Learning with Rounding, Revisited

New Reduction, Properties and Applications

Joël Alwen[1]*, Stephan Krenn[2], Krzysztof Pietrzak[3]** Daniel Wichs[4]***

[1] ETH Zurich
alwenj@inf.ethz.ch
[2] IBM Research – Zurich, Rüschlikon
skr@zurich.ibm.com
[3] Institute of Science and Technology Austria
pietrzak@ist.ac.at
[4] Northeastern University
wichs@ccs.neu.edu

Abstract. The learning with rounding (LWR) problem, introduced by Banerjee, Peikert and Rosen [BPR12] at EUROCRYPT '12, is a variant of learning with errors (LWE), where one replaces random errors with deterministic rounding. The LWR problem was shown to be as hard as LWE for a setting of parameters where the modulus and modulus-to-error ratio are super-polynomial. In this work we resolve the main open problem of [BPR12] and give a new reduction that works for a larger range of parameters, allowing for a polynomial modulus and modulus-to-error ratio. In particular, a smaller modulus gives us greater efficiency, and a smaller modulus-to-error ratio gives us greater security, which now follows from the worst-case hardness of GapSVP with polynomial (rather than super-polynomial) approximation factors.

As a tool in the reduction, we show that there is a "lossy mode" for the LWR problem, in which LWR samples only reveal partial information about the secret. This property gives us several interesting new applications, including a proof that LWR remains secure with weakly random secrets of sufficient min-entropy, and very simple new constructions of deterministic encryption, lossy trapdoor functions and reusable extractors.

Our approach is inspired by a technique of Goldwasser et al. [GKPV10] from ICS '10, which implicitly showed the existence of a "lossy mode" for LWE. By refining this technique, we also improve on the parameters of that work to only requiring a polynomial (instead of super-polynomial) modulus and modulus-to-error ratio.

Keywords: Learning with Errors, Learning with Rounding, Lossy Trapdoor Functions, Deterministic Encryption.

1 Introduction

Learning With Errors. The Learning with Errors (LWE) assumption states that "noisy" inner products of a secret vector with random public vectors, look pseudorandom. In the last years many cryptosystems have been constructed whose security can be proven under LWE, including (identity-based, leakage-resilient, fully homomorphic, functional) encryption [Reg05,GPV08,AGV09,LPR10,AFV11,BV11,LP11,GKP+12], oblivious transfer [PVW08], (blind) signatures [GPV08,Lyu09,Rüc10,Lyu12], pseudo-random functions [BPR12], hash functions [KV09,PR06],etc.

* This research was conducted in part while the author was visiting IST Austria.

** Supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Starting Grant (259668-PSPC)

*** This research was conducted in part while the author was visiting IST Austria and at IBM Research, T.J. Watson.

The LWE assumption, with parameters $n, m, q \in \mathbb{N}$ and a "small" error distribution χ over $\mathbb{Z}$, states that for uniformly random $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$ and an error vector $\mathbf{e} \leftarrow \chi^m$

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) \text{ is computationally indistinguishable from } (\mathbf{A}, \mathbf{u}).$$

Sometimes it will be convenient to think of this distribution as consisting of m "LWE samples" of the form $(\mathbf{a}_i, \langle \mathbf{a}_i, s\rangle + e_i) \in \mathbb{Z}_q^{n+1}$. One of the main advantages of the LWE problem is that, for some settings of parameters, we can prove its security under certain worst-case hardness assumptions over lattices; see [Reg05,Pei09]. One important parameter is the "size" of the error terms $e \xleftarrow{\$} \chi$ which we denote by β.[5] As long as β exceeds some minimum threshold $\approx \sqrt{n}$, the concrete hardness of the LWE problem mainly depends on the dimension n and on the ratio of the modulus q to the error-size β. Therefore, we will often be unspecific about the exact distribution χ, and only focus on the error-size β.

Learning With Rounding. The Learning with Rounding (LWR) problem was introduced in [BPR12]. Instead of adding a random small error to various samples $\langle \mathbf{a}, \mathbf{s} \rangle \in \mathbb{Z}_q$ so as to hide their exact value, we release a *deterministically rounded* version of $\langle \mathbf{a}, \mathbf{s} \rangle$. In particular, for some $p < q$, we divide up the elements of $\mathbb{Z}_q$ into p contiguous intervals of roughly q/p elements each and define the *rounding function* $\lfloor \cdot \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$ that maps $x \in \mathbb{Z}_q$ into the index of the interval that x belongs to. For example if q, p are both powers of 2, than this could correspond to outputting the $\log(p)$ most significant bits of x. We can extend the rounding function to vectors by applying it componentwise. The LWR assumption states that:

$$(\mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p) \text{ is computationally indistinguishable from } (\mathbf{A}, \lfloor \mathbf{u} \rceil_p).$$

If p divides q, than $\lfloor \mathbf{u} \rceil_p$ is itself uniform over $\mathbb{Z}_p^m$. The main advantage of LWR is that one does not need to sample any additional "errors", therefore requiring fewer random bits. The assumption has been used to construct simple and efficient pseudorandom generators and functions in [BPR12], and deterministic encryption in [XXZ12].

The work of [BPR12] shows a beautifully simple reduction proving the hardness of the LWR problem under the LWE assumption for some range of parameters. In particular, they observe that if the error size β is sufficiently small and the ratio q/p is sufficiently big, then $\lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rceil_p = \lfloor \langle \mathbf{a}, \mathbf{s} \rangle + e \rceil_p$ with overwhelming probability over random $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q$ and $e \xleftarrow{\$} \chi$. In particular, the only way that the two values differ is if $\langle \mathbf{a}, \mathbf{s} \rangle$ ends up within a distance of $|e|$ from a boundary between two different intervals; but since the intervals are of size q/p and the ball around the boundary is only of size $2|e|$ this is unlikely to happen when q/p is super-polynomially bigger than $2|e|$. Therefore, one can show that:

$$(\mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p) \overset{\text{stat}}{\approx} (\mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \rceil_p) \overset{\text{comp}}{\approx} (\mathbf{A}, \lfloor \mathbf{u} \rceil_p)$$

where the first modification is statistically close and the second follows immediately from the hardness of LWE.

Unfortunately, the argument only goes through, when (q/p) is bigger than the error size β by a super-polynomial factor. In fact, if we want statistical distance $2^{-\lambda}$ we would need to set $q \geq 2^{\lambda} \beta p$, where λ is a security parameter. This has three important consequences: (1) the modulus q has to be super-polynomial, which makes all of the computations less efficient, (2) the modulus-to-error ratio

[5] We will be informal for now; we can think of β as the the standard deviation or the expected/largest absolute value of the errors.

q/β is super-polynomial which makes the LWE problem easier and only gives us a reduction if we assume the hardness of the lattice problem GapSVP with super-polynomial approximation factors (a stronger assumption), (3) the ratio of the input-to-output modulus q/p is super-polynomial, meaning that we must "throw away" a lot of information when rounding and therefore get fewer bits of output per LWR sample. The work of [BPR12] conjectured that the LWR problem should be hard even for a polynomial modulus q, but left it as the main open problem to give a reduction.

1.1 The New Reduction and Properties of LWR

LWR with Polynomial Modulus. In this work, we resolve the open problem of [BPR12] and give a new reduction showing the hardness of LWR from that of LWE for a more general setting of parameters, including when the modulus q is only polynomial. In particular, instead of requiring $q \geq 2^{\lambda}\beta p$, where λ is a security parameter as in [BPR12], we only require $q \geq nm\beta p$, where we recall that n is the dimension of the secret, and m is the number of LWR samples that we output, β is the size of the LWE errors, and p is the new modulus we round to. In particular, as long as the number of LWR samples m is fixed a-priori by some arbitrary polynomial, we can allow the modulus q to be polynomial. As mentioned, this setting provides greater efficiency (computation with smaller q) and greater security (smaller ratio q/β) allowing for a reduction from the worst-case hardness of the lattice problem GapSVP with polynomial approximation factors. In particular, the above efficiency and security improvements for LWR directly translate into improvements of the PRG and the basic synthesizer-based PRF constructions of [BPR12] (but not to their optimized "degree-k synthesizer" construction, which doesn't directly rely on LWR, and requires a super-polynomial modulus for different reasons).

To be even more precise, our reduction shows the hardness of LWR with parameters n, m, q, p assuming the hardness of LWE with parameters n', m, q, β (note: different dimension n' vs. n) as long as:

$$n \geq \frac{\log(q)}{\log(2\gamma)} \cdot n' \qquad \text{and} \qquad q \geq \gamma(nm\beta p) \tag{1}$$

for some flexible parameter $\gamma \geq 1$. For example, setting $\gamma = 1$ allows for the smallest modulus $q \approx nm\beta p$, but requires a larger dimension $n \approx n' \log(q)$ in the LWR problem than the dimension n' of the underlying LWE assumption. On the other hand, setting $\gamma = q^{\delta}$ for some constant $\delta \in (0, 1)$ gives a bigger polynomial modulus $q \approx (nm\beta p)^{1/(1-\delta)}$ but allow us to set the LWR dimension $n \approx (1/\delta)n' = O(n')$ to be closer to that of the underlying LWE assumption.

Note that, for our reduction to work, the modulus q must always be sufficiently larger than the number of LWR samples m. Therefore, we can only set q to polynomial in a setting where the total number of samples m given out is known ahead of time. This is the case for all of the applications of [BPR12] as well as the new applications that we will describe in this work. In settings where m is not known ahead of time (e.g., the attacker can decide how many samples it will get) we would need to make the modulus q and the modulus-to-error ratio q/β super-polynomial, but our reduction still provides tighter exact security than that of [BPR12]. It remains as an interesting open problem to improve the parameters of the reduction further, and especially to remove the dependence between the modulus q and the number of LWR samples m that we give out.

LWR with Weak and Leaky Secrets. Another advantage of our reduction is that we prove the security of the LWR problem even when the secret $\mathbf{s}$ is not necessarily uniform over $\mathbb{Z}_q^n$. Indeed, our proof also works when $\mathbf{s}$ is uniform over a smaller integer interval $\mathbf{s} \xleftarrow{\$} \{-\gamma, \ldots, \gamma\}^n \subseteq \mathbb{Z}_q^n$, where the relation of $\gamma \geq 1$ to the other parameters is given by equation (1). Moreover, our reduction works when the secret $\mathbf{s}$ is not even truly uniform over this interval (say, because the attacker observed some leakage on $\mathbf{s}$, or $\mathbf{s}$ was sampled using a weak random source) as long as $\mathbf{s}$ retains some sufficiently high amount of *min-entropy* $k \approx n' \log(q)$, where n' is the dimension of the underlying LWE assumption. Notice that, no matter how small the entropy k is, we can still prove security under some LWE assumption with correspondingly smaller dimension n'.

The work of Goldwasser et al. [GKPV10] shows similar results for the hardness of LW$\underline{\text{E}}$ with a weak and leaky secret, at least as long as the modulus q and the modulus-to-error ratio q/β are super-polynomial. Indeed, we will use a refinement of the technique from their work as the basis of our LWE to LWR reduction. Our refinement will also allow us to improve the parameters of [GKPV10], and show the hardness of LWE with a weak and leaky secret when the modulus q and the ratio q/β are polynomial.

The Reduction. As discussed above, the original reduction of [BPR12] required us to choose parameters so that rounded samples with and without error are almost always identical: $\Pr[\lfloor\langle \mathbf{a}, \mathbf{s}\rangle\rfloor_p \neq \lfloor\langle \mathbf{a}, \mathbf{s}\rangle + e\rfloor_p] \leq \mathsf{negl}$. Therefore LWR outputs do not provide any more information than LWE outputs. In contrast, in our setting of parameters, when q is polynomial, there is a noticeable probability that the two values are different and therefore we will need a completely different proof strategy.

Surprisingly, our strategy does *not* try to directly convert an LWE instance with a secret $\mathbf{s}$ into an LWR instance with secret $\mathbf{s}$. Instead, we will rely on the LWE problem to change the distribution of the coefficient matrix $\mathbf{A}$. In particular, we show that there is a "lossy" method of sampling a matrix $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}()$ such that:

(a) Under the LWE assumption, $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}()$ is computationally indistinguishable from $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m\times n}$.

(b) When $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}()$, the values $\tilde{\mathbf{A}}, \lfloor\tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p$ do not reveal too much information about $\mathbf{s}$. In particular, $\mathbf{s}$ maintains a large fraction of its statistical entropy given $\tilde{\mathbf{A}}, \lfloor\tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p$.

Before we describe how the $\mathsf{Lossy}()$ sampler works in the next paragraph, let us show that the above two properties allow us to prove the hardness of LWR problem. We can do so via a hybrid argument where, given many LWR samples, we replace one sample at a time from being an LWR sample to being uniformly random. In particular, assume we have $m+1$ LWR samples and let the matrix $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m\times n}$ denote the coefficient vectors of the first m samples, and let $\mathbf{a} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$ be the coefficient vector of the last sample. Then we can show:

$$\left(\begin{bmatrix}\mathbf{A}\\ \mathbf{a}\end{bmatrix}, \begin{bmatrix}\lfloor\mathbf{A}\cdot\mathbf{s}\rfloor_p\\ \lfloor\langle\mathbf{a},\mathbf{s}\rangle\rfloor_p\end{bmatrix}\right) \overset{\text{comp}}{\approx} \left(\begin{bmatrix}\tilde{\mathbf{A}}\\ \mathbf{a}\end{bmatrix}, \begin{bmatrix}\lfloor\tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p\\ \lfloor\langle\mathbf{a},\mathbf{s}\rangle\rfloor_p\end{bmatrix}\right) \overset{\text{stat}}{\approx} \left(\begin{bmatrix}\tilde{\mathbf{A}}\\ \mathbf{a}\end{bmatrix}, \begin{bmatrix}\lfloor\tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p\\ \lfloor u\rfloor_p\end{bmatrix}\right) \overset{\text{comp}}{\approx} \left(\begin{bmatrix}\mathbf{A}\\ \mathbf{a}\end{bmatrix}, \begin{bmatrix}\lfloor\mathbf{A}\cdot\mathbf{s}\rfloor_p\\ \lfloor u\rfloor_p\end{bmatrix}\right)$$

In the first step, we use the LWE assumption to replace a uniformly random $\mathbf{A}$ by a lossy matrix $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}()$, but still choose the last row $\mathbf{a} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$ at random. In the second step, we use the fact that *inner product* is a strong extractor, where we think of the secret $\mathbf{s}$ as the source and the vector $\mathbf{a}$ as a seed. In particular, by the properties of the lossy sampler, we know that $\mathbf{s}$ maintains entropy conditioned on seeing $\tilde{\mathbf{A}}, \lfloor\tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p$ and therefore the "extracted value" $\langle\mathbf{a},\mathbf{s}\rangle$ is statistically close to a uniformly random and independent $u \overset{\$}{\leftarrow} \mathbb{Z}_q$. In the last step, we simply replace the lossy matrix $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}()$ back by a uniformly random $\mathbf{A}$. This shows that, given the first m LWR samples the last one looks uniform and independent. We can then repeat the above steps m more times to replace each of the remaining LWR samples (rows) by uniform, one-by-one.

