A Lower Bound for Proving Hardness of Learning with Rounding with Polynomial Modulus

Parker Newton* Silas Richelson†

Abstract

Regev's Learning with Errors (LWE) problem (STOC 2005) is a fundamental hardness assumption for modern cryptography. The Learning with Rounding (LWR) Problem was put forth by Banerjee, Peikert and Rosen (Eurocrypt 2012) as an alternative to LWE, for use in cryptographic situations which require determinism. The only method we currently have for proving hardness of LWR is the so-called "rounding reduction" which is a specific reduction from an analogous LWE problem. This reduction works whenever the LWE error is small relative to the noise introduced by rounding, but it fails otherwise. For this reason, all prior work on establishing hardness of LWR forces the LWE error to be small, either by setting other parameters extremely large (which hurts performance), or by limiting the number of LWR samples seen by the adversary (which rules out certain applications). Hardness of LWR is poorly understood when the LWE modulus (q) is polynomial and when the number of LWE samples (m) seen by the adversary is an unbounded polynomial. This range of parameters is the most relevant for practical implementations, so the lack of a hardness proof in this situation is not ideal.

In this work, we identify an obstacle for proving the hardness of LWR from LWE in the above framework when q is polynomial and m is an unbounded polynomial. Specifically, we show that any "pointwise" reduction from LWE to LWR (*i.e.*, any reduction which maps LWE samples independently to LWR samples) admits an efficient algorithm which directly solves LWE (without the use of an LWR solver). Consequently, LWE cannot be reduced to LWR in our pointwise reduction model with our setting of q and m, unless LWE is easy. Our model of a pointwise reduction from LWE to LWR captures all prior reductions from LWE to LWR except the rejection sampling reduction of Bogdanov *et al.* (TCC 2016); while their reduction still operates in a pointwise manner, it can reject an LWE sample instead of mapping it to an LWR sample. However we conjecture that our result still holds in this setting.

Our argument proceeds roughly as follows. First, we show that any pointwise reduction from LWE to LWR must have good agreement with some affine map. Then, we use the affine agreement of a pointwise reduction together with a type of Goldreich-Levin "prediction-implies-inversion" argument to extract the LWE secret from LWE input samples. Both components may be of independent interest.

*University of California, Riverside. Email: pnewt001@ucr.edu. Work completed prior to joining Amazon.
†University of California, Riverside. Email: silas@cs.ucr.edu.

1 Introduction

Regev's learning with errors (LWE) problem [Reg05] is fundamental for modern cryptography due to its versatility and strong security guarantees. LWE asks an algorithm to solve a random noisy linear system of equations mod q: given integers n, q, m, an "error" distribution χ on $\mathbb{Z}_q$ and a uniform $\mathbf{s} \sim \mathbb{Z}_q^n$, recover $\mathbf{s}$ given samples

$$\left\{(\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i)\right\} \subset \left(\mathbb{Z}_q^n \times \mathbb{Z}_q\right)^m, \tag{1}$$

where the $\mathbf{a}_i$ are drawn uniformly from $\mathbb{Z}_q^n$ and the e_i are drawn according to χ. It is known that when q is sufficiently large compared to n, there are error distributions which make solving LWE efficiently given any number of samples as hard as solving computational problems on lattices in the worst case [Reg05, Pei09, BLP+13]; such problems are conjectured to be hard even for quantum computers. In addition to the strong hardness guarantees, LWE has proven to be extremely useful for cryptography. Since its introduction, an immense research effort has established LWE-based constructions for most known cryptographic primitives (*e.g.*, [GPV08, ACPS09, BGV11, MP12, GSW13, PS19] and many, many more).

The randomness inherent to the LWE problem (*i.e.*, the randomness used to draw the $e_i \sim \chi$) precludes straightforward constructions of certain cryptographic primitives which require determinism, such as PRFs. Banarjee, Peikert and Rosen [BPR12] introduced the learning with rounding (LWR) problem in order to overcome this obstacle. LWR asks an algorithm to solve a random linear system with "deterministic noise": given n, p, q, m with $p < q$ and a uniform $\mathbf{s} \sim \mathbb{Z}_q^n$, recover $\mathbf{s}$ from

$$\left\{(\mathbf{a}_i, b_i = \lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle \rceil_p)\right\} \subset \left(\mathbb{Z}_q^n \times \mathbb{Z}_p\right)^m, \tag{2}$$

where each $\mathbf{a}_i \sim \mathbb{Z}_q^n$ and where $\lfloor \cdot \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$ is the function which, given $x \in \mathbb{Z}_q$, outputs the nearest integer to px/q. Since its introduction, LWR has been used in numerous works to give cryptographic constructions where determinism is mandatory (*e.g.*, [BPR12, BLL+15, BV15], and more).

Hardness of LWR is established via the following reduction from LWE: given an LWE sample $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, round the second value and output $(\mathbf{a}, \lfloor b \rceil_p) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$. In [BPR12], it is shown that this reduction is valid whenever $q/p = n^{\omega(1)}$ (n the security parameter), and so establishes hardness of LWR for this parameter regime. In practice we would like to be able to use small q as this lends itself better to efficient implementations. So establishing hardness for LWR in the "polynomial modulus" setting, where $q = \mathsf{poly}(n)$, was an important open problem left by [BPR12]. This direction was pursued in the follow-up works [AKPW13, BGM+16, AA16] where it is shown that if the number of LWR samples given to the solver (*i.e.*, m) is bounded, then the correctness proof of the above reduction goes through and one can establish hardness of LWR with polynomial modulus in the "bounded sample" setting. This is good enough for some cryptographic applications [AKPW13], but not for all, *e.g.*, PRFs.

The problem with the above reduction when q/p is small is that the error in the LWE sample might cause the rounding function to make a mistake. The reason for this is that the "threshold points" of the rounding function[1] $\lfloor \cdot \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$ have density p/q in $\mathbb{Z}_q$, and so when $q/p \ll m$, some of the $\mathbf{a}_i$'s chosen will be such that their secret inner product $\langle \mathbf{a}_i, \mathbf{s} \rangle$ is close to a threshold point. Whenever this occurs, the reduction will make an error if $\langle \mathbf{a}_i, \mathbf{s} \rangle + e_i$ is on the opposite side of the threshold from $\langle \mathbf{a}_i, \mathbf{s} \rangle$. Prior work handles this issue by forcing q/p to be large relative to m (either by setting q/p to be superpolynomial, or by bounding m).

[1] By threshold points we mean the half integer multiples of q/p where the rounding function switches from rounding to adjacent values in $\mathbb{Z}_p$.

Getting a version of the above reduction to yield a hardness proof for LWR in the case when m is large compared to q/p is challenging because it requires dealing with situations where the LWE error creates a rounding problem. By definition, a reduction from LWE to LWR is an oracle algorithm which solves LWE when instantiated with access to any LWR solver, *including the pathological LWR solver who aborts whenever it sees a rounding error.* Specifically, suppose S is an algorithm which takes m LWR samples $\{(\mathbf{a}_i, b'_i)\} \subset \mathbb{Z}_q \times \mathbb{Z}_p$, (somehow) recovers the hidden secret $\mathbf{s}$, then scans the m samples to make sure that $b'_i = \left\lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle \right\rceil_p$ for all i, aborting if it finds an error, outputting $\mathbf{s}$ otherwise. It is clear that S will solve LWR when it is given true LWR samples, however in order for the reduction to make use of S's solving power to solve LWE, it must produce m LWR samples without making an error. This is the core challenge in proving hardness of LWR with polynomial modulus and unbounded samples.

1.1 Our Contribution

In this work we convert the above difficulty into a lower bound for proving hardness of LWR with polynomial modulus and an unbounded number of samples via reductions from LWE. Our barrier applies to any "pointwise" reduction from LWE to LWR, *i.e.*, any function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$. This includes and broadly extends the reduction $(\mathbf{a}, b) \mapsto (\mathbf{a}, \lfloor b \rceil_p)$ mentioned above. The starting observation for our work is that any pointwise reduction f which works in this parameter regime must implicitly be able to handle the "problematic" LWE pairs which are close to a rounding threshold. What we prove is essentially that f's understanding of how to handle these threshold samples can be *extracted* in the form of knowledge about the LWE secret. Our main theorem is the following.

Theorem 1 (Informal). *Let $n, q, p \in \mathbb{N}$ be integers such that $q = \mathsf{poly}(n)$ is prime and such that $q^{2/3+c} < p < q$ for a small constant $c > 0$. Let χ be an error distribution on $\mathbb{Z}_q$. Suppose an efficient function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Then f can be used to design an efficient algorithm which solves $\mathsf{LWE}_{n,q,\chi}$.*

The Hypotheses of our Theorem. We view the requirements that q be prime and especially that $q^{2/3+c} < p$ as shortcomings of our work, and we believe it should be possible to improve our result to remove these extra hypotheses. Our proof requires q to be prime so that linear algebra works on the set $\mathbb{Z}_q^n$. The lower bound on p comes from one place in the proof where we use two LWE samples $(\mathbf{a}_0, b_0), (\mathbf{a}_1, b_1) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ to generate three LWR samples:

$$(\mathbf{a}'_0, b'_0) = f(\mathbf{a}_0, b_0);\ (\mathbf{a}'_1, b'_1) = f(\mathbf{a}_1, b_1);\ (\mathbf{a}'_2, b'_2) = f(\mathbf{a}_0 + \mathbf{a}_1, b_0 + b_1) \in \mathbb{Z}_q^n \times \mathbb{Z}_p,$$

and we require essentially that the three output values $b'_0, b'_1, b'_2 \in \mathbb{Z}_p$ contain more information than the input values $b_0, b_1 \in \mathbb{Z}_q$. We suspect that a different proof technique could be used to improve the lower bound required of p or remove it altogether. We note however that our result does not require the amount of LWR "noise" (*i.e.*, q/p) to be small relative to the amount of LWE noise. In particular, our theorem applies in situations where q/p is much larger than the standard deviation of the discrete Gaussian used for the LWE noise.

Aborting Pointwise Reductions. Another way to relax the hypotheses of our main theorem would be to allow f to abort. In this case, the reduction works by applying the aborting pointwise function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \left(\mathbb{Z}_q^n \times \mathbb{Z}_p\right) \cup \{\perp\}$ to all LWE samples, and then invoking the LWR solver on all "non-bot" outputs. Such "aborting pointwise reductions" were considered in prior work [BGM+16] as a way to prove hardness of LWR in the polynomial modulus setting assuming hardness of LWE with

uniform errors (uLWE). The key intuition behind the proof of our main theorem is that it is impossible for the reduction to correctly produce the LWR distribution given LWE samples because doing this would require converting the LWE error into the "rectangular" LWR error. In uLWE, the errors are already rectangular and the only difference between uLWE and LWR is that in uLWE the rectangles are centered around $\langle \mathbf{a}, \mathbf{s} \rangle$ while in LWR they are centered around the rounding points. Bogdanov *et al.* [BGM+16] showed that it is possible to fix this "rectangle center discrepancy" using rejection sampling and obtained an aborting pointwise reduction from uLWE to LWR. An interesting question is: *how much power do we get by allowing the reduction to abort?* Does aborting allow transforming Gaussian LWE errors to the rectangular LWR errors? Or does aborting *only* allow us to reposition the centers of the error distribution, *and not* convert non-rectangular errors to rectangular errors? We tend to believe that aborting reductions can only translate the errors, and cannot convert non-rectangular errors into rectangular errors. However, several parts of our proof break down if we allow the function to abort. We state the following conjecture.

Conjecture 1. *Let $n, p, q \in \mathbb{N}$ be integers such that such that $q = \mathsf{poly}(n)$ is prime and $2 \leqslant p < q$. Let χ be a discrete Gaussian distribution on $\mathbb{Z}_q$. Suppose an efficient function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \left(\mathbb{Z}_q^n \times \mathbb{Z}_p\right) \cup \{\perp\}$ is part of an aborting pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Then f can be used to build an efficient algorithm $\mathcal{B}$ which solves $\mathsf{LWE}_{n,q,\chi}$.*

Compared to Theorem 1, Conjecture 1 removes the hypothesis that $q^{2/3+c} < p$ and allows f to abort, though focuses in on the case when χ is a discrete Gaussian. We tend to believe Conjecture 1 is true for any non-rectangular error distribution χ, in which case it combines with (a slight extension of) Theorem 5 of [BGM+16] to give a dichotomy: if $q = \mathsf{poly}(n)$ is prime and if there is an aborting pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ for $2 \leqslant p < q$ then either 1) χ is rectangular; or 2) there is an efficient algorithm which solves $\mathsf{LWE}_{n,q,\chi}$.

Extensions of our Reduction Model. One can ask whether our reduction holds for other extensions of our reduction model. For example, does our theorem hold for pointwise reductions between problems with different dimensions and moduli (*i.e.*, reductions from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n',q',p'}$)? Furthermore, our notion of pointwise reductions does not allow the reduction to use two or more LWE samples to produce an LWR sample. One might hope that a similar theorem to ours would hold for any "$k-$to$-$one" function $f : \left(\mathbb{Z}_q^n \times \mathbb{Z}_q\right)^k \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ as long as k is small enough to ensure that $\mathbf{s}$ has sufficient entropy given k LWE samples. Note that if k is large enough so that k LWE samples determine $\mathbf{s}$ information theoretically, then one could imagine a function f which takes k LWE samples, (somehow) recovers $\mathbf{s}$, and outputs a single LWR sample with secret $\mathbf{s}$. While it feels like such a function is breaking LWE, it would be hard to prove a theorem like the above since it seems that in order to extract any knowledge about the LWE secret, one would have to solve LWR.

Interpreting our Result. Our main theorem identifies a barrier to proving the hardness of LWR in certain practical parameter regimes via reductions from LWE. This explains, to some extent, why this problem has remained open for so long. Our result **does not** suggest that LWR is easy. Rather, it speaks to the fact that the current techniques we have available for deriving hardness from worst-case lattice problems are inherently probabilistic. Our work indicates that a reduction from a hard lattice problem to LWR with these parameter settings would be extremely interesting as it would likely contain significant new ideas.

1.2 Technical Overview

We now give a summary of our proof of Theorem 1 which says that a pointwise reduction from LWE to LWR can be used to design an algorithm which solves LWE. The proof consists of three main parts. First, we derive some basic combinatorial structure about the pointwise function which it must satisfy if it is to be part of a reduction. Next, building on this basic structure we show that in fact the pointwise function must very close to an affine function. Finally, we show how to use a pointwise function which has good affine agreement and which is part of a reduction to directly solve LWE.

Notation. Let $n, q, p \in \mathbb{N}$ such that $q = \mathsf{poly}(n)$ is prime, and $q^{2/3+c} < p < q$, for a small constant $c > 0$. Let χ be an error distribution on $\mathbb{Z}_q$. Let $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ be part of a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ (formal definition is in Section 3). If $\mathbf{s} \in \mathbb{Z}_q^n$, then let $\mathsf{LWE}_{\mathbf{s}}$ denote the distribution which chooses $\mathbf{a} \sim \mathbb{Z}_q^n, e \sim \chi$, and outputs $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. Likewise, if $\mathbf{s}' \in \mathbb{Z}_q^n$, then $\mathsf{LWR}_{\mathbf{s}'}$ draws $\mathbf{a}' \sim \mathbb{Z}_q^n$ and outputs $\left(\mathbf{a}', b' = \lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \rceil_p\right) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$. We write $(\mathbf{a}', b') \in \mathsf{LWR}_{\mathbf{s}'}$ if $b' = \left\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \right\rceil_p$. Finally, for $m \in \mathbb{N}$, let LWE_m (resp. LWR_m) be the distribution which draws $\mathbf{s} \sim \mathbb{Z}_q^n$ (resp. $\mathbf{s}' \sim \mathbb{Z}_q^n$) and outputs m samples from LWE_s (resp. $\mathsf{LWR}_{\mathbf{s}'}$).

Establishing Basic Combinatorial Structure of f. The key observation which allows us to get started imposing structure on f is the following: *all statistics of the LWR distribution and the output distribution of f (given LWE samples as input) must be the same*. Indeed, if there is a statistic which differs between LWR_m and $f(\mathsf{LWE}_m)$, we can conceive of a "pathological LWR solver" which draws enough samples to approximate the statistic, aborting if it decides it is being fed with mapped LWE samples, solving if it decides it is being fed with true LWR samples.

For example, for all $\mathbf{s}' \in \mathbb{Z}_q^n$, clearly, $\Pr_{(\mathbf{a}', b') \sim \mathsf{LWR}_{\mathbf{s}'}}\left[(\mathbf{a}', b') \in \mathsf{LWR}_{\mathbf{s}'}\right] = 1$. Thus, if f is a reduction then the following *correctness condition* must hold: with non-negligible probability over $\mathbf{s} \sim \mathbb{Z}_q^n$ there must exist some $\mathbf{s}' \in \mathbb{Z}_q^n$ such that

$$\Pr_{(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}}\left[f(\mathbf{a}, b) \in \mathsf{LWR}_{\mathbf{s}'}\right] = 1 - \mathsf{negl}(n).$$

If not, then consider the "pathological LWR solver" S which, given $(\mathbf{a}'_1, b'_1), \ldots, (\mathbf{a}'_m, b'_m) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$ statistically recovers $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $(\mathbf{a}'_i, b'_i) \in \mathsf{LWR}_{\mathbf{s}'}$ for all $i = 1, \ldots, m$ and outputs $\mathbf{s}'$ (outputting $\perp$ if no such $\mathbf{s}' \in \mathbb{Z}_q^n$ exists). Note, S does indeed solve LWR when fed with samples from LWR_m, however S outputs $\perp$ with high probability when fed with samples from $f(\mathsf{LWE}_m)$. This means that f is not a reduction since it is unable to make use of S's LWR solving power. The fact that f induces a mapping on secrets $\mathbf{s} \mapsto \mathbf{s}'$ (meaning that $f(\mathbf{a}, b) \in \mathsf{LWR}_{\mathbf{s}'}$ holds with high probability over $(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}$) turns out to be immensely useful, as we will already see throughout the remainder of this overview. In Section 4, we use analogous "pathological solver" arguments to establish this and several other combinatorial properties of f which will be useful throughout the remainder of the paper.

Establishing High Affine Agreement of f. After establishing some basic statistics of f, the technical core of our paper involves proving that f has high agreement with an affine function. More specifically, we algorithmically recover a matrix $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ of rank $n - \mathcal{O}(1)$ and a constant dimensional vector space $\mathbf{V} \subset \mathbb{Z}_q^n$ such that

$$\Pr_{(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q}\left[\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}\right] \geqslant 1 - \eta,$$

for a small parameter $\eta > 0$, where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. We prove this in two stages. First, we recover $\mathbf{V} \subset \mathbb{Z}_q^n$ of constant dimension such that f passes the following test with high probability.

- Draw $(\mathbf{a}_0, b_0), (\mathbf{a}_1, b_1) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$.
- Set $(\mathbf{a}_2, b_2) = (\mathbf{a}_0 + \mathbf{a}_1, b_0 + b_1)$ and $(\mathbf{a}_i', b_i') = f(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$ for $i = 0, 1, 2$.
- Pass if $\mathbf{a}_2' \in \mathrm{Span}\big(\{\mathbf{a}_0', \mathbf{a}_1'\}\big) + \mathbf{V}$, fail if not.

The idea here is that when f does not pass this test, it is using three linearly dependent relations about the LWE secret to generate three linearly independent relations about the LWR secret. Either this behavior must be extremely unlikely, or it must be that the map $\mathbf{s} \mapsto \mathbf{s}'$ mapping $\mathbf{s} \in \mathbb{Z}_q^n$ to $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\Pr_{(\mathbf{a},b)\sim \mathsf{LWE}_{\mathbf{s}}}\big[f(\mathbf{a}, b) \in \mathsf{LWR}_{\mathbf{s}'}\big] = 1 - \mathsf{negl}(n)$ is many-to-one, which we show is impossible in Section 4.

