Title
Novel Linearity Tests with Applications to Lattices and Learning Problems

Author
Newton, Parker

Publication Date
2022

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Novel Linearity Tests with Applications to Lattices and Learning Problems

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

by

Parker Newton

June 2022

Dissertation Committee:

Dr. Silas Richelson, Chairperson
Dr. Tao Jiang
Dr. Jiasi Chen
Dr. Samet Oymak

The Dissertation of Parker Newton is approved:

Committee Chairperson

University of California, Riverside

Acknowledgments

First and foremost, I would like to thank my advisor, Silas Richelson, for all of his mentorship, instruction, and patience, as well as for opening my eyes to the expansive world of cryptography and theoretical computer science. Without his guidance and support, I would not be where I am today. I would also like to acknowledge the other members of my dissertation committee, namely Tao Jiang, Jiasi Chen, and Samet Oymak, for their guidance and support as well.

Next, I thank my family, including my parents, Scott and Cynthia Newton, my sister, Courtney Newton, and my godmother, Susan Leavitt, for their infinite love and support throughout the course of my academic career.

I would also like to acknowledge our cryptography and theory group at U.C. Riverside, specifically Sourya Roy, Irem Ergun, and Yugarshi Shashwat. Also, special thanks to Chase Wilson for his contributions to our group. I'd also like to acknowledge our Algorithms and Theory Lab in the Computer Science department at U.C. Riverside, specifically Abid Hasan, Huong Luu, Miguel Coviello, Dipankar Baisya, Qihua Liang, Weihua Pan, and Abbas Roayaei Ardakany.

Finally, I'd like to give a special acknowledgement to Arman Elahi, Immanuel Amirtharaj, Eric Bonilla, Shachi Kakkar, and Alicia Chan for their unconditional friendship and support over the course of my graduate studies; truly, the best friends one could have.

ABSTRACT OF THE DISSERTATION

Novel Linearity Tests with Applications to
Lattices and Learning Problems

by

Parker Newton

Doctor of Philosophy, Graduate Program in Computer Science
University of California, Riverside, June 2022
Dr. Silas Richelson, Chairperson

Linearity testing is an area of theoretical computer science which has been widely studied over the last few decades, and lends itself to numerous applications in cryptography and coding theory. Beginning with the seminal work of [BLR93], the goal of linearity testing is to design an algorithm which distinguishes functions that are linear from those that are far from linear. Linearity tests are characterized by two parameters: the test acceptance parameter ε and agreement parameter δ. Linearity tests have the property that they should always accept when the input function is linear, but if the function is $(1-\delta)-$far from linear should reject with probability $1-\varepsilon$. We build on previous work in this space to construct two novel high dimensional linearity tests, each of which operate in what we refer to as a "hybrid" mode between the high and low agreement regimes. Additionally, we provide concrete applications of each to cryptography. Specifically, we use our first novel linearity test to aid us in proving a lattice hardness result; we use our second novel linearity test in the course of proving an approximation variant of the celebrated Goldreich-Levin Theorem in high dimension with low agreement.

Lattice-based cryptography is the theory of constructing cryptosystems whose security is based on the assumed hardness of lattice problems. Several hard lattice problems, such as the renowned problem of Learning with Errors (LWE), are conjectured to be hard even in the quantum model of computation. As such, many lattice-based cryptosystems serve as the most promising candidates to offer robust security guarantees against quantum adversaries. Beginning with the seminal work of Oded Regev in 2005, a plethora of cryptosystems have been constructed from LWE, including symmetric key encryption, public key encryption, digital signatures, trapdoor functions, fully homomorphic encryption, and many more. However, for many years we had been unable to construct certain types of "deterministic" cryptosystems directly from LWE, such as pseudorandom function (PRF) families. In 2012, Banerjee, Peikert, and Rosen (EUROCRYPT'12) introduced the problem of Learning with Rounding (LWR), reduced LWE to LWR when the modulus parameter q is super-polynomial in the lattice dimension n, and constructed a PRF family from LWR. The assumption that q is super-polynomial in n, though, does not provide for practical instantiations of LWR-based cryptosystems; in order to achieve these, we need q to be polynomial in n. Many subsequent works provided improved reductions from LWE to LWR for polynomial q, but all such works impose an *a priori* bound on the number m of input samples to the reduction.

We show that when q is polynomial and m is unbounded, there does not exist a certain type of reduction from LWE to LWR. In particular, every prior reduction from LWE to LWR is of this type. Hence no prior reduction from LWE to LWR will work for polynomial q and unbounded m. Our result *does not* imply that LWR is easy; instead it

asserts that any prior reduction from LWE to LWR for polynomial q and unbounded m cannot operate like any prior reduction in the literature.

In the course of proving our result, we introduce a high-dimensional approximate version of the celebrated Goldreich-Levin Theorem (STOC'89). The Goldreich-Levin Theorem at its core solves a learning problem; specifically, it states that any function which predicts the inner product of a secret vector and a uniformly random vector with any advantage over randomly guessing admits an algorithm which recovers the secret vector. Over the last few decades, the Goldreich-Levin Theorem has found numerous applications in cryptography and coding theory, such as to symmetric key cryptography and error correction codes. In our work, we introduce a novel conditional linearity test, and reduce our approximate version of the Goldreich-Levin Theorem in higher dimension to proving such a conditional linearity testing theorem.

Contents

List of Figures

Chapter 1

Introduction

Linearity testing is the problem of algorithmically distinguishing between functions which are linear and those which are far from linear. Beginning with the seminal work of [BLR93], the goal of linearity testing is to design an algorithm $\mathcal{T}$ which gets oracle access to a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ $(m, n, q \in \mathbb{N})$, and outputs 1 if f is linear, and 0 if f is far from linear. Linearity tests have found a plethora of applications in theoretical computer science over the last few decades, such as to coding theory (locally testable codes [FS95], [RS96], [GS02], [DEL$^+$21]), program checking and PCP systems ([ALM$^+$92], [BGS95], [Hås96], [Hås97]), and cryptography (linear cryptanalysis [Mat93], [SS22]).

In slightly more detail, if $f, g : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ are functions, then we define $\operatorname{agree}(f, g) := \Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\left[f(\mathbf{x}) = g(\mathbf{x})\right]$, and $\operatorname{dist}(f, g) := 1 - \operatorname{agree}(f, g)$. We say that an oracle algorithm $\mathcal{T}$ is an $(\varepsilon, \delta)-$linearity test $(\varepsilon, \delta > 0)$ if for every function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$, if f is linear, then $\Pr\left[\mathcal{T}^f = 1\right] = 1$, and if $\operatorname{dist}(f, \phi) \geq 1 - \delta$ for every linear map $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$, then $\Pr\left[\mathcal{T}^f = 0\right] \geq 1 - \varepsilon$. Henceforth, we will consider the completeness condition in the definition of an $(\varepsilon, \delta)-$linearity test $\mathcal{T}$ (*i.e.*, $\mathcal{T}^f = 1$ with probability 1 when $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ is linear) as implicit in the definition. So, we will equivalently define an $(\varepsilon, \delta)-$linearity test $\mathcal{T}$ as follows: If $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ is a function such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\left[\mathcal{T}^f = 1\right] \geq \varepsilon$, then there exists a linear map $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ such that $\mathrm{agree}(f, \phi) \geq \delta$. We will call ε the *test acceptance* parameter and δ the *agreement* parameter.

The original work of [BLR93] in this space constructed an $(1-\varepsilon, 1-9\varepsilon/2)-$linearity test when $q = 2$ and $n = 1$, and the follow-up works of [BGLR93, BS94, BCH$^+$96, BCLR04] improved this result by obtaining an $(\varepsilon, \varepsilon/3 + \mathcal{O}(\varepsilon^2))-$linearity test. In [SW06], the authors generalized the notion of linearity testing to affine homomorphism testing, and constructed an $(1 - \varepsilon, 1 - 5\varepsilon)-$affine homomorphism test for general q, n.

A natural question which arises in this space is can we construct a high-dimensional linearity test (*i.e.*, for general n) in prime characteristic q for low ε. In particular, it would be desirable for such a high-dimensional low agreement linearity test to have $\delta \geq \mathsf{poly}(\varepsilon)$. In [Sam07], the author constructed an $(\varepsilon, q^{-\Omega(n)}\mathsf{poly}(\varepsilon))-$linearity test; note in their construction that δ is exponentially small in n. The breakthrough result of [San12] improved this result by constructing an $(\varepsilon, q^{-\mathsf{polylog}(1/\varepsilon)})-$linearity test, in which δ is quasi-polynomial in ε. Indeed, the construction of a high-dimensional $(\varepsilon, \mathsf{poly}(\varepsilon))-$linearity test has since evaded the community. However, it was shown that there exists a high-dimensional $(\varepsilon, \mathsf{poly}(\varepsilon))-$linearity test assuming the Polynomial Freiman-Ruzsa (PRF) Conjecture, a major conjecture in additive combinatorics. Unfortunately, a proof of the Polynomial Freiman-Ruzsa Conjecture today remains an open problem in additive combinatorics. Interestingly enough, the fact that the PFR Conjecture is sufficient for constructing an

$(\varepsilon, \mathsf{poly}(\varepsilon))-$linearity test provides a data point on the difficulty of, and quite possibly exposes a barrier in, constructing such a high dimensional low agreement linearity test.

In this dissertation, we prove two novel linearity testing theorems, and demonstrate a concrete application of each to lattices and learning problems, respectively. Both linearity tests are high dimensional in prime characteristic q, and can be viewed as a “hybrid” between the high and low test acceptance parameter regimes; we expand on what we mean by this hybrid model in the next section. The application of our first novel linearity test is to the hardness of lattice problems. Specifically, we continue a line of work on the hardness of a lattice problem called Learning with Rounding (LWR) from the renowned lattice problem of Learning with Errors (LWE). We show that in a specific parameter regime, there does not exist a type of reduction from LWE to LWR. Our application of our second novel linearity test is to a celebrated learning problem in theoretical computer science called the Goldreich-Levin Theorem. Specifically, we continue a line of work on list-decoding group homomorphism codes by applying our second novel linearity test to proving a high dimensional Goldreich-Levin Theorem; again, we expand on our result in the next section.

1.1 Our Contributions

Here we elaborate on our contributions in this dissertation, which include two novel linearity testing theorems and concrete applications of them to lattice-based cryptography and the Goldreich-Levin Theorem in high dimension, respectively.

1.1.1 A Novel High Dimensional Linearity Test

We introduce a novel high dimensional linearity test, captured by the following theorem.

Theorem 1 (Informal) *Let $n, q \in \mathbb{N}$ such that $q = \mathsf{poly}(n)$ is prime, and $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ be a function. If $\forall(\alpha_1, \alpha_2) \in \mathbb{Z}_q^2$,*

$$\Pr_{\mathbf{a}_1,\mathbf{a}_2 \sim \mathbb{Z}_q^n}\big[h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2) \in \mathrm{Span}(h(\mathbf{a}_1), h(\mathbf{a}_2))\big] \geq 1 - \varepsilon,$$

for $\varepsilon > 0$, then there exists a linear map $\phi : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ such that

$$\Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathrm{Span}(\phi(\mathbf{a}))\big] \geq 1 - \mathcal{O}(nq^2\varepsilon).$$

Strictly speaking, this test isn't exactly a linearity test, and as such we call it the $(1-\varepsilon)-$*quasi-linearity test.* However, note that if we apply an averaging argument to the conclusion of Theorem 1, we obtain the following corollary.

Corollary 2 (Informal) *Let $n, q \in \mathbb{N}$ such that $q = \mathsf{poly}(n)$ is prime, and $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ be a function. If $\forall(\alpha_1, \alpha_2) \in \mathbb{Z}_q^2$,*

$$\Pr_{\mathbf{a}_1,\mathbf{a}_2 \sim \mathbb{Z}_q^n}\big[h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2) \in \mathrm{Span}(h(\mathbf{a}_1), h(\mathbf{a}_2))\big] \geq 1 - \varepsilon,$$

for $\varepsilon > 0$, then there exists a linear map $\phi' : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ such that

$$\Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}\big[h(\mathbf{a}) = \phi'(\mathbf{a})\big] \geq (1 - \mathcal{O}(nq^2\varepsilon))/q.$$

Moreover, if we examine our quasi-linearity test more closely, we discover that by similarly applying an averaging argument to the hypothesis of the corollary, it follows that $\exists(\alpha_1, \alpha_2) \in \mathbb{Z}_q^2$

$$\Pr_{\mathbf{a}_1,\mathbf{a}_2 \sim \mathbb{Z}_q^n}\big[h(\mathbf{a}_1 + \mathbf{a}_2) = \alpha_1 h(\mathbf{a}_1) + \alpha_2 h(\mathbf{a}_2)\big] \geq (1 - \varepsilon)/q^2, \tag{1.1}$$

and viewed through this lens, our quasi-linearity test appears to resemble a traditional high dimensional *low test acceptance* linearity test. Of course in reality, (1.1) is not equivalent to the hypothesis in Theorem 1 and Corollary 2, and thus our result is not equivalent to a high dimensional low test acceptance linearity test. However, our result can be viewed as somewhat of a "hybrid" high dimensional linearity test between the high and low agreement regimes. Finally, while most traditional linearity testing theorems ([BLR93], [BCH$^+$96], [SW06], etc.) use combinatorial and Fourier analytic techniques, our proof of Theorem 1 uses purely algebraic techniques. Indeed, our proof follows that of the Fundamental Theorem of Projective Geometry ([Art57]).

1.1.2 Application to Lattices

Lattice-based cryptography is the theory of constructing cryptosystems whose security is based on the assumed hardness of lattice problems. A lattice is defined as a free $\mathbb{Z}-$module, and can be thought of as the $\mathbb{Z}-$span of a set of linearly independent vectors in $\mathbb{R}^n$. Several computational problems defined on lattices, such as Shortest Vector Problem (SVP), Closest Vector Problem (CVP), Short Integer Solution (SIS), Short Basis Problem, etc., have been studied thoroughly for the past few decades. In particular, SVP is widely conjectured to be hard even in the quantum model of computation. As such, cryptosystems whose security can be based on the assumed hardness of SVP serve as candidates for post-quantum cryptography. The seminal works of [Ajt96] and [Ajt99] showed that many of these lattice problems are all equivalent in hardness to SVP. The breakthrough work of Oded Revev in 2005 ([Reg05]) introduced the problem of Learning with Errors (LWE), which can be intuitively thought of as the problem of solving a noisy random linear system, in which each linear equation is perturbed by a small error term sampled from a discrete Gaussian distribution defined on a lattice. Regev provided a quantum reduction from SVP to LWE and constructed symmetric key encryption (SKE) and public key encryption (PKE) from LWE. A line of follow-up work ([Pei09], [BLP$^+$13]) provided improved reductions from SVP to LWE, which culminated in a classical reduction from SVP to LWE, as well as improved constructions of SKE and PKE from LWE.

In the years since, the community has constructed an extensive list of cryptographic primitives from LWE which enjoy robust security guarantees against quantum adversaries, such as digital signatures ([Pei09]), trapdoor functions ([GPV08], [MP12]), fully homomorphic encryption ([BGV11], [Bra12], [FV12], [GSW13]), non-interactive zero knowledge proofs ([PS19]), and many more. However, one cryptographic primitive whose construction from LWE evaded cryptographers for many years was pseudorandom function (PRF) families. A PRF family is a family $\mathcal{F}$ of functions $f : \mathbb{Z}_2^m \to \mathbb{Z}_2^n$ that is efficiently sampleable and pseudorandom (*i.e.*, the distribution which samples a function $f \sim \mathcal{F}$ and outputs f is computationally indistinguishable from that which samples a uniformly random function $f : \mathbb{Z}_q^m \to \mathbb{Z}_2^n$ and outputs f). An explicit construction of PRF's from LWE evaded the community for so long because of the deterministic nature of PRF's; that is, a PRF family $\mathcal{F}$ as the property that randomness is used to sample $f \sim \mathcal{F}$, but no randomness should be employed thereafter to evaluate $f(\mathbf{x})$ on any input $\mathbf{x}$. The distribution of LWE, which is inherently dependent on sampling error terms from a probabilistic distribution (specifically, a discrete Gaussian distribution on a lattice), lends itself well to constructing encryption schemes and many public key cryptographic primitives, but not to constructing "deterministic" cryptosystems such as PRF's. However, in 2012, the work of Banerjee, Peikert, and Rosen ([BPR12]) introduced a novel learning problem called Learning with Rounding (LWR), which is almost exactly the same as LWE, but instead of perturbing each linear equation by an error term sampled from a probabilistic distribution, the error in each linear equation is introduced by rounding the equation (an element in $\mathbb{Z}_q$, for some modulus q which is polynomial in the security parameter n), down to the nearest integer (mod p), for an integer $p < q$. Banerjee et al. showed that LWE reduces to LWR when $q = n^{\omega(1)}$ and constructed PRF's from LWR in this parameter regime.

However, the assumption that $q = n^{\omega(1)}$ does not admit very computationally practical instantiations of PRF's; ideally, q should be a small polynomial in n. A line of follow-up work ([AKPW13, BGM$^+$16, AA16]) provided improved reductions from LWE to LWE with $q = \mathsf{poly}(n)$, but all such reductions imposed an *a priori* bound on the number m of input samples to the reduction. Also, every known reduction from LWE to LWR ([BPR12, AKPW13, BGM$^+$16, AA16]) operates in the follow general method: Take as input LWE samples $\{(\mathbf{a}_i, b_i)\}_{i=1}^m \in (\mathbb{Z}_q^n \times \mathbb{Z}_q)^m$, map these to LWR samples $\{(\mathbf{a}'_i, b'_i)\}_{i=1}^m$ point-wise, according to some function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ $((\mathbf{a}'_i, b'_i) = f(\mathbf{a}_i, b_i), \forall i)$, invoke a solver $\mathcal{S}$ for LWR on the resulting LWR samples, and output the secret $\mathbf{s} \in \mathbb{Z}_q^n$ computed by $\mathcal{S}$. Therefore, these reductions all construct an efficient function f which maps LWE samples computed under an LWE secret $\mathbf{s}$ to LWR samples under $\mathbf{s}$. We call reductions from LWE to LWR which conform to the aforementioned template a *pointwise reduction* parameterized by the function f. We show that when $q = \mathsf{poly}(n)$ is prime, m is unbounded, and $p > q^{2/3}$, then there cannot exist a reduction from LWE to LWR, unless

LWE is tractable. We state our result informally in the following theorem; see Chapter 3 for the formal theorem statement.