The Lossy Sampler. The basic idea of our Lossy sampler is taken from the work of Goldwasser et al. [GKPV10]. We sample the lossy matrix $\tilde{\mathbf{A}} \in \mathbb{Z}_q^{m\times n}$ as

$$\tilde{\mathbf{A}} \overset{\text{def}}{=} \mathbf{BC} + \mathbf{F} \qquad \text{where } \mathbf{B} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m\times n'} \;\;,\;\; \mathbf{C} \overset{\$}{\leftarrow} \mathbb{Z}_q^{n'\times n} \;\;,\;\; \mathbf{F} \overset{\$}{\leftarrow} \chi^{m\times n}$$

where $n' < n$ is some parameter and χ is a "small" LWE error distribution. We now need to show that this satisfies the properties (a) and (b) described above.

It is easy to see that $\tilde{\mathbf{A}}$ is computationally indistinguishable from a uniformly random matrix under the LWE assumption with parameters n', m, q, χ. In particular, each column i of the matrix $\tilde{\mathbf{A}}$ can be thought of as an LWE distribution $\mathbf{B}\cdot\mathbf{c}_i + \mathbf{f}_i$ with coefficient matrix $\mathbf{B}$, secret $\mathbf{c}_i$ which is the ith column of the matrix $\mathbf{C}$, and error vector $\mathbf{f}_i$ which is the ith column of $\mathbf{F}$. Therefore, using n hybrid arguments, we can replace each column i of $\tilde{\mathbf{A}}$ by a uniformly random and independent one. This part of the argument is the same as in [GKPV10].

Next, we need to show that the secret $\mathbf{s}$ retains entropy even conditioned on seeing $\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p$. Let us first prove this property in the case when $\mathbf{s} \overset{\$}{\leftarrow} \{-1,0,1\}^n$ is itself a random "short" vector.[6] All of the information that we give out about $\mathbf{s}$ can be reconstructed from:

- The matrices $\mathbf{B}, \mathbf{C}, \mathbf{F}$ which define $\tilde{\mathbf{A}} = \mathbf{BC}+\mathbf{F}$ and are independent of $\mathbf{s}$ on their own.
- The value $\mathbf{C}\cdot\mathbf{s}$ whose bit-length is $n'\log(q)$.
- A set Z consisting of all pairs $(i, v_i) \in [m]\times\mathbb{Z}_p$ such that $\lfloor(\mathbf{BC}\cdot\mathbf{s})_i\rfloor_p \neq \lfloor(\tilde{\mathbf{A}}\cdot\mathbf{s})_i\rfloor_p$ along with the value $v_i = \lfloor(\tilde{\mathbf{A}}\cdot\mathbf{s})_i\rfloor_p$. The subscript i denotes the ith component of a vector.

Given the three pieces of information above, we can reconstruct $\tilde{\mathbf{A}}, \lfloor\tilde{\mathbf{A}}\cdot\mathbf{s}\rfloor_p$ by setting $\lfloor(\tilde{\mathbf{A}}\cdot\mathbf{s})_i\rfloor_p := \lfloor(\mathbf{BC}\cdot\mathbf{s})_i\rfloor_p$ for every index i not contained in Z, and setting $\lfloor(\tilde{\mathbf{A}}\cdot\mathbf{s})_i\rfloor_p := v_i$ for every i which is in Z. Therefore, we just need to show that the three pieces of information above do not reveal too much about $\mathbf{s}$. First, we show that the set Z is small with overwhelming probability. In particular, an index i is contained in Z if and only if

$$\lfloor(\mathbf{BC}\cdot\mathbf{s})_i\rfloor_p \neq \lfloor(\mathbf{BC}\cdot\mathbf{s})_i + (\mathbf{F}\cdot\mathbf{s})_i\rfloor_p. \tag{2}$$

Assume that the entries of the error matrix $\mathbf{F}$ are all bounded by β in absolute value with overwhelming probability, and therefore $(\mathbf{F}\cdot\mathbf{s})_i$ is bounded by $n\beta$ in absolute value.[7] Then the event (2) can only occur if the value $(\mathbf{BC}\cdot\mathbf{s})_i$ falls within distance $n\beta$ of a boundary between two different intervals. Since each interval is of size $\approx q/p$ and the ball around each boundary is of size $2n\beta$, this happens with (noticeable but small) probability $\leq 2n\beta p/q \leq 1/m$, when $q \geq 2nm\beta p$ (which gives us the bound of (1)). Therefore, the probability of any index i being in Z is at most $1/m$, the expected size of Z is at most 1, and because these probabilities are independent, we can use Chernoff to bound $|Z| \leq n'$ with overwhelming probability $1 - 2^{-n'}$. So in total, Z can be described by $|Z|(\log m + \log p) \leq n'\log q$ bits with overwhelming probability. Therefore, together, $Z, \mathbf{Cs}$ reveal only $O(n'\log q)$ bits of information about $\mathbf{s}$, even given $\mathbf{B}, \mathbf{C}, \mathbf{F}$. We can summarize the above as:

$$H_\infty(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor\tilde{\mathbf{A}}\mathbf{s}\rfloor_p) \geq H_\infty(\mathbf{s}|\mathbf{B},\mathbf{C},\mathbf{F},\mathbf{C}\cdot\mathbf{s}, Z) \geq H_\infty(\mathbf{s}|\mathbf{B},\mathbf{C},\mathbf{F}) - O(n'\log q) \geq n - O(n'\log q).$$

Hence, if n is sufficiently larger than some $O(n'\log q)$, the LWR secret maintains a large amount of entropy given the LWR samples with a lossy $\tilde{\mathbf{A}}$. The above analysis also extends to the case where $\mathbf{s}$ is not uniformly random, but only has a sufficient amount of entropy.

[6] This proof generalizes to larger intervals $\{-\gamma,\ldots,\gamma\}$ and corresponds to the parameter γ in equation (1). Here we set $\gamma = 1$.

[7] Our actual proof is more refined and only requires us to bound the *expected* absolute value of the entries.

We can also extend the above analysis to the case where $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ is uniformly random over the entire space (and not short), by thinking of $\mathbf{s} = \mathbf{s}_1 + \mathbf{s}_2$ where $\mathbf{s}_1 \xleftarrow{\$} \mathbb{Z}_q^n$ is uniformly random and $\mathbf{s}_2 \xleftarrow{\$} \{-1, 0, 1\}^n$ is random and short. Using the same argument as above, we can show that, even given $\mathbf{s}_1$, $\tilde{\mathbf{A}}$ and $\lfloor \tilde{\mathbf{A}} \cdot \mathbf{s} \rfloor_p$, the value $\mathbf{s}_2$ (and therefore also $\mathbf{s}$) maintains entropy.

Our analysis of lossiness as described above is inspired by [GKPV10] but differs from it significantly. In particular that work considered LW$\underline{\text{E}}$ (not LWR) samples with the matrix $\tilde{\mathbf{A}}$, didn't explicitly analyze lossiness, and required super-polynomial modulus and modulus-to-error ratio. Indeed, we can use the ideas from the above analysis to also improve the parameters of that work, showing the robustness of the LW$\underline{\text{E}}$ problem to weak and leaky secrets for a polynomial modulus and modulus-to-error ratio. See Appendix B.

1.2 Applications

Reusable Computational Extractor. Recall that, by the leftover-hash lemma, the function $\mathsf{Ext}(\mathbf{s}; \mathbf{a}) := \langle \mathbf{s}, \mathbf{a} \rangle$ is a good randomness extractor which can take any (secret) source $\mathbf{s} \in \mathbb{Z}_q^n$ of sufficient min-entropy $k \geq \log(q) + 2\log(1/\varepsilon)$ and a random public seed $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n$, and its output will be ε-close to uniformly random in $\mathbb{Z}_q$. But let's say we want to extract many different mutually (pseudo-)random values from the source $\mathbf{s}$ without keeping any long term state: each time we want to extract a new output we choose a fresh seed and apply the extractor. It's easy to see that the above inner-product extractor is completely insecure after at most n applications, and each successive output is easy to predict from the previous ones. The work of [DKL09] introduced the notion of a *reusable computational extractor* that remains secure even after m applications, where m can be an arbitrary polynomial, and gave a construction under a non-standard "learning-subspaces with noise" assumption. Our results immediately give us a new simple construction of reusable extractors defined by $\mathsf{Ext}(\mathbf{s}; \mathbf{a}) := \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rfloor_p$. That is, we just round the output of the standard inner product extractor! We show that, as long as the LWE assumption holds with some parameters n', m, q, β, the source $\mathbf{s}$ is distributed over $\{0, 1\}^n$ and has entropy $k \geq O(n' \log(q))$, and the modulus satisfies $q \geq 2\beta nmp$, the above extractor is secure for m uses. In particular, we can have $m \gg n \gg k$.

Lossy Trapdoor Functions. Lossy trapdoor functions (LTDFs) [PW08,PW11] are a family of functions $f_{pk}(\cdot)$ keyed by some public key pk, which can be sampled in one of two indistinguishable modes: `injective` and `lossy`. In the `injective` mode the function $f_{pk}(\cdot)$ is an injective function and we can even sample pk along with a secret trapdoor key sk that allows us to invert it efficiently. In the `lossy` mode, the function $f_{pk}(\cdot)$ is "many-to-one" and $f_{pk}(\mathbf{s})$ statistically loses information about the input $\mathbf{s}$. LTDFs have many amazing applications in cryptography, such as allowing us to output many hardcore bits, construct CCA-2 public-key encryption [PW11,MY10], and deterministic encryption [FOR12]. We construct very simple and efficient LTDFs using the LWR problem: the public key is a matrix $pk = \mathbf{A}$ and the function is defined as $f_{\mathbf{A}}(\mathbf{s}) = \lfloor \mathbf{A} \cdot \mathbf{s} \rfloor_p$. We can sample an injective $\mathbf{A}$ with a trapdoor using the techniques of Ajtai [Ajt99] or subsequent improvements [AP11,MP12], and one can sample a lossy $\mathbf{A}$ using our lossy sampler. Although prior constructions of LTDFs based on LWE are known [PW11,BKPW12], our construction is extremely simple to describe and implement and has the advantage that our lossy mode loses "almost all" of the information contained in $\mathbf{s}$. We also construct very simple and efficient "all-but-one" (ABO) lossy trapdoor functions based on LWR, which are useful for building efficient CCA-2 encryption.

Deterministic Encryption. Deterministic encryption [BBO07,BFOR08,BFO08,BS11,FOR12] is intended to guarantee security as long as the messages have sufficient entropy. Although there are black-box constructions of deterministic encryption using LTDFs [BFO08], we get a very simple direct construction from the LWR problem: the public key is a matrix $pk = \mathbf{A} \in \mathbb{Z}_q^{m\times n}$, and to encrypt a message $\mathbf{s} \in \{0,1\}^n$, we simply output $\lfloor \mathbf{A}\cdot\mathbf{s}\rfloor_p$. We can sample $\mathbf{A}$ with a decryption trapdoor using the standard techniques [Ajt99,AP11,MP12] mentioned previously. Our analysis here is essentially the same as for our reusable extractor – we simply note that whenever $\mathbf{s}$ has sufficient entropy, the output $\lfloor \mathbf{A}\cdot\mathbf{s}\rfloor_p$ is pseudorandom. We note that the same construction was proposed by Xie et al. [XXZ12], but because the analysis there was similar to [BPR12,GKPV10], they required a super-polynomial modulus and modulus-to-error ratio. The main advantage of this scheme over other deterministic encryption schemes is that we do not need any fixed threshold on the entropy of the message $\mathbf{s}$: no matter how low it is we can still prove security under an LWE assumption with correspondingly degraded parameters.

2 Preliminaries

Notation. Throughout, we let λ denote the *security parameter*. We use bold lower-case letters (e.g., $\mathbf{s},\mathbf{e}$) to denote vectors, and bold upper-case letters (e.g., $\mathbf{A},\mathbf{B}$) to denote matrices. If X is a distribution or a random variable, we write $x \overset{\$}{\leftarrow} X$ to denote the process of sampling x according to X. If X is a set, we write $x \overset{\$}{\leftarrow} X$ to denote the process of sampling x *uniformly* at random over X. For two distribution ensembles $X = \{X_\lambda\}, Y = \{Y_\lambda\}$, we write $X \overset{\text{comp}}{\approx} Y$ if for all probabilistic polynomial time (PPT) distinguishers D there is a negligible function $\mathsf{negl}(\cdot)$ such that: $|\Pr[D(1^\lambda, X_\lambda) = 1] - \Pr[D(1^\lambda, Y_\lambda)] = 1| \le \mathsf{negl}(\lambda)$.

Bounded Distribution. For a distribution χ over the reals, and a bound β, we say that χ is *β-bounded* if the average absolute value of $x \overset{\$}{\leftarrow} \chi$ is less then β, i.e., if $\mathbb{E}[|\chi|] \le \beta$.

Probabilistic Notions. We recap some definitions and results from probability theory.

Definition 2.1 (Statistical Distance). *Let X, Y be random variables with supports S_X, S_Y, respectively. We define their* statistical difference *as* $\Delta(X,Y) = \frac{1}{2}\sum_{u\in S_X\cup S_Y} |\Pr[X=u] - \Pr[Y=u]|$. *We write $X \overset{\text{stat}}{\approx} Y$ to denote that $\Delta(X,Y)$ is negligible in the security parameter.*

The *min-entropy* of a random variable X is $H_\infty(X) \overset{\text{def}}{=} -\log(\max_x \Pr[X=x])$, and measures the "best guess" for X. The *conditional min-entropy* of X given Z, defined by Dodis et al. [DORS08], is $H_\infty(X|Z) \overset{\text{def}}{=} -\log\left(\mathbb{E}_{z\leftarrow Z}\left[\ \max_x \Pr[X=x|Z=z]\ \right]\right) = -\log\left(\mathbb{E}_{z\leftarrow Z}\left[2^{-H_\infty(X|Z=z)}\right]\right)$. This measures the "best guess" for X by an adversary that observes a correlated variable Z. That is, for all (potentially inefficient) functions $\mathcal{A}$, we have $\Pr[\mathcal{A}(Z) = X] \le 2^{-H_\infty(X|Z)}$ and there exists some $\mathcal{A}$ which achieves equality. The following lemma says that conditioning on ℓ bits of information, the min-entropy drops by at most ℓ bits.

Lemma 2.2 ([DORS08]). *Let X, Y, Z be arbitrary (correlated) random variables where the support of Z is of size at most 2^ℓ. Then $H_\infty(X|Y,Z) \ge H_\infty(X|Y) - \ell$.*

We will rely on the following definition of *smooth min-entropy*, which was first introduced by Renner and Wolf [RW04]. Intuitively, a random variable has high smooth min-entropy, if it is statistically close to a random variable with high min-entropy.

Definition 2.3 (Smooth Entropy). *We say that a random variable X has* ε-smooth min-entropy at least k, *denoted by $H_\infty^\varepsilon(X) \geq k$, if there exists some variable X' such that $\Delta(X, X') \leq \varepsilon$ and $H_\infty(X') \geq k$. Similarly, we say that the* ε-smooth conditional min-entropy of X given Y is at least k, *denoted $H_\infty^\varepsilon(X|Y) \geq k$ if there exist some variables (X', Y') such that $\Delta((X, Y), (X', Y')) \leq \varepsilon$ and $H_\infty(X'|Y') \geq k$.*

In the remainder of this paper, we will write $H_\infty^{\mathsf{smooth}}(\cdot)$ to denote $H_\infty^\varepsilon(\cdot)$ for some (unspecified) negligible ε. We also prove a variant of Lemma 2.2 for smooth min entropy, which works when Z takes on at most 2^ℓ values with overwhelming probability.

