How to Sample a Discrete Gaussian (and more) from a Random Oracle

George Lu^{*} and Brent Waters^{**}

University of Texas at Austin, NTT Research

Abstract. The random oracle methodology is central to the design of many practical cryptosystems. A common challenge faced in several systems is the need to have a random oracle that outputs from a structured distribution \mathcal{D} , even though most heuristic implementations such as SHA-3 are best suited for outputting bitstrings.

Our work explores the problem of sampling from discrete Gaussian (and related) distributions in a manner that they can be programmed into random oracles. We make the following contributions:

- We provide a definitional framework for our results. We say that a sampling algorithm Sample for a distribution is explainable if there exists an algorithm Explain which, when given an x in the support of \mathcal{D} , outputs an $r \in \{0,1\}^n$ such that Sample(r) = x. Moreover, if x is sampled from \mathcal{D} the explained distribution is statistically close to choosing r uniformly at random. We consider a variant of this definition that allows the statistical closeness to be a "precision parameter" given to the Explain algorithm. We show that sampling algorithms which satisfy our 'explainability' property can be programmed as a random oracle.
- We provide a simple algorithm for explaining *any* sampling algorithm that works over distributions with polynomial sized ranges. This includes discrete Gaussians with small standard deviations.
- We show how to transform a (not necessarily explainable) sampling algorithm Sample for a distribution into a new Sample' that is explainable. The requirements for doing this is that (1) the probability density function is efficiently computable (2) it is possible to efficiently uniformly sample from all elements that have a probability density above a given threshold p, showing the equivalence of random oracles to these distributions and random oracles to uniform bitstrings. This includes a large class of distributions, including all discrete Gaussians.
- A potential drawback of the previous approach is that the transformation requires an additional computation of the density function.
 We provide a more customized approach that shows the Miccancio-Walter discrete Gaussian sampler is explainable as is. This suggests that other discrete Gaussian samplers in a similar vein might also be explainable as is.

^{*} University of Texas at Austin. Email: gclu@cs.utexas.edu.

^{**} University of Texas at Austin and NTT Research. Email: bwaters@cs.utexas.edu. Supported by NSF CNS-1908611, CNS-1414082, Packard Foundation Fellowship, and Simons Investigator Award.

1 Introduction

The random oracle methodology proposed by Bellare and Rogaway [5] allows one to develop a cryptosystem under the premise that all users have access to an oracle that outputs a random string for each queried input. In practice, when deploying said systems, the calls to the oracle are heuristically replaced with calls to an appropriate hash function, such as SHA-3. While the use of such a heuristic comes with some controversy [13], the methodology has been leveraged for a broad spectrum of problems such as chosen-ciphertext security [17] and non-interactive zero knowledge proofs [16], to name just a few. In addition, it has been key to the development of many practical and deployed cryptosystems.

Although the heuristic of replacing a random oracle with a hash function such as SHA-3 is naturally aligned with oracles that output random bitstrings, there are many examples of cryptosystems that require random oracles to output from other distributions. For instance, the seminal identity-based encryption (IBE) scheme of Boneh-Franklin [7] uses a random oracle that outputs a bilinear group element, as does the Boneh-Lynn-Shacham signature scheme [8] and the multi-authority Attribute-Based Encryption (ABE) systems of [26]. Other examples include the GPV IBE scheme [19], which needs a random oracle to output a vector over \mathbb{Z}_p for some prime p as well as RSA-based full domain hash signatures [6], which need an element over \mathbb{Z}_N for a composite N. Other works explored constructing random oracles which hash into elliptic curves [20, 23].

If one delves deeper into the deployment of such cryptosystems, we can see that there is no specialized hash function for each of these different domains. Instead, to create a random oracle scheme for a certain distribution \mathcal{D} (e.g. random elements over a particular bilinear group), one will utilize a sampling function Sample. The function Sample will take in a string r and outputs something in the desired domain. The distribution of calling Sample on a random string should be statistically close to that of a given distribution \mathcal{D} .

While achieving statistical (or computational) closeness to a given distribution is a necessary property of a sampling function, it is not sufficient, as the sampling function may not allow for the "programmability" of a random oracle necessary in a security proof. For example, suppose we ran the BLS signature scheme over a bilinear group \mathbb{G} of prime order p with generator g, public key g^a and secret key $a \in \mathbb{Z}_p$. In the scheme, a signature on a message m is created as $H(m)^a$ where $H(\cdot)$ is an oracle function that outputs a bilinear group element. Suppose we implement H by employing a random oracle H' that outputs bitstrings alongside a Sample algorithm that computes q^r (interpreting r as an integer) so that $H(m) = q^{H'(m)}$. Such an instantiation will indeed output elements statistically close to random bilinear group elements so long as r is sufficiently long. However, this results in a completely broken cryptosystem. To see this, observe that if an attacker can obtain a signature σ on message m, the attacker can then create a signature $\tilde{\sigma}$ on any other message \tilde{m} by computing $\tilde{\sigma} = \sigma^{H'(\tilde{m})/H'(m)}$. Similar counterexamples exist for the other cryptosystems mentioned.

The more general goal of defining sufficient conditions of replacing cryptographic functionalities with each other has been explored in the line of work on indifferentiability initiated in [29], and further expanded on in papers such as [32]. Indeed, the application of the indifferentiability framework with regard to sampling from particular elliptic curve groups was explored in [11].

We define an important property for any sampler to also have a property we call explainability. This can be viewed as a relaxation of indifferentiability specific to sampling functionalities. Roughly, given an element x in the domain, it should be possible to efficiently "reverse sample" an r such that $\mathsf{Sample}(r) = x$. Moreover, the distribution of receiving an x from \mathcal{D} and then outputting r from reverse sampling should be close to just choosing r uniformly at random. Prior works dealt with this issue with various levels of formality. For many of the aforementioned works, 'ad-hoc' workarounds often invoking specific cryptographic assumptions are used to obtain a proof from the plain random oracle model. In bilinear groups, one often calls this a "hash-to-point" function which is present for many bilinear groups, but not necessarily guaranteed. In general there can exist distributions where sampling cannot be explained; consider if the function Sample were a one way function.

Sampling and Explaining Discrete Gaussian Distributions

In this work we explore the problem of sampling and explaining discrete Gaussian distributions. And consequentially, the problem of programming discrete Gaussian distributions into random oracles. Discrete Gaussian distributions are heavily utilized in the design and analysis of lattice-based cryptosystems, a flour-ishing area of research over the last several years. The problem of sampling discrete Gaussians has been well studied. While it is possible to sample such distributions using a very basic form of rejection sampling [19], further works have both improved on the efficiency of the rejection sampling method [10, 12, 25], as well as explore other techniques to sample from discrete Gaussians [15, 31, 2, 30, 18, 24, 34, 4, 22, 33] such as computing the cumulative density function, or taking convolutions of smaller standard deviation discrete Gaussians. The goal of most of these works focus on making these samplers more secure and usable, allowing for features such as being constant time, providing better memory-time tradeoffs, or supporting a greater deal of offline precomputation.

To the best of our knowledge, however, the problem of explaining and programming random oracles with discrete Gaussian distributions has received little attention to date, with a few exceptions. One exception is the universal sampler work of Hofheinz et al. [21] that implicitly shows how indistinguishability obfuscation can be used to obtain a computational form of explainability for any efficiently sampleable distribution from a random oracle. However, all current indistinguishability obfuscation candidates are highly impractical and at best invoke further computational assumptions that go beyond those typically used in lattice-based cryptosytems. Our solutions will be both statistical and significantly more efficient. In a more recent work [1], Agrawal, Wichs and Yamada sketch how the rejection sampling algorithm of Gentry-Peikert and Vaikuntanathan [19] is explainable. Interestingly, the problem of explaining such distributions has come up in multiple contexts. Brakerski, Cash, Tsabary, and Wee [9] gave a homomorphic ABE scheme provably secure from the LWE assumption in the random oracle model. At one point, their construction required a random oracle that outputs a discrete Gaussian. Since no such solution were available, the authors worked around this by using a specialized sampler due to Lyubashevsky and Wichs [28] which is a blend of a standard discrete Gaussian and a binary string. Another example is in a recent multi-authority scheme of Datta, Komargodski and Waters [14] where the need for a random oracle that outputs discrete Gaussians arises. In this case the authors compensate by using a random oracle that outputs an integer over a subexponentially large integer range that can hide a smaller discrete Gaussian by smudging [3]. In this case the workaround resulted in a subexponentially large modulus. (It should be noted that smudging was used elsewhere in their analysis as well).

Finally, we want to emphasize that the need to explainably sample distributions also arises outside of the random oracle model. An interesting example comes up in the aforementioned work of [1], which requires public parameters that can generate discrete Gaussians, but should look like uniform bitstrings. To prove security the authors require the distribution to be explainable.

We advocate for this importance of studying the problem of explaining and programming discrete Gaussian distributions. Ideally, such solutions will match the performance of the prior works on discrete Gaussian sampling (e.g. [15, 31, 2, 30, 18, 24, 34, 4, 22, 33]) that were focused on performance, but not explainability. Pursuing this goal is a natural and fundamental property given the important role of discrete Gaussians in lattice-based cryptography.

1.1 Our Contributions

Our work consists of the following contributions.

Definitional:

We begin by providing a definitional framework for describing our results. We define an explainable sampling system for distribution $\{\mathcal{D}_{\lambda}\}_{\lambda}$ to have two efficiently computable algorithms. The first is a $\mathsf{Sample}(1^{\lambda}; r)$ algorithm which is parameterized by a security parameter and takes in random coins r. This algorithm should output a distribution statistically close (in λ) to \mathcal{D}_{λ} when r is chosen randomly. The second algorithm is a randomized algorithm $\mathsf{Explain}(1^{\lambda}, 1^{\kappa}, x)$ that takes as input the security parameter, a "precision" parameters κ and an element $x \in \mathcal{D}_{\lambda}$. Its job is to output an r such that $\mathsf{Sample}(1^{\lambda}; r) = x$. In addition, calling $\mathsf{Explain}$ on an x sampled from \mathcal{D} should have a statistically close distribution to that of simply choosing r at random.

A particular feature of our definition is the use of a tunable precision parameter κ where we only require a statistical distance of $\frac{1}{\kappa}$, as opposed to requiring the statistical distance (or computational advantage) to be negligibly close like in indifferentiability. We show that this is sufficient to allow for proofs to go through in programming a random oracle. Intuitively, the parameter κ will in a reduction will be tuned to an attacker's advantage. We note that the parameter κ is only used in the Explain algorithm and is not a priori set in the Sample algorithm used in a construction. Thus it can be adjusted to fit a particular attacker with a particular advantage in a reduction. This relaxation allows us to prove explainability, and hence show equivalence of random oracles, to a broader range of distributions than otherwise.

We note that our use of a tunable precision parameter is the main definitional difference between our framework and other works that explored reverse sampling.

Sampling over Small Ranges:

We next show a simple algorithm for explainable sampling over small ranges. Suppose that \mathbb{R}_{λ} is the range of distribution \mathcal{D}_{λ} where $|\mathbb{R}_{\lambda}|$ grows polynomially in λ . We show that *any* sampling algorithm Sample for $\{\mathcal{D}_{\lambda}\}_{\lambda}$ is explainable. One simply calls Sample repeatedly until it outputs a desired x or until $|\mathbb{R}_{\lambda}| \cdot \kappa$ attempts occur without success. This simple process illustrates the flexibility given by the precision parameter in our framework.