We then prove a property testing-type result showing that any function f which passes this test with high probability must be close to some linear map in the above sense. For this we use techniques from the proof of the following "fundamental theorem of projective geometry" which says that any function $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ which maps lines to lines must be affine.

Proposition 1 (FTPG − [Art57], Section 2.10). *Let q be a prime and $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ be a function such that for any one-dimensional line $\ell \subset \mathbb{Z}_q^n$, the set $h(\ell) := \big\{f(\mathbf{x}) : \mathbf{x} \in \ell\big\} \subset \mathbb{Z}_q^n$ is also a line. Then h is affine.*

To see the connection between Proposition 1 and our setting, note that high probability of passing the above test in the simplified setting where $\mathbf{V} = \{\mathbf{0}\}$ means that f is mapping $2-$planes to $2-$planes with high probability. The techniques used in this part of our analysis may be of independent interest.

Solving LWE Using an Affine Reduction. Finally, once we know that f has good affine agreement, we can use f to recover the LWR secret of the output samples using a Goldreich-Levin-type argument. Assume for simplicity that $\mathbf{a}' = \mathbf{H}\mathbf{a}$, rather than $\mathbf{a}' \in \mathrm{Span}(\mathbf{H}\mathbf{a}) + \mathbf{V}$ holds with high probability over $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$ where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. The key point is that if $(\mathbf{a}', b') \in \mathsf{LWR}_{\mathbf{s}'}$ also holds with high probability over $(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}$, then

$$b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \big\rceil_p = \big\lfloor \langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle \big\rceil_p,$$

and so b' allows us to predict the inner product $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$ with non-negligible advantage over guessing randomly (simply by drawing $x \sim \mathbb{Z}_q$ such that $\lfloor x \rceil_p = b'$). The Goldreich-Levin machinery can then be used to recover $\mathbf{H}^{\mathsf{t}}\mathbf{s}'$, which will allow recovering $\mathbf{s}'$ with non-negligible probability since $\mathbf{H}$ has nearly full rank.

Putting Everything Together. Suppose when playing the distinguishing game for LWE we are given samples $(\mathbf{a}_1, b_1), \ldots, (\mathbf{a}_m, b_m) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ which are either drawn from $\mathsf{LWE}_{\mathbf{s}}$ for a uniform $\mathbf{s} \sim \mathbb{Z}_q^n$, or from $\mathbb{Z}_q^n \times \mathbb{Z}_q$. We can use f to distinguish as follows.

- Let $(\mathbf{a}_i', b_i') = f(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$ for $i = 1, \ldots, m$.
- If the LWR secret reconstruction procedure succeeds in obtaining $\mathbf{s}' \in \mathbb{Z}_q^n$ proceed, if not output a random bit.

- Now check whether $(\mathbf{a}'_i, b'_i) \in \mathsf{LWR}_{\mathbf{s}'}$ holds for all $i = 1, \ldots, m$; if so output 0 (corresponding to LWE samples), otherwise output a random bit.

As discussed above, if we have been fed with LWE samples, then the LWR recovery procedure will work with non-negligible probability. On the other hand, a simple statistical argument (proven in Section 4) shows that there cannot exist $\mathbf{s}' \in \mathbb{Z}_q^n$ such that f maps uniform samples into $\mathsf{LWE}_{\mathbf{s}'}$ with high probability.

2 Preliminaries

Throughout this work, the integer n will denote the security parameter. We use boldface lower case for vectors, and boldface capitals for matrices (*e.g.*, $\mathbf{v}$ or $\mathbf{M}$). Given a distribution χ on a set X, we write $x \sim \chi$ to indicate that $x \in X$ is drawn according to χ; we write $x \sim X$ as shorthand for $x \sim \mathsf{Unif}(X)$, the uniform distribution on X.

2.1 Learning with Errors/Rounding

Definition 1 (**The LWE/LWR Distributions).** *Let $n, q \in \mathbb{N}$, $\mathbf{s} \in \mathbb{Z}_q^n$, and χ be a distribution on $\mathbb{Z}_q$.*

- **The LWE Distribution:** *The* learning with errors distribution $\mathsf{LWE}_{n,q,\mathbf{s},\chi}$ *works as follows:*
  - *draw $\mathbf{a} \sim \mathbb{Z}_q^n$, $e \sim \chi$, set $b = \langle \mathbf{a}, \mathbf{s} \rangle + e$ and output $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.*
- **The LWR Distribution:** *If $p \in \mathbb{N}$ such that $2 \leqslant p < q$, then the* learning with rounding distribution $\mathsf{LWR}_{n,q,\mathbf{s},p}$ *is:*
  - *draw $\mathbf{a} \sim \mathbb{Z}_q^n$, set $b = \left\lfloor \langle \mathbf{a}, \mathbf{s} \rangle \right\rceil_p$, and output $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$, where $\lfloor \cdot \rceil_p$ maps $x \in \mathbb{Z}_q$ to $x \cdot (p/q)$ rounded to the nearest integer* $(\mathrm{mod}\ p)$.

Given $m \in \mathbb{N}$, the distributions $\mathsf{LWE}_{n,q,m,\chi}$ (resp. $\mathsf{LWR}_{n,q,m,p}$) work by drawing $\mathbf{s} \sim \mathbb{Z}_q^n$ and outputting m independent samples from $\mathsf{LWE}_{n,q,\mathbf{s},\chi}$ (resp. $\mathsf{LWR}_{n,q,\mathbf{s},p}$).

Definition 2 (**The LWE/LWR Problems).** *Let $n, q, m \in \mathbb{N}$ and χ be a distribution on $\mathbb{Z}_q$. The* search/decisional *version of the* learning with errors/rounding problems *refer to the following computational tasks.*[2]

- **Search LWE:** *Given $\left\{(\mathbf{a}_1, b_1), \ldots, (\mathbf{a}_m, b_m)\right\} \sim \mathsf{LWE}_{n,q,m,\chi}$, output* $\mathbf{s}$.
- **Decisional LWE:** *Distinguish $\mathsf{LWE}_{n,q,m,\chi}$ from $\mathsf{Unif}\left(\mathbb{Z}_q^n \times \mathbb{Z}_q\right)^m$.*
- **Search LWR:** *If $p \in \mathbb{N}$ such that $2 \leqslant p < q$, then given $\left\{(\mathbf{a}_1, b_1), \ldots, (\mathbf{a}_m, b_m)\right\} \sim \mathsf{LWR}_{n,q,m,p}$, output* $\mathbf{s}$.

Error Distributions and Rounding Subsets. The most common choice for the error distribution χ is a discrete Gaussian on $\mathbb{Z}_q$, centered at 0 with standard deviation αq for some $\alpha = 1/\mathsf{poly}(n)$. Hardness of decisional LWE with this error distribution is known assuming worst-case hardness of computational problems on lattices which are believed to be hard even for quantum computers [Reg05, Pei09, BLP+13]. The arguments in this work will apply equally well to any bounded error distribution which gives output in $\{-B, \ldots, B\} \subset \mathbb{Z}_q$ for $B \ll q$ with overwhelming probability $1 - 2^{-n}$.

[2] We will not need the decisional version of LWR in this work, so we do not give the definition.

Solvers and Distinguishers. Given $\varepsilon > 0$ and $m \in \mathbb{N}$, we say an algorithm S is an $(\varepsilon, m)-$*solver* for $\mathsf{LWE}_{n,q,\chi}$ (resp. $\mathsf{LWR}_{n,q,p}$) if it solves search LWE (resp. search LWR) with probability at least ε, given m samples:

$$\Pr_{\{(\mathbf{a}_i,b_i)\}_{i=1}^m \sim \mathsf{LWE}_{n,q,m,\chi}}\Big[\mathsf{S}\big(\{(\mathbf{a}_i,b_i)\}_{i=1}^m\big) = \mathbf{s}\Big] \geqslant \varepsilon,$$

and similarly for $\mathsf{LWR}_{n,q,m,p}$ except the probability is over $\{(\mathbf{a}_i,b_i)\}_{i=1}^m \sim \mathsf{LWR}_{n,q,m,p}$. Likewise, we say that an algorithm D is an $(\varepsilon, m)-$*distinguisher* for $\mathsf{LWE}_{n,q,\chi}$ if

$$\left|\Pr_{\{(\mathbf{a}_i,b_i)\}_{i=1}^m \sim \mathsf{LWE}_{n,q,m,\chi}}\Big[\mathsf{D}\big(\{(\mathbf{a}_i,b_i)\}_i\big) = 1\Big] - \Pr_{\{(\mathbf{a}_i,b_i)\}_{i=1}^m \sim (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m}\Big[\mathsf{D}\big(\{(\mathbf{a}_i,b_i)\}\big) = 1\Big]\right| \geqslant \varepsilon.$$

We write the inputs to solvers and distinguishers as sets even though technically speaking they are lists: they have an ordering and they can contain duplicated elements (this distinction will not matter for us).

Definition 3 (Reduction from LWE to LWR). *Let $n, q, p \in \mathbb{N}$ be integers with $p < q$, and let χ be a distribution on $\mathbb{Z}_q$, and let $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{R}_{>0}$ and $\ell_{\mathsf{samp}} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{N}$ be functions. We say that a PPT oracle algorithm $\mathcal{A}$ is an* $(\ell_{\mathsf{err}}, \ell_{\mathsf{samp}})-$reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ *if the following holds: if S is an $(\varepsilon', m')-$ solver for $\mathsf{LWR}_{n,q,p}$, then $\mathcal{A}^{\mathsf{S}}$ (*i.e.*, $\mathcal{A}$ instantiated with oracle access to S) is an $(\varepsilon, m)-$solver for $\mathsf{LWE}_{n,q,\chi}$, where $(\varepsilon, m) = \big(\ell_{\mathsf{err}}(\varepsilon', m'), \ell_{\mathsf{samp}}(\varepsilon', m')\big)$.*

Remark. We are interested in noticeable solvers which run in polynomial time; *i.e.*, $(\varepsilon', m')-$solvers for $\varepsilon' = \mathsf{poly}\big(1/n\big)$ and $m' = \mathsf{poly}(n)$. In order to preserve this, our reductions will always have $\ell_{\mathsf{err}}(\varepsilon', m') = \mathsf{poly}\big(1/n, \varepsilon', 1/m'\big)$ and $\ell_{\mathsf{samp}}(\varepsilon', m') = \mathsf{poly}(n, 1/\varepsilon', m')$. Thus, our reduction model requires $\mathcal{A}^{\mathsf{S}}$ to be a polynomial time noticeable solver for LWE whenever S is a polynomial time noticeable solver for LWR. As mentioned in the introduction, several prior works [AKPW13, BLL+15, BGM+16] prove hardness results for LWR with $q = \mathsf{poly}(n)$ via LWE hardness as long as there is a bound B on the overall number of samples given to the LWR solver. In the above language, these works give a reduction $\mathcal{A}$ such that $\mathcal{A}^{\mathsf{S}}$ is a polytime noticeable solver for LWE whenever S is a polytime noticeable solver for LWR which uses $m' \leqslant B$ samples.

2.2 Pseudorandomness

Definition 4 (Statistical Distance). *Let X and Y be random variables, both supported on the same set Ω. The* statistical distance between X and Y*, denoted $\Delta(X, Y)$, is equal to both of the following expressions:*

$$\max_{T \subset \Omega} \Big|\Pr_{x \sim X}\big[x \in T\big] - \Pr_{y \sim Y}\big[y \in T\big]\Big| = \frac{1}{2} \cdot \sum_{z \in \Omega} \Big|\Pr_{x \sim X}\big[x = z\big] - \Pr_{y \sim Y}\big[y = z\big]\Big|.$$

We will use a version of the fact that the inner product mod q is a good two-source extractor. The original proof of this fact when $q = 2$ is in [CG88]; see [LLTT05] for the following generalization to larger prime q.

Fact 1. *Let $n, q \in \mathbb{N}$ be such that q is prime, let $X \subset \mathbb{Z}_q^n$ be a subset, and let $\mathcal{D}$ be the distribution on $\mathbb{Z}_q^{n+1}$ which draws $\mathbf{a} \sim \mathbb{Z}_q^n$, $\mathbf{x} \sim X$ and outputs $\big(\mathbf{a}, \langle \mathbf{a}, \mathbf{x} \rangle\big)$. Then*

$$\Delta\big(\mathcal{D}, \mathsf{Unif}(\mathbb{Z}_q^{n+1})\big)^2 \leqslant \frac{q}{4|X|}.$$

The following corollary will be used several times throughout the paper. Intuitively, it says that any property which holds with good probability over $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$ holds with similar probability over $(\mathbf{a}, b) \sim \mathsf{LWE}_{n,q,\mathbf{s},\chi}$ for almost all $\mathbf{s} \in \mathbb{Z}_q^n$.

Corollary 1 (Sampling of LWE). *For any test set $T \subset \mathbb{Z}_q^n \times \mathbb{Z}_q$ of size $|T| = \tau \cdot q^{n+1}$, and any $e \in \mathbb{Z}_q$,*

$$\Pr_{\mathbf{s}\sim\mathbb{Z}_q^n}\left[\left|\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\left[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in T\right] - \tau\right| > q^{-n/4}\right] = q^{-\Omega(n)}.$$

In particular,

$$\Pr_{\mathbf{s}\sim\mathbb{Z}_q^n}\left[\left|\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\left[(\mathbf{a}, b) \in T\right] - \tau\right| > q^{-n/4}\right] = q^{-\Omega(n)}.$$

Proof. Fix $T \subset \mathbb{Z}_q^n \times \mathbb{Z}_q$ of size $|T| = \tau \cdot q^{n+1}$, and let $S \subset \mathbb{Z}_q^n$ be the set of $\mathbf{s} \in \mathbb{Z}_q^n$ such that $\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\left[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in T\right] > \tau + q^{-n/4}$ for some $e \in \mathbb{Z}_q$. We will prove $|S| < q^{n/2+3} = q^{-(n/2-3)} \cdot q^n$; the result follows since we can argue similarly for the set of $\mathbf{s} \in \mathbb{Z}_q^n$ such that for some $e \in \mathbb{Z}_q$, $\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\left[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in T\right] < \tau - q^{-n/4}$. For the part of the claim about LWE samples, note that if $\mathbf{s} \notin S$ then

$$\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\left[(\mathbf{a}, b) \in T\right] = \sum_{e\in\mathbb{Z}_q} \Pr\left[\chi = e\right] \cdot \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\left[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in T\right] \leqslant \tau + q^{-n/4}.$$

So it suffices to bound $|S|$. Let $S_e \subset S$ be the $\mathbf{s} \in S$ such that $\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\left[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in T\right] > \tau + q^{-n/4}$. For all $e \in \mathbb{Z}_q$, we have

$$\tau + q^{-n/4} < \Pr_{\mathbf{s}\sim S_e, \mathbf{a}\sim\mathbb{Z}_q^n}\left[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle) + (0, e) \in T\right] \leqslant \tau + \sqrt{\frac{q}{4|S_e|}},$$

where the inequality on the second line is Fact 1. Thus, $|S_e| \leqslant q^{n/2+1}/4$ holds for all $e \in \mathbb{Z}_q$, and so $|S| = \left|\bigcup_e S_e\right| \leqslant q^{n/2+2}$. The result follows. □

3 Our Reduction Model and Main Theorem

3.1 Pointwise Reductions and Main Theorem Statement

In this section we define *pointwise reductions from LWE to LWR*, which are the reductions ruled out by our main theorem. To say that $\mathcal{A}$ is a pointwise reduction is to require that the LWE solver $\mathcal{A}^{\mathsf{S}}$ uses its oracle access to S in a precise way. First, $\mathcal{A}^{\mathsf{S}}$ must map its input LWE samples to LWR samples in a pointwise fashion (*i.e.*, using $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$, applied pointwise on each of the input samples). Then $\mathcal{A}^{\mathsf{S}}$ invokes S on the outputs obtaining an LWR secret. Finally, $\mathcal{A}^{\mathsf{S}}$ outputs an LWE secret computed using the original LWE samples and the LWR secret.

Definition 5 (Point-Wise Reduction from LWE to LWR). *Let $n, p, q \in \mathbb{N}$ be integers such that $p < q$, let χ be a distribution on $\mathbb{Z}_q$, and let $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{R}_{>0}$ and $\ell_{\mathsf{samp}} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{N}$ be functions. We say the PPT oracle algorithm $\mathcal{A}$ is an $(\ell_{\mathsf{err}}, \ell_{\mathsf{samp}})-$*pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ *if it is a reduction per Definition 3 and, moreover, if there exists an efficiently computable function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ and a PPT algorithm $\mathcal{B}$ such that for any $(\varepsilon', m')-$solver S for $\mathsf{LWR}_{n,q,p}$, the $(\varepsilon, m)-$solver $\mathcal{A}^{\mathsf{S}}$ for $\mathsf{LWE}_{n,q,\chi}$ works as follows where $(\varepsilon, m) = \left(\ell_{\mathsf{err}}(\varepsilon', m'), \ell_{\mathsf{samp}}(\varepsilon', m')\right)$.*

1. *Given* $\{(\mathbf{a}_i, b_i)\}_{i=1}^m \subset \mathbb{Z}_q^n \times \mathbb{Z}_q$*, compute* $(\mathbf{a}'_i, b'_i) = f(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$ *for* $i = 1, \ldots, m$*.*

2. *Call* $\mathsf{S}\big(\{(\mathbf{a}'_i, b'_i)\}\big)$ *obtaining* $\mathbf{s}' \in \mathbb{Z}_q^n \cup \{\perp\}$ *(*S *reads only the first* m' *pairs; if fewer than* m' *pairs are given,* S *outputs* $\perp$*).*

3. *Compute* $\mathcal{B}\big(\{(\mathbf{a}_i, b_i)\}, \mathbf{s}'\big)$ *obtaining* $\mathbf{s} \in \mathbb{Z}_q^n \cup \{\perp\}$*; output* $\mathbf{s}$*.*

Note that in a pointwise reduction, $m = \ell_{\mathsf{samp}}(\varepsilon', m') = m'$, since each LWE sample is mapped to an LWR sample which is then used by the LWR solver. For this reason, we usually ignore ℓ_{samp} when dealing with pointwise reductions.

Theorem 2 (Main). *Let* $n, p, q \in \mathbb{N}$ *be integers with* $q = \mathsf{poly}(n)$ *prime and* $q^{2/3+c} < p < q = \mathsf{poly}(n)$ *for a universal constant* $c > 0$*, and let* χ *be a distribution on* $\mathbb{Z}_q$*. Let* $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{R}_{>0}$ *be a function so* $\ell_{\mathsf{err}}(\varepsilon', m') = \mathsf{poly}\big(1/n, 1/m', \varepsilon'\big)$*. Then any* $\ell_{\mathsf{err}}-$*pointwise reduction* $\mathcal{A} = (f, \mathcal{B})$ *from* $\mathsf{LWE}_{n,q,\chi}$ *to* $\mathsf{LWR}_{n,q,p}$ *can be used to build an efficient* $(\varepsilon, m)-$*distinguisher for* $\mathsf{LWE}_{n,q,\chi}$ *for some non-negligible* $\varepsilon > 0$ *and some* $m = \mathsf{poly}(n)$*.*

If the error distribution χ on $\mathbb{Z}_q$ is such that $\mathsf{LWE}_{n,q,m,\chi}$ is hard for all $m = \mathsf{poly}(n)$ (*e.g.*, if χ is a discrete Gaussian), then Theorem 2 says that it is impossible to reduce $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ in a pointwise fashion.

3.2 The LWR Secret Recovery Algorithm and Proof of Theorem 2

Notation. Let $n, p, q \in \mathbb{N}$ be integers such that q is prime such that $q^{2/3+c} < p < q$ for a small constant $c > 0$. Let $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ be part of a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Since n, p, q, χ are fixed throughout the remainder of the paper, we write $\mathsf{LWE}_{\mathbf{s}}$ and $\mathsf{LWR}_{\mathbf{s}'}$ instead of $\mathsf{LWE}_{n,q,\mathbf{s},\chi}$ and $\mathsf{LWR}_{n,q,\mathbf{s}',p}$, respectively. The lemmas in this section make reference to non-negligible quantities $\eta, \delta > 0$ which will be specified in the next section.