Theorem 3 (Informal) *Let $n, p, q \in \mathbb{N}$ be integers such that such that $q = \mathsf{poly}(n)$ is prime and $q^{2/3+c} < p < q = \mathsf{poly}(n)$ for a small constant $c > 0$, and let χ be a discrete Gaussian distribution on $\mathbb{Z}_q$. If there exists a pointwise reduction $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ from LWE to LWR, then there exists an efficient algorithm which directly solves LWE.*

Briefly, our proof works by first showing if there exists a pointwise reduction $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ from LWE to LWR, then f satisfies our quasi-linearity test of Theorem 1. Then, by Corollary 2, f agrees with an linear map with non-negligible probability. We use the linear agreement of f together with statistical properties of a pointwise reduction to efficiently extract the LWE secret $\mathbf{s}$ from m input LWE samples $\{(\mathbf{a}_i, b_i)\}_{i=1}^m$, thereby efficiently solving LWE. See Chapter 3 for further details.

1.1.3 A Novel High Dimensional Conditional Affine Linearity Test

We prove a second novel linearity testing theorem, which is a high dimensional conditional affine linearity test. We state our result informally below.

Theorem 4 (Informal) *Let $m, d, q \in \mathbb{N}$ such that q is prime. Let $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^d$ be a function, and $T \subset \mathbb{Z}_q^m$ a subset of density λ. If*

$$\Pr_{\mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\left[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y}) \;\middle|\; \mathbf{x}, \mathbf{y}, \mathbf{x}+\mathbf{y} \in T\right] \geq 1 - \frac{(q-1)\lambda^3}{8q^{3d-2}},$$

then there exists an affine map $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{d \times m} \times \mathbb{Z}_q^d$ such that

$$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\left[f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} \;\middle|\; \mathbf{x} \in T\right] \geq \left(1/q + \mathsf{poly}(\lambda, q^{-d})\right)^d.$$

Our theorem proves a high dimensional conditional affine linearity test in characteristic q. Although the parameters suggest an affine linearity test with high test acceptance and low agreement, we note that the hypothesis of the theorem includes a conditional test acceptance property; if we remove the condition, we obtain a low test acceptance. This low test acceptance condition is clearly not sufficient for our conclusion, but instead demonstrates how our conditional affine linearity test can be viewed as a "hybrid" between the high and low test acceptance regimes.

1.1.4 Application to the Goldreich-Levin Theorem in High Dimension

We apply our conditional affine linearity test to solving an approximate version of the Goldreich-Levin Theorem in high dimension. The Goldreich-Levin Theorem ([GL89]) is a cornerstone theorem in theoretical computer science with numerous applications to, for example, cryptography ([GL89]) and coding theory (list-decoding group homomorphism codes [GKS06],[DGKS08],[BBW18]). At its core, the Goldreich-Levin Theorem is a learning result, which states that any function $f : \mathbb{Z}_2^m \to \mathbb{Z}_2$ ($m \in \mathbb{N}$) which predicts the inner product of a secret vector $\mathbf{s} \in \mathbb{Z}_2^m$ and a uniformly random vector with non-zero advantage over guessing randomly admits an algorithm which recovers $\mathbf{s}$. The original application of the Goldreich-Levin Theorem in [GL89] was to symmetric key cryptography; the authors used their learning result to show that one-way functions imply pseudorandom generations, which, when combined with the work of [GGM84], concludes that one-way functions are sufficient for symmetric key encryption. An immediate consequence of the Goldreich-Levin Theorem in coding theory is that the Hadamard code is list-decodable. A line of work

([GKS06],[DGKS08],[BBW18]) generalized the Goldreich-Levin Theorem's "prediction implies inversion" techniques to work for general finite abelian groups.

We continue this line of work by investigating the Goldreich-Levin Theorem in high dimension when the prediction guarantee is low (that is, less than guessing randomly). Suppose $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ ($q, n \in \mathbb{N}$ such that q is prime) is a function and $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ is a matrix such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon$, for $0 < \varepsilon \leq 1/q$. Note that when $n = 1$, it can be shown that this problem is impossible, since a function which guesses randomly will agree with the inner product of the secret and a uniformly random vector with probability $1/q$. For general n, this probability becomes q^{-n}. So, we will restrict our attention to values of ε such that $q^{-n} < \varepsilon \leq q^{-1}$.

As a simple example, consider the function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ which, on input $\mathbf{x} \in \mathbb{Z}_q^m$, computes $\mathbf{y} = \mathbf{A}\mathbf{x} \in \mathbb{Z}_q^n$, chooses $\alpha \sim \mathbb{Z}_q$, and outputs $(y_1, \ldots, y_{n-1}, \alpha) \in \mathbb{Z}_q^n$. That is, f agrees with $\mathbf{A}$ in the first $n-1$ components, and outputs a random guess for the last component. It follows that f agrees with $\mathbf{A}$ with probability $1/q$, but also agrees with probability $1/q$ with any other matrix $\mathbf{A}^* \in \mathbb{Z}_q^{n\times m}$ which has the same first $n-1$ rows as $\mathbf{A}$. It follows that f agrees with at least q^m such matrices $\mathbf{A}^*$, and so we can't hope to exactly recover any particular one. However, we can hope to recover an *approximation* of $\mathbf{A}$ (*i.e.*, a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{A}'x = \mathbf{A}\mathbf{x}\big] \geq \varepsilon'$). Indeed, we prove the following theorem, stated informally below.

Theorem 5 (Informal) *Let $m, n, q \in \mathbb{N}$ such that q is prime. Suppose there exists a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ and a matrix $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ such that*

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon,$$

for $\varepsilon > 0$ *such that* $q^{-n} < \varepsilon \leq q^{-1}$*. Then, there exists an algorithm* $\mathcal{A}$*, running in time* $\mathsf{poly}(m, n, q, 1/\varepsilon)$*, which gets oracle access to* f*, and outputs with probability* $\mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$ *over its randomness a matrix* $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ *such that* $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[\mathbf{A}'\mathbf{x} = \mathbf{A}\mathbf{x}\right] \geq q^{-\mathcal{O}(1)}$*, unless* f *satisfies a conditional linearity test with high probability.*

Briefly, the proof of our theorem shows that either we can use the function f to compute an approximation of $\mathbf{A}$, or there exists a subset on which f conditionally passes the linearity test of Theorem 4 with high probability. In this informal theorem statement, for ease of exposition we have vaguely described the conditions under which f satisfies the conditional linearity test; see Theorem 4 for the formal theorem statement.

Chapter 2

Two Novel High Dimensional Linearity Tests

In this chapter, we prove two novel high dimensional linearity testing theorems. These high dimensional linearity tests are both in prime characteristic q and can be viewed as a "hybrid" between the high and low test acceptance regimes. To be clear, these *are not* low test acceptance linearity tests. However, a careful analysis of the hypothesis of each theorem reveals that both of these tests can be viewed through a certain lens as a "hybrid" between the high and low test acceptance regimes. We elaborate more in the sequel. But first, we establish some preliminaries.

2.1 Preliminaries

Basic Notation. If $n \in \mathbb{N}$, then we denote by $[n]$ the set $\{1, \ldots, n\}$. For a prime $q \in \mathbb{N}$, we denote by $\mathbb{Z}_q$ the field of integers modulo q. We will denote scalars, vectors and matrices with lowercase italic, lowercase bold, and uppercase bold respectively (*e.g.*, $z \in \mathbb{Z}_q$, $\mathbf{z} \in \mathbb{Z}_q^n$ and $\mathbf{Z} \in \mathbb{Z}_q^{n\times m}$). For a distribution $\mathcal{D}$ (resp. set D), we write $r \sim \mathcal{D}$ (resp. $r \sim D$) to indicate that the random variable r is drawn according to $\mathcal{D}$ (resp. the uniform distribution on D). For an event $\mathbf{E}$, we denote by $\mathbb{1}_{\mathbf{E}}$ the indicator random variable corresponding to $\mathbf{E}$. Namely, $\mathbb{1}_{\mathbf{E}} = 1$ (resp. $\mathbb{1}_{\mathbf{E}} = 0$) when $\mathbf{E}$ occurs (resp. does not occur). Unless otherwise specified, every vector space in this work is over the field $\mathbb{Z}_q$, for a prime $q \in \mathbb{N}$. If $n \in \mathbb{N}$, then a quantity $f(n)$ is said to be negligible if $f(n) \leq n^{-c}$ for every constant $c \in \mathbb{N}$. On the other hand, we say that $f(n)$ is non-negligible if there exists a constant $c \in \mathbb{N}$ such that $f(n) \geq n^{-c}$.

Roots of Unity. Let $q \in \mathbb{N}$ be prime. A q^{th} *root of unity in* $\mathbb{C}$ is a complex root of the monic degree q polynomial $f(X) = X^q - 1$. It is a well known fact that there exist exactly q roots of f, which are the elements $\{e^{-2\pi ik/q}\}_{k=0}^{q-1}$, where i is the imaginary unit. The q^{th} roots of unity in $\mathbb{C}$ form a multiplicative cyclic group of order q, which is isomorphic to $\mathbb{Z}_q$. A generator of the group is called a *primitive* q^{th} *root of unity in* $\mathbb{C}$. The subset of primitive q^{th} roots of unity in $\mathbb{C}$ form a subgroup of order $q-1$, which is isomorphic to $\mathbb{Z}_q^*$.

We'll prove two claims regarding q^{th} roots of unity which will be useful in this section. Our proof of the first claim utilizes the following well-known fact: If $\omega \in \mathbb{C}$ is a primitive q^{th} root of unity, then $\sum\limits_{k=0}^{q-1} \omega^k = 0$.

Claim 1 *Let $m, q \in \mathbb{N}$ such that q is prime, and $\omega \in \mathbb{C}$ be a primitive q^{th} root of unity. Then, $\forall \mathbf{x} \in \mathbb{Z}_q^m$,*

$$\mathbb{E}_{\mathbf{y}\sim\mathbb{Z}_q^m}\left[\omega^{\langle \mathbf{x},\mathbf{y}\rangle}\right] = \mathbb{1}_{\mathbf{x}=\mathbf{0}}.$$

Proof. *Let* $\mathbf{x} \in \mathbb{Z}_q^m$. *If* $\mathbf{x} = \mathbf{0}$, *then* $\mathbb{E}_{\mathbf{y} \sim \mathbb{Z}_q^m}\left[\omega^{\langle \mathbf{x}, \mathbf{y} \rangle}\right] = 1$. *Now, assume* $\mathbf{x} \neq \mathbf{0}$. *Then,*

$$\mathbb{E}_{\mathbf{y} \sim \mathbb{Z}_q^m}\left[\omega^{\langle \mathbf{x}, \mathbf{y} \rangle}\right] = q^{-m} \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \omega^{\langle \mathbf{x}, \mathbf{y} \rangle} = q^{-m} \sum_{\beta \in \mathbb{Z}_q} \sum_{\substack{\mathbf{y} \in \mathbb{Z}_q^m \ s.t. \\ \langle \mathbf{x}, \mathbf{y} \rangle = \beta}} \omega^{\beta}$$

$$= \sum_{\beta \in \mathbb{Z}_q} \omega^{\beta} \cdot \Pr_{\mathbf{y} \sim \mathbb{Z}_q^m}\left[\langle \mathbf{x}, \mathbf{y} \rangle = \beta\right] = q^{-1} \sum_{\beta \in \mathbb{Z}_q} \omega^{\beta} = 0.$$

■

Claim 2 *Let* $m, q \in \mathbb{N}$ *such that* q *is prime, and* $\omega \in \mathbb{C}$ *be a primitive* q^{th} *root of unity. Let* $\{\alpha_k\}_{k \in \mathbb{Z}_q} \subset [0,1]$ *such that* $\sum_{k \in \mathbb{Z}_q} \alpha_k = 1$, *and let* $z = \sum_{k \in \mathbb{Z}_q} \alpha_k \omega^k \in \mathbb{C}$. *If* $|z| \geq 1/q + \gamma$, *for some* $\gamma > 0$, *then* $\exists k \in \mathbb{Z}_q$ *such that* $\alpha_k \geq 1/q + (1/q^3 + \gamma/q^2)$.

Proof. Let $\eta = 1/q^3 + \gamma/q^2$, and suppose that $\forall k \in \mathbb{Z}_q$, $\alpha_k < 1/q + \eta$. We have that each $\alpha_k = 1 - \sum_{\ell \in \mathbb{Z}_q \setminus \{k\}} \alpha_\ell > 1 - (q-1)(1/q + \eta) = 1/q - (q-1)\eta$. So, $1/q - \eta(q-1) \leq \alpha_k \leq 1/q + \eta \leq 1/q + \eta(q-1)$, which implies that $|\alpha_k - 1/q| \leq \eta(q-1)$. Hence

$$|z| = \left|(1/q) \sum_{k \in \mathbb{Z}_q} \omega^k + \sum_{k \in \mathbb{Z}_q} (\alpha_k - 1/q)\omega^k\right| \leq \sum_{k \in \mathbb{Z}_q} |\alpha_k - 1/q| \leq q(q-1)\eta < q^2\eta = 1/q + \gamma,$$

a contradiction. For the first inequality above, we have used ■

Discrete Fourier Analysis on Finite Abelian Groups. Here we briefly summarize several key results in discrete Fourier analysis on finite abelian groups. Actually, we specialize our results to the group $\mathbb{Z}_q^m$ ($m, q \in \mathbb{N}$ such that q is prime), which is also a vector space over $\mathbb{Z}_q$ of dimension n. See [Luo09] for a complete treatment on general finite abelian groups. Let $\mathcal{F} = \{f : \mathbb{Z}_q^m \to \mathbb{C}\}$ be the set of all functions $f : \mathbb{Z}_q^m \to \mathbb{C}$, which is a $\mathbb{C}-$vector space. A *character* of $\mathbb{Z}_q^m$ is a group homomorphism $\chi : \mathbb{Z}_q^m \to \mathbb{C}^*$. It turns out that each character χ of $\mathbb{Z}_q^m$ is characterized by an element $\mathbf{u} \in \mathbb{Z}_q^m$, and is defined by $\chi_{\mathbf{u}}(\mathbf{x}) = \omega^{\langle \mathbf{u},\mathbf{x}\rangle}$, where $\omega \in \mathbb{C}$ is a primitive q^{th} root of unity in $\mathbb{C}$. Moreover, the set $\{\chi_{\mathbf{u}} : \mathbf{u} \in \mathbb{Z}_q^m\} \subset \mathcal{F}$ of characters of $\mathbb{Z}_q^m$ form a $\mathbb{C}-$basis for $\mathcal{F}$, hence $\mathcal{F}$ is a $\mathbb{C}-$vector space of finite dimension q^m. If $f \in \mathcal{F}$ and $\mathbf{u} \in \mathbb{Z}_q^m$, then we define the *Fourier coefficient of f at* $\mathbf{u}$ as $\hat{f}(\mathbf{u}) := \mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[f(\mathbf{x})\bar{\chi}_{\mathbf{u}}(\mathbf{x})\right] = \mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[f(\mathbf{x})\omega^{-\langle \mathbf{u},\mathbf{x}\rangle}\right]$. It turns out that for every $f \in \mathcal{F}$, f can be expressed as a unique linear combination of the characters of $\mathbb{Z}_q^m$ over $\mathbb{C}$ of the form $f = \sum\limits_{\mathbf{u}\in\mathbb{Z}_q^m} \hat{f}(\mathbf{u})\chi_{\mathbf{u}}$. Parseval's identity, an extremely useful tool in application, relates the expectation of the norm of $|f(\cdot)|^2$ to the sum over its Fourier coefficients of their squared norm:

$$\mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[\left|f(\mathbf{x})\right|^2\right] = \sum_{\mathbf{u}\in\mathbb{Z}_q^m} \left|\hat{f}(\mathbf{u})\right|^2.$$

2.2 A High Dimensional Linearity Test

In this section, we prove a novel high dimensional linearity test. Our linearity testing theorem operates on a function $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ ($n, q \in \mathbb{N}$ such that $q = \mathsf{poly}(n)$ is prime) which has a non-degeneracy property, defined below.

Definition 6 (Non-Degeneracy) *Let $n, q \in \mathbb{N}$ such that $q = \mathsf{poly}(n)$ is prime, and $\zeta, \rho > 0$ such that $\zeta > \rho$. A function $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ is said to be $(\zeta,\rho)-$degenerate if there exists a subset $T \subset \mathbb{Z}_q^n$ of density $|T|q^{-n} = \rho$ such that $\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\left[h(\mathbf{a}) \in T\right] \geq \zeta$. If h is not $(\zeta,\rho)-$degenerate then we say that h is $(\zeta,\rho)-$non-degenerate. If h is $(\zeta,\rho)-$non-degenerate for every non-negligible $\zeta = \zeta(n)$ and negligible $\rho = \rho(n)$, then we say that h is non-degenerate.*

Remark. Suppose $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ is a non-degenerate function, and let $\mathbf{V} \subset \mathbb{Z}_q^n$ be a subspace of constant dimension $d = \mathcal{O}(1)$. Then, $|\mathbf{V}|q^{-n} = q^{-(n-d)}$, which is negligible in n, and since h is non-degenerate, it follows that $\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathbf{V}\big] = \mathsf{negl}(n)$. We will use this fact heavily throughout the proofs in this section.

We now state our high dimensional linearity testing theorem.

Theorem 7 (Theorem 1 (Formal)) *Let $n, q \in \mathbb{N}$ such that $q = \mathsf{poly}(n)$ is prime. If $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ is a non-degenerate function such that $\forall(\alpha_1, \alpha_2) \in \mathbb{Z}_q^2$,*

$$\Pr_{\mathbf{a}_1,\mathbf{a}_2\sim\mathbb{Z}_q^n}\big[h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2)) \in \mathrm{Span}(h(\mathbf{a}_1), h(\mathbf{a}_2))\big] \geq 1 - \varepsilon,$$

for $\varepsilon = \varepsilon(n) > 0$ non-negligible, then $\exists\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ such that

$$\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathrm{Span}(\mathbf{H}\mathbf{a})\big] \geq 1 - \mathcal{O}(n^2 q\sqrt{q\varepsilon}).$$

We remark that our result is not a traditional linearity test. However, by averaging our conclusion, it follows that $\exists\mathbf{H}' \in \mathbb{Z}_q^{n\times n}$ such that $\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[h(\mathbf{a}) = \mathbf{H}'\mathbf{a}\big] \geq (1-\mathcal{O}(nq^2\varepsilon))/q$. Although this appears to give us a high test acceptance and low agreement linearity test, a closer analysis of the hypothesis of our theorem yields that if we again average, $\exists(\alpha_1, \alpha_2) \in \mathbb{Z}_q^2$ such that our hypothesis becomes:

$$\Pr_{\mathbf{a}_1,\mathbf{a}_2\sim\mathbb{Z}_q^n}\big[h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2) = \alpha_1 h(\mathbf{a}_1) + \alpha_2 h(\mathbf{a}_2)\big] \geq (1 - \varepsilon)/q^2,$$

which resembles more traditional linearity tests. Note in this case we obtain a low test acceptance. Even though this low test acceptance condition is not sufficient for our theorem, it provides an interesting data point to support our belief that our linearity test can be thought of as a "hybrid" between the high and low test acceptance regimes. With this in mind, we now prove Theorem 7.