A basic corollary extends this lemma to any distribution which is statistically close to a distribution with polynomial support. This implies sampling for discrete Gaussians with poly-sized standard deviations.

Generic Sampling over Conforming Distributions:

A drawback of the previous approach is that it is only applicable when the distribution range is small, and thus cannot be used to explain, say, a discrete Gaussian with a super-polynomial sized standard deviation. We show an approach to generically sample from a broad class of distributions that includes discrete Gaussians with large standard deviations. Unlike the prior solution we will not be able to use a Sample algorithm as is and explain it. Instead we will transform it into a new Sample' algorithm that is explainable. For our transformation we will require a distribution with:

- 1. A (not necessarily explainable) Sample algorithm for the distribution.
- 2. A probability density function PDens where $\mathsf{PDens}(1^{\lambda}, x)$ returns a value proportional to the probability x occurs for distribution \mathcal{D}_{λ} .
- 3. A "heavy element sampler" SampleUniform where SampleUniform $(1^{\lambda}, p; r)$ is an explainable sampling algorithm that samples uniformly from all elements in the range of \mathcal{D}_{λ} that have a probability density above p.

Given these we can build an explainable Sample' algorithm that operates by first sampling x' from Sample. Then computing the probability density p' of x'. Next, it randomly "scales" p' by choosing a random $s_0 \in [0, 1]$ and letting $p = s_0 p'$. Finally, it outputs $x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p; r)$. We go on to show that this distribution is explainable.

We additionally give an explicit heavy element sampler for the case of discrete Gaussians, and in general observe that this primitive can be computed from the probability density function for many natural distributions. This gives us an explainable sampler for all discrete Gaussians with exponentially bounded centers and standard deviations. Finally, we show that something akin to the heavy element sampler is needed for a truly generic transformation. We describe oracles which describe a distribution which has the first two properties above, but not the third and show that it is impossible to create an explainable sampler.

Explaining the Miccancio-Walter '17[30] sampling algorithm:

While the previous technique can create a sampler for discrete Gaussians with large standard deviations, it requires creating a new sampler rather than using one as is. If such a sampler is in a critical path, the additional introduced overhead of a high precision computation of the probability distribution function may be undesirable.

As our final contribution, we provide a tailored explain algorithm to the MW sampler. With any sampler, proving its explainability is vital to enabling its use in random oracle based applications, and allows us to securely instantiate such cryptosystems while carrying over the performance benefits of the sampler in question. While we focus on the MW sampler in particular, we believe that the ideas we demonstrate can extend to similar works [31] and provide techniques for showing explainability of other classes of discrete Gaussian samplers.

2 Preliminaries

We say a function $\operatorname{negl}(x)$ is negligible if for all polynomials p(x), there exists some N such that $\forall x > N$, $\operatorname{negl}(x) < \frac{1}{p(x)}$. A function is noticeable if it is not negligible. The notation $\operatorname{poly}(x)$ will be used to refer to a polynomial function in x, and $\operatorname{EXP}(x)$ will refer to a function $\leq 2^{\operatorname{poly}(x)}$. We will indicate sampling an element from a probability distribution \mathcal{D} as $x \stackrel{R}{\leftarrow} \mathcal{D}$. Similarly, we will use $x \stackrel{R}{\leftarrow} S$ to indicate sampling over the uniform distribution on a set S. We will use [a, b] and (a, b) to denote the closed and open interval from a to b respectively on \mathbb{R} . We will subscript the brackets with \mathbb{Z} to denote the same interval on \mathbb{Z} . We will use range(\mathcal{D}) to refer to the set of elements in \mathcal{D} which occur with probability > 0.

We will use log as logarithms base 2 if no base is explicitly specified. We use $\lfloor x \rfloor$, $\lceil x \rceil$ to refer to the usual floor and ceiling rounding operations to the integers, and $\lfloor x \rceil$ to mean a randomized rounding to $\lceil x \rceil$ with probability $x \mod 1$ and otherwise $\lfloor x \rfloor$. We subscript rounding operations (e.g. $\lfloor x \rfloor_k$) to round to $\mathbb{Z}/2^k$ instead of \mathbb{Z} .

We use statistical distance between two probability distributions \mathcal{D}_1 and \mathcal{D}_2 to refer to

$$\max_{A \subseteq \operatorname{range}(\mathcal{D}_1) \cup \operatorname{range}(\mathcal{D}_2)} \left| \Pr[x_1 \xleftarrow{R} \mathcal{D}_1, x_1 \in A] - \Pr[x_2 \xleftarrow{R} \mathcal{D}_2, x_2 \in A] \right|$$

We will denote a family of distributions indexed by λ as $\{\mathcal{D}_{\lambda}\}_{\lambda}$. We say two distribution families $\{\mathcal{D}_{\lambda}^{1}\}_{\lambda}$, $\{\mathcal{D}_{\lambda}^{2}\}_{\lambda}$ are statistically close if exists some negligible function $\mathsf{negl}(\lambda)$ such that the statistical distance between $\mathcal{D}_{\lambda}^{1}$ and $\mathcal{D}_{\lambda}^{2}$ is $< \mathsf{negl}(\lambda)$.

We notate a randomized function f which takes as input x and randomness r by f(x;r). For brevity, the function may be written as f(x) when fresh randomness r is used but does not need to be referenced.

Definition 1. We say a function $\mathsf{PDens}(1^{\lambda}, x)$ with domain $x \in D_{\lambda}$ computes the probability density of a distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ if it runs in $\mathsf{poly}(\lambda)$ time, returns a nonnegative integer, and the distribution where element x' is selected with probability

$$\frac{\mathsf{PDens}(1^{\lambda}, x')}{\sum_{x \in D_{\lambda}} \mathsf{PDens}(1^{\lambda}, x)}$$

is statistically close to $\mathcal{D}_{\lambda}^{1}$.

Since PDens runs in $poly(\lambda)$, it's output length is at most polynomial, so we can bound the maximum value with some $PDF_{max} = PDF_{max}(\lambda) \leq EXP(\lambda)$.²

The discrete Gaussian distribution on a set $S \subseteq \mathbb{R}$ with center c and standard deviation σ is defined as the distribution where an element $i \in S$ is picked with probability

$$\frac{e^{-\frac{(i-c)^2}{2\sigma^2}}}{\sum_{x\in S} e^{-\frac{(x-c)^2}{2\sigma^2}}}$$

When the set S is not specified, assume $S = \mathbb{Z}$. This distribution is of particular interest to cryptography due to its presence in the learning with errors assumption and other lattice-based cryptosystems. For simplicity, we will mostly consider univariate discrete gaussians, as, much like their continuous counterparts, discrete gaussians over arbitrary multivariate lattices can be generated via a linear transformation on a set of independent univariate discrete gaussians, where the linear transformation is derived from the covariance of the target discrete gaussian.

3 Explainable Sampling

In this section we define our notion of explainable sampling. Intuitively a distribution $\{\mathcal{D}_{\lambda}\}_{\lambda}$ can be sampled by a function $\mathsf{Sample}(1^{\lambda}; r)$ if $\mathsf{Sample}(1^{\lambda}; r)$ gives a distribution that is statistically close (in λ) to $\{\mathcal{D}_{\lambda}\}_{\lambda}$ for when the string r is chosen uniformly at random. We will further say that such a sampling algorithm is explainable if given an element x in the domain of the distribution there is a function $\mathsf{Explain}(1^{\lambda}, x) \to r'$ that will output an r' such that $\mathsf{Sample}(1^{\lambda}; r') \to x$. Moreover, the process of picking random coins to sample an element compared

¹ While this definition is not the most general for probability distributions over infinite sets, it will suffice for the cases we consider

² Note that we only require the probability density to be proportional to the probability of an element being sampled, rather than equal to. As such, any PDens function which outputs fixed precision reals can be converted to one which outputs integers as above by simply multiplying a sufficiently large constant. We choose to define our output to be integers to give a convenient fixed granularity

to first sampling a random element and then calculating the coins should be statistically close. This final property is what allows one to program in a random oracle.

Below we give our formal definitions of a distribution being sampleable and a Sample algorithm being explainable. Then we sketch how a pair of algorithms meeting this criteria can be used in a cryptographic game to sample from a random oracle.

Definition 2. We say an algorithm $\mathsf{Sample}(1^{\lambda}; r \in \{0, 1\}^n)$ is a sampler for probability distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ if Sample runs in $\mathsf{poly}(\lambda)$ time and the output of $\mathsf{Sample}(1^{\lambda}; r)$ is statistically close to $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

Definition 3. We say a Sample algorithm using $n = n(\lambda)$ bits of randomness is explainable if there exists a (likely randomized) algorithm $\text{Explain}(1^{\lambda}, 1^{\kappa}, x)$ such that $\text{Explain runs in poly}(\lambda, \kappa)$ time, and there exists a negligible function $\text{negl}(\lambda)$ such that the statistical distance between the following two distributions is at most $\frac{1}{\kappa} + \text{negl}(\lambda)^{-3}$.

Distribution A	Distribution B
$- r \xleftarrow{R} \{0, 1\}^n.$ - $x \leftarrow Sample(1^\lambda; r).$ - Return r, x .	$ \begin{array}{l} - r' \xleftarrow{R} \{0,1\}^n \\ - x \leftarrow Sample(1^\lambda;r'). \\ - r \leftarrow Explain(1^\lambda,1^\kappa,x). \\ - \operatorname{Return} r, x. \end{array} $

We make a few brief remarks on our definition. First, notice that a call to $\mathsf{Explain}(1^{\lambda}, 1^{\kappa}, x)$ algorithm is not explicitly required to even return an r such that $\mathsf{Sample}(1^{\lambda}; r) = x$. However, the definition implies that if it does not do so with sufficiently high probability, it will not meet our requirements.

Next, our Explain definition takes in a "fidelity" parameter κ in unary. Here we only require that the explain algorithm is within $1/\kappa$ statistical difference in the above game. While also requiring the explain algorithm to run in time polynomial in κ and λ .

Our motivation is to allow for greater flexibility in the case where it might be difficult to design a polytime explain algorithm where the statistical difference is negligibly close, but that there is a natural running time versus precision tradeoff in the explain algorithm. As we will see below, the latter is sufficient for proving security in a game which uses said sampler to programs a random oracle. Suppose there exists an attacker that wins with non-negligible probability ϵ in a cryptographic game that samples using a random oracle. In proving security we will "tune" κ , so we can switch from sampling from the random oracle to 'reverse sampling' using explain such that the statistical distance between these two games is still some nonnegligible fraction (say $\epsilon/2$). We again remark that

³ One could consider a computational analogue of this definition.

this relaxation to a tunable precision parameter is the main definitional difference between our framework and other works that explored reverse sampling.

Finally, we note that it will often be convenient to interpret the randomness r as being drawn from the uniform distribution on an interval of \mathbb{Z} or an exponentially precise element of \mathbb{R} rather than over uniformly random bits. It is easy to see we can interpret a bitstring as a binary representation either of the aforementioned domains of values, and so we will directly show that r is uniform on said domain rather than on the underlying bit representation.