Lemma 2.4. *Let X, Y, Z be correlated random variables and $\mathcal{Z}$ be some set such that* $\Pr[Z \in \mathcal{Z}] \geq 1 - \varepsilon$ *and $|\mathcal{Z}| \leq 2^\ell$. Then, for any $\varepsilon' > 0$, $H_\infty^{\varepsilon+\varepsilon'}(X|Y, Z) \geq H_\infty^{\varepsilon'}(X|Y) - \ell$.*

Proof. Let $k = H_\infty^{\varepsilon'}(X|Y)$ and let (X', Y') be the random variables such that $\Delta((X, Y), (X', Y')) \leq \varepsilon'$ and $H_\infty(X'|Y') = k$. Let $f(x, y)$ be a randomized function which outputs a sample from the conditional distribution $(Z \mid X = x, Y = y)$ or outputs $\perp$ if $\Pr[X = x, Y = y] = 0$. In other words the joint distribution of (X, Y, Z) is the same as $(X, Y, f(X, Y))$. Let $g(z)$ be a function with range $\mathcal{Z}$ which outputs the input z if $z \in \mathcal{Z}$ and else outputs some fixed element $z_{fxd} \in \mathcal{Z}$. Then $\Delta((X, Y, Z), (X, Y, g(f(X, Y))) \leq \varepsilon$ and $\Delta((X, Y, g(f(X, Y)), (X', Y', g(f(X', Y'))) \leq \varepsilon'$. Therefore:

$$H_\infty^{\varepsilon+\varepsilon'}(X|Y, Z) \geq H_\infty(X'|Y', g(f(X', Y'))) \geq H_\infty(X'|Y') - \ell \geq k - \ell.$$

where the second inequality follows from Lemma 2.2. □

We will use the following version of the Chernoff bound from [MU05].

Lemma 2.5. *Let $X_1, \ldots, X_n$ be independent Poisson trials, and let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. Then, for $R \geq 6\mu$, the following Chernoff bound holds:* $\Pr[X \geq R] \leq 2^{-R}$.

2.1 Learning with Errors and Learning with Rounding

Learning With Errors (LWE). The *decisional learning with errors* (*LWE*) problem was first introduced by Regev [Reg05]. Informally, the problem asks to distinguish slightly perturbed random linear equations from truly random ones. (Since we will only talk about the decisional problem in this work, we will make this the default notion.)

Definition 2.6 (LWE Assumption [Reg05]). *Let λ be the security parameter, $n = n(\lambda), m = m(\lambda), q = q(\lambda)$ be integers and let $\chi = \chi(\lambda)$ be a distribution over $\mathbb{Z}_q$. The* $\mathrm{LWE}_{n,m,q,\chi}$ *assumption says that, if we choose $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, $\mathbf{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$ then the following distributions are computationally indistinguishable:*

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) \overset{\mathrm{comp}}{\approx} (\mathbf{A}, \mathbf{u}).$$

It has been shown that the LWE-assumption holds for certain error distributions χ, assuming the worst-case hardness of certain lattice problems. In particular, this is the case if χ is a discrete Gaussian distribution with appropriate variance, see, e.g., [Pei09,Reg05] for precise statements.

Learning with Rounding (LWR). The *learning with rounding* (*LWR*) problem was introduced by Banerjee et al. [BPR12]. It can, in some sense, be seen as a de-randomized version of the LWE-problem. The idea is to compute the error terms deterministically: instead of perturbing the answer by adding a small error, we simply round the answer – in both cases we are intuitively hiding the low order bits.

More formally, the LWR-problem is defined via the following *rounding function* for integers $q \geq p \geq 2$:

$$\lfloor \cdot \rceil_p \ : \ \mathbb{Z}_q \to \mathbb{Z}_p \ : \ x \mapsto \lfloor (p/q) \cdot x \rceil,$$

where we naturally identify elements of $\mathbb{Z}_k$ with the integers in the interval $\{0, \ldots, k-1\}$.[8] More intuitively, $\lfloor . \rceil_p$ partitions $\mathbb{Z}_q$ into intervals of length $\approx \frac{q}{p}$ which it maps to the same image. We naturally extend the rounding function to vectors over $\mathbb{Z}_q$ by applying it component-wise.

In the presentation of our results we will make use that the probability that a random element in $\mathbb{Z}_q$ is close to a step in the rounding function is small. We therefore define, for any integer $\tau > 0$:

$$\mathsf{border}_{p,q}(\tau) \overset{\text{def}}{=} \{x \in \mathbb{Z}_q \ : \ \exists y \in \mathbb{Z}, |y| \leq \tau, \lfloor x \rceil_p \neq \lfloor x + y \rceil_p\}\,.$$

We can easily bound the probability of a random element being on the border.

Lemma 2.7. *We have* $|\mathsf{border}_{p,q}(\tau)| \leq 2\tau p$ *and therefore* $\Pr_{x \overset{\$}{\leftarrow} \mathbb{Z}_q}[x \in \mathsf{border}_{p,q}(\tau)] \leq \frac{2\tau p}{q}$.

Proof. Let $I = \bigcup_{i \in \{0,\ldots,p-1\}} [i \cdot q/p - \tau\ ,\ i \cdot q/p + \tau)$ be a subset over the reals. It is easy to see that $\mathsf{border}_{p,q}(\tau) = \mathbb{Z} \cap I$ and therefore $|\mathsf{border}_{p,q}(\tau)| \leq 2\tau p$. The claim follows. □

The learning with rounding problem is now defined as follows:

Definition 2.8 (LWR [BPR12]). *Let λ be the security parameter, $n = n(\lambda), m = m(\lambda), q = q(\lambda), p = p(\lambda)$ be integers. The* $\mathrm{LWR}_{n,m,q,p}$ *problem states that for* $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$, $\mathbf{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$ *the following distributions are computationally indistinguishable:* $(\mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p) \overset{\text{comp}}{\approx} (\mathbf{A}, \lfloor \mathbf{u} \rceil_p)$.

Notice that when p divides q, the distribution $\lfloor u \rceil_p \ : \ u \overset{\$}{\leftarrow} \mathbb{Z}_q$ is just the uniform over $\mathbb{Z}_p$. Otherwise, the distribution is slightly skewed with some values in $\mathbb{Z}_p$ having probability $\frac{\lfloor q/p \rfloor}{q}$ and others $\frac{\lceil q/p \rceil}{q}$. However, it is easy to deterministically extract random bits from such independent samples with an asymptotic rate of $O(\log(p))$ bits per sample. Therefore, independent samples from the skewed distribution are often "good enough" in practice.

We also define a variant of the LWR assumption where the secret $\mathbf{s}$ can come from some *weak source of entropy* and the attacker may observe some *partial leakage* about $\mathbf{s}$.

Definition 2.9 (LWR with Weak and Leaky Secrets). *Let λ be the security parameter and n, m, q, p be integer parameters as in Definition 2.8. Let $\gamma = \gamma(\lambda) \in (0, q/2)$ be an integer and $k = k(\lambda)$ be a real. The* $\mathrm{LWR}_{n,m,q,p}^{\mathsf{WL}(\gamma,k)}$ *problem says that for any efficiently samplable correlated random variables* $(\mathbf{s}, \mathsf{aux})$*, where the support of* $\mathbf{s}$ *is the integer interval* $[-\gamma, \gamma]^n$ *and* $H_\infty(\mathbf{s}|\mathsf{aux}) \geq k$*, the following distributions are computationally indistinguishable:*

$$(\mathsf{aux}, \mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p) \overset{\text{comp}}{\approx} (\mathsf{aux}, \mathbf{A}, \lfloor \mathbf{u} \rceil_p)$$

where $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m \times n}$, $\mathbf{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$ *are chosen randomly and independently of* $\mathbf{s}, \mathsf{aux}$.

[8] The choice of the floor function rather than ceiling or nearest integer is arbitrary and unimportant.

3 Lossy Mode for LWR

We now show that, under the LW$\underline{\text{E}}$ assumption, the LW$\underline{\text{R}}$ problem has a '*lossy mode*': we can sample a matrix $\tilde{\mathbf{A}}$ which is computationally indistinguishable from a uniformly random $\mathbf{A}$ such that the tuple $(\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p)$ does not reveal too much information about the secret $\mathbf{s}$.

Definition 3.1 (Lossy Sampler). *Let $\chi = \chi(\lambda)$ be an efficiently samplable distribution over $\mathbb{Z}_q$. We define an efficient* lossy sampler $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$ *via:*

$$\mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi) \;:\; \textit{Sample } \mathbf{B} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m\times\ell}, \mathbf{C} \overset{\$}{\leftarrow} \mathbb{Z}_q^{\ell\times n}, \mathbf{F} \overset{\$}{\leftarrow} \chi^{m\times n} \textit{ and output } \tilde{\mathbf{A}} = \mathbf{B}\cdot\mathbf{C} + \mathbf{F}.$$

Although the matrix $\tilde{\mathbf{A}}$ computed by the Lossy algorithm is *statistically* far from a uniformly random matrix, it is easy to show that it is computationally indistinguishable from one under the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption, where the dimension of the secret is now ℓ instead of n. In particular, we can think of each column of $\mathbf{C}$ as an LWE secrets, the matrix $\mathbf{B}$ as the coefficients, and each column of $\tilde{\mathbf{A}}$ as the corresponding LWE output. Therefore, the following lemma from [GKPV10] follows by a simple hybrid argument.

Lemma 3.2 ([GKPV10]). *Let $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m\times n}$, and let $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$. Then, under the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption, the following two distributions are computationally indistinguishable: $\mathbf{A} \overset{\text{comp}}{\approx} \tilde{\mathbf{A}}$.*

We now prove the following lemma, which states that for appropriate parameters, the secret $\mathbf{s}$ maintains a high level of *smooth min-entropy* (see Definition 2.3) given $\tilde{\mathbf{A}}$ and $\lfloor \tilde{\mathbf{A}}\cdot\mathbf{s} \rfloor_p$. The proof is sketched in Section 1.1, and the full derails are below.

Lemma 3.3. *Let n, m, ℓ, p, γ be positive integers, χ be some β-bounded distribution (i.e., $\mathbb{E}[|\chi|] \leq \beta$), and $q \geq 2\beta\gamma nmp$ be a prime. Then the following holds:*

(i) (Uniform Secret) For $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$, $\mathbf{s} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$ we have, for $\varepsilon = 2^{-\lambda} + q^{-\ell}$:

$$H_\infty^\varepsilon(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p) \geq n\log(2\gamma) - (\ell+\lambda)\log(q).$$

(ii) (High-Entropy Secret) Let $(\mathbf{s}, \mathsf{aux})$ be correlated random variables with $\mathbf{s} \in [-\gamma, \gamma]^n \subseteq \mathbb{Z}^n$, and let $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$ be chosen independently. Then, for $\varepsilon = 2^{-\lambda} + q^{-\ell}$ and any $\varepsilon' > 0$ we have:

$$H_\infty^{\varepsilon'+\varepsilon}(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p, \mathsf{aux}) \geq H_\infty^{\varepsilon'}(\mathbf{s}|\mathsf{aux}) - (\ell+\lambda)\log(q).$$

Both parts above also holds when q is not *prime, as long as the largest prime divisor of q, denoted p_{max}, satisfies $GCD(q, q/p_{max}) = 1$, $p_{max} \geq 2\beta\gamma nmp$. In this case we get $\varepsilon = (2^{-\lambda} + (p_{max})^{-\ell} + \Pr[\mathbf{s} = 0^n \mod p_{max}])$.*

Proof. We start by proving part *(ii)* of the Lemma, for a prime q. Recall that for $\tilde{\mathbf{A}}$ chosen via the Lossy function, we can write $\tilde{\mathbf{A}} = \mathbf{BC} + \mathbf{F}$ as described above. Let us define the set $I \overset{\text{def}}{=} \{i \in [m] \;:\; \lfloor (\mathbf{BCs})_i \rfloor_p \neq \lfloor (\tilde{\mathbf{A}}\mathbf{s})_i \rfloor_p\}$, where the subscript i denotes the i^{th} component of the vector. Let $Z = \{(i, \lfloor (\tilde{\mathbf{A}}\mathbf{s})_i \rfloor_p) \;:\; i \in I\}$. It is easy to see that we can reconstruct $\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p$ completely given $\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{Cs}, Z$. Therefore:

$$H_\infty^{\varepsilon+\varepsilon'}(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p, \mathsf{aux}) \geq H_\infty^{\varepsilon+\varepsilon'}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{Cs}, Z, \mathsf{aux}).$$

Now we claim that $|I| \le \lambda$ with overwhelming probability. This is trivially the case if $\mathbf{s} = 0$ since we get $\tilde{\mathbf{A}}\mathbf{s} = \mathbf{BCs} = 0$ and so $I = \emptyset$. Therefore, let us fix any choice of $\mathbf{s} \neq 0$. Then we also have $\Pr_{\mathbf{C}}[\mathbf{Cs} = 0] \le q^{-\ell}$ and so $\Pr[|I| \le \lambda] \le q^{-\ell} + \Pr[|I| \le \lambda \mid \mathbf{Cs} \neq 0]$. To compute $\Pr[|I| \le \lambda \mid \mathbf{Cs} \neq 0]$, let us fix any choice of $\mathbf{C}, \mathbf{s}$ such that $\mathbf{Cs} \neq 0$. For $i \in [m]$, let $\mathbf{b}_i$ and $\mathbf{f}_i$ be the ith row of the matrices $\mathbf{B}$ and $\mathbf{F}$ respectively. Then:

$$\begin{aligned}
\Pr[i \in I] &= \Pr_{\mathbf{b}_i,\mathbf{f}_i}[\ \lfloor \langle \mathbf{b}_i, \mathbf{Cs}\rangle \rfloor_p \neq \lfloor \langle \mathbf{b}_i, \mathbf{Cs}\rangle + \langle \mathbf{f}_i, \mathbf{s}\rangle \rfloor_p\ ] \\
&\le \Pr_{\mathbf{b}_i,\mathbf{f}_i}[\ \langle \mathbf{b}_i, \mathbf{Cs}\rangle \in \mathsf{border}_{p,q}(|\langle \mathbf{f}_i, \mathbf{s}\rangle|)\ ] && (3)\\
&= \sum_j \Pr_{\mathbf{f}_i}[\ |\langle \mathbf{f}_i, \mathbf{s}\rangle| = j\ ] \Pr_{\mathbf{b}_i}[\ \langle \mathbf{b}_i, \mathbf{Cs}\rangle \in \mathsf{border}_{p,q}(j)\ ] && (4)\\
&\le \sum_j \Pr_{\mathbf{f}_i}[|\langle \mathbf{f}_i, \mathbf{s}\rangle| = j] \left(\frac{2jp}{q}\right) && (5)\\
&= \mathbb{E}[|\langle \mathbf{f}_i, \mathbf{s}\rangle|] \left(\frac{2p}{q}\right) \\
&\le \frac{2\beta\gamma np}{q} \le \frac{1}{m} && (6)
\end{aligned}$$

where (3) follow from the definition of a 'border', (4) follows by conditioning, (5) follows by Lemma 2.7 and the fact that $\langle \mathbf{b}_i, \mathbf{Cs}\rangle$ is uniform over $\mathbb{Z}_q$, and (6) follows since each entry of $\mathbf{s}$ is bounded by γ in absolute value and each entry of $\mathbf{f}_i$ is by expectation less than β in absolute value. Furthermore, the events $i \in I$ are mutually independent since the above probabilities are over independent choices of $\mathbf{b}_i, \mathbf{f}_i$. Therefore, for a fixed $\mathbf{C}, \mathbf{s}$ such that $\mathbf{Cs} \neq 0$, we have $\mathbb{E}[|I|] \le 1$ and, by Chernoff (Lemma 2.5), $\Pr[|I| \ge \lambda] \le 2^{-\lambda}$. Finally, taking into account the probability that $\mathbf{Cs} = 0$, we get $\Pr[|I| \ge \lambda] \le 2^{-\lambda} + q^{-\ell} = \varepsilon$.