Proof of Theorem 7. We prove the theorem through a sequence of claims.

Claim 3 *It holds with probability* $1 - 4nq\varepsilon$ *over* $\mathbf{a}_1, \dots, \mathbf{a}_n \sim \mathbb{Z}_q^n$ *that* $\exists\{\mathbf{a}'_1, \dots, \mathbf{a}'_n\} \subset \mathbb{Z}_q^n$ *such that:*

1. $h(\alpha\mathbf{a}_i) \in \mathrm{Span}(\alpha\mathbf{a}'_i), \forall\alpha \in \mathbb{Z}_q^*$.

2. $h(\mathbf{a}_1 + \mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1 + \mathbf{a}'_i), \forall i \geq 2$.

Proof. Let $\mathbf{a}'_1 = h(\mathbf{a}_1) \in \mathbb{Z}_q^n$, and fix $i \in \{2, \dots, n\}$. By hypothesis, except with probability ε, $\mathbf{y}_i := h(\mathbf{a}_1 + \mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1, h(\mathbf{a}_i))$, and so $\exists(\beta_1, \beta_2) \in \mathbb{Z}_q^2$ such that $\mathbf{y}_i = \beta_1\mathbf{a}'_1 + \beta_2 h(\mathbf{a}_i)$. Since h is non-degenerate, $\beta_1 \neq 0$ except with negligible probability. So, $\mathbf{y}_i = \beta_1(\mathbf{a}'_1 + \beta_1^{-1}\beta_2 h(\mathbf{a}_i))$. Let $\lambda_i := \beta_1^{-1}\beta_2 \in \mathbb{Z}_q$, and note that λ_i must be unique, except with negligible probability, since h is non-degenerate. Let $\mathbf{a}'_i := \lambda_i h(\mathbf{a}_i) \in \mathbb{Z}_q^n$, and note that $\mathbf{y}_i \in \mathrm{Span}(\mathbf{a}'_1 + \mathbf{a}'_i)$ with probability $1 - (\varepsilon + \mathsf{negl}(n)) \geq 1 - 2\varepsilon$. By the union bound, the second point holds with probability $1 - 2n\varepsilon$.

To prove the first point, let $i, j \in [n]$ be distinct, and $\alpha \in \mathbb{Z}_q^*$. We have $\mathbf{z} := h(\alpha\mathbf{a}_i) = h(\alpha\mathbf{a}_i + 0\mathbf{a}_j) \in \mathrm{Span}(\mathbf{a}'_i, \mathbf{a}'_j)$, except with probability ε. So, $\exists(\gamma_1, \gamma_2) \in \mathbb{Z}_q^2$ such that $\mathbf{z} = \gamma_1\mathbf{a}'_i + \gamma_2\mathbf{a}'_j$. Since h is non-degenerate, then $\gamma_2 = 0$, except with negligible probability. Hence $\mathbf{z} \in \mathrm{Span}(\mathbf{a}'_i) = \mathrm{Span}(\alpha\mathbf{a}'_i)$, except with probability $\varepsilon + \mathsf{negl}(n) = 2\varepsilon$. By the union bound, the first point holds with probability $1 - 2nq\varepsilon$.

In total, we've shown that with probability $1 - 4nq\varepsilon$, $\exists\{\mathbf{a}'_1, \dots, \mathbf{a}'_n\} \subset \mathbb{Z}_q^n$ such that conditions (1) and (2) in the statement of the claim hold. ■

If $\mathbf{a}_1, \dots, \mathbf{a}_n \sim \mathbb{Z}_q^n$, then Claim 3 guarantees with high probability the existence of a set $\{\mathbf{a}'_1, \dots, \mathbf{a}'_n\} \subset \mathbb{Z}_q^n$ such that conditions (1) and (2) of the claim hold; we say such a set $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\}$ which satisfies conditions (1) and (2) is *good* for $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$. If $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n$, then let $\mathcal{D}(\{\mathbf{a}_i\}_i)$ denote the distribution which computes a good subset $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\}$ with probability $1 - 4nq\varepsilon$. Throughout the remainder of this proof, unless otherwise stated, all probabilities are over $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n$, $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\} \sim \mathcal{D}(\{\mathbf{a}_i\}_{i=1}^n)$.

Claim 4 *It holds with probability $1 - 6n^2q^2\varepsilon$ that $\forall i \in \{2, \ldots, n\}$, there exists a map $\pi_i : \mathbb{Z}_q \to \mathbb{Z}_q$ defined by $\pi_i(\alpha) = \beta$ such that $h(\mathbf{a}_1 + \alpha\mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1 + \beta\mathbf{a}'_i)$.*

Proof. Let $i \in \{2, \ldots, n\}$ and $\alpha \in \mathbb{Z}_q$. Except with probability $4nq\varepsilon + \varepsilon \leq 5nq\varepsilon$, we have $\mathbf{y} := h(\mathbf{a}_1 + \alpha\mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1, \mathbf{a}'_i)$, and so $\exists(\beta_1, \beta_2) \in \mathbb{Z}_q^2$ such that $\mathbf{y} = \beta_1\mathbf{a}'_1 + \beta_2\mathbf{a}'_2 = \beta_1(\mathbf{a}'_1 + \beta_1^{-1}\beta_2\mathbf{a}'_2)$ ($\beta_1 \neq 0$ except with negligible probability since h is non-degenerate). We define $\pi_i(\alpha) = \beta_1^{-1}\beta_2 \in \mathbb{Z}_q$ such that $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha)\mathbf{a}'_i)$ (note that $\pi_i(\alpha)$ is indeed well defined, except with negligible probability, since h is non-degenerate). Our total probability loss is $5nq\varepsilon + \mathsf{negl}(n) \leq 6nq\varepsilon$, and so the conclusion follows by the union bound over all $\alpha \in \mathbb{Z}_q$ and $i \in \{2, \ldots, n\}$. ■

Claim 5 *For all distinct $i, j \in \{2, \ldots, n\}$ and for all $(\alpha_1, \alpha_i, \alpha_j) \in \{0, 1\} \times \mathbb{Z}_q \times \mathbb{Z}_q$ such that $(\alpha_i, \alpha_j) \neq (0, 0)$, it holds with probability $1 - 30n^2q^2\varepsilon$ that $h(\alpha_1\mathbf{a}_1 + \alpha_2\mathbf{a}_2 + \alpha_3\mathbf{a}_3) \in \mathrm{Span}(\alpha_1\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_i(\alpha_j)\mathbf{a}'_j)$.*

Proof. Fix distinct $i, j \in \{2, \ldots, n\}$, and $(\alpha_i, \alpha_j) \in \mathbb{Z}_q \times \mathbb{Z}_q$. First, let $\alpha_1 = 1$. Let $\mathbf{y} = h(\mathbf{a}_1 + \alpha_i\mathbf{a}_i + \alpha_j\mathbf{a}_j)$. If either $\alpha_i = 0$ or $\alpha_j = 0$, then WLOG write $\mathbf{y} = h(\mathbf{a}_1 + \alpha_i\mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\mathbf{a}_i)\mathbf{a}'_i) = \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\mathbf{a}_i)\mathbf{a}'_i + \pi_j(0)\mathbf{a}'_j)$, except with probability $6n^2q^2\varepsilon$ by Claim 4. So, assume $\alpha_i \neq 0$ and $\alpha_j \neq 0$. We have $\mathbf{y} = h((\mathbf{a}_1 + \alpha_i\mathbf{a}_i) + \alpha_j\mathbf{a}_j) \in \mathrm{Span}(h(\mathbf{a}_1 + \alpha_i\mathbf{a}_i), h(\alpha_j\mathbf{a}_j))$, except with probability ε, by hypothesis. We have that $h(\mathbf{a}_1 +$

$\alpha_i \mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i)$, except with probability $6n^2q^2\varepsilon$ by Claim 4. Also, $h(\alpha_j \mathbf{a}_j) \in \mathrm{Span}(\alpha_j h(\mathbf{a}_j)) = \mathrm{Span}(\mathbf{a}'_j)$, except with probability $4nq\varepsilon$ by Claim 3. Thus $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i, \mathbf{a}'_j)$, except with probability $\varepsilon + 6n^2q^2\varepsilon + 4nq\varepsilon \leq 11n^2q^2\varepsilon$. Similarly, we also have that $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_1 + \pi_j(\alpha_j)\mathbf{a}'_j, \mathbf{a}'_i)$, except with probability $11n^2q^2\varepsilon$. So, $\exists(\beta_1, \beta_2), (\gamma_1, \gamma_2) \in \mathbb{Z}_q^2$ such that

$$\beta_1(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i) + \beta_2 \mathbf{a}'_j = \mathbf{y} = \gamma_1(\mathbf{a}'_1 + \pi_j(\alpha_j)\mathbf{a}'_j) + \gamma_2 \mathbf{a}'_i,$$

and we must have $\beta_1 = \gamma_1, \beta_2 = \gamma_1 \pi_j(\alpha_j) = \beta_1 \pi_j(\alpha_j)$, except with negligible probability, since h is non-degenerate. It thus follows that $\mathbf{y} = \beta_1(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j)$, except with probability $22n^2q^2\varepsilon + \mathsf{negl}(n) \leq 23n^2q^2\varepsilon$.

Now, let $\alpha_1 = 0$. If either $\alpha_i = 0$ or $\alpha_j = 0$ then WLOG write $\mathbf{y} = h(\alpha_i \mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_i)$, except with probability $4nq\varepsilon$ by Claim 3. Also, $\mathbf{y} = h(\alpha_i \mathbf{a}_i) = h((\mathbf{a}_1 + \alpha_i \mathbf{a}_i) - \mathbf{a}_1) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i, \mathbf{a}'_1)$, except with probability $\varepsilon + 6n^2q^2\varepsilon \leq 7n^2q^2\varepsilon$ by the hypothesis of the theorem and Claim 4. Hence $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i, \mathbf{a}'_1) \bigcap \mathrm{Span}(\mathbf{a}'_i) \subset \mathrm{Span}(\pi_i(\alpha_i)\mathbf{a}'_i)$, except with negligible probability, since h is non-degenerate. In total, $\mathbf{y} \in \mathrm{Span}(\pi_i(\alpha_i)\mathbf{a}'_i)$, except with probability $4nq\varepsilon + 7n^2q^2\varepsilon + \mathsf{negl}(n) \leq 12n^2q^2\varepsilon$.

We may now assume that $\alpha_i \neq 0$ and $\alpha_j \neq 0$. We have $\mathbf{y} = h(\alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) = h((\mathbf{a}_1 + \alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) - \mathbf{a}_1) \in \mathrm{Span}(h(\mathbf{a}_1 + \alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j), \mathbf{a}'_1)$, except with probability ε. Also, by the above argument, $h(\mathbf{a}_1 + \alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j)$, except with probability $23n^2q^2\varepsilon$. So, $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j, \mathbf{a}'_1)$, except with probability $23n^2q^2\varepsilon + \varepsilon \leq 24n^2q^2\varepsilon$. We also have that $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_i, \mathbf{a}'_j)$, except with probability $\varepsilon + 4nq\varepsilon \leq 5nq\varepsilon$ by the hypothesis of the theorem and Claim 3. Hence, except with probability $24n^2q^2\varepsilon + 5nq\varepsilon \leq 29n^2q^2\varepsilon$, $\mathbf{y} \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j, \mathbf{a}'_1) \bigcap \mathrm{Span}(\mathbf{a}'_i, \mathbf{a}'_j) \subset$

$\mathrm{Span}(\pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j)$, except with negligible probability since h is non-degenerate. In total, we have that $\mathbf{y} \in \mathrm{Span}(\pi_i(\alpha_i)\mathbf{a}'_i + \pi_j(\alpha_j)\mathbf{a}'_j)$, except with probability $29n^2q^2\varepsilon + \mathsf{negl}(n) \leq 30n^2q^2\varepsilon$. ■

Claim 6 *It holds with probability* $1 - 48n^3q^3\varepsilon$ *that* $\forall i \in \{2, \ldots, n\}$, $\pi_i(\alpha) = \alpha, \forall \alpha \in \mathbb{Z}_q$.

Proof. Let $i \in \{2, \ldots, n\}$. If $\alpha = 0$ then $h(\mathbf{a}_1 + \alpha\mathbf{a}_i) = h(\mathbf{a}_1) \in \mathrm{Span}(\mathbf{a}'_1) = \mathrm{Span}(\mathbf{a}'_1 + \alpha\mathbf{a}'_i)$, except with probability $4nq\varepsilon \leq 48n^2q^2\varepsilon$ by Claim 3. We now proceed by induction to show that $\pi_i(\alpha) = \alpha$ with probability at least $1 - \alpha(48n^2q^2\varepsilon), \forall \alpha \in \mathbb{Z}_q^*$. If $\alpha_1 = 1$, then $h(\mathbf{a}_1 + \alpha\mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha)\mathbf{a}'_i)$, except with probability $6n^2q^2\varepsilon \leq 48n^2q^2\varepsilon$ by Claim 4.

For the induction step, assume that $\pi_i(\alpha - 1) = \alpha - 1$ with probability $1-(\alpha-1)(48n^2q^2\varepsilon)$. Let $\mathbf{y} := h(\mathbf{a}_1+\alpha\mathbf{a}_i+\mathbf{a}_j) = h((\mathbf{a}_1+(\alpha-1)\mathbf{a}_i)+(\mathbf{a}_i+\mathbf{a}_j)) \in \mathrm{Span}(\mathbf{a}'_1+(\alpha-1)\mathbf{a}'_i, \mathbf{a}'_i + \mathbf{a}'_j)$. Here we have used the hypothesis of the theorem, the induction hypothesis, and Claim 5, which incurs an additional probability loss of $\varepsilon + (\alpha-1)(48n^2q^2\varepsilon) + 30n^2q^2\varepsilon \leq (\alpha - 1)(48n^2q^2\varepsilon) + 31n^2q^2\varepsilon$. Also, we can write $\mathbf{y} = h((\mathbf{a}_1 + \alpha\mathbf{a}_i) + \mathbf{a}_j) \in \mathrm{Span}(\mathbf{a}'_1 + \pi_i(\alpha)\mathbf{a}'_i, \mathbf{a}'_j)$, except with probability $\varepsilon + 6n^2q^2\varepsilon + 4nq\varepsilon \leq 11n^2q^2\varepsilon$ by the hypothesis of the theorem, Claim 4, and Claim 3. Finally, we have $\mathbf{y} = h((\mathbf{a}_1 + \mathbf{a}_j) + \alpha\mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_1 + \mathbf{a}'_j, \mathbf{a}'_i)$, except with probability $\varepsilon + 4nq\varepsilon \leq 5nq\varepsilon$ by hypothesis and Claim 3. In total, except with probability $(\alpha - 1)(48n^2q^2\varepsilon) + 31n^2q^2\varepsilon + 11n^2q^2\varepsilon + 5nq\varepsilon \leq (\alpha - 1)(48n^2q^2\varepsilon) + 47n^2q^2\varepsilon$, we can express $\mathbf{y}$ as:

1. $\mathbf{y} = \beta_1(\mathbf{a}'_1 + (\alpha - 1)\mathbf{a}'_i) + \beta_2(\mathbf{a}'_i + \mathbf{a}'_j)$ $(\beta_1, \beta_2 \in \mathbb{Z}_q)$.

2. $\mathbf{y} = \gamma_1(\mathbf{a}'_1 + \pi_i(\alpha)\mathbf{a}'_i) + \gamma_2(\mathbf{a}'_j)$ $(\gamma_1, \gamma_2 \in \mathbb{Z}_q)$.

3. $\mathbf{y} = \delta_1(\mathbf{a}'_1 + \mathbf{a}'_j) + \delta_2\mathbf{a}'_i$ $(\delta_1, \delta_2 \in \mathbb{Z}_q)$.

By rearranging, and since h is non-degenerate, it holds with overwhelming probability that $\beta_1 = \gamma_1, \gamma_1(\alpha-1)+\gamma_2 = \beta_1\pi_i(\alpha)$, and $\gamma_2 = \delta_1 = \beta_1$. Hence $\beta_1\pi_i(\alpha) = \beta_1(\alpha-1)+\beta_1 = \beta_1\alpha$. Since h is non-degenerate, then $\beta_1 \neq 0$ except with negligible probability, hence $\pi_i(\alpha) = \alpha$. Our total probability loss is $(\alpha-1)(48n^2q^2\varepsilon) + 47n^2q^2\varepsilon + \mathsf{negl}(n) \leq (\alpha-1)(48n^2q^2\varepsilon) + 48n^2q^2\varepsilon = \alpha(48n^2q^2\varepsilon)$, which completes the induction step. The conclusion follows by induction and by the union bound. ■

Claim 7 *For all distinct $i, j \in \{2,\dots,n\}$ and $(\alpha_1,\alpha_i,\alpha_j) \in \mathbb{Z}_q^3$ such that $(\alpha_i,\alpha_j) \neq (0,0)$, it holds with probability $1-82n^3q^3\varepsilon$ that $h(\alpha_1\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) \in \mathrm{Span}(\alpha_1\mathbf{a}'_1+\alpha_i\mathbf{a}'_i+\alpha_j\mathbf{a}'_j)$.*

Proof. Fix distinct $i, j \in \{2,\dots,n\}$ and $(\alpha_1,\alpha_i,\alpha_j) \in \mathbb{Z}_q^3$ such that $(\alpha_i,\alpha_j) \neq (0,0)$. If $\alpha_1 = 0$, then the conclusion follows by Claims 5 and 6 with probability $1-(30n^2q^2\varepsilon+48n^3q^3\varepsilon) \geq 1-78n^3q^3\varepsilon$. So, assume $\alpha_1 \neq 0$. Then, $\mathbf{y} := h(\alpha_1\mathbf{a}_1+\alpha_i\mathbf{a}_i+\alpha_j\mathbf{a}_j) = h(\alpha_1(\mathbf{a}_1+\alpha_1^{-1}\alpha_i\mathbf{a}_i+\alpha_1^{-1}\alpha_j\mathbf{a}_j)) \in \mathrm{Span}(\alpha_1 h(\mathbf{a}_1+\alpha_1^{-1}\alpha_i\mathbf{a}_i+\alpha_1^{-1}\alpha_j\mathbf{a}_j))$, except with probability $4nq\varepsilon$ by Claim 3. Also, $h(\mathbf{a}_1+\alpha_1^{-1}\alpha_i\mathbf{a}_i+\alpha_1^{-1}\alpha_j\mathbf{a}_j) \in \mathrm{Span}(\mathbf{a}'_1+\alpha_1^{-1}\alpha_i\mathbf{a}'_i+\alpha_1^{-1}\alpha_j\mathbf{a}'_j)$, except with probability $30n^2q^2\varepsilon+48n^3q^3\varepsilon \leq 78n^3q^3\varepsilon$ by Claims 5 and 6. Thus $\mathbf{y} = \mathrm{Span}(\alpha_1(\mathbf{a}'_1+\alpha_1^{-1}\alpha_i\mathbf{a}'_i+\alpha_1^{-1}\alpha_j\mathbf{a}'_j) = \mathrm{Span}(\alpha_1\mathbf{a}'_1+\alpha_i\mathbf{a}'_i+\alpha_j\mathbf{a}'_j)$, except with probability $4nq\varepsilon+78n^3q^3\varepsilon \leq 82n^3q^3\varepsilon$.