3.1 Explainability in Cryptographic Games

We will use $\mathsf{Game}^{\mathsf{R}(\cdot)}(1^{\lambda}, \mathcal{A})$ to refer to a series of cryptographic game parameterized by λ where parties are permitted oracle access to some $\mathsf{R}(\cdot)$ against an adversary \mathcal{A} consisting of one or more algorithms (also with access to $\mathsf{R}(\cdot)$. We say $\mathsf{Game}^{\mathsf{R}(\cdot)}(1^{\lambda}, \cdot)$ is secure, if for all adversaries \mathcal{A} which run in $\mathsf{poly}(\lambda)$ time, $\Pr[\mathsf{Game}^{\mathsf{R}(\cdot)}(1^{\lambda}, \mathcal{A}) = 1] = \mathsf{negl}(\lambda)$. We will refer to the event of Game returning 1 as 'winning'.

Theorem 1. Suppose Sample is an explainable sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$ with corresponding Explain algorithm. Let $\mathsf{R}(\cdot)$ be a random oracle to distribution $\{\mathcal{D}_{\lambda}\}_{\lambda}$, and $\mathsf{R}'(\cdot)$ be a random oracle to (r, x) where $r \xleftarrow{R} \{0, 1\}^n$ and $x = \mathsf{Sample}(1^{\lambda}; r)$. Then if $\mathsf{Game}^{\mathsf{R}}(1^{\lambda}, \cdot)$ is secure, then so is $\mathsf{Game}^{\mathsf{R}'}(1^{\lambda}, \cdot)$.

Proof. Lemma 1. Suppose Sample is an explainable sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$ with corresponding Explain algorithm. Let $\mathsf{R}(\cdot)$ be a random oracle to distribution $\{\mathcal{D}_{\lambda}\}_{\lambda}$, and $\mathsf{R}_{\kappa}(\cdot)$ be a random oracle to (r, x) where $x \xleftarrow{\mathsf{R}} \{\mathcal{D}_{\lambda}\}_{\lambda}$ and $r \xleftarrow{\mathsf{Explain}(1^{\lambda}, 1^{\kappa}, x)}$. Then if $\mathsf{Game}^{\mathsf{R}}(1^{\lambda}, \cdot)$ is secure, then $\mathsf{Game}^{\mathsf{R}_{\kappa}}(1^{\lambda}, \cdot)$ is secure for all $\kappa \in \mathsf{poly}(\lambda)$.

Proof. Assume there is some PPT adversary $\mathcal{A}^{\mathsf{R}_{\kappa}}$ for which there exists $\kappa'(\lambda)$ such that \mathcal{A} wins $\mathsf{Game}^{\mathsf{R}_{\kappa}}(1^{\lambda}, \cdot)$ with noticeable probability. Then define \mathcal{A}' to be an adversary for $\mathsf{Game}^{\mathsf{R}}(1^{\lambda}, \cdot)$ which runs $\mathcal{A}^{\mathsf{R}_{\kappa}}$ and simulates oracle calls to R_{κ} by taking calling oracle R and running $\mathsf{Explain}(1^{\lambda}, 1^{\kappa}, \cdot)$ on the output. Since this is exactly the same game, we conclude \mathcal{A}' has a noticeable probability of winning $\mathsf{Game}^{\mathsf{R}}(1^{\lambda}, \cdot)$. Since $\kappa \in \mathsf{poly}(\lambda)$, this is efficient.

Lemma 2. Let $\mathsf{R}_{\kappa}(\cdot)$ be a random oracle to (r, x) where $x \xleftarrow{R} \{\mathcal{D}_{\lambda}\}_{\lambda}$ and $r \leftarrow \mathsf{Explain}(1^{\lambda}, 1^{\kappa}, x)$. and $\mathsf{R}'(\cdot)$ be a random oracle to (r, x) where $r \xleftarrow{R} \{0, 1\}^n$ and $x = \mathsf{Sample}(1^{\lambda}; r)$. Suppose $\mathsf{Game}^{\mathsf{R}_{\kappa}}(1^{\lambda}, \cdot)$ is secure for all $\kappa \in \mathsf{poly}(\lambda)$, then $\mathsf{Game}^{\mathsf{R}'(\cdot)}(1^{\lambda}, \cdot)$ is secure.

Proof. Again, assume for sake of contradiction there is some adversary PPT $\mathcal{A}^{\mathsf{R}'}$ which wins $\mathsf{Game}^{\mathsf{R}'}(1^{\lambda}, \cdot)$ with noticeable probability. Specifically, since it is poly time, let's suppose there exists constants a, b such that $\mathcal{A}^{\mathsf{R}'}$ wins $\mathsf{Game}^{\mathsf{R}'}(1^{\lambda}, \cdot)$ makes at most λ^a queries and wins with probability $> \lambda^{-b}$ infinitely often. However, notice that by Definition 3, the statistical distance between queries to a query to R' and a query to R_{κ} is $\frac{1}{\kappa}$. Thus, if we set $\kappa = 2\lambda^{a+b}$, we can union

bound the total statistical difference of all queries with $\frac{1}{2\lambda^b} + \operatorname{negl}(\lambda)$. Observe that this means we can bound

$$\left|\Pr[\mathsf{Game}^{\mathsf{R}'}(1^{\lambda},\mathcal{A})=1] - \Pr[\mathsf{Game}^{\mathsf{R}_{2\lambda^{a+b}}}(1^{\lambda},\mathcal{A})=1]\right| \leq \frac{1}{2\lambda^{b}} + \mathsf{negl}(\lambda)$$

However, Since \mathcal{A} wins $\mathsf{Game}^{\mathsf{R}_{2\lambda^{a+b}}}(1^{\lambda}, \mathcal{A})$ with probability $\frac{1}{\lambda^{b}}$ infinitely often, that means it wins $\mathsf{Game}^{\mathsf{R}_{2\lambda^{a+b}}}(1^{\lambda}, \mathcal{A})$ with probability $\geq \frac{1}{2\lambda^{b}}$ infinitely often, contradicting the assumption that $\mathsf{Game}^{\mathsf{R}_{\kappa}}(1^{\lambda}, \cdot)$ is secure for all $\kappa \in \mathsf{poly}(\lambda)$. \Box

Taking Lemma 1 and Lemma 2 together gives us the theorem statement.

4 Explaining Sampling over Small Ranges with respect to Discrete Gaussian Samplers

We begin by showing that any efficient Sample algorithm over a polynomial sized range, $\{R_{\lambda}\}_{\lambda}$, is explainable. The core idea is rather simple. A call to Explain $(1^{\lambda}, 1^{\kappa}, x)$ will simply call the Sample algorithm up to $\kappa \cdot |R_{\lambda}|$ times until x is output.

We show an immediate corollary to the theorem where if a distribution family has a super-polynomial size range $\{R_{\lambda}\}_{\lambda}$, but there exists subsets of the range $\{S_{\lambda}\}_{\lambda}$ where $S_{\lambda} \subseteq R_{\lambda}$ and $\{S_{\lambda}\}_{\lambda}$ are polynomially sized, then our sampling algorithm also works for these distribution. In particular, this covers a discrete Gaussian where the standard deviation σ grows polynomially with λ . We will see that this simple procedure can serve for explaining the "base case" for the Micciancio-Walter algorithm. Below we formally give our theorem and proof.

Theorem 2. Let Sample be an sampler for some distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ on range $\{R_{\lambda}\}_{\lambda}$. If $|R_{\lambda}| \leq \text{poly}(\lambda)$, then Sample is explainable.

Proof. The idea here is to simply brute force the sampler output to find a valid randomness to a given element x. We use the fact that the range is polynomial to bound the amount of probability mass which can be contained by 'infrequent' elements. Consider the algorithm below:

Explain $(1^{\lambda}, 1^{\kappa}, x)$ - Repeat $\kappa \cdot |R_{\lambda}|$ times • Select fresh rep

- Select fresh randomness r
- $x' \leftarrow \mathsf{Sample}(1^{\lambda}; r)$
- If x' = x, stop and return r.
- Return \perp

Claim 3. Explain runs in $poly(\lambda, \kappa)$ time

Proof. First, by the efficiency requirement of Definition $2 \kappa \cdot |R_{\lambda}|$, Sample $(1^{\lambda}; r)$ has runtime $poly(\lambda)$. Explain simply calls $Sample(1^{\lambda}; r)$ up to $\kappa \cdot |R_{\lambda}|$ many times (along with some other minor efficient computation). Thus, we can bound the runtime with $\kappa \cdot |R_{\lambda}| \cdot poly(\lambda)$. Since $|R_{\lambda}| \leq poly(\lambda)$ by assumption, this bounds the runtime with $\kappa \cdot poly(\lambda) \in poly(\lambda, \kappa)$.

Claim 4. Explain returns \perp in Game B with probability $\leq \frac{1}{\kappa}$

Proof. Observe that over the course of Game B, Sample is called on independent randomness up to $\kappa \cdot |R_{\lambda}| + 1$ times (once directly from the Game and $\kappa \cdot |R_{\lambda}|$ times in the execution of Game B). Let us call these outputs x_0 and $x_1, x_2, \ldots, x_{\kappa \cdot |R_{\lambda}|}$ respectively. Let the set $\mathcal{X} = \{x_i : \forall j \neq i \ x_j \neq x_i\}$ By pigeonhole, we know that $|\mathcal{X}| \leq |R_{\lambda}| - 1$. Thus, since the calls are independent, the probability that the single call directly from Game B (x_0) is in this set \mathcal{X} is $\frac{|\mathcal{X}|}{\kappa |R_{\lambda}| + 1} \leq \frac{|R_{\lambda}| - 1}{\kappa |R_{\lambda}| + 1} < \frac{1}{\kappa}$. On the other hand, if this call is not unique, we can see Explain finds j > 0: $x_j = x_0$ and so does not return \perp .

Claim 5. The statistical distance of Game A and Game B in Definition 3 using Explain is $\leq \frac{1}{r}$

Proof. We first note that if $\text{Explain}(1^{\lambda}, 1^{\kappa}, x)$ returned $r \neq \bot$, then $x = x' = \text{Sample}(1^{\lambda}; r)$. Thus, since x is the same function of r in both games, it suffices to show that r in Game B has statistical distance $\leq \frac{1}{\kappa}$ to uniform.

Consider an alternate Explain' which, rather than sampling x' at most $\kappa \cdot |R_{\lambda}|$ times, samples x' until it finds an appropriate r. Here, we can compute the probability that any particular r is chosen as the probability the x chosen for Explain is equal to Sample $(1^{\lambda}; r)$ multiplied by the probability that r is chosen conditional on $x' \leftarrow \text{Sample}(1^{\lambda}; r)$. This is equal to

$$\Pr[\mathsf{Sample}(1^{\lambda}) = x] \cdot \frac{2^{-n}}{\Pr[\mathsf{Sample}(1^{\lambda}) = x']} = 2^{-n}$$

so r is uniform on $\{0,1\}^n$. Now note that Explain' only differs from Explain when Explain returns \perp before finding such an r. By Claim 4, this happens with probability $<\frac{1}{\kappa}$, bounding the statistical distance.

Corollary 6. Let Sample be a sampler for some distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ with range $\{R_{\lambda}\}_{\lambda}$. If there exists sets $\{S_{\lambda}\}_{\lambda}$ where $S_{\lambda} \subseteq R_{\lambda}$, $|S_{\lambda}| \leq \mathsf{poly}(\lambda)$, and

$$\Pr\left[\begin{matrix} x \xleftarrow{R} \mathcal{D}_{\lambda} \\ x \notin S_{\lambda} \end{matrix} \right] < \mathsf{negl}(\lambda),$$

then Sample is explainable.

