Rounded Gaussians are Discrete Gaussians

Corentin Jeudy
corentin.jeudy@orange.com

Orange Research, Applied Crypto Group, Cesson-Sévigné, France

Abstract. Many lattice-based cryptographic schemes still require sampling integers following a Gaussian-like distribution. Depending on the application, this sampling should be generic to tackle variable (and possibly large) standard deviations and centers, with a low memory footprint, easily implementable in constant-time, and with the highest throughput. Combining all these features is very challenging, as shown by the numerous works on optimizing discrete Gaussian samplers. In contrast, the rounded Gaussian (and associated samplers), formalized by Hülsing, Lange and Smeets (PKC'18), seems to satisfy these criteria, but its formal definition makes security analyses of the schemes using it more complex, thus explaining why they are rarely used in practice.
In this paper, we show that a small correction on the parameter σ of the rounded Gaussian yields a distribution that is close to a discrete Gaussian of parameter s. More formally, defining $\sigma = s - 1/(24s)$ yields a Rényi divergence between the two distributions evolving as $1 + O(1/s^8)$. As a result, for mild to large parameters s, one can use the corrected rounded Gaussian sampler, while still analyzing security with an ideal discrete Gaussian at a very mild loss. It then readily gives a discrete Gaussian sampler that benefits from all the features of rounded Gaussian samplers.

Keywords: Lattice-Based Cryptography · Rounded Gaussians · Discrete Gaussians · Sampling

1 Introduction

Sampling integers following a Gaussian distribution is one of the key ingredients of lattice cryptography, either in the hardness analysis of lattice problems [Reg05,GFH26], or directly in actual cryptographic schemes. Gaussian distributions feature interesting convolution properties which make them easy to manipulate, concentration properties which are particularly relevant when sampling short lattice vectors, and naturally arise in many contexts. Yet, the recently published post-quantum cryptography standards ML-KEM [NIS24b] and ML-DSA [NIS24a], which are urged to be used in the next generation of security systems across all sectors to migrate to quantum-safe infrastructures, seem to avoid Gaussians altogether. Even in the case of the Falcon signature scheme [PFH+20] selected for standardization, its reliance of discrete Gaussian sampling seems to be the main source of debate on whether it should be used as

a multi-purpose algorithm, despite its impressive compactness. Practical deployment considerations are then of primary interest, for these standards but also for other cryptographic designs serving use-cases that may require more demanding security assurances than the first standards can provide (e.g., advanced encryption like IBE [DLP14,GPV08] or ABE [AKY24], privacy-enhanced primitives like zero-knowledge proofs [LNP22], anonymous credentials [BLNS23,JS26b,JS26a], group signatures [LNPS21,DFJS26,Jeu26], etc.). Among them, many still have to rely on Gaussian distributions, which shifts the question from whether to use Gaussians to how we can use them efficiently and securely. This integer Gaussian sampling problem has then been studied in a number of ways to improve the practicality and security of integer Gaussian samplers, from which we can extract two main distributions: the discrete Gaussian and the rounded Gaussian.

Discrete Gaussians. The most widely used of the two is the discrete Gaussian distribution, which is defined by a probability mass function that is, by definition, proportional to the Gaussian function $\rho_s(\mathbf{x}) = \exp(-\|\mathbf{x}\|_2^2/2s^2)$. For not too small standard deviations s, the discrete Gaussian distribution behaves almost exactly as its continuous counterparts, benefitting from similar convolution properties, tail bounds, and so on. It is particularly useful to easily analyze the security of cryptographic schemes, which is why discrete Gaussians have been favored over rounded ones. One of their main applications lies in digital signatures, ever since [GPV08] formalized the security of trapdoor Gaussian sampling over arbitrary lattices. Their framework sparked several others, e.g., [Pei10,MP12] (and variants/optimizations thereof), but all eventually rely on sampling discrete Gaussian integers.

Despite reducing the problem of high-dimensional preimage sampling to sampling a discrete Gaussian over $\mathbb{Z}$, the latter still poses a handful of practical issues. One needs to sample such integers with possibly variable parameters s and centers c, with possibly large s (depending on the application), with a small memory footprint, in constant-time, and with the highest possible speed as many such integers generally need to be sampled. This gives the four most important criteria of integer Gaussian samplers: genericity, storage, resistance against timing attacks, and speed. Meeting all these requirements at once is a rather complicated task, which is why there currently is a plethora of methods to sample discrete Gaussian integers. We can categorize them in three main families: rejection samplers, inversion samplers, and hybrid samplers.

Rejection Samplers. First formalized in [GPV08], rejection-based methods sample a uniform integer x in the domain of the tailcutted discrete Gaussian and accept it with probability $\rho_s(x)$. This has the advantage of being fully generic, avoiding any stored precomputations, but by nature requires several iterations each computing the transcendental function ρ_s. Improvements of the latter idea, such as the discrete Ziggurat sampler [BCG+13] or Karney's sampler [Kar16], suffer the same limitations albeit with different trade-offs between speed, memory, center genericity, etc.

Inversion Samplers. On the opposite direction, inversion samplers aim at removing rejections and improving implementation protection by relying on cumula-

tive distribution tables (CDT) and simple look-up sampling. First proposed by Peikert [Pei10], inversion Gaussian sampling based on CDT consists in precomputing a table T with $\mathsf{T}[z] = \mathbb{P}_x[x \leq z]$ for x following a discrete Gaussian of parameter s. Sampling then simply consists in drawing a uniform $u \in [0, 1)$ and searching for $z \in \mathbb{Z}$ such that $u \in [\mathsf{T}[z-1], \mathsf{T}[z])$. One obviously has to truncate the table to entries in the tailcutted domain $[-\tau s, \tau s]$, but it gives high-speed sampling for small s. This method's limitations arise for larger values of s. The table indeed becomes very cumbersome to store when s or the desired precision grows, thus hindering memory. In addition to not being able to efficiently tackle large deviations s, the table is set for a single pair (s, c) which affects genericity in applications where s or c might change between samples. Several optimizations exist to handle more generic centers [AMAR17], reduce the table precision and size [AMR18,KAA22], speed [DN12,DB15]. Other inversion samplers based on tree-traversal methods like Knuth-Yao [DG14,RVV13] (and improvements thereof) show remarkable throughput, but suffer similar genericity and storage issues. Alternatively, the sampler of [MW17] relies on inversion samplers but combines them using convolution. It designs a generic, rejection-free sampler by convoluting many small samples, but unfortunately only tackles arbitrary centers by using several base samplers for different cosets. These are instantiated with inversion samplers [Pei10,DG14] which requires storing sampling material for these different coests.

Hybrid Samplers. Finally, hybrid samplers can be seen as a combination of the latter too, combining an inversion sampler for a small fixed parameter s_0, uniform sampling, and rejection sampling to produce correctly distributed integers. This is the approach of the Bernoulli sampler [DDLL13] and later optimizations, or the SamplerZ algorithm [HPRR20] used in Falcon. They provide interesting trade-offs between speed, memory, security and genericity. The original sampler from [DDLL13] typically scales very well to large deviations, with a very small CDT, and only a small constant rejection rate. However, the native design lacks the genericity to handle arbitrary centers. The SamplerZ algorithm [HPRR20] builds upon the Bernoulli sampler to accomodate for variable centers and simplifying the constant-time protection of the algorithm. It leverages a CDT for a small parameter $s_{\max}$ and a polynomial approximation of the exponential, but can only sample a discrete Gaussian with variable parameter $s \leq s_{\max}$. The genericity on s (which is necessary in Falcon as the parameter s depend on the secret key and changes between samples) is therefore limited to relatively small parameters, which may be unsuited for other applications.

Rounded Gaussians. Overall, despite numerous and impressive improvements in discrete Gaussian integer samplers, all come with a certain set of limitations which have to be assessed when deciding which sampler to use. The other main distribution is the rounded Gaussian distribution [HLS18]. As opposed to the discrete Gaussian, the latter is obtained by sampling a continuous Gaussian of parameter σ and center c[1], and rounding it to the nearest integer. Its use in lattice cryptography was motivated in [HLS18] by it satisfying the main criteria

[1] The latter actually reduces to sampling it for $\sigma = 1$ and $c = 0$.

we discussed above. The authors indeed provide a simple constant-time sampler, based on the Box-Muller transform [BM58] for sampling continuous Gaussians, that is rejection-free, memory-less, and for generic parameters σ and centers c. Despite its reliance on transcendental functions (cos, sin, sqrt), a deeper analysis of its alleged constant-time behavior was carried in [BOKR20], which proposed a fully detailed constant-time portable implementation for Box-Muller and the rounded Gaussian sampler. They show the sampler is competitive with the Micciancio-Walter sampler [MW17][2].

Despite seemingly checking all the boxes, rounded Gaussians are rarely used in lattice cryptography, which is actually explained by their very definition. Although it keeps a Gaussian-like shape, its probability mass function is defined through the integral of ρ_σ over the domain $\mathbf{x} + [-1/2, 1/2)^n$. This non-analytic description naturally makes it much more complex to handle in cryptographic security analyses, as can be seen in the example conducted in [HLS18] on BLISS [DDLL13]. Using rounded Gaussians in actual schemes then requires assessing their security on a case-by-case basis, which is far from trivial. To get the relative best of both worlds, another approach, explored for example in [ZSS20], is instead to use the rounded Gaussian to construct a more efficient uncentered discrete Gaussian sampler, by using additional rejection sampling. It benefits from some of the nice features of rounded Gaussians, but also inherit some of the drawbacks of discrete Gaussian samplers (rejection, and evaluation of exp in addition to the continuous Gaussian sampler).

1.1 Our Contributions

Instead of trying to devise a new discrete Gaussian sampler that would simultaneously be fast, rejection-free, memory-less, generic, handling large deviations, and with a constant-time implementation, we start from the rounded Gaussian distribution, which already satisfies these properties, and wonder the following: can rounded Gaussians be used in place of discrete Gaussians without drastically changing security analyses? We answer positively, albeit with a mild disclaimer. We indeed show that a very small correction on the rounded Gaussian parameter σ yields a distribution that is extremely close to a discrete Gaussian. How close however depends on the size of the target discrete Gaussian parameter s, and is likely insufficient for applications needing many samples of very small deviations.

More formally, it is well-known, and stated in [HLS18], that rounding the continuous Gaussian distribution of parameter σ into $\mathbb{Z}$ inflates the variance to roughly $\sigma'^2 = \sigma^2 + 1/12$. Aiming for the discrete Gaussian of parameter s would then require $\sigma' \approx s$ and therefore $\sigma \approx \sqrt{s^2 - 1/12}$. In the context of fast, generic, and secure discrete Gaussian sampling, computing $\sqrt{s^2 - 1/12}$ is to be avoided because of the square root. Hence, for dynamic samplers where s may change from sample to sample, this simple variance matching is unfortunately out of scope. Nevertheless, taking the expansion of the square root at the first order,

[2] Actually, they compare their sampler to the total time of the calls to the Knuth-Yao base samplers, which does not account for convolution steps.

yielding $\sigma \approx s - 1/(24s) + O(1/s^3)$, may do the trick. Our idea is therefore to truncate this expansion to define the rounded Gaussian parameter, i.e., $\sigma(s) = s - 1/(24s)$, in the hope that it gives a distribution close to the discrete Gaussian of parameter s. We show that the Rényi divergence between the two distributions turns out to be rapidly decreasing as s grows, leading to our main theorem which we state here informally. The formal statement is provided in Theorem 3.1.

Theorem 1.1 (Informal). *The Rényi divergence of order a from the rounded Gaussian over $\mathbb{Z}^n$ of parameter $\sigma(s)$ to the discrete Gaussian over $\mathbb{Z}^n$ of parameter s is $1 + \frac{73an}{33177600} \cdot \frac{1}{s^8} + o(s^{-9})$.*

This rapidly decreasing polynomial dependency in s shows that, for medium to large values of s, the distributions can be indistinguishable in (search) cryptographic applications. Rounded Gaussians can then be used in place of discrete Gaussians, without much affecting the existing analysis based on the latter. One would first replace, in the security proof, the real rounded Gaussian by an ideal discrete Gaussian distribution using Theorem 1.1, and pursue the proof as before. The disclaimer hinted above about our result however lies in the edge regime where s is very small, e.g., s close to 1. In that case, our result still yields a constant Rényi divergence that is not much larger than 1, but that is not close enough to 1 for typical security arguments. In particular, for applications requiring a lot of sampling for such small parameters, e.g., some lattice signature schemes like Falcon [PFH+20], the Rényi loss would impact the security analysis too much. Using rounded Gaussians would then require a dedicated analysis where all the discrete Gaussians would be replaced by rounded ones, in the spirit of what was done for the BLISS signature scheme [DDLL13] in [HLS18]. Nevertheless, our result sharpens our theoretical understanding of rounded Gaussians and how they can be used safely in lattice cryptography.

Beyond the theoretical reach, we can interpret Theorem 1.1 as giving a direct discrete Gaussian sampler for $\mathbb{Z}$ (or $\mathbb{Z}^n$) for not too small parameters s, which builds upon a corrected rounded Gaussian sampler. Sampling a rounded Gaussian of parameter $\sigma(s) = s - 1/(24s)$ would indeed produce a distribution close to the discrete Gaussian of parameter s, with all the advantages of rounded Gaussian samplers. It is then generic enough to easily handle on-the-fly parameters s and centers c, does not deteriorate with large s beyond the natural floating-point precision requirements, and it is memory-less, rejection-free, and can easily be implemented. Additionally, building upon [HLS18,BOKR20], a constant-time implementation would simply require computing $\sigma(s)$ online and in constant-time, which we detail for completeness.

