the distribution $D_{\Lambda/\Lambda', s, \mathbf{c}} = D_{\Lambda, s, \mathbf{c}} \bmod \Lambda'$ within SD $2^{-\lambda}$ of the uniform distribution $U(\Lambda/\Lambda')$. This translates via Lemma 2.1 to use $s = \Omega(\sqrt{\lambda + \log n} \cdot \lambda_n(\Lambda'))$. Whereas if using an RD bound $R_\infty(D_{\Lambda/\Lambda', s, \mathbf{c}}\|U_{\Lambda/\Lambda'}) = O(1)$ suffices for the application, one can take $\varepsilon = O(1)$ in the corollary below, which translates to just $s = \Omega(\sqrt{\log n} \cdot \lambda_n(\Lambda'))$, saving a factor $\Theta(\sqrt{\lambda})$.

Lemma 2.9. *Let Λ, Λ' be n -dimensional lattices with $\Lambda' \subseteq \Lambda$ and $\varepsilon \in (0, 1/2)$. Let $D_{\Lambda/\Lambda', s, \mathbf{c}}$ denote the distribution on Λ/Λ' induced by sampling from $D_{\Lambda, s, \mathbf{c}}$ and reducing modulo Λ' , and let $U_{\Lambda/\Lambda'}$ denote the uniform distribution on Λ/Λ' . Then for any $\mathbf{c} \in \mathbb{R}^n$ and $s \geq \eta_\varepsilon(\Lambda')$ and any $x \in \Lambda/\Lambda'$ we have*

$$R_\infty(D_{\Lambda/\Lambda', s, \mathbf{c}}\|U_{\Lambda/\Lambda'}) \leq \frac{1 + \varepsilon}{1 - \varepsilon}.$$

3 Application to lattice-based signature schemes

In this section, we use the RD to improve the security proofs of the GPV and BLISS signature schemes [GPV08, DDLL13], allowing to take smaller parameters for any fixed security level.

3.1 Sampling Discrete Gaussians and the BLISS Signature Scheme

We show that the use of RD in place of SD leads significant savings in the required precision of integers sampled according to a discrete Gaussian distribution in the security analysis of lattice-based signature schemes. These savings consequently lower the precomputed table storage for sampling discrete Gaussians with the method described in [DDLL13, PDG14].

Discrete Gaussian Sampling. In the BLISS signature scheme [DDLL13] (and similarly in earlier variants [Lyu12]), each signature requires the signing algorithm to sample $O(n)$ independent integers from the 1-dimensional discrete

Gaussian distribution $D_{\mathbb{Z},s}$, where $s = O(m)$ is the standard deviation parameter (here the variable m denotes a parameter related to the underlying lattice dimension, and is typically in the order of several hundreds).⁵

In [DDLL13], a particularly efficient sampling algorithm for $D_{\mathbb{Z},s}$ is presented. To produce a sample from $D_{\mathbb{Z},s}$, this algorithm samples about $\ell = \lceil \log(0.22s^2(1 + 2\tau s)) \rceil$ Bernoulli random variables of the form $B_{\exp(-\pi 2^i/s^2)}$, where $i = 0, \dots, \ell - 1$ and $\tau = O(\sqrt{\lambda})$ is the tail-cut factor for the Gaussian. To sample those Bernoulli random variables, the authors of [DDLL13] use a precomputed table of the probabilities $c_i = \exp(-\pi 2^i/s^2)$, for $i = 1, \dots, \ell$. Since these probabilities are real numbers, they must be truncated to some bit precision p in the precomputed table, so that truncated values $\tilde{c}_i = c_i + \varepsilon_i$ are stored, where $|\varepsilon_i| \leq 2^{-p}c_i$ are the truncation errors.

In previous works, the precision was determined by an analysis either based on the statistical distance (SD) [DDLL13] or the Kullback-Leibler divergence (KLD) [PDG14]. In this section, we review and complete these methods, and we propose an RD-based analysis that leads to bigger savings, asymptotically and in practice (see Table 1). More precisely, we give sufficient lower bounds on the precision p to ensure security on the scheme implemented with truncated values against adversaries succeeding with probability $\geq \varepsilon$ and making $\leq q_s$ signing queries. For any adversary, the distributions Φ' and Φ denote the signatures in the view of the adversary in the untruncated (resp. truncated) cases.

SD-based analysis [DDLL13]. Any forging adversary \mathcal{A} with success probability $\geq \varepsilon$ on the scheme implemented with truncated Gaussian has a success probability $\varepsilon' \geq \varepsilon - \Delta(\Phi, \Phi')$ against the scheme implemented with perfect Gaussian sampling. We select parameters to handle adversaries with success probabilities $\geq \varepsilon/2$ against the untruncated scheme; we can set the required precision p so that $\Delta(\Phi, \Phi') \leq \varepsilon/2$. Each signature requires $\ell \cdot m$ samples from the Bernoulli random variables $(B_{\tilde{c}_i})_i$. To ensure security against q_s signing queries, each of the truncated Bernoulli random variables $B_{\tilde{c}_i}$ should be within SD $\Delta(\Phi, \Phi')/(\ell \cdot m \cdot q_s)$ of the desired B_{c_i} (by the union bound). Using $\Delta(B_{\tilde{c}_i}, B_{c_i}) = |\varepsilon_i| \leq 2^{-p}c_i \leq 2^{-p-1}$ leads to a precision requirement

$$p \geq \log(\ell \cdot m \cdot q_s / \Delta(\Phi, \Phi')) \geq \log(\ell \cdot m \cdot q_s / \varepsilon). \quad (1)$$

The overall precomputed table is hence of bit-size $L_{\text{SD}} = p \cdot \ell \geq \log(\ell \cdot m \cdot q_s / \varepsilon) \cdot \ell$.

Note that in [DDLL13], the authors omitted the term $\ell \cdot m \cdot q_s$ in their analysis: they only ensured that $\Delta(B_{\tilde{c}_i}, B_{c_i}) \leq \varepsilon$, leading to the requirement that $p \geq \log(1/\varepsilon)$.

One may also set the precision p_i depending on i for $0 \leq i \leq \ell - 1$. It is sufficient to set

$$2^{-p_i}c_i = 2^{-p_i} \exp(-\pi 2^i/s^2) \leq (\varepsilon/2)/(\ell \cdot m \cdot q_s).$$

⁵ Note that [Lyu12, DDLL13] consider the unnormalized Gaussian function $\rho'_{\sigma,c}(\mathbf{x}) = \exp(-\|\mathbf{x} - \mathbf{c}\|/(2\sigma^2))$ instead of $\rho_{s,c}$. We have $\rho_{s,c} = \rho'_{\sigma,c}$ when $\sigma = s/\sqrt{2\pi}$.