Proof. Let $\{\mathcal{D}'_{\lambda}\}_{\lambda}$ be the distribution $\{\mathcal{D}_{\lambda}\}_{\lambda}$ conditional on the output being $\in \{S_{\lambda}\}_{\lambda}$. It is easy to see that $\{\mathcal{D}'_{\lambda}\}_{\lambda}$ satisfies the conditions for Theorem 2, and so any sampler for $\{\mathcal{D}'_{\lambda}\}_{\lambda}$ is explainable. Since $\{\mathcal{D}'_{\lambda}\}_{\lambda}$ is statistically close to $\{\mathcal{D}_{\lambda}\}_{\lambda}$, any sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$ is a sampler for $\{\mathcal{D}'_{\lambda}\}_{\lambda}$, and so is explainable. \Box

5 Sampling and Explaining Conforming Distributions

The previous section showed how one could explain arbitrary sampler so long as the corresponding distributions grew polynomially in the security parameter. In this section we explore a class of distributions for which we can build explainable samplers. This includes distributions with superpolynomial sized ranges. To perform this we require the distribution family to have:

- 1. A Sample algorithm for the distribution. (It is not necessarily explainable).
- 2. A probability density function PDens where $\mathsf{PDens}(1^{\lambda}, x)$ returns the probability density function of x for distribution \mathcal{D}_{λ} .
- 3. An "heavy element sampler" algorithm SampleUniform which explainably samples from the uniform distribution over elements in the support of \mathcal{D}_{λ} that have probability density above p.

We show that if a distribution has all three elements, then there exists another pair of algorithms Sample', Explain' that comprise an explainable sampler for the family $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

We argue for the utility of this transformation by observing that many natural distributions, including the discrete Gaussian, have an easily computable heavy element sampler. If we consider the discrete Gaussian in particular, then a heavy element sampler simply needs to calculate the values x_0, x_1 such that $\mathsf{PDens}(1^\lambda, x_0) = \mathsf{PDens}(1^\lambda, x_1) = p$, which can be done by explicitly solving for p using the probability density function of $e^{-\frac{(x-c)^2}{2\sigma^2}} = p$. We can then choose integers uniformly at random in the interval $[x_0, x_1]$.

In fact, in general for monotonic distributions (or distributions which can be partitioned into a polynomial number of monotone segments) on ordered sets, we can easily compute an explainable heavy element sampler by binary searching for the endpoints of the ranges of heavy element with only polynomially many calls to the probability density function, then uniformly sampling on the interval found. As long as the underlying domain has an explainable representation (such as \mathbb{Z} or \mathbb{R}), this sampler too is explainable. This condition alone encompasses almost all frequently seen distributions such as (discrete) Gaussians, binomial, geometric, Poisson, etc.

We conclude by arguing that having such a heavy element sampler is necessary for a generic transformation. To do this we provide a distribution for which oracle access to the first two properties is not sufficient to construct an explainable sampler to said distribution.

5.1 Explainable Sampling through heavy element samplers

Definition 4. We say a function SampleUniform $(1^{\lambda}, p; r)$ is a heavy element sampler for a distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ with probability density function PDens if it runs in $\operatorname{poly}(\lambda)$ time, and for all $p \in [0, \operatorname{PDF}_{\max})$, $S_p(1^{\lambda}; r) = \operatorname{SampleUniform}(1^{\lambda}, p; r)$ is a sampler for the uniform distribution on elements $x \in R_{\lambda}$ with $\operatorname{PDens}(1^{\lambda}, x) \geq$ p (we will notate said set as R_{λ}^{p} . We say a heavy element sampler is explainable if there exists a $\operatorname{poly}(1^{\lambda}, 1^{\kappa})$ algorithm ExplainUniform $(1^{\lambda}, 1^{\kappa}, p, x)$ such that for all $p = p(\lambda) \in \mathsf{EXP}(\lambda), \ E_p(1^{\lambda}, 1^{\kappa}, x) = \mathsf{ExplainUniform}(1^{\lambda}, 1^{\kappa}, p, x) \ is \ an \ explainer for \ S_p.$

Theorem 7. Let Sample, PDens, SampleUniform be a sampler, probability density function, and explainable heavy element sampler for the probability distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ on range $\{R_{\lambda}\}_{\lambda}$. Then there exists a Sample' which is an explainable sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

Proof. The underlying idea of the construction of this Sample', Explain' is very similar to the generic Explain algorithm in Theorem 2 - we can use a brute force approach to polynomially approximate $\{\mathcal{D}_{\lambda}\}_{\lambda}$. However, because the domain can now be superpolynomial, we can no longer hope polynomially approximate the exact elements of $\{\mathcal{D}_{\lambda}\}_{\lambda}$ itself. Instead, we categorize elements into 'buckets' of similar approximate probability density. However, ordinarily, finding an element of a similar probability density would not suffice to satisfy the 'correctness' of an Explain algorithm. Here, we modify our sampler Sample' so that it uses Sample to produce an initial sample in the range, but this is then 'smudged' to an element of the range with similar probability density which is the actual output of Sample'.

By increasing the precision parameter κ , we decrease the size of each bucket, decreasing the statistical distance of Explain' but simultaneously increasing the expected number of tries to find an element in the same bucket.

Let Sample' be as follows:

$$\begin{split} \mathsf{Sample}'(1^{\lambda}; r = r_0 \in \{0, 1\}^n, s_0 \in [0, 1], t_0 \in \{0, 1\}^n) \\ &- x' \leftarrow \mathsf{Sample}(1^{\lambda}; r_0) \\ &- p' \leftarrow \mathsf{PDens}(1^{\lambda}, x') \\ &- p \leftarrow p' \cdot s_0 \\ &- x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p; t_0) \\ &- \operatorname{Return} x \end{split}$$

As noted in Section 3, we can interpret a uniform bitstring as the binary expansion of a real number $\in [0, 1]$. Since randomness s_0 is only used to compute the probability density on SampleUniform, which is integral, it suffices to use only $\log(\mathsf{PDF}_{\max}) \in \log(\mathsf{EXP}(\lambda)) = \mathsf{poly}(\lambda)$ bits of randomness.

Proof of Sampleability We first prove that Sample' is a good sampler per Definition 2.

Definition 5. We define the distribution $\{\mathsf{PDF}(\mathcal{D}_{\lambda})\}_{\lambda}$ be defined as the joint distribution on two variables (a, b) such that the distribution of a is \mathcal{D}_{λ} and b is uniform from $[0, \mathsf{PDens}(1^{\lambda}, a))$ where PDens is a probability density function of \mathcal{D}_{λ} .

Lemma 3. The following distributions are statistically close

Distribution $1 = \mathcal{D}_1$	Distribution $2 = \mathcal{D}_2$
$ - (a,b) \xleftarrow{R} PDF(\mathcal{D}_{\lambda}) - \text{Output} (a,b) $	$- (a', b) \xleftarrow{R} PDF(\mathcal{D}_{\lambda}) - a \xleftarrow{R} SampleUniform(1^{\lambda}, b) - Output (a, b)$

Proof. Consider some fixed (a^*, b^*) in the support of $\mathsf{PDF}(\mathcal{D}_{\lambda})$. We can compute the explicit probability this element is picked in distribution 1 as the probability a^* is picked - which is $\frac{\mathsf{PDens}(a^*)}{\sum_{a \in R_{\lambda}} \mathsf{PDens}(a)}$, multiplied by the probability that b^* is picked given that a^* - which is $\frac{1}{\mathsf{PDens}(a^*)}$, as this is a uniform distribution of size $\mathsf{PDens}(a^*)$. Together, we get

$$\Pr[(a^*, b^*) \leftarrow \mathcal{D}_1] = \frac{\mathsf{PDens}(a^*)}{\sum_{a \in R_\lambda} \mathsf{PDens}(a)} \cdot \frac{1}{\mathsf{PDens}(a^*)} = \frac{1}{\sum_{a \in R_\lambda} \mathsf{PDens}(a)}$$

In distribution 2, we can write the probability (a^*, b^*) occurs as

$$\Pr[(_, b^*) \leftarrow \mathsf{PDF}(\mathcal{D}_{\lambda})] \cdot \Pr[a^* = \mathsf{SampleUniform}(1^{\lambda}, b^*)|(_, b^*) \leftarrow \mathsf{PDF}(\mathcal{D}_{\lambda})]$$

Note that since SampleUniform and $\mathsf{PDF}(\mathcal{D}_{\lambda})$ are invoked independently, we can ignore the conditional. Using our analysis from distribution 1, we can compute the first probability as

$$\begin{split} \Pr[(_, b^*) \leftarrow \mathsf{PDF}(\mathcal{D}_{\lambda})] &= \sum_{a: \mathsf{PDens}(a) > b^*} \Pr[(a, b^*) \leftarrow \mathsf{PDF}(\mathcal{D}_{\lambda})] \\ &= \sum_{a: \mathsf{PDens}(a) > b^*} \frac{1}{\sum_{a \in R_{\lambda}} \mathsf{PDens}(a)} \end{split}$$

Meanwhile, since by definition SampleUniform outputs a uniform element of sufficiently high probability density, we have that

$$\Pr[\mathsf{SampleUniform}(1^{\lambda}, b^*) = a^*] = \frac{1}{|\{a : \mathsf{PDens}(a) > b^*\}|}$$

Which brings the total probability of $(a^*, b^*) \leftarrow \mathcal{D}_2$ as

$$\sum_{a:\mathsf{PDens}(a) > b^*} \frac{1}{\sum_{a \in R_{\lambda}} \mathsf{PDens}(a)} \cdot \frac{1}{|\{a:\mathsf{PDens}(a) > b^*\}|} = \frac{1}{\sum_{a \in R_{\lambda}} \mathsf{PDens}(a)}$$

the same as in \mathcal{D}_1

14

Lemma 4. Sample' is an sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

Proof.

Claim 8. Sample' produces a distribution statistically close to $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

Proof. Since Sample is a sampler to $\{\mathcal{D}_{\lambda}\}_{\lambda}$, the distribution of (x', p) as defined in Sample' is statistically close to $\{\mathsf{PDF}(\mathcal{D}_{\lambda})\}_{\lambda}$. By Lemma 3, the distribution of (x, p) must also be statistically close to $\{\mathsf{PDF}(\mathcal{D}_{\lambda})\}_{\lambda}$. By Definition 5, the distribution of x must be statistically close to $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

Claim 9. Sample' runs in $poly(\lambda)$ time.

Proof. Sample' makes a single call to each of Sample, PDens, SampleUniform, which, by assumption, are $poly(\lambda)$ time algorithms along with a single multiplication, and so is also $poly(\lambda)$ time.

By Claim 8 and Claim 9, Sample' fulfills the definition of a sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$

Proof of Explainability We next prove that Sample' is explainable per Definition 3. The general idea here is that because x is chosen randomly using the value p generated by PDens rather than directly from Sample, elements x' with similar PDens values are similarly likely to be explanation for x. We can then use our precision parameter κ to define intervals of PDens($1^{\lambda}, x'$) and sample an interval of 'acceptable' p values relative to the correct conditional density. By restricting the number and size of such intervals, we can guarantee that a polynomial number of calls to Sample finds an x' in the chosen interval while still ensuring the p is approximated fairly closely.