2 Preliminaries

For integers $m \leq n$, we define $[m, n] = \{m, \ldots, n\}$, and $[n] = [1, n]$. We let $\mathbb{N}, \mathbb{Z}, \mathbb{R}, \mathbb{C}$ respectively be the sets of natural integers, integers, reals, and complex numbers, and define $\mathbb{N}^\times = \mathbb{N} \setminus \{0\}$. We also denote by $\mathcal{P}_\pm$ the continuous interval $[-1/2, 1/2)$. We use bold font and lower case for vectors, e.g., $\mathbf{a}$. Finally, $\|\mathbf{a}\|_2 = (\sum_i a_i^2)^{1/2}$ denotes the standard Euclidean norm of $\mathbf{a} = [a_i]_{i \in [n]} \in \mathbb{C}^n$.

2.1 Probabilities

For a set S of positive finite measure, we define $U(S)$ to be the uniform probability distribution over S. For a family of probability distributions $(P_i)_{i\in[n]}$, we note $\otimes_{i\in[n]}P_i$ their joint distribution. In this work, we use the Rényi divergence [Rén61] to measure the closeness of two probability distributions. The definition is recalled below, as well as its multiplicativity and probability preservation properties [vEH14].

Definition 2.1. *Let P, Q be two discrete probability distributions over a countable set S such that $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q) \subseteq S$. For any $a \in (1, +\infty]$, we define the Rényi divergence of order a from P to Q by*

$$\mathrm{RD}_a(P\|Q) = \left(\sum_{\mathbf{x}\in\mathrm{Supp}(P)} \frac{P(\mathbf{x})^a}{Q(\mathbf{x})^{a-1}}\right)^{\frac{1}{a-1}}.$$

Lemma 2.1 ([vEH14, Thm. 28]). *Let $n \in \mathbb{N}^\times$, $a \in (1, +\infty]$, and $(P_i)_{i\in[n]}$, $(Q_i)_{i\in[n]}$ be families of probability distributions over countable sets $(S_i)_{i\in[n]}$, such that $\mathrm{Supp}(P_i) \subseteq \mathrm{Supp}(Q_i) \subseteq S_i$. It then holds that*

$$\mathrm{RD}_a\left(\otimes_{i\in[n]}P_i \| \otimes_{i\in[n]} Q_i\right) = \prod_{i\in[n]} \mathrm{RD}_a(P_i\|Q_i).$$

Additionally, for P, Q two probability distributions over a countable set such that $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$, and any event $E \subseteq \mathrm{Supp}(P)$, it holds that

$$P(E) \leq Q(E)^{\frac{a-1}{a}} \cdot \mathrm{RD}_a(P\|Q)^{\frac{a-1}{a}}.$$

Discrete Gaussian Distributions. Let $n \in \mathbb{N}^\times$. We define the discrete Gaussian distribution over $\mathbb{Z}^n$ with parameter $s > 0$ and center $\mathbf{c} \in \mathbb{R}^n$ by its probability mass function $\mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}} = \rho_{s,\mathbf{c}}\mathbf{1}_{\mathbb{Z}^n}/\rho_{s,\mathbf{c}}(\mathbb{Z}^n)$, where

$$\forall \mathbf{x} \in \mathbb{R}^n, \rho_{s,\mathbf{c}}(\mathbf{x}) = \exp\left(-\frac{1}{2s^2}\|\mathbf{x}-\mathbf{c}\|_2^2\right),$$

and $\mathbf{1}_A$ is the indicator function of A. When $\mathbf{c} = \mathbf{0}$, we omit the subscript. Notice that $\rho_{s,\mathbf{c}}(\mathbf{x}) = \rho_s(\mathbf{x}-\mathbf{c})$. We also define the *continuous* multivariate Gaussian density $D_{s,\mathbf{c}} = \rho_{s,\mathbf{c}}/(\sqrt{2\pi}s)^n$.

Rounded Gaussian Distributions. The rounded Gaussian distribution formalized in [HLS18] can be defined compactly by rounding the continuous Gaussian, i.e., by $\lfloor D_{s,\mathbf{c}} \rceil$, where rounding is performed coefficientwise. It corresponds to averaging the continuous distribution $D_{s,\mathbf{c}}$ over the volume-1 hypercube $\mathcal{P}_\pm^n$ centered at each point of $\mathbb{Z}^n$. More formally, the rounded Gaussian distribution over $\mathbb{Z}^n$ of parameter $\sigma > 0$ and center $\mathbf{c} \in \mathbb{R}^n$ is defined by the probability mass function $\mathcal{R}_{\mathbb{Z}^n,\sigma,\mathbf{c}}$ as follows

$$\forall \mathbf{x} \in \mathbb{Z}^n, \mathcal{R}_{\mathbb{Z}^n,s,\mathbf{c}}(\mathbf{x}) = \int_{\mathbf{x}+\mathcal{P}_\pm^n} D_{\sigma,\mathbf{c}}(\mathbf{y})d\mathbf{y}$$

$$= \int_{\mathbf{x}+\mathcal{P}_{\pm}^n} \frac{1}{(\sqrt{2\pi}\sigma)^n} \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y}-\mathbf{c}\|_2^2\right) d\mathbf{y}.$$

We also omit the subscript when $\mathbf{c} = \mathbf{0}$.

2.2 Calculus Results

In our analysis of the divergence between these distributions, we often manipulate functions with Gaussian decays. In particular, we need a few technical lemmas on such functions. First, the following straightforward lemma bounds the mass of $\mathbb{Z} \setminus \{0\}$ for a function with Gaussian decay.

Lemma 2.2. *Let $f : \mathbb{R} \to \mathbb{R}$ be a function such that there exists $\alpha, \beta > 0$ so that for all $k \in \mathbb{Z} \setminus \{0\}$, $|f(k)| \leq \alpha e^{-\beta k^2}$. Then, $f(\mathbb{Z} \setminus \{0\})$ exists and satisfies $f(\mathbb{Z} \setminus \{0\}) \leq 2\alpha e^{-\beta}/(1-e^{-\beta})$.*

Proof. Because of the bound on f for integer inputs, it holds that the series $\sum_{k\in\mathbb{N}^\times} f(k)$ and $\sum_{k\in\mathbb{N}^\times} f(-k)$ are absolutely convergent and thus convergent, giving the existence of $f(\mathbb{Z} \setminus \{0\})$. In particular, we have $f(\mathbb{Z} \setminus \{0\}) \leq 2\alpha \sum_{k\in\mathbb{N}^\times} e^{-\beta k^2}$. Using that for all $k \in \mathbb{N}^\times$, $k^2 \geq k$, it yields

$$\sum_{k\in\mathbb{N}^\times} e^{-\beta k^2} \leq \sum_{k\in\mathbb{N}^\times} e^{-\beta k} = e^{-\beta} \sum_{k\in\mathbb{N}} e^{-\beta k} = \frac{e^{-\beta}}{1-e^{-\beta}},$$

as $e^{-\beta} \in (0,1)$, which yields the claim. □

As we target an exact taylor expansion of the divergence, a mere bound on the mass $f(\mathbb{Z} \setminus \{0\})$ is not sufficient. It is then helpful to express the mass of $\mathbb{Z}$ (or a shift $\mathbb{Z} - c$) of a function h with Gaussian decay in terms of its integral over $\mathbb{R}$. The integral later facilitates the exact computation of the expansion. For that we use the following technical lemma, which we prove below. The idea is to use the Poisson summation formula, and use the Gaussian decay condition to bound the Fourier mass of $\mathbb{Z} \setminus \{0\}$. In this paper, we denote by $f^{(m)}$ the m-th derivative of a function $f \in \mathcal{D}^m(\mathbb{R})$. We also recall that a complex-valued function f is called *entire* if it is holomorphic over the whole complex field $\mathbb{C}$, i.e., it is complex differentiable over $\mathbb{C}$.

Lemma 2.3. *Let $h : \mathbb{R} \to \mathbb{C}$ be a function that extends to an entire function over $\mathbb{C}$. We assume there exists $K, A, B > 0$ such that*

$$\forall x, y \in \mathbb{R}, |h(x+iy)| \leq K \exp\left(-\frac{A}{2}x^2 + \frac{B}{2}y^2\right). \qquad (*)$$

Then, we have for any $c \in \mathbb{R}$,

$$h(\mathbb{Z}-c) = \int_{\mathbb{R}} h(x)dx + R_h, \quad \text{with } |R_h| \leq 2K\sqrt{\frac{2\pi}{A}} \cdot \frac{e^{-2\pi^2/B}}{1-e^{-2\pi^2/B}}.$$

Proof. Because the extension of h is entire, it is holomorphic over $\mathbb{C}$. As being complex-differentiable over $\mathbb{C}$ implies being infinitely differentiable over $\mathbb{C}$, we get that its restriction to $\mathbb{R}$, i.e., h, is therefore in $\mathcal{C}^\infty(\mathbb{R})$. Using Equation $(*)$ for $y = 0$, it holds that $\forall x \in \mathbb{R}$, $|h(x)| \leq Ke^{-Ax^2/2}$, which implies $h \in L^1(\mathbb{R})$, proving h admits a Fourier transform and that the integral $\int_{\mathbb{R}} h(x)dx$ from the statement exists. We now prove h is a Schwartz function, i.e.,

$$\forall \alpha, \beta \in \mathbb{N},\ \sup_{x \in \mathbb{R}} \left| x^\alpha h^{(\beta)}(x) \right| < \infty,$$

to be able to apply the Poisson summation formula. To bound $\left|h^{(\beta)}(x)\right|$, we use Cauchy's estimate on a circle of radius 1. Let $x \in \mathbb{R}$ and define the complex circle around x by $C = \{z \in \mathbb{C} : |z - x| = 1\}$. For all $z = x' + iy' \in C$, we have $|x' - x|^2 + y'^2 = 1$. As such, it holds that $y'^2 \leq 1$ and $|x'| \geq |x| - 1$. The last inequality stems from the reverse triangle inequality as $|x| - |x'| \leq ||x| - |x'|| \leq |x - x'| \leq 1$. Using Equation $(*)$, we have

$$\begin{aligned} |h(x' + iy')| &\leq Ke^{-Ax'^2/2 + By'^2/2} \\ &\leq Ke^{B/2} e^{-A(|x|-1)^2/2} \\ &\leq Ke^{(B-A)/2} e^{A|x| - Ax^2/2}. \end{aligned}$$

Then, by Cauchy's estimate, it holds that for all $\beta \in \mathbb{N}^\times$

$$\left|h^{(\beta)}(x)\right| \leq \frac{\beta!}{1^\beta} \sup_{z \in C} |h(z)| \leq \beta! \cdot Ke^{(B-A)/2} e^{A|x| - Ax^2/2}.$$

Multiplying by x^α, and using the continuity of the right-hand side and the fact that it tends to 0 in $\pm\infty$, we directly get that $\forall \alpha, \beta \in \mathbb{N}, \sup_{x \in \mathbb{R}} \left|x^\alpha h^{(\beta)}(x)\right| < \infty$. As a result, h is a Swchartz function. Let $c \in \mathbb{R}$. Then, the constant translation $g : x \mapsto h(x - c)$ is also a Schwartz function. The Poisson summation formula therefore gives

$$h(\mathbb{Z} - c) = g(\mathbb{Z}) = \widehat{g}(\mathbb{Z}) = \widehat{g}(0) + \widehat{g}(\mathbb{Z} \setminus \{0\}).$$

Recalling that $\widehat{g}(k) = e^{-2i\pi kc}\widehat{h}(k)$, we first have $\widehat{g}(0) = \widehat{h}(0) = \int_{\mathbb{R}} h(x)dx$. It remains to bound the Fourier mass of $\mathbb{Z} \setminus \{0\}$, but by the phase-shift identity between $\widehat{g}$ and $\widehat{h}$, we can actually study the $\widehat{h}(k)$. Now, let $k > 0$ and $y_0 > 0$. Because $z \in \mathbb{C} \mapsto h(z)e^{-2i\pi kz}$ is holomorphic with Gaussian decay, we can use Cauchy's integral theorem to argue the following complex change of variable

$$\begin{aligned} \widehat{h}(k) = \int_{\mathbb{R}} h(x)e^{-2i\pi kx}dx &= \int_{\mathbb{R}} h(x - iy_0)e^{-2i\pi k(x - iy_0)}dx \\ &= e^{-2\pi ky_0} \int_{\mathbb{R}} h(x - iy_0)e^{-2i\pi kx}dx \end{aligned}$$

Taking the magnitude and using Equation $(*)$, we get

$$\left|\widehat{h}(k)\right| \leq e^{-2\pi ky_0} \int_{\mathbb{R}} Ke^{-\frac{A}{2}x^2 + \frac{B}{2}y_0^2}dx = e^{-2\pi ky_0 + \frac{B}{2}y_0^2} \cdot K \cdot \sqrt{2\pi/A},$$

We can then optimize over $y_0 > 0$, where the right-hand side is minimized at $y_0 = 2\pi k/B > 0$. It then gives $\left|\widehat{h}(k)\right| \leq K\sqrt{2\pi/A}e^{-2\pi^2k^2/B}$. The exact same reasoning for $k < 0$ with $y_0 < 0$ gives the exact same bound. This then gives $|\widehat{g}(k)| = \left|\widehat{h}(k)\right| \leq K\sqrt{2\pi/A}e^{-2\pi^2k^2/B}$ for all $k \neq 0$. Then, Lemma 2.2 with $\alpha = K\sqrt{2\pi/A} > 0$ and $\beta = 2\pi^2/B > 0$ gives

$$|\widehat{g}(\mathbb{Z} \setminus \{0\})| \leq 2K\sqrt{\frac{2\pi}{A}} \cdot \frac{e^{-2\pi^2/B}}{1 - e^{-2\pi^2/B}},$$

Combining it all with $R_h = \widehat{g}(\mathbb{Z} \setminus \{0\})$ gives the claim. □

3 First-Order Sheppard's Correction for Rounded Gaussians

When rounded into $\mathbb{Z}$, the continuous Gaussian distribution inflates the variance by the term $\mathbb{V}[U(\mathcal{P}_\pm)] = 1/12$ known as Sheppard's correction [She97], i.e., $\mathbb{V}[\mathcal{R}_{\mathbb{Z},\sigma}] \approx \sigma^2 + 1/12$. Forcing this variance to be some s^2, to target the discrete Gaussian $\mathcal{D}_{\mathbb{Z},s}$, entails $\sigma = \sqrt{s^2 - 1/12}$. Although variance matching is insufficient to argue distribution closeness, it intuitively seems that $\mathcal{R}_{\mathbb{Z}^n,\sqrt{s^2-1/12}}$ and $\mathcal{D}_{\mathbb{Z}^n,s}$ will turn out to be close. We indeed empirically observe they are, with a quickly decreasing Rényi divergence. However, in applications where s may vary from one sample to the next, using such corrected rounded Gaussians would require computing many square roots, which would both slow down the sampler and make it harder to implement securely [KHTA26]. To avoid that, we can instead expand σ into

$$\sigma = \sqrt{s^2 - 1/12} = s - \frac{1}{24s} - \frac{1}{1152s^3} + O\left(\frac{1}{s^5}\right).$$

Truncating at the first order in $1/s$ avoids the square root and already gives an extremely close distribution. The goal of this work is then to determine how close they are. We show that $\mathcal{R}_{\mathbb{Z}^n,s-1/(24s),\mathbf{c}}$ is very close to $\mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}$, with a Rényi divergence in $1 + \gamma an/s^8 + o(s^{-9})$, where a is the Rényi divergence order and γ is an absolute constant we determine. We formalize it in the following theorem.