Hence the precision p_i is

$$p_i \geq \log \left(\frac{\ell \cdot m \cdot q_s}{\epsilon} \cdot \exp(-\pi 2^i / s^2) \right) + 1. \quad (2)$$

The bit-size of the overall precomputed table can be computed as a sum of the above p_i 's. Using the symmetry of the Bernoulli variable, we can further drop the bit-size of the precomputed table.

KLD-based analysis [PDG14]. In [PDG14], Pöppelman, Ducas and Güneysu replace the SD-based analysis by a KLD-based analysis (i.e., using the RD of order $a = 1$) to reduce the precision p needed in the precomputed table. They show that any forging adversary \mathcal{A} with success probability ϵ on the scheme implemented with truncated Gaussian has a success probability $\epsilon' \geq \epsilon - \sqrt{\ln R_1(\Phi \|\Phi')}/2$ on the scheme implemented with perfect Gaussian (see remark at the end of Subsection 2.3). By the multiplicative property of the RD over the $\ell \cdot m \cdot q_s$ independent samples needed for signing q_s times, we get that $R_1(\Phi \|\Phi') \leq (\max_{i=1, \dots, \ell} R_1(B_{\tilde{c}_i} \| B_{c_i}))^{\ell \cdot m \cdot q_s}$. Now, we have:

$$\begin{aligned} \ln R_1(B_{\tilde{c}_i} \| B_{c_i}) &= (1 - c_i - \epsilon_i) \ln \frac{1 - c_i - \epsilon_i}{1 - c_i} + (c_i + \epsilon_i) \ln \frac{c_i + \epsilon_i}{c_i} \\ &\leq -(1 - c_i - \epsilon_i) \frac{\epsilon_i}{1 - c_i} + (c_i + \epsilon_i) \frac{\epsilon_i}{c_i} = \frac{\epsilon_i^2}{(1 - c_i)c_i}. \end{aligned}$$

Using $|\epsilon_i| \leq 2^{-p}c_i$ and $1 - c_i \geq 1/2$, we obtain $\ln R_1(B_{\tilde{c}_i} \| B_{c_i}) = 2^{-2p+1} \frac{c_i}{1 - c_i} \leq 2^{-2p}$. Therefore, we obtain $\epsilon' \geq \epsilon - \sqrt{\ell \cdot m \cdot q_s \cdot 2^{-2p}}$. We can select parameters such that $\sqrt{\ell \cdot m \cdot q_s \cdot 2^{-2p+1}} \leq \epsilon/2$. This leads to a precision requirement

$$p \geq \frac{1}{2} \log \left(\frac{\ell \cdot m \cdot q_s}{\epsilon^2} \right) + \frac{1}{2}. \quad (3)$$

The overall precomputed table is hence of bit-size $L_{\text{KLD}} \geq (\log(\ell \cdot m \cdot q_s / \epsilon^2) / 2 + 1/2) \cdot \ell$. This KLD-based analysis may save some storage if ϵ is not too small.

Note that in [PDG14], the authors selected $\epsilon = 1/2$ and $\ell \cdot m \cdot q_s = 2^\lambda$ where λ is the desired bit-security, and hence obtained $p \geq \lambda/2 + 1$.

One may also set the precision p_i depending on i . It is sufficient to set

$$p_i \geq \frac{1}{2} \log \left(\frac{\ell \cdot m \cdot q_s}{\epsilon^2} \cdot \frac{c_i}{1 - c_i} \right) + 1. \quad (4)$$

Using symmetry, we may assume $c_i \leq 1/2$.

R_∞ -based analysis. The probability preservation property of the Rényi divergence from Lemma 2.7 is multiplicative for $a > 1$ (rather than additive for $a = 1$). Here we use the order $a = \infty$. This property gives that any forging adversary \mathcal{A} having success probability ϵ on the scheme implemented with truncated Gaussian sampling has a success probability $\epsilon' \geq \epsilon / R_\infty(\Phi \|\Phi')$ on the scheme implemented with perfect Gaussian. If $R = R_\infty(\Phi \|\Phi') \leq O(1)$, then $\epsilon' = \Omega(\epsilon)$. By

the multiplicative property of the RD (over the $\ell \cdot m \cdot q_s$ samples needed for signing q_s times), we have $R_\infty(\Phi\|\Phi') \leq R_\infty(B_{\tilde{c}_i}\|B_{c_i})^{\ell \cdot m \cdot q_s}$. By our assumption that $c_i \leq 1/2$, we have $R_\infty(B_{\tilde{c}_i}\|B_{c_i}) = 1 + |\varepsilon_i|/c_i \leq 1 + 2^{-p}$. Therefore, we get $\varepsilon' \geq \varepsilon/(1 + 2^{-p})^{\ell \cdot m \cdot q_s}$. We select parameters to get adversaries with success probabilities $\geq \varepsilon/2$ against the untruncated scheme and set the precision so that $(1 + 2^{-p})^{\ell \cdot m \cdot q_s} \leq 2$. This yields an approximated precision requirement

$$p \geq \log(\ell \cdot m \cdot q_s). \quad (5)$$

Note above estimate may not be accurate unless $\ell \cdot m \cdot q_s \ll 2^p$. Hence we may also require that $p \geq \log(\ell \cdot m \cdot q_s) + C$ for some constant C . This condition essentially eliminates the term ε from the precision needed by the SD-based and KLD-based analyses.⁶ Overall, we get a precomputed table of bit-size $L_{RD} = \log(\ell \cdot m \cdot q_s) \cdot \ell$.