Since the bit-length of Z is $|Z| = |I|(\log(m) + \log(p))$ we have $|Z| \le \lambda(\log(m) + \log(p))$ with overwhelming probability. Therefore

$$\begin{aligned}
H_\infty^{\varepsilon+\varepsilon'}(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p, Z, \mathsf{aux}) &\ge H_\infty^{\varepsilon+\varepsilon'}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{Cs}, Z, \mathsf{aux}) \\
&\ge H_\infty^{\varepsilon'}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{Cs}, \mathsf{aux}) - \lambda(\log(m) + \log(p)) \\
&\ge H_\infty^{\varepsilon'}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathsf{aux}) - \ell \log(q) - \lambda(\log(m) + \log(p)) \\
&\ge H_\infty^{\varepsilon'}(\mathbf{s}|\mathsf{aux}) - (\ell + \lambda)\log(q)
\end{aligned}$$

where the second and third line follows by Lemma 2.4, and the last line follows since $q \ge mp$. This proves part *(ii)* of the lemma for a prime q.

We now prove part *(i)* of the lemma for a prime q using the same techniques. Let us write $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$ as a sum $\mathbf{s} = \mathbf{s}_1 + \mathbf{s}_2$ where $\mathbf{s}_1 \xleftarrow{\$} \mathbb{Z}_q^n$ and $\mathbf{s}_2 \xleftarrow{\$} [-\gamma, \gamma]^n$. Let $I = \{i \in [m], \lfloor(\tilde{\mathbf{A}}\mathbf{s}_1 + \mathbf{BCs}_2)_i\rfloor_p \neq \lfloor(\tilde{\mathbf{A}}\mathbf{s})_i\rfloor_p\}$. Then we can completely reconstruct $\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p$ given $\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{Cs}_2, \tilde{\mathbf{A}}\mathbf{s}_1, Z$ where $Z = \{(i, \lfloor(\tilde{\mathbf{A}}\mathbf{s})_i\rfloor_p)\ :\ i \in I\}$. We can prove that $|I| \le \lambda$ with overwhelming probability the exact same way as for part *(i)* and therefore $|Z| \le \lambda(\log(m) + \log(p))$ w.o.p. This tells us that

$$\begin{aligned}
H_\infty^{\varepsilon}(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p, Z) &\ge H_\infty^{\varepsilon}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{Cs}_2, \tilde{\mathbf{A}}\mathbf{s}_1, Z) \\
&\ge H_\infty(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \tilde{\mathbf{A}}\mathbf{s}_1) - \ell \log(q) - \lambda(\log(m) + \log(p)) \\
&\ge H_\infty(\mathbf{s}|\mathbf{s}_1) - (\ell + \lambda)\log(q) \\
&\ge n\log(2\gamma) - (\ell + \lambda)\log(q).
\end{aligned}$$

This concludes the proof of part *(i)* of the lemma when q is prime.

Finally, it is easy to extend the analysis to a composite q as long as the largest prime factor p_{max} of q is sufficiently large and relatively prime toq/p_{max}. The only place in our analysis where we used that q is prime is in equation (5) to argue that, for any $\mathbf{Cs} \neq 0$ and uniform $\mathbf{b}_i \xleftarrow{\$} \mathbb{Z}_q$, the value $\langle \mathbf{b}_i, \mathbf{Cs} \rangle$ is uniform over $\mathbb{Z}_q$. Unfortunately, this is no true for composite q. However, if $\mathbf{Cs} \neq 0^n \pmod{p_{max}}$ than $\langle \mathbf{b}_i \mathbf{Cs} \rangle$ is (at the very least) uniform in the p_{max} component and therefore

$$\Pr_{\mathbf{b}_i \xleftarrow{\$} \mathbb{Z}_q^n} [\langle \mathbf{b}_i, \mathbf{Cs} \rangle \in \mathsf{border}_{p,q}(j)] \leq \frac{|\mathsf{border}_{p,q}(j)|}{p_{max}} \leq \frac{2jp}{p_{max}}.$$

This gives us a modified version of equation (6): $\frac{2\beta\gamma np}{p_{max}} \leq \frac{1}{m}$ which is satisfied as long as $p_{max} \geq 2\beta\gamma nmp$. Lastly, we now also need the modified bound:

$$\begin{aligned} \Pr[\mathbf{Cs} = 0^n \mod p_{max}] &\leq \Pr[\mathbf{s} = 0^n \mod p_{max}] + \Pr[\mathbf{Cs} = 0^\ell \mod p_{max} \mid \mathbf{s} \neq 0^n \mod p_{max}] \\ &\leq \Pr[\mathbf{s} = 0^n \mod p_{max}] + (p_{max})^{-\ell} \end{aligned}$$

The rest of the proof is exactly the same. □

4 New "LWR from LWE" Reduction

In the following section we present the main result of this paper, namely sufficient conditions under which the LWR-assumption holds. As discussed in the introduction, the main difference to the work of Banerjee et al. [BPR12] is that our reduction works for a strictly larger range of parameters. In particular, in contrast to their result, our result does not require a super-polynomial modulus or modulus-to-error ratio, at least as long as the number of samples m that will be given out is known ahead of time. The proof of the theorem is sketched in Section 1.1, and the full details are given below.

Theorem 4.1. *Let k, ℓ, n, m, p, γ be positive integers and q be a prime. Further, let χ be a β-bounded distribution for some real-valued β (all parameters are functions of λ) such that $q \geq 2\beta\gamma nmp$. Assuming that the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption holds:*

(i) If $n \geq (\ell + \lambda + 1)\frac{\log(q)}{\log(2\gamma)} + 2\lambda$, then the $\mathrm{LWR}_{n,m,q,p}$-assumption holds.

(ii) If $k \geq (\ell + \lambda + 1)\log(q) + 2\lambda$, then the weak and leaky $\mathrm{LWR}^{\mathsf{WL}(\gamma,k)}_{n,m,q,p}$-assumption holds.

For exact security, if the above LWE assumption is (t, ε)-secure and $\ell \geq \lambda$, then in both cases the corresponding LWR-problem is (t', ε')-secure, where $t' = t - \mathsf{poly}(\lambda)$, $\varepsilon' = m(2 \cdot n\varepsilon + 3 \cdot 2^{-\lambda}) = \mathsf{poly}(\lambda)(\varepsilon + 2^{-\lambda})$. Both parts of the above theorem also hold if q is not prime as long as the largest prime divisor of q, denoted p_{max}, satisfies $GCD(q, q/p_{max}) = 1$, $p_{max} \geq 2\beta\gamma nmp$. In this case we still get $t' = t - \mathsf{poly}(\lambda)$, $\varepsilon' = \mathsf{poly}(\lambda)(\varepsilon + 2^{-\lambda})$.

Proof. We will only prove part *(i)* of the theorem here using part *(i)* of Lemma 3.3. The proof of part *(ii)* of the theorem works in complete analogy using part (ii) of the lemma.

Let us first prove the following: under the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption with the parameters as in the theorem, we have

$$\left(\begin{bmatrix} \lfloor \mathbf{A} \rceil_p \\ \lfloor \mathbf{a} \rceil_p \end{bmatrix}, \begin{bmatrix} \lfloor \mathbf{As} \rceil_p \\ \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rceil_p \end{bmatrix} \right) \overset{\text{comp}}{\approx} \left(\begin{bmatrix} \lfloor \mathbf{A} \rceil_p \\ \lfloor \mathbf{a} \rceil_p \end{bmatrix}, \begin{bmatrix} \lfloor \mathbf{As} \rceil_p \\ \lfloor \mathbf{u} \rceil_p \end{bmatrix} \right) \tag{7}$$

where $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}, \mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, u \xleftarrow{\$} \mathbb{Z}_q$.

Firstly, by Lemma 3.3, we can replace $\mathbf{A}$ with $\tilde{\mathbf{A}} \xleftarrow{\$} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$ to get

$$\left(\begin{bmatrix} \lfloor \mathbf{A} \rfloor_p \\ \lfloor \mathbf{a} \rfloor_p \end{bmatrix}, \begin{bmatrix} \lfloor \mathbf{As} \rfloor_p \\ \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rfloor_p \end{bmatrix} \right) \overset{\text{comp}}{\approx} \left(\begin{bmatrix} \lfloor \tilde{\mathbf{A}} \rfloor_p \\ \lfloor \mathbf{a} \rfloor_p \end{bmatrix}, \begin{bmatrix} \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p \\ \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rfloor_p \end{bmatrix} \right)$$

Now, by part *(i)* of Lemma 3.3, we have $H_\infty^{2^{-\lambda}+q^{-\ell}}(\mathbf{s}|\tilde{\mathbf{A}}, \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p) \geq n\log(2\gamma) - (\ell+\lambda)\log(q) \geq 2\lambda + \log(q)$. Therefore, using the fact that inner product is a good (average-case) strong extractor, we have:

$$\left(\begin{bmatrix} \lfloor \tilde{\mathbf{A}} \rfloor_p \\ \lfloor \mathbf{a} \rfloor_p \end{bmatrix}, \begin{bmatrix} \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p \\ \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rfloor_p \end{bmatrix} \right) \overset{\text{stat}}{\approx} \left(\begin{bmatrix} \lfloor \tilde{\mathbf{A}} \rfloor_p \\ \lfloor \mathbf{a} \rfloor_p \end{bmatrix}, \begin{bmatrix} \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p \\ \lfloor u \rfloor_p \end{bmatrix} \right)$$

where $u \leftarrow \mathbb{Z}_q$ is uniformly random and independent. The exact statistical distance is bounded by $2 \cdot 2^{-\lambda} + q^{-\ell}$. Finally, we can replace the lossy $\tilde{\mathbf{A}}$ with a uniformly random $\mathbf{A}$ to get:

$$\left(\begin{bmatrix} \lfloor \tilde{\mathbf{A}} \rfloor_p \\ \lfloor \mathbf{a} \rfloor_p \end{bmatrix}, \begin{bmatrix} \lfloor \tilde{\mathbf{A}}\mathbf{s} \rfloor_p \\ \lfloor u \rfloor_p \end{bmatrix} \right) \overset{\text{comp}}{\approx} \left(\begin{bmatrix} \lfloor \mathbf{A} \rfloor_p \\ \lfloor \mathbf{a} \rfloor_p \end{bmatrix}, \begin{bmatrix} \lfloor \mathbf{As} \rfloor_p \\ \lfloor u \rfloor_p \end{bmatrix} \right)$$

Combining the above three hybrids proves (7).

To prove the statement of Theorem 4.1, we need to prove

$$(\mathbf{A}, \lfloor \mathbf{As} \rfloor_p) \overset{\text{comp}}{\approx} (\mathbf{A}, \lfloor \mathbf{u} \rfloor_p)$$

where $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$. This follows by a simple sequence of hybrid arguments. Define the hybrid distribution $H_i = (\mathbf{A}, \lfloor \mathbf{h}^{(i)} \rfloor_p)$ where the first i components of the vector $\mathbf{h}^{(i)}$ are taken from $\mathbf{As}$ and the rest are chosen uniformly at random. Then $H_m = (\mathbf{A}, \lfloor \mathbf{As} \rfloor_p)$, $H_0 = (\mathbf{A}, \lfloor \mathbf{u} \rfloor_p)$. Moreover, equation (7) immediately shows the indistinguishability of hybrids H_i and H_{i+1}.

For exact security, assume that the underlying LWE assumption is (t, ε) secure. Notice that our argument consists of m hybrids in which we rely on equation (7). Furthermore, to prove (7) we relied on three sub-steps: the first and the third rely on Lemma 3.3, which itself consists of n hybrids using LWE, and the middle step using having statistical distance $2 \cdot 2^{-\lambda} + q^{-\ell}$. Therefore, the total distance is $m(2n\varepsilon + 2 \cdot 2^{-\lambda} + q^{-\ell})$. □

Remark on β-bounded distributions. In the theorem, we require that the distribution χ is β-bounded meaning that $\mathbb{E}[|\chi|] \leq \beta$. A different definition, which also would have been sufficient for us, would be to require that $\Pr_{x \xleftarrow{\$} \chi}[|x| > \beta] \leq \mathsf{negl}(\lambda)$. The latter notion of boundedness is used in the work of Banerjee et al. [BPR12]. Although the two notions are technically incomparable (one does not imply the other) for natural distributions, such as the discrete Gaussian, it is easier to satisfy out notion. In particular, the discrete Gaussian distribution Ψ_σ with standard deviation σ satisfies $\mathbb{E}[|\Psi_\sigma|] \leq \sigma$ but we can only get the weaker bound $\Pr_{x \xleftarrow{\$} \Psi_\sigma}[|x| > \sqrt{\omega(\log(\lambda))}\sigma] \leq \mathsf{negl}(\lambda)$. Therefore, we find it advantageous to work with our definition.

Remark on Parameters. Notice that in the above theorem, the parameter γ offers a tradeoff between the size of the modulus q and the secret vector length n: for a bigger γ we need a bigger modulus q but can allow smaller secret length n. The following corollary summarizes two extreme cases of small and large γ.

Corollary 4.2. *Let Ψ_σ denote a discrete Gaussian distribution over $\mathbb{Z}_q$ with standard deviation σ, and assume that the* LWE$_{\ell,m,q,\Psi_\sigma}$*-assumption holds. Then the* LWR$_{n,m,q,p}$*-assumption holds in either of the following settings of parameters:*

- *(Minimize Modulus/Error Ratio.) If $q \geq 2\sigma nmp$ is a prime, and $n \geq (\ell+\lambda+1)\log(q)+2\lambda$. By setting $p = O(1)$, we can get a modulus-to-error ratio as small as $q/\sigma = O(m \cdot n)$.*
- *(Maximize Efficiency.) If $q \geq (2\sigma nm)^3$ is a prime, $p = \sqrt[3]{q}$ and $n \geq 3\ell + 5\lambda + 3$. The efficiency of LWR is now similar to the LWE assumption with $n = O(\ell)$ and $\log(p) = O(\log q)$.*

Proof. The corollary follows directly from part *(i)* of Theorem 4.1. In part *(i)* we set $\gamma = 1$ and in part *(ii)* we set $\gamma = p = 2\sigma nm = \sqrt[3]{q}$. We also rely on the fact that the discrete Gaussian distribution satisfies $\mathbb{E}[|\Psi_\sigma|] \leq \sigma$ and is therefore σ-bounded. □

5 Reusable Extractors

The notion of a 'computational reusable extractor' was defined by Dodis et al. [DKL09]. Intuitively, this is a tool that allows us to take some weak secret $\mathbf{s}$ that has a sufficient amount of entropy, and to use it to repeatedly extract fresh pseudorandomness $\mathsf{Ext}(\mathbf{s}; \mathbf{a}_i)$ using multiple public random seeds $\mathbf{a}_i$. Each extracted output should look random and independent. Equivalently, we can think of a reusable extractor as a weak PRF $f_{\mathbf{s}}(\cdot)$, which should be indistinguishable from a random function when evaluated on random inputs $\mathbf{a}_i$ and remain secure even if the secret key $\mathbf{s}$ is not chosen uniformly at random, as long as it has entropy. The work of [DKL09] constructed such reusable extractors under a new assumption called "Learning Subspaces with Noise (LSN)". Reusable extractors were also implicitly constructed based on the DDH assumption in the work of Naor and Segev [NS09].[9] Here we give a new construction based on the LWR problem, with a security reduction from the LWE assumption.