Theorem 3.1. *Let $n \in \mathbb{N}^\times$ and $a \in \mathbb{N} \setminus \{0, 1\}$ be fixed. For any $s > 1/\sqrt{24}$, we define $\sigma(s) = s - 1/(24s) > 0$. It then holds that for any $\mathbf{c} \in \mathbb{R}^n$,*

$$\mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z}^n,\sigma(s),\mathbf{c}} \| \mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}\right) \underset{s\to+\infty}{=} 1 + \frac{73an}{33177600} \cdot \frac{1}{s^8} + o\left(\frac{1}{s^9}\right).$$

Proof. As the proof is quite long, we split into several subsections for clarity of presentation. We later refer to the constant $\gamma_a = 73a(a-1)/33177600$, so that $n\gamma_a/(a-1)$ corresponds to the factor of $1/s^8$ from the statement above.

3.1 Dimension Reduction

The first step is to reduce the problem to the case $n = 1$. As $\mathcal{R}_{\mathbb{Z}^n,\sigma(s),\mathbf{c}}$ and $\mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}$ are both product distributions, we can resort to the multiplicativity of the Rényi divergence to focus on the one-dimension case.

Lemma 3.1. *Let $n \in \mathbb{N}^\times$, $a > 1$, $\sigma, s > 0$, and $\mathbf{c} = [c_i]_{i\in[n]} \in \mathbb{R}^n$. It holds that* $\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z}^n,\sigma,\mathbf{c}} \| \mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}) = \prod_{i\in[n]} \mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma,c_i} \| \mathcal{D}_{\mathbb{Z},s,c_i})^n$.

Proof (Lemma 3.1). We first show that both distributions are tensor distributions. We indeed have that for any $r > 0$, $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{x} \in \mathbb{Z}^n$, $\rho_{r,\mathbf{a}}(\mathbf{x}) = \prod_{i\in[n]} \rho_r(x_i - a_i)$, where x_i, a_i are the i-th entry of $\mathbf{x}$ and $\mathbf{a}$ respectively. We also have

$$\begin{aligned}
\rho_{s,\mathbf{c}}(\mathbb{Z}^n) &= \sum_{\mathbf{y}\in\mathbb{Z}^n} \prod_{i\in[n]} \rho_s(y_i - c_i) \\
&= \sum_{y_1\in\mathbb{Z}} \rho_s(y_1 - c_1) \left(\sum_{y_2\in\mathbb{Z}} \rho_s(y_2 - c_2) \left(\dots \sum_{y_n\in\mathbb{Z}} \rho_s(y_n - c_n) \dots \right) \right) \\
&= \prod_{i\in[n]} \left(\sum_{y_i\in\mathbb{Z}} \rho_s(y_i - c_i) \right) \\
&= \prod_{i\in[n]} \rho_{s,c_i}(\mathbb{Z}).
\end{aligned}$$

We then get that $\mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}(\mathbf{x}) = \prod_{i\in[n]} \rho_{s,c_i}(x_i)/\rho_{s,c_i}(\mathbb{Z}) = \prod_{i\in[n]} \mathcal{D}_{\mathbb{Z},s,c_i}(x_i)$ for any $\mathbf{x} \in \mathbb{Z}^n$. Hence, $\mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}} = \otimes_{i\in[n]} \mathcal{D}_{\mathbb{Z},s,c_i}$. We can similarly factor the integrals appearing in the rounded Gaussian as

$$\begin{aligned}
&\int_{\mathbf{x}+\mathcal{P}_\pm^n} \rho_{\sigma,\mathbf{c}}(\mathbf{y}) d\mathbf{y} \\
&= \int_{x_1+\mathcal{P}_\pm} \rho_\sigma(y_1 - c_1) \left(\int_{x_2+\mathcal{P}_\pm} \rho_\sigma(y_2 - c_2) \left(\dots \int_{x_n+\mathcal{P}_\pm} \rho_\sigma(y_n - c_n) dy_n \dots \right) dy_2 \right) dy_1 \\
&= \prod_{i\in[n]} \int_{x_i+\mathcal{P}_\pm} \rho_\sigma(y_i - c_i) dy_i.
\end{aligned}$$

It yields $\mathcal{R}_{\mathbb{Z}^n,\sigma,\mathbf{c}}(\mathbf{x}) = \prod_{i\in[n]} \int_{x_i+\mathcal{P}_\pm} D_\sigma(y_i - c_i) dy_i = \prod_{i\in[n]} \mathcal{R}_{\mathbb{Z},\sigma,c_i}(x_i)$, and therefore $\mathcal{R}_{\mathbb{Z}^n,\sigma,\mathbf{c}} = \otimes_{i\in[n]} \mathcal{R}_{\mathbb{Z},\sigma,c_i}$. By the multiplicativity of the Rényi divergence of Lemma 2.1, we get

$$\begin{aligned}
\mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z}^n,\sigma,\mathbf{c}} \| \mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}\right) &= \mathrm{RD}_a\left(\otimes_{i\in[n]} \mathcal{R}_{\mathbb{Z},\sigma,c_i} \| \otimes_{i\in[n]} \mathcal{D}_{\mathbb{Z},s,c_i}\right) \\
&= \prod_{i\in[n]} \mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z},\sigma,c_i} \| \mathcal{D}_{\mathbb{Z},s,c_i}\right).
\end{aligned}$$

□

We are now left with studying $\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c}\|\mathcal{D}_{\mathbb{Z},s,c})$ for some $c \in \mathbb{R}$. We can further express this quantity as follows.

$$\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c}\|\mathcal{D}_{\mathbb{Z},s,c}) = \left(\rho_s(\mathbb{Z}-c)^{a-1}\sum_{k\in\mathbb{Z}}\frac{\mathcal{R}_{\mathbb{Z},\sigma(s),c}(k)^a}{\rho_s(k-c)^{a-1}}\right)^{\frac{1}{a-1}}.$$

We define $\mathcal{R}'_{\sigma(s)} : x \in \mathbb{R} \mapsto \int_{x+\mathcal{P}_\pm}\rho_{\sigma(s)}(y)dy = \int_{\mathcal{P}_\pm}\rho_{\sigma(s)}(x+u)du$, and $f = {\mathcal{R}'_{\sigma(s)}}^a/\rho_s^{a-1}$. We then have $\mathcal{R}_{\mathbb{Z},\sigma(s),c}(k) = \mathcal{R}'_{\sigma(s)}(k-c)/\sqrt{2\pi\sigma(s)^2}$ for $k \in \mathbb{Z}$, which means we can re-express the quantity above as

$$\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c}\|\mathcal{D}_{\mathbb{Z},s,c})^{a-1} = \frac{\rho_s(\mathbb{Z}-c)^{a-1}}{(\sqrt{2\pi}\sigma(s))^a}\cdot f(\mathbb{Z}-c), \tag{1}$$

where we define $f(\mathbb{Z}-c) = \sum_{k\in\mathbb{Z}} f(k-c)$. The idea onwards to is express $f(\mathbb{Z}-c)$ using $\int_{\mathbb{R}} f(y)dy$, showing the remaining term has a Gaussian decay in s. Then, by expanding f and bounding the remainder terms from said expansion, we show that $\rho_s(\mathbb{Z}-c)^{a-1}(\sqrt{2\pi}\sigma(s))^{-a}\int_{\mathbb{R}} f(y)dy$ has an expansion in $1+\gamma_a/s^8+o(s^{-8})$.

3.2 Smoothing the Discrete Sums

In order to get an exact expansion of $f(\mathbb{Z}-c)$, we express it in terms of its integral $\int_{\mathbb{R}} f(x)dx$ to facilitate subsequent computations. We then wish to apply Lemma 2.3 to the function f. First, because $\mathcal{R}'_{\sigma(s)}$ and ρ_s naturally extend to entire functions over $\mathbb{C}$, we directly have that f also does using the fact that the set of entire functions is commutative ring. We must now derive a bound on f similar to Equation $(*)$.

Let x, y be in $\mathbb{R}$. First, we directly obtain that $\left|\rho_s(x+iy)^{1-a}\right| = \exp(\frac{a-1}{2s^2}(x^2-y^2))$, meaning we can just focus on bounding the numerator. We have the following inequalities.

$$\begin{aligned}
\left|\mathcal{R}'_{\sigma(s)}(x+iy)\right| &\le \int_{\mathcal{P}_\pm}\left|\exp(-(x+iy+u)^2/2\sigma(s)^2)\right|du\\
&= e^{y^2/2\sigma(s)^2}\int_{\mathcal{P}_\pm} e^{-(x+u)^2/2\sigma(s)^2}du\\
&\le e^{y^2/2\sigma(s)^2}\max_{u\in\mathcal{P}_\pm} e^{-(x+u)^2/2\sigma(s)^2}\\
&= e^{y^2/2\sigma(s)^2}e^{-\max(|x|-1/2,0)^2/2\sigma(s)^2}.
\end{aligned}$$

We define $C = a/\sigma(s)^2-(a-1)/s^2$. Notice that for $s > 1/\sqrt{24}$ (so that $\sigma(s) > 0$), it holds that $C > 0$. We now distinguish two cases, whether or not $|x| \le 1/2$. First, if $|x| > 1/2$, we can distribute $(|x|-1/2)^2$ to get

$$|f(x+iy)| \le e^{-a/8\sigma(s)^2}e^{-Cx^2/2+Cy^2/2}e^{a|x|/2\sigma(s)^2}$$

Using that $\alpha\beta \le (\alpha^2+\beta^2)/2$, we have that

$$\frac{a}{2\sigma(s)^2}|x| = \sqrt{\frac{C}{2}}|x|\cdot\frac{a}{\sqrt{2C}\sigma(s)^2} \le \frac{C}{4}x^2 + \frac{a^2}{4C^2\sigma(s)^4},$$

which, when plugged into the equation above, gives

$$|f(x+iy)| \leq e^{-\frac{a}{8\sigma(s)^2}+\frac{a^2}{4C\sigma(s)^4}} e^{-Cx^2/4+Cy^2/2} = K_1(s)\cdot e^{-Cx^2/4+Cy^2/2}.$$

Next, when $|x| \leq 1/2$, we get

$$\begin{aligned}|f(x+iy)| \leq e^{(a-1)x^2/2s^2+Cy^2/2} &= e^{(C/4+(a-1)/2s^2)x^2} e^{-Cx^2/4+Cy^2/2} \\ &= e^{(\frac{a}{4\sigma(s)^2}+\frac{a-1}{4s^2})x^2} e^{-Cx^2/4+Cy^2/2} \\ &\leq e^{(\frac{a}{4\sigma(s)^2}+\frac{a-1}{4s^2})\frac{1}{4}} e^{-Cx^2/4+Cy^2/2}.\end{aligned}$$

Defining $K_2(s) = \exp(a/16\sigma(s)^2 + (a-1)/16s^2)$, and $K = \max(K_1, K_2)$, we have $\forall x, y \in \mathbb{R}, |f(x+iy)| \leq K \cdot e^{-Cx^2/4+Cy^2/2}$. Plugging the expression of C shows that $K_1 \geq K_2$ and therefore $K = K_1$. We can then apply Lemma 2.3 with K, $A = C/2$ and $B = C$ and get

$$f(\mathbb{Z}-c) = \int_{\mathbb{R}} f(x)dx + R_f, \quad \text{with } |R_f| \leq 2K\sqrt{\frac{4\pi}{C}} \cdot \frac{e^{-2\pi^2/C}}{1-e^{-2\pi^2/C}}. \tag{2}$$

We insist that K and C are constants with respect to x, y, c which allows us to use Lemma 2.3, but they naturally depend on s. As we want to track the asymptotic behavior in s, we need to have their exact expression as functions of s.