■

Claim 8

$$\Pr_{\substack{\mathbf{a}_1,\dots,\mathbf{a}_n\sim\mathbb{Z}_q^n\\ \{\mathbf{a}'_1,\dots,\mathbf{a}'_n\}\sim\mathcal{D}(\{\mathbf{a}_i\}_{i=1}^n)\\ (\alpha_1,\dots,\alpha_n)\sim\mathbb{Z}_q^n\backslash\{\mathbf{0}\}}}\left[h\Big(\sum_{i=1}^n \alpha_i\mathbf{a}_i\Big) \in \mathrm{Span}\Big(\sum_{i=1}^n \alpha_i\mathbf{a}'_i\Big)\right] \geq 1-89n^4q^3\varepsilon.$$

Proof. We proceed by induction on $r \in \{3,\dots,n\}$. Assume all probabilities are over randomness in the hypothesis of the claim. Since $(\alpha_1,\dots,\alpha_r) \neq \mathbf{0}$ then assume WLOG that $\alpha_1 \neq 0$.

The base case in which $r = 3$ follows from Claim 7. For the induction step, assume that $h\Big(\sum_{i=1}^{r-1} \alpha_i \mathbf{a}_i\Big) \in \mathrm{Span}\Big(\sum_{i=1}^{r-1} \alpha_i \mathbf{a}'_i\Big)$ with probability $1-(r-1)(89n^3q^3\varepsilon)$. Let $\mathbf{y} = h(\sum_{i=1}^{r} \alpha_i \mathbf{a}_i)$. Then, $\mathbf{y} = h((\alpha_1 \mathbf{a}_1 + \ldots, \alpha_{r-1}\mathbf{a}_{r-1}) + \alpha_r \mathbf{a}_r) \in \mathrm{Span}(\alpha_1 \mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}, \mathbf{a}'_r)$, except with probability $(r-1)(89n^3q^3\varepsilon) + \varepsilon + 4nq\varepsilon \leq (r-1)(89n^3q^3\varepsilon) + 5nq\varepsilon$. Here we have used the hypothesis of the theorem, the induction hypothesis, and Claim 3. Also, we have $\mathbf{y} = h((\alpha_1\mathbf{a}_1 + \alpha_r\mathbf{a}_r) + (\alpha_2\mathbf{a}_2 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1})) \in \mathrm{Span}(\alpha_1\mathbf{a}'_1 + \alpha_r\mathbf{a}'_r, \mathbf{z})$ $\big(\mathbf{z} = h(\alpha_2\mathbf{a}_2 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1})\big)$, except with probability $\varepsilon + 82n^3q^3\varepsilon$. Here we have used the hypothesis of the theorem and Claim 7. In total, except with probability $(r-1)(89n^3q^3\varepsilon) + 5nq\varepsilon + \varepsilon + 82n^3q^3\varepsilon \leq (r-1)(89n^3q^3\varepsilon) + 88n^3q^3\varepsilon$, we can write

$$\beta_1(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}) + \beta_2\mathbf{a}'_r = \mathbf{y} = \gamma_1(\alpha_1\mathbf{a}'_1 + \alpha_r\mathbf{a}'_r) + \gamma_2\mathbf{z},$$

for $\beta_1, \beta_2, \gamma_1, \gamma_2 \in \mathbb{Z}_q$. Note that the randomness of $\mathbf{z}$ is independent of $\mathbf{a}_1$ and $\mathbf{a}_r$. Then, since h is non-degenerate, except with negligible probability we must have $\beta_1\alpha_1 = \gamma_1\alpha_1$ and $\beta_2 = \gamma_1\alpha_r$. Since $\alpha_1 \neq 0$ then $\beta_1 = \gamma_1$ and $\beta_2 = \beta_1\alpha_r$. Hence $\mathbf{y} = \beta_1(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1} + \alpha_r\mathbf{a}'_r) \in \mathrm{Span}(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1} + \alpha_r\mathbf{a}'_r)$. Our total probability loss is $(r-1)(89n^3q^3\varepsilon) + 88n^3q^3\varepsilon + \mathsf{negl}(n) \leq (r-1)(89n^3q^3\varepsilon) + 89n^3q^3\varepsilon = r(89n^3q^3\varepsilon)$. The claim follows by induction. ■

Now, by averaging Claim 8, we have that with probability $1 - \sqrt{89n^4q^3\varepsilon}$ over $\mathbf{a}_1, \ldots, \mathbf{a}_n \sim \mathbb{Z}_q^n, \{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\} \sim \mathcal{D}(\{\mathbf{a}_i\}_i)$, it holds that

$$\mathrm{P}(\{\mathbf{a}_i\}_i, \{\mathbf{a}'_i\}_i) := \Pr_{(\alpha_1,\ldots,\alpha_n)\sim\mathbb{Z}_q^n}\left[h\Big(\sum_{i=1}^{n}\alpha_i\mathbf{a}_i\Big) \in \mathrm{Span}\Big(\sum_{i=1}^{n}\alpha_i\mathbf{a}'_i\Big)\right] \geq 1 - \sqrt{89n^4q^3\varepsilon} - q^{-n}.$$

Hence $\Pr_{\mathbf{a}_1,\ldots,\mathbf{a}_n\sim\mathbb{Z}_q^n}\Big[\{\mathbf{a}_i\}_i \text{ L.I. } \& \ \exists\{\mathbf{a}'_i\}_i \text{ s.t. } \mathrm{P}(\{\mathbf{a}_i\}_i, \{\mathbf{a}'_i\}_i) \geq 1 - (\sqrt{89n^4q^3\varepsilon} + q^{-n})\Big] \geq 1 - (n/q + \sqrt{89n^4q^3\varepsilon})$. So, if $\{\mathbf{a}_1, \ldots, \mathbf{a}_n\} \subset \mathbb{Z}_q^n$ is a basis for $\mathbb{Z}_q^n$, $\{\mathbf{a}'_1, \ldots, \mathbf{a}'_n\}$ is good, and

$\mathrm{P}(\{\mathbf{a}_i\}_i, \{\mathbf{a}'_i\}_i) \geq 1-(\sqrt{89n^4q^3\varepsilon}+q^{-n})$, then letting $(\mathbf{A}, \mathbf{A}') \in \mathbb{Z}_q^{n\times n} \times \mathbb{Z}_q^{n\times n}$ be the matrices whose i^{th} rows are $\mathbf{a}_i, \mathbf{a}'_i$, respectively, and $\mathbf{H} = \mathbf{A}'\mathbf{A}^{-1} \in \mathbb{Z}_q^{n\times n}$, it follows that

$$\begin{aligned}
&\Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathrm{Span}(\mathbf{Ha})\big] \\
&= \Pr_{\boldsymbol{\alpha}\sim\mathrm{Span}(\mathbf{a}_1,\ldots,\mathbf{a}_n)}\Bigg[h\Big(\sum_{i=1}^n \alpha_i\mathbf{a}_i\Big) \in \mathrm{Span}\Bigg(\mathbf{H}\Big(\sum_{i=1}^n \alpha_i\mathbf{a}_i\Big)\Bigg)\Bigg] \\
&= \Pr_{\boldsymbol{\alpha}\sim\mathrm{Span}(\mathbf{a}_1,\ldots,\mathbf{a}_n)}\Bigg[h\Big(\sum_{i=1}^n \alpha_i\mathbf{a}_i\Big) \in \mathrm{Span}\big(\mathbf{HA}\boldsymbol{\alpha}\big)\Bigg] \\
&= \Pr_{\boldsymbol{\alpha}\sim\mathrm{Span}(\mathbf{a}_1,\ldots,\mathbf{a}_n)}\Bigg[h\Big(\sum_{i=1}^n \alpha_i\mathbf{a}_i\Big) \in \mathrm{Span}\big(\mathbf{A}'\boldsymbol{\alpha}\big)\Bigg] \\
&= \Pr_{\boldsymbol{\alpha}\sim\mathrm{Span}(\mathbf{a}_1,\ldots,\mathbf{a}_n)}\Bigg[h\Big(\sum_{i=1}^n \alpha_i\mathbf{a}_i\Big) \in \mathrm{Span}\Big(\sum_{i=1}^n \alpha_i\mathbf{a}'_i\Big)\Bigg] \\
&\geq 1-(\sqrt{89n^4q^3\varepsilon}+q^{-n}) \geq 1-\mathcal{O}(n^2q\sqrt{q\varepsilon}).
\end{aligned}$$

Hence $\Pr_{\mathbf{a}_1,\ldots,\mathbf{a}_n\sim\mathbb{Z}_q^n}\Big[\exists\mathbf{H} \text{ s.t. } \Pr_{\mathbf{a}\sim\mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathrm{Span}(\mathbf{Ha})\big] \geq 1-\mathcal{O}(n^2q\sqrt{q\varepsilon})\Big] \geq 1-(n/q+\mathcal{O}(n^2q\sqrt{q\varepsilon}))$. ■

2.3 A High Dimensional Conditional Affine Linearity Test

In this section, we prove the following theorem.

Theorem 8 *Let $m, d, q \in \mathbb{N}$ such that q is prime. Let $T \subset \mathbb{Z}_q^m$ be a subset of density $\lambda := |T|q^{-m} \leq 1/\sqrt{2}$ such that for every subset $T' \subset T$ such that $|T'|/|T| \geq q^{-i}$ ($i \in \{0,1,\ldots,d-1\}$), $\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[\mathbf{x},\mathbf{y},\mathbf{x}+\mathbf{y} \in T'\big] \geq \lambda^3q^{-3i}/4$. If there exists a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^d$ such that*

$$\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x})+f(\mathbf{y}) \;\big|\; \mathbf{x},\mathbf{y},\mathbf{x}+\mathbf{y} \in T\big] \geq 1-\frac{(q-1)\lambda^3}{8q^{3d-2}},$$

then there exists an affine map $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{d \times m} \times \mathbb{Z}_q^d$ *such that*

$$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} \;\big|\; \mathbf{x} \in T\big] \geq \left(\frac{1}{q} + \left(\frac{1}{q^3} + \frac{\lambda}{8q^{d+1}}\right)\right)^d.$$

Before proving Theorem 8, we prove the following lemma.

Lemma 9 *Let* $m, q \in \mathbb{N}$ *such that* q *is prime, and* $T \subset \mathbb{Z}_q^m$ *be a subset of density* $\lambda := |T|q^{-m} \leq 1/\sqrt{2}$ *such that* $\Pr_{\mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\big[\mathbf{x}, \mathbf{y}, \mathbf{x}+\mathbf{y} \in T\big] \geq \lambda^3/4$. *If* $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q$ *is a function such that*

$$\Pr_{\mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\big[\phi(\mathbf{x}+\mathbf{y}) = \phi(\mathbf{x}) + \phi(\mathbf{y}) \;\big|\; \mathbf{x}, \mathbf{y}, \mathbf{x}+\mathbf{y} \in T\big] \geq 1 - \gamma,$$

for some $\gamma > 0$, *then* $\exists(\mathbf{a}, b) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$ *such that*

$$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\phi(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle + b \;\big|\; \mathbf{x} \in T\big] \geq \frac{1}{q} + \left(\frac{1}{q^3} + \frac{\lambda}{4q^2}\left(1 - \frac{q\gamma}{q-1}\right)\right).$$

Proof. Define $\psi : \mathbb{Z}_q^m \to \mathbb{Z}_q$ by $\psi(\mathbf{x}) = \phi(\mathbf{x})$, if $\mathbf{x} \in T$, and $\alpha \sim \mathbb{Z}_q$ otherwise. Let $\mathrm{P} := \Pr_{\mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\big[\mathbf{x}, \mathbf{y}, \mathbf{x}+\mathbf{y} \in T\big]$. Then, we have

$$\Pr_{\mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\big[\psi(\mathbf{x}+\mathbf{y}) = \psi(\mathbf{x}) + \psi(\mathbf{y})\big] = \mathrm{P}(1-\nu) + (1-\mathrm{P}) \cdot \frac{1}{q} = \frac{1}{q} + \mathrm{P}\left(1 - \left(\gamma + \frac{1}{q}\right)\right).$$

Now, for all $\alpha \in \mathbb{Z}_q$, define $\tau_\alpha : \mathbb{Z}_q^m \to \mathbb{C}$ by $\tau_\alpha(\mathbf{x}) = \omega^{\alpha \cdot \psi(\mathbf{x})}$, where $\omega \in \mathbb{C}$ is a primitive q^{th} root of unity in $\mathbb{C}$. We have

$$\begin{aligned}
&\frac{1}{q} + \mathrm{P}\left(1 - \left(\gamma + \frac{1}{q}\right)\right) = \Pr_{\mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\big[\psi(\mathbf{x}+\mathbf{y}) = \psi(\mathbf{x}) + \psi(\mathbf{y})\big] \\
&= \mathbb{E}_{\alpha \sim \mathbb{Z}_q, \mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\big[\tau_\alpha(\mathbf{x}+\mathbf{y})\bar{\tau}_\alpha(\mathbf{x})\bar{\tau}_\alpha(\mathbf{y})\big] \\
&= \mathbb{E}_{\alpha \sim \mathbb{Z}_q, \mathbf{x},\mathbf{y} \sim \mathbb{Z}_q^m}\left[\sum_{\mathbf{u},\mathbf{v}_1,\mathbf{v}_2 \in \mathbb{Z}_q^m} \hat{\tau}_\alpha(\mathbf{u}) \cdot \bar{\hat{\tau}}_\alpha(\mathbf{v}_1) \cdot \bar{\hat{\tau}}_\alpha(\mathbf{v}_2) \cdot \omega^{\langle \mathbf{x}, \mathbf{u}-\mathbf{v}_1 \rangle} \cdot \omega^{\langle \mathbf{x}, \mathbf{u}-\mathbf{v}_2 \rangle}\right] \\
&= \frac{1}{q} \sum_{\alpha \in \mathbb{Z}_q} \sum_{\mathbf{u} \in \mathbb{Z}_q^m} \hat{\tau}_\alpha(\mathbf{u}) \cdot \bar{\hat{\tau}}_\alpha(\mathbf{u})^2.
\end{aligned}$$

We prove Claim 9 below, from which the lemma follows. To see this, observe that we have

$$\frac{1}{q} + \mathrm{P}\Big(1 - \Big(\gamma + \frac{1}{q}\Big)\Big) = \frac{1}{q}\Big|\sum_{\alpha \in \mathbb{Z}_q}\sum_{\mathbf{u} \in \mathbb{Z}_q^m} \hat{\tau}_\alpha(\mathbf{u})\bar{\hat{\tau}}_\alpha(\mathbf{u})^2\Big| \le \frac{\lambda^3}{q}\sum_{\alpha \in \mathbb{Z}_q}\Big|\hat{\tau}_\alpha|_T(\mathbf{u}_\alpha)\Big|\sum_{\mathbf{u} \in \mathbb{Z}_q^m}\Big|\hat{\tau}_\alpha|_T(\mathbf{u})\Big|^2$$
$$= \frac{\lambda^2}{q}\sum_{\alpha \in \mathbb{Z}_q}\Big|\hat{\tau}_\alpha|_T(\mathbf{u}_\alpha)\Big| \le \frac{\lambda^2}{q} \cdot \Big(1 + (q-1)\Big|\hat{\tau}_{\alpha^*}|_T(\mathbf{u}_{\alpha^*})\Big|\Big),$$

where the first inequality follows from the first point of Claim 9 and the triangle inequality $\Big(\forall \alpha \in \mathbb{Z}_q, \mathbf{u}_\alpha := \arg\max_{\mathbf{u} \in \mathbb{Z}_q^m}\Big\{\Big|\hat{\tau}_\alpha|_T(\mathbf{u}_\alpha)\Big|\Big\}\Big)$, the second equality follows from the second point of Claim 9, and the last inequality follows from $\Big|\hat{\tau}_0|_T(\mathbf{u}_0)\Big| = 1$. Here we have defined $\alpha^* := \arg\max_{\alpha \in \mathbb{Z}_q^*}\Big\{\Big|\hat{\tau}_\alpha|_T(\mathbf{u}_\alpha)\Big|\Big\}$. It thus follows that

$$\Big|\hat{\tau}_{\alpha^*}|_T(\mathbf{u}_{\alpha^*})\Big| \ge \frac{1-\lambda^2}{\lambda^2} \cdot \frac{1}{q-1} + \frac{q}{q-1} \cdot \frac{\mathrm{P}}{\lambda^2} \cdot \Big(1 - \Big(\gamma + \frac{1}{q}\Big)\Big) \ge \frac{1}{q} + \eta,$$

for $\eta = \frac{\lambda}{4}\big(1 - \frac{q\gamma}{q-1}\big)$. Then,

$$\frac{1}{q} + \eta \le \Big|\mathbb{E}_{\mathbf{x} \sim T}\big[\omega^{\alpha^*\phi(\mathbf{x}) - \langle \mathbf{u}_{\alpha^*}, \mathbf{x}\rangle}\big]\Big| = \Big|\sum_{b \in \mathbb{Z}_q}\omega^b \cdot \Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\alpha^*\phi(\mathbf{x}) = \langle \mathbf{u}_{\alpha^*}, \mathbf{x}\rangle + b \;\big|\; \mathbf{x} \in T\big]\Big|,$$

and so by Claim 2, $\exists b \in \mathbb{Z}_q$ such that

$$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\phi(\mathbf{x}) = \langle (\alpha^*)^{-1}\mathbf{u}_{\alpha^*}, \mathbf{x}\rangle + (\alpha^*)^{-1}b\big] \ge \frac{1}{q} + \Big(\frac{1}{q^3} + \frac{\eta}{q^2}\Big),$$

which completes the proof.