Definition 5.1 (Reusable Extractor). *Let $\mathcal{S}, \mathcal{D}, \mathcal{U}$ be some domains, parameterized by the security parameter λ. A function $\mathsf{Ext} : \mathcal{S} \times \mathcal{D} \to \mathcal{U}$ is a (k,m)-reusable-extractor if for any efficiently samplable correlated random variables $\mathbf{s}, \mathsf{aux}$ such that the support of $\mathbf{s}$ is $\mathcal{S}$ and $H_\infty(\mathsf{s}|\mathsf{aux}) \geq k$, we have:*

$$(\mathsf{aux}, \mathbf{a}_1, \ldots, \mathbf{a}_m, \mathsf{Ext}(\mathbf{s}; \mathbf{a}_1), \ldots, \mathsf{Ext}(\mathbf{s}; \mathbf{a}_m)) \overset{\text{comp}}{\approx} (\mathsf{aux}, \mathbf{a}_1, \ldots, \mathbf{a}_m, u_1, \ldots, u_m)$$

where the values $\{\mathbf{a}_j \overset{\$}{\leftarrow} \mathcal{D}\}, \{u_j \overset{\$}{\leftarrow} \mathcal{U}\}$ are sampled independently.

Theorem 5.2. *Let n, p, γ be integers, p' be a prime, and define $q = p \cdot p'$. Then the function*

$$\mathsf{Ext} : [-\gamma, \gamma]^n \times \mathbb{Z}_q^n \to \mathbb{Z}_p \quad \text{defined by} \quad \mathsf{Ext}(\mathbf{s}; \mathbf{a}) \overset{\text{def}}{=} \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rfloor_p$$

is a (k,m)-reusable extractor assuming that the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption holds for some β-bounded distribution χ such that $p' > 2\beta\gamma nmp$ and $k \geq (\ell+\lambda+1)\log(q)+2\lambda$.

Proof. Follow directly from part *(ii)* of Theorem 4.1. In particular, we want to show that

$$(\mathsf{aux}, \mathbf{a}_1, \ldots, \mathbf{a}_m, \mathsf{Ext}(\mathsf{s}; \mathbf{a}_1), \ldots, \mathsf{Ext}(\mathbf{s}; \mathbf{a}_m)) \overset{\text{comp}}{\approx} (\mathsf{aux}, \mathbf{a}_1, \ldots, \mathbf{a}_m, u_1, \ldots u_m).$$

Let $\mathbf{A}$ be the matrix with rows $\mathbf{a}_1, \ldots, \mathbf{a}_m$. Then showing this is the same as showing $(\mathsf{aux}, \mathbf{A}, \lfloor \mathbf{A}\mathbf{s} \rfloor_p) \overset{\text{comp}}{\approx} (\mathsf{aux}, \mathbf{A}, \lfloor \mathbf{u} \rfloor_p)$ where $\mathbf{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$. □

[9] The function $\mathsf{Ext}(\mathbf{s}; \mathbf{a}) = \prod \mathbf{a}_i^{\mathbf{s}_i}$ is a reusable extractor if $\mathbf{s} \in \mathbb{Z}_q^n$, and the $\mathbf{a} \in \mathbb{G}^n$ for some DDH group of prime order q.

Notice that one nice property of the above reusable extractor is that it has a *graceful degradation of security* as the min-entropy k of the source drops. In particular, there is no hard threshold on the entropy k determined by the parameters that define the scheme: γ, n, q, p. Instead, as the entropy k drops we can still reduce security from a correspondingly less secure LWE assumption with smaller secret size ℓ. In other words, the scheme designer does not need to know the actual entropy k of the secret - but the scheme gets gradually less/more secure as the entropy of the secret shrinks/grows. A similar notion of graceful security degradation was noted in the work of Goldwasser et al. [GKPV10].

6 Lossy Trapdoor Functions

Lossy trapdoor functions (LTDFs) [PW08,PW11], are a family of functions $f_{pk}(\cdot)$ keyed by some public key pk, which can be sampled in one of two indistinguishable modes: `injective` and `lossy`. In the `injective` mode the function $f_{pk}(\cdot)$ is an injective function and we can even sample pk along with a secret trapdoor key sk that allows us to invert it efficiently. In the `lossy` mode, the function $f_{pk}(\cdot)$ is "many-to-one" and $f_{pk}(\mathbf{s})$ statistically loses information about the input $\mathbf{s}$. LTDFs have many amazing applications in cryptography, such as allowing us to output many hardcore bits, construct CCA-2 public-key encryption [PW11,MY10], and deterministic encryption [FOR12]. In this section, we construct very simple and efficient LTDFs using the LWR problem, with security based on standard LWE. Our LTDF function is incredibly simple: the public key is a matrix $pk = \mathbf{A}$ and the function is defined as $f_{\mathbf{A}}(\mathbf{s}) = \lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p$. As we will describe, one can sample an injective $\mathbf{A}$ with a trapdoor using the techniques of Ajtai [Ajt99] or subsequent improvements [AP11,MP12], and one can sample a lossy $\mathbf{A}$ using the techniques we developed in Section 3. Although prior constructions of LTDFs from LWE are known [PW11,BKPW12], our construction here has several advantages. Firstly, our scheme is extremely simple to describe and implement (we find it much simpler than the alternatives). Secondly, in contrast to both [PW08,BKPW12], our lossy mode loses "almost all" of the information contained in $\mathbf{s}$. In fact, the amount of "lossiness" in our LTDF construction is flexible and not determined by the parameters of the scheme itself. Even after we fix the parameters that allow us to sample the injective mode, we have an additional free parameter that allows us to make the lossy mode progressively more lossy under under a progressively stronger variant of the LWE assumption (this is similar to the "graceful security degradation" property of our reusable extractor).

We start by giving formal definitions of LTDFs and a more complex variant called "all-but-one lossy TDFs" (ABO-TDFs). Then, we construct both variants using the LWR problem.

6.1 Definitions of LTDFs

We now define lossy trapdoor function (LTDFs). Our notion differs somewhat from that of [PW11] in how we define the "lossy" property. Instead of requiring that, for a lossy pk, the range of $f_{pk}(\cdot)$ is small, we require that very little *entropy* is lost from observing $f_{pk}(\cdot)$. As far as we know, our version can be used interchangeably in all of the applications of LTDFs to date. As an example of how existing proofs can be extended to work with this new notion of lossiness we describe the necessary modifications to the proof of security for the CCA2 encryption scheme of [PW11] in Appendix A. To avoid confusion, we call our notion *entropic* LTDF (eLTDF).

Definition 6.1 (eLTDF). *A family of* $l(\lambda)$-entropic lossy trapdoor functions (eLTDF) *with security parameter* λ *and domain* $\mathcal{D}_\lambda$ *consists of a* PPT *sampling algorithms* Gen *and two deterministic polynomial-time algorithms* F, F^{-1} *such that:*

Injective Functions: *For any* (pk, sk) *in the support of* $\mathsf{Gen}(1^\lambda, \texttt{injective})$*, any* $\mathbf{s} \in \mathcal{D}_\lambda$ *we require that* $F^{-1}(sk, F(pk, \mathbf{s})) = \mathbf{s}$*. In particular, the function* $f_{pk}(\cdot) = F(pk, \cdot)$ *is injective.*

Lossy Functions: *When* $pk \overset{\$}{\leftarrow} \mathsf{Gen}(1^\lambda, \texttt{lossy})$*, the function* $F(pk, \cdot)$ *is lossy. In particular, for any mutually correlated random variables* $(\mathbf{s}, \mathsf{aux})$ *where the domain of* $\mathbf{s}$ *is* $\mathcal{D}_\lambda$ *and for an independently sampled* $pk \overset{\$}{\leftarrow} \mathsf{Gen}(1^\lambda, \texttt{lossy})$*, we have:* $H_\infty^{\mathsf{smooth}}(\mathbf{s}|pk, F(pk, \mathbf{s}), \mathsf{aux}) \geq H_\infty^{\mathsf{smooth}}(\mathbf{s}|\mathsf{aux}) - l(\lambda)$*. We call the parameter* $l = l(\lambda)$ *the* residual leakage *of the LTDF.*

Indistinguishability: *The distributions of* pk *as sampled by* $\mathsf{Gen}(1^\lambda, \texttt{injective})$ *and* $\mathsf{Gen}(1^\lambda, \texttt{lossy})$ *are computationally indistinguishable.*

We also define a variant called an "all-but-one" trapdoor function (ABO-TDF) [PW11]. In this variant, the function $F(pk, b, \cdot)$ also takes in a "branch" b in some (large) domain $\mathcal{B}$. The sampling algorithm chooses (pk, sk) in such a way that there is a special branch b^* for which $F(pk, b^*, \cdot)$ is lossy, but for all other $b \neq b^*$, the function $F(pk, b, \cdot)$ is injective. Furthermore, the value of the special branch b^* is computationally hidden by the public key pk.

Definition 6.2 (eABO-TDF). *A family of* l-entropic all-but-one trapdoor functions (eABO-TDF) *with security parameter* λ*, branch collection* $\mathcal{B}_\lambda$ *and domain* $\mathcal{D}_\lambda$ *consists of a* PPT *algorithm* Gen *and two deterministic polynomial-time algorithms* G, G^{-1}*. For* $b^* \in \mathcal{B}_\lambda$*, algorithm* $\mathsf{Gen}(1^\lambda, b^*)$ *outputs a function description* pk *and trapdoor* sk *such that the following holds.*

Injective Branches: *For all* $b \in \mathcal{B}_\lambda \setminus \{b^*\}$*, all* $\mathbf{s} \in \mathcal{D}_\lambda$ *we have* $G^{-1}(sk, b, G(pk, b, \mathbf{s})) = \mathbf{s}$*. In particular the function* $G(pk, b, \cdot)$ *is injective.*

Lossy Branch: *The function* $G(pk, b^*, \cdot)$ *is lossy in the following sense: for any fixed* $b^* \in \mathcal{B}_\lambda$*, any mutually correlated random variables* $(\mathbf{s}, \mathsf{aux})$ *where the domain of* $\mathbf{s}$ *is* $\mathcal{D}_\lambda$*, and for an independently sampled* $(pk, sk) \overset{\$}{\leftarrow} \mathsf{Gen}(1^\lambda, b^*)$*, we have:* $H_\infty^{\mathsf{smooth}}(\mathbf{s}|pk, G(pk, b^*, \mathbf{s}), \mathsf{aux}) \geq H_\infty(\mathbf{s}|\mathsf{aux}) - l$*.*

Indistinguishability: *For any pair* b_0^**,* $b_1^* \in \mathcal{B}_\lambda$ *the resulting distributions of the function descriptions sampled by* $\mathsf{Gen}(1^\lambda, b_0^*)$ *and* $\mathsf{Gen}(1^\lambda, b_1^*)$ *are computationally indistinguishable.*

6.2 Construction of LTDFs

We now show how to construct eLTDFs from the LWR problem (assuming standard LWE).

Tools. As a tool in our construction, we will rely on the fact that we can sample a random LWE matrix $\mathbf{A}$ along with an *inversion* trapdoor that allows us to recover $\mathbf{s}, \mathbf{e}$ given an LWE sample $\mathbf{A}\mathbf{s} + \mathbf{e}$ where the error $\mathbf{e}$ is "sufficiently" short. The first example of such algorithms was given by Ajtai in [Ajt99], and was subsequently improved in [AP11]. More recently [MP12] significantly improved the efficiency of these results, by using a "qualitatively" different type of trapdoor. We describe the properties that we need abstractly, and can use any of the above algorithms in a black-box manner. In particular we need the following algorithms for some range of parameters (m, n, q, β):

$\mathsf{GenTrap}(1^n, 1^m, q)$**:** A PPT algorithm which on input positive integers n, q and sufficiently large m samples a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and trapdoor T such that $\mathbf{A}$ is statistically close to uniform (in $n \log q$).

$\mathsf{Invert}(T, \mathbf{A}, \mathbf{c})$**:** An algorithm which receives as input $(\mathbf{A}, T)$ in the support of $\mathsf{GenTrap}(1^n, 1^m, q)$ and some value $\mathbf{c} \in \mathbb{Z}_q^m$ such that $\mathbf{c} = \mathbf{A}\mathbf{s} + \mathbf{e}$ for some $\mathbf{s} \in \mathbb{Z}_q^n$ and some error satisfying $||\mathbf{e}||_2 \leq \beta$. The algorithm outputs $\mathbf{s}$.

$\mathsf{LWRInvert}(T, \mathbf{A}, \mathbf{c})$ Takes as input $(\mathbf{A}, T)$ in the support of $\mathsf{GenTrap}(1^n, 1^m, q)$ and some value $\mathbf{c} \in \mathbb{Z}_p^m$ such that $\mathbf{c} = \lfloor \mathbf{A}\mathbf{s} \rceil_p$ for some $\mathbf{s} \in \mathbb{Z}_q^n$ and outputs $\mathbf{s}$.

In particular [MP12] shows that there are algorithms $(\mathsf{GenTrap}, \mathsf{Invert})$ which work for $n \geq 1$, $q \geq 2$, sufficiently large $m = O(n \log q)$ and sufficiently small $\beta < q/O(\sqrt{n \log q})$. Since we can convert LWR samples $\lfloor \mathbf{A}\mathbf{s} \rceil_p$ into samples $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ for some reasonable short error $||\mathbf{e}||_2 \leq \sqrt{m}q/p$, this also implies the following.

Lemma 6.3 (Trapdoors for LWR). *There exist efficient* $(\mathsf{GenTrap}, \mathsf{LWRInvert})$ *as above for any* $n \geq 1$, $q \geq 2$, *sufficiently large* $m \geq O(n \log q)$ *and* $p \geq O(\sqrt{mn \log q})$.

Proof. We can take an LWR̲ sample $\mathbf{A}, \mathbf{c} = \lfloor \mathbf{A}\mathbf{s} \rceil_p$ and transforms it into an LWE̲ sample $\mathbf{A}, \mathbf{A}\mathbf{s}+\mathbf{e}$ for some "short" $\mathbf{e}$ as follows:

$\mathsf{Transform}_q(\mathbf{c})$**:** Takes as input $\mathbf{c} \in \mathbb{Z}_p^m$ and outputs $\lceil (q/p) \cdot \mathbf{c} \rfloor \in \mathbb{Z}_q^m$.