Before pursuing, we also need to express the factor $\rho_s(\mathbb{Z}-c)^{a-1}/(\sqrt{2\pi}\sigma(s))^a$ from Equation (1). Because ρ_s naturally extend to an entire function over $\mathbb{C}$ and because $|\rho_s(x+iy)| = e^{-x^2/2s^2+y^2/2s^2}$, we can directly use Lemma 2.3 with $(K, A, B) = (1, 1/s^2, 1/s^2)$ to get $\rho_s(\mathbb{Z}-c) = \int_{\mathbb{R}} \rho_s(x)dx + R_{\rho_s} = \sqrt{2\pi}s(1 + R_{\rho_s}/(\sqrt{2\pi}s))$. The lemma also gives $\left|R_{\rho_s}/(\sqrt{2\pi}s)\right| \leq 2e^{-2\pi^2 s^2}/(1-e^{-2\pi^2 s^2})$ which has Gaussian decay in s. As a result, we obtain

$$\frac{\rho_s(\mathbb{Z}-c)^{a-1}}{(\sqrt{2\pi}\sigma(s))^a} \leq \frac{1}{\sqrt{2\pi}s}\left(\frac{s}{\sigma(s)}\right)^a \left(1+\frac{2e^{-2\pi^2 s^2}}{1-e^{-2\pi^2 s^2}}\right)^{a-1}. \tag{3}$$

Actually, we have the exact asymptotic behavior as

$$\frac{\rho_s(\mathbb{Z}-c)^{a-1}}{(\sqrt{2\pi}\sigma(s))^a} = \frac{1}{\sqrt{2\pi}s}\left(\frac{s}{\sigma(s)}\right)^a + o(s^{-N}), \text{ for any } N \in \mathbb{N} \tag{4}$$

3.3 Gaussian Decay of the Remainder

Before diving into the expression of $\int_{\mathbb{R}} f(x)dx$, we can already bound the summand $R = \rho_s(\mathbb{Z}-c)^{a-1}/(\sqrt{2\pi}\sigma(s))^a \cdot R_f$. More precisely, we show that R has a Gaussian decay in s, thanks to R_f.

Lemma 3.2. *It holds that for all $N \in \mathbb{N}$*

$$\frac{\rho_s(\mathbb{Z}-c)^{a-1}}{(\sqrt{2\pi}\sigma(s))^a} \cdot R_f = O(e^{-s^2/2}) = o(s^{-N})$$

For the proof of Lemma 3.2, and also the rest of the main proof, we require a lower and upper bound on $1/C$. Because $s > \sigma(s)$, we directly get $C > a/s^2 - (a-1)/s^2 = 1/s^2$, thus proving $1/C < s^2$. For the lower bound, we can first rewrite $1/C$ as $1/C = s^2/(as^2/\sigma(s)^2 - a + 1)$, from which a straightforward chain of inequalities shows that for any $\alpha < 1$, we have $1/C \geq \alpha s^2$ if and only if

$$s \geq \frac{1}{\sqrt{24}} \cdot \frac{\sqrt[4]{a\alpha + 1 - \alpha}}{\sqrt{\sqrt{a\alpha + 1 - \alpha} - \sqrt{a\alpha}}} =: s_0^\star.$$

We later take $\alpha = 1/4\pi^2$ (which gives the condition $s \gtrsim 1.082$ for $1 < a \leq 512$ which is the typical range we need for a). Hence, for $s \geq s_0^\star$, $s^2/4\pi^2 \leq 1/C \leq s^2$.

Proof (Lemma 3.2). First, using the upper bound $1/C < s^2$, we can bound $K(s)$ for s away from $1/\sqrt{24}$, e.g., $s \geq s_0 > 1/\sqrt{24}$. We directly obtain that for any $s_0 > 1/\sqrt{24}$,

$$\forall s \geq s_0, K(s) \leq \exp\left(\frac{a}{4\sigma(s_0)^2}\left(\frac{a}{(1 - 1/(24s_0^2))^2} - \frac{1}{2}\right)\right).$$

We now look at the other factors in the bound of $|R_f|$ from Equation (2). The factor $\sqrt{4\pi/C}$ can directly be bounded by $\sqrt{4\pi}s$. For the exponential terms, we use the lower bound $1/C \geq s^2/4\pi^2$ that is valid for all $s \geq s_0^\star$. Then,

$$\frac{e^{-2\pi^2/C}}{1 - e^{-2\pi^2/C}} \leq \frac{e^{-s^2/2}}{1 - e^{-s^2/2}} \leq \frac{e^{-s^2/2}}{1 - e^{-s_0^{\star 2}/2}}.$$

Similarly, if we look at Equation (3), we can obtain the following bound for all $s \geq s_0^\star$.

$$\frac{\rho_s(\mathbb{Z} - c)^{a-1}}{(\sqrt{2\pi}\sigma(s))^a} \leq \frac{1}{\sqrt{2\pi}s} \frac{1}{(1 - 1/(24s_0^{\star 2}))^a} \left(1 + \frac{2e^{-2\pi^2 s_0^{\star 2}}}{1 - e^{-2\pi^2 s_0^{\star 2}}}\right)^{a-1}$$

Combined with the bound on $|R_f|$, we get that for all $s \geq s_0^\star$

$$|R| \leq \frac{2\sqrt{2}e^{\frac{a}{4\sigma(s_0^\star)^2}\left(\frac{a}{(1-1/(24s_0^{\star 2}))^2} - \frac{1}{2}\right)}}{(1 - 1/(24s_0^{\star 2}))^a \cdot (1 - e^{-s_0^{\star 2}/2})} \left(1 + \frac{2e^{-2\pi^2 s_0^{\star 2}}}{1 - e^{-2\pi^2 s_0^{\star 2}}}\right)^{a-1} e^{-s^2/2}.$$

Said differently, we have $R = O(e^{-s^2/2})$, which proves R has Gaussian decay. In particular, it yields $R = o(s^{-N})$ for all $N \in \mathbb{N}$. □

We notice that at this stage, we have completely smoothed out the center c in the terms with exponential decay. This is rather intuitive as we already know that for parameters s exceeding the smoothing parameter of $\mathbb{Z}$ (which is relatively small), the shift c has no noticeable impact on the discrete Gaussian mass of $\mathbb{Z}$.

3.4 Expanding the function f

The objective now is to determine the asymptotic behavior in s of the integral, as Equation (4) already gives that of the factor in front of the integral. We first give another expression of $\mathcal{R}'_{\sigma(s)}(x)$ without the integral to facilitate the expansion of $\int_{\mathbb{R}} f(x)dx$.

Lemma 3.3. *For any $k \in \mathbb{N}$, we define the k-th probabilist's Hermite polynomial* He_k *as* $\mathrm{He}_k(x) = (-1)^k e^{x^2/2}\rho_1^{(k)}(x)$. *It then holds that for any $N \in \mathbb{N}^\times$,*

$$\forall x \in \mathbb{R}, \mathcal{R}'_{\sigma(s)}(x) = \rho_{\sigma(s)}(x)\left(\sum_{k=0}^{N} \frac{\mathrm{He}_{2k}\left(\frac{x}{\sigma(s)}\right)}{\sigma(s)^{2k}\cdot 2^{2k}\cdot(2k+1)!} + R_{2N+2}(x,\sigma(s))\right),$$

such that for any $s_0 > 1/\sqrt{24}$, it holds that

$$\forall t > 0, \forall x \in [-t\sigma(s), t\sigma(s)],$$
$$|R_{2N+2}(x,\sigma(s))| \leq \frac{e^{\frac{t}{2\sigma(s)}+\frac{1}{8\sigma(s)^2}}\cdot \max_{|z|\leq t+1/2\sigma(s_0)} |\mathrm{He}_{2N+2}(z)|}{\sigma(s)^{2N+2}} \cdot \frac{1}{(2N+3)!2^{2N+2}}.$$

Proof (Lemma 3.3). Let x be in $\mathbb{R}$. We let $h_x : u \in \mathcal{P}_\pm \mapsto x+u$. By a simple induction relying on the chain rule for derivatives, we directly have that $(\rho_{\sigma(s)} \circ h_x)^{(k)} = \rho_{\sigma(s)}^{(k)} \circ h_x$ for any $k \in \mathbb{N}$. Additionally, another induction based on the chain rule shows that for all $k \in \mathbb{N}$, $\rho_{\sigma(s)}^{(k)}(y) = \sigma(s)^{-k}\rho_1^{(k)}(y/\sigma(s))$ for all $y \in \mathbb{R}$. As a result, we have that for all $y \in \mathbb{R}$,

$$\begin{aligned}
\rho_{\sigma(s)}^{(k)}(y) &= (-1)^k \cdot (-1)^k e^{(y/\sigma(s))^2/2}\rho_{\sigma(s)}(y)\rho_{\sigma(s)}^{(k)}(y) \\
&= (-1)^k \cdot (-1)^k e^{(y/\sigma(s))^2/2}\rho_{\sigma(s)}(y)\rho_1^{(k)}(y/\sigma(s))\cdot\sigma(s)^{-k} \\
&= (-1)^k \sigma(s)^{-k}\mathrm{He}_k(y/\sigma(s))\cdot\rho_{\sigma(s)}(y).
\end{aligned}$$

Taylor's theorem then gives that for all $M \in \mathbb{N}$,

$$\forall u \in \mathcal{P}_\pm, \exists \xi(u) \in [u,0]\cup[0,u],$$
$$\rho_{\sigma(s)}(x+u) = \sum_{k=0}^{M} \frac{(\rho_{\sigma(s)}\circ h_x)^{(k)}(0)}{k!}u^k + \frac{(\rho_{\sigma(s)}\circ h_x)^{(M+1)}(\xi(u))}{(M+1)!}u^{M+1}$$

Using the previous identities, we have

$$\begin{aligned}
\rho_{\sigma(s)}(x+u) = &\sum_{k=0}^{M} \frac{(-1)^k \mathrm{He}_k\left(\frac{x}{\sigma(s)}\right)\rho_{\sigma(s)}(x)}{\sigma(s)^k\cdot k!}u^k \\
&+ \frac{(-1)^{M+1}\mathrm{He}_{M+1}\left(\frac{x+\xi(u)}{\sigma(s)}\right)\rho_{\sigma(s)}(x+\xi(u))}{\sigma(s)^{M+1}\cdot(M+1)!}u^{M+1}
\end{aligned}$$

We now integrate over $u \in \mathcal{P}_\pm$. The first term reduces to computing $\int_{\mathcal{P}_\pm} u^k du$. A straightfoward integration gives that $\int_{\mathcal{P}_\pm} u^k du = \frac{1}{2^k(k+1)} \mathbf{1}_{2\mathbb{N}}(k)$, meaning it is 0 when k is odd and $1/(2^k(k+1))$ otherwise. We can then take $M = 2N+1$ for any $N \in \mathbb{N}$, making the first term equal to

$$\sum_{k=0}^{2N+1} \frac{(-1)^k \mathrm{He}_k\left(\frac{x}{\sigma(s)}\right) \rho_{\sigma(s)}(x)}{\sigma(s)^k \cdot k! \cdot (k+1)2^k} \mathbf{1}_{2\mathbb{N}}(k) = \rho_{\sigma(s)}(x) \sum_{k=0}^{N} \frac{\mathrm{He}_{2k}\left(\frac{x}{\sigma(s)}\right)}{2^{2k}\sigma(s)^{2k}(2k+1)!},$$

as desired. We then define

$$R_{2N+2}(x, \sigma(s)) = \int_{\mathcal{P}_\pm} \frac{\mathrm{He}_{2N+2}\left(\frac{x+\xi(u)}{\sigma(s)}\right) \rho_{\sigma(s)}(x+\xi(u))}{\sigma(s)^{2N+2} \cdot (2N+2)! \cdot \rho_{\sigma(s)}(x)} u^{2N+2} du$$

which by construction gives the first claimed identity. We now only need to prove the second statement on this remainder $R_{2N+2}(x, \sigma(s))$. Let $t > 0$ and $x \in [-t\sigma(s), t\sigma(s)]$. First, let us deal with the term $\mathrm{He}_{2N+2}((x+\xi(u))/\sigma(s))$. For any $u \in \mathcal{P}_\pm$, it holds that $|\xi(u)| \leq |u| \leq 1/2$. Hence $|(x+\xi(u))/\sigma(s)| \leq t + 1/2\sigma(s)$. As $1/\sigma(s)$ is a decreasing function of s, it holds that for any $s_0 > 1/\sqrt{24}$, $|(x+\xi(u))/\sigma(s)| \leq t + 1/2\sigma(s_0)$. Typically, for $s_0 = 1/2$, the bound is $t + 6/5$. Then, we have

$$\left|\mathrm{He}_{2N+2}\left(\frac{x+\xi(u)}{\sigma(s)}\right)\right| \leq \max_{|z| \leq t+1/2\sigma(s_0)} |\mathrm{He}_{2N+2}(z)|.$$

Later on, we take t that grows with s, in which case the bound is asymptotically $O(t^{2N+2})$ due to He_{2N+2} being a monic polynomial of degree $2N+2$. We then look at the Gaussian ratio. We have

$$\frac{\rho_{\sigma(s)}(x+\xi(u))}{\rho_{\sigma(s)}(x)} = e^{-\frac{2x\xi(u)+\xi(u)^2}{2\sigma(s)^2}} \leq e^{\frac{|x|}{2\sigma(s)^2} + \frac{1}{8\sigma(s)^2}} \leq e^{\frac{t}{2\sigma(s)} + \frac{1}{8\sigma(s)^2}}.$$

Then, using that $\int_{\mathcal{P}_\pm} |u|^{2N+2} du = \int_{\mathcal{P}_\pm} u^{2N+2} du = 1/((2N+3)2^{2N+2})$, we get

$$|R_{2N+2}(x, \sigma(s))| \leq \frac{e^{\frac{t}{2\sigma(s)} + \frac{1}{8\sigma(s)^2}} \cdot \max_{|z| \leq t+1/2\sigma(s_0)} |\mathrm{He}_{2N+2}(z)|}{\sigma(s)^{2N+2}} \cdot \frac{1}{(2N+3)!2^{2N+2}},$$

as claimed. □

By Lemma 3.3, we can express $f(x)$ as

$$\begin{aligned}
\forall x \in \mathbb{R}, f(x) &= \frac{\rho_{\sigma(s)}(x)^a}{\rho_s(x)^{a-1}} \left(\sum_{k=0}^{N} \frac{\mathrm{He}_{2k}(x/\sigma(s))}{\sigma(s)^{2k} \cdot 2^{2k} \cdot (2k+1)!} + R_{2N+2}(x, \sigma(s)) \right)^a \\
&= e^{-Cx^2/2} \left(\sum_{k=0}^{N} \frac{\mathrm{He}_{2k}(x/\sigma(s))}{\sigma(s)^{2k} \cdot 2^{2k} \cdot (2k+1)!} + R_{2N+2}(x, \sigma(s)) \right)^a
\end{aligned}$$

Given we want to study the order 8 in $1/s$, it suffices to select $N = 4$. We denote $P(x, \sigma(s)) = \sum_{k=1}^{4} \mathrm{He}_{2k}(x/\sigma(s))/(\sigma(s)^{2k} \cdot 2^{2k} \cdot (2k+1)!)$ and $Q(x, \sigma(s)) = 1 + P(x, \sigma(s))$. We can further develop the expression of f as follows.

$$\begin{aligned}
f(x)e^{Cx^2/2} &= (Q(x,\sigma(s)) + R_{10}(x,\sigma(s)))^a \\
&= \sum_{m=0}^{a} \tbinom{a}{m} Q(x,\sigma(s))^m R_{10}(x,\sigma(s))^{a-m} \\
&= Q(x,\sigma(s))^a + \sum_{m=0}^{a-1} \tbinom{a}{m} Q(x,\sigma(s))^m R_{10}(x,\sigma(s))^{a-m} \\
&= \sum_{j=0}^{a} \tbinom{a}{j} P(x,\sigma(s))^j + \sum_{m=0}^{a-1} \tbinom{a}{m} Q(x,\sigma(s))^m R_{10}(x,\sigma(s))^{a-m} \\
&= \sum_{j=0}^{4} \tbinom{a}{j} P(x,\sigma(s))^j + \sum_{j=5}^{a} \tbinom{a}{j} P(x,\sigma(s))^j \\
&\qquad + \sum_{m=0}^{a-1} \tbinom{a}{m} Q(x,\sigma(s))^m R_{10}(x,\sigma(s))^{a-m} \qquad (5)
\end{aligned}$$

Note that Equation (5) splits the sum to separate powers of P with exponent less than 4 as they will contribute to the terms $1/s^k$ with $k \leq 8$ in the expansion. Following the convention that empty sums are 0 and that $\binom{k}{\ell} = 0$ for $k < \ell$, the last splitting of the sum remains valid even if $a \leq 4$.