R $_\alpha$ -based analysis. We may also consider R_α -based analysis for general $\alpha > 1$. The reductions here are not tight however. For R_α -based analysis with $\alpha > 1$, the probability preservation we want to use is $\varepsilon' > \varepsilon^{\alpha/(\alpha-1)}/R_\alpha(\Phi\|\Phi')$. The Rényi divergence can be computed by

$$\begin{aligned} (R_\alpha(\Phi\|\Phi'))^{\alpha-1} &= \frac{(1 - c_i - \varepsilon_i)^\alpha}{(1 - c_i)^{\alpha-1}} + \frac{(c_i + \varepsilon_i)^\alpha}{c_i^{\alpha-1}} \\ &= (1 - c_i - \varepsilon_i) \left(1 - \frac{\varepsilon_i}{1 - c_i}\right)^{\alpha-1} + (c_i + \varepsilon_i) \left(1 + \frac{\varepsilon_i}{c_i}\right)^{\alpha-1}. \end{aligned}$$

If $\alpha \ll 2^p$, we can see above

$$\begin{aligned} &\approx (1 - c_i - \varepsilon_i) \left(1 - \frac{(\alpha-1)\varepsilon_i}{1 - c_i}\right) + (c_i + \varepsilon_i) \left(1 + \frac{(\alpha-1)\varepsilon_i}{c_i}\right) \\ &= 1 + \frac{\varepsilon_i^2(\alpha-1)}{c_i(1 - c_i)} \leq 1 + 2^{-2p}(\alpha-1) \frac{c_i}{1 - c_i} \leq 1 + 2^{-2p}(\alpha-1). \end{aligned}$$

For instance if we take $\alpha = 2$, we have $R_2(\Phi\|\Phi') \leq 1 + 2^{-2p}$ and hence $\varepsilon' \geq \varepsilon^2/R_2(\Phi\|\Phi')$. To get a success probability lower bound $\varepsilon^2/2$, it is sufficient to set

$$p \geq \frac{1}{2} \log(\ell \cdot m \cdot q_s). \quad (6)$$

On the other hand, if $\alpha \gg 2^p$, then we have

$$\begin{aligned} (R_\alpha(\Phi\|\Phi'))^{\alpha-1} &= (1 - c_i - \varepsilon_i) \left(1 - \frac{\varepsilon_i}{1 - c_i}\right)^{\alpha-1} + (c_i + \varepsilon_i) \left(1 + \frac{\varepsilon_i}{c_i}\right)^{\alpha-1} \\ &\approx (c_i + \varepsilon_i) \exp\left(\frac{(\alpha-1)\varepsilon_i}{c_i}\right). \end{aligned}$$

⁶ Note that the resulting precision is not independent of ε . The parameters $m = m(\varepsilon)$ and $\ell = \ell(\varepsilon)$ are chosen in [DDLL13] so that any forging adversary has success probability at most ε on the scheme implemented with perfect Gaussian sampling.

Hence the Rényi divergence

$$R_\alpha(\Phi\|\Phi') \approx (c_i + \varepsilon_i)^{1/(\alpha-1)} \exp\left(\frac{\varepsilon_i}{c_i}\right) \approx 1 + \frac{\varepsilon_i}{c_i}.$$

As $\alpha \rightarrow \infty$, $R_\alpha(\Phi\|\Phi') \rightarrow 1 + 2^{-p}$.

Thus if the tightness of the reduction is not a concern, using R_α with small α reduces the precision requirement. Furthermore, we can amplify the success probability of the forger on the truncated Gaussian from ϵ' to some $\epsilon'' > \epsilon'$.

Numerical examples In Table 1, we consider a numerical example which gives the lower bound on the precision for the scheme BLISS-I (with $\varepsilon = 2^{-128}$) when allowing up to $q_s = 2^{64}$ sign queries to the adversary. For the BLISS-I parameters, we use $m = 1024$, $\ell = 29$, $s = \lceil \sqrt{2\pi} \cdot 254 \cdot \sqrt{1/(2 \ln 2)} \rceil = 541$ and $\tau = 13.4/\sqrt{2\pi} \approx 5.4$. The reductions in the table are tight, except for R_2 (and we are a little bit loose for the R_∞).

Table 1. Comparison of the precision to handle adversaries with success probability $\geq \varepsilon$ making $\leq q_s$ sign queries to BLISS.

| | Lower bound on the precision p | Example p | Table bit-size |
|----------------------|---|--------------|----------------|
| SD (Eq. (1)) | $p \geq \log(\ell \cdot m \cdot q_s / \varepsilon)$ | $p \geq 207$ | 6003 |
| SD (Eq. (2)) | $p_i \geq \log(\ell \cdot m \cdot q_s \cdot e^{-\pi 2^i / s^2} / \varepsilon) + 1$ | – | 4598 |
| KLD (Eq. (3)) | $p \geq \log(\ell \cdot m \cdot q_s / \varepsilon^2) / 2 + 1/2$ | $p \geq 168$ | 4872 |
| KLD (Eq. (4)) | $p \geq \log(\ell \cdot m \cdot q_s / \varepsilon^2 \cdot c_i / (1 - c_i)) / 2 + 1$ | – | 3893 |
| R_∞ (Eq. (5)) | $p \geq \log(\ell \cdot m \cdot q_s)$ | $p \geq 79$ | 2291 |
| R_2 (Eq. (6)) | $p \geq \log(\ell \cdot m \cdot q_s) / 2$ | $p \geq 40$ | 1160 |

3.2 GPV Signature scheme

The RD can also be used to reduce the parameters obtained via the SD-based analysis of the GPV signature scheme in [GPV08].

In summary, the signature and the security proof from [GPV08] work as follows. The signature public key is a matrix $A \in \mathbb{Z}_q^{n \times m}$ with n linear in the security parameter λ , $q = \text{poly}(n)$, and $m = O(n \log q)$. The private signing key is a short basis matrix T for the lattice $A^\perp = \{\mathbf{x} \in \mathbb{Z}^m : A \cdot \mathbf{x} = \mathbf{0} \bmod q\}$, whose last successive minimum satisfies $\lambda_m(A^\perp) \leq O(1)$ when $m = \Omega(n \log q)$ (see [GPV08]). A signature $(\boldsymbol{\sigma}, s)$ on a message M is a short vector $\boldsymbol{\sigma} \in \mathbb{Z}^m$ and a random salt $s \in \{0, 1\}^\lambda$, such that $A \cdot \boldsymbol{\sigma} = H(M, s) \bmod q$, where H is a random oracle hashing into \mathbb{Z}_q^n . The short vector $\boldsymbol{\sigma}$ is sampled by computing an arbitrary vector \mathbf{t} satisfying $A \cdot \mathbf{t} = H(M, s) \bmod q$ and using T along with a Gaussian sampling algorithm (see [GPV08, BLP⁺13]) to produce a sample from $\mathbf{t} + D_{A^\perp, r, -\mathbf{t}}$.