It is easy to see that, if $\mathbf{c} = \lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p$ for some $\mathbf{s}$, and $q > p$ then:

$$\mathsf{Transform}_q(\mathbf{c}) = \lceil (q/p) \cdot \mathbf{c} \rfloor = \lceil (q/p) \lfloor (p/q)\mathbf{A} \cdot \mathbf{s} \rceil \rfloor = \lceil (q/p)((p/q)\mathbf{A} \cdot \mathbf{s} + \mathbf{e}') \rfloor = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$$

where $\mathbf{e}' \in (-1, 0]^m$ and $\mathbf{e} \in (-q/p, 0]^m$. Therefore $||\mathbf{e}||_\infty \leq q/p$, and $||\mathbf{e}||_2 \leq \sqrt{m}q/p$. As long as $p \geq O(\sqrt{mn \log q})$ is sufficiently big, we have $||\mathbf{e}||_2 \leq q/O(\sqrt{n \log q})$ is sufficiently small so that $\mathsf{Invert}(T, \mathbf{A}, \mathsf{Transform}_q(\mathbf{c}))$ outputs $\mathbf{s}$. Therefore we just define $\mathsf{LWRInvert}(\mathbf{c}) \overset{\text{def}}{=} \mathsf{Invert}(T, \mathbf{A}, \mathsf{Transform}_q(\mathbf{c}))$.

The Construction. We now describe our construction of eLTDFs based LWR. We will rely on the algorithms $\mathsf{GenTrap}$ and $\mathsf{LWRInvert}$ described above. We also rely on the lossy sampling algorithm Lossy and its properties developed in Section 3. The construction is parameterized by integers n, m, q, p (all functions of the security parameter λ). Furthermore, there will be two additional parameters ℓ and χ which are only needed by the lossy sampler.

$\mathsf{Gen}(1^\lambda, \texttt{injective})$**:** Sample $(\mathbf{A}, T) \xleftarrow{\$} \mathsf{GenTrap}(1^n, 1^m, q)$. Output $pk = \mathbf{A}$ and trapdoor $sk = (\mathbf{A}, T)$.
$\mathsf{Gen}(1^\lambda, \texttt{lossy})$**:** Sample $\mathbf{A} \xleftarrow{\$} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$. Output description $pk = \mathbf{A}$.
$F(pk, \mathbf{s})$**:** On input $\mathbf{s} \in \{0,1\}^n$ and matrix $pk = \mathbf{A} \in \mathbb{Z}_q^{m \times n}$ output $\lfloor \mathbf{A}\mathbf{s} \rceil_p$.
$F^{-1}(pk, \mathbf{c})$**:** On input $\mathbf{c} \in \mathbb{Z}_p^m$ and $t = (\mathbf{A}, T)$ output $\mathsf{LWRInvert}(T, \mathbf{A}, \mathbf{c})$.

For the following theorem summarizes the properties of this construction.

Theorem 6.4. *Let* χ *be an efficiently samplable* β*-bounded distribution and* λ *be the security parameter. For any positive integers* $n \geq \lambda$, *sufficiently large* $m \geq O(n \log q)$, $p \geq O(\sqrt{mn \log q})$ *and a prime* $q \geq 2\beta nmp$, *if the* $\mathrm{LWE}_{\ell,m,q,\chi}$ *assumption holds then the above construction is an* l*-LTDF with where* $l = (\ell + \lambda) \log q$.

Proof. Firstly, the correctness of inversion follows directly from Lemma 6.3.

Secondly, the indistinguishability property follows since: (1) in injective mode we set $pk = \mathbf{A}$ as the output of $\mathsf{GenTrap}$ which is statistically close to uniform, and (2) in lossy mode we set $pk = \mathbf{A}$ as the output of Lossy which is computationally indistinguishable from uniform by Lemma 3.2. Therefore, the two distributions are computationally indistinguishable.

Finally the second part Lemma 3.3 using $\gamma = 1$ directly implies the bound on the lossiness of the construction. □

7 Construction of “All-But-One” Lossy Trapdoor Functions

We now show how to construct eABO-TDFs from the LWR problem. Our construction relies on the ideas of [ABB10] used to construct identity based encryption from LWE. It also bears some similarity to the identity-based lossy-trapdoor functions of [BKPW12].

Technical Tools. As part of the construction we make use of the following full rank difference mappings for which we have adapted the definition of [ABB10].

Definition 7.1 (FRD Mapping). *Let q positive integers. A collection of* full rank difference *(FRD) mappings is a sequence of functions $H = \{H_{n,q}\}$ with domain $\mathcal{B}_{n,q}$ and range $\mathbb{Z}_q^{n\times n}$ such that for all distinct $x_0, x_1 \in \mathcal{B}_n$ the matrix $H(x_0) - H(x_1) \in \mathbb{Z}_q^{n\times n}$ has full rank. Moreover $H_{n,q}$ is computable in time polynomial in $n \log q$.*

A construction of an FRD mappings for prime q and any n is given in [CD09] with domain $\mathcal{B}_{n,q} = \mathbb{Z}_q^n$.

As a second tool, we also use the following lemma (similar to one in [ABB10]) which shows how to extend a trapdoor that allows us to solve LWE with some matrix $\mathbf{A}$ into a trapdoor that allows us to solve LWE for some larger matrix derived from $\mathbf{A}$.

Lemma 7.2. *For $n \geq 1$, $q \geq 2$, and sufficiently large $m = O(n \log q)$, $p = O(m\sqrt{n \log q})$ there exists* PPT *algorithms* $\mathsf{BigInvert}$ *and* $\mathsf{GenTrap}$ *such that for $(\mathbf{A}_1, T) \xleftarrow{\$} \mathsf{GenTrap}(1^n, 1^m, q)$, any $\mathbf{A}_0 \in Z_q^{m\times n}$, $\mathbf{R} \in \{-1,1\}^{m\times m}$, $\mathbf{s} \in \mathbb{Z}_q^n$ and invertible $\mathbf{G} \in \mathbb{Z}_q^{n\times n}$ we have $\mathbf{s} =$ $\mathsf{BigInvert}(\mathbf{R}, T, \mathbf{A}_1, \mathbf{G}, \lfloor \mathbf{Ds} \rfloor_p)$ where*

$$\mathbf{D} = \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{R}\mathbf{A}_0 + \mathbf{A}_1\mathbf{G} \end{bmatrix}.$$

Proof. We use the construction of [MP12] to instantiate $\mathsf{GenTrap}$ and Invert as described at the beginning of this section, and the algorithm $\mathsf{Transform}_q$ described in the proof of Lemma 6.3. First, we describe algorithm $\mathsf{BigInvert}$.

Let $\mathbf{c} = \mathsf{Transform}_q(\lfloor \mathbf{Ds} \rfloor_p) = \mathbf{Ds} + \mathbf{e} \in \mathbb{Z}_q^{2m}$ where $||\mathbf{e}||_\infty \leq q/p$. Denote the first half of $\mathbf{c}$ as $\mathbf{c}_0 \in \mathbb{Z}_q^m$ and the second half as $\mathbf{c}_1 \in \mathbb{Z}_q^m$. Algorithm $\mathsf{BigInvert}$ first computes $\mathbf{c}' = \mathbf{c}_1 - \mathbf{R}\mathbf{c}_0$. Then it computes $\bar{\mathbf{s}} \xleftarrow{\$} \mathsf{Invert}(T, \mathbf{A}_1, \mathbf{c}')$ and outputs $\mathbf{G}^{-1}\bar{\mathbf{s}}$.

To see why this works we can write $\mathbf{c}_0 := \mathbf{A}_0\mathbf{s} + \mathbf{e}_0$ and $\mathbf{c}_1 := (\mathbf{R}\mathbf{A}_0 + \mathbf{A}_1\mathbf{G})\mathbf{s} + \mathbf{e}_1$ for error vectors $\mathbf{e}_0$ and $\mathbf{e}_1$ with $||\mathbf{e}_0||_\infty, ||\mathbf{e}_1||_\infty \leq q/p$. Therefor we have

$$\mathbf{c}' = \mathbf{c}_1 - \mathbf{R}\mathbf{c}_0 = \mathbf{R}\mathbf{A}_0\mathbf{s} + \mathbf{A}_1\mathbf{G}\mathbf{s} + \mathbf{e}_1 - \mathbf{R}\mathbf{A}_0\mathbf{s} - \mathbf{R}\mathbf{e}_0 = \mathbf{A}_1\mathbf{s}' + \mathbf{e}'$$

where $\mathbf{s}' := \mathbf{G}\mathbf{s}$ and $\mathbf{e}' := \mathbf{e}_1 - \mathbf{R}\mathbf{e}_0$. To show that $\mathsf{Invert}(T, \mathbf{A}_1, \mathbf{c}')$ returns $\bar{\mathbf{s}} = \mathbf{s}'$ (and so that $\mathbf{G}^{-1}\bar{\mathbf{s}} = \mathbf{s}$) it suffices to upper-bound $||\mathbf{e}'||_2$ such that the algorithms of [MP12] can be used for $\mathsf{GenTrap}$ and Invert. In particular by definition of $\mathbf{R}$ and p we see that $||\mathbf{e}'||_\infty \leq (m+1)(q/p)$ and therefore $||\mathbf{e}'||_2 \leq \sqrt{m}(m+1)(q/p) < O(q/(\sqrt{n \log q}))$ as required. □

Construction. The following construction uses the algorithm Lossy, $\mathsf{GenTrap}$ and $\mathsf{BigInvert}$ described above. Moreover it is parameterized by integers n, m, q, ℓ, p and a distribution χ, all functions of security parameter λ. Finally the construction also makes use of an FRD collection $H = H_{n,q}$ with domain $\mathbb{Z}_q^n$.

Parameters: For security parameter λ define branch set $\mathcal{B} := \mathbb{Z}_q^n$ and the domain $\mathcal{D} = \mathbb{Z}_2^n$.

$\mathsf{Gen}(1^\lambda, b^*)$**:** Function Sampling

1. Sample $\mathbf{A}_0 \xleftarrow{\$} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$ and uniform $\mathbf{R} \xleftarrow{\$} \{-1, 1\}^{m \times m}$.
2. Sample $(\mathbf{A}, T) \xleftarrow{\$} \mathsf{GenTrap}(1^n, 1^m, q)$ and set $\mathbf{A}_1 := \mathbf{R}\mathbf{A}_0 + \mathbf{A}H(b^*)$.
3. Output $pk = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A})$ and $sk = (\mathbf{R}, T, \mathbf{A}, H(b^*))$.

$G(pk, b, \mathbf{s})$**:** Function Evaluation

1. Set
$$\bar{\mathbf{A}} := \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{A}_1 - \mathbf{A}H(b) \end{bmatrix}$$
and output $\lfloor \bar{\mathbf{A}} \cdot \mathbf{s} \rfloor_p \in \mathbb{Z}_q^{2m}$.

$G^{-1}(sk, b, \mathbf{c})$**:** Function Inversion

1. Parse $sk = (\mathbf{R}, T, \mathbf{A}, H(b^*))$.
2. Set $\mathbf{G} := H(b^*) - H(b)$ and output vector $\mathbf{s} = \mathsf{BigInvert}(\mathbf{R}, T, \mathbf{A}, \mathbf{G}, \mathbf{c})$.

We summarize the properties of this construction in the following theorem.

Theorem 7.3. *Let λ be a security parameter. Let ℓ, m, n, p be integers, q prime, β real (all functions of λ) such that $n \geq \lambda$, the values $m \geq O(n \log q)$, $p \geq O(m\sqrt{n \log q})$ are sufficiently large and $q \geq 4\beta n m^2 p$. Let χ be some distribution such that $\Pr_{x \xleftarrow{\$} \chi}[|x| \geq \beta] \leq \mathsf{negl}(\lambda)$. Then, assuming that the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption holds, the above construction is an l-ABO TDF where $l = (\ell + \lambda)\log q$.*

Proof. Firstly, Lemma 7.2 immediately implies that the algorithm G^{-1} inverts $\mathbf{c}$ correctly when $b \neq b^*$.

Next, we show that, no matter what branch b^* is chosen, the output of $pk \xleftarrow{\$} \mathsf{Gen}(1^\lambda, b^*)$ is computationally indistinguishable from sampling three uniformly random and independent matrices $\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}$ and setting $pk = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A})$. Firstly, in the computation of $\mathsf{Gen}(1^\lambda, b^*)$, we can switch $\mathbf{A}_0$ to being sampled uniformly at random instead of in its lossy mode, and this is indistinguishable by Lemma 3.2. Therefore, now $\mathbf{A}, \mathbf{A}_0$ are random and mutually independent. Secondly, we can view multiplication by $\mathbf{A}_0$ as a strong extractor [DORS08] applied to $\mathbf{R}$ (think of each row of $\mathbf{R}$ as weak randomness, and the matrix $\mathbf{A}_0$ as a seed), and therefore $\mathbf{R}\mathbf{A}_0$ is close to being uniform and independent of $\mathbf{A}$ and $\mathbf{A}_0$. Thus $\mathbf{A}_1 = \mathbf{R}\mathbf{A}_0 + \mathbf{A}H(b^*)$ looks random and independent of $\mathbf{A}_0$ and $\mathbf{A}$. This proves the indistinguishability property of eABO-TDFs.

It remains to show that the residual leakage of the lossy branch is at most l. When evaluating $G(pk, b^*, \mathbf{s})$ we have $\mathbf{c} = \lfloor \bar{\mathbf{A}}\mathbf{s} \rfloor_p$ where $\bar{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{R}\mathbf{A}_0 \end{bmatrix}$ and $\mathbf{A}_0 \xleftarrow{\$} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$. Unfortunately, since the matrix $\bar{\mathbf{A}}$ itself is not sampled from the lossy algorithm, we cannot use Lemma 3.3 in a black-box way. Instead, recall that we can write $\mathbf{A}_0 = \mathbf{B}\mathbf{C} + \mathbf{F}$ where $\mathbf{B}, \mathbf{C}$ are random and $\mathbf{F} \leftarrow \chi^{m \times n}$. Therefore we can write

$$\bar{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{R}\mathbf{A}_0 \end{bmatrix} = \mathbf{B}^*\mathbf{C} + \mathbf{F}^* \quad \text{where} \quad \mathbf{B}^* = \begin{bmatrix} \mathbf{B} \\ \mathbf{R}\mathbf{B} \end{bmatrix} \quad \text{and} \quad \mathbf{F}^* = \begin{bmatrix} \mathbf{F} \\ \mathbf{R}\mathbf{F} \end{bmatrix}.$$

Secondly, we claim that $\mathbf{R}\mathbf{B}$ is statistically close to being uniformly random and independent of $\mathbf{B}, \mathbf{R}\mathbf{F}$. Here we again rely on the fact that matrix multiplication is a good extractor: we can think of $\mathbf{R}\mathbf{F}$ as leaking $n\log(q)$ bits of information on each *row* of $\mathbf{R}$, and $\mathbf{R}\mathbf{B}$ as extracting $n\log(q)$ bits of information – therefore the statistical distance is negligible if $m > 2n\log(q) + O(\lambda)$. Therefore, $\bar{\mathbf{A}} = \mathbf{B}^*\mathbf{C} + \mathbf{F}^*$ has the almost the same distribution as an output of the lossy algorithm $\mathsf{Lossy}(1^n, 1^{2m}, 1^\ell, q, \chi)$ *except* that the distribution of $\mathbf{F}^*$ is different. Nevertheless, we can bound each entry of $\mathbf{F}^*$ by $\beta^* = m\beta$ in absolute value with overwhelming probability. Since this was the only property that Lemma 3.3 relied on, we can apply its conclusions with $m^* = 2m$ and $\beta^* = m\beta$. □

8 Deterministic Encryption

Deterministic public-key encryption [BBO07,BFOR08,BFO08,BS11,FOR12] is intended to guarantee security as long as the messages have sufficient entropy. Although there are black-box constructions of deterministic encryption using LTDFs [BFO08], here we present a very simple direct construction from the LWR problem. There are several definitions of deterministic encryption which can be proven equivalent; see [BFOR08,BFO08]. Here, we will use one such simple definition based on indistinguishability of encrypting messages from two different distributions.