Claim 9 *For all $\alpha \in \mathbb{Z}_q$:*

1. *For all $\mathbf{u} \in \mathbb{Z}_q^m$, let $\hat{\tau}_\alpha|_T(\mathbf{u}) := \mathbb{E}_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\tau_\alpha(\mathbf{x})\omega^{-\langle \mathbf{u}, \mathbf{x}\rangle} \;\big|\; x \in T\big]$. Then, each $\hat{\tau}_\alpha(\mathbf{u}) = \lambda \cdot \hat{\tau}_\alpha|_T(\mathbf{u})$.*

2. $\sum_{\mathbf{u} \in \mathbb{Z}_q^m}\big|\hat{\tau}_\alpha|_T(\mathbf{u})\big|^2 = \lambda^{-1}$.

Proof. Let $\alpha \in \mathbb{Z}_q$. We begin by proving the first point. We have

$$\hat{\tau}_\alpha(\mathbf{u}) = \mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[\omega^{\alpha\cdot\psi(\mathbf{x})-\langle\mathbf{u},\mathbf{x}\rangle}\right] = q^{-m}\cdot \sum_{\substack{(\beta,\mathbf{x})\in\mathbb{Z}_q\times\mathbb{Z}_q^m \text{ s.t.}\\ \alpha\cdot\psi(\mathbf{x})-\langle\mathbf{u},\mathbf{x}\rangle=\beta}} \omega^\beta$$

$$= \sum_{\beta\in\mathbb{Z}_q} \omega^\beta \cdot \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[\alpha\cdot\psi(\mathbf{x}) - \langle\mathbf{u},\mathbf{x}\rangle = \beta\right]$$

$$= \sum_{\beta\in\mathbb{Z}_q} \omega^\beta\Big(\lambda\cdot\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[\alpha\cdot\phi(\mathbf{x}) - \langle\mathbf{u},\mathbf{x}\rangle = \beta \;\middle|\; \mathbf{x}\in T\right] + (1-\lambda)q^{-1}\Big)$$

$$= \lambda\sum_{\beta\in\mathbb{Z}_q} \omega^\beta \cdot \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[\alpha\cdot\phi(\mathbf{x}) - \langle\mathbf{u},\mathbf{x}\rangle = b \;\middle|\; \mathbf{x}\in T\right]$$

$$= (\lambda/|T|)\sum_{(\mathbf{x},\beta)\in T\times\mathbb{Z}_q} \mathbb{1}_{\alpha\cdot\phi(\mathbf{x})-\langle\mathbf{u},\mathbf{x}\rangle=\beta}\cdot\omega^\beta = (\lambda/|T|)\sum_{\mathbf{x}\in T}\omega^{\alpha\cdot\phi(\mathbf{x})-\langle\mathbf{u},\mathbf{x}\rangle} = \lambda\cdot\hat{\tau}_\alpha|_T(\mathbf{u}).$$

where the third line follows from the definition of ψ and the fourth line follows from Claim 1.

Now, we prove the second point. We have

$$\sum_{\mathbf{u}\in\mathbb{Z}_q^m}\left|\hat{\tau}_\alpha|_T(\mathbf{u})\right|^2 = \sum_{\mathbf{u}\in\mathbb{Z}_q^m}\hat{\tau}_\alpha|_T(\mathbf{u})\cdot\bar{\hat{\tau}}_\alpha|_T(\mathbf{u}) = |T|^{-2}\sum_{\mathbf{u}\in\mathbb{Z}_q^m}\sum_{\mathbf{x},\mathbf{x}'\in T}\hat{\tau}_\alpha(\mathbf{x})\cdot\bar{\hat{\tau}}_\alpha(\mathbf{x}')\cdot\omega^{\langle\mathbf{u},\mathbf{x}'-\mathbf{x}\rangle}$$

$$= |T|^{-2}\sum_{\mathbf{x}\in T}\left|\hat{\tau}_\alpha(\mathbf{x})\right|^2\cdot q^m = q^m/|T| = \lambda^{-1},$$

where the third equality follows from Claim 1. ■ ■

Now, we can prove Theorem 8.

Proof of Theorem 8. For each $i \in [d]$, we'll denote the i^{th} projection of f by $f_i : \mathbb{Z}_q^m \to \mathbb{Z}_q$. We use induction to show that for each $i \in [d]$, $\exists((\mathbf{a}_j, b_j))_{j=1}^i \in (\mathbb{Z}_q^m \times \mathbb{Z}_q)^i$ such that

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[f_j(\mathbf{x}) = \langle\mathbf{a}_j,\mathbf{x}\rangle + b_j, \forall j\in[i] \;\middle|\; \mathbf{x}\in T\right]$$
$$\geq \prod_{j=1}^{i}\left(\frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{4q^{i+1}}\left(1 - \frac{q}{q-1}\cdot\frac{4q^{3(i-1)}\gamma}{\lambda^3}\right)\right),$$

for $\gamma = \frac{(q-1)\lambda^3}{8q^{3d-2}}$. Thus

$$\begin{aligned}
&\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f_i(\mathbf{x}) = \langle \mathbf{a}_i, \mathbf{x}\rangle + b_i, \forall i \in [d] \;\big|\; \mathbf{x} \in T\big] \\
&\geq \left(\frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{4q^{d+1}}\left(1 - \frac{q}{q-1} \cdot \frac{4q^{3(d-1)}\gamma}{\lambda^3}\right)\right)^d \geq \left(\frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{8q^{d+1}}\right)^d,
\end{aligned}$$

as desired.

For the base case in which $i = 1$, observe that

$$\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f_1(\mathbf{x}+\mathbf{y}) = f_1(\mathbf{x}) + f_1(\mathbf{y}) \;\big|\; \mathbf{x} \in T\big] \geq \Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y}) \;\big|\; \mathbf{x} \in T\big]$$

$\geq 1 - \gamma$.

So, by Lemma 9, $\exists(\mathbf{a}_1, b_1) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$ such that

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f_1(\mathbf{x}) = \langle \mathbf{a}_1, \mathbf{x}\rangle + b_1\big] \geq \frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{4q^2}\left(1 - \frac{q\gamma}{(q-1)\lambda^3}\right).$$

Now, for the induction step, assume $\exists((\mathbf{a}_j, b_j))_{j=1}^{i-1} \in (\mathbb{Z}_q^m \times \mathbb{Z}_q)^{i-1}$ such that

$$\begin{aligned}
&\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f_j(\mathbf{x}) = \langle \mathbf{a}_j, \mathbf{x}\rangle + b_j, \forall j \in [i-1] \;\big|\; \mathbf{x} \in T\big] \\
&\geq \prod_{j=1}^{i-1}\left(\frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{4q^{j+1}}\left(1 - \frac{q}{q-1} \cdot \frac{4q^{3(j-1)}\gamma}{\lambda^3}\right)\right).
\end{aligned}$$

Let $T_i = \{\mathbf{x} \in T : f_j(\mathbf{x}) = \langle \mathbf{a}_j, \mathbf{x}\rangle + b_j, \forall j \in [i-1]\}$, and note that $\lambda_i := |T_i|q^{-m} \geq \lambda q^{-(i-1)}$.

Also, we have

$$\begin{aligned}
&\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f_i(\mathbf{x}+\mathbf{y}) = f_i(\mathbf{x}) + f_i(\mathbf{y}) \;\big|\; \mathbf{x} \in T_i\big] \\
&\geq \Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y}) \;\big|\; \mathbf{x} \in T_i\big] \\
&= \frac{\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y}) \text{ and } \mathbf{x} \in T_i \;\big|\; \mathbf{x} \in T\big]}{\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[\mathbf{x} \in T_i \;\big|\; \mathbf{x} \in T\big]} \\
&\geq \frac{\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[\mathbf{x} \in T_i \;\big|\; \mathbf{x} \in T\big] - \gamma}{\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[\mathbf{x} \in T_i \;\big|\; \mathbf{x} \in T\big]} \geq 1 - \frac{4\gamma}{\lambda_i^3} = 1 - \frac{4q^{3(i-1)}\gamma}{\lambda^3}.
\end{aligned}$$

Thus by Lemma 9, $\exists(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$ such that

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[f_i(\mathbf{x}) = \langle \mathbf{a}_i, \mathbf{x}\rangle + b_i \;\middle|\; \mathbf{x} \in T_i\right] \geq \frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{4q^{i+1}}\left(1 - \frac{q}{q-1} \cdot \frac{4q^{3(i-1)}\gamma}{\lambda^3}\right),$$

and so $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\left[f_j(\mathbf{x}) = \langle \mathbf{a}_j, \mathbf{x}\rangle + b_j, \forall j \in [i] \;\middle|\; \mathbf{x} \in T\right]$ is at least

$$\prod_{j=1}^{i}\left(\frac{1}{q} + \frac{1}{q^3} + \frac{\lambda}{4q^{j+1}}\left(1 - \frac{q}{q-1} \cdot \frac{4q^{3(j-1)}\gamma}{\lambda^3}\right)\right),$$

which completes the induction step. ■

Chapter 3

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

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 Banarjee, Peikert and Rosen (Eurocrypt'12) 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 via a reduction from LWE in the above parameter regime. Specifically, we show that any "point-wise" reduction from LWE to LWR can be used to directly break the corresponding LWE problem. A reduction is "point-wise" if it maps LWE samples to LWR samples one at a time. Our argument goes roughly as follows: first we show that any point-wise reduction from LWE to LWR must have good agreement with some affine map; then we use a Goldreich-Levin-type theorem to extract the LWE secret given oracle access to a point-wise reduction with good affine agreement. Both components may be of independent interest.

3.1 Introduction

Regev's learning with errors (LWE) problem [Reg05] is fundamental for modern cryptography due to its versitility 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

$$\big\{(\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s}\rangle + e_i)\big\} \subset \big(\mathbb{Z}_q^n \times \mathbb{Z}_q\big)^m, \tag{3.1}$$

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 15 years ago an immense research effort has established LWE-based constructions for most known cryptographic primitives (*e.g.*, [GPV08, ACPS09, BGV11, CHKP12, MP12, BNS13, GSW13, GVW15, GKW18, 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 constructing 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

$$\big\{(\mathbf{a}_i, b_i = \lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle \rceil_p)\big\} \subset \big(\mathbb{Z}_q^n \times \mathbb{Z}_p\big)^m, \tag{3.2}$$

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).

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

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

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 $\big\{(\mathbf{a}_i, b_i')\big\} \subset \mathbb{Z}_q \times \mathbb{Z}_p$, (somehow) recovers the hidden secret $\mathbf{s}$, then scans the m samples to make sure that $b_i' = \big\lfloor \langle \mathbf{a}_i, \mathbf{s} \rangle \big\rceil_p$ 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.

3.1.1 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.

Our Reduction Model. A natural question is whether our theorem holds for relaxations 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 $\mathsf{LWR}_{n',q',p'}$)? Moreover, we might hope that our main result would hold even for pointwise reductions which are allowed to abort on some inputs. We actually consider such reductions and note that part of the proof of our main theorem goes through even when the pointwise reduction is allowed to abort. However, we were only able to prove some of the steps for non-aborting pointwise reductions so our main theorem inherits this restriction. We believe that it should be possible to prove our main theoem even for pointwise reductions which are allowed to abort.

In a similar vein, 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. 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 and (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.

3.2 Preliminaries

Throughout this work, the integer n will denote the security parameter. If $m \in \mathbb{N}$, then we denote by $[m]$ the set $\{1, \ldots, m\}$. 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.

3.2.1 Learning with Errors/Rounding

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

- **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:** *The* learning with rounding distribution $\mathsf{LWR}_{n,q,\mathbf{s},X}$ *is:*
  - *draw $\mathbf{a} \sim \mathbb{Z}_q^n$, set $b = \mathrm{argmin}_{x \in X}\big\{|\langle \mathbf{a}, \mathbf{s} \rangle - x|\big\}$ (breaking ties arbitrarily) and output $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times X$.*[2]

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

[2] Here $|\alpha - \beta|$ for $\alpha, \beta \in \mathbb{Z}_q$ denotes $\min\{|\hat{\alpha} - \hat{\beta}| : \hat{\alpha}, \hat{\beta} \in \mathbb{Z} \text{ st } (\hat{\alpha}, \hat{\beta}) \equiv (\alpha, \beta) \pmod{q}\}$; $|\cdot|$ the real absolute value.

Definition 2 (The LWE/LWR Problems) *Let $n, q, m \in \mathbb{N}$ be positive integers, χ be a distribution on $\mathbb{Z}_q$, and $X \subsetneq \mathbb{Z}_q$ be a proper subset. The* search/decisional *version of the* learning with errors/rounding problems *refer to the following computational tasks.*[3]

- **Search LWE/LWR:** *Given $(\mathbf{a}_1, b_1), \ldots, (\mathbf{a}_m, b_m) \sim \mathsf{LWE}_{n,q,m,\chi}/\mathsf{LWR}_{n,q,m,X}$, output* $\mathbf{s}$.
- **Decisional LWE:** *Distinguish $\mathsf{LWE}_{n,q,m,\chi}$ from $\mathsf{Unif}\big(\mathbb{Z}_q^n \times \mathbb{Z}_q\big)^m$.*

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 $1 - 2^{-n}$. The rounding set for LWR will be $X = \mathbb{Z}_p$, the set of nearest integers to the multiples of q/p in $\mathbb{Z}_q$. We write $\lfloor b \rceil_p$ instead of $\mathrm{argmin}_{x \in X}\big\{|b - x|\big\}$, and we write $\mathsf{LWR}_{n,q,p}$ instead of $\mathsf{LWR}_{n,q,\mathbb{Z}_p}$.

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. $\mathsf{LWR}_{n,q,X}$) if it solves search LWE (resp. search LWR) with probability at least ε, given m samples:

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

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

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

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

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 $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ and $\ell_{\mathsf{samp}} : \mathbb{N} \to \mathbb{N}$ 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 $(\varepsilon, m) = \big(\ell_{\mathsf{err}}(\varepsilon'), \ell_{\mathsf{samp}}(m')\big)$.*

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 $\ell_{\mathsf{err}}(\varepsilon') = \mathsf{poly}\big(1/n, \varepsilon'\big)$ and $\ell_{\mathsf{samp}}(m') = \mathsf{poly}(n, m')$. 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 $m' \leq B$ samples.

3.2.2 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 the fact that the inner product mod q is a good two-source extractor. Results of this type originated with the work of Goldreich and Chor [CG88], the proof of this next claim is similar.

We will use the mod q version of the the fact that the inner product is a good two-source extractor. Results of this type originated with the work of Goldreich and Chor [CG88].

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 $\mathbf{a}\sim\mathbb{Z}_q$, $\mathbf{x}\sim X$ and outputs $\big(\mathbf{a},\langle\mathbf{a},\mathbf{x}\rangle\big)$. Then*

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

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 $e \in \mathbb{Z}_q$,*

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

In particular,

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

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 $\Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}\big[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in T\big] > \tau + q^{-n/4}$ 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 $e \in \mathbb{Z}_q$, $\Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}\big[(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in T\big] < \tau - q^{-n/4}$. For the part of the claim about LWE samples, note that if $\mathbf{s} \notin S$ then

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

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

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

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

3.3 Our Reduction Model and Main Theorem

3.3.1 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 $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to (\mathbb{Z}_q^n \times \mathbb{Z}_p) \cup \{\perp\}$, applied pointwise on each of the input samples). Then $\mathcal{A}^{\mathsf{S}}$ invokes S on the "non-bot" outputs obtaining an LWR secret. Finally, $\mathcal{A}^{\mathsf{S}}$ outputs an LWE secret computed using the original LWE samples and the LWR secret. All LWE to LWR reductions in the literature fit into this pointwise model.

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 $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ and $\ell_{\mathsf{samp}} : \mathbb{N} \to \mathbb{N}$ 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 $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \big(\mathbb{Z}_q^n \times \mathbb{Z}_p\big) \cup \{\perp\}$ 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 $(\varepsilon, m) = \big(\ell_{\mathsf{err}}(\varepsilon'), \ell_{\mathsf{samp}}(m')\big)$.*

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

2. *Call $\mathsf{S}\big(\{(\mathbf{a}_i', b_i')\} \setminus \{\perp\}\big)$ 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* $\mathbf{s}$.

We say $\mathcal{A} = (f, \mathcal{B})$ *is a* $\nu-$non-aborting *pointwise reduction if* $\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[f(\mathbf{a}, b) \neq \perp\big] \geq \nu$*. We say* $\mathcal{A}$ *is a* non-aborting *pointwise reduction if it is a* $1-$*non-aborting pointwise reduction;* i.e., *if* $f(\mathbf{a}, b) \neq \perp$ *for all* $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$*.*

Theorem 2 (Main) *Let* $n, p, q \in \mathbb{N}$ *be integers such that such that* $q = \mathsf{poly}(n)$ *is prime and such that* $q^{2/3+c} < p < q = \mathsf{poly}(n)$ *for a constant* $c > 0$*, and let* χ *be a distribution on* $\mathbb{Z}_q$*. Let* $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ *and* $\ell_{\mathsf{samp}} : \mathbb{N} \to \mathbb{N}$ *be functions so* $\ell_{\mathsf{err}}(\varepsilon') = \mathsf{poly}\big(1/n, \varepsilon'\big)$ *and* $\ell_{\mathsf{samp}}(m') = \mathsf{poly}(n, m')$*. Then any non-aborting* $(\ell_{\mathsf{err}}, \ell_{\mathsf{samp}})-$*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* $m = \mathsf{poly}(n)$*.*

We also state as a conjecture, our main theorem without the lower bound requirement on p, and where the pointwise reduction is allowed to abort.

Conjecture 1 *Let* $n, p, q \in \mathbb{N}$ *be integers such that such that* $q = \mathsf{poly}(n)$ *is prime. Let* $\nu = \nu(n) > 0$ *be non-negligible in* n*, and let* χ *be a distribution on* $\mathbb{Z}_q$*. Let* $\ell_{\mathsf{err}} : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ *and* $\ell_{\mathsf{samp}} : \mathbb{N} \to \mathbb{N}$ *be functions such that* $\ell_{\mathsf{err}}(\varepsilon') = \mathsf{poly}\big(1/n, \varepsilon'\big)$ *and* $\ell_{\mathsf{samp}}(m') = \mathsf{poly}(n, m')$*. Then any* $\nu-$*non-aborting* $(\ell_{\mathsf{err}}, \ell_{\mathsf{samp}})-$*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 these results say that it is impossible to reduce $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$ in a pointwise fashion. The only difference between Theorem 3 and Conjecture 1 is that Theorem 3 makes two additional assumptions about the parameters:

- $q^{2/3+c} < p$;

- $\nu = 1$ (*i.e.*, f is non-aborting).