The main idea in the security proof from the SIS problem [GPV08] is based on simulating signatures without T , by sampling σ from $D_{\mathbb{Z}^m, r}$ and then programming the random oracle H at (M, s) according to $H(M, s) = A \cdot \sigma \bmod q$. As shown in [GPV08, Le. 5.2], the conditional distribution of σ given $A \cdot \sigma \bmod q$ is exactly the same in the simulation and in the real scheme. Therefore, the SD between the simulated signatures and the real signatures is bounded by the SD between the marginal distribution D_1 of $A \cdot \sigma \bmod q$ for $\sigma \leftarrow D_{\mathbb{Z}^m, r}$ and $U(\mathbb{Z}_q^m)$. This SD for one signature is bounded by ε if $r \geq \eta_\varepsilon(A^\perp)$. This leads, over the q_s sign queries of the attacker, in the SD-based analysis of [GPV08], to take $\varepsilon = O(2^{-\lambda} q_s^{-1})$ and thus $r = \Omega(\sqrt{\lambda + \log q_s})$ (using Lemma 2.2), in order to handle attackers with success probability $2^{-o(\lambda)}$.

Now, by Lemma 2.9, we have that the RD $R_\infty(D_1 \| U)$ is bounded by $1 + c \cdot \varepsilon$ for one signature, for some constant c . By the multiplicativity property of Lemma 2.7, over q_s queries, it is bounded by $(1 + c\varepsilon)^{q_s}$. By taking $\varepsilon = O(q_s^{-1})$, we obtain overall an RD bounded as $O(1)$ between the view of the attacker in the real attack and simulation, leading to a security proof with respect to SIS but with a smaller $r = \Omega(\sqrt{\log \lambda + \log(nq_s)})$. When the number of sign queries q_s allowed to the adversary is much smaller than 2^λ , this leads to significant parameter savings, because SIS's β is reduced and hence n, m, q may be set smaller for the same security parameter λ .

4 Rényi divergence and distinguishing problems

In this section, we prove Theorem 4.1 which allows to use the RD for distinguishing problems and we show how to apply it to the dual-Regev encryption scheme.

4.1 Problems with public sampleability

A general setting one comes across in analyzing the security of cryptographic schemes has the following form. Let P denote a decision problem that asks to distinguish whether a given x was sampled from distribution X_0 or X_1 , defined as follows:

$$X_0 = \{x : r \leftarrow \Phi, x \leftarrow D_0(r)\}, \quad X_1 = \{x : r \leftarrow \Phi, x \leftarrow D_1(r)\}.$$

Here r is some parameter that is sampled from the same distribution Φ in both X_0 and X_1 . The parameter r then determines the conditional distributions $D_0(r)$ and $D_1(r)$ from which x is sampled in X_0 and X_1 , respectively, given r . Now, let P' denote another decision problem that is defined similarly to P , except that in P' the parameter r sampled from a different distribution Φ' (rather than Φ). Given r , the conditional distributions $D_0(r)$ and $D_1(r)$ are the same in P' as in P . Let X'_0 (resp. X'_1) denote the resulting marginal distributions of x in problem P' . Now, in the applications we have in mind, the distributions Φ' and Φ are “close” in some sense, and we wish to show that this implies an

efficient reduction between problems P' and P , in the usual sense that every distinguisher with efficient run-time T and non-negligible advantage ε against P implies a distinguisher for P' with efficient run-time T' and non-negligible advantage ε' . In the classical situation, if the SD $\Delta(\Phi, \Phi')$ between Φ' and Φ is negligible, then the reduction is immediate. Indeed, for $b \in \{0, 1\}$, if p_b (resp. p'_b) denotes the probability that a distinguisher algorithm \mathcal{A} outputs 1 on input distribution X_b (resp. X'_b), then we have, from the SD probability preservation property, that $|p'_b - p_b| \leq \Delta(\Phi, \Phi')$. As a result, the advantage $\varepsilon' = |p'_1 - p'_0|$ of \mathcal{A} against P' is bounded from below by $\varepsilon - 2\Delta(\Phi, \Phi')$ which is non-negligible (here $\varepsilon = |p_1 - p_0|$ is the assumed non-negligible advantage of \mathcal{A} against P).

Unfortunately, for general decision problems P, P' of the above form, it seems difficult to obtain an RD-based analogue of the above SD-based argument, in the weaker setting when the SD $\Delta(\Phi, \Phi')$ is non-negligible, but the RD $R = R(\Phi \| \Phi')$ is small. Indeed, the probability preservation property of the RD in Lemma 2.7 does not seem immediately useful in the case of general decision problems P, P' . With the above notations, it can be used to conclude that $p'_b \geq p_b^2/R$ but this does not allow us to usefully relate the advantages $|p'_1 - p'_0|$ and $|p_1 - p_0|$.

Nevertheless, we now make explicit a special class of “publicly sampleable” problems P, P' for which such a reduction can be made. In such problems, it is possible to efficiently sample from both distributions $D_0(r)$ (resp. $D_1(r)$) given the single sample x from the unknown $D_b(r)$. This technique is implicit in the application of RD in the reductions of [LPR13], and we abstract it and make it explicit in the following.

Theorem 4.1. *Let Φ, Φ' denote two distributions with $\text{Supp}(\Phi) \subseteq \text{Supp}(\Phi')$, and $D_0(r)$ and $D_1(r)$ denote two distributions determined by some parameter $r \in \text{Supp}(\Phi')$. Let P, P' be two decision problems defined as follows:*

- *Problem P : Distinguish whether input x is sampled from distribution X_0 or X_1 , where*

$$X_0 = \{x : r \leftarrow \Phi, x \leftarrow D_0(r)\}, \quad X_1 = \{x : r \leftarrow \Phi, x \leftarrow D_1(r)\}.$$

- *Problem P' : Distinguish whether input x is sampled from distribution X'_0 or X'_1 , where*

$$X'_0 = \{x : r \leftarrow \Phi', x \leftarrow D_0(r)\}, \quad X'_1 = \{x : r \leftarrow \Phi', x \leftarrow D_1(r)\}.$$

Assume that $D_0(\cdot)$ and $D_1(\cdot)$ satisfy the following public sampleability property: there exists a sampling algorithm \mathcal{S} with run-time T_S such that for all (r, b) , given any sample x from $D_b(r)$:

- $\mathcal{S}(0, x)$ *outputs a fresh sample distributed as $D_0(r)$ over the randomness of \mathcal{S} ,*
- $\mathcal{S}(1, x)$ *outputs a fresh sample distributed as $D_1(r)$ over the randomness of \mathcal{S} .*