Definition 8.1 (Deterministic Encryption). *A* deterministic encryption scheme *with message length $n = n(\lambda)$ consists of a PPT procedure $(pk, sk) \xleftarrow{\$} \mathsf{Gen}(1^\lambda)$ along with two deterministic poly-time functions $\mathsf{Enc}, \mathsf{Dec}$. For correctness, we require that for all (pk, sk) in the support of $\mathsf{Gen}(1^\lambda)$, all messages $\mathbf{s} \in \{0,1\}^n$, we have $\mathsf{Dec}_{sk}(\mathsf{Enc}_{pk}(\mathbf{s})) = \mathbf{s}$. We say that the scheme is secure for all $k(\lambda)$-sources if for any two distribution ensembles $\{S_\lambda^{(0)}\}_{\lambda\in\mathbb{N}}, \{S_\lambda^{(1)}\}_{\lambda\in\mathbb{N}}$ over $\{0,1\}^{n(\lambda)}$ which are efficiently sampleable in $\mathsf{poly}(\lambda)$-times and have sufficient entropy $H_\infty(S_\lambda^0) \geq k$, $H_\infty(S_\lambda^1) \geq k$, we have $(pk, \mathsf{Enc}_{pk}(\mathbf{s}_0)) \overset{\text{comp}}{\approx} (pk, \mathsf{Enc}_{pk}(\mathbf{s}_1))$, where $\mathbf{s}_0 \xleftarrow{\$} S_\lambda^{(0)}$ and $\mathbf{s}_1 \xleftarrow{\$} S_\lambda^{(1)}$ and $(pk, sk) \xleftarrow{\$} \mathsf{Gen}(1^\lambda)$.*

Construction. We give a very simple construction of deterministic encryption based on the LWR assumption. This construction is the same as one given by Xie et al. [XXZ12], except for the setting of parameters. Whereas they required a super-polynomial modulus and modulus to error ratio by relying on variants of the analysis of [GKPV10,BPR12] we use our improved analysis from Section 4. We will rely on the LWR trapdoor generation and inversion algorithms $\mathsf{GenTrap}, \mathsf{LWRInvert}$ described in Section 6.2 and Lemma 6.3. Our scheme is parameterized by some n, m, q, p, all functions of the security parameter λ, and has message length n.

$(pk, sk) \xleftarrow{\$} \mathsf{Gen}(1^\lambda)$: Choose $(\mathbf{A}, T) \xleftarrow{\$} \mathsf{GenTrap}(1^n, 1^m, q)$. Output $pk = \mathbf{A}$, $sk = T$.
$\mathsf{Enc}_{pk}(\mathbf{s})$: For a message $\mathbf{s} \in \{0,1\}^n$, output $\lfloor \mathbf{A} \cdot \mathbf{s} \rceil_p$.
$\mathsf{Dec}_{sk}(\mathbf{c})$: For a ciphertext $\mathbf{c} \in \mathbb{Z}_p^m$, output $\mathsf{LWRInvert}(T, \mathbf{A}, \mathbf{c})$.

Theorem 8.2. *Let λ be the security parameter, $n \geq \lambda, \ell, m, p$ be an integers, q be a prime, and χ be an efficiently samplable β-bounded distribution (all parameters are functions of λ) such that $m \geq O(n \log q)$, $p \geq O(\sqrt{mn \log q})$ are sufficiently large and $q \geq 2\beta nmp$. If the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption holds then the above construction with parameters n, m, q, p is a deterministic encryptions secure for all k sources where $k \geq (\ell + \Omega(\lambda)) \log(q)$.*

Proof. Correctness of decryption follows directly from Lemma 6.3. On the other hand, indistinguishability follows by part (ii) of Theorem 4.1. In particular, we chose our parameters such that the $\mathrm{LWR}_{n,m,q,p}^{\mathsf{WL}(1,k)}$ assumption with weak (and leaky) secrets holds. This means that for any sources $S_\lambda^{(0)}, S_\lambda^{(1)}$ over $\{0,1\}^n$ such that $H_\infty(S_\lambda^0) \geq k$, $H_\infty(S_\lambda^1) \geq k$, we have:

$$(pk, \mathsf{Enc}_{pk}(\mathbf{s}_0)) \overset{\text{stat}}{\approx} (\mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s}_0 \rceil_p) \overset{\text{comp}}{\approx} (\mathbf{A}, \lfloor \mathbf{u} \rceil_p) \overset{\text{comp}}{\approx} (\mathbf{A}, \lfloor \mathbf{A} \cdot \mathbf{s}_1 \rceil_p) \overset{\text{stat}}{\approx} (pk, \mathsf{Enc}_{pk}(\mathbf{s}_1))$$

where $\mathbf{s}_0 \xleftarrow{\$} S_\lambda^{(0)}, \mathbf{s}_1 \xleftarrow{\$} S_\lambda^{(1)}$, $(pk, sk) \xleftarrow{\$} \mathsf{Gen}(1^\lambda)$, $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m\times n}$ and $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^m$. In particular, the first step follows by noting that $\mathsf{GenTrap}$ produces $pk = \mathbf{A}$ statistically close to uniform, the second step follows by $\mathrm{LWR}_{n,m,q,p}^{\mathsf{WL}(1,k)}$ and the last two steps mimic the first two in reverse.

One big advantage of our scheme is that the parameters n, m, q, p do not determine the minimal entropy k. Instead for any k, we can prove security under a corresponding LWE assumption with dimension $\ell < k$. This is similar to the property of our reusable extractors and LTDF, allowing for a graceful degradation of security as the entropy of the message decreases.

9 Conclusion

In summary, we give an improved security reduction showing the hardness of LWR under the LWE assumption for a wider setting of parameters. In doing so, we also show that the LWR problem has a "lossy mode". Together, these results lead to several interesting applications: security with weak/leaky secrets, reusable extractors, lossy trapdoor functions, and deterministic encryption. We conclude with several interesting open problems. Firstly, can we improve the reduction further and get rid of the dependence between the modulus q and the number of samples m? Secondly, can we use the techniques from this work to also improve the parameters of the "degree-k synthesizer" PRF construction of [BPR12]? Lastly, can we use the techniques from this work to also get a reduction for *Ring LWR* from *Ring LWE*? This does not seem to follow in a straight-forward manner.

10 Acknowledgements

We would like to acknowledge David Cash and Shai Halevi for initial discussions on the LWR problem and getting us interested in the possibility of finding a better reduction.

References

ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient Lattice (H)IBE in the Standard Model. In Gilbert [Gil10], pages 553–572.

AFV11. Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. Functional Encryption for Inner Product Predicates from Learning with Errors. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 21–40. Springer, 2011.

AGV09. Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous Hardcore Bits and Cryptography against Memory Attacks. In Omer Reingold, editor, *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 474–495. Springer, 2009.

Ajt99. Miklós Ajtai. Generating Hard Instances of the Short Basis Problem. In Jirí Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *ICALP*, volume 1644 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 1999.

AJW11. Gilad Asharov, Abhishek Jain, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold fhe. *IACR Cryptology ePrint Archive*, 2011:613, 2011.

AP11. Joël Alwen and Chris Peikert. Generating Shorter Bases for Hard Random Lattices. *Theory Comput. Syst.*, 48(3):535–553, 2011.

BBO07. Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. Deterministic and efficiently searchable encryption. In Alfred Menezes, editor, *CRYPTO*, volume 4622 of *Lecture Notes in Computer Science*, pages 535–552. Springer, 2007.

BFO08. Alexandra Boldyreva, Serge Fehr, and Adam O'Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In Wagner [Wag08], pages 335–359.

BFOR08. Mihir Bellare, Marc Fischlin, Adam O'Neill, and Thomas Ristenpart. Deterministic encryption: Definitional equivalences and constructions without random oracles. In Wagner [Wag08], pages 360–378.

BKPW12. Mihir Bellare, Eike Kiltz, Chris Peikert, and Brent Waters. Identity-Based (Lossy) Trapdoor Functions and Applications. In Pointcheval and Johansson [PJ12], pages 228–245.

BPR12. Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom Functions and Lattices. In Pointcheval and Johansson [PJ12], pages 719–737.

BS11. Zvika Brakerski and Gil Segev. Better security for deterministic public-key encryption: The auxiliary-input setting. In Rogaway [Rog11], pages 543–560.

BV11. Zvika Brakerski and Vinod Vaikuntanathan. Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages. In Rogaway [Rog11], pages 505–524.

CD09. Ronald Cramer and Ivan Damgård. On the Amortized Complexity of Zero-Knowledge Protocols. In Halevi [Hal09], pages 177–191.

DKL09. Yevgeniy Dodis, Yael Tauman Kalai, and Shachar Lovett. On Cryptography with Auxiliary Input. In Mitzenmacher [Mit09], pages 621–630.

DORS08. Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM J. Comput.*, 38(1):97–139, 2008.

Dwo08. Cynthia Dwork, editor. *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*. ACM, 2008.

FOR12. Benjamin Fuller, Adam O'Neill, and Leonid Reyzin. A Unified Approach to Deterministic Encryption: New Constructions and a Connection to Computational Entropy. In Ronald Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 582–599. Springer, 2012.

Gil10. Henri Gilbert, editor. *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*. Springer, 2010.

GKP$^+$12. Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Succinct functional encryption and applications: Reusable garbled circuits and beyond. Cryptology ePrint Archive, Report 2012/733, 2012. `http://eprint.iacr.org/`.

GKPV10. Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the Learning with Errors Assumption. In Andrew Chi-Chih Yao, editor, *ICS*, pages 230–240. Tsinghua University Press, 2010.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for Hard Lattices and New Cryptographic Constructions. In Dwork [Dwo08], pages 197–206.

Hal09. 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*. Springer, 2009.

KV09. Jonathan Katz and Vinod Vaikuntanathan. Smooth Projective Hashing and Password-Based Authenticated Key Exchange from Lattices. In Matsui [Mat09], pages 636–652.

LP11. Richard Lindner and Chris Peikert. Better Key Sizes (and Attacks) for LWE-Based Encryption. In Aggelos Kiayias, editor, *CT-RSA*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On Ideal Lattices and Learning with Errors over Rings. In Gilbert [Gil10], pages 1–23.

Lyu09. Vadim Lyubashevsky. Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures. In Matsui [Mat09], pages 598–616.

Lyu12. Vadim Lyubashevsky. Lattice Signatures without Trapdoors. In Pointcheval and Johansson [PJ12], pages 738–755.

Mat09. Mitsuru Matsui, editor. *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*. Springer, 2009.

Mit09. Michael Mitzenmacher, editor. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*. ACM, 2009.

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller. In Pointcheval and Johansson [PJ12], pages 700–718.

MU05. Michael Mitzenmacher and Eli Upfal. *Probability and Computing – Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.

MY10. Petros Mol and Scott Yilek. Chosen-Ciphertext Security from Slightly Lossy Trapdoor Functions. In Phong Q. Nguyen and David Pointcheval, editors, *Public Key Cryptography*, volume 6056 of *Lecture Notes in Computer Science*, pages 296–311. Springer, 2010.

NS09. Moni Naor and Gil Segev. Public-Key Cryptosystems Resilient to Key Leakage. In Halevi [Hal09], pages 18–35.

Pei09. Chris Peikert. Public-Key Cryptosystems from the Worst-Case Shortest Vector Problem: Extended Abstract. In Mitzenmacher [Mit09], pages 333–342.

PJ12. 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*. Springer, 2012.

PR06. Chris Peikert and Alon Rosen. Efficient Collision-Resistant Hashing from Worst-Case Assumptions on Cyclic Lattices. In Shai Halevi and Tal Rabin, editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006.

PVW08. Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A Framework for Efficient and Composable Oblivious Transfer. In Wagner [Wag08], pages 554–571.

PW08. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Dwork [Dwo08], pages 187–196.

PW11. Chris Peikert and Brent Waters. Lossy Trapdoor Functions and Their Applications. *SIAM J. Comput.*, 40(6):1803–1844, 2011.

Reg05. Oded Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005.

Rog11. Phillip Rogaway, editor. *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*. Springer, 2011.

Rüc10. Markus Rückert. Lattice-Based Blind Signatures. In Masayuki Abe, editor, *ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 413–430. Springer, 2010.

RW04. Renato Renner and Stefan Wolf. Smooth Rényi Entropy and Applications. In *ISIT 04*, page 233, 2004.

Wag08. David Wagner, editor. *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, volume 5157 of *Lecture Notes in Computer Science*. Springer, 2008.

XXZ12. Xiang Xie, Rui Xue, and Rui Zhang. Deterministic Public Key Encryption and Identity-Based Encryption from Lattices in the Auxiliary-Input Setting. In Ivan Visconti and Roberto De Prisco, editors, *SCN*, volume 7485 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2012.

A Using Entropic Lossiness for CCA-2 Encryption.

We briefly (informally) argue that our *entropic* notions of eLTDF and eABO-TDF can be used to securely instantiate the CCA2 encryption construction of [PW11]. Let us first recall the construction. It uses a pairwise independent hash function family $\mathcal{H}$ and one-time signature scheme. Public key generation consists of sampling $h \xleftarrow{\$} \mathcal{H}$, a public key for LTDF $F(pk_F, \cdot)$ and for an ABO-TDF $G(pk_G, \cdot, \cdot)$, all with the same input domain $\mathcal{D}$. The branch-set of ABO-TDF contains the space of verification keys for the signature scheme. The public key is $pk = (h, pk_F, pk_G)$ and secret key is the trapdoor sk_F for F. To encrypt a message m sample fresh signature key pair (vk, sk) and uniform $\mathbf{s} \xleftarrow{\$} \mathcal{D}$. Then compute components $c_1 = f(\mathbf{s})$, $c_2 = g(vk, \mathbf{s})$ and $c_3 = m + h(\mathbf{s})$ and output ciphertext $c = (c_1, c_2, c_3, vk, \sigma)$ where σ is a signature of (c_1, c_2, c_3). To decrypt, we recover $\mathbf{s}$ from c_1, recompute and verify that $c_2 \overset{?}{=} G(pk_G, vk, \mathbf{s})$ and that the signature σ is good, and if so recover m from c_3.