The first assumption is needed in one specific point of the proof of Theorem 3; we will indicate this point when we get to it. We make use of the second assumption throughout. Occasionally, it is possible to rework the proofs to some of our supporting lemmas to allow f to abort, but since there is more than one point where we require it, we just assume it everywhere; this will simplify our overall proof. Nevertheless, as mentioned in the introduction, we believe it should be possible to remove the dependence on these extra hypotheses.

3.3.2 The LWR Secret Recovery Algorithm and Proof of Theorem 3

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 $\mathsf{LWR}_{\mathbf{s}'}$, respectively, instead of $\mathsf{LWE}_{n,q,\mathbf{s},\chi}$ and $\mathsf{LWE}_{n,q,\mathbf{s},p}$. The lemmas in this section make reference to non-negligible quantities $\eta, \delta > 0$ which will be specified in the next section.

Lemma 10 (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 $\left(\mathbb{Z}_q^n \times \mathbb{Z}_q\right)-$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:*

$$\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\left[b' = \left\lfloor\langle \mathbf{a}', \mathbf{s}'\rangle\right\rceil_p\right] \geq 1 - \eta. \tag{3.3}$$

Lemma 11 *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*

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

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

We now prove Theorem 3 assuming Lemmas 10 and 11.

Proof of Theorem 3. Let $\mathcal{A}$ denote the algorithm promised by Lemma 10. Consider the following distinguishing algorithm $\mathcal{D}$, which gets oracle access to a $\left(\mathbb{Z}_q^n \times \mathbb{Z}_q\right)-$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 a random bit.

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

$$\mathsf{P} := \Pr_{(\mathbf{a},b)\sim\mathcal{O}}\Big[b' = \big\lfloor\langle\mathbf{a}', \mathbf{s}'\rangle\big\rceil_p\Big].$$

   If $\hat{\mathsf{P}} \geq 1 - 3\eta/4$, D outputs 1, otherwise D outputs a random bit.

Note D either outputs 1 or a random bit. We show that it outputs a random bit 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 $|\hat{\mathsf{P}} - \mathsf{P}| < \eta/4$ holds with probability $1 - 2^{-\Omega(n)}$. Thus by Lemma 11, $\hat{\mathsf{P}} < 1 - 3\eta/4$ 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 10 ensures that with non-negligible probability, $\mathcal{A}$ outputs $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\mathsf{P} \geq 1 - \eta$. In this case, $\hat{\mathsf{P}}$ is again accurate to within $\pm\eta/4$ by the Chernoff bound, and so $\hat{\mathsf{P}} \geq 1 - 3\eta/4$ and D outputs 1 with non-negligible probability.

■

3.4 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.

3.4.1 Non-Degeneracy

We prove that the distribution which draws $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$ 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 $\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[\mathbf{a}' \in T\big] \geq \zeta$, 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 10 (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 $m = \max\big\{qn(1 + \log q), \rho n/\varepsilon^2\big\}$.*

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 $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$ and outputs $\mathbf{a}'$. By definition, there exists $T \subset \mathbb{Z}_q^n$ of density ρ such that $\Pr_{\mathcal{D}}\big[\mathbf{a}' \in T\big] \geq \rho + \varepsilon$. 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 $t \geq \big(\rho + \varepsilon/2\big)m'$; 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 $m' \geq \rho n/\varepsilon^2$). As $m' \geq nq(1 + \log q)$, 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 $m \geq 2m'/\nu$ LWE samples are chosen and S is fed with $\big\{f(\mathbf{a}_i, b_i)\big\}$, $t \geq (\rho + \varepsilon)m'$ in expectation, and so by the Chernoff-Hoeffding inequality, $t \geq \big(\rho + \varepsilon/2\big)m'$ holds with probability $1 - 2^{-\Omega(n)}$. 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. ■

3.4.2 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 $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} := \Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_{\mathbf{s}}}\big[b' = \lfloor\langle\mathbf{a}',\mathbf{s}'\rangle\rceil_p\big]$. We write $Y_\varepsilon(\mathbf{s}) = \{\mathbf{s}' \in Y : \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1-\varepsilon\}$ for $\mathbf{s} \in X$ and $\varepsilon > 0$. Likewise, given $\mathbf{s}' \in Y$ and $\varepsilon > 0$, $X_\varepsilon(\mathbf{s}') = \{\mathbf{s} \in X : \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1-\varepsilon\}$. 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, $2\delta = \ell_{\mathsf{err}}(1/3)$, so that if S is a $\frac{1}{3}-$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 $\eta \leq \min\big\{\sigma, \big(1/3nq\big)^3\big\}$. The reader is encouraged on a first pass to 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:*

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

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

Note that as $\eta \leq \sigma$, points (1) and (2) combine to imply that for every $\mathbf{s} \in \mathsf{G}$ there is a unique $\mathbf{s}' \in \mathbb{Z}_q^n$ such that $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1-\eta$. Thus, point (3) additionally says 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 11 *Suppose* $(f, \mathcal{B})$ *is a pointwise reduction from* $\mathsf{LWE}_{n,q,\chi}$ *to* $\mathsf{LWR}_{n,q,p}$. *Then either* $|\mathsf{G}| \geq \delta \cdot q^n$, *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 $\big\{(\mathbf{a}_i', b_i')\big\}_{i=1}^m$, does the following:

(i) it looks at the first $nq(1+\log q)$ samples (this is less than m since $\eta \leq 1/q$) and checks whether there exist distinct $\mathbf{s}', \mathbf{s}'' \in \mathbb{Z}_q^n$ such that $\big\lfloor \langle \mathbf{a}_i', \mathbf{s}' \rangle \big\rceil_p = b_i' = \big\lfloor \langle \mathbf{a}_i', \mathbf{s}'' \rangle \big\rceil_p$ 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 $b_i' = \big\lfloor \langle \mathbf{a}_i', \mathbf{s}' \rangle \big\rceil_p$ 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 ${}^{\#}\{\mathbf{s} \in \mathbb{Z}_q^n : |Y_\eta(\mathbf{s})| = 1 \ \&\ \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1 - \eta\} \geq 2$; 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 10, these two points suffice. Point 1 says that S is a $\big(\frac{1}{3}, m\big)$−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 10, 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

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

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

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

(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 $\mathsf{p}_{(\mathbf{s},\mathbf{s}')}, \mathsf{p}_{(\mathbf{s},\mathbf{s}'')} \geq 1 - \sigma$ both hold. In this case,

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

(using $\sigma \leq \delta/2nq(1 + \log q)$) 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 $|X_\eta(\mathbf{s}')| \geq 1$ since $\mathbf{s} \in X_\eta(\mathbf{s}')$, thus if (3) does not hold then it must be that $|X_\eta(\mathbf{s}')| \geq 2$. 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. ■

3.4.3 Proof of Lemma 11

Claim 11 imposes quite a lot of structure on a pointwise reduction. We will refer to Claim 11 repeatedly throughout the remainder of the paper. Additionally, we can already derive Lemma 11 as a corollary.

Lemma 11 (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*

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

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 $\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[b' = \lfloor\langle\mathbf{a}', \mathbf{s}'\rangle\rceil_p\big] \geq 1 - \eta/2$. Then by Corollary 1, $\Pr_{(\mathbf{a},b)\sim\mathsf{LWE}_\mathbf{s}}\big[b' = \lfloor\langle\mathbf{a}', \mathbf{s}'\rangle\rceil_p\big] \geq 1 - \eta/2 - q^{-n/4} \geq 1 - \eta$ holds for all but a $q^{-\Omega(n)}-$fraction of $\mathbf{s} \in \mathbb{Z}_q^n$. In other words, $|X_\eta(\mathbf{s}')| \geq (1 - q^{-\Omega(n)}) \cdot q^n$, so the degree of $\mathbf{s}'$ is way too high to have any neighbors in G. However, this means that $\mathsf{G} \subset \mathbb{Z}_q^n \setminus X_\eta(\mathbf{s}')$, and so $|\mathsf{G}| \leq q^{-\Omega(n)} \cdot q^n$ and so by Claim 11, $\mathcal{B}$ is a $(\delta, m)-$solver for $\mathsf{LWE}_{n,q,\chi}$. ■

3.5 Outline of the Rest of the Paper

At this point we have reduced our main result (Theorem 3) to proving Lemma 10; 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 11, 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

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

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 10 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 10: $n, p, q \in \mathbb{N}$ 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 11.

Lemma 12 *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:*

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

where $\tau = 8q^2n^4\eta^{1/3t}$, and $t \in \mathbb{N}$ minimal such that $t \geq \frac{\log_q(1/\delta)+2}{3c}$ holds.

Using Lemma 12 to Prove Lemma 10. Once we know that $\mathbf{a}'$ has good agreement with $\mathbf{H}\mathbf{a}$, we can recover $\mathbf{s}'$ using a Goldreich-Levin-type argument. Let us assume for simplicity in this discussion that $\mathbf{a}' = \mathbf{H}\mathbf{a}$ occurs with good probability, rather than $\mathbf{a}' \in \mathrm{Span}(\mathbf{H}\mathbf{a}) + \mathbf{V}$. The key point is that when $\mathbf{s} \in \mathsf{G}$ is good,

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

occurs with high probability. Thus, if we simply output a random $x \sim \mathbb{Z}_q$ such that $\lfloor x \rceil_p = b'$ we will be predicting the inner product $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}} \mathbf{s}' \rangle$ with non-negligible advantage over guessing. The Goldreich-Levin machinery can then be used to recover $\mathbf{H}^{\mathsf{t}} \mathbf{s}'$, and this will be good enough to prove Lemma 10.

Proving Lemma 12. The proof of Lemma 12 is broken into two parts. In the first part of the proof of Lemma 12, we prove that for any pointwise reduction from $\mathsf{LWE}_{n,q,\chi}$ to $\mathsf{LWR}_{n,q,p}$, there exists a constant dimensional $\mathbf{V} \subset \mathbb{Z}_q^n$ such that the following property test accepts with good probability:

- choose $(\mathbf{a}_0, b_0), (\mathbf{a}_1, b_1) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$ and non-zero $\alpha, \beta \sim \mathbb{Z}_q \setminus \{0\}$;
- compute $(\mathbf{a}'_0, b'_0) = f(\mathbf{a}_0, b_0)$, $(\mathbf{a}'_1, b'_1) = f(\mathbf{a}_1, b_1)$, and $(\mathbf{a}'_2, b'_2) = f(\alpha \mathbf{a}_0 + \beta \mathbf{a}_1, \alpha b_0 + \beta b_1)$;
- output 1 if $\mathbf{a}'_2 \in \mathrm{Span}\big(\{\mathbf{a}'_0, \mathbf{a}'_1\}\big) + \mathbf{V}$; output 0 if not.

The logic behind this property test is the following. Let us pretend for this discussion that $\mathbf{V} = \{\mathbf{0}\}$, in which case the property tests whether $\{\mathbf{a}'_0, \mathbf{a}'_1, \mathbf{a}'_2\}$ is linearly independent or not. If $\{\mathbf{a}'_0, \mathbf{a}'_1, \mathbf{a}'_2\}$ were linearly independent, then $\{b'_0, b'_1, b'_2\}$ would represent three different linear relations about the LWR secret $\mathbf{s}'$. Since $\{\mathbf{a}_0, \mathbf{a}_1, \mathbf{a}_2\}$ is linearly dependent (writing $\mathbf{a}_2 = \alpha \mathbf{a}_0 + \beta \mathbf{a}_1$), $\{b_0, b_1, b_2\}$ represents only two linear relations about the LWE secret $\mathbf{s}$. The key point is that a pointwise reduction cannot allow you to generate many linear relations about $\mathbf{s}'$ using only a few linear relations about a good $\mathbf{s} \in \mathsf{G}$, since otherwise it would mean that there would be many good $\mathbf{s} \in \mathsf{G}$ which correspond to the same LWR secret $\mathbf{s}'$, contradicting that good LWE secrets form a perfect matching with their corresponding LWR secrets. It is here that we need the bound $q^{2/3+c} < p$, since each b_i' does not decrease the number of possible secrets by $1/q$, but rather by $1/p$, since there are q/p different possibilities for $\langle \mathbf{a}', \mathbf{s}' \rangle$ which satisfy $b' = \lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \rceil_p$. Thus, when $\{\mathbf{a}_0', \mathbf{a}_1', \mathbf{a}_2'\}$ is linearly independent, only $p^{-3}-$fraction of the LWR secrets will satisfy the linear constraints, whereas $q^{-2}-$fraction of the LWE secrets will satisfy the linear constraints corresponding to $\{\mathbf{a}_0, \mathbf{a}_1, \mathbf{a}_2\}$. We need $p^{-3} \ll q^{-2}$ to ensure that the set of remaining LWR secrets is shrinking faster than the set of remaining LWE secrets.

The final part of the proof of Lemma 12 involves proving that any function which passes the above property test with good probability must have good agreement with a linear function. This part of the proof follows the proof of the fundamental theorem of projective geometry (see *e.g.* Section 2.10 of [Art57]).

Proposition 1 (Fundamental Theorem of Projective Geometry) *Let q be a prime and $f : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ be a function such that for any one-dimensional line $\ell \subset \mathbb{Z}_q^n$, the set $f(\ell) := \{f(\mathbf{x}) : \mathbf{x} \in \ell\} \subset \mathbb{Z}_q^n$ is also a line. Then f is affine.*

In our case, the hypothesis that $f(\ell) \subset \mathbb{Z}_q^n$ is a line for all lines $\ell \subset \mathbb{Z}_q^n$ is replaced by the property test passing with good probability over $(\mathbf{a}_0, b_0), (\mathbf{a}_1, b_1), (\mathbf{a}_2, b_2) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$, and $\alpha, \beta \sim \mathbb{Z}_q \setminus \{0\}$. Likewise, the conclusion is replaced by $\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}$ with high probability over $(\mathbf{a}, b) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$.

3.6 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 11, and a set of "good" LWE secrets $\mathsf{G} \subset \mathbb{Z}_q^n$ from Section 3.4.2. 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

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

For $\mathbf{s} \in \mathbb{Z}_q^n$ and $e \in \mathbb{Z}_q$, let us define $\mathsf{P}_{\mathbf{s},e}(\mathbf{H}, \mathbf{V}) := \Pr_{\mathbf{a} \sim \mathbb{Z}_q^n}\big[\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}\big]$, 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:

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

3.6.1 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 $\mathbf{H}^{\mathbf{t}}\mathbf{s}'$ given oracle access to $\mathsf{LWE}_{\mathbf{s}}$ for unknown $\mathbf{s}$.

Lemma 13 (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 $\Pr\big[\chi = e\big] \geq \frac{4\zeta}{5qn^2}$, and non-zero $\gamma \in \mathbb{Z}_q^*$, $\mathrm{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, 0) \geq \mathrm{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, \gamma) - \zeta$;*

  · *for all non-zero $\gamma \in \mathbb{Z}_q^*$, $\mathrm{P}_{\mathbf{s}}(\bar{\mathbf{s}}, 0) \geq \mathrm{P}_{\mathbf{s}}(\bar{\mathbf{s}}, \gamma) + 10\zeta$.*

  Then

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

Remark 14 *Intuitively, the requirement $\mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}}, 0) \geq \mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}}, \gamma) + 10\zeta$ means that the most likely output of the predictor on samples from $\mathsf{LWE}_{\mathbf{s}}$ is $\bar{\mathbf{s}}$. The additional requirement that $\mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, 0) \geq \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}}, \gamma) - \zeta$ 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 $\hat{\mathsf{p}}_i(\beta) \geq \hat{\mathsf{p}}_i(\beta') + 3\zeta$ for all $\beta' \neq \beta$, set $w_i = \beta$; otherwise set $w_i = \perp$.

3. Finally, let $W = \big\{ i \in \{1, \ldots, m\} : w_i \neq \perp \big\}$, 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), $\big|\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}} \rangle + \gamma) - \mathsf{P}_{\mathbf{s}, e_i}(\bar{\mathbf{s}}, \gamma)\big| < \zeta$ 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 $\Pr\big[\chi = e_i\big] \geq \frac{1}{5qm}$; 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

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

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 $\mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}}, 0) - \mathsf{P}_{\mathbf{s}}(\bar{\mathbf{s}}, \gamma) \geq 10\zeta$ for non-zero $\gamma \in \mathbb{Z}_q^*$ means that $\mathsf{P}_{\mathbf{s}, e}(\bar{\mathbf{s}}, 0) - \mathsf{P}_{\mathbf{s}, e}(\bar{\mathbf{s}}, \gamma) \geq 5\zeta$ holds with probability at least 5ζ over $e \sim \chi$. By Chernoff, the probability that $\mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},0) - \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma) \geq 5\zeta$ 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 $\mathrm{Span}\big(\{\mathbf{a}_{i_1}, \ldots, \mathbf{a}_{i_n}\}\big) = \mathbb{Z}_q^n$ and such that each error term e satisfies $\mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},0) - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},\gamma) \geq 5\zeta$ for all non-zero $\gamma \in \mathbb{Z}_q^*$. For each $i \in \{i_1, \ldots, i_n\}$,

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

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 $4/5 - 2^{-\Omega(n)} \geq 3/4$, 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 $\big|\hat{\mathsf{p}}_i(\langle \mathbf{a}_i, \bar{\mathbf{s}}\rangle + \gamma) - \mathsf{P}_{\mathbf{s},e_i}(\bar{\mathbf{s}},\gamma)\big| < \zeta$ 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 $\mathsf{Q} := \Pr\big[|\hat{\mathsf{p}}(\langle \mathbf{a}, \bar{\mathbf{s}}\rangle + \gamma) - \mathsf{P}_{\mathbf{s},e}(\bar{\mathbf{s}},\gamma)| > \zeta\big]$ be shorthand. We have

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

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

3.6.2 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 pointise 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 $\mathsf{P}(\mathbf{H}, \mathbf{V}) \geq 1 - \tau$. 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 $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$ well enough so that it is possible to use Lemma 13 to recover $\mathbf{H}^{\mathsf{t}}\mathbf{s}'$.

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 $\mathbf{a}' = \alpha \mathbf{H}^{\mathsf{t}}\mathbf{a} + \mathbf{v}$ both hold, $b' = \big\lfloor \alpha \langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle + \langle \mathbf{v}, \mathbf{s}' \rangle \big\rceil_p$ 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 $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$ 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 13, and so can be used to recover $\mathbf{H}^{\mathsf{t}}\mathbf{s}'$.