Lemma 1 (Main Technical Lemma). *Let notations be as above. There exists an efficient algorithm* $\mathcal{A}$ *with the following syntax and correctness guarantees.*

- **Syntax:** $\mathcal{A}$ *takes no input, gets oracle access to a* $\big(\mathbb{Z}_q^n \times \mathbb{Z}_q\big)-$*oracle and to* f*, and outputs a vector* $\mathbf{s}' \in \mathbb{Z}_q^n$*.*

- **Correctness:** *If* $\mathcal{A}$ *is run when given oracle access to* $\mathsf{LWE}_{\mathbf{s}}$ *for a random* $\mathbf{s} \sim \mathbb{Z}_q^n$*, then with non-negligible probability (over* $\mathbf{s} \sim \mathbb{Z}_q^n$ *and the random coins of* $\mathcal{A}$*),* $\mathcal{A}$ *outputs* $\mathbf{s}' \in \mathbb{Z}_q^n$ *such that:*

$$\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\Big[b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \big\rceil_p\Big] \geqslant 1 - \eta, \tag{3}$$

  where $(\mathbf{a}', b') = f(\mathbf{a}, b)$*.*

Lemma 2. *Assume* $(f, \mathcal{B})$ *is a pointwise reduction from* $\mathsf{LWE}_{n,q,\chi}$ *to* $\mathsf{LWR}_{n,q,p}$*. If there exists* $\mathbf{s}' \in \mathbb{Z}_q^n$ *such that*

$$\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\Big[b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \big\rceil_p\Big] \geqslant 1 - 2\eta,$$

where $(\mathbf{a}', b') = f(\mathbf{a}, b)$*, then* $\mathcal{B}$ *is a* $(\delta, m)-$*solver for* $\mathsf{LWE}_{n,q,\chi}$ *for* $m = n(1 + \log q)/\eta$*.*

Proof of Theorem 2 Assuming Lemmas 1 and 2. Let $\mathcal{A}$ denote the algorithm promised by Lemma 1. Consider the following distinguishing algorithm $\mathcal{D}$, which gets oracle access to a $\big(\mathbb{Z}_q^n \times \mathbb{Z}_q\big)-$oracle $\mathcal{O}$ and works as follows.

1. D instantiates $\mathcal{A}$ with oracle access to $\mathcal{O}$, obtaining output $\mathbf{s}' \in \mathbb{Z}_q^n$. If $\mathcal{A}$ fails to give output of the proper type, D outputs 0.

2. Now D draws samples $(\mathbf{a}_1, b_1), \ldots, (\mathbf{a}_N, b_N) \sim \mathcal{O}$ for $N = n/\eta$, and computes an approximation $\hat{\mathsf{P}}$ of the probability
$$\mathsf{P} := \Pr_{(\mathbf{a},b)\sim\mathcal{O}}\Big[b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \big\rceil_p\Big],$$
where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. If $\hat{\mathsf{P}} \geqslant 1 - 3\eta/4$, D outputs 1, otherwise D outputs 0.

We show that D outputs 0 with probability $1 - 2^{-\Omega(n)}$ when $\mathcal{O}$ is a random oracle, and outputs 1 with non-negligible probability when $\mathcal{O}$ is an LWE oracle. The theorem follows.

Uniform Samples. Consider the execution of D when $\mathcal{O}$ is a random oracle, and let $\mathbf{s}' \in \mathbb{Z}_q^n$ be the vector obtained by $\mathcal{A}$ in Step 1 (if $\mathcal{A}$ outputs $\perp$ during this step then D outputs a random bit). In this case, the Chernoff-Hoeffding inequality ensures that $|\hat{\mathsf{P}} - \mathsf{P}| < \eta/2$ holds with probability $1 - 2^{-\Omega(n)}$. Thus by Lemma 2, $\hat{\mathsf{P}} < 1 - 3\eta/2$ occurs with probability $1 - 2^{-\Omega(n)}$, and so D outputs a random bit with high probability.

LWE Samples. Now consider the execution of D when instantiated with a $\mathsf{LWE}_{\mathbf{s}}-$oracle for a random $\mathbf{s} \sim \mathbb{Z}_q^n$. In this case, Lemma 1 ensures that with non-negligible probability, $\mathcal{A}$ outputs $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\mathsf{P} \geqslant 1 - \eta$. In this case, $\hat{\mathsf{P}}$ is again accurate to within $\pm\eta/2$ by the Chernoff bound, and so $\hat{\mathsf{P}} \geqslant 1 - 3\eta/2$ and D outputs 1 with non-negligible probability. □

4 The Statistics of a Pointwise Reduction

In this section we begin to impose structure on $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ which is part of a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. The fundamental intuition of this section is the following "meta" statement: *all statistics of the LWR distribution and the output distribution of f (given LWE samples as input) must be the same*. The reason for this is that any statistic which differs can be used to build a "pathological solver" which solves LWR but which will be useless for solving LWE via f. The solver simply draws enough samples to approximate the statistic, aborting if it decides it is being fed with mapped LWE samples, solving if it decides it is being fed with true LWR samples.

4.1 Non-Degeneracy

We prove that the distribution which draws $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$, computes $(\mathbf{a}', b') = f(\mathbf{a}, b)$ and outputs $\mathbf{a}' \in \mathbb{Z}_q^n$ cannot give non-negligible weight to any set $T \subset \mathbb{Z}_q^n$ with negligible density.

Definition 6. *Let $\zeta, \rho > 0$ be such that $\zeta > \rho$, and let $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ be a function. We say f is $(\zeta, \rho)-$degenerate if there exists $T \subset \mathbb{Z}_q^n$ of density $|T|/q^n = \rho$ such that $\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[\mathbf{a}' \in T\big] \geqslant \zeta$, where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. We say that f is $(\zeta, \rho)-$non-degenerate if it is not $(\zeta, \rho)-$degenerate.*

Claim 1 (Non-Degeneracy). *Let $n, q, p \in \mathbb{N}$ such that $p < q$ and χ be a distribution on $\mathbb{Z}_q$. Suppose $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is part of a pointwise reduction $(f, \mathcal{B})$ from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Suppose f is $(\rho + \varepsilon, \rho)-$degenerate for $\rho, \varepsilon > 0$ with ε non-negligible. Then $\mathcal{B}$ is an $(\varepsilon, m)-$solver of $\mathsf{LWE}_{n,q,\chi}$, for $m = \max\big\{qn(1 + \log q), n/(\rho\varepsilon^2)\big\}$.*

Proof. Let $\varepsilon > 0$ be non-negligible and suppose $(f, \mathcal{B})$ is a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ which is $(\rho+\varepsilon, \rho)-$degenerate. Let $\mathcal{D}$ be the distribution on $\mathbb{Z}_q^n$ which draws $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$, computes $(\mathbf{a}', b') = f(\mathbf{a}, b)$, and outputs $\mathbf{a}'$. By definition, there exists $T \subset \mathbb{Z}_q^n$ of density ρ such that $\Pr_{\mathcal{D}}\big[\mathbf{a}' \in T\big] \geqslant \rho + \varepsilon$. Let S be the pathological $(1 - 2^{-\Omega(n)}, m)-$solver for $\mathsf{LWR}_{n,q,p}$ which, on input $\{(\mathbf{a}_i', b_i')\}_{i=1}^m \subset \mathbb{Z}_q^n \times \mathbb{Z}_p$, computes $t := \#\{i : \mathbf{a}_i' \in T\}$ and outputs $\perp$ if $t \geqslant \big(\rho + \varepsilon/2\big)m$; otherwise if $t < \big(\rho + \varepsilon/2\big)m$, S outputs the unique $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $b_i' = \big\lfloor \langle \mathbf{a}_i', \mathbf{s}' \rangle \big\rceil_p$ for all $i = 1, \ldots, m$ (if no such $\mathbf{s}'$ exists or if more than one such $\mathbf{s}'$ exists, S outputs $\perp$). Note that when S is fed with LWR samples $t = \rho m$ in expectation as the $\mathbf{a}_i' \sim \mathbb{Z}_q^n$ are uniform. By the Chernoff-Hoeffding inequality, $t < \big(\rho + \varepsilon/2\big)m$ holds with probability $1 - 2^{-\Omega(n)}$ (since $m \geqslant n/(\rho\varepsilon^2)$). As $m \geqslant nq(1 + \log q)$, with probability at least $1 - 2^{-\Omega(n)}$, there exists exactly one $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $b_i' = \big\lfloor \langle \mathbf{a}_i', \mathbf{s}' \rangle \big\rceil_p$ for all $i = 1, \ldots, m$. Therefore, when S is fed with LWR samples it outputs the LWR secret $\mathbf{s}'$ with high probability.

On the other hand, when m LWE samples are chosen and S is fed with $\big\{f(\mathbf{a}_i, b_i)\big\}$, $t \geqslant (\rho + \varepsilon)m$ in expectation, and so by the Chernoff-Hoeffding inequality, $t \geqslant \big(\rho + \varepsilon/2\big)m$ holds with probability $1 - 2^{-\Omega(n)}$ (since $m \geqslant n/(\rho\varepsilon^2) \geqslant n/(\rho + \varepsilon)$). Therefore, S outputs $\perp$ with high probability when fed with mapped LWE samples. As $(f, \mathcal{B})$ is a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$, $\mathcal{B}$ outputs the LWE secret with non-negligible probability when fed with $\big(\{(\mathbf{a}_i, b_i)\}, \perp\big)$, where the $(\mathbf{a}_i, b_i)$ are LWE samples and the $\perp$ is the output of S on their images under f. Thus $\mathcal{B}$ solves $\mathsf{LWE}_{n,q,m,\chi}$ with non-negligible probability. □

4.2 Good LWE Secrets

We now identify a non-negligible subset $\mathsf{G} \subset \mathbb{Z}_q^n$ of *good* LWE secrets, where $\mathbf{s} \in \mathsf{G}$ guarantees some good behavior from f when fed with samples from $\mathsf{LWE}_{n,q,\mathbf{s},\chi}$.

The Secret Graph. The secret graph is a weighted complete bipartite graph whose left and right vertex sets (X and Y, respectively) are both $\mathbb{Z}_q^n$, and where the weight of the edge $(\mathbf{s}, \mathbf{s}') \in X \times Y$ is $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} := \Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[b' = \lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \rceil_p\big]$, where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. For $\mathbf{s} \in X$ and $\varepsilon > 0$, we write $Y_\varepsilon(\mathbf{s}) = \{\mathbf{s}' \in Y : \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1-\varepsilon\}$. Likewise, given $\mathbf{s}' \in Y$ and $\varepsilon > 0$, $X_\varepsilon(\mathbf{s}') = \{\mathbf{s} \in X : \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1-\varepsilon\}$. So intuitively, $Y_\varepsilon(\mathbf{s})$ is the subset of $\mathbf{s}$'s neighborhood which is connected to $\mathbf{s}$ by an edge with weight at least $1 - \varepsilon$; and similarly for $X_\varepsilon(\mathbf{s}')$.

Parameters. In addition to the parameters mentioned above, the good secrets are defined in terms of three non-negligible values $\delta, \eta, \sigma > 0$. The quantity δ is defined using the error loss function ℓ_{err} of the pointwise reduction $(f, \mathcal{B})$. Specifically, $2\delta = \ell_{\mathsf{err}}(1/3, m)$ where $m = 2n(1 + \log q)/\eta$, so that if S is a $\big(\frac{1}{3}, m\big)-$solver for $\mathsf{LWR}_{n,q,p}$, $\mathcal{B}^{\mathsf{S}}$ is a $2\delta-$solver for $\mathsf{LWE}_{n,q,\chi}$. Given δ, we set $\sigma = \delta/2nq(1 + \log q)$ and $\eta = \min\big\{\sigma, \big(1/3nq\big)^3\big\}$. The reader is encouraged on a first pass to just think of δ, η, σ all as arbitrarily small, but non-negligible, quantities.

Definition 7 (Good LWE Secrets). *With the above notation and conventions, we say that* $\mathbf{s} \in \mathbb{Z}_q^n$ *is* good*, and write* $\mathbf{s} \in \mathsf{G}$*, if the following three conditions hold:*

$$(1)\ |Y_\eta(\mathbf{s})| = 1; \quad (2)\ |Y_\sigma(\mathbf{s})| \leqslant 1; \quad (3)\ |X_\eta(\mathbf{s}')| = 1.$$

In point (3), $\mathbf{s}' \in \mathbb{Z}_q^n$ *is the LWR secret for which* $Y_\eta(\mathbf{s}) = \{\mathbf{s}'\}$.

Note that Points (1) and (3) together establish that the edges in the secret graph with weight above $1-\eta$ induce a matching between good LWE secrets and (a subset of) LWR secrets.

Claim 2. *Suppose $(f, \mathcal{B})$ is a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Then either $|\mathsf{G}| \geqslant \delta \cdot q^n$, or $\mathcal{B}$ is a $(\delta, m)-$solver for $\mathsf{LWE}_{n,q,\chi}$ for $m = 2n(1+\log q)/\eta$.*

Proof. Let $m = n(1+\log q)/\eta$, and let S be the pathological solver for $\mathsf{LWR}_{n,q,p}$ which, on input $\{(\mathbf{a}'_i, b'_i)\}_{i=1}^m$, does the following:

(i) it looks at the first $nq(1+\log q)$ samples (this is less than m since $\eta \leqslant 1/q$) and checks whether there exist distinct $\mathbf{s}', \mathbf{s}'' \in \mathbb{Z}_q^n$ such that $\lfloor \langle \mathbf{a}'_i, \mathbf{s}' \rangle \rceil_p = b'_i = \lfloor \langle \mathbf{a}'_i, \mathbf{s}'' \rangle \rceil_p$ holds for all $i = 1, \ldots, nq(1+\log q)$; if so, S outputs $\perp$;

(ii) S computes the unique $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $b'_i = \lfloor \langle \mathbf{a}'_i, \mathbf{s}' \rangle \rceil_p$ holds for all $i = 1, \ldots, m$, if no such $\mathbf{s}'$ exists, S outputs $\perp$;

(iii) using the $\mathbf{s}' \in \mathbb{Z}_q^n$ just computed, S checks if $^{\#}\{\mathbf{s} \in \mathbb{Z}_q^n : |Y_\eta(\mathbf{s})| = 1 \;\&\; \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1-\eta\} \geqslant 2$; if so S outputs $\perp$;

(iv) if it has not already aborted, S outputs $\mathbf{s}' \in \mathbb{Z}_q^n$ recovered in Step (ii).

Assume $|\mathsf{G}| < \delta \cdot q^n$. We will prove the following two points.

1. if S is called on $\{(\mathbf{a}'_i, b'_i)\} \sim \mathsf{LWR}_{n,q,m,p}$, then S outputs the secret $\mathbf{s}'$ with probability at least $1/3$;
2. if S is called on $\{(\mathbf{a}'_i, b'_i)\}$ for $\{(\mathbf{a}_i, b_i)\} \sim \mathsf{LWE}_{n,q,m,\chi}$ and $(\mathbf{a}'_i, b'_i) = f(\mathbf{a}_i, b_i)$, then S outputs $\perp$ with probability at least $1-\delta$.

Just as in Claim 1, these two points suffice. Point 1 says that S is a $\left(\frac{1}{3}, m\right)-$solver for $\mathsf{LWR}_{n,q,m,p}$. As $(f, \mathcal{B})$ is a pointwise reduction, with probability at least $2\delta = \ell_{\mathsf{err}}(1/3)$ over $\{(\mathbf{a}_i, b_i)\} \sim \mathsf{LWE}_{n,q,m,\chi}$, $\mathcal{B}$ outputs the LWE secret given $\{(\mathbf{a}_i, b_i)\}$ and $\mathsf{S}\big(\{(\mathbf{a}'_i, b'_i)\}\big)$. By point 2, the probability that $\mathcal{B}$ recovers the LWE secret without the second argument is at least δ. It remains to establish the two points.

Point 1 $-$ S on LWR samples: If S is fed with LWR instances, then certainly there exists $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $b'_i = \lfloor \langle \mathbf{a}'_i, \mathbf{s}' \rangle \rceil_p$ for all i (namely, the LWR secret). So S will solve LWR in step (ii) and give correct output as long as it does not abort in steps (i) or (iii). Just as in the proof of Claim 1, the probability that S outputs $\perp$ in Step (i) because it finds distinct $\mathbf{s}' \neq \mathbf{s}''$ such that $\lfloor \langle \mathbf{a}'_i, \mathbf{s}' \rangle \rceil_p = b'_i = \lfloor \langle \mathbf{a}'_i, \mathbf{s}'' \rangle \rceil_p$ for $i = 1, \ldots, m$ is $2^{-\Omega(n)}$. Moreover, note that

$$^{\#}\{\mathbf{s} \in \mathbb{Z}_q^n : |Y_\eta(\mathbf{s})| = 1 \;\&\; \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1-\eta\} \geqslant 2$$

holds for at most half of the $\mathbf{s}' \in \mathbb{Z}_q^n$. Therefore S aborts given LWR samples with probability at most $1/2 + 2^{-\Omega(n)} \leqslant 2/3$, and otherwise solves LWR.

Point 2 $-$ S on mapped LWE samples: If S is fed with mapped LWE instances, then some $\mathbf{s} \sim \mathbb{Z}_q^n$ is chosen, $\{(\mathbf{a}_i, b_i)\}_{i=1}^m \sim \mathsf{LWE}_{n,q,\mathbf{s},\chi}$ are drawn, and $(\mathbf{a}'_i, b'_i) = f(\mathbf{a}_i, b_i)$ are computed and fed to S. With probability at least $1-\delta$, $\mathbf{s} \notin \mathsf{G}$ in which case one of the properties (1), (2) and (3) does not hold. If (1) does not hold, then $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} < 1-\eta$ for all $\mathbf{s}' \in \mathbb{Z}_q^n$ and so

$$\Pr_{\{(\mathbf{a}_i,b_i)\}_{i=1}^m \sim \mathsf{LWE}_{n,q,\mathbf{s},\chi}}\Big[\exists\, \mathbf{s}' \in \mathbb{Z}_q^n \text{ st } b'_i = \big\lfloor \langle \mathbf{a}'_i, \mathbf{s}' \rangle \big\rceil_p \;\forall\; i = 1,\ldots,m\Big] < q^n \cdot \big(1-\eta\big)^m \leqslant 2^{-n},$$

(since $m = n(1+\log q)/\eta$) and so S outputs $\perp$ in Step (ii) with high probability $1-2^{-n}$. On the other hand, if (2) does not hold then there exist distinct $\mathbf{s}', \mathbf{s}'' \in \mathbb{Z}_q^n$ such that $\mathsf{p}_{(\mathbf{s},\mathbf{s}')}, \mathsf{p}_{(\mathbf{s},\mathbf{s}'')} \geqslant 1-\sigma$ both hold. In this case,

$$\Pr_{\{(\mathbf{a}_i,b_i)\}_{i=1}^m \sim \mathsf{LWE}_{n,q,\mathbf{s},\chi}}\Big[\big\lfloor \langle \mathbf{a}'_i, \mathbf{s}' \rangle \big\rceil_p = b'_i = \big\lfloor \langle \mathbf{a}'_i, \mathbf{s}'' \rangle \big\rceil_p \;\forall\; i\Big] \geqslant 1 - 2nq(1+\log q)\sigma \geqslant 1-\delta,$$

(using $\sigma \leqslant \delta/2nq(1+\log q)$) and so S outputs $\perp$ in Step (i) with probability $1-\delta$. Finally, suppose that (1) and (2) both hold and that S does not abort in Steps (i) or (ii) but that (3) does not hold. Note that $|X_\eta(\mathbf{s}')| \geqslant 1$ since $\mathbf{s} \in X_\eta(\mathbf{s}')$, thus if (3) does not hold then it must be that $|X_\eta(\mathbf{s}')| \geqslant 2$. In this case S simply outputs $\perp$ in Step (iii). So we have shown that when $\mathbf{s} \notin \mathsf{G}$, S outputs $\perp$ with probability at least $1-\delta$, as desired. □

4.3 Proof of Lemma 2

Claim 2 imposes quite a lot of structure on a pointwise reduction. We will refer to Claim 2 repeatedly throughout the remainder of the paper. Additionally, we can already derive Lemma 2 as a corollary.