The CCA2 security of the scheme is argued in several steps. First, we choose a random vk^* for the challenger ciphertext and, when generating pk_G, we set the lossy branch to be $b^* = vk^*$. Second, we decrypt all decryption queries using the trapdoor to sk_G (instead of sk_F) and then check that c_1 was computed correctly. Lastly, we choose pk_F in its lossy mode. At this point we can use an information theoretic argument to argue that, for the secret $\mathbf{s}^*$ used in the challenge ciphertext $c^* = (c_1^* = F(pk_F, \mathbf{s}^*), c_2^* = G(pk_G, vk^*, \mathbf{s}^*), c_3^* = m + h(\mathbf{s}^*))$, the min-entropy of $\mathbf{s}^*$ given c_1^*, c_2^* is large. In the original proof, we could argue that $H_\infty(\mathbf{s}^* \mid c_1^*, c_2^*) \geq H_\infty(\mathbf{s}^*) - 2l$ where the range of the lossy functions $f(\cdot)$, $g(vk^*, \cdot)$ is at most 2^l. Then since h is a strong (average case) extractor [DORS08], its output looks essentially uniform when evaluated on $\mathbf{s}^*$ even given c_1^*, c_2^*, and therefore the message is information theoretically hidden.

We notice that the same conclusion can be reached when F is an eLTDF and G is an eABO TDF. Assume that the leakage of each function is l. Then, for the challenge ciphertext, we can show

$$H_\infty^{\mathsf{smooth}}(\mathbf{s}^* \mid pk_F, pk_G, vk^*, c_1^*, c_2^*) \geq H_\infty^{\mathsf{smooth}}(\mathbf{s}^* | pk_F, c_1^*) - l \geq H_\infty^{\mathsf{smooth}}(\mathbf{s}^*) - 2l$$

where, in the first step, we think of $\mathsf{aux} = (pk_F, c_1^*)$ as auxiliary information. The rest of the proof is exactly the same.

B Robustness of LWE for Non-Uniform Secrets

As mentioned in the main part of this document, an analogue to Theorem 4.1 also holds for the LWE-assumption with weak and leaky keys, i.e., the case where the LWE-secret $\mathbf{s}$ is not drawn uniformly at random, but only from a high min-entropy distribution.

This is formalized by the following definition:

Definition B.1 (LWE with Weak and Leaky Secrets). *Let λ be the security parameter and n, m, q be integer parameters and χ be a distribution as in Definition 2.6. Let $\gamma = \gamma(\lambda) \in (0, q/2)$ be an integer and $k = k(\lambda)$ be a real. The $\mathrm{LWE}_{n,m,q,\chi}^{\mathsf{WL}(\gamma,k)}$ problem says that for any efficiently samplable correlated random variables $(\mathbf{s}, \mathsf{aux})$, where the support of $\mathbf{s}$ is the integer interval $[-\gamma, \gamma]^n$ and $H_\infty(\mathbf{s}|\mathsf{aux}) \geq k$, the following distributions are computationally indistinguishable:*

$$(\mathsf{aux}, \mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) \overset{\text{comp}}{\approx} (\mathsf{aux}, \mathbf{A}, \mathbf{u})$$

where $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}_q^{m \times n}$, $\mathbf{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^m$, $\mathbf{e} \overset{\$}{\leftarrow} \chi^m$ are chosen randomly and independently of $\mathbf{s}, \mathsf{aux}$.

We will need the following lemma, which intuitively says that if two distributions are close, then with high probability a sample from the first distribution can also be used as a sample for the second distribution:

Lemma B.2. *Let D_X, D_Y be probability distributions such that $\Delta(D_X, D_Y) \leq \delta$. Then there exist distributions D_W, D'_X, D'_Y such that for $X \overset{\$}{\leftarrow} D_X, Y \overset{\$}{\leftarrow} D_Y$ the following holds:*

$$\tilde{X} \approx X \qquad \textit{and} \qquad \tilde{Y} \approx Y\,,$$

are identically distributed, where we define

$$\tilde{X} = \begin{cases} W \overset{\$}{\leftarrow} D_W \text{ with probability } 1-\delta \\ X' \overset{\$}{\leftarrow} D'_X \text{ with probability } \delta \end{cases} \qquad \text{and} \qquad \tilde{Y} = \begin{cases} W \overset{\$}{\leftarrow} D_W \text{ with probability } 1-\delta \\ Y' \overset{\$}{\leftarrow} D'_Y \text{ with probability } \delta\,. \end{cases}$$

Proof. We prove the lemma for the discrete case, the continuous case can be shown analogously. Assume, without loss of generality, that $\Delta(D_X, D_Y) = \delta$, and that D_X and D_Y are both defined over the same set $\mathcal{S}$. For $s \in \mathcal{S}$ and distribution D_A we write $D_A(s)$ to denote $\Pr_{A \overset{\$}{\leftarrow} D_A}[A = s]$.

We first note that:

$$\begin{aligned}
\delta = \Delta(D_X, D_Y) &= \frac{1}{2} \sum_{a \in \mathcal{S}} |(D_X(a) - D_Y(a))| \\
&= \frac{1}{2} \left(\sum_{\substack{a \in \mathcal{S} \\ D_X(a) \geq D_Y(a)}} (D_X(a) - D_Y(a)) + \sum_{\substack{a \in \mathcal{S} \\ D_Y(a) > D_X(a)}} (D_Y(a) - D_X(a)) \right) \\
&= \frac{1}{2} \left(\sum_{a \in \mathcal{S}} (D_X(a) - \min(D_X(a), D_Y(a))) + \sum_{a \in \mathcal{S}} (D_Y(a) - \min(D_X(a), D_Y(a))) \right) \\
&= \frac{1}{2} \left(2 - 2 \sum_{a \in \mathcal{S}} (\min(D_X(a), D_Y(a))) \right)
\end{aligned}$$

and thus $\sum_{a \in \mathcal{S}}(\min(D_X(a), D_Y(a))) = 1 - \delta$.

We now define D_W as follows:

$$\forall s \in \mathcal{S} \,:\, D_W(s) \overset{\text{def}}{=} \frac{\min(D_X(s), D_Y(s))}{\sum_{a \in \mathcal{S}} \min(D_X(a), D_Y(a))} = \frac{\min(D_X(s), D_Y(s))}{1 - \delta} \,.$$

Similarly, we define D'_X by:

$$\forall s \in \mathcal{S} \,:\, D'_X(s) \overset{\text{def}}{=} \frac{D_X(s) - D_W(s)}{1 - \sum_{a \in \mathcal{S}} D_W(a)} = \frac{D_X(s) - D_W(s)}{\delta} = \begin{cases} 0 & \text{if } D_X(s) < D_Y(s) \\ \frac{D_X(s) - D_Y(s)}{\delta} & \text{otherwise,} \end{cases}$$

and similarly for D'_Y.

We now get that

$$D_{\tilde{X}}(s) = \left\{ \begin{array}{ll} (1-\delta)\frac{\min(D_X(s), D_Y(s))}{1-\delta} + \delta \cdot 0 & \text{if } D_X(s) < D_Y(s) \\ (1-\delta)\frac{\min(D_X(s), D_Y(s))}{1-\delta} + \delta\frac{D_X(s) - D_Y(s)}{\delta} & \text{otherwise} \end{array} \right\} = D_X(s) \,,$$

and similar for $D_{\tilde{Y}}(s)$. □

Using this Lemma, we can prove the following proposition which states that samples from some distribution X can be deterministically turned into samples from another (unknown) distribution Y using some extra auxiliary information whose length depends on the statistical distance of X and Y.

Proposition B.3. *There exists some fixed deterministic function f such that the following holds. For $i \in \{1, \ldots, m\}$, let D_{X_i}, D_{Y_i} be any distributions whose support is of size at most 2^t such that $\Delta(D_{X_i}, D_{Y_i}) \leq \delta$. Define the random variables $\boldsymbol{X} = (X_1, \ldots, X_m), \boldsymbol{Y} = (Y_1, \ldots, Y_m)$ where $X_i \overset{\$}{\leftarrow} D_{X_i}, Y_i \overset{\$}{\leftarrow} D_{Y_i}$ are sampled independently. Then there exists some random variable Z correlated with $\boldsymbol{X}$ such that:*

1. *The output $f(\boldsymbol{X}, Z)$ has the same distribution as $\boldsymbol{Y}$.*
2. *Z is short with high probability. More precisely, if $\delta \leq 1/m$ then the expected bit-length of Z is $\mathbb{E}[|Z|] = (t + \log m)$ and for any $\lambda \geq 6$, we have $\Pr[|Z| \geq \lambda(t + \log m)] \leq 2^{-\lambda}$.*

Proof. For every $i \in \{1, \ldots, m\}$, let $D_{W_i}, D_{X'_i}, D_{Y'_i}$ be the distributions satisfying Lemma B.2.

Following the lemma, we can think of each component X_i of $\boldsymbol{X} = (X_1, \ldots, X_m)$ as being sampled independently as follows: flip a coin $c_i \in \{0, 1\}$ with bias $\Pr[c_i = 0] = \delta$. If $c_i = 0$ sample $X_i \overset{\$}{\leftarrow} D_{X'_i}$, otherwise sample $X_i \overset{\$}{\leftarrow} D_{W_i}$. Notice that the function f does not "know" (depend on) the distributions $D_{W_i}, D_{X'_i}$.

We can define the correlated random variable Z as follows. For every i where $c_i = 0$ the variable Z contains the tuple (i, Y_i) where $Y_i \overset{\$}{\leftarrow} D_{Y'_i}$. The function $f(\boldsymbol{X}, Z)$ simply takes $\boldsymbol{X}$ and, for all indices i contained in Z, replaces X_i with the corresponding Y_i. It is easy to see that, by Lemma B.2 the output of f is identically distributed to $\boldsymbol{Y}$.

It remains to show that Z is short. Each such tuple contained in Z consists of $(\log m + t)$ bits of information. The expected number of such tuples is $\delta \cdot m \leq 1$, and since the probabilities of i being in Z are independent, we can use the Chernoff bound (Lemma 2.5) to prove that the number of tuples in Z is greater than λ with probability only $\leq 2^{-\lambda}$. □

The following lemma shows that, when $\tilde{\mathbf{A}} \overset{\$}{\leftarrow} \mathsf{Lossy}()$ is chosen via the lossy sampler than the LWE samples $\tilde{\mathbf{A}}, \tilde{\mathbf{A}} \cdot \mathbf{s} + \mathbf{e}$ with some sufficiently large noise $\mathbf{e}$ do not reveal too much information about $\mathbf{s}$. This is a direct analogue to Lemma 3.3, which showed the above property for LWR samples.

Lemma B.4. *Let $\ell, n, m, q, \beta, \gamma, \sigma$ be integer parameters and χ a distribution (all parameterized by λ) such that $\Pr_{x \xleftarrow{\$} \chi}[|x| \geq \beta] \leq \mathsf{negl}(\lambda)$ and $\sigma \geq \beta\gamma nm$. Let $\mathit{\Psi}_\sigma$ be either (1) the discrete Gaussian distribution with standard deviation σ, or (2) the uniform distribution over the integer interval $[-\sigma, \sigma]$. Then, for any random variable $\mathbf{s}$ over $[-\gamma, \gamma]^n$ and independently sampled $\tilde{\mathbf{A}} \xleftarrow{\$} \mathsf{Lossy}(1^n, 1^m, 1^\ell, q, \chi)$, $\mathbf{e} \xleftarrow{\$} \mathit{\Psi}_\sigma$, we have:*

$$H_\infty^{\mathsf{smooth}}(\mathbf{s}|\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}) \geq H_\infty(\mathbf{s}) - (\ell + 2\lambda)\log(q)\,.$$

(Recall that $\tilde{\mathbf{A}}$ is computationally indistinguishable from $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$ under the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption.)

Proof. Recall that, by the definition of Lossy, we can write $\tilde{\mathbf{A}} = \mathbf{BC} + \mathbf{F}$ as described in Section 3. Firstly, we can assume that all entries of $\mathbf{F}$ are bounded by β in absolute value, as this modification is statistically close to the original distribution. We can also write $\tilde{\mathbf{A}} \cdot \mathbf{s} + \mathbf{e} = \mathbf{BC} \cdot \mathbf{s} + \mathbf{F} \cdot \mathbf{s} + \mathbf{e}$, where each entry of the vector $\mathbf{F} \cdot \mathbf{s}$ is bounded by $\beta\gamma n$ in absolute value. For any *fixed choice* of the vector $\mathbf{v} = \mathbf{Fs}$, the components of $\mathbf{v} + \mathbf{e}$ and $\mathbf{e}$ are independently distributed and the statistical distance $\mathit{\Delta}(\mathbf{e}_i, \mathbf{v}_i + \mathbf{e}_i) \leq \beta\gamma n/\sigma \leq 1/m$ (see e.g., [GKPV10] for a proof of this when $\mathit{\Psi}_\sigma$ is the discrete Gaussian, and [AJW11] when $\mathit{\Psi}_\sigma$ is uniform over $[-\sigma, \sigma]$). Therefore, using Proposition B.3, there is some universal function f and some random variable $Z(\mathbf{v})$ that is correlated with $\mathbf{e}$ such that $f(\mathbf{e}, Z(\mathbf{v})) \approx \mathbf{v} + \mathbf{e}$ are identically distributed and the bit-length of $Z(\mathbf{v})$ is $\leq \lambda(\log(m) + \log(q))$ with overwhelming probability. This gives us the distributional equivalence:

$$(\tilde{\mathbf{A}}, \tilde{\mathbf{A}} \cdot \mathbf{s} + \mathbf{e}) \approx (\mathbf{BC} + \mathbf{F}, \mathbf{BC} \cdot \mathbf{s} + \mathbf{F} \cdot \mathbf{s} + \mathbf{e}) \approx (\mathbf{BC} + \mathbf{F}, \mathbf{BC} \cdot \mathbf{s} + f(\mathbf{e}, Z(\mathbf{F} \cdot \mathbf{s})))$$

Therefore:

$$\begin{aligned} H_\infty^{\mathsf{smooth}}(\mathbf{s}|\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}) &\geq H_\infty^{\mathsf{smooth}}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{e}, \mathbf{C} \cdot \mathbf{s}, Z(\mathbf{F} \cdot \mathbf{s})) \\ &\geq H_\infty^{\mathsf{smooth}}(\mathbf{s}|\mathbf{B}, \mathbf{C}, \mathbf{F}, \mathbf{e}) - \ell \log q - \lambda(\log(m) + \log(q)) \\ &\geq H_\infty^{\mathsf{smooth}}(\mathbf{s}) - (\ell + 2\lambda)\log q. \end{aligned}$$

This concludes the proof. □

We can now formulate an analogue to Theorem 4.1, stating that the LWE-assumption is also satisfied for weak and leaky keys:

Theorem B.5. *Let $k, \ell, m, n, \beta, \gamma, \sigma, q$ be integer parameters and χ a distribution (all parameterized by λ) such that $\Pr_{x \xleftarrow{\$} \chi}[|x| \geq \beta] \leq \mathsf{negl}(\lambda)$ and $\sigma \geq \beta\gamma nm$. Let $\mathit{\Psi}_\sigma$ be either (1) the discrete Gaussian distribution with standard deviation σ or (2) the uniform distribution over the integer interval $[-\sigma, \sigma]$. Assuming that the $\mathrm{LWE}_{\ell,m,q,\chi}$ assumption holds, the weak and leaky $\mathrm{LWE}^{\mathsf{WL}(\gamma,k)}_{n,m,q,\mathit{\Psi}_\sigma}$-assumption holds if $k \geq (\ell + \mathit{\Omega}(\lambda))\log(q)$.*

We omit the proof as it follows directly from Lemma B.4, in the exact same way that Theorem 4.1 follows from Lemma 3.3.