Claim 12 *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* $\beta \in \mathbb{Z}_q$, $\Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \geq \frac{1}{q^2}$*, where* $\mathcal{D}_{\mathbf{s}}$ *is the distribution which draws* $(\mathbf{a}, b) \sim \mathsf{LWE}_{\mathbf{s}}$ *such that* $\mathbf{a}' \in \mathrm{Span}(\mathbf{Ha}) + \mathbf{V}$*, and outputs* $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$*. 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 13 are satisfied for* $\bar{\mathbf{s}} = \mathbf{H}^{\mathsf{t}}\mathbf{s}'$*.*

Proof. Fix $\zeta = \frac{1-2\tau-q^2\eta}{10q^3}$. We must show two points:

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

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 13:

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

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

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

So if we shorthand the second term by $\mathsf{P}_e^{(2)}(\gamma)$, then $\mathsf{P}_e^{(1)}(0) - \mathsf{P}_e^{(1)}(\gamma) = \mathsf{P}_e^{(2)}(0) - \mathsf{P}_e^{(2)}(\gamma)$. Now let

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

Note that when $e \in \mathbb{Z}_q$ is such that $\Pr\big[\chi = e\big] \geq \frac{4\zeta}{5qn^2}$, $\mathsf{P}_3^{(2)} - \frac{5qn^2\eta}{4\zeta} \leq \mathsf{P}_e^{(3)}(\gamma) \leq \mathsf{P}_e^{(2)}(\gamma)$, since $\mathsf{s} \in \mathsf{G}$ and so $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1 - \eta$. Therefore, $\mathsf{P}_e^{(2)}(0) - \mathsf{P}_e^{(2)}(\gamma) \geq \mathsf{P}_e^{(3)}(0) - \mathsf{P}_e^{(3)}(\gamma) - \zeta$, using

$\eta \le \frac{4\zeta^2}{5qn^2}$. 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 $\alpha^{-1}\bigl(x - \sum_i c_i\alpha_i\bigr)$ for a random $x \sim \mathbb{Z}_q$ such that $\lfloor x \rceil_p = b'$. Therefore, when $b' = \bigl\lfloor \langle \mathbf{a}', \mathbf{s}' \rangle \bigr\rceil_p = \bigl\lfloor \alpha\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle + \langle \mathbf{v}, \mathbf{s}' \rangle \bigr\rceil_p$, Pred outputs $\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle$ with probability roughly p/q when $\bigl\lfloor \alpha(\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle + \gamma) + \langle \mathbf{v}, \mathbf{s}' \rangle \bigr\rceil_p = \bigl\lfloor \alpha\langle \mathbf{a}, \mathbf{H}^{\mathsf{t}}\mathbf{s}' \rangle + \langle \mathbf{v}, \mathbf{s}' \rangle \bigr\rceil_p$, and with probability 0 otherwise. It follows that $\mathsf{P}_e^{(3)}(0) - \mathsf{P}_e^{(3)}(\gamma)$ is roughly

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

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

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

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

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 $\bigl\lfloor \alpha(\beta + \gamma) + \sum_i c_i\alpha_i \bigr\rceil_p = \bigl\lfloor \alpha\beta + \sum_i c_i\alpha_i \bigr\rceil_p$, and is zero otherwise. The second term is $\Pr\bigl[\mathcal{D}_{\mathbf{s}} = \beta\bigr]$, 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 $\gamma \in \mathbb{Z}_q^*$,

$$\begin{aligned}\mathsf{P}^{(2)}(0) - \mathsf{P}^{(2)}(\gamma) &\geq \Big(\mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V}) - q^2\eta\Big) \cdot \sum_{\beta \in \mathbb{Z}_q} \Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \cdot \big(\mathsf{Q}_\beta(0) - \mathsf{Q}_\beta(\gamma)\big) \\ &\geq \Big(\frac{\mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V})}{q^2} - \eta\Big) \cdot \sum_{\beta : \mathbb{1}(\beta,\gamma)=0} \frac{1}{q} \geq \Big(\frac{1 - 2\tau - q^2\eta}{q^3}\Big) = 10\zeta,\end{aligned}$$

where the second inequality on the second line holds since when $\gamma \neq 0$ there exists at least one β such that $\mathbb{1}(\beta,\gamma) = 0$. The second point follows. ■

3.6.3 Proving Lemma 10 Assuming Lemma 12

Lemma 10 (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 $\mathsf{P}(\mathbf{H},\mathbf{V}) \geq 1 - \tau$. 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 12 and Lemma 13, it suffices simply to show that for an overwhelming fraction of the $\mathbf{s} \in \mathsf{G}$ have $\Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \geq \frac{1}{q^2}$ 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 $\mathsf{P}_{\mathbf{s}}(\mathbf{H},\mathbf{V}) \geq 1 - 2\tau$, $\mathcal{D}_{\mathbf{s}}$ 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

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

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, $\Pr\big[\mathcal{D}_{\mathbf{s}} = \beta\big] \geq \frac{1}{q^2}$ holds for all $\beta \in \mathbb{Z}_q$ for an overwhelming fraction of the $\mathbf{s} \in \mathsf{G}$. Lemma 10 follows. ■

3.7 Proving Lemma 12

Notations. 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 3.4.2, we have a set $\mathsf{G} \subset \mathbb{Z}_q^n$ of "good secrets"; this set has size at least $|\mathsf{G}| \geq \delta q^n$ 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 $\mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1 - \eta$ for non-negligible $\eta > 0$. It was also shown in Claim 10 that for all subset $S \subset \mathbb{Z}_q^n$ of size $|S| = \rho q^n$, and non-negligible $\nu > 0$, $\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\big[\mathbf{a}' \in S\big] \leq \rho + \nu$. 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

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

for $\tau = 8n^4q^2\eta^{1/3t}$, where $t \in \mathbb{N}$ is a new parameter; it is the minimal integer such that $t \geq \frac{\log_q(1/\delta)+2}{3c}$ 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 derived from f as follows. Most of the time, if given $\mathbf{a} \in \mathbb{Z}_q^n$, h simply draws $b \sim \mathbb{Z}_q$ uniformly, computes $(\mathbf{a}', b') = f(\mathbf{a}, b)$ and outputs $\mathbf{a}'$. However, we will occasionally need to assume that h uses previously drawn values of b to produce a new b, rather than drawing $b \sim \mathbb{Z}_q$ fresh each time. For example, in this section we will be interested in the experiment which draws $\mathbf{a}_0, \mathbf{a}_1 \sim \mathbb{Z}_q^n$, $(\alpha_0, \alpha_1) \sim \mathbb{Z}_q^2 \setminus \{(0,0)\}$, then sets $\mathbf{a}_2 = \alpha_0\mathbf{a}_0 + \alpha_1\mathbf{a}_1$ and computes $\mathbf{a}'_j = h(\mathbf{a}_j)$ for $j = 0, 1, 2$. The computations of h in this context will draw $b_0, b_1 \sim \mathbb{Z}_q$ and then set $b_2 = \alpha_0 b_0 + \alpha_1 b_1$, rather than drawing $b_2 \sim \mathbb{Z}_q$. It will be considerably simpler to work with h rather than f. The non-degeneracy property for h says that for all $S \subset \mathbb{Z}_q^n$ of size $|S| = \rho q^n$, and non-negligible $\nu > 0$, $\Pr_{\mathbf{a} \sim \mathbb{Z}_q}\big[h(\mathbf{a}) \in S\big] \leq \rho + \nu$.

3.7.1 Recovering V

The Algorithm to Recover 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 $(\alpha_{i,0}, \alpha_{i,1}) \sim \mathbb{Z}_q^2 \setminus \{(0,0)\}$;

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

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

Proof. Let $\nu > 0$ be such that $\nu^{3t} = \eta$. Consider an execution of $\mathcal{D}_t$; for $i = 0, \ldots, t$, 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 $r \in \{1, \ldots, t-1\}$ be maximal such that $\Pr\big[d_r = 3r\big] \geq \nu^r$. We have

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

and so $\Pr\big[d_{r+1} < 3(r+1)\big|d_r = 3r\big] \geq 1 - \nu$. Let $\mathbf{a}_0, \mathbf{a}_1 \in \mathbb{Z}_q^n$ and $(\alpha_0, \alpha_1) \in \mathbb{Z}_q^2 \setminus \{(0,0)\}$ 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 $1 - 3\nu - q^{-\Omega(n)} \geq 1 - 4\nu$, and so

$$\mathsf{P}(\mathbf{V}_r) = \Pr_{\substack{\mathbf{a}_0,\mathbf{a}_1 \sim \mathbb{Z}_q^n \\ (\alpha_0,\alpha_1)\sim\mathbb{Z}_q^2\setminus\{(0,0)\}}}\Big[h(\alpha_0\mathbf{a}_0 + \alpha_1\mathbf{a}_1) \in \mathrm{Span}\big(\{h(\mathbf{a}_0), h(\mathbf{a}_1)\}\big) + \mathbf{V}_r\Big] \geq 1 - 4\nu.$$

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

Claim 14 *Let notations be as above. Then* $\Pr_{\mathcal{D}_t}\big[\dim(\mathbf{V}) = 3t\big] < \eta^{1/3}$.

Remark 15 *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 $i = 1, \dots, t$:

- $(\mathbf{a}_{i,0}, b_{i,0}), (\mathbf{a}_{i,1}, b_{i,1}) \sim \mathbb{Z}_q^n \times \mathbb{Z}_q$;
- $(\alpha_{i,0}, \alpha_{i,1}) \sim \mathbb{Z}_q^2 \setminus \{(0,0)\}$ 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:

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

Note that all three events cannot occur simultaneously. Indeed, the events $\mathbf{E}_2$ and $\mathbf{E}_3$ together imply that $^{\#}\{\mathbf{s} \in S : \mathbf{s}' \in S'\} \geq (1 - \sqrt{3tq\eta}) \cdot q^{-2t-1} \cdot |\mathsf{G}| \geq \frac{1}{2} \cdot q^{-2t-1} \cdot |\mathsf{G}|$, 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

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

which violates property 3 of G since it means some $\mathbf{s}' \in S'$ has $^{\#}\{\mathbf{s} \in S : \mathsf{p}_{(\mathbf{s},\mathbf{s}')} \geq 1 - \eta\} \geq 2$. 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, the following two claims complete the proof.

Claim 15 $\Pr_{\mathcal{D}}\big[\mathbf{E}_2\big] > 1 - q^{-n/3}$.

Proof. Recall $\mathbf{E}_2$ is the event that $|S| \geq q^{-2t-1} \cdot |\mathsf{G}|$. 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

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

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

$$\Pr_{(\mathbf{a},b)\sim\mathbb{Z}_q^n\times\mathbb{Z}_q}\Big[\Pr_{\mathbf{s}\sim\mathsf{G}_{r-1}}\big[b = \langle \mathbf{a}, \mathbf{s}\rangle\big] \geq q^{-1-1/2t}\Big] \geq 1 - q^{-n/2} \tag{3.4}$$

holds. This proves the claim as it gives $\Pr_{\mathcal{D}}\big[\mathbf{E}_2\big] \geq \big(1 - q^{-n/2}\big)^{2t} > 1 - q^{-n/3}$, so it remains to prove (3.4). For $b \in \mathbb{Z}_q$, let

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

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

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

We have

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

proving (3.4). ■

Claim 16 $\Pr_{\mathcal{D}}\big[\mathbf{E}_3\big] \geq 1 - \sqrt{3tq\eta}$.

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

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

by the union bound. ■ ■

3.7.2 Recovering H.

In the previous section we showed how to recover a constant dimensional subspace $\mathbf{V} \subset \mathbb{Z}_q^n$ such that $\mathsf{P}(\mathbf{V}) \geq 1 - 4\nu$, where $\nu = \eta^{1/3t}$. Here, we show how to use h such that $\mathsf{P}(\mathbf{V}) \geq 1 - 4\nu$ holds, to recover $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$ such that $\mathsf{P}(\mathbf{H}, \mathbf{V}) \geq 1 - \tau$ holds where $\tau = 8n^4q^2\nu$. This completes the proof of Lemma 12, and thus also the proof of Theorem 3.

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

$$\Pr_{\alpha_1,\dots,\alpha_n\sim\mathbb{Z}_q}\Big[h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_n\mathbf{a}_n) \in \mathrm{Span}(\alpha_1\mathbf{a}_1' + \cdots + \alpha_n\mathbf{a}_n') + \mathbf{V}\Big] \geq 1 - \tau. \qquad (3.5)$$

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 $i = 1, \dots, n$; $\mathsf{P}(\mathbf{H}, \mathbf{V}) \geq 1 - \tau$ follows from (3.5).

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 $\mathbf{a}_1, \dots, \mathbf{a}_n \sim \mathbb{Z}_q^n$ uniformly such that $\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ is linearly independent.
2. For $i = 1, \dots, n$, set $\mathbf{a}_i' = \lambda_i h(\mathbf{a}_i)$ for scalars $\{\lambda_i\}_{i=1}^n$ computed as follows:
   - set $\lambda_1 = 1$;
   - for $i \geq 2$, 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 $i \in \{2, \dots, n\}$ with probability at least $1 - 4(n-1)q^2\nu$, since $\mathsf{P}(\mathbf{V}) \geq 1 - 4\nu$. 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 $1 - 4nq^2\nu$.

Establishing (3.5). Given $\{\mathbf{a}_i\}_{i=1}^n$ which are linearly independent, define the quantities $\mathsf{P}_r\big(\{\mathbf{a}_i\}\big)$ for $r = 3, \dots, n$ as

$$\mathsf{P}_r\big(\{\mathbf{a}_i\}\big) := \mathrm{Pr}_{\alpha_1,\dots,\alpha_r \sim \mathbb{Z}_q}\Big[h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_r\mathbf{a}_r) \in \mathrm{Span}(\alpha_1\mathbf{a}_1' + \cdots + \alpha_r\mathbf{a}_r') + \mathbf{V}\Big].$$

It remains to show that with good probability over $\{\mathbf{a}_i\}$, $\mathsf{P}_n\big(\{\mathbf{a}_i\}\big) \geq 1 - \tau$ holds. We will prove this using induction on r. The following claim is key to this argument; it gives us our base case and will also be crucial to our induction step. We prove this claim in Section 3.7.3.

Claim 17 *For all distinct $i, j \in \{2, \dots, n\}$, 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 $1 - 4q^2n^2\nu$ over $\{\mathbf{a}_i\}_{i=1}^n$.

Let us now see how to use Claim 17 to establish (3.5). We will show that $\mathsf{P}_r \geq 1 - 8r^2n^2q^2$ for all $r = 3, \dots, n$. We use induction; the base case of $r = 3$ follows immediately from Claim 17, so fix $r > 3$ and assume that $\mathsf{P}_{r-1} \geq 1 - 8(r-1)^2n^2q^2\nu$. Draw linearly independent $\{\mathbf{a}_i\}_{i=1}^n$ from $\mathbb{Z}_q^n$. Additionally, draw a non-zero $\vec{\alpha} = (\alpha_1, \dots, \alpha_r) \sim \mathbb{Z}_q^r \setminus \{\mathbf{0}\}$. We group the sum $\alpha_1\mathbf{a}_1 + \cdots + \alpha_n\mathbf{a}_n$ in two ways:

$$(\alpha_1\mathbf{a}_1 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1}) + \alpha_r\mathbf{a}_r = (\alpha_1\mathbf{a}_1 + \alpha_r\mathbf{a}_r) + (\alpha_2\mathbf{a}_2 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1}).$$

Consider what happens if the following things occur:

- $h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1}) \in \text{Span}(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}) + \mathbf{V}$;
- $h(\alpha_1\mathbf{a}_1 + \alpha_r\mathbf{a}_r) \in \text{Span}(\alpha_1\mathbf{a}'_1 + \alpha_r\mathbf{a}'_r) + \mathbf{V}$.
- $h(\alpha_r\mathbf{a}_r) \in \text{Span}(\alpha_r\mathbf{a}'_r) + \mathbf{V}$;
- $h(\alpha_2\mathbf{a}_2 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1}) \in \text{Span}(\alpha_2\mathbf{a}'_2 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}) + \mathbf{V}$.

Note the first and last events occur with probability $\mathsf{P}_{r-1}\big(\{\mathbf{a}_i\}\big)$ and $\mathsf{P}_{r-2}\big(\{\mathbf{a}_i\}\big)$ by the induction hypothesis; the middle two events occur with probability $1 - 8q^2n^2\nu$ by Claim 17. Moreover, note that when all four of these events occur $h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_r\mathbf{a}_r)$ is contained in

$$\Bigg(\text{Span}\Big(\big\{\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}, \alpha_r\mathbf{a}'_r\big\}\Big) + \mathbf{V}\Bigg) \cap \Bigg(\text{Span}\Big(\big\{\alpha_1\mathbf{a}'_1 + \alpha_r\mathbf{a}'_r, \mathbf{z}\big\}\Big) + \mathbf{V}\Bigg),$$

where $\mathbf{z} = h(\alpha_2\mathbf{a}_2 + \cdots + \alpha_{r-1}\mathbf{a}_{r-1})$. It follows that there exist scalars $A, B, A', B' \in \mathbb{Z}_q$ such that

$$A \cdot \big(\alpha_1\mathbf{a}'_1 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}\big) + B \cdot \alpha_r\mathbf{a}'_r \in A' \cdot (\alpha_1\mathbf{a}'_1 + \alpha_r\mathbf{a}'_r) + B' \cdot \mathbf{z} + \mathbf{V}.$$

Thus either $A' = A$ or else $\mathbf{a}'_1 \in \text{Span}\big(\big\{\alpha_2\mathbf{a}'_2 + \cdots + \alpha_{r-1}\mathbf{a}'_{r-1}, \mathbf{a}'_r, \mathbf{z}\big\}\big) + \mathbf{V}$, which happens only with negligible probability by non-degeneracy. Similarly, $A' = B$ except with negligible probability. It follows that except with probability $1 - 8rq^2n^4\nu$, $A = B$ and so

$$h(\alpha_1\mathbf{a}_1 + \cdots + \alpha_r\mathbf{a}_r) \in \text{Span}\big(\{\alpha_1\mathbf{a}'_1 + \cdots + \alpha_r\mathbf{a}'_r\}\big) + \mathbf{V}$$

as desired.

3.7.3 Proof of Claim 17

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 $1 - 4nq\nu$ (since $\mathsf{P}(\mathbf{V}) \geq 1 - 4\nu$). 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. We complete the proof of Claim 17 by showing the following both occur with good probability over $\{\mathbf{a}_i\}$

Point 1: for all $(\alpha_1, \alpha_i, \alpha_j) \in \mathbb{Z}_q^3 \setminus \{\mathbf{0}\}$, and for all $i, j \in \{2, \ldots, n\}$,

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

Point 2: every π_i is the identity function.

Point 1 when $\alpha_1 = \alpha_j = 0$. 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 $1 - 4\nu$. This means that either

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

The latter happens with negligible probability since h is non-degenerate. Thus, $h(\alpha_i \mathbf{a}_i) \in \mathrm{Span}(\mathbf{a}'_i) + \mathbf{V}$ holds simultaneously for all $i \in \{2, \dots, n\}$ and $\alpha_i \in \mathbb{Z}_q$ with probability at least $1 - 4qn\nu$ over $\{\mathbf{a}_i\}$.