Lemma 2 (Restated). *Assume $(f, \mathcal{B})$ is a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. If there exists $\mathbf{s}' \in \mathbb{Z}_q^n$ such that*

$$\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\Big[b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \big\rceil_p\Big] \geqslant 1-2\eta,$$

where $(\mathbf{a}', b') = f(\mathbf{a}, b)$, then $\mathcal{B}$ is a $(\delta, m)-$solver for $\mathsf{LWE}_{n,q,\chi}$ for $m = n(1+\log q)/\eta$.

Proof. Suppose there exists $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[b' = \lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \rceil_p\big] \geqslant 1-2\eta$, where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. Then by Corollary 1, $\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_\mathbf{s}}\big[b' = \lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \rceil_p\big] \geqslant 1-2\eta-q^{-n/4} \geqslant 1-3\eta$ holds for all but a $q^{-\Omega(n)}-$fraction of $\mathbf{s} \in \mathbb{Z}_q^n$. In other words, $|X_\eta(\mathbf{s}')| \geqslant (1-q^{-\Omega(n)})\cdot q^n$, so the degree of $\mathbf{s}'$ is way too high to have any neighbors in G. However, this means that $\mathsf{G} \subset \mathbb{Z}_q^n \backslash X_\eta(\mathbf{s}')$, and so $|\mathsf{G}| \leqslant q^{-\Omega(n)} \cdot q^n$ and so by Claim 2, $\mathcal{B}$ is a $(\delta, m)-$solver for $\mathsf{LWE}_{n,q,\chi}$. □

4.4 Outline of the Proof of Lemma 1

At this point we have reduced our main result (Theorem 2) to proving Lemma 1; namely we must design an algorithm which, given oracle access to $\mathsf{LWE}_\mathbf{s}$ for some uniform secret $\mathbf{s} \sim \mathbb{Z}_q^n$, reconstructs the LWR secret $\mathbf{s}' \in \mathbb{Z}_q^n$ of the mapped LWE pairs. We have also already proved a key claim, Claim 2, which specifies a notion of "good" behavior from an LWE secret $\mathbf{s}$ and proves that the set of good secrets $\mathsf{G} \subset \mathbb{Z}_q^n$ comprises a non-negligible fraction of the entire space. Intuitively, $\mathbf{s} \in \mathsf{G}$ if there exists a unique $\mathbf{s}' \in \mathbb{Z}_q^n$ such that

$$\mathsf{p}_{(\mathbf{s},\mathbf{s}')} := \Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_\mathbf{s}}\Big[b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \big\rceil_p\Big] \geqslant 1-\eta,$$

and, moreover, if this $\mathbf{s}'$ is unique to $\mathbf{s}$ (*i.e.*, so $\mathsf{p}_{(\mathbf{s}^*,\mathbf{s}')} < 1-\eta$ for all $\mathbf{s}^* \neq \mathbf{s}$). The algorithm of Lemma 1 will aim to recover $\mathbf{s}'$ whenever $\mathbf{s} \in \mathsf{G}$.

The bulk of the technical work of the remainder of the paper will go into proving the following lemma. Recall the notation of Lemma 1: $n, p, q \in \mathbb{N}$ are integers such that q is prime and $q^{2/3+c} < p < q$; $\nu = \nu(n) > 0$ is non-negligible and $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is part of a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Recall also that we inherited the non-negligible parameters $\delta, \eta, \sigma > 0$ from Claim 2.

Lemma 3. *Assume the above setup. There exists an efficient algorithm $\mathcal{A}_{\mathsf{AffRec}}$ which takes no input, gets oracle access to f, and outputs a pair $(\mathbf{H}, \mathbf{V})$ where $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ and $\mathbf{V} \subset \mathbb{Z}_q^n$ is a constant dimensional vector space such that with non-negligible probability (over the random coins of $\mathcal{A}_{\mathsf{AffRec}}$) the following holds:*

$$\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\Big[\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}\Big] \geqslant 1-\tau,$$

where $(\mathbf{a}', b') = f(\mathbf{a}, b)$ and $\tau = nq^2\eta^{1/12t}\sqrt{178n}$, and $t \in \mathbb{N}$ minimal such that $t \geqslant \frac{\log_q(1/\delta)+2}{3c}$ holds.

As mentioned in Section 1.2, once we know that $\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}$ holds with high probability, we can recover $\mathbf{s}'$ using a Goldreich-Levin-type argument. This part of our proof is in Section 5. Also as mentioned in Section 1.2, the proof of Lemma 3 consists of two separate pieces. First, we show that f passes a certain "property test" with high probability, then we show that any function which passes the test must have good agreement with an affine function. See Section 6 for a detailed overview and the formal proofs.

5 Recovering the LWR Secret via Goldreich-Levin Inversion

In this section we show how to use the Goldreich-Levin (GL) inversion technique [GL89] to recover the LWR secret. We begin by recalling the parameters and notations which we will use in this section.

Notations. We have integers $n, p, q \in \mathbb{N}$ such that q is prime and $q^{2/3+c} < p < q$ for some small constant $c > 0$. Additionally, $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is part of a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. We have non-negligible parameters $\delta, \eta, \sigma > 0$ from Claim 2, and a set of "good" LWE secrets $\mathsf{G} \subset \mathbb{Z}_q^n$ from Section 4.2. Additionally, we have an additional non-negligible $\tau > 0$ and $(\mathbf{H}, \mathbf{V})$ where $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ and $\mathbf{V} \subset \mathbb{Z}_q^n$ is a constant dimensional subspace such that

$$\mathsf{P}(\mathbf{H}, \mathbf{V}) := \Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}\big] \geqslant 1-\tau,$$

where $(\mathbf{a}', b') = f(\mathbf{a}, b)$. For $\mathbf{s} \in \mathbb{Z}_q^n$ and $e \in \mathbb{Z}_q$, let $\mathsf{P}_{\mathbf{s},e}(\mathbf{H}, \mathbf{V}) := \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}\big]$, where $(\mathbf{a}', b') = f(\mathbf{a}, \langle\mathbf{a}, \mathbf{s}\rangle + e)$. It follows immediately from Corollary 1 that for at most a $q^{-\Omega(n)}$–fraction of $\mathbf{s} \in \mathbb{Z}_q^n$, there exists an $e \in \mathbb{Z}_q$ such that $\mathsf{P}_{\mathbf{s},e}(\mathbf{H}, \mathbf{V}) < 1 - 2\tau$. So let us remove all such $\mathbf{s}$ from G; G will still comprise a non-negligible fraction of $\mathbb{Z}_q^n$. At this point what we will need from $\mathbf{s} \in \mathsf{G}$ is that the following points both hold:

$$(1)\ \exists \text{ unique } \mathbf{s}' \in \mathbb{Z}_q^n \text{ st } \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1-\eta; \quad (2)\ \mathsf{P}_{\mathbf{s},e}(\mathbf{H}, \mathbf{V}) \geqslant 1 - 2\tau\ \forall\ e$$

5.1 A Goldreich-Levin Theorem for LWE Samples

In this section, we state and prove a Goldreich-Levin-type theorem which will allow us to recover $\mathbf{H}^{\mathsf{t}}\mathbf{s}'$ given oracle access to $\mathsf{LWE}_{\mathbf{s}}$ for unknown $\mathbf{s}$.

Lemma 4 (A Goldreich-Levin Theorem for LWE Samples). *Let $n, q \in \mathbb{N}$ be such that $q = \mathsf{poly}(n)$ is prime, $\zeta \in (0,1)$. For a function $\mathsf{Pred} : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q$, and quantities $(\mathbf{s}, e, \bar{\mathbf{s}}, \gamma) \in \mathbb{Z}_q^n \times \mathbb{Z}_q \times \mathbb{Z}_q^n \times \mathbb{Z}_q$, let*

$$\mathrm{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, \gamma) := \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[\mathsf{Pred}(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) = \langle \mathbf{a}, \bar{\mathbf{s}}\rangle + \gamma\big];\ \mathrm{P}_{\mathbf{s}}(\bar{\mathbf{s}}, \gamma) := \Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[\mathsf{Pred}(\mathbf{a}, b) = \langle \mathbf{a}, \bar{\mathbf{s}}\rangle + \gamma\big].$$

Then there exists a randomized algorithm Inv which takes $\{(\mathbf{a}_i, b_i)\}_{i=1}^m \in (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m$ as input, outputs $\bar{\mathbf{s}}^* \in \mathbb{Z}_q^n$, runs in time $\mathsf{poly}(n, q, 1/\zeta, \mathsf{T}_{\mathsf{Pred}})$ where $\mathsf{T}_{\mathsf{Pred}}$ is the running time of Pred, and has the following correctness guarantee.

- **Correctness:** *Suppose that $\mathbf{s}, \bar{\mathbf{s}} \in \mathbb{Z}_q^n$ are such that both of the following hold:*
  - *for all $e \in \mathbb{Z}_q$ such that $\Pr\big[\chi = e\big] \geqslant \frac{4\zeta}{5qn^2}$, and non-zero $\gamma \in \mathbb{Z}_q^*$, $\mathrm{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, 0) \geqslant \mathrm{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, \gamma) - \zeta$;*
  - *for all non-zero $\gamma \in \mathbb{Z}_q^*$, $\mathrm{P}_{\mathbf{s}}(\bar{\mathbf{s}}, 0) \geqslant \mathrm{P}_{\mathbf{s}}(\bar{\mathbf{s}}, \gamma) + 10\zeta$.*

  Then

$$\Pr_{\{(\mathbf{a}_i,b_i)\}_{i=1}^m\sim\mathsf{LWE}_{\mathbf{s},\chi}}\big[\mathsf{Inv}\big(\{(\mathbf{a}_i, b_i)\}\big) = \bar{\mathbf{s}}\big] \geqslant \frac{8\zeta^6}{9n^4q^6}.$$

Remark. *Intuitively, the requirement $\mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}}, 0) \geqslant \mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}}, \gamma) + 10\zeta$ means that the most likely output of the predictor on samples from $\mathsf{LWE}_{\mathbf{s}}$ is $\bar{\mathbf{s}}$. The additional requirement that $\mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, 0) \geqslant \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, \gamma) - \zeta$ means that the predictor performs pretty well regardless of the LWE error. Note that the most likely output of the "trivial" predictor $\mathsf{Pred}(\mathbf{a}, b) = b$ is $\langle \mathbf{a}, \mathbf{s}\rangle$ (assuming $e = 0$ is the most likely LWE error, which is standard). However, as soon as $e \neq 0$, the trivial predictor starts performing extremely badly, always outputting the wrong value. Clearly if $\mathbf{s}$ could be recovered from the trivial predictor then LWE would be efficiently solvable. Thus the requirement that the predictor perform well for all errors is a critical hypothesis for the above lemma.*

Proof. Let $m = n^2/4\zeta$ and $k = 1 + \lceil \log_q(3mq/\zeta^2)\rceil$; Inv works as follows.

Input: Inv gets input $\{(\mathbf{a}_i, b_i)\}_{i=1}^m \in (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m$ and uses an algorithm for Pred as a subroutine.

Output: Inv outputs $\bar{\mathbf{s}}^* \in \mathbb{Z}_q^n$.

1. Choose $\mathbf{x}_1, \ldots, \mathbf{x}_k \sim \mathbb{Z}_q^n$, $g_1, h_1, \ldots, g_k, h_k \sim \mathbb{Z}_q$. For all $\mathbf{u} = (u_1, \ldots, u_k) \in \mathbb{Z}_q^k$, let

$$\mathbf{x}_{\mathbf{u}} := \sum_{j=1}^k u_j \mathbf{x}_j \in \mathbb{Z}_q^n;\ g_{\mathbf{u}} := \sum_{j=1}^k u_j g_j \in \mathbb{Z}_q;\ \text{and } h_{\mathbf{u}} := \sum_{j=1}^k u_j h_j \in \mathbb{Z}_q.$$

2. For all $i = 1, \ldots, m$, do the following:
   - for each $\beta \in \mathbb{Z}_q$, compute $\hat{\mathsf{p}}_i(\beta) := \Pr_{\mathbf{u}\sim\mathbb{Z}_q^k\setminus\{\mathbf{0}\}}\Big[\mathsf{Pred}(\mathbf{a}_i + \mathbf{x}_{\mathbf{u}}, b_i + g_{\mathbf{u}}) - h_{\mathbf{u}} = \beta\Big]$;
   - if there exists $\beta \in \mathbb{Z}_q$ such that $\hat{\mathsf{p}}_i(\beta) \geqslant \hat{\mathsf{p}}_i(\beta') + 3\zeta$ for all $\beta' \neq \beta$, set $w_i = \beta$; otherwise set $w_i = \perp$.

3. Finally, let $W = \{i \in \{1,\ldots,m\} : w_i \neq \perp\}$, and let $\{i_1,\ldots,i_n\} \subset W$ be such that $\{\mathbf{a}_{i_1},\ldots,\mathbf{a}_{i_n}\}$ is linearly independent (if no such subset exists, output the failure symbol $\perp$). Let $(\mathbf{A},\mathbf{w}) \in \mathbb{Z}_q^{n\times n} \times \mathbb{Z}_q^n$ be such that the $t-$th row (resp., coordinate) of $\mathbf{A}$ (resp., $\mathbf{w}$) is $\mathbf{a}_{i_t}$ (resp., w_{i_t}). Output $\bar{\mathbf{s}}^* = \mathbf{A}^{-1}\mathbf{w} \in \mathbb{Z}_q^n$.

It is clear that Inv runs in time $\mathsf{poly}(n,q,1/\zeta,\mathsf{T}_{\mathsf{Pred}})$. Assume that $\mathbf{s},\bar{\mathbf{s}} \in \mathbb{Z}_q^n$ are such that both correctness hypotheses hold. We will show that Inv outputs $\bar{\mathbf{s}}^* = \bar{\mathbf{s}}$ with probability at least $1/2q^{2k}$. Consider first the random choices $(\mathbf{x}_j, g_j, h_j) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q \times \mathbb{Z}_q$ drawn during Step 1. Let us say that these random choices are *correct* if:

$$g_j = \langle \mathbf{x}_j, \mathbf{s}\rangle \text{ and } h_j = \langle \mathbf{x}_j, \bar{\mathbf{s}}\rangle \;\forall\; j = 1,\ldots,k.$$

Note these random choices are correct with probability q^{-2k}. When the random choices are correct, we have $g_{\mathbf{u}} = \langle \mathbf{x}_{\mathbf{u}}, \mathbf{s}\rangle$ and $h_{\mathbf{u}} = \langle \mathbf{x}_{\mathbf{u}}, \bar{\mathbf{s}}\rangle$ for all $\mathbf{u} \in \mathbb{Z}_q^k$. Consider now the values $\hat{\mathsf{p}}_i(\beta)$ for $\beta \in \mathbb{Z}_q$ and $i \in \{1,\ldots,m\}$ computed in Step 2, and let us interpret the $\hat{\mathsf{p}}_i(\beta)$ as random variables over $\mathbf{x}_j \sim \mathbb{Z}_q^n$. Note that if the choices are correct, then $(\mathbf{a}_i + \mathbf{x}_{\mathbf{u}}, b_i + g_{\mathbf{u}})$ is a random $\mathsf{LWE}_{\mathbf{s}}$ pair with the same error as $(\mathbf{a}_i, b_i)$; thus the expectation of $\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle + \gamma)$ is $\mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma)$ for all $\gamma \in \mathbb{Z}_q$ and $i \in \{1,\ldots,m\}$, where $e_i = b_i - \langle \mathbf{a}_i, \mathbf{s}\rangle$. We will prove a concentration bound using the pairwise independence of $(\mathbf{x}_{\mathbf{u}}, \mathbf{x}_{\mathbf{u}'})$ for $\mathbf{u} \neq \mathbf{u}' \in \mathbb{Z}_q^k$ which will guarantee that with probability at least $2/3$ (conditioned on correctness), $\left|\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle + \gamma) - \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma)\right| < \zeta$ holds for all $i = 1,\ldots,m$ and $\gamma \in \mathbb{Z}_q$. Let us first show how this completes the analysis of Inv.

Assume that the error term e_i is such that $\Pr\left[\chi = e_i\right] \geqslant \frac{1}{5qm}$; by the union bound the probability that this holds for all $i = 1,\ldots,m$ is at least $4/5$. The first observation is that for all $i \in \{1,\ldots,m\}$ and non-zero $\gamma \in \mathbb{Z}_q^*$, we have

$$\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle) > \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},0) - \zeta \geqslant \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma) - 2\zeta > \hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle + \gamma) - 3\zeta.$$

This means that Step 2 never sets w_i to be any value other than $\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle$. Likewise, we have the bound $\mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}},0) - \mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}},\gamma) \geqslant 10\zeta$ for non-zero $\gamma \in \mathbb{Z}_q^*$ means that $\mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},0) - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},\gamma) \geqslant 5\zeta$ holds with probability at least 5ζ over $e \sim \chi$. By Chernoff, the probability that $\mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},0) - \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma) \geqslant 5\zeta$ holds for at least $4\zeta m = n^2$ of the input LWE pairs $(\mathbf{a}_i, b_i)$ is $1 - 2^{-\Omega(n)}$. The probability that n^2 random vectors in $\mathbb{Z}_q^n$ span a proper subspace is at most $q^{-\Omega(n)}$; thus with probability at least $1 - 2^{-\Omega(n)}$, there exist n input samples $(\mathbf{a}_{i_1}, b_{i_1}), \ldots, (\mathbf{a}_{i_n}, b_{i_n})$ such that $\mathrm{Span}\left(\{\mathbf{a}_{i_1},\ldots,\mathbf{a}_{i_n}\}\right) = \mathbb{Z}_q^n$ and such that each error term e satisfies $\mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},0) - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},\gamma) \geqslant 5\zeta$ for all non-zero $\gamma \in \mathbb{Z}_q^*$. For each $i \in \{i_1,\ldots,i_n\}$,

$$\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle) > \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},0) - \zeta \geqslant \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma) + 4\zeta > \hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \mathbf{s}\rangle + \gamma) + 3\zeta,$$

and so Inv sets $w_i = \langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle$ during Step 2. So we have shown that, conditioned on the random choices in Step 1 being correct, Inv never sets w_i equal to anything but $\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle$ in Step 2, and furthermore, with probability at least $4/5 - 2^{-\Omega(n)} \geqslant 3/4$, Inv sets $w_i = \langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle$ for at least n values of $i \in \{1,\ldots,m\}$ such that the corresponding $\mathbf{a}_i$'s span $\mathbb{Z}_q^n$. Thus, once we show that $\left|\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle + \gamma) - \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma)\right| < \zeta$ holds simultaneously for all $i = 1,\ldots,m$ and $\gamma \in \mathbb{Z}_q$ with probability at least $2/3$, we will have shown that Inv recovers $\bar{\mathbf{s}}$ with probability at least $q^{-2k}/2$, as desired.