Lemma 5. Sample' is explainable.

Proof. Consider the following Explain' algorithm for Sample'.

Explain' $(1^{\lambda}, 1^{\kappa}, x)$ $\mathsf{PIdx}(x,\rho)$ - Let $b = 1 + \frac{1}{2}$ $-p \xleftarrow{R} [0, \mathsf{PDens}(x)]$ - If $\mathsf{PDens}(x) = 0$, $-x_0 \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p)$ return \perp - Run 9 ln(PDF_{max}) $\cdot \kappa^2 \cdot \lambda$ times: – Return • Generate fresh randomness r'. $\left|\log_{b}(\mathsf{PDens}(x))\right|$ • $x' \leftarrow \mathsf{Sample}(1^{\lambda}; r').$ • If $\mathsf{Pldx}(x_0, 3\kappa) = \mathsf{Pldx}(x', 3\kappa)$. * Set $r_0 = r'$ * Set $s_0 = \frac{p}{\mathsf{PDens}(x')}$ * Set $t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, p, x)$ * Return $r = (r_0, s_0, t_0)$ - Return \perp .

Claim 10. Explain' runs in $poly(\lambda, \kappa)$ time.

Proof. Explain' utilizes some efficient (polynomial time) subprocedures in Sample, SampleUniform, PDens, and Pldx. Moreover, since $\mathsf{PDF}_{\max} \in \mathsf{EXP}(\lambda)$, $\log(\mathsf{PDF}_{\max}) \in \mathsf{poly}(\lambda)$, so the algorithm loops only a $\mathsf{poly}(\lambda,\kappa)$ amount of times. \Box

Claim 11. The statistical distance of Game A and Game B in Definition 3 using Explain' is $<\frac{1}{\kappa} + \operatorname{negl}(\lambda)$

Proof. We will proceed with a sequence of games argument, where Game 0 is an execution of Game A using Sample' and Game 11 is an execution of Game B using Sample', Explain'.

Game 0

Game 0 (Sample' Expanded)

$- r = (r_0, s_0, t_0) \xleftarrow{R} \{0, 1\}^n \times \{0, 1\} \times \{0, 1\}^n$ $- x \leftarrow Sample'(x; r)$ $- \operatorname{Return} r, x.$	$\begin{array}{cccc} & - & r_0 & \stackrel{R}{\leftarrow} & \{0,1\}^n, \ s_0 & \stackrel{R}{\leftarrow} & [0,1], \ t_0 & \stackrel{R}{\leftarrow} \\ & & \{0,1\}^n \\ & - & x' \leftarrow Sample(1^\lambda;r_0) \\ & - & p' \leftarrow PDens(1^\lambda,x') \\ & - & p = p' \cdot s_0 \\ & - & x \leftarrow SampleUniform(1^\lambda,p;t_0) \end{array}$
	- Return r_0, s_0, t_0, x .

Game 1

 $\begin{array}{l} - \begin{array}{c} r_0 \xleftarrow{R} \{0,1\}^n, s_0 \xleftarrow{R} [0,1], t_0 \xleftarrow{R} \{0,1\}^n \\ - \begin{array}{c} x'' \xleftarrow{R} \mathcal{D}_\lambda \\ \end{array} \\ - \begin{array}{c} \text{Run until break}^4 \\ \bullet \begin{array}{c} \text{Generate fresh randomness } r' \\ \bullet \begin{array}{c} x' \leftarrow \text{Sample}(1^\lambda, r') \\ \bullet \begin{array}{c} \text{If } x' = x'' \\ & \ast \begin{array}{c} \text{Set } r_0 = r' \end{array} \text{ and break} \\ - \begin{array}{c} p' \leftarrow \text{PDens}(1^\lambda, x'') \\ - \begin{array}{c} p = p' \cdot s_0 \\ - \begin{array}{c} x \leftarrow \text{SampleUniform}(1^\lambda, p; t_0) \\ \end{array} \\ - \begin{array}{c} \text{Return } r_0, s_0, t_0, x. \end{array}$

Game 2

 $\begin{array}{l} - \begin{array}{c} s_{0} \xleftarrow{R} [0,1], t_{0} \xleftarrow{R} \{0,1\}^{n} \\ - \begin{array}{c} x'' \xleftarrow{R} \mathcal{D}_{\lambda} \\ - \text{ Run until break} \\ \bullet \text{ Generate fresh randomness } r' \\ \bullet \begin{array}{c} x' \leftarrow \text{Sample}(1^{\lambda}, r') \\ \bullet \begin{array}{c} \text{ If } x'' = x' \\ & * \begin{array}{c} \text{ Set } r_{0} = r' \text{ and break} \end{array} \\ - \begin{array}{c} p'' \xleftarrow{R} [0, \text{PDens}(1^{\lambda}, x'')] \end{array}$

 $^{^4}$ this process could potentially take unbounded time, but simply act as 'bridging' steps to make the change in distribution easier to see. The final game will be efficient

- Set $s_0 = \frac{p''}{\mathsf{PDens}(1^\lambda, x'')}$
- $-x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p''; t_0)$
- Return r_0, s_0, t_0, x .

Game 3

 $- t_0 \xleftarrow{R} \{0,1\}^n$ $- (x'',p'') \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda})$ - Run until break $\bullet \text{ Generate fresh randomness } r'$ $\bullet x' \leftarrow \mathsf{Sample}(1^{\lambda},r')$ $\bullet \text{ If } x' = x''$ $* \text{ Set } r_0 = r' \text{ and break}$ $- \text{ Set } s_0 = \frac{p''}{\mathsf{PDens}(1^{\lambda},x_0)}$ $- x \leftarrow \mathsf{SampleUniform}(1^{\lambda},p'';t_0)$ $- \text{ Return } r_0, s_0, t_0, x.$

Game 4

 $\begin{array}{l} - \underbrace{t_0 \leftarrow^R \quad \{0,1\}^n}_{- \quad (x'',p'') \leftarrow^R \quad \mathsf{PDF}(\mathcal{D}_\lambda) \\ - & \operatorname{Run \ until \ break} \\ \bullet & \operatorname{Generate \ fresh \ randomness \ r'} \\ \bullet & x' \leftarrow \mathsf{Sample}(1^\lambda,r') \\ \bullet & \operatorname{If \ } x' = x'' \\ & * \ \operatorname{Set \ } r_0 = r' \ \operatorname{and \ break} \\ - & \operatorname{Set \ } s_0 = \frac{p''}{\operatorname{PDens}(1^\lambda,x'')} \\ - & x \leftarrow \mathsf{SampleUniform}(1^\lambda,p'') \\ - & \operatorname{Set \ } t_0 = \mathsf{ExplainUniform}(1^\lambda,1^{3\kappa},p'',x) \\ - & \operatorname{Return \ } r_0,s_0,t_0,x. \end{array}$

Game 5

 $\begin{array}{l} - (x'',p'') \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) \\ - x \leftarrow \mathsf{SampleUniform}(1^{\lambda},p'') \\ - & \mathsf{Run until break} \\ \bullet & \mathsf{Generate fresh randomness } r' \\ \bullet x' \leftarrow \mathsf{Sample}(1^{\lambda},r') \\ \bullet & \mathsf{If } x' = x'' \\ & * \; \mathsf{Set } r_0 = r' \\ & * \; \mathsf{Set } r_0 = r' \\ & * \; \mathsf{Set } s_0 = \frac{p''}{\mathsf{PDens}(1^{\lambda},x')} \\ & * \; \mathsf{Set } t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, p'', x) \\ & * \; \mathsf{Break} \\ - \; \mathsf{Return } r_0, s_0, t_0, x. \end{array}$

Game 6

 $- (x'', p'') \stackrel{R}{\leftarrow} \mathsf{PDF}(\mathcal{D}_{\lambda})$ $- \operatorname{Set} x_{0} = x''$ $- p \stackrel{R}{\leftarrow} [0, \mathsf{PDens}(1^{\lambda}, x_{0})]$ $- x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p)$ $- \operatorname{Run until break}$ $\bullet \operatorname{Generate fresh randomness } r'$ $\bullet x' \leftarrow \mathsf{Sample}(1^{\lambda}, r')$ $\bullet \operatorname{If} x' = x''$ $* \operatorname{Set} r_{0} = r'$ $* \operatorname{Set} s_{0} = \frac{p}{\mathsf{PDens}(1^{\lambda}, x')}$

- * Set $t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, \mathbf{p}, x)$
- * Break
- Return r_0, s_0, t_0, x .

Game 7

 $\begin{array}{l} - (x'',p'') \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) \\ - \operatorname{Sample} x_{0} \xleftarrow{R} \mathcal{D}_{\lambda} \text{ conditional on } \mathsf{Pldx}(x_{0},3\kappa) = \mathsf{Pldx}(x'',3\kappa) \\ - p \xleftarrow{R} [0, \mathsf{PDens}(1^{\lambda}, x_{0})] \\ - x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p) \\ - \operatorname{Run until break} \\ \bullet \text{ Generate fresh randomness } r' \\ \bullet x' \leftarrow \mathsf{Sample}(1^{\lambda}, r') \\ \bullet \text{ If } x' = x'' \\ &* \operatorname{Set} r_{0} = r' \\ &* \operatorname{Set} s_{0} = \frac{p}{\mathsf{PDens}(1^{\lambda}, x')} \end{array}$

- * Set $t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, p, x)$
- * Break
- Return r_0, s_0, t_0, x .

Game 8

- $(x_0, p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda})$

- Sample $x'' \xleftarrow{R} \mathcal{D}_{\lambda}$ conditional on $\mathsf{Pldx}(x_0, 3\kappa) = \mathsf{Pldx}(x'', 3\kappa)$
- $x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p)$
- Run until break
 - Generate fresh randomness r^\prime
 - $x' \leftarrow \mathsf{Sample}(1^{\lambda}, r')$
 - If x' = x''
 - * Set $r_0 = r'$
 - * Set $s_0 = \frac{p}{\mathsf{PDens}(1^\lambda, x')}$
 - * Set $t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, p, x)$
 - * Break
- Return r_0, s_0, t_0, x .

Game 9

- $\begin{array}{l} -(x_0,p) \stackrel{R}{\leftarrow} \mathsf{PDF}(\mathcal{D}_{\lambda}) \\ -x \leftarrow \mathsf{SampleUniform}(1^{\lambda},p) \\ -\operatorname{Run} 9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2 \cdot \lambda \text{ times} \\ \bullet \text{ Generate fresh randomness } r' \\ \bullet x' \leftarrow \mathsf{Sample}(1^{\lambda},r') \\ \bullet \text{ If } \mathsf{Pldx}(x_0,3\kappa) = \mathsf{Pldx}(x',3\kappa) \\ &* \operatorname{Set} r_0 = r' \\ &* \operatorname{Set} s_0 = \frac{p}{\mathsf{PDens}(1^{\lambda},x')} \\ &* \operatorname{Set} t_0 = \mathsf{ExplainUniform}(1^{\lambda},1^{3\kappa},p,x) \\ &* \operatorname{Break} \end{array}$
 - Return r_0, s_0, t_0, x .