Then, given a T -time distinguisher \mathcal{A} for problem P with advantage ε , we can construct a distinguisher \mathcal{A}' for problem P' with run-time and distinguishing advantage respectively bounded from above and below by (for any $a \in (1, +\infty]$):

$$O\left(\frac{1}{\varepsilon^2} \log\left(\frac{R_a(\Phi \| \Phi')}{\varepsilon^{a/(a-1)}}\right) \cdot (T_S + T)\right) \quad \text{and} \quad \frac{\varepsilon}{4 \cdot R_a(\Phi \| \Phi')} \cdot \left(\frac{\varepsilon}{2}\right)^{\frac{a}{a-1}}.$$

Proof. Distinguisher \mathcal{A}' is given an input x sampled from $D_b(r)$ for some r sampled from Φ' and some unknown $b \in \{0, 1\}$. For an ε' to be determined later, it runs distinguisher \mathcal{A} on $N = O(\varepsilon^{-2} \log(1/\varepsilon'))$ independent inputs sampled from $D_0(r)$ and $D_1(r)$ calling algorithm S on $(0, x)$ and $(1, x)$ to obtain estimates \hat{p}_0 and \hat{p}_1 for the acceptance probabilities $p_0(r)$ and $p_1(r)$ of \mathcal{A} given as inputs samples from $D_0(r)$ and $D_1(r)$ (with the r fixed to the value used to sample the input x of \mathcal{A}'). By the choice of N and the Hoeffding bound, the estimation errors $|\hat{p}_0 - p_0|$ and $|\hat{p}_1 - p_1|$ are $< \varepsilon/8$ except with probability $< \varepsilon'$ over the randomness of S . Then, if $\hat{p}_1 - \hat{p}_0 > \varepsilon/4$, distinguisher \mathcal{A}' runs \mathcal{A} on input x and returns whatever \mathcal{A} returns, else distinguisher \mathcal{A}' returns a uniformly random bit. This completes the description of distinguisher \mathcal{A}' .

Let \mathcal{S}_1 denote the set of r 's such that $p_1(r) - p_0(r) \geq \varepsilon/2$, \mathcal{S}_2 denote the set of r 's that are not in \mathcal{S}_1 and such that $p_1(r) - p_0(r) \geq 0$, and \mathcal{S}_3 denote all the remaining r 's. Then:

- If $r \in \mathcal{S}_1$, then except with probability $< \varepsilon'$ over the randomness of S , we will have $\hat{p}_1 - \hat{p}_0 > \varepsilon/4$ and thus \mathcal{A}' will output $\mathcal{A}(x)$. Thus, in the case $b = 1$, we have $\Pr[\mathcal{A}'(x) = 1 | r \in \mathcal{S}_1] \geq p_1(r) - \varepsilon'$ and in the case $b = 0$, we have $\Pr[\mathcal{A}'(x) = 1 | r \in \mathcal{S}_1] \leq p_0(r) + \varepsilon'$.
- Assume that $r \in \mathcal{S}_2$. Let $u(r)$ be the probability over the randomness of S that $\hat{p}_1 - \hat{p}_0 > \varepsilon/4$. Then \mathcal{A}' will output $\mathcal{A}(x)$ with probability $u(r)$ and a uniform bit with probability $1 - u(r)$. Thus, in the case $b = 1$, we have $\Pr[\mathcal{A}'(x) = 1 | r \in \mathcal{S}_2] = u(r) \cdot p_1(r) + (1 - u(r))/2$, and in the case $b = 0$, we have $\Pr[\mathcal{A}'(x) = 1 | r \in \mathcal{S}_2] = u(r) \cdot p_0(r) + (1 - u(r))/2$.
- If $r \in \mathcal{S}_3$, except with probability $< \varepsilon'$ over the randomness of S , we have $\hat{p}_1 - \hat{p}_0 < \varepsilon/4$ and \mathcal{A}' will output a uniform bit. Thus, in the case $b = 1$, we have $\Pr[\mathcal{A}'(x) = 1 | r \in \mathcal{S}_3] \geq 1/2 - \varepsilon'$, and in the case $b = 0$, we have $\Pr[\mathcal{A}'(x) = 1 | r \in \mathcal{S}_3] \leq 1/2 + \varepsilon'$.

Overall, the advantage of \mathcal{A}' is bounded from below by:

$$\begin{aligned} \sum_{r \in \mathcal{S}_1} \Phi'(r) (p_1(r) - p_0(r) - 2\varepsilon') + \sum_{r \in \mathcal{S}_2} \Phi'(r) u(r) (p_1(r) - p_0(r)) - \sum_{r \in \mathcal{S}_3} \Phi'(r) 2\varepsilon' \\ \geq \Phi'(\mathcal{S}_1) \cdot \frac{\varepsilon}{2} - 2\varepsilon'. \end{aligned}$$

Without loss of generality, we may assume that the advantage of \mathcal{A} is positive. By an averaging argument, the set \mathcal{S}_1 has probability $\Phi(\mathcal{S}_1) \geq \varepsilon/2$ under distribution Φ . Hence, by the RD probability preservation property (see Lemma 2.7), we have $\Phi'(\mathcal{S}_1) \geq (\varepsilon/2)^{\frac{\alpha}{\alpha-1}} / R_\alpha(\Phi \| \Phi')$. The proof may be completed by setting $\varepsilon' = (\varepsilon/8) \cdot (\varepsilon/2)^{\frac{\alpha}{\alpha-1}} / R_\alpha(\Phi \| \Phi')$. \square

4.2 Application to dual-Regev encryption

As an application to Theorem 4.1, we show that LWE with non-uniform and possibly statistically correlated \mathbf{a}_i 's of the samples (\mathbf{a}_i, b_i) 's (with b_i either independently sampled from $U(\mathbb{T})$ or close to $\langle \mathbf{a}_i, \mathbf{s} \rangle$ for a secret vector \mathbf{s}) remains

at least as hard as standard LWE, as long as the RD $R(U\|\Phi)$ remains small, where Φ is the joint distribution of the given \mathbf{a}_i 's and U denotes the uniform distribution.