So fix an LWE sample $(\mathbf{a}, b)$ and $\gamma \in \mathbb{Z}_q$, and let $\mathbb{1}(\mathbf{u})$ be the $0/1$ random variable which outputs 1 if $\mathsf{Pred}(\mathbf{a}+\mathbf{x}_{\mathbf{u}}, b+g_{\mathbf{u}}) - h_{\mathbf{u}} = \langle \mathbf{a}, \bar{\mathbf{s}}\rangle + \gamma$ and 0 otherwise. Let $\mathsf{Q} := \Pr\left[|\hat{\mathsf{p}}(\langle \mathbf{a}, \bar{\mathbf{s}}\rangle + \gamma) - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},\gamma)| > \zeta\right]$

be shorthand. We have

$$\begin{aligned}\zeta^2\mathsf{Q} &\leqslant \mathbb{E}\Big[\hat{\mathsf{p}}(\langle \mathbf{a}, \bar{\mathbf{s}}\rangle + \gamma)^2\Big] - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, \gamma)^2 \\ &= \frac{1}{(q^k-1)^2} \cdot \sum_{\mathbf{u}\neq\mathbf{u}'\in\mathbb{Z}_q^k\setminus\{\mathbf{0}\}} \mathbb{E}\big[\mathbb{1}(\mathbf{u})\cdot\mathbb{1}(\mathbf{u}')\big] - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},\gamma)^2 + \frac{1}{(q^k-1)} \\ &\leqslant \frac{1}{(q^k-1)},\end{aligned}$$

and so $\mathsf{Q} \leqslant \frac{1}{\zeta^2(q^k-1)} \leqslant \frac{1}{3mq}$. So the concentration bound holds simultaneously for all $i \in \{1,\ldots,m\}$ and $\gamma \in \mathbb{Z}_q$ with probability at least $2/3$ by the union bound. The result follows. □

5.2 The Natural Predictor

Let notations be as specified at the beginning of this section. So, $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is part of a pointwise reduction, and $(\mathbf{H}, \mathbf{V})$ are such that $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ and $\mathbf{V} \subset \mathbb{Z}_q^n$ is a constant dimensional vector space such that $\mathsf{P}(\mathbf{H},\mathbf{V}) \geqslant 1-\tau$. Let $\{\mathbf{v}_1,\ldots,\mathbf{v}_d\}$ be a basis for $\mathbf{V}$. Given such setup, we now describe the "natural predictor", which given samples $(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}$ for sufficiently good $\mathbf{s} \in \mathsf{G}$, predicts the inner product $\langle \mathbf{a}, \mathbf{H}^{\mathrm{t}}\mathbf{s}'\rangle$ well enough so that it is possible to use Lemma 4 to recover $\mathbf{H}^{\mathrm{t}}\mathbf{s}'$.

The Natural Predictor. The predictor function $\mathsf{Pred} : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q$ works as follows.

- The natural predictor is parametrized by $\alpha_1,\ldots,\alpha_d \in \mathbb{Z}_q$.
- Given $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, Pred computes $(\mathbf{a}', b') = f(\mathbf{a}, b)$; if $\mathbf{a}' = \alpha\mathbf{H}\mathbf{a} + \mathbf{v}$ for $\alpha \in \mathbb{Z}_q^*$ and $\mathbf{v} = \sum_{i=1}^d c_i\mathbf{v}_i \in \mathbf{V}$, then output $\alpha^{-1}\big(x - \sum_{i=1}^d c_i\alpha_i\big)$ where $x \sim \mathbb{Z}_q$ is random such that $\lfloor x \rceil_p = b'$.
- If $\mathbf{a}' \notin \mathrm{Span}(\mathbf{H}\mathbf{a}) + \mathbf{V}$, output a random $x \sim \mathbb{Z}_q$.

Note that whenever $b' = \big\lfloor \langle \mathbf{a}', \mathbf{s}'\rangle \big\rceil_p$ and $\mathbf{a}' = \alpha\mathbf{H}^{\mathrm{t}}\mathbf{a} + \mathbf{v}$ both hold, $b' = \big\lfloor \alpha\langle \mathbf{a}, \mathbf{H}^{\mathrm{t}}\mathbf{s}'\rangle + \langle \mathbf{v}, \mathbf{s}'\rangle \big\rceil_p$ also holds; so when the natural predictor draws x, a random rounding preimage of b' and outputs $\alpha^{-1}\big(x - \sum_i c_i\alpha_i\big)$, it has probability roughly $p/q \gg 1/q$ of outputting $\langle \mathbf{a}, \mathbf{H}^{\mathrm{t}}\mathbf{s}'\rangle$ as long as $\alpha_i = \langle \mathbf{v}_i, \mathbf{s}'\rangle$ for all $i = 1,\ldots,d$. The following claim proves that this predictor satisfies the hypotheses of Lemma 4, and so can be used to recover $\mathbf{H}^{\mathrm{t}}\mathbf{s}'$.

Claim 3. *Let notations be as above. Suppose that the natural predictor is fed with inputs from an* $\mathsf{LWE}_{\mathbf{s}}$*−oracle for some unknown* $\mathbf{s} \in \mathsf{G}$ *such that for all* $\beta \in \mathbb{Z}_q$, $\Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \geqslant \frac{1}{q^2}$, *where* $\mathcal{D}_{\mathbf{s}}$ *is the distribution which draws* $(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}$ *such that* $\mathbf{a}' \in \mathrm{Span}(\mathbf{H}\mathbf{a}) + \mathbf{V}$, *and outputs* $\langle \mathbf{a}, \mathbf{H}^{\mathrm{t}}\mathbf{s}'\rangle$. *Assume furthermore that the parameters of the predictor are* $\alpha_i = \langle \mathbf{v}_i, \mathbf{s}'\rangle$ *for all* $i = 1,\ldots,d$. *Then both of the correctness hypotheses of Lemma 4 are satisfied for* $\bar{\mathbf{s}} = \mathbf{H}^{\mathrm{t}}\mathbf{s}'$.

Proof. Fix $\zeta = \frac{1-2\tau-q^2\eta}{11q^3}$. We must show two points:

- for all $e \in \mathbb{Z}_q$ with $\Pr\big[\chi = e\big] \geqslant \frac{4\zeta}{5qn^2}$ and all non-zero $\gamma \in \mathbb{Z}_q^*$, $\mathsf{P}_{\mathbf{s},e}(\mathbf{H}^{\mathrm{t}}\mathbf{s}', 0) \geqslant \mathsf{P}_{\mathbf{s},e}(\mathbf{H}^{\mathrm{t}}\mathbf{s}', \gamma) - \zeta$;
- for all non-zero $\gamma \in \mathbb{Z}_q^*$, $\mathsf{P}_{\mathbf{s}}(\mathbf{H}^{\mathrm{t}}\mathbf{s}', 0) - \mathsf{P}_{\mathbf{s}}(\mathbf{H}^{\mathrm{t}}\mathbf{s}', \gamma) \geqslant 10\zeta$;

where $\mathsf{P}_{\mathbf{s},e}(\mathbf{H}^{\mathsf{t}}\mathbf{s}',\gamma)$ and $\mathsf{P}_{\mathbf{s}}(\mathbf{H}^{\mathsf{t}}\mathbf{s}',\gamma)$ are the notations from Lemma 4:

$$\mathsf{P}_{\mathbf{s},e}(\mathbf{H}^{\mathsf{t}}\mathbf{s}',\gamma) := \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[\mathsf{Pred}(\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+e) = \langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\gamma\big],$$

and $\mathsf{P}_{\mathbf{s}}(\mathbf{H}^{\mathsf{t}}\mathbf{s}',\gamma)$ is the same except the probability is over $(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}$. Let us simplify the shorthand by writing $\mathsf{P}_e^{(1)}(\gamma)$ and $\mathsf{P}^{(1)}(\gamma)$ instead of $\mathsf{P}_{\mathbf{s},e}(\mathbf{H}^{\mathsf{t}}\mathbf{s}',\gamma)$ and $\mathsf{P}_{\mathbf{s}}(\mathbf{H}^{\mathsf{t}}\mathbf{s}',\gamma)$. Note

$$\mathsf{P}_e^{(1)}(\gamma) = \big(1-\mathsf{P}_{\mathbf{s},e}(\mathbf{H},\mathbf{V})\big)\cdot\frac{1}{q} + \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[\mathsf{Pred}(\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+e) = \langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\gamma \;\&\; \mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\big],$$

where $(\mathbf{a}',b') = f(\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+e)$. So if we shorthand the second term by $\mathsf{P}_e^{(2)}(\gamma)$, then we get that $\mathsf{P}_e^{(1)}(0)-\mathsf{P}_e^{(1)}(\gamma) = \mathsf{P}_e^{(2)}(0)-\mathsf{P}_e^{(2)}(\gamma)$. Now let

$$\mathsf{P}_e^{(3)}(\gamma) := \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[\mathsf{Pred}(\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+e) = \langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\gamma \;\&\; b' = \big\lfloor\langle\mathbf{a}',\mathbf{s}'\rangle\big\rceil_p \;\&\; \mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\big],$$

where $(\mathbf{a}',b') = f(\mathbf{a},\langle\mathbf{a},\mathbf{s}\rangle+e)$. If $\Pr\big[\chi=e\big]\geqslant\frac{4\zeta}{5qn^2}$, $\mathsf{P}_3^{(2)}-\frac{5qn^2\eta}{4\zeta}\leqslant\mathsf{P}_e^{(3)}(\gamma)\leqslant\mathsf{P}_e^{(2)}(\gamma)$, since $\mathbf{s}\in\mathsf{G}$ and so $\mathsf{p}_{(\mathbf{s},\mathbf{s}')}\geqslant 1-\eta$. Therefore, $\mathsf{P}_e^{(2)}(0)-\mathsf{P}_e^{(2)}(\gamma)\geqslant\mathsf{P}_e^{(3)}(0)-\mathsf{P}_e^{(3)}(\gamma)-\zeta$, using $\eta\leqslant\frac{4\zeta^2}{5qn^2}$. To bound the $\mathsf{P}^{(3)}$ terms, recall that when $\mathbf{a}'=\alpha\mathbf{H}\mathbf{a}+\mathbf{v}$ for $\mathbf{v}=\sum_i c_i\mathbf{v}_i\in\mathbf{V}$, Pred outputs $\alpha^{-1}\big(x-\sum_i c_i\alpha_i\big)$ for a random $x\sim\mathbb{Z}_q$ such that $\lfloor x\rceil_p = b'$. Therefore, when $b' = \big\lfloor\langle\mathbf{a}',\mathbf{s}'\rangle\big\rceil_p = \big\lfloor\alpha\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\langle\mathbf{v},\mathbf{s}'\rangle\big\rceil_p$, Pred outputs $\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle$ with probability roughly p/q when $\big\lfloor\alpha(\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\gamma)+\langle\mathbf{v},\mathbf{s}'\rangle\big\rceil_p = \big\lfloor\alpha\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\langle\mathbf{v},\mathbf{s}'\rangle\big\rceil_p$, and with probability 0 otherwise. It follows that $\mathsf{P}_e^{(3)}(0)-\mathsf{P}_e^{(3)}(\gamma)$ is roughly

$$\frac{p}{q}\cdot\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\Big[\big\lfloor\alpha(\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\gamma)+\langle\mathbf{v},\mathbf{s}'\rangle\big\rceil_p \neq \big\lfloor\alpha\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\langle\mathbf{v},\mathbf{s}'\rangle\big\rceil_p \;\&\; \mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\Big]\geqslant 0.$$

Thus, $\mathsf{P}_e(0)\geqslant\mathsf{P}_e(\gamma)-\zeta$ for all non-zero $\gamma\in\mathbb{Z}_q^*$, which establishes the first point.

For the second point, we can define $\mathsf{P}^{(2)}(\gamma),\mathsf{P}^{(3)}(\gamma)$ analogously to how we defined $\mathsf{P}_e^{(2)}(\gamma),\mathsf{P}_e^{(3)}(\gamma)$, respectively (except probability is over $(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}$), and we get $\mathsf{P}^{(1)}(0)-\mathsf{P}^{(1)}(\gamma) = \mathsf{P}^{(2)}(0)-\mathsf{P}^{(2)}(\gamma)\geqslant\mathsf{P}^{(3)}(0)-\mathsf{P}^{(3)}(\gamma)-\eta\geqslant\mathsf{P}^{(3)}(0)-\mathsf{P}^{(3)}(\gamma)-\zeta$. Now, let us write $\mathsf{P}^{(3)}(\gamma)=\sum_{\beta\in\mathbb{Z}_q}S_\beta(\gamma)$ where each $S_\beta(\gamma)$ is the product of the following four terms:

- $\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[\mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\big] =: \mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V})$;
- $\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle=\beta\big|\mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\big]$;
- $\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[b'=\big\lfloor\langle\mathbf{a}',\mathbf{s}'\rangle\big\rceil_p\big|\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle=\beta\;\&\;\mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\big]$;
- $\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[\mathsf{Pred}(\mathbf{a},b)=\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle+\gamma\big|b'=\big\lfloor\langle\mathbf{a}',\mathbf{s}'\rangle\big\rceil_p\;\&\;\langle\mathbf{a},\mathbf{H}^{\mathsf{t}}\mathbf{s}'\rangle=\beta\;\&\;\mathbf{a}'\in\mathrm{Span}(\mathbf{H}\mathbf{a})+\mathbf{V}\big]$.

Let $\mathsf{Q}_\beta(\gamma)$ be shorthand for the fourth term; as noted above, $\mathsf{Q}_\beta(\gamma)$ is roughly equal to $\frac{p}{q}\cdot\mathbb{1}(\beta,\gamma)$ where $\mathbb{1}(\beta,\gamma)=1$ if $\big\lfloor\alpha(\beta+\gamma)+\sum_i c_i\alpha_i\big\rceil_p = \big\lfloor\alpha\beta+\sum_i c_i\alpha_i\big\rceil_p$, and is zero otherwise. The second term is $\Pr\big[\mathcal{D}_{\mathbf{s}}=\beta\big]$, where $\mathcal{D}_{\mathbf{s}}$ is the distribution defined in the claim statement. Finally, note that the third term is at least $1-\frac{q^2\eta}{\mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V})}$. Thus, for non-zero $\gamma\in\mathbb{Z}_q^*$,

$$\begin{aligned}\mathsf{P}^{(3)}(0)-\mathsf{P}^{(3)}(\gamma) &\geqslant \Big(\mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V})-q^2\eta\Big)\cdot\sum_{\beta\in\mathbb{Z}_q}\Pr\big[\mathcal{D}_{\mathbf{s}}=\beta\big]\cdot\big(\mathsf{Q}_\beta(0)-\mathsf{Q}_\beta(\gamma)\big)\\ &\geqslant \left(\frac{\mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V})}{q^2}-\eta\right)\cdot\sum_{\beta:\mathbb{1}(\beta,\gamma)=0}\frac{1}{q}\geqslant\left(\frac{1-2\tau-q^2\eta}{q^3}\right)=11\zeta,\end{aligned}$$

where the second inequality on the second line holds since when $\gamma\neq 0$ there exists at least one β such that $\mathbb{1}(\beta,\gamma)=0$. Hence we have that $\mathsf{P}^{(1)}(0)-\mathsf{P}^{(1)}(\gamma)\geqslant(1-(2\tau+q^2\eta))/q^3-\zeta\geqslant 10\zeta$, which completes the proof of the second point. □

5.3 Proving Lemma 1 Assuming Lemma 3

Lemma 1 (Restated). *Assume the notations described in the beginning of the section. So specifically, $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is part of a pointwise reduction and $(\mathbf{H}, \mathbf{V})$ are such that $\mathsf{P}(\mathbf{H}, \mathbf{V}) \geqslant 1 - \tau$. Then there exists an algorithm which, given oracle access to an $\mathsf{LWE}_{\mathbf{s}}$−oracle for a random $\mathbf{s} \sim \mathsf{G}$, outputs $\mathbf{H}^{\mathsf{t}}\mathbf{s}'$ with non-negligible probability over $\mathbf{s} \sim \mathsf{G}$ and the random coins.*

Proof. By Claim 3 and Lemma 4, it suffices simply to show that for an overwhelming fraction of the $\mathbf{s} \in \mathsf{G}$ have $\Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \geqslant \frac{1}{q^2}$ for all $\beta \in \mathbb{Z}_q$ where $\mathcal{D}_{\mathbf{s}}$ is the distribution which draws $(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}$ such that $\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}$ and outputs $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$. Since $\mathsf{P}_{\mathbf{s}}(\mathbf{H}, \mathbf{V}) \geqslant 1 - 2\tau$, $\mathcal{D}_{\mathbf{s}}$ is within statistical distance 2τ of the distribution $\hat{\mathcal{D}}_{\mathbf{s}}$ which simply draws $\mathbf{a} \sim \mathbb{Z}_q^n$ and outputs $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$. For $\beta \in \mathbb{Z}_q$, define the sets:

$$X_\beta := \big\{\mathbf{s} \in \mathsf{G} : \Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}[\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle = \beta] < q^{-2}\big\}; \text{ and } Y_\beta := \big\{\mathbf{H}^{\mathsf{t}}\mathbf{s}' : \mathbf{s} \in X_\beta\big\},$$

and consider the distribution $\mathcal{D}_\beta$, which draws $\mathbf{a} \sim \mathbb{Z}_q^n$, $\mathbf{s} \sim X_\beta$ and outputs $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$. We have

$$\frac{1}{q} - \frac{1}{q^2} - 2\tau < \Delta\big(\mathcal{D}_\beta, \mathsf{Unif}(\mathbb{Z}_q)\big) \leqslant q^c \Delta\big(\langle \mathsf{Unif}(\mathbb{Z}_q^n), \mathsf{Unif}(Y_\beta) \rangle, \mathsf{Unif}(\mathbb{Z}_q)\big) \leqslant \sqrt{\frac{q}{4|Y_\beta|}}.$$

The first inequality used the definition of X_β; the second used that $\mathbf{H}$ has rank $n - c$ for some constant c (since otherwise f would be degenerate), and that G induces a perfect matching between LWE secrets and LWR secrets; and the last inequality is Fact 1. It follows that $|Y_\beta| = q^{\mathcal{O}(1)}$, and thus so are $|X_\beta|$, and $\bigcup_\beta X_\beta$. Therefore, $\Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \geqslant \frac{1}{q^2}$ holds for all $\beta \in \mathbb{Z}_q$ for an overwhelming fraction of the $\mathbf{s} \in \mathsf{G}$. Lemma 1 follows. □

6 Proof of Lemma 3

Notation. Recall we have integers $n, p, q \in \mathbb{N}$ such that q is prime and $q^{2/3+c} < p < q$ for some small constant $c > 0$. Additionally, $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ is part of a pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$. Recall from Section 4.2, we have a set $\mathsf{G} \subset \mathbb{Z}_q^n$ of "good secrets"; this set has size at least $|\mathsf{G}| \geqslant \delta q^n$ for non-negligible $\delta > 0$ and for each $\mathbf{s} \in \mathsf{G}$ there exists a unique $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1 - \eta$ for non-negligible $\eta > 0$. It was also shown in Claim 1 that for all subset $S \subset \mathbb{Z}_q^n$ of size $|S| = \rho q^n$, and non-negligible $\nu > 0$, $\Pr_{(\mathbf{a},b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q}\big[\mathbf{a}' \in S\big] \leqslant \rho + \nu$. We have been calling this the "non-degenerate" property of f; this will play a major role in this section. Our goal in this section is to algorithmically recover $(\mathbf{H}, \mathbf{V})$ such that $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ and $\mathbf{V} \subset \mathbb{Z}_q^n$ is a constant dimensional vector subspace such that

$$\mathsf{P}(\mathbf{H}, \mathbf{V}) := \Pr_{(\mathbf{a},b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q}\big[\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}\big] \geqslant 1 - \tau,$$

where $(\mathbf{a}', b') = f(\mathbf{a}, b)$, and $\tau = nq^2\eta^{1/12t}\sqrt{178n}$, where $t \in \mathbb{N}$ is a new parameter; it is the minimal integer such that $t \geqslant \frac{\log_q(1/\delta)+2}{3c}$ holds. Note $t = \mathcal{O}(1)$.