3.5 Splitting Integration at the Tail

Even with the Gaussian weight $e^{-Cx^2/2}$, the remainder $R_{10}(x, \sigma(s))$ makes it non-trivial to integrate f, especially because Lemma 3.3 only provides a bound for $|x| \leq t\sigma(s)$. Because the Hermite polynomial would go to infinity with x, a general bound for all $x \in \mathbb{R}$ is not possible. The idea is then to carefully select some tailcut $t(s)$ which may depend on $s, \sigma(s)$, and split the integral at this tailcut. On $[-t(s)\sigma(s), t(s)\sigma(s)]$, we can use the bound on $R_{10}(x, \sigma(s))$ to facilitate integration, while outside $[-t(s)\sigma(s), t(s)\sigma(s)]$, $e^{-Cx^2/2}$ is already sufficiently small to make it scale as $o(s^{-8})$.

We then set $t(s) = b\sqrt{\ln \sigma(s)}$ for $b = \sqrt{112}$, and we define the interval $\mathcal{I}_s = [-t(s)\sigma(s), t(s)\sigma(s)]$ and $\overline{\mathcal{I}_s} = \mathbb{R} \setminus \mathcal{I}_s$. Note that this choice of $t(s)$ is sufficient for our subsequent asymptotic arguments to guarantee that all the parasitic terms are indeed $o(s^{-8})$. Proving a closed-form bound on the $o(s^{-8})$ may require a different choice, where the value of the constant b may likely impact said bound. In that case, a smaller b would be better. In our case, we only need b to be a constant. We then have

$$\int_{\mathbb{R}} f(x)dx = \int_{\mathcal{I}_s} f(x)dx + \int_{\overline{\mathcal{I}_s}} f(x)dx.$$

We first look at the second integral in the following lemma.

Lemma 3.4. *It holds that* $\int_{\overline{\mathcal{I}_s}} f(x)dx = o(s^{-8})$.

Proof (Lemma 3.4). Here, we do not need the expansion of f from Section 3.4. Indeed, recall from Section 3.2 that we proved the inequality

$$\forall x, y \in \mathbb{R}, |f(x+iy)| \leq K(s)e^{-Cx^2/4+Cy^2/2},$$

with $K(s) \leq \exp(\frac{a}{4\sigma(s_0)^2}(\frac{a}{(1-1/(24s_0^2))^2} - \frac{1}{2})) = K^\star(s_0)$ for any $s \geq s_0 > 1/\sqrt{24}$, and $C = a/\sigma(s)^2 - (a-1)/s^2$. Applied with $y = 0$ and with the positivity of f, we get $f(x) \leq K^\star(s_0)e^{-Cx^2/4}$ for any fixed constant $s_0 > 1/\sqrt{24}$ and $s \geq s_0$.

We also have that for $x \in \overline{\mathcal{I}_s}$, it holds that $Cx^2/4 = Cx^2/8 + Cx^2/8 \geq Cx^2/8 + C(t(s)\sigma(s))^2/8$. Then, we can bound the integral as

$$\begin{aligned}\left|\int_{\overline{\mathcal{I}_s}} f(x)dx\right| &\leq K^\star(s_0)e^{-Ct(s)^2\sigma(s)^2/8}\int_{\overline{\mathcal{I}_s}} e^{-Cx^2/8}dx \\ &\leq K^\star(s_0)e^{-Ct(s)^2\sigma(s)^2/8}\int_{\mathbb{R}} e^{-Cx^2/8}dx \\ &\leq K^\star(s_0)e^{-Ct(s)^2\sigma(s)^2/8}\sqrt{8\pi/C}.\end{aligned}$$

We have that $C\sigma(s)^2 = a - (a-1)(1-1/(24s^2))^2$ defines a continuous and decreasing function over $(1/\sqrt{24}, \infty)$, and tends to $a-(a-1) = 1$ when s goes to infinity. As a result, we have that $C\sigma(s)^2 \geq 1$ for any $s > 1/\sqrt{24}$. Also, recall that $1/C < s^2$. Combining both gives that for $s > s_0$,

$$\left|\int_{\overline{\mathcal{I}_s}} f(x)dx\right| \leq K^\star(s_0)\cdot\sqrt{8\pi}s\cdot e^{-b^2\ln(\sigma(s))/8} = K^\star(s_0)\sqrt{8\pi}\frac{s}{\sigma(s)^{b^2/8}}$$

Because $\sigma(s) \sim s$, and that $b^2/8 - 1 \geq 9$ (as $b \geq \sqrt{80}$), we then obtain that

$$\int_{\overline{\mathcal{I}_s}} f(x)dx = O(s^{-9}) = o(s^{-8}),$$

as claimed. Actually, for the chosen constant b, we obtain $\int_{\overline{\mathcal{I}_s}} f(x)dx = O(s^{-13}) = o(s^{-12})$. We did not choose a smaller b, e.g., $\sqrt{80}$, because it is also constrained by other arguments later on, requiring it to be larger than $\sqrt{80}$. □

3.6 Bounding the Remainder Terms over $\mathcal{I}_s$

We are now left with studying $\int_{\mathcal{I}_s} f(x)dx$. This is where we use the expansion of f from Section 3.4, which splits $f(x)$ into three terms defined in Equation (5), each with a Gaussian weight. By using the expression of $P(x, \sigma(s))$ and the bound on $R_{10}(x, \sigma(s))$ over $\mathcal{I}_s$, we show that the second and third integrals are indeed negligible with respect to s^{-8}. For that, we use the following asymptotics on $P(x, \sigma(s))$ and $R_{10}(x, \sigma(s))$. Note that the bound on the remainder derived in Lemma 3.3 over $[-t(s)\sigma(s), t(s)\sigma(s)]$ depends on $t(s)$ and $\sigma(s)$. We therefore have to show that this part does not contribute to the terms in $1/s^k$ for $k \leq 8$.

Lemma 3.5. *There exists $K_R(\sigma(s))$ and $K_P(\sigma(s))$ functions of $\sigma(s)$ such that for any $s_0 > 1/\sqrt{24}$, any $s \geq s_0$, and any $x \in \mathcal{I}_s$, $|R_{10}(x, \sigma(s))| \leq K_R(\sigma(s))$ and $|P(x, \sigma(s))| \leq K_P(\sigma(s))$. We additionally have that*

$$K_R(\sigma(s)) = O(\ln(\sigma(s))^5/\sigma(s)^{10}), \quad \text{and } K_P(\sigma(s)) = O(\ln(\sigma(s))/\sigma(s)^2).$$

Proof (Lemma 3.5). Let $s_0 > 1/\sqrt{24}$. For any $s \geq s_0$ and $x \in \mathcal{I}_s$, Lemma 3.3 with $t = t(s) = b\sqrt{\ln(\sigma(s))}$ yields

$$|R_{10}(x, \sigma(s))| \leq \frac{e^{\frac{b\sqrt{\ln(\sigma(s))}}{2\sigma(s)} + \frac{1}{8\sigma(s)^2}} \cdot \max_{|z| \leq b\sqrt{\ln(\sigma(s))} + 1/2\sigma(s_0)} |\mathrm{He}_{10}(z)|}{\sigma(s)^{10}} \cdot \frac{1}{11! \cdot 2^{10}}.$$

The bound is indeed a function of $\sigma(s)$, call it $K_R(\sigma(s))$. We then have that $b\sqrt{\ln(\sigma(s))}/2\sigma(s) + 1/8\sigma(s)^2$ goes to 0 as s goes to infinity, which means

$$e^{\frac{b\sqrt{\ln(\sigma(s))}}{2\sigma(s)} + \frac{1}{8\sigma(s)^2}} = O(1).$$

As He_{10} is a monic degree-10 polynomial, we have that for a sufficiently large s, $\max_{|z| \leq t(s) + 1/2\sigma(s_0)} |\mathrm{He}_{10}(z)| = \mathrm{He}_{10}(t(s) + 1/2\sigma(s_0))$. Yet because $\mathrm{He}_{10}(z) \sim z^{10}$ when z goes to infinity, it yields $\mathrm{He}_{10}(t(s) + 1/2\sigma(s_0)) \sim (t(s) + 1/2\sigma(s_0))^{10} \sim b^{10} \ln(\sigma(s))^5$. We then conclude that $\max_{|z| \leq t(s) + 1/2\sigma(s_0)} |\mathrm{He}_{10}(z)| = O(\ln(\sigma(s))^5)$. It then yields that $K_R(\sigma(s)) = O(\ln(s)^5/\sigma(s)^{10})$, as desired.

Then, for P, we directly have that for all $x \in \mathcal{I}_s$,

$$|P(x, \sigma(s))| \leq \sum_{k=1}^{4} \frac{\max_{|z| \leq b\sqrt{\ln(\sigma(s))}} |\mathrm{He}_{2k}(z)|}{\sigma(s)^{2k} \cdot 2^{2k} \cdot (2k+1)!} =: K_P(\sigma(s))$$

Using again that He_{2k} is a monic degree-$2k$ polynomial, for a sufficiently large s, we have $\max_{|z| \leq t(s)} |\mathrm{He}_{2k}(z)| = \mathrm{He}_{2k}(t(s)) \sim b^{2k} \ln(\sigma(s))^k$. As a result, the k-th term is in $O(\ln(\sigma(s))^k/\sigma(s)^{2k})$. This means that $K_P(\sigma(s)) = O(\ln(\sigma(s))/\sigma(s)^2)$ as the dominant term is the first one for $k = 1$. □

By Lemma 3.5, we directly get that, over $\mathcal{I}_s$, $|Q(x, \sigma(s))| \leq 1 + K_P(\sigma(s)) = O(1)$ as $\ln(\sigma)/s^2 \to 0$. We now look at the second and third integrals issued from Equation (5). More concretely, we define I_2 and I_3 as

$$I_2 = \sum_{j=5}^{a} \binom{a}{j} \int_{\mathcal{I}_s} e^{-Cx^2/2} P(x, \sigma(s))^j dx$$

$$I_3 = \sum_{m=0}^{a-1} \binom{a}{m} \int_{\mathcal{I}_s} e^{-Cx^2/2} Q(x, \sigma(s))^m R_{10}(x, \sigma(s))^{a-m} dx.$$

We show that both of these integrals are negligible compared to $1/s^8$.

Lemma 3.6. *It holds that $I_2 = o(s^{-8})$ and $I_3 = o(s^{-8})$.*

Proof (Lemma 3.6). Let $s_0 > 1/\sqrt{24}$ and fix $s \geq s_0$. Lemma 3.5 gives that for all $x \in \mathcal{I}_s$, $\left|P(x, \sigma(s))^j\right| \leq K_P(\sigma(s))^j$. Hence, we have

$$\begin{aligned}
|I_2| &\leq \sum_{j=5}^{a} \binom{a}{j} K_P(\sigma(s))^j \int_{\mathcal{I}_s} e^{-Cx^2/2} dx \\
&= \left(\int_{\mathcal{I}_s} e^{-Cx^2/2} dx\right) \cdot K_P(\sigma(s))^5 \sum_{j=0}^{a-5} \binom{a}{j+5} K_P(\sigma(s))^j \\
&\leq \sqrt{\frac{2\pi}{C}} \cdot K_P(\sigma(s))^5 \sum_{j=0}^{a-5} \binom{a}{j+5} K_P(\sigma(s))^j,
\end{aligned}$$

where the last inequality uses that $\int_{\mathcal{I}_s} e^{-Cx^2/2} dx \leq \int_{\mathbb{R}} e^{-Cx^2/2} dx = \sqrt{2\pi/C}$. Then, we have $\sqrt{2\pi/C} \leq \sqrt{2\pi} s$. By Lemma 3.5, we also have $K_P(\sigma(s))^5 = O(\ln(\sigma(s))^5/\sigma(s)^{10})$. Additionally, the sum is dominated by the first term, which means $\sum_{j=0}^{a-5} \binom{a}{j+5} K_P(\sigma(s))^j = O(1)$. As a result, we obtain that the bound on $|I_2|$ is $O(s \ln(\sigma(s))^5/\sigma(s)^{10}) = O(\ln(s)^5/s^9)$. Hence $I_2 = O(\ln(s)^5/s^9)$. Because $\ln(s)^5/s \to 0$, we indeed get $I_2 = o(s^{-8})$. Note that all the above reasoning remain valid when $a \leq 4$ as $I_2 = 0$ due to the empty sum, which is indeed negligible before s^{-8}.