Game 10

 $\begin{array}{l} - (x,p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) \\ - x_0 \leftarrow \mathsf{SampleUniform}(1^{\lambda},p) \\ - \operatorname{Run} 9 \ln(\mathsf{PDF}_{\max}) \cdot \kappa^2 \cdot \lambda \text{ times} \\ \bullet \text{ Generate fresh randomness } r' \\ \bullet x' \leftarrow \mathsf{Sample}(1^{\lambda},r') \end{array}$

- If $\mathsf{Pldx}(x_0, 3\kappa) = \mathsf{Pldx}(x', 3\kappa)$
- * Set $r_0 = r'$
 - * Set $s_0 = \frac{p}{\mathsf{PDens}(1^\lambda, x')}$
 - * Set $t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, p, x)$
 - * Break
- Return r_0, s_0, t_0, x .

Game 11

- $x \leftarrow \mathsf{Sample}'(1^{\lambda})$ $- p \xleftarrow{R} [0, \mathsf{PDens}(x)]$ $- x_0 \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p)$
- $x_0 \leftarrow \text{SampleOmorm}(1, p)$ - Run $9\ln(\text{PDF}_{max}) \cdot \kappa^2 \cdot \lambda$ times
- Run 9 ln (PDF_{max}) $\cdot \kappa \cdot \lambda$ times
 - Generate fresh randomness r'
 - $x' \leftarrow \mathsf{Sample}(1^{\lambda}, r')$
 - If $\mathsf{Pldx}(x_0, 3\kappa) = \mathsf{Pldx}(x', 3\kappa)$
 - * Set $r_0 = r'$
 - * Set $s_0 = \frac{p}{\mathsf{PDens}(1^\lambda, x')}$
 - * Set $t_0 = \mathsf{ExplainUniform}(1^{\lambda}, 1^{3\kappa}, p, x)$
 - * Break
- Return r_0, s_0, t_0, x .

Game 11 (Shortened)

 $-x \leftarrow \mathsf{Sample}'(1^{\lambda})$

 $-r = (r_0, s_0, t_0) \leftarrow \mathsf{Explain}(1^{\lambda}, 1^{\kappa}, x)$

- Return r, x.

Claim 12. The distributions output by Game 0 and Game 1 have statistical distance $negl(\lambda)$.

Proof. By Definition 2, the distribution of x' in Game 0 (from Sample) and Game 1 (from \mathcal{D}_{λ}) are statistically close. Now since r_0 is uniform, the conditional probability of r_0 on any fixed x' is still uniform on all r_0 such that Sample $(1^{\lambda}; r_0) = x'$ in both games. The only other changes are notational between x' and x'', which are equal in this game.

Claim 13. The distributions output by Game 1 and Game 2 have statistical distance 0.

Proof. This game only changes the way s_0 is generated. In Game 1, the distribution of s_0 is uniform on [0, 1]. In Game 2, it is the quotient of p'' which is uniform on $[0, \mathsf{PDens}(1^{\lambda}, x'')]$ and $\mathsf{PDens}(1^{\lambda}, x'')$, which is simply uniform on [0, 1]. \Box

Claim 14. The distributions output by Game 2 and Game 3 have statistical distance 0.

Proof. By definition, $(x'', p'') \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda})$ is defined to be $x'' \xleftarrow{R} \mathcal{D}_{\lambda}$ and $p'' \xleftarrow{R} [0, \mathsf{PDens}(1^{\lambda}, x'')]$, which is exactly how it is generated in Game 3.

Claim 15. The distributions output by Game 3 and Game 4 have statistical distance $\frac{1}{3\kappa} + \operatorname{negl}(\lambda)$.

Proof. By Definition 3, the distributions

Game A	Game B
$- t_0 \xleftarrow{R} \{0, 1\}^n.$ - $x \leftarrow SampleUniform(1^{\lambda}, p; t_0).$ - Return t_0, x .	$\begin{array}{l} - x \leftarrow SampleUniform(1^{\lambda}, p). \\ - r \leftarrow ExplainUniform(1^{\lambda}, 1^{3\kappa}, p, x). \\ - \operatorname{Return} t_0, x. \end{array}$

have statistical distance $\frac{1}{3\kappa} + \operatorname{negl}(\lambda)$. Note this corresponds exactly to how t_0, x are generated in Games 3 and 4 respectively.

Claim 16. The distributions output by Game 4 and Game 5 have statistical distance 0.

Proof. The changes in Game 5 only change the order some variables are generated, but not the way they are generated. \Box

Claim 17. The distributions output by Game 5 and Game 6 have statistical distance 0.

Proof. In this game, we substitute all uses of p'' with p. By Definition 1, $x'' \leftarrow \mathcal{D}_{\lambda}$, so (x'', p) is distributed according to $\mathsf{PDF}(\mathcal{D}_{\lambda})$, so are statistically identical.

Claim 18. The distributions output by Game 6 and Game 7 have statistical distance $\frac{1}{3\kappa}$.

Proof. In this game, rather than setting $x_0 = x''$, we set x_0 to be an element with the same Pldx as x''. Observe that x_0 is only used to sample p, and by definition of Pldx, $\frac{|\mathsf{PDens}(1^{\lambda}, x_0) - \mathsf{PDens}(1^{\lambda}, x'')|}{\mathsf{PDens}(1^{\lambda}, x'')}$ is at most $\frac{1}{3\kappa}$, and so we can conclude that the uniform distribution on $[0, \mathsf{PDens}(1^{\lambda}, x_0)]$ and $[0, \mathsf{PDens}(1^{\lambda}, x'')]$ have statistical distance at most $\frac{1}{3\kappa}$.

Claim 19. The distributions output by Game 7 and Game 8 have statistical distance 0.

Proof. By this game, note that p'' is unused, so we can examine the difference in how the joint distribution on x_0, x'', p is generated. Let x_0^*, x''^*, p^* be some set of values taken by x_0, x'', p . We can see the probability density of this is in Game 7 proportional to

$$\Pr_{x'' \leftarrow \mathcal{D}_{\lambda}} [x''^* = x''] \cdot \Pr_{x_0 \leftarrow \mathcal{D}_{\lambda}} [x_0^* = x_0 | \mathsf{Pldx}(x''^*, 3\kappa) = \mathsf{Pldx}(x_0, 3\kappa)] \cdot \frac{1}{\mathsf{PDens}(1^{\lambda}, x_0)}$$
$$= \Pr_{x'' \leftarrow \mathcal{D}_{\lambda}} [x''^* = x''] \cdot \frac{\Pr_{x_0 \leftarrow \mathcal{D}_{\lambda}} [x_0^* = x_0]}{\mathsf{Pr}_{x_0 \leftarrow \mathcal{D}_{\lambda}} [\mathsf{Pldx}(x''^*, 3\kappa) = \mathsf{Pldx}(x_0, 3\kappa)]} \cdot \frac{1}{\mathsf{PDens}(1^{\lambda}, x_0)}$$

by definition of conditional probability

$$= \frac{\Pr_{x'' \xleftarrow{R} \mathcal{D}_{\lambda}}[x''^* = x'']}{\Pr_{x_0 \xleftarrow{R} \mathcal{D}_{\lambda}}[\mathsf{Pldx}(x''^*, 3\kappa) = \mathsf{Pldx}(x_0, 3\kappa)]} \cdot \Pr_{x_0 \xleftarrow{R} \mathcal{D}_{\lambda}}[x_0^* = x_0] \cdot \frac{1}{\mathsf{PDens}(1^{\lambda}, x_0)}$$

via simple algebraic manipulation

$$= \frac{\Pr_{x'' \xleftarrow{R} \mathcal{D}_{\lambda}}[x''^* = x'']}{\Pr_{x'' \xleftarrow{R} \mathcal{D}_{\lambda}}[\mathsf{Pldx}(x''^*, 3\kappa) = \mathsf{Pldx}(x'', 3\kappa)]} \cdot \Pr_{x_0 \xleftarrow{R} \mathcal{D}_{\lambda}}[x_0^* = x_0] \cdot \frac{1}{\mathsf{PDens}(1^{\lambda}, x_0)}$$

by simply renaming x_0 to x''

$$= \Pr_{x'' \xleftarrow{R} \mathcal{D}_{\lambda}} [x''^* = x'' | \mathsf{Pldx}(x''^*, 3\kappa) = \mathsf{Pldx}(x'', 3\kappa)] \cdot \Pr_{x_0 \xleftarrow{R} \mathcal{D}_{\lambda}} [x_0^* = x_0] \cdot \frac{1}{\mathsf{PDens}(1^{\lambda}, x_0)}$$

again by definition of conditional probability. We can see the final line is the probability density of x_0^*, x''^*, p^* in Game 8.

Claim 20. The distributions output by Game 8 and Game 9 have statistical distance $\frac{1}{3\kappa} + \operatorname{negl}(\lambda)$.

Proof. Game 9 contains 2 changes. One is purely notational, where x'' is eliminated and we directly test if $\mathsf{Pldx}(x_0, 3\kappa) = \mathsf{Pldx}(x', 3\kappa)$. Since x'' was drawn from \mathcal{D}_{λ} and x' is drawn from $\mathsf{Sample}(1^{\lambda})$, the distribution of x' is statistically close from this change. The other change made here is that we restrict the loop to $9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2 \cdot \lambda$ iterations rather than an unbounded number. Note that the loop terminating is entirely dependent on sampling x' such that $\mathsf{Pldx}(x_0, 3\kappa) = \mathsf{Pldx}(x', 3\kappa)$. So we will show that in Game 8, such an x' is found in the first $9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2 \cdot \lambda$ iterations with probability $\leq \frac{1}{3\kappa} + \mathsf{negl}(\lambda)$.

To see this, we first want to observe that the total number of possible values of Pldx is bounded by

$$\log_{1+\frac{1}{3\kappa}}(\mathsf{PDF}_{\max}) = \ln(\mathsf{PDF}_{\max}) / \ln(((1+\frac{1}{3\kappa})^{3\kappa\cdot\frac{1}{3\kappa}}))$$
$$= \ln(\mathsf{PDF}_{\max}) \cdot 3\kappa / \ln((1+\frac{1}{3\kappa})^{3\kappa}) \approx 3\ln(\mathsf{PDF}_{\max}) \cdot \kappa$$

as the PDens function is nonnegative and integral, so Pldx returns \perp or an integer in the range $[0, \ln(\mathsf{PDF}_{\max}) \cdot \kappa)]$. Let the set $A = \{a_1, a_2, \ldots a_q\}$ denote the range of values Pldx can take.

We partition A into sets A_0 and A_1 such that

$$a \in A_0 \Leftrightarrow \Pr\left[\begin{array}{c} x_0 \leftarrow \mathcal{D}_{\lambda} \\ \mathsf{Pldx}(x_0, 3\kappa) = a \end{array}
ight] \leq rac{1}{9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2}$$

and similarly

$$a \in A_1 \Leftrightarrow \Pr\left[\begin{matrix} x_0 \leftarrow \mathsf{Sample}(1^{\lambda}) \\ \mathsf{Pldx}(x_0, 3\kappa) = a \end{matrix} \right] > \frac{1}{9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2}$$

Since A has $3\ln(\mathsf{PDF}_{\max}) \cdot \kappa$ elements, we can bound

$$\Pr\left[x_0 \leftarrow \mathcal{D}_{\lambda}\mathsf{Pldx}(x_0, 3\kappa) \in A_0\right] \le \frac{3\ln(\mathsf{PDF}_{\max}) \cdot \kappa}{9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2} = \frac{1}{3\kappa}$$

If we suppose $\mathsf{Pldx}(x_0, 3\kappa) \in A_1$, then we can see the probability that $9\ln(\mathsf{PDF}_{\max}) \cdot \kappa^2 \cdot \lambda$ fail to find an x' is lower bounded by

$$\left(1-\frac{1}{9\ln(\mathsf{PDF}_{\max})\cdot\kappa^2}\right)^{9\ln(\mathsf{PDF}_{\max})\cdot\kappa^2\cdot\lambda}\approx e^{-\lambda}\in\mathsf{negl}(\lambda)$$

Since we know $\mathsf{Pldx}(x_0, 3\kappa) \in A_1$ with probability at least $1 - \frac{1}{3\kappa}$, we can lower bound the probability that this loop terminates within $9\ln(\mathsf{PDF}_{\max})\cdot\kappa^2\cdot\lambda$ iterations as $1 - \frac{1}{3\kappa} - \mathsf{negl}(\lambda)$, which bounds the statistical distance between Games 8 and 9 with $\frac{1}{3\kappa} + \mathsf{negl}(\lambda)$.