Let m, n, q, χ be as in Definition 2.5 and Φ denote a distribution over $\mathbb{Z}_q^{m \times n}$. We define the LWE variant $\text{LWE}_{n,q,\chi,m}(\Phi)$ as follows: Sample $A \leftarrow \Phi$, $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, $\mathbf{e} \leftarrow \chi^m$ and $\mathbf{u} \leftarrow U(\mathbb{T}^m)$; The goal is to distinguish between the distributions $(A, \frac{1}{q}A\mathbf{s} + \mathbf{e})$ and (A, \mathbf{u}) over $\mathbb{Z}_q^{m \times n} \times \mathbb{T}^m$. Note that standard LWE is obtained by taking $\Phi = U(\mathbb{Z}_q^{m \times n})$.

Corollary 4.2. *Let Φ and Φ' be two distributions over $\mathbb{Z}_q^{m \times n}$ with $\text{Supp}(\Phi) \subseteq \text{Supp}(\Phi')$. If there exists a distinguisher \mathcal{A} against $\text{LWE}_{n,q,\chi,m}(\Phi)$ with run-time T and advantage $\varepsilon = o(1)$, then there exists a distinguisher \mathcal{A}' against $\text{LWE}_{n,q,\chi,m}(\Phi')$ with run-time $T' = O(\varepsilon^{-2} \log \frac{R_a(\Phi\|\Phi')}{\varepsilon^{a/(a-1)}} \cdot (T + \text{poly}(m, \log q)))$ and advantage $\Omega\left(\frac{\varepsilon^{1+a/(a-1)}}{R_a(\Phi\|\Phi')}\right)$, for any $a \in (1, +\infty]$.*

Proof. Apply Theorem 4.1 with $r = A \in \mathbb{Z}_q^m$, $x = (A, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{T}^m$, $D_0(r) = (A, A \cdot \mathbf{s} + \mathbf{e})$ with $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$ and $\mathbf{e} \leftarrow \chi^m$, and $D_1(r) = (A, \mathbf{u})$ with $\mathbf{u} \leftarrow U(\mathbb{Z}_q^m)$. The sampling algorithm \mathbf{S} is such that $\mathbf{S}(0, x)$ outputs $(A, A \cdot \mathbf{s}' + \mathbf{e}')$ for $\mathbf{s}' \leftarrow U(\mathbb{Z}_q^n)$ and $\mathbf{e}' \leftarrow \chi^m$, while $\mathbf{S}(1, x)$ outputs (A, \mathbf{u}') with $\mathbf{u}' \leftarrow U(\mathbb{Z}_q^m)$. \square

The latter result allows us to save an $\Omega(\sqrt{\lambda/\log \lambda})$ factor in the Gaussian deviation parameter r used for secret key generation in the dual-Regev encryption scheme [GPV08], where λ refers to the security parameter.

We recall that the dual-Regev encryption scheme has a general public parameter $A \in \mathbb{Z}_q^{m \times n}$, a secret key of the form $sk = \mathbf{e}$ with $\mathbf{x} \leftarrow D_{\mathbb{Z}_q^m, r}$ and a public key of the form $\mathbf{u} = A^t \mathbf{x} \bmod q$. A ciphertext for a message $m \in \{0, 1\}$ is obtained as follows: Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$, $\mathbf{e}_1 \leftarrow \chi^m$ and $e_2 \leftarrow \chi$; return ciphertext $(c_1, c_2) = (\frac{1}{q}A\mathbf{s} + \mathbf{e}_1, \frac{1}{q}(\mathbf{u}, \mathbf{s}) + x_2 + \frac{m}{2}) \in \mathbb{T} \times \mathbb{T}$.

Corollary 4.3. *Suppose that q is prime, $m \geq 2n \log q$ and $r \geq 4\sqrt{\log(12m)/\pi}$. If there exists an adversary against the IND-CPA security of the dual-Regev encryption scheme with run-time T and advantage ε , then there exists a distinguishing algorithm for $\text{LWE}_{n,q,\chi,m+1}$ with run-time $O((\varepsilon')^{-2} \log(\varepsilon')^{-1} \cdot (T + \text{poly}(m)))$ and advantage $\Omega((\varepsilon')^2)$, where $\varepsilon' = \varepsilon - 2q^{-n}$.*

Proof. The IND-CPA security of the dual-Regev encryption scheme as described above is at least as hard as $\text{LWE}_{n,q,\chi,m+1}(\Phi)$ where Φ is obtained by sampling $A \leftarrow U(\mathbb{Z}_q^{m \times n})$, $\mathbf{u} \leftarrow A^t D_{\mathbb{Z}_q^m, r} \bmod q$ and returning the $(m+1) \times n$ matrix obtained by appending \mathbf{u}^t at the bottom of A . We apply Corollary 4.2 with $\Phi' = U(\mathbb{Z}_q^{(m+1) \times n})$.

Since q is prime, if A is full rank, then the multiplication by A^t induces an isomorphism between the quotient group \mathbb{Z}^m/A^\perp and \mathbb{Z}_q^n , where $A^\perp = \{\mathbf{x} \in \mathbb{Z}^m : A^t \cdot \mathbf{x} = \mathbf{0} \bmod q\}$. By Lemma 2.2, we have $\eta_{1/3}(A^\perp) \leq 4\sqrt{\log(12m)/\pi} \leq r$, except for a fraction $\leq q^{-n}$ of the A 's. Let BAD denote the union of such bad A 's and the A 's that are not full rank. We have $\Pr[BAD] \leq 2q^{-n}$.

By the multiplicativity property of Lemma 2.7, we have:

$$R_\infty(\Phi\|\Phi') \leq \max_{A \notin BAD} R_\infty(D_{\mathbb{Z}^m, r} \bmod A^\perp \| U_{\mathbb{Z}^m/A^\perp}).$$

Thanks to Lemma 2.9, we know that the latter is ≤ 2 . The result now follows from Corollary 4.2. \square

In all applications we are aware of, the parameters satisfy $m \leq \text{poly}(\lambda)$ and $q^{-n} \leq 2^{-\lambda}$, where λ refers to the security parameter. The $r = \Omega(\sqrt{\log \lambda})$ bound of our Corollary 4.3, that results from using $\delta = 1/3$ in the condition $r \geq \eta_\delta(A^\perp)$ in the RD-based smoothing argument of the proof above, improves on the corresponding bound $r = \Omega(\sqrt{\lambda})$ that results from the requirement to use $\delta = O(2^{-\lambda})$ in the condition $r \geq \eta_\delta(A^\perp)$ in the SD-based smoothing argument of the proof of [GPV08, Theorem 7.1], in order to handle adversaries with advantage $\varepsilon = 2^{-o(\lambda)}$ in both cases. Thus our RD-based analysis saves a factor $\Omega(\sqrt{\lambda/\log \lambda})$ in the choice of r , and consequently of α^{-1} and q . (The authors of [GPV08] specify a choice of $r = \omega(\sqrt{\log \lambda})$ for their scheme because they use in their analysis the classical “no polynomial attacks” security requirement, corresponding to assuming attacks with advantage $\varepsilon = \lambda^{-O(1)}$, rather than the stronger $\varepsilon = \omega(2^{-\lambda})$ but more realistic setting we take.)