A similar strategy based on Lemma 3.5 allows us to bound I_3. For all $x \in \mathcal{I}_s$, we have $|Q(x, \sigma(s))^m| \leq (1 + K_P(\sigma(s))^m$ and $|R_{10}(x, \sigma(s))|^{a-m} \leq K_R(\sigma(s))^{a-m}$. Then, we get

$$\begin{aligned}
|I_3| &\leq \left(\int_{\mathcal{I}_s} e^{-Cx^2/2} dx\right) \sum_{m=0}^{a-1} \binom{a}{j} (1 + K_P(\sigma(s)))^m K_R(\sigma(s))^{a-m} \\
&= \left(\int_{\mathcal{I}_s} e^{-Cx^2/2} dx\right) \cdot K_R(\sigma(s)) \sum_{m=0}^{a-1} \binom{a}{j} (1 + K_P(\sigma(s)))^m K_R(\sigma(s))^{a-1-m} \\
&\leq \sqrt{\frac{2\pi}{C}} \cdot K_R(\sigma(s)) \sum_{m=0}^{a-1} \binom{a}{j} (1 + K_P(\sigma(s)))^m K_R(\sigma(s))^{a-1-m},
\end{aligned}$$

Because $1 + K_P(\sigma(s)) = O(1)$, the sum is asymptotically dominated by the term of index $m = a - 1$, which is $a(1 + K_P(\sigma(s)))^{a-1} = O(1)$. Then, combined with $1/\sqrt{C} \leq s$ and $K_R(\sigma(s)) = O(\ln(\sigma(s))^5/\sigma(s)^{10})$ by Lemma 3.5, we get the bound on $|I_3|$ is $O(s \ln(\sigma(s))^5/\sigma(s)^{10}) = O(\ln(s)^5/s^9)$. Hence, we have $I_3 = O(\ln(s)^5/s^9) = o(s^{-8})$ as desired. □

3.7 Computing the Main Integral

We are then left with computing the integral I_1 defined by

$$I_1 = \int_{\mathcal{I}_s} e^{-Cx^2/2} \sum_{j=0}^{4} \binom{a}{j} P(x, \sigma(s))^j dx$$

or rather its asymptotic behavior so that $\frac{1}{\sqrt{2\pi}} I_1 s^{a-1}/\sigma(s)^a = 1 + \gamma_a/s^8 + o(s^{-8})$. Indeed, at this stage, Lemma 3.4 and 3.6 have shown that

$$\int_{\mathbb{R}} f(x)dx = \int_{\mathcal{I}_s} f(x)dx + \int_{\overline{\mathcal{I}_s}} f(x)dx = I_1 + I_2 + I_3 + \int_{\overline{\mathcal{I}_s}} f(x)dx = I_1 + o(s^{-8}).$$

Combined with Equations (1), (2), and (4), and Lemma 3.2, it yields

$$\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c} \| \mathcal{D}_{\mathbb{Z},s,c})^{a-1} = \frac{1}{\sqrt{2\pi}s} \left(\frac{s}{\sigma(s)}\right)^a I_1 + o(s^{-9}). \tag{6}$$

We can first expand the polynomial $T = \sum_{j=0}^{4} \binom{a}{j} P(x, \sigma(s))^j$. As each He_{2k} is a polynomial of degree $2k$, it holds $P(x, \sigma(s))$ is a polynomial of degree at most 8 in x. Hence, T is a polynomial of degree at most 32. We can then write $T = \sum_{\ell=0}^{32} T_\ell(\sigma(s)) x^\ell$. As the He_{2k} only have monomials of even degree, it directly holds that T does too. So $T = \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) x^{2\ell}$. Also, due to the expression of P, each of its coefficients is itself a polynomial in $1/\sigma(s)^2$. As a result, the $T_{2\ell}(\sigma(s))$ are polynomials in $1/\sigma(s)^2$ as well. When $a \leq 4$, we also have $T_{2\ell}(\sigma(s)) = 0$ for $\ell \in (4a, 16]$, but we won't need to distinguish this case. Let us now give an expression of I_1 in terms of the $T_{2\ell}(\sigma(s))$.

Lemma 3.7. *It holds that* $I_1 = \sqrt{\frac{2\pi}{C}} \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) \frac{(2\ell-1)!!}{C^\ell} + o(s^{-8})$.

Proof (Lemma 3.7). By distributing the Gaussian weight and by linearity of the integral, we first have

$$I_1 = \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) \int_{\mathcal{I}_s} x^{2\ell} e^{-Cx^2/2} dx.$$

Then, it holds that $\int_{\mathcal{I}_s} x^{2\ell} e^{-Cx^2/2} dx = \int_{\mathbb{R}} x^{2\ell} e^{-Cx^2/2} dx - \int_{\overline{\mathcal{I}_s}} x^{2\ell} e^{-Cx^2/2} dx$. The first term directly equates to $\int_{\mathbb{R}} x^{2\ell} e^{-Cx^2/2} dx = \sqrt{2\pi/C} \cdot (2\ell-1)!!/C^\ell$, with the convention that $(-1)!!$ is an empty product and thus $(-1)!! = 1$. Indeed, for $\ell = 0$, we have $\int_{\mathbb{R}} x^{2\ell} e^{-Cx^2/2} dx = \int_{\mathbb{R}} e^{-Cx^2/2} dx = \sqrt{2\pi/C}$. We then have

$$I_1 = \sqrt{\frac{2\pi}{C}} \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) \frac{(2\ell-1)!!}{C^\ell} - \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) \int_{\overline{\mathcal{I}_s}} x^{2\ell} e^{-Cx^2/2} dx.$$

It only remains to show the second term is negligible in s^{-8}.

We first factor $x^{2\ell} e^{-Cx^2/2} = (x^{2\ell} e^{-Cx^2/4}) e^{-Cx^2/4}$. For any $\ell \in [0, 16]$, let us define $g_{\ell,s}(x) = x^{2\ell} e^{-Cx^2/4}$. A naive argument is to argue that $g_{\ell,s}$ goes to 0 when x goes to infinity, and is eventually bounded by 1. Then for a sufficiently large s, we would have that $g_{\ell,s}(x) \leq 1$ for all $x \in \overline{\mathcal{I}_s}$. This would reduce the problem to bounding $\int_{\overline{\mathcal{I}_s}} e^{-Cx^2/4} dx$. Unfortunately, this fails as C goes to 0 with s. To avoid this double limit issue, we give an explicit bound on $g_{\ell,s}(x)$

over $[t(s)\sigma(s), \infty)$. By a simple differentiation, it holds that $g_{\ell,s}$ increases over $[0, x^\star(s))$ and decreases over $[x^\star(s), \infty)$, where

$$x^\star(s) = 2\sqrt{\frac{\ell}{C}}$$

We then have $x^\star(s) = 2\sigma(s)\sqrt{\ell/(C\sigma(s)^2)} \leq 2\sigma(s)\sqrt{\ell}$ as $C\sigma(s)^2 \geq 1$. Then, if $t(s) \geq 2\sqrt{\ell}$, we get that $t(s)\sigma(s) \geq x^\star(s)$. The former condition is equivalent to $\sigma(s) \geq e^{4\ell/b^2}$. Hence, there exists some $s_0 > 1/\sqrt{24}$ (for $\ell \leq 16$, we have $s_0 \leq 1.73$) such that for all $s \geq s_0$, $g_{\ell,s}$ decreases over $[t(s)\sigma(s), \infty)$ and therefore

$$\sup_{x \in [t(s)\sigma(s), \infty)} g_{\ell,s}(x) = g_{\ell,s}(t(s)\sigma(s)) = t(s)^{2\ell}\sigma(s)^{2\ell} e^{-Ct(s)^2\sigma(s)^2/4}.$$

By parity of $g_{\ell,s}$, the same bound holds over all $\overline{\mathcal{I}_s}$. We then get

$$\int_{\overline{\mathcal{I}_s}} x^{2\ell} e^{-Cx^2/2} dx \leq t(s)^{2\ell}\sigma(s)^{2\ell} e^{-Ct(s)^2\sigma(s)^2/4} \int_{\overline{\mathcal{I}_s}} e^{-Cx^2/4} dx$$

Using the same computation as in the proof of Lemma 3.4, we have the bound $\int_{\overline{\mathcal{I}_s}} e^{-Cx^2/4} dx \leq e^{-Ct(s)^2\sigma(s)^2/8}\sqrt{8\pi/C}$. This then shows

$$\begin{aligned}
\int_{\overline{\mathcal{I}_s}} x^{2\ell} e^{-Cx^2/2} dx &\leq t(s)^{2\ell}\sigma(s)^{2\ell}\sqrt{\frac{8\pi}{C}} e^{-3Ct(s)^2\sigma(s)^2/8} \\
&\leq t(s)^{2\ell}\sigma(s)^{2\ell} \cdot \sqrt{8\pi} s \cdot e^{-3t(s)^2/8} \\
&= \sqrt{8\pi} b^{2\ell} \frac{\ln(\sigma(s))^\ell s}{\sigma(s)^{3b^2/8 - 2\ell}}
\end{aligned}$$

As each $T_{2\ell}(\sigma(s))$ is a polynomial in $1/\sigma(s)^2$, we have that $T_{2\ell}(\sigma(s)) = O(1)$ at worse (as explained above, some of them may even be 0 when $a \leq 4$, but this has no incidence on the rest of the argument). As a result, we get

$$\begin{aligned}
\sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) \int_{\overline{\mathcal{I}_s}} x^{2\ell} e^{-Cx^2/2} dx &= O\left(\frac{s \cdot \ln(\sigma(s))^{16}}{\sigma(s)^{3b^2/8-32}}\right) \\
&= O\left(\frac{\ln(s)^{16}}{s^{3b^2/8-33}}\right).
\end{aligned}$$

As $b = \sqrt{112}$, we get that the exponent of s is $3b^2/8 - 33 = 9$, and thus

$$\sum_{\ell=0}^{16} T_{2\ell}(\sigma(s)) \int_{\overline{\mathcal{I}_s}} x^{2\ell} e^{-Cx^2/2} dx = O\left(\frac{\ln(s)^{16}}{s^9}\right) = o(s^{-8}),$$

which finishes the proof. □

3.8 Formal Taylor Expansion

Combining Lemma 3.7 with Equation (6) gives that

$$\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c}\|\mathcal{D}_{\mathbb{Z},s,c})^{a-1} = \frac{1}{\sqrt{C}s}\left(\frac{s}{\sigma(s)}\right)^a \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s))\frac{(2\ell-1)!!}{C^\ell} + o(s^{-9}).$$

It means the Taylor expansion of order 9 of $\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c}\|\mathcal{D}_{\mathbb{Z},s,c})^{a-1}$ in $1/s$ at 0 is that of $\frac{1}{\sqrt{C}s}\left(\frac{s}{\sigma(s)}\right)^a \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s))\frac{(2\ell-1)!!}{C^\ell}$. As the computation of such a high order Taylor expansion is rather tedious, especially in tracking all the coefficients, we can instead do it using rational symbolic computations in Sagemath [The25] directly from this expression. Indeed, because the Hermite Polynomials $\mathrm{He}_0, \mathrm{He}_2, \mathrm{He}_4, \mathrm{He}_6, \mathrm{He}_8$ have the following closed-form expressions

$$\begin{aligned}
\mathrm{He}_0(x) &= 1\\
\mathrm{He}_2(x) &= x^2 - 1\\
\mathrm{He}_4(x) &= x^4 - 6x^2 + 3\\
\mathrm{He}_6(x) &= x^6 - 15x^4 + 45x^2 - 15\\
\mathrm{He}_8(x) &= x^8 - 28x^6 + 210x^4 - 420x^2 + 105,
\end{aligned}$$

we can compute closed-form expressions of the $T_{2\ell}(\sigma(s))$. We note that if one wants to compute the Taylor expansion manually, we can already observe (and verify using symbolic computations) that $T_{2\ell}(\sigma(s)) = O(1/\sigma(s)^{4\ell})$. This means that for $\ell \in [5, 16]$, $T_{2\ell}(\sigma(s))(2\ell-1)!!/C^\ell = O(s^{2\ell}/s^{4\ell}) = O(1/s^{10}) = o(1/s^9)$. One can therefore cut the sum and only consider the terms of index $\ell \in [0, 4]$. For completeness, we avoid this truncation when using symbolic computations.