The Function h. We introduce the function $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ which is defined from f as follows. First, if $b \in \mathbb{Z}_q$, then define the function $h_b : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ by $h_b(\mathbf{a}) = \mathbf{a}'$ such that $(\mathbf{a}', b') = f(\mathbf{a}, b)$. If $\mathbf{a} \sim \mathbb{Z}_q^n$, then $h(\mathbf{a})$ chooses $b \sim \mathbb{Z}_q$ and outputs $h_b(\mathbf{a}) \in \mathbb{Z}_q^n$. However, if $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n$, and $\alpha_1, \ldots, \alpha_n \in \mathbb{Z}_q$, then we define $h\Big(\sum_i \alpha_i \mathbf{a}_i\Big) = h_{\sum_i \alpha_i b_i}\Big(\sum_i \alpha_i \mathbf{a}_i\Big)$, where each $b_i \in \mathbb{Z}_q$ is the randomness chosen in the

computation of $h(\mathbf{a}_i)$. In this section, it will be considerably simpler to work with h rather than f. The non-degeneracy property framed in terms of h asserts that for all $S \subset \mathbb{Z}_q^n$ of size $|S| = \rho q^n$, and non-negligible $\nu > 0$, $\Pr_{\mathbf{a} \sim \mathbb{Z}_q}\big[h(\mathbf{a}) \in S\big] \leqslant \rho + \nu$.

6.1 Proof Overview

We first provide a brief high-level overview of the proof of Lemma 3. We prove Lemma 3 in two parts. First, we show there exists an efficiently computable subspace $\mathbf{V} \subset \mathbb{Z}_q^n$ of constant dimension such that

$$\Pr_{\substack{\mathbf{a}_1, \mathbf{a}_2 \sim \mathbb{Z}_q^n \\ (\alpha_1, \alpha_2) \sim \mathbb{Z}_q^2 \\ \text{s.t. } (\alpha_1, \alpha_2) \neq (0,0)}} \Big[h(\alpha_1 \mathbf{a}_1 + \alpha_2 \mathbf{a}_2) \in \mathrm{Span}\big(\{h(\mathbf{a}_1), h(\mathbf{a}_2)\}\big) + \mathbf{V}\Big] \geqslant 1 - 2\sqrt{\nu},$$

for a non-negligible quantity $\nu = \nu(n)$. Next, we prove an affine linearity testing theorem to show there exists an efficiently computable matrix $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ such that $\mathsf{P}\big(\mathbf{H}, \mathbf{V}\big) \geqslant 1 - \tau$, as desired.

6.1.1 Part I: h Passes an Affine Linearity Test with High Probability

Towards proving the first point, we consider the experiment which, for all $i \in [t]$:

1. Chooses $\mathbf{a}_{i,0}, \mathbf{a}_{i,1} \sim \mathbb{Z}_q^n$, $\alpha_{i,0}, \alpha_{i,1} \sim \mathbb{Z}_q$;
2. Sets $\mathbf{a}_{i,2} = \alpha_0 \mathbf{a}_0 + \alpha_1 \mathbf{a}_1 \in \mathbb{Z}_q^n$;
3. Computes $\mathbf{a}'_{i,j} = h(\mathbf{a}_{i,j})$, $\forall j \in \{0,1,2\}$;

and then outputs $\{\mathbf{a}'_{i,j}\}_{\substack{i \in [t], \\ j \in \{0,1,2\}}} \subset \mathbb{Z}_q^n$. Suppose that $d := \dim\Big(\mathrm{Span}\big(\{\mathbf{a}'_{i,j}\}_{\substack{i \in [t], \\ j \in \{0,1,2\}}}\big)\Big) = 3t$. For ease of presentation, suppose furthermore that $\dim\Big(\mathrm{Span}\big(\{\mathbf{a}_{i,j}\}_{\substack{i \in [t], \\ j \in \{0,1\}}}\big)\Big) = 2t$. Since $\{(\mathbf{a}_{i,j}, b_{i,j})\}_{\substack{i \in [t] \\ j \in \{0,1\}}}$ is statistically close to $2t$ $\mathsf{LWE}_{\mathbf{s}}$ samples, for some $\mathbf{s} \in \mathbb{Z}_q^n$, by the correctness of f it follows that f generates $3t$ $\mathsf{LWR}_{\mathbf{s}'}$ samples from the $2t$ $\mathsf{LWE}_{\mathbf{s}}$ samples, where $\mathbf{s}' \in \mathbb{Z}_q^n$ is the unique right neighbor of $\mathbf{s}$ in G. But, letting $e_{i,j} \in \mathbb{Z}_q$ be the LWE error term of each sample $(\mathbf{a}_{i,j}, b_{i,j})$ $(i \in [t], j \in \{0,1\})$, observe that

$$\frac{\#\big\{\mathbf{s} \in \mathbb{Z}_q^n : b_{i,j} - e_{i,j} = \langle \mathbf{a}_{i,j}, \mathbf{s} \rangle \in \mathbb{Z}_q \ \forall i \in [t], j \in \{0,1\}\big\}}{\#\Big\{\mathbf{s}' \in \mathbb{Z}_q^n : b'_{i,j} = \big\lfloor \langle \mathbf{a}_{i,j}, \mathbf{s}' \rangle \big\rceil_p \in \mathbb{Z}_p \ \forall i \in [t], j \in \{0,1,2\}\Big\}} \geqslant \frac{q^{n-2t}}{(q/p)^{3t} q^{n-3t}} = \frac{p^{3t}}{q^{2t}} \geqslant q^{3ct} > 1,$$

hence at least two distinct good LWE secrets in G must map to the same right-vertex $\mathbf{s}'$, contradicting the definition of G. Hence $d < 3t$.

In reality, our full proof actually shows that $d < 3t$ with probability at least $1 - \nu^t$ over our experiment. A routine counting argument then shows $\exists r \in \{0, 1, \ldots, t-1\}$ such that

$$\begin{aligned}
1 - \nu &\leqslant \Pr\Bigg[\dim\Big(\mathrm{Span}\big(\{\mathbf{a}'_{i,0}, \mathbf{a}'_{i,1}, \mathbf{a}'_{i,2}\}_{i \leqslant r+1}\big)\Big) < 3(r+1) \ \Bigg| \ \dim\Big(\mathrm{Span}\big(\{\mathbf{a}'_{i,0}, \mathbf{a}'_{i,1}, \mathbf{a}'_{i,2}\}_{i \leqslant r}\big)\Big) = 3r\Bigg] \\
&\leqslant \Pr\Big[\mathbf{a}'_{r+1,0} \in \mathbf{V}\Big] + \Pr\Big[\mathbf{a}'_{r+1,1} \in \mathrm{Span}\big(\{\mathbf{a}'_{r+1,0}\}\big) + \mathbf{V}\Big] + \Pr\Big[\mathbf{a}'_{r+1,2} \in \mathrm{Span}\big(\{\mathbf{a}'_{r+1,0}, \mathbf{a}'_{r+1,1}\}\big) + \mathbf{V}\Big] \\
&\leqslant 2\nu + q^{-\Omega(n)} + \Pr\Big[\mathbf{a}'_{r+1,2} \in \mathrm{Span}\big(\{\mathbf{a}'_{r+1,0}, \mathbf{a}'_{r+1,1}\}\big) + \mathbf{V}\Big],
\end{aligned}$$

where the probabilities are all over the randomness of our experiment, $\mathbf{V} = \mathrm{Span}\big(\{\mathbf{a}'_{i,0}, \mathbf{a}'_{i,1}, \mathbf{a}'_{i,2}\}_{i \leqslant r}\big)$, and the last inequality follows from the non-degeneracy of h. Hence

$$\Pr\Big[\mathbf{a}'_{r+1,2} \in \mathrm{Span}\big(\{\mathbf{a}'_{r+1,0}, \mathbf{a}'_{r+1,1}\}\big) + \mathbf{V}\Big] \geqslant 1 - 4\nu,$$

which implies that with probability at least $1 - 2\sqrt{\nu}$ over $\mathbf{V}$ it holds that

$$\Pr_{\substack{\mathbf{a}_1,\mathbf{a}_2 \sim \mathbb{Z}_q^n \\ (\alpha_1,\alpha_2) \in \mathbb{Z}_q^2 \\ \text{s.t. } (\alpha_1,\alpha_2) \neq (0,0)}} \Big[h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2) \in \mathrm{Span}\big(h(\mathbf{a}_1), h(\mathbf{a}_2)\big) + \mathbf{V}\Big] \geqslant 1 - 2\sqrt{\nu}.$$

Our algorithm then simply chooses $r^* \sim \{0, 1, \ldots, t-1\}$ as a guess for r, repeats our experiment substituting r^* for t, and outputs $\mathbf{V} := \mathrm{Span}\big(\{\mathbf{a}'_{i,0}, \mathbf{a}'_{i,1}, \mathbf{a}'_{i,2}\}_{i \leqslant r^*}\big)$. See Section 6.2 for the full proof.

6.1.2 Part II: Recovering H via an Affine Linearity Testing Theorem

In Part I of this proof overview, we showed that the function h satisfies a type of affine linearity test. Specifically, we showed there exists an efficiently computable subspace $\mathbf{V} \subset \mathbb{Z}_q^n$ of constant dimension such that

$$\Pr_{\substack{\mathbf{a}_1,\mathbf{a}_2 \sim \mathbb{Z}_q^n \\ (\alpha_1,\alpha_2) \sim \mathbb{Z}_q^2 \\ \text{s.t. } (\alpha_1,\alpha_2) \neq (0,0)}} \Big[h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2) \in \mathrm{Span}\big(\{h(\mathbf{a}_1), h(\mathbf{a}_2)\} + \mathbf{V}\big)\Big] \geqslant 1 - 2\sqrt{\nu}. \tag{4}$$

Now, we outline an affine linearity testing theorem which concludes there exists an efficiently computable matrix $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ such that

$$\Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathrm{Span}\big(\{\mathbf{H}\mathbf{a}\}\big) + \mathbf{V}\big] \geqslant 1 - \tau.$$

The high-level idea is that we compute a random basis $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\} \subset \mathbb{Z}_q^n$ of $\mathbb{Z}_q^n$, and then compute $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\} \subset \mathbb{Z}_q^n$ that

$$\Pr_{\boldsymbol{\alpha} \sim \mathbb{Z}_q^n}\left[h\Big(\sum_{i=1}^n \alpha_i \mathbf{a}_i\Big) \in \mathrm{Span}\left(\Big\{\sum_{i=1}^n \alpha_i \mathbf{a}'_i\Big\}\right) + \mathbf{V}\right] \geqslant 1 - \tau.$$

Then, we construct an algorithm which simply computes the matrices $\mathbf{A}, \mathbf{A}' \in \mathbb{Z}_q^{n \times n}$ where the i^{th} column of $\mathbf{A}$ (resp., $\mathbf{A}'$) is $\mathbf{a}_i$ (resp., $\mathbf{a}'_i$), and outputs the matrix $\mathbf{H} = \mathbf{A}'\mathbf{A}^{-1} \in \mathbb{Z}_q^{n \times n}$.

For the purpose of this proof overview, we assume that the hypothesis (4) of the affine linearity holds with probability 1, and that $\mathbf{V} = \{\mathbf{0}\}$. Let $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\} \subset \mathbb{Z}_q^n$ is a basis for $\mathbb{Z}_q^n$. We additionally assume that $\{h(\mathbf{a}_1), \ldots, h(\mathbf{a}_n)\} \subset \mathbb{Z}_q^n$ is linearly independent. Finally, for simplicity, here we'll actually show there exists $\{\mathbf{a}'_1, \mathbf{a}'_2, \mathbf{a}'_3\} \subset \mathbb{Z}_q^n$ such that $\forall (\alpha_1, \alpha_2, \alpha_3) \in \big(\mathbb{Z}_2 \times \mathbb{Z}_q \times \mathbb{Z}_q\big) \backslash \{(0,0,0)\}$,

$$h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3) \in \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1 + \alpha_2\mathbf{a}'_2 + \alpha_3\mathbf{a}'_3\}\big).$$

Our techniques follow those in the proof of a major result in algebraic geometry called the Fundamental Theorem of Projective Geometry [Art57]. As we will see, our argument assumes a polynomially bounded number of events, each of which in reality occurs with high probability. So, in our full proof, we can assume all of these events hold simultaneously, allowing us to adapt these techniques while applying the union bound to obtain the conclusion with probability $1 - \tau$. For the full proof, see Section 6.3.

We begin by setting $\mathbf{a}'_1 = h(\mathbf{a}_1) \in \mathbb{Z}_q^n$. Observe that since $h(\mathbf{a}_1 + \mathbf{a}_2) \in \mathrm{Span}\big(\{\mathbf{a}'_1, h(\mathbf{a}_2)\}\big)$, then we can define $\mathbf{a}'_2 \in \mathbb{Z}_q^n$ such that $h(\mathbf{a}_1 + \mathbf{a}_2) \in \mathrm{Span}\big(\{\mathbf{a}'_1 + \mathbf{a}'_2\}\big)$ ($\mathbf{a}'_2$ is well-defined since $\{h(\mathbf{a}_1), h(\mathbf{a}_2)\}$ are linearly independent). We similarly define $\mathbf{a}'_3 \in \mathbb{Z}_q^n$ such that $h(\mathbf{a}_1 + \mathbf{a}_3) \in \mathrm{Span}\big(\{\mathbf{a}'_1 + \mathbf{a}'_3\}\big)$. Note that $\{\mathbf{a}'_1, \mathbf{a}'_2, \mathbf{a}'_3\} \subset \mathbb{Z}_q^n$ are linearly independent. Now, for each $i \in \{2, 3\}$, we can similarly define a map $\pi_i : \mathbb{Z}_q \to \mathbb{Z}_q$ by $\pi_i(\alpha) = \beta \in \mathbb{Z}_q$ such that $h(\mathbf{a}_1 + \alpha\mathbf{a}_i) \in \mathrm{Span}\big(\{\mathbf{a}'_1 + \beta\mathbf{a}'_i\}\big)$. Observe that $\pi_i(0) = 0$ and $\pi_i(1) = 1$. Moreover, if $\alpha_2, \alpha_3 \in \mathbb{Z}_q$, then by rewriting $\mathbf{a}_1 + \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3 = (\mathbf{a}_1 + \alpha_2\mathbf{a}_2) + \alpha_3\mathbf{a}_3 = (\mathbf{a}_1 + \alpha_3\mathbf{a}_3) + \alpha_2\mathbf{a}_2 \in \mathbb{Z}_q^n$, we see that

$$\begin{aligned} h(\mathbf{a}_1 + \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3) &\in \mathrm{Span}\big(\{\mathbf{a}'_1 + \pi_2(\alpha_2)\mathbf{a}'_2, \mathbf{a}'_3\}\big) \cap \mathrm{Span}\big(\{\mathbf{a}'_1 + \pi_3(\alpha_3)\mathbf{a}'_3, \mathbf{a}'_2\}\big) \\ &= \mathrm{Span}\big(\{\mathbf{a}'_1 + \pi_2(\alpha_2)\mathbf{a}'_2 + \pi_3(\alpha_3)\mathbf{a}'_3\}\big). \end{aligned}$$

Next, let $(\alpha_2, \alpha_3) \in \mathbb{Z}_q^2 \backslash \{(0,0)\}$, and let $\mathbf{x} = \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3 \in \mathbb{Z}_q^n$. We'll show that $h(\mathbf{x}) \in \mathrm{Span}\big(\{\pi_2(\alpha_2)\mathbf{a}'_2 + \pi_3(\alpha_3)\mathbf{a}'_3\}\big)$. Note that we have $h(\mathbf{x}) \in \mathrm{Span}\big(\{\mathbf{a}'_2, \mathbf{a}'_3\}\big)$. On the other hand, we can write $\mathbf{x} = (\mathbf{a}_1 + \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3) - \mathbf{a}_1$, hence $h(\mathbf{x}) \in \mathrm{Span}\big(\{h(\mathbf{a}_1 + \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3), \mathbf{a}'_1\}\big) \subset \mathrm{Span}\big(\{\mathbf{a}'_1 + \pi_2(\alpha_2)\mathbf{a}'_2 + \pi_3(\alpha_3)\mathbf{a}'_3, \mathbf{a}'_1\}\big)$. So, we have

$$\begin{aligned} h(\mathbf{x}) &\in \mathrm{Span}\big(\{\mathbf{a}'_2, \mathbf{a}'_3\}\big) \cap \mathrm{Span}\big(\{\mathbf{a}'_1 + \pi_2(\alpha_2)\mathbf{a}'_2 + \pi_3(\alpha_3)\mathbf{a}'_3, \mathbf{a}'_1\}\big) \\ &= \mathrm{Span}\big(\{\pi_2(\alpha_2)\mathbf{a}'_2 + \pi_3(\alpha_3)\mathbf{a}'_3\}\big). \end{aligned}$$

Finally, we show that each map π_i is actually the identity map on $\mathbb{Z}_q$, which completes the proof. Consider the map π_2 (a similar symmetric argument holds for π_3). We proceed by induction on $\alpha \in \mathbb{Z}_q$. The base cases in which $\alpha \in \{0, 1\}$ follow from the above discussion. Suppose that $\pi_2(\alpha - 1) = \alpha - 1$. Let $\mathbf{x} = \mathbf{a}_1 + \alpha\mathbf{a}_2 + \mathbf{a}_3 \in \mathbb{Z}_q^n$. We can first write $\mathbf{x} = (\mathbf{a}_1 + \pi_2(\alpha)\mathbf{a}_2) + \mathbf{a}_3 \in \mathbb{Z}_q^n$, hence $h(\mathbf{x}) \in \mathrm{Span}\big(\{\mathbf{a}'_1 + \pi_2(\alpha)\mathbf{a}'_2, \mathbf{a}'_3\}\big)$, so $\exists(\beta_1, \beta_2) \in \mathbb{Z}_q^2$ such that $h(\mathbf{x}) = \beta_1(\mathbf{a}'_1 + \pi_2(\alpha)\mathbf{a}'_2) + \beta_2\mathbf{a}'_3 \in \mathbb{Z}_q^n$. On the other hand, we can write also write $\mathbf{x} = (\mathbf{a}_1 + (\alpha - 1)\mathbf{a}_2) + (\mathbf{a}_2 + \mathbf{a}_3)$, hence $h(\mathbf{x}) \in \mathrm{Span}\big(\{\mathbf{a}'_1 + (\alpha - 1)\mathbf{a}'_2, \mathbf{a}'_2 + \mathbf{a}'_3\}\big)$, where we have used the induction hypothesis, the conclusion of the previous paragraph, and the base case. So, $\exists(\gamma_1, \gamma_2) \in \mathbb{Z}_q^2$ such that $h(\mathbf{x}) = \gamma_1(\mathbf{a}'_1 + (\alpha - 1)\mathbf{a}'_2) + \gamma_2(\mathbf{a}'_2 + \mathbf{a}_3) \in \mathbb{Z}_q^n$. Finally, we can write $\mathbf{x} = (\mathbf{a}_1 + \mathbf{a}_3) + \alpha_2\mathbf{a}_2 \in \mathbb{Z}_q^n$, and so $h(\mathbf{x}) \in \mathrm{Span}\big(\{\mathbf{a}'_1 + \mathbf{a}'_3, \mathbf{a}'_2\}\big)$. Then $\exists(\delta_1, \delta_2) \in \mathbb{Z}_q^2$ such that $h(\mathbf{x}) = \delta_1(\mathbf{a}'_1 + \mathbf{a}'_3) + \delta_2\mathbf{a}_2 \in \mathbb{Z}_q^n$. By equating these three representations of $h(\mathbf{x})$, and by the linear independence of $\{\mathbf{a}'_1, \mathbf{a}'_2, \mathbf{a}'_3\}$, it follows that $\gamma_1(\alpha - 1) + \gamma_2 = \beta_1\pi_2(\alpha)$, $\gamma_1 = \beta_1$, and $\gamma_2 = \delta_1 = \gamma_1$. Hence $\gamma_1\alpha = \gamma_1\pi_2(\alpha)$, and so $\alpha = \pi_2(\alpha)$ when $\gamma_1 \neq 0$. Indeed, it can be shown that $\gamma_1 \neq 0$ (with overwhelming probability) since h is non-degenerate.

6.2 Recovering $\mathbf{V}$

The Algorithm to Recover $\mathbf{V}$. Let notations be as above. We recover $\mathbf{V}$ as follows.