5 Application to LWE with uniform noise

We show the hardness of LWE with a uniform noise by providing a reduction from the $\text{LWE}_{n,q,D_\alpha,m}$ distinguishing problem to the $\text{LWE}_{n,q,U([- \beta, \beta]),m}$ distinguishing problem.

Theorem 5.1. *Let $m \geq n \geq 1$ and with $q \leq \text{poly}(m, n)$ prime. Let $\alpha, \beta > 0$ be real numbers with $\beta = \Omega(m\alpha)$. Then there is a polynomial-time reduction from $\text{LWE}_{n,q,D_\alpha,m}$ to $\text{LWE}_{n,q,\phi,m}$, with $\phi = \frac{1}{q}[qU([- \beta, \beta])]$.*

Proof. In the proof, we let U_β denote the distribution $U([- \beta, \beta])$, to ease notations. Our reduction relies on four steps:

- A reduction from $\text{LWE}_{n,q,D_\alpha,m}$ to $\text{LWE}_{n,q,\psi,m}$ with $\psi = D_\alpha + U_\beta$,
- A reduction from $\text{LWE}_{n,q,\psi,m}$ to $\text{sLWE}_{n,q,\psi,m}$,
- A reduction from $\text{sLWE}_{n,q,\psi,m}$ to $\text{sLWE}_{n,q,U_\beta,m}$,
- A reduction from $\text{sLWE}_{n,q,U_\beta,m}$ to $\text{LWE}_{n,q,U_\beta,m}$.

First step. The reduction is given m elements $(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{T}$, all drawn from $A_{\mathbf{s}, D_\alpha}$ (for some \mathbf{s}), or all drawn from $U(\mathbb{Z}_q^n \times \mathbb{T})$. The reduction consists in adding independent samples from U_β to each b_i . The reduction maps the uniform distribution to itself, and $A_{\mathbf{s}, D_\alpha}$ to $A_{\mathbf{s}, \psi}$.

Second step. Reducing the distinguishing variant of LWE to its search variant is direct.

Third step. The reduction from $\text{sLWE}_{n,q,\psi,m}$ to $\text{sLWE}_{n,q,U_\beta,m}$ is vacuous: by using the RD (and in particular the probability preservation property of Lemma 2.7), we show that an oracle solving $\text{sLWE}_{n,q,U_\beta,m}$ also solves $\text{sLWE}_{n,q,\psi,m}$.

Lemma 5.2. *Let α, β be real numbers with $\alpha \in (0, 1/e)$ and $\beta \geq \alpha$. Let $\psi = D_\alpha + U_\beta$. Then*

$$R_2(U_\beta \|\psi) = \frac{\alpha}{\beta} \int_0^\beta \frac{1}{\int_{-\beta}^\beta e^{-\frac{\pi(x-y)^2}{\alpha^2}} dy} dx \leq 1 + 16 \frac{\alpha}{\beta} \sqrt{\ln(1/\alpha)/\pi}.$$

Proof. The density function of ψ is the convolution of the density functions of D_α and U_β :

$$f_\psi(x) = \frac{1}{2\alpha\beta} \int_{-\beta}^\beta e^{-\frac{\pi(x-y)^2}{\alpha^2}} dy.$$

Using Rényi of order 2, we have:

$$R_2(U_\beta \|\psi) = \int_{-\beta}^\beta \frac{\frac{1}{(2\beta)^2}}{\frac{1}{2\alpha\beta} \int_{-\beta}^\beta e^{-\frac{\pi(x-y)^2}{\alpha^2}} dy} dx = \frac{\alpha}{\beta} \int_0^\beta \frac{1}{\int_{-\beta}^\beta e^{-\frac{\pi(x-y)^2}{\alpha^2}} dy} dx.$$

The denominator in the integrand is a function for $x \in [0, \beta]$.

$$\phi(x) = \alpha - \int_{\beta+x}^\infty \exp\left(\frac{-\pi y^2}{\alpha^2}\right) dy - \int_{\beta-x}^\infty \exp\left(\frac{-\pi y^2}{\alpha^2}\right) dy.$$

For standard Gaussian, we use the following tail bound [CDS03]:

$$\frac{1}{\sqrt{2\pi}} \int_z^\infty e^{-x^2/2} dx \leq \frac{1}{2} e^{-z^2/2}.$$

Then we have

$$\phi(x) \geq \alpha \left(1 - \frac{1}{2} \exp\left(\frac{-\pi(\beta+x)^2}{\alpha^2}\right) - \frac{1}{2} \exp\left(\frac{-\pi(\beta-x)^2}{\alpha^2}\right) \right).$$

Taking the reciprocal of above, we use the first-order Taylor expansion. Note here

$$t(x) = \frac{1}{2} \exp\left(\frac{-\pi(\beta+x)^2}{\alpha^2}\right) + \frac{1}{2} \exp\left(\frac{-\pi(\beta-x)^2}{\alpha^2}\right).$$

We want to bound the function $t(x)$ by a constant $c \in (0, 1)$. Here $t(x)$ is not monotonic. We take the maximum of the first-half and the maximum of the second-half of $t(x)$. An upper-bound ($\beta \geq \alpha$) is:

$$t(x) \leq \frac{1}{2} e^{-\pi\beta^2/\alpha^2} + \frac{1}{2} := \sigma_{\alpha,\beta} + \frac{1}{2} < 1.$$

We can then bound the Rényi of order 2.