As we want the result for an arbitrary integer order $a \geq 2$, we keep a as a symbolic variable in our computation. This however complicates the Taylor expansion of $(s/\sigma(s))^a$ in Sage (which stops before computing the a-th power). We therefore simplify the expression manually here. First, we have that

$$\frac{s}{\sigma(s)} = \frac{1}{1 - \frac{1}{24s^2}} = 1 + \frac{1}{24s^2} + \frac{1}{24^2 s^4} + \frac{1}{24^3 s^6} + \frac{1}{24^4 s^8} + o(s^{-9})$$

Writing this expansion as $1 + F(s)$, we can take the a-th power and compute $\sum_{k=0}^{a} \binom{a}{k} F(s)^k$. Because $F(s) = O(1/s^2)$, then $F(s)^k = O(1/s^{10})$ for $k \in [5, a]$. Hence, we get that $(s/\sigma(s))^a = \sum_{k=0}^{4} \binom{a}{k} F(s)^k + o(s^{-9})$. This can then be handled with symbolic computations. Additionally, when a is a symbolic variable, it may not cancel the terms $\binom{a}{k}$ for $k > a$. Following our proof above, this is only relevant when $a \leq 4$. Our formal computation thus treats this case $a \leq 4$ separately. The algorithm of Figure 3.1 then computes the desired Taylor expansion of order 9 in $1/s$ when s goes to infinity. In particular, the script's output shows

$$\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z},\sigma(s),c}\|\mathcal{D}_{\mathbb{Z},s,c})^{a-1} = 1 + \frac{73a(a-1)}{33177600} \cdot \frac{1}{s^8} + o\left(\frac{1}{s^9}\right) = 1 + \frac{\gamma_a}{s^8} + o\left(\frac{1}{s^9}\right).$$

```
low_a = False # change for low Renyi orders 2 <= a <= 4
a = 2 if low_a else var("a") # can change a to 2,3,4
var("isig,invs,x")
He0 = 1
He2 = (x*isig)^2 - 1
He4 = (x*isig)^4 - 6*(x*isig)^2 + 3
He6 = (x*isig)^6 - 15*(x*isig)^4 + 45*(x*isig)^2 - 15
He8 = (x*isig)^8 - 28*(x*isig)^6 + 210*(x*isig)^4 - 420*(
    x*isig)^2 + 105
He = [He0, He2, He4, He6, He8]
P = sum([He[k]*isig^(2*k) / (2^(2*k) * factorial(2*k+1))
    for k in range(1,5)])
T = sum([binomial(a,j) * P^j for j in range(5)])
T = T.collect(x)

for i in range(16):
    # Checking odd-degree monomials in x are 0
    assert T.coefficient(x, 2*i+1) == 0

C = a*isig^2 - (a-1)*invs^2
res = T.coefficient(x,0).collect(isig)
for l in range(1, 17):
    tmp = (2*l - 1).multifactorial(2) / C^l
    res += T.coefficient(x,2*l).collect(isig) * tmp
res *= invs / sqrt(C)
if low_a:
    res *= (1/(1 - invs^2/24))^a
else: # Manual expansion of (1 + F_s)^a
    F_s = (1/(1 - invs^2/24)).taylor(invs,0,9) - 1
    res *= sum([binomial(a,i) * F_s^i for i in range(5)])

RD = res.subs(isig=1/(1/invs - invs/24)).taylor(invs,0,9)
var("s")
print("RD_a^(a-1) =", RD.subs(invs=1/s), "+ o(1/s^9)")
```

Fig. 3.1. Sage script for the symbolic computation of the Taylor expansion of $\frac{1}{\sqrt{C}s}\left(\frac{s}{\sigma(s)}\right)^a \sum_{\ell=0}^{16} T_{2\ell}(\sigma(s))\frac{(2\ell-1)!!}{C^\ell}$. The variables `isig` and `invs` respectively represent $1/\sigma(s)$ and $1/s$.

Composing with the $1/(a-1)$-th power, and taking the product over all $i \in [n]$ each with $c = c_i$, we then deduce by Lemma 3.1 that

$$\mathrm{RD}_a(\mathcal{R}_{\mathbb{Z}^n,\sigma(s),\mathbf{c}} \| \mathcal{D}_{\mathbb{Z}^n,s,\mathbf{c}}) = 1 + \frac{n\gamma_a}{(a-1)s^8} + o\left(\frac{1}{s^9}\right) = 1 + \frac{73an}{33177600} \cdot \frac{1}{s^8} + o\left(\frac{1}{s^9}\right),$$

which concludes the proof of the main theorem. □

4 A New Discrete Gaussian Integer Sampler

Our result can be interpreted as giving a new Gaussian sampler over $\mathbb{Z}$ (or $\mathbb{Z}^n$) that produces a distribution close to a discrete Gaussian. In particular, we can build upon the sampler of [HLS18] with a simple additional step that computes the Sheppard correction. For completeness, we give the sampler's description in Algorithm 4.1, which we generalize from [HLS18] to also account for arbitrary centers. As the Box-Muller component for continuous Gaussian sampling is exactly the same, we can readily reach constant-time behavior by using the constant-time rounded Gaussian sampler of [BOKR20]. One only needs to protect the first step, which is to compute

$$\sigma(s) = s - \frac{1}{24s},$$

on-the-fly. Although this is fairly classic, we detail a method to do so in Section 4.1.

Algorithm 4.1: SampleZn$(n, s, \mathbf{c})$

Input: Dimension n, Gaussian parameter $s > 1/\sqrt{24}$, center vector $\mathbf{c} \in \mathbb{R}^n$

1. $\sigma \leftarrow s - 1/(24s)$ | Sheppard's correction
2. **For** $i = 0, \ldots, \lceil n/2 \rceil$ **do**
3. $\quad u_1, u_2 \hookleftarrow U((0,1))$
4. $\quad a \leftarrow \sqrt{-2\ln(u_1)}$ | Box-Muller
5. $\quad b \leftarrow 2\pi u_2$ | Sampling
6. $\quad (x_1, x_2) \leftarrow (a\cos(b), a\sin(b))$
7. $([\mathbf{z}]_{2i}, [\mathbf{z}]_{2i+1}) \leftarrow (\lfloor \sigma x_1 + [\mathbf{c}]_{2i} \rceil, \lfloor \sigma x_1 + [\mathbf{c}]_{2i+1} \rceil)$

Output: $\mathbf{z} \in \mathbb{Z}^n$ (following $\mathcal{R}_{\mathbb{Z}^n, s-1/(24s), \mathbf{c}}$)

4.1 Sheppard's Correction Without Inverses

The implementation and precision analysis of [HLS18] can then be used as is, meaning we only need to detail the first step of SampleZn, i.e., computing $\sigma(s) = s - 1/(24s)$ efficiently and in constant-time. We can avoid the direct division over floating-point numbers (as the `divsd` instruction is not constant-time on modern CPU architectures) by using convergence methods to only rely on additions, substractions and multiplications. Inversion can indeed be done using a few iterations of Newton-Raphson's method for division [OF97]. Concretely, to compute $1/y$ (for $y = 24s$), one would first compute $d \in [0.5, 1)$ and $k \in \mathbb{Z}$ such that $y = d2^k$, and then compute $1/y = 2^{-k} \cdot 1/d$, reducing inversion to a known interval. The latter is done by starting with an initial guess x_0 and iteratively computing $x_{i+1} = x_i - f(x_i)/f'(x_i)$, where $f(x) = 1/x - d$. The update rules becomes $x_{i+1} = x_i(2 - x_i d)$.

The relative error $\varepsilon_i = |1/d - x_i|/(1/d) = |1 - dx_i|$ verifies $\varepsilon_{i+1} = \varepsilon_i^2$, which means convergence is quadratic and independent of the value d to invert. It now remains to choose an initial guess x_0. For that, we simply choose the linear approximation $x_0 = c_0 + c_1 d$ that minimizes the relative error $\varepsilon_0 = \left|1 - c_0 d - c_1 d^2\right|$.

Using Chebyshev's equioscillation theorem, the optimal constants are $c_0 = 48/17$ and $c_1 = -32/17$. These constants are not new and have been used to seed Newton-Raphson's method for division in many software and hardware implementations, as they provide faster convergence regardless of the input. In particular, they ensure $\varepsilon_0 \leq 1/17$ and therefore that $\varepsilon_4 \leq 1/17^{16} \approx 2^{-65.4}$, thus ensuring exact double precision computation. We formalize the computation in Algorithm 4.2.

Algorithm 4.2: Sheppard(s)

1. $y \leftarrow 24s$
2. Compute $(d, k) \in [1/2, 1) \times \mathbb{Z}$ s.t. $y = d2^k$ | Done with bit operations
3. $x_0 \leftarrow c_0 + c_1 d$
4. **For** $i = 0, \ldots, 3$ **do** $x_{i+1} \leftarrow x_i(2 - dx_i)$
5. $\sigma \leftarrow s - x_4 2^{-k}$

Output: $\sigma \approx s - 1/(24s)$

Constant-Time Sheppard's Correction. Following the "binary64" format specified in the IEEE-754 standard for floating-point arithmetic [IEE19], and as per the analysis of [ANB+18], additions, substractions and multiplications are fully constant-time on recent x86 CPUs, except when the operands or result are subnormal numbers (i.e., numbers strictly below 2^{-1022}). We show below that, for $s \in [1/\sqrt{24} + 2^{-1022}, 2^{1017}]$, subnormal numbers never arise in the computation of $\sigma(s)$ described above. A similar, yet much more involved, analysis is typically required for Falcon [Por19] (also excluding `NaN`s and `Inf`s) to ensure floating-point operations performed by the hardware Floating-Point Unit are indeed constant-time.

First, in Sheppard.1, we compute $y = 24s$. Because we look at $s > 1/\sqrt{24}$, it naturally holds that neither s nor y is subnormal. As Sheppard.2 is done with bit operations on the representation of y, we are not doing floating-point operations. We can thus focus on the Newton-Raphson method and show all the intermediate values are not subnormal.

For the initial guess in Sheppard.3, it holds that c_0, c_1, d are not subnormal. Then, we have $c_1 d \in (-32/17, -16/17]$ and $c_0 + c_1 d \in (16/17, 32/17]$, which are out of the subnormal range. Next, let $i \in [0, 3]$. By definition of ε_i, we have

$$\begin{aligned} dx_i &\in [1 - \varepsilon_i, 1 + \varepsilon_i] \\ 2 - dx_i &\in [1 - \varepsilon_i, 1 + \varepsilon_i] \\ x_i(2 - dx_i) = x_{i+1} &\in [d^{-1}(1 - \varepsilon_{i+1}), d^{-1}(1 + \varepsilon_{i+1})] \end{aligned}$$

Recalling that $\varepsilon_i \leq 1/17^{2^i}$ yields that all the dx_i and $2 - dx_i$ in Sheppard.4 actually lie in $[1 - \varepsilon_0, 1 + \varepsilon_0] \subseteq [16/17, 18/17]$ and are then normal values. Because $1/d \in [1, 2]$, we have that all the x_{i+1} belong to $[1 - \varepsilon_1, 2(1 + \varepsilon_1)] \subseteq [288/289, 580/289]$ and are thus normal values as well.

It then remains to analyze the last two operations in Sheppard.5, which may yield subnormal numbers for extreme values of s. Scaling back by 2^{-k} may indeed cause the result to underflow if s was too large to begin with. As $x_4 2^{-k} \in 2^{-k} d^{-1}[1 - \varepsilon_4, 1 + \varepsilon_4] = 1/(24s) \cdot [1 - \varepsilon_4, 1 + \varepsilon_4]$, requiring $(1 - \varepsilon_4)/24s \geq 2^{-1022}$

ensures this stays out of the subnormal range. This translates to $s \leq 2^{1019} \cdot \frac{1-\varepsilon_4}{3} \approx 2^{1017.42}$. Then, the final substraction can cause issues if s is too close to $1/\sqrt{24}$. To avoid this regime, we need $s - 1/(24s) \geq 2^{-1022}$ which results in $s \geq 2^{-1023} + (1 + 2^{-2041})^{1/2}/\sqrt{24}$. The right-hand side can be bounded by $2^{-1023} + (1+2^{-2042})/\sqrt{24} \leq 2^{-1022} + 1/\sqrt{24}$. Combined with the above, the *safe* range for s is $[1/\sqrt{24} + 2^{-1022}, 2^{1017.42}]$. As typical applications do not require such extreme values of s, this shows that the computation of $\sigma(s)$ can be done in constant-time.

4.2 Security Arguments for Rounded Gaussians

The result of Theorem 3.1 shows that one can replace the ideal rounded Gaussian of parameter $\sigma(s)$ by an ideal discrete Gaussian of parameter s, at the expense of a Rényi divergence loss that is very close to 1. It means a scheme designed and analyzed with discrete Gaussians can be implemented by using SampleZn instead. For a search problem, like the unforgeability of a signature scheme, one can easily obtain a condition on the minimal parameter s to maintain an unnoticeable security loss.

One can indeed follow the same reasoning as [Pre17, Sec. 3.3]. Consider a scheme that has been analyzed with ideal discrete Gaussian distributions. We assume the scheme makes at most Q queries to ideal discrete Gaussians $\mathcal{D}_{\mathbb{Z}^n,s_i,\mathbf{c}_i}$ for $i \in [Q]$, and we let $\mathrm{Adv}_{\mathcal{D}}[\mathcal{A}]$ be the advantage of an adversary in breaking the search security property. We then consider the same scheme but where the discrete Gaussians are replaced by ideal rounded Gaussians $\mathcal{R}_{\mathbb{Z}^n,\sigma_i,\mathbf{c}_i}$ with $\sigma_i = \sigma(s_i) = s_i - 1/(24s_i)$. Denoting by $\mathrm{Adv}_{\mathcal{R}}[\mathcal{A}]$ the advantage of $\mathcal{A}$ for this new scheme, the probability preservation and multiplicativity properties of the Rényi divergence from Lemma 2.1 yield

$$\mathrm{Adv}_{\mathcal{R}}[\mathcal{A}] \leq \mathrm{Adv}_{\mathcal{D}}[\mathcal{A}]^{\frac{a-1}{a}} \cdot \prod_{i\in[Q]} \mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z}^n,\sigma_i,\mathbf{c}_i} \| \mathcal{D}_{\mathbb{Z}^n,s_i,\mathbf{c}_i}\right)^{\frac{a-1}{a}}$$

Taking $s = \min s_i$, this roughly gives a loss factor bounded by $1+\gamma(a-1)nQ/s^8+QR$, where $\gamma = 73/33177600$ and R is the $o(s^{-9})$ in Theorem 3.1. Then, if $s \geq (\gamma(a-1)nQ/\alpha)^{1/8}$ for a small constant α, using rounded Gaussians only yields a constant security loss factor of $1+\alpha$. Following the suggestion of [Pre17] of $a = 2\lambda$, with example parameters $\lambda = 128$, $n = 1024$, $Q = 2^{64}$, and $\alpha = 1$, the minimal condition becomes $s \gtrsim 238.87$, while yielding a loss of only 1 bit in security[3]. Many applications actually use Gaussians with much larger standard deviations, with possibly much fewer queries, making the security loss unnoticeable.