1. Initialize $\mathbf{V} = \{\mathbf{0}\}$; choose $r \sim \{1, \ldots, t\}$; for $i = 1, \ldots, r$, do the following:
   - choose $\mathbf{a}_{i,0}, \mathbf{a}_{i,1} \sim \mathbb{Z}_q^n$ and $(\alpha_{i,0}, \alpha_{i,1}) \sim \mathbb{Z}_q^2 \backslash \{(0,0)\}$;
   - compute $\mathbf{a}'_{i,j} = h(\mathbf{a}_{i,j})$ for $j = 0, 1, 2$, where $\mathbf{a}_{i,2} = \alpha_{i,0}\mathbf{a}_{i,0} + \alpha_{i,1}\mathbf{a}_{i,1}$;
   - update $\mathbf{V} := \mathbf{V} + \mathrm{Span}\big(\{\mathbf{a}'_{i,0}, \mathbf{a}'_{i,1}, \mathbf{a}'_{i,2}\}\big)$.
2. Output $\mathbf{V}$.

Claim 4. *Let $\mathcal{D}_r$ denote the random procedure used to generate the vectors $\{\mathbf{a}'_{i,0}, \mathbf{a}'_{i,1}, \mathbf{a}'_{i,2}\}_{i=1,\dots,r}$. Suppose the function $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ is such that $\Pr_{\mathcal{D}_t}\big[\dim \mathrm{Span}\big(\{\mathbf{a}'_{i,j}\}_{i,j}\big) = 3t\big] < \eta^{1/3}$. Then with non-negligible probability, the vector space $\mathbf{V}$ output above satisfies $\mathsf{P}(\mathbf{V}) \geqslant 1 - 4\eta^{1/6t}$, where*

$$\mathsf{P}(\mathbf{V}) := \Pr_{\substack{\mathbf{a}_1,\mathbf{a}_2 \sim \mathbb{Z}_q^n \\ (\alpha_1,\alpha_2) \sim \mathbb{Z}_q^2 \backslash \{(0,0)\}}} \Big[h(\alpha_1 \mathbf{a}_1 + \alpha_2 \mathbf{a}_2) \in \mathrm{Span}\big(\{h(\mathbf{a}_1), h(\mathbf{a}_2)\}\big) + \mathbf{V} \Big].$$

Proof. Let $\nu > 0$ be such that $\nu^{3t} = \eta$. Consider an execution of $\mathcal{D}_t$; for $i = 0, \dots, t$, let $\mathbf{V}_i$ denote the vector space $\mathbf{V}$ after the $i-$th iteration, and let $d_i = \dim(\mathbf{V}_i)$. We are given that $\Pr\big[d_t = 3t\big] < \nu^t$; let $r \in \{0, 1, \dots, t-1\}$ be maximal such that $\Pr\big[d_r = 3r\big] \geqslant \nu^r$. We have

$$\begin{aligned} \nu^{r+1} &> \Pr\big[d_{r+1} = 3(r+1)\big] = \Pr\big[d_{r+1} = 3(r+1) \big| d_r = 3r\big] \cdot \Pr\big[d_r = 3r\big] \\ &\geqslant \Pr\big[d_{r+1} = 3(r+1) \big| d_r = 3r\big] \cdot \nu^r, \end{aligned}$$

and so $\Pr\big[d_{r+1} < 3(r+1) \big| d_r = 3r\big] \geqslant 1 - \nu$. Let $\mathbf{a}_0, \mathbf{a}_1 \in \mathbb{Z}_q^n$ and $(\alpha_0, \alpha_1) \in \mathbb{Z}_q^2 \backslash \{(0,0)\}$ be the vectors and scalars drawn during the $(r+1)-$th round of $\mathcal{D}_t$. Note if $d_{r+1} < 3(r+1)$ then it must be that at least one of the following occurs:

$$(1)\ \mathbf{a}'_0 \in \mathbf{V}_r; \quad (2)\ \mathbf{a}'_1 \in \mathbf{V}_r + \mathrm{Span}(\mathbf{a}'_0); \quad (3)\ \mathbf{a}'_2 \in \mathbf{V}_r + \mathrm{Span}\big(\{\mathbf{a}'_0, \mathbf{a}'_1\}\big).$$

By non-degeneracy, the first two points happen with probability at most $\nu + q^{-\Omega(n)}$. Thus, the third point holds with probability at least $1 - 3\nu - q^{-\Omega(n)} \geqslant 1 - 4\nu$, and so it holds with probability $1 - 2\sqrt{\nu}$ over $\mathbf{V}_r$ that

$$\mathsf{P}(\mathbf{V}_r) = \Pr_{\substack{\mathbf{a}_0,\mathbf{a}_1 \sim \mathbb{Z}_q^n \\ (\alpha_0,\alpha_1) \sim \mathbb{Z}_q^2 \backslash \{(0,0)\}}} \Big[h(\alpha_0 \mathbf{a}_0 + \alpha_1 \mathbf{a}_1) \in \mathrm{Span}\big(\{h(\mathbf{a}_0), h(\mathbf{a}_1)\}\big) + \mathbf{V}_r \Big] \geqslant 1 - 2\sqrt{\nu}.$$

The probability that the above algorithm chooses this r is $1/t$. The claim follows. □

Claim 5. *Let notations be as above. Then $\Pr_{\mathcal{D}_t}\big[\dim(\mathbf{V}) = 3t\big] < \eta^{1/3}$.*

Remark. *This is the only place in the paper where we need to use the assumption that $q^{2/3+c} < p < q$.*

Proof. Let $\mathcal{D}$ be the distribution which runs the same random procedure as in $\mathcal{D}_t$ except which also outputs the $\{\mathbf{a}_{i,j}\}$, and additionally which outputs the $\{b_{i,j}\}$ and $\{b'_{i,j}\}$ used to compute h. So specifically, $\mathcal{D}$ outputs

$$\Big\{ (\mathbf{a}_{i,j}, b_{i,j}), (\mathbf{a}'_{i,j}, b'_{i,j}) \Big\}_{\substack{i=1,\dots,t \\ j=0,1,2}} \subset \big(\mathbb{Z}_q^n \times \mathbb{Z}_q\big)^3 \times \big(\mathbb{Z}_q^n \times \mathbb{Z}_p\big)^3$$

where for all $i = 1, \dots, t$:

- $(\mathbf{a}_{i,0}, b_{i,0}), (\mathbf{a}_{i,1}, b_{i,1}) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$;
- $(\alpha_{i,0}, \alpha_{i,1}) \sim \mathbb{Z}_q^2 \backslash \{(0,0)\}$ and $(\mathbf{a}_{i,2}, b_{i,2}) = (\alpha_{i,0}\mathbf{a}_{i,0} + \alpha_{i,1}\mathbf{a}_{i,1}, \alpha_{i,0} b_{i,0} + \alpha_{i,1} b_{i,1})$;
- $(\mathbf{a}'_{i,j}, b'_{i,j}) = f(\mathbf{a}_{i,j}, b_{i,j})$.

Consider a draw $\big(\{(\mathbf{a}_{i,j}, b_{i,j})\}, \{(\mathbf{a}'_{i,j}, b'_{i,j})\}\big) \sim \mathcal{D}$, let $d := \dim\big(\mathrm{Span}\big(\{\mathbf{a}'_{i,j}\}\big)\big)$, and let $S, S' \subset \mathbb{Z}_q^n$ be the following subsets of LWE and LWR secrets:

$$S := \big\{ \mathbf{s} \in \mathsf{G} : b_{i,j} = \langle \mathbf{a}_{i,j}, \mathbf{s} \rangle \ \forall\ i, j \big\}; \text{ and } S' := \big\{ \mathbf{s}' \in \mathbb{Z}_q^n : b'_{i,j} = \big\lfloor \langle \mathbf{a}'_{i,j}, \mathbf{s}' \rangle \big\rceil_p \ \forall\ i, j \big\}.$$

Consider the following three events:

- $\mathbf{E}_1$**:** $d = 3t$;
- $\mathbf{E}_2$**:** $|S| \geqslant q^{-2t-1} \cdot |\mathsf{G}|$;
- $\mathbf{E}_3$**:** $\Pr_{\mathbf{s}\sim S}\big[\mathbf{s}' \in S'\big] \geqslant 1 - \sqrt{3tq\eta}$, where $\mathbf{s}' \in \mathbb{Z}_q^n$ is the unique LWR secret st $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1 - \eta$.

Note that all three events cannot occur simultaneously. Indeed, the events $\mathbf{E}_2$ and $\mathbf{E}_3$ together imply that $^{\#}\{\mathbf{s} \in S : \mathbf{s}' \in S'\} \geqslant (1 - \sqrt{3tq\eta}) \cdot q^{-2t-1} \cdot |\mathsf{G}| \geqslant \frac{1}{2} \cdot q^{-2t-1} \cdot |\mathsf{G}|$, while $\mathbf{E}_1$ implies that $|S'| = \big(q/p\big)^{3t} \cdot q^{-3t} \cdot q^n = p^{-3t} \cdot q^n$. If all three hold then

$$\frac{^{\#}\{\mathbf{s} \in S : \mathbf{s}' \in S'\}}{|S'|} \geqslant \frac{q^{-2t-1} \cdot \delta}{2 \cdot p^{-3t}} > \frac{q^{3tc-1} \cdot \delta}{2} > 1,$$

which violates property 3 of G since it means some $\mathbf{s}' \in S'$ has $^{\#}\{\mathbf{s} \in S : \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1 - \eta\} \geqslant 2$. We finish by showing that both $\mathbf{E}_2$ and $\mathbf{E}_3$ occur with high probability. Specifically, we show that $\Pr_{\mathcal{D}}\big[\mathbf{E}_2 \;\&\; \mathbf{E}_3\big] > 1 - \eta^{1/3}$. Since all three events cannot occur simultaneously, $\Pr_{\mathcal{D}}\big[\mathbf{E}_1\big] < \eta^{1/3}$ must hold. So, Points 1 and 2 below complete the proof.

Claim 6. $\Pr_{\mathcal{D}}\big[\mathbf{E}_2\big] > 1 - q^{-n/3}$.

Proof. Recall $\mathbf{E}_2$ is the event that $|S| \geqslant q^{-2t-1} \cdot |\mathsf{G}|$. In this proof, it will be more convenient to label the $2t$ pairs in $\mathbb{Z}_q^n \times \mathbb{Z}_q$ drawn during $\mathcal{D}$ as $(\mathbf{a}_1, b_1), \ldots, (\mathbf{a}_{2t}, b_{2t})$, rather than $(\mathbf{a}_{i,j}, b_{i,j})$, $i = 1, \ldots, t$ and $j = 0, 1$. Given a draw $\{(\mathbf{a}_i, b_i)\}_{i=1}^{2t}$ during $\mathcal{D}$, let $\mathsf{G}_r = \{\mathbf{s} \in \mathsf{G} : b_i = \langle \mathbf{a}_i, \mathbf{s}\rangle \;\forall\; i = 1, \ldots, r\}$. So $\mathsf{G} = \mathsf{G}_0$ and $S = \mathsf{G}_{2t}$. We have

$$\begin{aligned}
\Pr_{\mathcal{D}}\big[\mathbf{E}_2\big] &= \Pr_{\mathcal{D}}\Big[|S| \geqslant q^{-2t-1} \cdot |\mathsf{G}|\Big] \geqslant \Pr_{\mathcal{D}}\Big[|\mathsf{G}_r| \geqslant q^{-1-1/2t} \cdot |\mathsf{G}_{r-1}| \;\forall\; r = 1, \ldots, 2t\Big] \\
&= \prod_{r=1}^{2t} \Pr_{\mathcal{D}}\Big[|\mathsf{G}_r| \geqslant q^{-1-1/2t} \cdot |\mathsf{G}_{r-1}| \Big| |\mathsf{G}_i| \geqslant q^{-1-1/2t} \cdot |\mathsf{G}_{i-1}| \;\forall\; i = 1, \ldots, r-1\Big].
\end{aligned}$$

We will show that for all $r = 1, \ldots, 2t$, as long as $|\mathsf{G}_{r-1}| \geqslant q^{-r} \cdot |\mathsf{G}|$, then

$$\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\Big[\Pr_{\mathbf{s}\sim\mathsf{G}_{r-1}}\big[b = \langle \mathbf{a}, \mathbf{s}\rangle\big] \geqslant q^{-1-1/2t}\Big] \geqslant 1 - q^{-n/2} \tag{5}$$

holds. This proves the claim as it gives $\Pr_{\mathcal{D}}\big[\mathbf{E}_2\big] \geqslant \big(1 - q^{-n/2}\big)^{2t} > 1 - q^{-n/3}$, so it remains to prove (5). For $b \in \mathbb{Z}_q$, let

$$X_b := \big\{\mathbf{a} \in \mathbb{Z}_q^n : \Pr_{\mathbf{s}\sim\mathsf{G}_{r-1}}\big[\langle \mathbf{a}, \mathbf{s}\rangle = b\big] < q^{-1-1/2t}\big\}.$$

Clearly $\Delta\big(\langle X_b, \mathsf{G}_{r-1}\rangle, \mathsf{Unif}(\mathbb{Z}_q)\big) > q^{-1} \cdot (1 - q^{-1/2t}) \geqslant q^{-2}$. Therefore, by Fact 1,

$$|X_b| \leqslant \frac{q^{n+1}}{|\mathsf{G}_{r-1}| \cdot q^{-4}} \leqslant \frac{q^{n+5}}{q^{-r} \cdot |\mathsf{G}|} \leqslant \frac{q^{n+5+2t}}{\delta \cdot q^n} = q^{2t+5} \cdot \delta^{-1}.$$

We have

$$\begin{aligned}
\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\Big[\Pr_{\mathbf{s}\sim\mathsf{G}_{r-1}}\big[b = \langle \mathbf{a}, \mathbf{s}\rangle\big] < q^{-1-1/2t}\Big] &\leqslant \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\Big[\exists\; b \in \mathbb{Z}_q \text{ st } \mathbf{a} \in X_b\Big] \\
&\leqslant q^{2t+6} \cdot \delta^{-1} \cdot q^{-n} < q^{-n/2},
\end{aligned}$$

proving (5). □

Claim 7. $\Pr_{\mathcal{D}}\big[\mathbf{E}_3\big] \geqslant 1 - \sqrt{3tq\eta}$.

Proof. Recall $\mathbf{E}_3$ is the event that $\Pr_{\mathbf{s}\sim S}\big[\mathbf{s}' \in S'\big] \geqslant 1 - \sqrt{3tq\eta}$, where $\mathbf{s}' \in \mathbb{Z}_q^n$ is the unique $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geqslant 1 - \eta$. We prove $\Pr_{\mathcal{D},\mathbf{s}\sim S}\big[\mathbf{s}' \in S'\big] \geqslant 1 - 3tq\eta$; the claim then follows by averaging. Note that $\Pr_{(\mathbf{a},b)\sim \mathsf{LWE}_{\mathbf{s}}}\big[b' = \big\lfloor\langle \mathbf{a}', \mathbf{s}'\rangle\big\rceil_p \big| b = \langle \mathbf{a}, \mathbf{s}\rangle\big] \geqslant 1 - q\eta$, since χ outputs $e = 0$ with probability at least $1/q$. It follows that

$$\begin{aligned}\Pr_{\mathcal{D},\mathbf{s}\sim S}\big[\mathbf{s}' \in S'\big] &= \Pr_{\mathcal{D},\mathbf{s}\sim \mathsf{G}}\Big[b'_{i,j} = \big\lfloor\langle \mathbf{a}'_{i,j}, \mathbf{s}'\rangle\big\rceil_p \ \forall\ i,j \Big| b_{i,j} = \langle \mathbf{a}_{i,j}, \mathbf{s}\rangle \ \forall\ i,j\Big] \\ &= \Pr_{\mathbf{s}\sim \mathsf{G},\{(\mathbf{a}_{i,j},b_{i,j})\}\sim \mathsf{LWE}_{\mathbf{s}}}\Big[b'_{i,j} = \big\lfloor\langle \mathbf{a}'_{i,j}, \mathbf{s}'\rangle\big\rceil_p \ \forall\ i,j \Big| b_{i,j} = \langle \mathbf{a}_{i,j}, \mathbf{s}\rangle \ \forall\ i,j\Big] \geqslant 1 - 3tq\eta,\end{aligned}$$

by the union bound. □

□

6.3 Recovering H.

In the previous section we showed how to recover a constant dimensional subspace $\mathbf{V} \subset \mathbb{Z}_q^n$ such that $\mathsf{P}(\mathbf{V}) \geqslant 1 - 4\gamma$, where $\gamma = \eta^{1/6t}$. Here, we show how to use h such that $\mathsf{P}(\mathbf{V}) \geqslant 1 - 4\gamma$ holds to recover $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ such that $\mathsf{P}(\mathbf{H}, \mathbf{V}) \geqslant 1 - \tau$ holds where $\tau = nq^2\sqrt{178n\gamma}$. This completes the proof of Lemma 3, and thus also the proof of Theorem 2. Rather than directly recovering $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$, our algorithm will recover vectors $\{\mathbf{a}_i, \mathbf{a}'_i\}_{i=1}^n \subset \mathbb{Z}_q^n$ such that $\{\mathbf{a}_i\}_i$ is linearly independent and such that

$$\Pr_{\alpha_1,\ldots,\alpha_n\sim\mathbb{Z}_q}\Big[h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_n\mathbf{a}_n) \in \mathrm{Span}(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_n\mathbf{a}'_n) + \mathbf{V}\Big] \geqslant 1 - \tau. \tag{6}$$

Given such $\{\mathbf{a}_i, \mathbf{a}'_i\}_i$, we let $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ be the linear map which sends $\mathbf{a}_i$ to $\mathbf{a}'_i$ for all $i = 1, \ldots, n$; $\mathsf{P}(\mathbf{H}, \mathbf{V}) \geqslant 1 - \tau$ follows from (6).

The Algorithm to Recover $\{\mathbf{a}_i, \mathbf{a}'_i\}_i$. Let notations be as above. We recover $\{\mathbf{a}_i, \mathbf{a}'_i\}_i$ as follows.

1. Choose $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n$ such that $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$ is linearly independent.
2. For $i = 1, \ldots, n$, set $\mathbf{a}'_i = \lambda_i h(\mathbf{a}_i) \in \mathbb{Z}_q^n$ for scalars $\{\lambda_i\}_{i=1}^n$ computed as follows:
   - set $\lambda_1 = 1$;
   - for $i \geqslant 2$, let $\lambda_i \in \mathbb{Z}_q$ be the unique scalar such that $h(\mathbf{a}_1 + \mathbf{a}_i) \in \mathrm{Span}\big(\mathbf{a}'_1 + \lambda_i h(\mathbf{a}_i)\big) + \mathbf{V}$; if no such λ_i exists, or if more than one such λ_i exists, halt and give no output.
3. Output $\{\mathbf{a}_i, \mathbf{a}'_i\}_{i=1}^n$.

Note that $h(\mathbf{a}_1 + \mathbf{a}_i) \in \mathrm{Span}\big(\{\mathbf{a}'_1, h(\mathbf{a}_i)\}\big) + \mathbf{V}$ holds for all $i \in \{2, \ldots, n\}$ with probability at least $1 - 4(n-1)q^2\gamma$, since $\mathsf{P}(\mathbf{V}) \geqslant 1 - 4\gamma$. In this case, for all i, there exist scalars (β_1, β_i) such that $h(\mathbf{a}_1 + \mathbf{a}_i) \in \beta_1\mathbf{a}'_1 + \beta_i h(\mathbf{a}_i) + \mathbf{V}$. If $\beta_1 = 0$ then $h(\mathbf{a}_1 + \mathbf{a}_i) \in \mathrm{Span}\big(h(\mathbf{a}_i)\big) + \mathbf{V}$; this happens only with negligible probability since h is non-degenerate. If $\beta_1 \neq 0$ then there exists some scalar $\lambda_i \in \mathbb{Z}_q$ such that $h(\mathbf{a}_1 + \mathbf{a}_i) \in \mathrm{Span}\big(\mathbf{a}'_1 + \lambda_i h(\mathbf{a}_i)\big) + \mathbf{V}$. Note, it is only possible for there to exist two such scalars, $\lambda_i \neq \lambda'_i$ such that

$$h(\mathbf{a}_1 + \mathbf{a}_i) \in \Big(\mathrm{Span}\big(\mathbf{a}'_1 + \lambda_i h(\mathbf{a}_i)\big) + \mathbf{V}\Big) \cap \Big(\mathrm{Span}\big(\mathbf{a}'_1 + \lambda'_i h(\mathbf{a}_i)\big) + \mathbf{V}\Big),$$

if $h(\mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1) + \mathbf{V}$. This also occurs with negligible probability since h is non-degenerate. Thus, the above algorithm completes and gives output without aborting with probability at least $1 - (n/q + 4nq^2\gamma)$.