Claim 21. The distributions output by Game 9 and Game 10 have statistical distance $negl(\lambda)$.

Proof. In Game 10, we alter the way which we generate the variables x, p, x_0 . From Lemma 3, the distributions of x_0, p below are statistically close

 $(x_0, p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) \qquad (_, p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda})$ $x_0 \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p)$

Using this, we can see that the following distributions of x_0, x, p are statistically close as well (note that the lefthand side is the distribution generated by Game 9)

 $\begin{array}{ll} (x_0,p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) & (_,p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) \\ x \leftarrow \mathsf{SampleUniform}(1^{\lambda},p) & x_0 \leftarrow \mathsf{SampleUniform}(1^{\lambda},p) \\ x \leftarrow \mathsf{SampleUniform}(1^{\lambda},p) \end{array}$

However, note in the latter distribution, x_0 and x are generated independently, so this is the same as the left side distribution below

 $(_, p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda}) \qquad (x, p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_{\lambda})$ $x \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p) \qquad x_0 \leftarrow \mathsf{SampleUniform}(1^{\lambda}, p).$

Applying Lemma 3 again, we can see the above 2 distributions are statistically close. We can see the final distribution above is the distribution of x_0, x, p in Game 10.

Claim 22. The distributions output by Game 10 and Game 11 have statistical distance $negl(\lambda)$.

 $\begin{array}{ccc} \mathcal{D}_1 & \mathcal{D}_2 & \mathcal{D}_3 \\ (x,p) \xleftarrow{R} \mathsf{PDF}(\mathcal{D}_\lambda) & x \xleftarrow{R} \mathcal{D}_\lambda & x \leftarrow \mathsf{Sample}'(1^\lambda) \\ & p \xleftarrow{R} [0, \mathsf{PDens}(1^\lambda, x)]. & p \xleftarrow{R} [0, \mathsf{PDens}(1^\lambda, x)]. \end{array}$

Proof. By Definition 5, the distribution of (x, p) in \mathcal{D}_1 and \mathcal{D}_2 are identical. From Lemma 4, \mathcal{D}_2 is statistically close to \mathcal{D}_3 . Note \mathcal{D}_1 is the distribution of (x, p) in Game 10 and \mathcal{D}_3 is the distribution of (x, p) in Game 11.

Combining Claim 12 through Claim 22, we get the total statistical distance of Game 0 and 11 is $\leq \frac{1}{\kappa} + \mathsf{negl}(\lambda)$

Combining Claim 11 and Claim 10, we get the explainability of Sample'.

Combining Lemma 4 and Lemma 5, we get that Sample' is an explainable sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$.

5.2 Instantiation on Discrete Gaussians

As an example, we can observe that the above gives us an explainable sampler for discrete Gaussians centered at $c(\lambda)$ with standard deviation $\sigma(\lambda)$. For the Sample algorithm, we can use any of the discrete Gaussian samplers present in literature. We can use $\mathsf{PDens}(1^{\lambda}, x) = e^{-\frac{(x-c)^2}{2\sigma^2}}$, which is efficiently computable, and we can let $\mathsf{SampleUniform}(1^{\lambda}, p)$ simply sample uniformly from integers on the interval $\left[\left[c - \sigma\sqrt{-2\ln(p)}\right], \left[c + \sigma\sqrt{-2\ln(p)}\right]\right]$. Since the only randomness here is sampling uniform integers on a fixed interval, this is easily explainable.

5.3 Impossibility of Generic Sampling without Heavy Element Samplers

We give some evidence on the tightness of the above result by showing an impossibility of a black box construction of an explainable sampler from only a sampler and probability density function to some distribution. This also highlights the inherent need of non-black-box techniques such as indistinguishability obfuscation used to construct universal samplers in [21].

Theorem 23. There exists a distribution family $\{\mathcal{D}_{\lambda}\}_{\lambda}$ such that there does not exist an efficient explainable sampler for $\{\mathcal{D}_{\lambda}\}_{\lambda}$ given oracle access to a sampler Sample and probability density function PDens.

We defer the proof of this theorem to the full version of the paper [27].

6 Explaining Discrete Gaussian Samplers

In this section we show an explanation algorithm for the Miccancio-Walter [30] Gaussian Sampling algorithm. This algorithm follows a series of works improving the practicality of discrete Gaussian sampling [10, 12, 15, 31, 25, 2, 24, 34, 4, 22, 33]. This usually comes in the form of some combination of decreased runtime (in either an offline or online phase), decreased memory usage, and decreased entropy usage. While we could simply apply our transformation of Section 5 to the MW sampler to get explainability, this may not in general preserve special properties of many samplers that could be leveraged, and would add an additional overhead for computing the probability density function. In this section we will show a tailored approach to make it explainable as is.

We begin by giving an adapted definition for explainable discrete Gaussian samplers that explicitly allows for the standard deviation and center of the distribution to be given as parameters. The syntax from the previous sections restricts these to be functions of the security parameter λ . We then describe the explainability algorithm for the Miccancio-Walter discrete Gaussian sampler.

24

Discrete Gaussians

To talk specifically about discrete Gaussians, we will modify need to expand the definition of distribution samplers to accomodate parameterization. The theorems and their proofs shown in the previous sections translate analogously to the definitions here. For completeness, the proofs will be available in the appendices.

Definition 6. We say an algorithm SampleDG $(1^{\lambda}, c, \sigma; r)$ is an $f_{\sigma}(\lambda)$ -discrete Gaussian sampler for a discrete Gaussians if for all $c \in [0, 1]$ and $\sigma \leq f_{\sigma}(\lambda)$, Sample runs in poly (λ) time and the output of SampleDG $(1^{\lambda}, c, \sigma; r)$ is statistically close to the discrete Gaussian on $\mathbb{Z}+c$ centered at 0 with standard deviation σ . Note that this is equivalent to producing discrete Gaussians samples on \mathbb{Z} centered at any $c' \in \exp(\lambda)$, as we can simply generate a sample on $\mathbb{Z}+(-c \mod 1)$ centered at 0 and add c.

Definition 7. Analogously, we say a SampleDG algorithm is explainable if there exists a (possibly randomized) algorithm $\text{ExplainDG}(1^{\lambda}, 1^{\kappa}, x, c, \sigma)$ such that ExplainDG runs in $\text{poly}(\lambda, \kappa)$ time, and there exists a negligible function $\text{negl}(\lambda)$ such that for all $c \in [0, 1]$, $\sigma \leq f_{\sigma}(\lambda)$, the statistical distance between the following 2 distributions is at most $\frac{1}{\kappa} + \text{negl}(\lambda)$.

Game A	Game B
$ \begin{array}{l} - \ r \xleftarrow{R} \{0,1\}^n. \\ - \ x \leftarrow SampleDG(1^{\lambda},c,\sigma;r). \\ - \ \mathrm{Return} \ r, x. \end{array} $	$\begin{array}{l} - x \leftarrow SampleDG(1^{\lambda}, c, \sigma). \\ - r \leftarrow ExplainDG(1^{\lambda}, 1^{\kappa}, x, c, \sigma). \\ - \operatorname{Return} r, x. \end{array}$

While Theorem 7 does imply the existence of an explainable discrete Gaussian sampler, oftentimes, it may be of interest whether existing implementations using particular sampling algorithms are explainable. We'll use this MW17 as an example of a practical discrete Gaussian sampler and prove its explainability, along the way hopefully demonstrating some techniques useful for proving explainability of other sampling algorithms.

6.1 Miccancio-Walter '17

Here, we look at a fairly recent method of sampling discrete Gaussians presented in [30]. Informally, this is done by taking a sampler SampleBase to discrete Gaussians with small standard deviation, for which there are easier and more efficient to compute, then taking particular linear combinations to create discrete Gaussian samples with large standard deviation in Samplel, before employing a randomized rounding technique in SampleC to sample from a specific coset of Z.

In comparison to previous works, this has the advantage of better performance, as well as taking time independent of the sample value. In addition, much of the computation can be done offline, without knowing the standard deviation or center an algorithm which calls the sampler may request. Samplel $(1^{\lambda}, i)$ SampleC $(1^{\lambda}, c, k)$ - If i = 0- If k = 0 $x \leftarrow \mathsf{SampleBase}(1^{\lambda}, 0, s_0)$ • Return 0. $-g \leftarrow 2^{-k+1} \cdot \mathsf{SampleBase}(1^{\lambda}, 2^{k-1}c, s_0)$ • Return x $x_1 \leftarrow \mathsf{Samplel}(1^\lambda, i-1)$ - Return $g + \mathsf{SampleC}(1^{\lambda}, c - g, k - 1)$ $x_2 \leftarrow \mathsf{Samplel}(1^{\lambda}, i-1)$ SampleDG $(1^{\lambda}, c, s)^{a}$ $y = z_i x_1 + \max(1, z_i - 1) x_2$ $-x \leftarrow \mathsf{Samplel}(m)$ - Return y $- K \leftarrow \sqrt{s^2 - \bar{s}_k^2} / s_m$ $-c' \leftarrow \dot{[c+Kx]_k}$ $- y \leftarrow \mathsf{SampleC}(1^{\lambda}, c', k)$ - Return y^{*a*} This was referred to as SampleZ in [30]

We remark that the notation of the above algorithms has been slightly altered to better fit with our definition of a discrete Gaussian sampler. To align with our definition of statistical closeness, we consider $k \in \Theta(\lambda)$. As in the original construction, we let s_i denote the exact standard deviation of Samplel, z_i denote the multipliers used in Samplel, and \bar{s}_k to be the standard deviation 'added' by SampleC, whose explicit values are determined recursively by the equations below. The exact formula is given as a function of the smoothing parameter $\eta_{\epsilon}(\mathbb{Z})$ of the discrete Gaussian, which, for our purposes, we will only bound as a value $\in O(\log(\lambda))$. Similarly, $m = m(\lambda)$ is a parameter controlling the precision controlling the maximum standard deviation SampleDG can generate, which we can think of as a value bounded by $\log(\lambda) + O(1)$. Readers of [30] may notice the lack of a base parameter b which controls the 'base' for (e.g. binary, decimal) which randomized rounding occurs. This has been fixed to 2 for simplicity.

$$s_0 = \sqrt{2}\eta_{\epsilon}(\mathbb{Z}), \quad z_i = \left\lfloor \frac{s_{i-1}}{\sqrt{2}\eta_{\epsilon}(\mathbb{Z})} \right\rfloor$$
$$s_i^2 = \left(z_i^2 + \max((z_i - 1)^2, 1)\right) s_{i-1}^2, \quad \bar{s}_k = s_0 \cdot \left(\sqrt{\sum_{i=0}^{k-1} 2^{-2i}}\right)$$

Theorem 24. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Then SampleDG $(1^{\lambda}, c, s; r)$ is an explainable $\exp(\lambda)$ -discrete Gaussian sampler for the family of discrete Gaussians centered at c with standard deviation s.