4.3 Experimental Estimation Beyond Order 8

Our main result gives an asymptotic estimation of the remainder past order 8 in $1/s$, and not an inequality uniformly in $1/s$. We must then ensure that

[3] The exponent $1 - 1/a$ in $\mathrm{Adv}_{\mathcal{D}}[\mathcal{A}]$ may also yield a small security loss, which is bounded by half a bit for $a = 2\lambda$ [Pre17]

this term has no noticeable impact on the security argument above. For that, we can empirically verify how close this asymptotic estimate $1 + \gamma an/s^8$ is to the actual Rényi divergence, for concrete values of s. As the term actually comes from the one-dimensional case, we can compute the Rényi divergence in high precision (for the tailcutted Gaussians, with a negligible tailcut probability). Even for small values of s, we observe the expansion is very close to the actual value. Figure 4.1 shows how the (log-)expansion error $\mathsf{logerr}_a(s) = \log_2(\left|\mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z},\sigma(s)} \| \mathcal{D}_{\mathbb{Z},s}\right) - 1 - \frac{73a}{33177600s^8}\right|)$ rapidly decreases with s for different values of a. For $s = 1$, we already have $\mathsf{logerr}_{256}(s) \leq 2^{-18}$. More interestingly, we observe that the scaled expansion error defined by

$$\mathsf{scerr}_a(s) = \frac{\left|\mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z},\sigma(s)} \| \mathcal{D}_{\mathbb{Z},s}\right) - 1 - \frac{73a}{33177600s^8}\right|}{\frac{73a}{33177600s^8}},$$

decreases as s grows, as expected, but is actually always bounded by $1/s^2$ for typical values of a (Figure 4.2). It essentially means

$$\mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z},\sigma(s)} \| \mathcal{D}_{\mathbb{Z},s}\right) \leq 1 + \frac{73a}{33177600s^8}\left(1 + \frac{1}{s^2}\right).$$

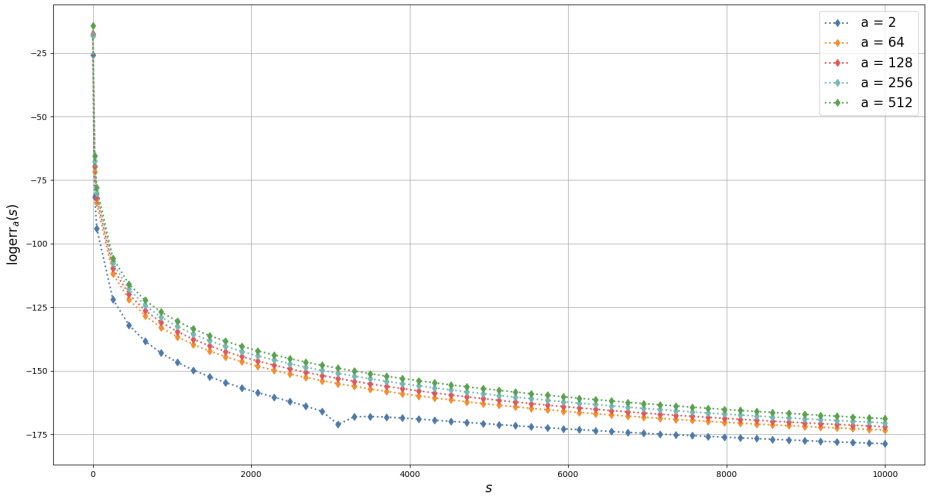

Fig. 4.1. Log expansion error $\mathsf{logerr}_a(s) = \log_2(\left|\mathrm{RD}_a\left(\mathcal{R}_{\mathbb{Z},\sigma(s)} \| \mathcal{D}_{\mathbb{Z},s}\right) - 1 - \frac{73a}{33177600s^8}\right|)$ for different values of a.

This is then sufficient to argue that the expansion error does not have a noticeable impact on the security argument sketched in Section 4.2, compared to the expansion $1 + \gamma a/s^8$ itself.

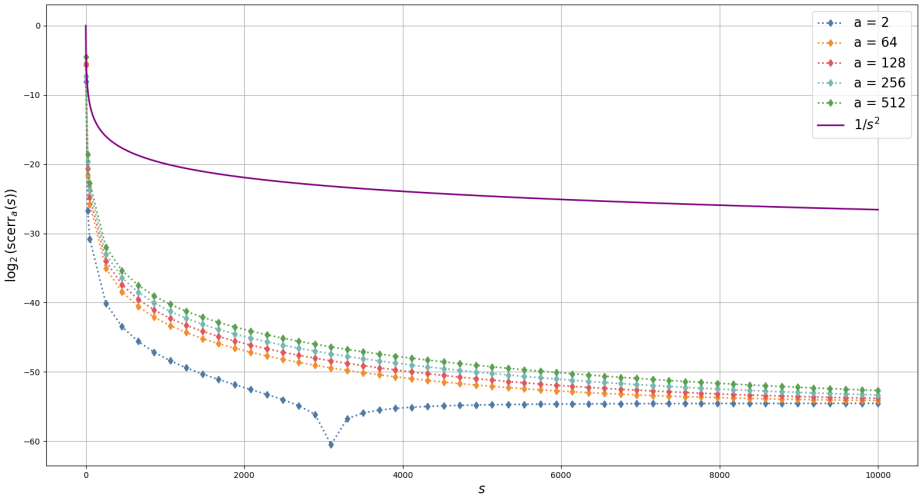

Fig. 4.2. Log scaled expansion error $\log_2 \mathsf{scerr}_a(s)$ for different values of a.

5 Conclusion

Despite the very interesting properties of rounded Gaussian samplers in practice [HLS18,BOKR20], their use was discarded in lattice cryptographic applications due to the complexity of the distribution itself, which is not particularly suited to security analyses. Our work shows that a small correction of rounded Gaussians makes them extremely close to much friendlier discrete Gaussian distributions. Not only does it better our theoretical understanding of the rounded Gaussian distribution, but it also has a direct practical application: any rounded Gaussian sampler can be transformed into a discrete Gaussian sampler, with a distribution divergence that rapidly gets closer to 1 when the target Gaussian parameter s grows. The only additional step lies in adjusting the underlying continuous Gaussian parameter to $\sigma = s - 1/(24s)$, which can be computed on-the-fly, efficiently, and in constant-time. This yields a discrete Gaussian sampler that is memory-less, rejection-free, generic, handles (and actually gets closer to the ideal distribution with) large deviations, and that is as fast and secure as any rounded Gaussian sampler. The main limitation however lies in the regimes where s is extremely small, in which case the rounded Gaussian visibly differs from a discrete one. These cases have fortunately been extensively studied, with already very efficient samplers, e.g., [HPRR20].

Acknowledgments

We thank Thomas Prest for helpful discussions on rounded Gaussians and discrete Gaussian samplers.

References

AKY24. S. Agrawal, S. Kumari, and S. Yamada. Attribute Based Encryption for Turing Machines from Lattices. In *CRYPTO*, 2024.

AMAR17. C. Aguilar-Melchor, M. R. Albrecht, and T. Ricosset. Sampling from Arbitrary Centered Discrete Gaussians for Lattice-Based Cryptography. In *ACNS*, 2017.

AMR18. C. Aguilar-Melchor and T. Ricosset. CDT-Based Gaussian Sampling: From Multi to Double Precision. *IEEE Trans. Computers*, 2018.

ANB+18. M. Andrysco, A. Nötzli, F. Brown, R. Jhala, and D. Stefan. Towards Verified, Constant-time Floating Point Operations. In *CCS*, 2018.

BCG+13. J. Buchmann, D. Cabarcas, F. Göpfert, A. Hülsing, and P. Weiden. Discrete Ziggurat: A Time-Memory Trade-Off for Sampling from a Gaussian Distribution over the Integers. In *SAC*, 2013.

BLNS23. J. Bootle, V. Lyubashevsky, N. K. Nguyen, and A. Sorniotti. A Framework for Practical Anonymous Credentials from Lattices. In *CRYPTO*, 2023.

BM58. G. E. P. Box and M. E. Muller. A Note on the Generation of Random Normal Deviates. *The Annals of Mathematical Statistics*, 1958.

BOKR20. S. Brannigan, M. O'Neill, A. Khalid, and C. Rafferty. A Secure Algorithm for Rounded Gaussian Sampling. In *CANS*, 2020.

DB15. C. Du and G. Bai. Towards Efficient Discrete Gaussian Sampling for Lattice-Based Cryptography. In *FPL*, 2015.

DDLL13. L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice Signatures and Bimodal Gaussians. In *CRYPTO*, 2013.

DFJS26. P. Delhom, P.-A. Fouque, C. Jeudy, and O. Sanders. Lattice Group Signatures, Revisited. In *ASIACRYPT*, 2026.

DG14. N. C. Dwarakanath and S. D. Galbraith. Sampling from Discrete Gaussians for Lattice-Based Cryptography on a Constrained Device. *Appl. Algebra Eng. Commun. Comput.*, 2014.

DLP14. L. Ducas, V. Lyubashevsky, and T. Prest. Efficient Identity-Based Encryption over NTRU Lattices. In *ASIACRYPT*, 2014.

DN12. L. Ducas and P. Q. Nguyen. Faster Gaussian Lattice Sampling Using Lazy Floating-Point Arithmetic. In *ASIACRYPT*, 2012.

GFH26. Y. Gao, Y. Feng, and H. Hu. Discrete Gaussian Sampling Meets BDGL Decoding: Solving the Shortest Vector Problem in $2^{0.5595n+o(n)}$ Time. *IACR Cryptol. ePrint Arch.*, page 1844, 2026.

GPV08. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for Hard Lattices and New Cryptographic Constructions. In *STOC*, 2008.

HLS18. A. Hülsing, T. Lange, and K. Smeets. Rounded Gaussians - Fast and Secure Constant-Time Sampling for Lattice-Based Crypto. In *PKC*, 2018.

HPRR20. J. Howe, T. Prest, T. Ricosset, and M. Rossi. Isochronous Gaussian Sampling: From Inception to Implementation. In *PQCrypto*, 2020.

IEE19. IEEE. 754-2019 - IEEE Standard for Floating-Point Arithmetic, 2019.

Jeu26. C. Jeudy. Practical Group Signatures from Tag-Based NTRU Sampler. *IACR Cryptol. ePrint Arch.*, page 2077, 2026.

JS26a. C. Jeudy and O. Sanders. Compact Lattice Anonymous Credentials from Tighter Approximate Range Proofs. *IACR Cryptol. ePrint Arch.*, page 1920, 2026.

JS26b. C. Jeudy and O. Sanders. Tag-Friendly Lattice Sampler and Applications. In *ASIACRYPT*, 2026.

KAA22. E. Karabulut, E. Alkim, and A. Aysu. Efficient, Flexible, and Constant-Time Gaussian Sampling Hardware for Lattice Cryptography. *IEEE Trans. Computers*, 2022.

Kar16. C. F. F. Karney. Sampling Exactly from the Normal Distribution. *ACM Trans. Math. Softw.*, 42(1):3:1–3:14, 2016.

KHTA26. H. Kaihara, C. A. Haidar, M. Tibouchi, and M. Abe. Square Root of All Evil: The Dangers of Falcon's Superfluous Square Roots. In *CCS*, 2026.

LNP22. V. Lyubashevsky, N. K. Nguyen, and M. Plançon. Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General. *CRYPTO*, 2022.

LNPS21. V. Lyubashevsky, N. K. Nguyen, M. Plançon, and G. Seiler. Shorter Lattice-Based Group Signatures via "Almost Free" Encryption and Other Optimizations. In *ASIACRYPT*, 2021.

MP12. D. Micciancio and C. Peikert. Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In *EUROCRYPT*, 2012.

MW17. D. Micciancio and M. Walter. Gaussian Sampling over the Integers: Efficient, Generic, Constant-Time. In *CRYPTO*, 2017.

NIS24a. NIST. Module-Lattice-Based Digital Signature Standard. *Federal Information Processing Standards Publication, NIST FIPS 204*, 2024.

NIS24b. NIST. Module-Lattice-Based Key-Encapsulation Mechanism Standard. *Federal Information Processing Standards Publication, NIST FIPS 203*, 2024.

OF97. S. F. Oberman and M. J. Flynn. Division Algorithms and Implementations. *IEEE Transactions on Computers*, 1997.

Pei10. C. Peikert. An Efficient and Parallel Gaussian Sampler for Lattices. In *CRYPTO*, 2010.

PFH+20. T. Prest, P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. *FALCON. Tech. rep.*, 2020. Available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

Por19. Thomas Pornin. New Efficient, Constant-Time Implementations of Falcon. *IACR Cryptol. ePrint Arch.*, page 893, 2019.

Pre17. T. Prest. Sharper Bounds in Lattice-Based Cryptography Using the Rényi Divergence. In *ASIACRYPT*, 2017.

Ré61. A. Rényi. On Measures of Entropy and Information. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. I*, 1961.

Reg05. O. Regev. On Lattices, Learning With Errors, Random Linear Codes, and Cryptography. In *STOC*, 2005.

RVV13. S. Sinha Roy, F. Vercauteren, and I. Verbauwhede. High Precision Discrete Gaussian Sampling on FPGAs. In *SAC*, 2013.

She97. W. F. Sheppard. On the Calculation of the most Probable Values of Frequency-Constants, for Data arranged according to Equidistant Division of a Scale. *Proceedings of The London Mathematical Society*, pages 353–380, 1897.

The25. The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 10.7)*, 2025. https://www.sagemath.org.

vEH14. T. van Erven and P. Harremoës. Rényi Divergence and Kullback-Leibler Divergence. *IEEE Trans. Inf. Theory*, 2014.

ZSS20. R. K. Zhao, R. Steinfeld, and A. Sakzad. COSAC: COmpact and Scalable Arbitrary-Centered Discrete Gaussian Sampling over Integers. In *PQCrypto*, 2020.