$$\begin{aligned}
R_2(U_\beta \|\psi) &= \frac{\alpha}{\beta} \int_0^\beta \frac{1}{\phi(x)} dx \\
&\leq \frac{1}{\beta} \int_0^\beta \frac{1}{1 - \frac{1}{2} \exp\left(\frac{-\pi(\beta+x)^2}{\alpha^2}\right) - \frac{1}{2} \exp\left(\frac{-\pi(\beta-x)^2}{\alpha^2}\right)} dx \\
&\leq \frac{1}{\beta} \int_0^\beta 1 + \frac{1}{1 - 2\sigma_{\alpha,\beta}} \exp\left(\frac{-\pi(\beta+x)^2}{\alpha^2}\right) + \frac{1}{1 - 2\sigma_{\alpha,\beta}} \exp\left(\frac{-\pi(\beta-x)^2}{\alpha^2}\right) dx \\
&= 1 + \frac{1}{(1 - 2\sigma_{\alpha,\beta})\beta} \int_0^{2\beta} \exp\left(\frac{-\pi x^2}{\alpha^2}\right) dx \\
&= 1 + \frac{1}{2(1 - 2\sigma_{\alpha,\beta})\beta} \int_{-2\beta}^{2\beta} \exp\left(\frac{-\pi x^2}{\alpha^2}\right) dx \\
&= 1 + \frac{\alpha}{(1 - 2\sigma_{\alpha,\beta})\beta} (1 - 2D_\alpha(2\beta)) \leq 1 + \frac{1}{1 - 2\sigma_{\alpha,\beta}} \frac{\alpha}{\beta}.
\end{aligned}$$

Hence we have the bound,

$$R_2(U_\beta \|\psi) \leq 1 + \frac{1}{1 - e^{-\pi\beta^2/\alpha^2}} \frac{\alpha}{\beta}.$$

□

We use Lemma 5.2 with m samples and $\beta = \Omega(m\alpha)$ to ensure that the m th power of the RD is ≤ 2 . The RD multiplicativity and probability preservation properties (see Lemma 2.7) imply that $\epsilon' \geq \epsilon^2/R_2^m(U_\beta \|\phi)$; hence if an oracle solves $\text{sLWE}_{n,q,U_\beta,m}$ with probability ϵ , then it also solves $\text{sLWE}_{n,q,\psi,m}$ with probability $\geq \epsilon^2/2$.

Fourth step. We reduce $\text{sLWE}_{n,q,U_\beta,m}$ with continuous noise U_β to $\text{sLWE}_{n,q,\phi,m}$ with discrete noise $\phi = \frac{1}{q} \lfloor qU_\beta \rfloor$ with support contained in \mathbb{T}_q , by rounding to the nearest multiple of $\frac{1}{q}$ any provided b_i (for $i \leq m$). We reduce $\text{sLWE}_{n,q,\phi,m}$ to $\text{LWE}_{n,q,\phi,m}$ by invoking Theorem 2.6. □

6 Open problems

Our results show the utility of the Rényi divergence in several areas of lattice-based cryptography. However, they also suggest some natural open problems, whose resolution could open up further applications. In particular, can we extend the applicability of RD to more general distinguishing problems than those satisfying our ‘public sampleability’ requirement? This may extend our results further. For instance, can we use RD-based arguments to prove the hardness of LWE with uniform noise without using the search to decision reduction of [MM11]? This may allow the proof to apply also to Ring-LWE with uniform noise, and to (Ring) Learning With Rounding [BPR12,AKPW13].

Acknowledgments. We thank Léo Ducas and Vadim Lyubashevsky for useful discussions. This work has been supported in part by ERC Starting Grant ERC-2013-StG-335086-LATTAC, an Australian Research Fellowship (ARF) from the Australian Research Council (ARC), and ARC Discovery Grants DP0987734 and DP110100628.

References

- [Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proc. of STOC*, pages 99–108. ACM, 1996.
- [AKPW13] J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs. Learning with rounding, revisited - new reduction, properties and applications. In *Proc. of CRYPTO (1)*, volume 8042 of *LNCS*, pages 57–74. Springer, 2013.
- [BLP⁺13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *Proc. of STOC*, pages 575–584. ACM, 2013.
- [BPR12] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 719–737. Springer, 2012.
- [BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proc. of FOCS*, pages 97–106. IEEE Computer Society Press, 2011.
- [CDS03] M. Chiani, D. Dardari, and M. K. Simon. New exponential bounds and approximations for the computation of error probability in fading channels. *IEEE Trans. Wireless. Comm.*, 2(4):840–845, July 2003.
- [DDLL13] L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *CRYPTO (1)*, volume 8042 of *LNCS*, pages 40–56. Springer, 2013.
- [DMQ13] N. Döttling and J. Müller-Quade. Lossy codes and a new variant of the learning-with-errors problem. In *Proc. of EUROCRYPT*, LNCS, pages 18–34. Springer, 2013.
- [Duc14] Léo Ducas. Accelerating Bliss: the geometry of ternary polynomials. Cryptology ePrint Archive, Report 2014/874, 2014. <http://eprint.iacr.org/>.
- [EH12] T. van Erven and P. Harremoës. Rényi divergence and Kullback-Leibler divergence. *CoRR*, abs/1206.2459, 2012.
- [GGH13] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *Proc. of EUROCRYPT*, volume 7881 of *LNCS*, pages 1–17. Springer, 2013.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
- [LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013.
- [LPSS14] S. Ling, D. H. Phan, D. Stehlé, and R. Steinfeld. Hardness of k -LWE and applications in traitor tracing. In *Proc. of CRYPTO (1)*, volume 8616 of *LNCS*, pages 315–334. Springer, 2014. Full version available at <http://eprint.iacr.org/2014/494>.

- [LSS14] A. Langlois, D. Stehlé, and R. Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In *Proc. of EUROCRYPT*, LNCS, pages 239–256. Springer, 2014. Full version available at <http://eprint.iacr.org/2014/487>.
- [Lyu12] V. Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 738–755. Springer, 2012.
- [MM11] D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of lwe search-to-decision reductions. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 465–484. Springer, 2011.
- [MR07] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
- [MR09] D. Micciancio and O. Regev. Lattice-based cryptography. In *Post-Quantum Cryptography*, D. J. Bernstein, J. Buchmann, E. Dahmen (Eds), pages 147–191. Springer, 2009.
- [PDG14] T. Pöppelmann, L. Ducas, and T. Güneysu. Enhanced lattice-based signatures on reconfigurable hardware, 2014. To appear in the proceedings of CHES. Full version available at <http://eprint.iacr.org/2014/254>.
- [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93, 2005.
- [Reg09a] O. Regev. Lecture notes of *lattices in computer science*, taught at the Computer Science Tel Aviv University. Available at <http://www.cims.nyu.edu/~regev>, 2009.
- [Reg09b] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [Rén61] A. Rényi. On measures of entropy and information. In *Proc. of the Fourth Berkeley Symposium on Math. Statistics and Probability*, volume 1, pages 547–561, 1961.