Henceforth, if $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n$, then we will implicitly condition on $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$ being linearly independent and computing $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\}$ as above. Specifically, we'll show that

$$\Pr_{\substack{\mathbf{a}_1,\ldots,\mathbf{a}_n \sim \mathbb{Z}_q^n \\ (\alpha_1,\ldots,\alpha_n) \sim \mathbb{Z}_q^n \setminus \{\mathbf{0}\}}} \left[h\Big(\sum_i \alpha_i \mathbf{a}_i \Big) \in \mathrm{Span}\left(\Big\{ \sum_i \alpha_i \mathbf{a}'_i \Big\} \right) + \mathbf{V} \right] \geqslant 1 - 178n^3q^4\gamma,$$

hence with non-negligible probability over $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n$, it follows that $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$ is linearly independent, our above algorithm outputs $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\}$ as desired, and (6) holds.

We'll use induction on $r \in \{3, \ldots, n\}$ to show that

$$\mathsf{P}_r := \Pr_{\substack{\mathbf{a}_1,\ldots,\mathbf{a}_n \sim \mathbb{Z}_q^n \\ (\alpha_1,\ldots,\alpha_n) \sim \mathbb{Z}_q^r \setminus \{\mathbf{0}\}}} \left[h\Big(\sum_{i=1}^{r} \alpha_i \mathbf{a}_i \Big) \in \mathrm{Span}\left(\Big\{ \sum_{i=1}^{r} \alpha_i \mathbf{a}'_i \Big\} \right) + \mathbf{V} \right] \geqslant 1 - (80n^2q^4\gamma + 89(r-3)n^2q^4\gamma),$$

hence $\mathsf{P}_n \geqslant 1 - 178n^3q^4\gamma$ as desired. We begin with the following key technical claim, which is proved in Section 6.4.

Claim 8. *For all distinct $i, j \in \{2, \ldots, n\}$, and $(\alpha_1, \alpha_i, \alpha_j) \in \mathbb{Z}_q^3 \setminus \{\mathbf{0}\}$,*

$$h(\alpha_1\mathbf{a}_1 + \alpha_i\mathbf{a}_i + \alpha_j\mathbf{a}_j) \in \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1 + \alpha_i\mathbf{a}'_i + \alpha_j\mathbf{a}'_j\}\big) + \mathbf{V},$$

holds with probability at least $1 - 80n^2q^4\gamma$ over $\{\mathbf{a}_i\}_{i=1}^n$.

The base case of $r = 3$ follows immediately from Claim 8. For the induction step, assume that $\mathsf{P}_{r-1} \geqslant 80n^2q^4\gamma + 89(r-4)n^2q^4\gamma$. Since the probability P_r over $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n, (\alpha_1, \ldots, \alpha_n) \sim \mathbb{Z}_q^n \setminus \{\mathbf{0}\}$, we will assume WLOG that $\alpha_1 \neq 0$. Observe that we can write $\mathbf{z} := h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_r\mathbf{a}_r) = h\big((\alpha_1\mathbf{a}_1 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1}) + \alpha_r\mathbf{a}_r\big) \in \mathrm{Span}\big(\{h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1}), h(\mathbf{a}_r)\}\big) + \mathbf{V} \subset \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}, \mathbf{a}'_r\}\big) + \mathbf{V}$, except with probability $4\gamma + (80n^2q^4\gamma + 89(r-4)n^2q^4\gamma)$ (here we have invoked the induction hypothesis). On the other hand, we can write $\mathbf{z} = h\big((\alpha_1\mathbf{a}_1 + \alpha_r\mathbf{a}_r) + (\alpha_2\mathbf{a}_2 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1})\big) \in \mathrm{Span}\big(\{h(\alpha_1\mathbf{a}_1 + \alpha_r\mathbf{a}_r), \mathbf{y}\}\big) + \mathbf{V} \subset \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1 + \alpha_r\mathbf{a}'_r, \mathbf{y}\}\big) + \mathbf{V}$, except with probability $4\gamma + 80n^2q^4\gamma$, by Claim 8. Hence $\exists \beta_1, \beta_2, \beta_3, \beta_4 \in \mathbb{Z}_q, \mathbf{v}_1, \mathbf{v}_2 \in \mathbf{V}$ such that

$$\beta_1(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}) + \beta_2\mathbf{a}'_r + \mathbf{v}_1 = \mathbf{z} = \beta_3(\alpha_1\mathbf{a}'_1 + \mathbf{a}_r\mathbf{a}'_r) + \beta_4\mathbf{y} + \mathbf{v}_2,$$

and since the distribution of $\mathbf{y}$ is independent of $\{\mathbf{a}_1, \mathbf{a}_r\}$, it follows from the non-degeneracy of h that with overwhelming probability $\beta_1\alpha_1 = \beta_3\alpha_1$ and $\beta_2 = \beta_3\alpha_r$. Since $\alpha_1 \neq 0$ we thus have that $\beta_1 = \beta_3$ and so $\beta_2 = \beta_1\alpha_r$, hence $\mathbf{z} \in \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1 + \cdots + \alpha_r\mathbf{a}'_r\}\big) + \mathbf{V}$ with probability $1 - (4\gamma + (80n^2q^4\gamma + 89(r-4)n^2q^4\gamma) + 4\gamma + 80n^2q^4\gamma + \mathsf{negl}(n)) \geqslant 1 - (80n^2q^4\gamma + 89(r-3)n^2q^4\gamma)$, which completes the induction step.

6.4 Proof of Claim 8

Proof. We must show that for all distinct $i, j \in \{2, \ldots, n\}$ and $(\alpha_1, \alpha_i, \alpha_j) \in \mathbb{Z}_q^3 \setminus \{\mathbf{0}\}$,

$$h(\alpha_1\mathbf{a}_1 + \alpha_i\mathbf{a}_i + \alpha_j\mathbf{a}_j) \in \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1 + \alpha_i\mathbf{a}'_i + \alpha_j\mathbf{a}'_j\}\big) + \mathbf{V}$$

holds with good probability over $\{\mathbf{a}_i\}$. We will build up to analyzing $h(\alpha_1\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j)$. To start out, we know that $h(\mathbf{a}_1) = \mathbf{a}'_1$ and $h(\mathbf{a}_1+\mathbf{a}_i) = \mathbf{a}'_1+\mathbf{a}'_i$ for all $i \in \{2,\ldots,n\}$; these are due to the algorithm specifications. So now consider $h(\mathbf{a}_1+\alpha_i\mathbf{a}_i)$ for $\alpha_i \neq 0,1$. Note $\mathbf{a}_1+\alpha_i\mathbf{a}_i = (1-\alpha_i)\mathbf{a}_1+\alpha_i(\mathbf{a}_1+\mathbf{a}_i)$, and so

$$h(\mathbf{a}_1+\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\{\mathbf{a}'_1,\mathbf{a}'_i\}\big)+\mathbf{V}$$

holds for all $i \in \{2,\ldots,n\}$ and $\alpha_i \in \mathbb{Z}_q$ with probability at least $1-4nq\gamma$ (since $\mathsf{P}(\mathbf{V}) \geqslant 1-4\gamma$). Now, if $h(\mathbf{a}_1+\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\{\mathbf{a}'_1,\mathbf{a}'_i\}\big)+\mathbf{V}$ holds for all (i,α_i), then we can define maps $\pi_i : \mathbb{Z}_q \to \mathbb{Z}_q$ so that $h(\mathbf{a}_1+\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\mathbf{a}'_1+\pi_i(\alpha_i)\mathbf{a}'_i\big)+\mathbf{V}$ always holds. Note $\pi_i(0)=0$ and $\pi_i(1)=1$ for all i.

Fix $i,j \in \{2,\ldots,n\}$ and $(\alpha_1,\alpha_i,\alpha_j) \in \mathbb{Z}_q^3\backslash\{\mathbf{0}\}$. We'll prove the following points, where all probabilities are implicitly over $\{\mathbf{a}_i\}$:

Point 1: $h(\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\{\alpha_i h(\mathbf{a}_i)\}\big)+\mathbf{V}$ with probability $1-5\gamma$.

Point 2: $h(\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \mathrm{Span}\big(\{\mathbf{a}'_1+\pi_i(\alpha_i)\mathbf{a}'_i+\pi_j(\alpha_j)\mathbf{a}'_j\}\big)+\mathbf{V}$ with probability $1-9\gamma$.

Point 3: $h(\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \mathrm{Span}\big(\{\pi_i(\alpha_i)\mathbf{a}'_i+\pi_j(\alpha_j)\mathbf{a}'_j\}\big)+\mathbf{V}$ with probability $1-18\gamma$.

Point 4: π_i and π_j are the identity maps with probability $1-62q\gamma$.

Point 5: $h(\alpha_1\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1+\alpha_i\mathbf{a}'_i+\alpha_j\mathbf{a}'_j\}\big)+\mathbf{V}$ with probability $1-76q\gamma$.

This will complete the proof of Claim 8, since by Point 5 and the union bound it then holds with probability $1-(4nq\gamma+76n^2q^4\gamma) \geqslant 1-80n^2q^4\gamma$ that $h(\alpha_1\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \mathrm{Span}\big(\{\alpha_1\mathbf{a}'_1+\alpha_i\mathbf{a}'_i+\alpha_j\mathbf{a}'_j\}\big)+\mathbf{V}$ $\forall i,j \in \{2,\ldots,n\}, (\alpha_1,\alpha_i,\alpha_j) \in \mathbb{Z}_q^3\backslash\{\mathbf{0}\}$.

Point 1. Note $\alpha_i\mathbf{a}_i = -\mathbf{a}_1+(\mathbf{a}_1+\alpha_i\mathbf{a}_i)$, and so $h(\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\{\mathbf{a}'_1,\mathbf{a}'_i\}\big)+\mathbf{V}$ holds with probability $1-4\gamma$. This means that either

$$h(\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\{\mathbf{a}'_i\}\big)+\mathbf{V};\ \text{or } \mathbf{a}'_1 \in \mathrm{Span}\big(\big\{h(\alpha_i\mathbf{a}_i),\mathbf{a}'_i\big\}\big)+\mathbf{V}.$$

The latter occurs only with negligible probability since h is non-degenerate. Hence $h(\alpha_i\mathbf{a}_i) \in \mathrm{Span}\big(\{\alpha_i h(\mathbf{a}_i)\}\big)+\mathbf{V}$ with probability $1-5\gamma$.

Point 2. Note $\alpha_j\mathbf{a}_j+(\mathbf{a}_1+\alpha_i\mathbf{a}_i) = \mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j = \alpha_i\mathbf{a}_i+(\mathbf{a}_1+\alpha_j\mathbf{a}_j)$, and so

$$h(\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \Big(\mathrm{Span}\big(\{\mathbf{a}'_i,\mathbf{a}'_1+\pi_j(\alpha_j)\mathbf{a}'_j\}\big)+\mathbf{V}\Big) \cap \Big(\mathrm{Span}\big(\{\mathbf{a}'_j,\mathbf{a}'_1+\pi_i(\alpha_i)\mathbf{a}'_i\}\big)+\mathbf{V}\Big)$$

holds with probability $1-8\gamma$. So $\exists A,B,A',B' \in \mathbb{Z}_q$ such that

$$A\mathbf{a}'_i+B\cdot\big(\mathbf{a}'_1+\pi_j(\alpha_j)\mathbf{a}'_j\big) \in A'\mathbf{a}'_j+B'\cdot(\mathbf{a}'_1+\pi_i(\alpha_i)\mathbf{a}'_i\big)+\mathbf{V}.$$

As we have seen a few times by now, either $B=B'$ or else $\mathbf{a}'_1 \in \mathrm{Span}\big(\{\mathbf{a}'_i,\mathbf{a}'_j\}\big)+\mathbf{V}$ and the latter happens with negligible probability by non-degeneracy. Therefore, $B=B'$ except with negligible probability. Similarly, $A=\pi_i(\alpha_i)B$, and so $h(\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \mathrm{Span}(\mathbf{a}'_1+\pi_i(\alpha_i)\mathbf{a}'_i+\pi_j(\alpha_j)\mathbf{a}'_j)+\mathbf{V}$ holds with probability at least $1-9\gamma$.

Point 3. Note $h(\alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) \in \mathrm{Span}\big(\{\mathbf{a}'_i, \mathbf{a}'_j\}\big) + \mathbf{V}$ with probability $1 - 4\gamma$. Additionally, we can write $\alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j = -\mathbf{a}_1 + (\mathbf{a}_1 + \alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j)$ and so

$$h(\alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) \in \mathrm{Span}\big(\{\mathbf{a}'_1, \mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j\}\big) + \mathbf{V}$$

holds with probability $1 - (4\gamma + 9\gamma) = 1 - 13\gamma$ by Point 2. Thus, with probability at least $1 - 17\gamma$, there exist scalars $A, B, A', B' \in \mathbb{Z}_q$ such that

$$A\mathbf{a}'_i + B\mathbf{a}'_j = A'\mathbf{a}'_1 + B'(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j).$$

By non-degeneracy, $A' = -B'$, $A = B'\pi_i(\alpha_i)$, and $B = B'\pi_j(\alpha_j)$ hold except with negligible probability. So $h(\alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) \in \mathrm{Span}\big(\{\pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j\}\big) + \mathbf{V}$ holds with probability $1 - 18\gamma$.

Point 4. We first prove that $\pi_i(\alpha) = \alpha$, $\forall \alpha \in \mathbb{Z}_q$, by induction on $\alpha \in \mathbb{Z}_q$. We have already seen that $\pi_i(0) = 0$ and $\pi_i(1) = 1$. So assume $\pi_i(\alpha - 1) = \alpha - 1$, and write $\mathbf{a}_1 + \alpha \mathbf{a}_i + \mathbf{a}_j$ in three different ways:

$$(\mathbf{a}_1 + \mathbf{a}_i) + ((\alpha - 1)\mathbf{a}_i + \mathbf{a}_j) = \mathbf{a}_j + (\mathbf{a}_1 + \alpha \mathbf{a}_i) = (\mathbf{a}_1 + \mathbf{a}_j) + \alpha \mathbf{a}_i.$$

By applying the union bound over three invocations of $\mathsf{P}(\mathbf{V}) \geqslant 1 - 4\gamma$ and Point 3, we have that with probability $1 - 30\gamma$, $h(\mathbf{a}_1 + \alpha \mathbf{a}_i + \mathbf{a}_j)$ is contained in:

$$\Big(\mathrm{Span}\big(\{\mathbf{a}'_1 + \mathbf{a}'_i, (\alpha - 1)\mathbf{a}'_i + \mathbf{a}'_j\}\big) \cap \mathrm{Span}\big(\{\mathbf{a}'_j, \mathbf{a}'_1 + \pi_i(\alpha)\mathbf{a}'_i\}\big) \cap \mathrm{Span}\big(\{\mathbf{a}'_1 + \mathbf{a}'_j, \mathbf{a}'_i\}\big)\Big) + \mathbf{V},$$

in which case $\exists A, B, A', B', A'', B'' \in \mathbb{Z}_q$ such that $h(\mathbf{a}_1 + \alpha \mathbf{a}_i + \mathbf{a}_j)$ is equal to

$$A(\mathbf{a}'_1 + \mathbf{a}'_i) + B((\alpha - 1)\mathbf{a}'_i + \mathbf{a}'_j) = A'\mathbf{a}'_j + B'(\mathbf{a}'_1 + \pi_i(\alpha)\mathbf{a}'_i) = A''(\mathbf{a}'_1 + \mathbf{a}'_j) + B''\mathbf{a}'_i.$$

Solving for $\mathbf{a}'_1$ gives $A'' = B' = A$. Solving for $\mathbf{a}'_j$ gives $A'' = A' = B$. In particular, $A = B = B'$. Solving for $\mathbf{a}'_i$ gives $\pi_i(\alpha) = \alpha$, as desired. We incurred a loss of $30\gamma + \mathsf{negl}(n) \leqslant 31\gamma$ to take a single step in the induction. Therefore, $\pi_i(\alpha) = \alpha$, $\forall \alpha \in \mathbb{Z}_q$, with probability at least $1 - 31q\gamma$, and the conclusion follows from the union bound.

Point 5. We may assume that $\alpha_1 \neq 0$ since the case in which $\alpha_1 = 0$ is handled by Points 3 and 4. By writing

$$\mathbf{x} := \alpha_1 \mathbf{a}_1 + \alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j = \alpha_1(\mathbf{a}_1 + \alpha_1^{-1}\alpha_i \mathbf{a}_i + \alpha_1^{-1}\alpha_j \mathbf{a}_j)$$

and applying Points 1, 2, and 4, we see that with probability at least $1 - (5\gamma + 9\gamma + 62q\gamma) \geqslant 1 - 76q\gamma$,

$$\begin{aligned} h(\mathbf{x}) \in \mathrm{Span}\big(\alpha_1 h(\mathbf{a}_1 + \alpha_1^{-1}\alpha_i \mathbf{a}_i + \alpha_1^{-1}\alpha_j \mathbf{a}_j)\big) + \mathbf{V} &\subset \mathrm{Span}\big(\alpha_1(\mathbf{a}'_1 + \alpha_1^{-1}\alpha_i \mathbf{a}'_i + \alpha_1^{-1}\alpha_j \mathbf{a}'_j)\big) + \mathbf{V} \\ &= \mathrm{Span}\big(\alpha_1 \mathbf{a}'_1 + \alpha_i \mathbf{a}'_i + \alpha_j \mathbf{a}'_j\big) + \mathbf{V}. \end{aligned}$$

□

References

[AA16] Jacob Alperin-Sheriff and Daniel Apon. Dimension-preserving reductions from LWE to LWR. *IACR Cryptol. ePrint Arch.*, page 589, 2016.

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.

[AKPW13] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited - new reduction, properties and applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 57–74. Springer, 2013.

[Art57] Emil Artin. *Geometric Algebra*. 1957.

[BGM+16] Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. On the hardness of learning with rounding over small modulus. In Eyal Kushilevitz and Tal Malkin, editors, *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 209–224. Springer, 2016.

[BGV11] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully homomorphic encryption without bootstrapping. *IACR Cryptol. ePrint Arch.*, 2011:277, 2011.

[BLL+15] Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: using the rényi divergence rather than the statistical distance. *IACR Cryptol. ePrint Arch.*, 2015:483, 2015.

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. *CoRR*, abs/1306.0281, 2013.

[BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737. Springer, 2012.

[BV15] Zvika Brakerski and Vinod Vaikuntanathan. Constrained key-homomorphic prfs from standard lattice assumptions - or: How to secretly embed a circuit in your PRF. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 1–30. Springer, 2015.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.

[GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In David S. Johnson, editor, *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washigton, USA*, pages 25–32. ACM, 1989.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.

[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92. Springer, 2013.

[LLTT05] Chia-Jung Lee, Chi-Jen Lu, Shi-Chun Tsai, and Wen-Guey Tzeng. Extracting randomness from multiple independent sources. *IEEE Trans. Inf. Theory*, 51(6):2224–2227, 2005.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer, 2012.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342. ACM, 2009.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 89–114. Springer, 2019.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.