Proof. Lemma 6. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Then for $i \leq \log(\lambda) + O(1)$, Samplel $(1^{\lambda}, i)$ is a sampler for the family of discrete Gaussians centered at 0 with standard deviation $\omega(2^{2^{i}})$.

See [30] for proof.

Lemma 7. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Then $y = \text{SampleC}(1^{\lambda}, c, k)$ is a sampler for the distribution of discrete Gaussians on $c + \mathbb{Z}$ with standard deviation \bar{s}_k .

See [30] for proof.

Corollary 25. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Let c' be a sample from a discrete Gaussians on $\mathbb{Z}/2^k$ with standard deviation $\sqrt{\sigma^2 - \bar{s}_k^2}$ and center c. Then for $k \in \Theta(\lambda)$, in SampleC $(1^{\lambda}, c', k)$, the output to the recursive call to SampleC $(1^{\lambda}, c'', i)$ is statistically close to the discrete Gaussian on $\mathbb{Z} + c''$ with standard deviation $\sqrt{\sigma^2 - \bar{s}_i^2}$.

Proof. This follows from a straightforward induction argument on i.

Lemma 8. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Then SampleDG $(1^{\lambda}, c, \sigma)$ is a $\exp(\lambda)$ -discrete Gaussian sampler.

See [30] for proof.

Lemma 9. Any $O(\log(\lambda))$ -discrete Gaussian sampler is explainable.

Proof. Note that we can bound the probability a sample is outside $\log(\lambda)$ standard deviations with $2 \cdot \sum_{i=\log(\lambda)\sigma}^{\infty} e^{-\frac{i^2}{2\sigma}}$ Since this decays exponentially, this is $\leq O(e^{-\log(\lambda)^2}) \in O(\lambda^{-\log(\lambda)})$. Thus, we can take the set S_{λ} to be the set of elements within $\log(\lambda)$ standard deviations of the center. By Corollary 6 from Section 4, this is explainable.

Lemma 10. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Then Samplel $(1^{\lambda}, i)$ is an **explainable** sampler for the family of discrete Gaussians centered at 0 with standard deviation $\omega(2^{2^{i}})$.

We defer the proof of this lemma to the full version of the paper [27].

Lemma 11. Suppose SampleBase $(1^{\lambda}, c)$ is an $O(\log(\lambda))$ -discrete Gaussian sampler. Let SampleDG_{2^k} $(1^{\lambda}, c, \sigma)$ be an explainable sampler for discrete Gaussians on $\mathbb{Z}/2^k$ with standard deviation σ and center c. Then for $k \in O(\lambda)$, SampleC $(1^{\lambda}, c, \sigma) =$ SampleC $(1^{\lambda}, \text{SampleDG}_{2^k}(1^{\lambda}, c, \sqrt{\sigma^2 - \bar{s}_k^2}), k)$ is an explainable discrete Gaussian sampler.

We defer the proof of this lemma to the full version of the paper [27].

Using Lemma 10, we know the Samplel algorithm is an explainable sampler for discrete Gaussians centered at 0 with standard deviation s_m . We know from this that c' is a sample from a discrete Gaussian on $\mathbb{Z}/2^k$ centered at c with standard deviation $\sqrt{s^2 - \overline{s}_k^2}$. Moreover, this is explainable (the only non-reversible step is the rounding to the nearest 2^{-k} , but since $k \in \Theta(\lambda)$, we can pick a uniform preimage to the rounding and still be statistically close). Thus, using Lemma 11, taking SampleDG to be the first 3 lines of SampleDG, we can conclude SampleDG is explainable. Combined with Lemma 8, this gives us the statement of Theorem 24.

References

- Shweta Agrawal, Daniel Wichs, and Shota Yamada. Optimal broadcast encryption from LWE and pairings in the standard model. In Rafael Pass and Krzysztof Pietrzak, editors, Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part I, volume 12550 of Lecture Notes in Computer Science, pages 149–178. Springer, 2020.
- Carlos Aguilar-Melchor, Martin R. Albrecht, and Thomas Ricosset. Sampling from arbitrary centered discrete gaussians for lattice-based cryptography. In Dieter Gollmann, Atsuko Miyaji, and Hiroaki Kikuchi, editors, *Applied Cryptography and Network Security*, pages 3–19, Cham, 2017. Springer International Publishing.
- 3. Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings, volume 7237 of Lecture Notes in Computer Science, pages 483–501. Springer, 2012.
- Gilles Barthe, Sonia Belaïd, Thomas Espitau, Pierre-Alain Fouque, Mélissa Rossi, and Mehdi Tibouchi. Galactics: Gaussian sampling for lattice-based constant- time implementation of cryptographic signatures, revisited. pages 2147–2164, 11 2019.
- Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993, pages 62–73. ACM, 1993.
- Mihir Bellare and Phillip Rogaway. The exact security of digital signatures how to sign with RSA and rabin. In Ueli M. Maurer, editor, Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding, volume 1070 of Lecture Notes in Computer Science, pages 399–416. Springer, 1996.
- Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, pages 213–229, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In Colin Boyd, editor, Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings, volume 2248 of Lecture Notes in Computer Science, pages 514–532. Springer, 2001.
- Zvika Brakerski, David Cash, Rotem Tsabary, and Hoeteck Wee. Targeted homomorphic attribute-based encryption. In Martin Hirt and Adam D. Smith, editors, *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II*, volume 9986 of Lecture Notes in Computer Science, pages 330–360, 2016.
- Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the Forty-Fifth Annual* ACM Symposium on Theory of Computing, STOC '13, page 575–584, New York, NY, USA, 2013. Association for Computing Machinery.
- 11. Eric Brier, Jean-Sébastien Coron, Thomas Icart, David Madore, Hugues Randriam, and Mehdi Tibouchi. Efficient indifferentiable hashing into ordinary elliptic curves.

In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, pages 237–254, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

- 12. Johannes Buchmann, Daniel Cabarcas, Florian Göpfert, Andreas Hülsing, and Patrick Weiden. Discrete ziggurat: A time-memory trade-off for sampling from a gaussian distribution over the integers. In Tanja Lange, Kristin Lauter, and Petr Lisoněk, editors, *Selected Areas in Cryptography – SAC 2013*, pages 402–417, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In Jeffrey Scott Vitter, editor, Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998, pages 209–218. ACM, 1998.
- Pratish Datta, Ilan Komargodski, and Brent Waters. Decentralized multi-authority abe for dnfs from lwe. In Kenneth G. Paterson, editor, Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zoagreb, Croatia, October 17-21, 2021. Proceedings, Lecture Notes in Computer Science. Springer, 2021.
- Nagarjun C. Dwarakanath and Steven D. Galbraith. Sampling from discrete gaussians for lattice-based cryptography on a constrained device. *Appl. Algebra Eng.*, *Commun. Comput.*, 25(3):159–180, June 2014.
- 16. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings, volume 263 of Lecture Notes in Computer Science, pages 186–194. Springer, 1986.
- Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael J. Wiener, editor, Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings, volume 1666 of Lecture Notes in Computer Science, pages 537–554. Springer, 1999.
- Nicholas Genise and Daniele Micciancio. Faster gaussian sampling for trapdoor lattices with arbitrary modulus. In Jesper Buus Nielsen and Vincent Rijmen, editors, Advances in Cryptology – EUROCRYPT 2018, pages 174–203, Cham, 2018. Springer International Publishing.
- Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Proceedings of the fortieth annual ACM symposium on Theory of computing, pages 197–206, 2008.
- Eu-Jin Goh, Stanislaw Jarecki, Jonathan Katz, and Nan Wang. Efficient signature schemes with tight reductions to the diffie-hellman problems. J. Cryptol., 20(4):493–514, 2007.
- 21. Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. How to generate and use universal samplers. In Jung Hee Cheon and Tsuyoshi Takagi, editors, Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part II, volume 10032 of Lecture Notes in Computer Science, pages 715–744, 2016.
- 22. James Howe, Thomas Prest, Thomas Ricosset, and Mélissa Rossi. Isochronous gaussian sampling: From inception to implementation. In Jintai Ding and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography*, pages 53–71, Cham, 2020. Springer International Publishing.
- 23. Thomas Icart. How to hash into elliptic curves. In Shai Halevi, editor, Advances in Cryptology CRYPTO 2009, 29th Annual International Cryptology Conference,

Santa Barbara, CA, USA, August 16-20, 2009. Proceedings, volume 5677 of Lecture Notes in Computer Science, pages 303–316. Springer, 2009.

- Angshuman Karmakar, Sujoy Sinha Roy, Oscar Reparaz, Frederik Vercauteren, and Ingrid Verbauwhede. Constant-time discrete gaussian sampling. *IEEE Trans*actions on Computers, 67(11):1561–1571, 2018.
- 25. Charles F. F. Karney. Sampling exactly from the normal distribution. ACM Transactions on Mathematical Software, 42(1):1–14, Mar 2016.
- 26. Allison B. Lewko and Brent Waters. Decentralizing attribute-based encryption. In Kenneth G. Paterson, editor, Advances in Cryptology - EUROCRYPT 2011 -30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings, volume 6632 of Lecture Notes in Computer Science, pages 568–588. Springer, 2011.
- 27. George Lu and Brent Waters. How to sample a discrete gaussian (and more) from a random oracle. Cryptology ePrint Archive, Paper 2022/1227, 2022. https://eprint.iacr.org/2022/1227.
- Vadim Lyubashevsky and Daniel Wichs. Simple lattice trapdoor sampling from a broad class of distributions. In Jonathan Katz, editor, Public-Key Cryptography -PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings, volume 9020 of Lecture Notes in Computer Science, pages 716–730. Springer, 2015.
- 29. Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *Theory of Cryptography*, pages 21–39, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- 30. Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II, volume 10402 of Lecture Notes in Computer Science, pages 455–485. Springer, 2017.
- 31. Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In Lejla Batina and Matthew Robshaw, editors, Cryptographic Hardware and Embedded Systems CHES 2014 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings, volume 8731 of Lecture Notes in Computer Science, pages 353–370. Springer, 2014.
- 32. Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework. In *Proceedings of the 30th Annual International Conference on Theory and Applications of Cryptographic Techniques: Advances in Cryptology*, EUROCRYPT'11, page 487–506, Berlin, Heidelberg, 2011. Springer-Verlag.
- 33. Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. Cosac: compact and scalable arbitrary-centered discrete gaussian sampling over integers. In Jintai Ding and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography*, Lecture Notes in Computer Science, pages 284–303. Springer, 2020. 11th International Conference on Post-Quantum Cryptography, PQCrypto 2020; Conference date: 15-04-2020 Through 17-04-2020.
- Raymond K. Zhao, Ron Steinfeld, and Amin Sakzad. Facct: Fast, compact, and constant-time discrete gaussian sampler over integers. *IEEE Transactions on Computers*, 69(1):126–137, 2020.