Point 1 when $\alpha_1 = 1$. 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 $1 - 8\nu$. In case $h(\mathbf{a}_1 + \alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j)$ is in the intersection, there exist scalars $A, B, A', B' \in \mathbb{Z}_q$ such that

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

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 for all $i, j \in \{2, \dots, n\}$ and $\alpha_i, \alpha_j \in \mathbb{Z}$ with probability at least $1 - 8q^2n^2\nu$ over $\{\mathbf{a}_i\}$.

Point 1 when $\alpha_1 = 0$. 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 $1 - 4\nu$ over $\{\mathbf{a}_i\}$. 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 $1 - 8\nu$ by the previous part. Thus, with probability at least $1 - 12\nu$, 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 $h(\alpha_i \mathbf{a}_i + \alpha_j \mathbf{a}_j) \in \operatorname{Span}\big(\{\pi_i(\alpha_i)\mathbf{a}_i' + \pi_j(\alpha_j)\mathbf{a}_j'\}\big) + \mathbf{V}$ holds for all $i, j \in \{2, \ldots, n\}$ and $\alpha_i, \alpha_j \in \mathbb{Z}_q$ with probability $1 - 12q^2n^2\nu$ over $\{\mathbf{a}_i\}$.

Point 2. We prove that $\pi_i(\alpha_i) = \alpha_i$ for all $i = 2, \ldots, n$ and $\alpha_i \in \mathbb{Z}_q$ by induction on α_i. We have already seen that $\pi_i(0) = 0$ and $\pi_i(1) = 1$ for all i. So assume $\pi_i(\alpha_i - 1) = \alpha_i - 1$, and write $\mathbf{a}_1 + \alpha_i\mathbf{a}_i + \mathbf{a}_j$ in three different ways:

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

With probability $1 - 12\nu$ over $\{\mathbf{a}_i\}$, $h(\mathbf{a}_1 + \alpha_i\mathbf{a}_i + \mathbf{a}_j)$ is contained in:

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

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

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

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 $\pi_i(\alpha_i) = \alpha_i$, as desired. We incurred a loss of 12ν to take a single step in the induction. Therefore, $\pi_i(\alpha_i) = \alpha_i$ for all $i \in \{2, \ldots, n\}$ and $\alpha_i \in \mathbb{Z}_q$ occurs with probability at least $1 - 12nq\nu$.

Point 1. Assume $\alpha_1 \neq 0$ since we have already handled this case above. Writing

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

we see that with probability at least $1 - 12\nu$ over $\{\mathbf{a}_i\}$, $h(\alpha_1\mathbf{a}_1 + \alpha_i\mathbf{a}_i + \alpha_j\mathbf{a}_j)$ is contained in

$$\mathrm{Span}\big(h(\mathbf{a}_1 + \alpha_1^{-1}\alpha_i\mathbf{a}_i + \alpha_1^{-1}\alpha_j\mathbf{a}_j)\big) + \mathbf{V} = \mathrm{Span}\big(\mathbf{a}'_1 + \alpha_1^{-1}\alpha_i\mathbf{a}'_i + \alpha_1^{-1}\alpha_j\mathbf{a}'_j\big) + \mathbf{V},$$

as desired. We have used point 2. ■

Chapter 4

A High Dimensional Goldreich-Levin Theorem with Low Agreement

In this work we prove a high dimensional analogue of the celebrated Goldreich-Levin Theorem (STOC'89). We consider the following algorithmic problem: given oracle access to a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ ($m, n, q \in \mathbb{N}$ such that q is prime) such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon$ for some matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\varepsilon > 0$, recover $\mathbf{A}$. We focus is on the case when $\varepsilon \leq 1/q$ since when $\varepsilon \geq 1/q + \delta$, the problem is solved by the original Goldreich-Levin Theorem. As stated, this problem cannot be efficiently solved when $\varepsilon \leq 1/q$ since the set of $\mathbf{A}$ with good agreement with f might be exponentially large. However, in this work we give an algorithm which efficiently recovers an *approximation* of $\mathbf{A}$; that is, a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n \times m}$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\mathbf{A}\mathbf{x} = \mathbf{A}'\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$. Our result extends a line of work relating to list-decoding group homomorphism codes, and also has applications to approximate list-decodable codes and effective algebraic property testing; we elaborate more in the next section.

4.1 Introduction

The celebrated Goldreich-Levin Theorem ([GL89]) is a cornerstone theorem in theoretical computer science. It yielded fundamental applications in cryptography ([Blu83], [HILL99]), led to the development of new categories of error-correcting codes ([Sud97a],[KT00]), and launched boolean learning theory ([KM93]). The technical core of the Goldreich-Levin theorem is the "prediction implies inversion" lemma which states that a function $f : \{0,1\}^n \to \{0,1\}$ which predicts random inner products with a secret $\mathbf{y} \in \{0,1\}^n$, with *any advantage at all over guessing randomly*, must "know" $\mathbf{y}$, in the sense that $\mathbf{y}$ can be recovered efficiently given oracle access to f. This lemma has been generalized in many different ways. One line of follow up work proves prediction implies inversion lemmas for general group homomorphisms $f : G \to H$ ([GKS06],[DGKS08],[BBW18]). Another work proves a degree 2 analogue using quadratic Fourier analysis ([TW14]). In this work we generalize the Goldreich-Levin theorem to higher dimensions.

4.1.1 Our Contributions

We consider a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ for integers $n, m, q \in \mathbb{N}$ with q prime which has the following linear agreement guarantee:

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon, \tag{4.1}$$

for some $\varepsilon > 0$ and matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. We ask whether, given oracle access to such a function, it is possible to efficiently recover $\mathbf{A}$. More precisely, and in the list decoding spirit of [GL89], we ask whether it is possible to efficiently output a list $\mathsf{L} = \{\mathbf{A}_1, \dots, \mathbf{A}_\ell\}$ such that any $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ which has good agreement with f is in L. When $\varepsilon \geq 1/q + \delta$, list decoding algorithms from prior work ([GL89, DGKS08]), indeed recover such a list. However, these algorithms fail when $\varepsilon \leq 1/q$. Actually, when $\varepsilon \leq 1/q$, the problem is not possible as stated, since the list might be exponentially large (and so cannot be efficiently recovered). For example, suppose that f always outputs $\mathbf{Ax}$ except for the first coordinate, which it chooses randomly. In other words,

$$f(\mathbf{x})_i = \begin{cases} \$ \sim \mathbb{Z}_q, & i = 1 \\ (\mathbf{Ax})_i, & i \geq 2 \end{cases}$$

where $(\mathbf{Ax})_i$ denotes the $i-$th coordinate of $\mathbf{Ax}$. Clearly, in this case $f(\mathbf{x}) = \mathbf{Ax}$ occurs with probability $1/q$. However, $f(\mathbf{x}) = \mathbf{A}'\mathbf{x}$ also holds with probability $1/q$ for any $\mathbf{A}'$ whose final $n-1$ rows are the same as those in $\mathbf{A}$. Indeed, the function f possesses no information about the first row of $\mathbf{A}$, so any matrix which equals $\mathbf{A}$ outside of the first row (there are q^m such matrices) will have just as good agreement with f as $\mathbf{A}$ does.

So to summarize, the algorithmic question above is solved by prior work when $\varepsilon > 1/q$ and is impossible when $\varepsilon \leq 1/q$. This is unfortunate because the setup is very natural when $\varepsilon \leq 1/q$. When $n = 1$, the barrier of $1/q + \delta$ makes sense conceptually since a random function (from which no secret can be extracted) will have agreement $1/q$. So the original (one-dimensional) Goldreich-Levin theorem promises that a secret can be extracted from any function which has a prediction advantage over guessing randomly. In higher dimension, a random function will agree with a linear function with probability

$q^{-n} \ll 1/q$ and so one might hope that some information about $\mathbf{A}$ would be recoverable from a function with agreement probability $\varepsilon > q^{-n}$.

The problem is that we have asked the wrong question. Rather than aiming to recover a matrix $\mathbf{A}$ which is guaranteed to have good agreement with f, we should try to *approximately* recover $\mathbf{A}$; that is, recover a matrix $\mathbf{A}'$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{A}'\mathbf{x} = \mathbf{A}\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$ when $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon$ ($q^{-n} < \varepsilon \leq q^{-1}$). In this work, we give an algorithm which recovers such an approximation of a matrix $\mathbf{A}$ that has low agreement with f, and characterize the conditions under which our algorithm fails. We informally state our main theorem below.

Theorem 3 (Informal) *Let $m, n, q \in \mathbb{N}$ be parameters with q prime. Let $\varepsilon > 0$ be such that $\varepsilon \geq q^{-c}$ for a constant $c > 0$. Let $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ be a function. There exists a randomized oracle algorithm $\mathcal{A}$ which runs in time $\mathsf{poly}\big(m, n, q, 1/\varepsilon\big)$ and with probability $\mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$ outputs a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{A}'\mathbf{x} = \mathbf{A}\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$, unless f satisfies a conditional affine linearity test.*

For ease of exposition, we have not described the conditional affine linearity test described above in the informal theorem statement (see Theorem 4 for the formal theorem statement). In slightly more detail, we'll show that the only way in which our algorithm fails to recover an approximation of the matrix $\mathbf{A}$ is if there exists a subset $S \subset \mathbb{Z}_q^m$ of density $|S|q^{-m} \geq \varepsilon$ and a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ of dimension $d = \mathcal{O}(1)$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W} \;\big|\; \mathbf{x} \in S\big] > 1/q$ and

$$\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[\phi(\mathbf{x}+\mathbf{y}) = \phi(\mathbf{x}) + \phi(\mathbf{y}) \;\big|\; \mathbf{x}, \mathbf{y}, \mathbf{x}+\mathbf{y} \in T\big] \geq 1 - \mathsf{poly}(1/m, 1/n, 1/q, \varepsilon),$$

where $T = \{\mathbf{x} \in S : f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}\}$ and $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^d$ is a function which, when restricted to T, is the projection of $f(T)$ onto $\mathbf{W} \cong \mathbb{Z}_q^d$. Again, for the purpose of this introduction we have presented an oversimplified narrative of this condition; see Theorem 4 for the formal theorem statement.

4.1.2 Applications

Approximate List Decodable Codes. *Code amplification schemes* are coding schemes which, when concatenated with an error-correcting code, yield a new code with better parameters. Code amplification schemes are relaxations of standard ECCs because they assume that their message space has a notion of distance (since the messages are codewords from the original ECC). Code amplification is an extremely useful technique; many of the best codes we have today are built in this way ([ABN$^+$92],[Ta-17]). Similarly, approximate list decoding assumes a notion of distance for the message space and given a corrupted codeword, an algorithm can recover a list of messages such that any valid codeword which is close to the corrupted codeword is encoding a message which is close to something in the list.

Approximate LDCs were introduced explicitly in [IJK09] and [IJKW10], where it is shown how to use an approximate list decoding algorithm for a direct product code to prove hardness amplification theorems. Our work falls under the jurisdiction of approximate LDCs, with our message space being $\mathbb{Z}_q^{n\times m}$ and the codeword being the truth table encoding.

Effective Algebraic Property Testing. Our work fits into the research landscape on effective property testing. The linearity tests of [BLR93] and [BCH$^+$96] promise that if

$f : \{0,1\}^n \to \{0,1\}$ satisfies $f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x})+f(\mathbf{y})$ with good probability over $\mathbf{x}, \mathbf{y} \sim \{0,1\}^n$ then there exists a linear map $\varphi : \{0,1\}^n \to \{0,1\}$ such that f has good agreement with φ. In the high dimensional and large modulus setting, linearity tests are much harder to prove. Samorodnitsky ([Sam07]) showed using methods from additive combinatorics that if $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ passes this linearity test with probability ε then it is ε'−close to a linear function where ε' depends exponentially on n. A breakthrough result of Sanders ([San12]) obtains a better (quasi-logarithmic) relationship between ε' and ε. The holy grail of this area would be a proof that ε' depends polynomially on ε. This is known to follow from the polynomial Freiman-Ruzsa conjecture in additive combinatorics. Our work makes any high dimensional linearity testing theorem effective by offering an algorithm which would recover the linear map which is close to f (whose existence would be ensured by the linearity testing theorem).

4.2 Preliminaries

Basic Notation. If $n \in \mathbb{N}$, then we denote by $[n]$ the set $\{1, \ldots, n\}$. For a prime $q \in \mathbb{N}$, we denote by $\mathbb{Z}_q$ the field of integers modulo q. We will denote scalars, vectors and matrices with lowercase italic, lowercase bold, and uppercase bold respectively (*e.g.*, $z \in \mathbb{Z}_q$, $\mathbf{z} \in \mathbb{Z}_q^n$ and $\mathbf{Z} \in \mathbb{Z}_q^{n \times m}$). For a distribution $\mathcal{D}$ (resp. set D), we write $r \sim \mathcal{D}$ (resp. $r \sim D$) to indicate that the random variable r is drawn according to $\mathcal{D}$ (resp. the uniform distribution on D). For an event $\mathbf{E}$, we denote by $\mathbb{1}_{\mathbf{E}}$ the indicator random variable corresponding to $\mathbf{E}$. Namely, $\mathbb{1}_{\mathbf{E}} = 1$ (resp. $\mathbb{1}_{\mathbf{E}} = 0$) when $\mathbf{E}$ occurs (resp. does not occur).

Linear Algebra. Unless otherwise specified, every vector space in this work is over the field $\mathbb{Z}_q$, for a prime $q \in \mathbb{N}$. If $\mathbf{W} \subset \mathbb{Z}_q^n$ is a subspace, then we denote by $\mathbf{W}^\perp \subset \mathbb{Z}_q^n$ the set $\{\mathbf{z} \in \mathbb{Z}_q^n : \langle \mathbf{z}, \mathbf{w} \rangle = 0 \; \forall \; \mathbf{w} \in \mathbf{W}\}$, where $\langle \mathbf{z}, \mathbf{w} \rangle = z_1 w_1 + \cdots + z_n w_n$ is the dot product. It is known that $\mathbf{W}^\perp \subset \mathbb{Z}_q^n$ is a subspace of dimension $n - d$, where $d = \dim(\mathbf{W})$. Given a matrix $\mathbf{Z} \in \mathbb{Z}_q^{n \times m}$, we denote by $\mathbf{Z}^{\mathrm{t}} \in \mathbb{Z}_q^{m \times n}$ its transpose.

Roots of Unity. Let $q \in \mathbb{N}$ be prime. A q^{th} *root of unity in* $\mathbb{C}$ is a complex root of the monic degree q polynomial $f(X) = X^q - 1$. It is a well known fact that there exist exactly q roots of f, which are the elements $\{e^{-2\pi i k/q}\}_{k=0}^{q-1}$, where i is the imaginary unit. The q^{th} roots of unity in $\mathbb{C}$ form a multiplicative cyclic group of order q, which is isomorphic to $\mathbb{Z}_q$. A generator of the group is called a *primitive* q^{th} *root of unity in* $\mathbb{C}$. The subset of primitive q^{th} roots of unity in $\mathbb{C}$ form a subgroup of order $q - 1$, which is isomorphic to $\mathbb{Z}_q^*$.

Recall that if $\omega \in \mathbb{C}$ is a primitive q^{th} root of unity, then $\sum\limits_{k=0}^{q-1} \omega^k = 0$. Also, it is a well-known fact that for any subspace $\mathbf{W} \subset \mathbb{Z}_q^n$, $\mathbb{E}_{\mathbf{z} \sim \mathbf{W}^\perp}\left[\omega^{\langle \mathbf{z}, \mathbf{w} \rangle}\right] = \mathbb{1}_{\mathbf{w} \in \mathbf{W}}$.

Discrete Fourier Analysis on Finite Abelian Groups. Here we briefly summarize several key results in discrete Fourier analysis on finite abelian groups. Actually, we specialize our results to the group $\mathbb{Z}_q^m$ ($m, q \in \mathbb{N}$ such that q is prime), which is also a vector space over $\mathbb{Z}_q$ of dimension n. See [Luo09] for a complete treatment on general finite abelian groups. Let $\mathcal{F} = \{f : \mathbb{Z}_q^m \to \mathbb{C}\}$ be the set of all functions $f : \mathbb{Z}_q^m \to \mathbb{C}$, which is a $\mathbb{C}-$vector space. A *character* of $\mathbb{Z}_q^m$ is a group homomorphism $\chi : \mathbb{Z}_q^m \to \mathbb{C}^*$. It turns out that each character χ of $\mathbb{Z}_q^m$ is characterized by an element $\mathbf{u} \in \mathbb{Z}_q^m$, and is defined by $\chi_{\mathbf{u}}(\mathbf{x}) = \omega^{\langle \mathbf{u}, \mathbf{x} \rangle}$, where $\omega \in \mathbb{C}$ is a primitive q^{th} root of unity in $\mathbb{C}$. Moreover, the set

$\{\chi_{\mathbf{u}} : \mathbf{u} \in \mathbb{Z}_q^m\} \subset \mathcal{F}$ of characters of $\mathbb{Z}_q^m$ form a $\mathbb{C}-$basis for $\mathcal{F}$, hence $\mathcal{F}$ is a $\mathbb{C}-$vector space of finite dimension q^m. If $f \in \mathcal{F}$ and $\mathbf{u} \in \mathbb{Z}_q^m$, then we define the *Fourier coefficient of f at* $\mathbf{u}$ as $\hat{f}(\mathbf{u}) := \mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x})\bar{\chi}_{\mathbf{u}}(\mathbf{x})\big] = \mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x})\omega^{-\langle\mathbf{u},\mathbf{x}\rangle}\big]$. It turns out that for every $f \in \mathcal{F}$, f can be expressed as a unique linear combination of the characters of $\mathbb{Z}_q^m$ over $\mathbb{C}$ of the form $f = \sum_{\mathbf{u}\in\mathbb{Z}_q^m} \hat{f}(\mathbf{u})\chi_{\mathbf{u}}$. Parseval's identity, an extremely useful tool in application, relates the expectation of the norm of $|f(\cdot)|^2$ to the sum over its Fourier coefficients of their squared norm:

$$\mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\Big[\big|f(\mathbf{x})\big|^2\Big] = \sum_{\mathbf{u}\in\mathbb{Z}_q^m} \big|\hat{f}(\mathbf{u})\big|^2.$$

4.2.1 The Goldreich-Levin Theorem

The Goldreich-Levin Theorem [GL89] gives an algorithmic procedure to efficiently recover a secret vector $\mathbf{y} \in \mathbb{Z}_q^m$ given oracle access to a function which predicts random inner products of $\mathbf{y}$ with probability noticeably better than guessing. The original result of [GL89] was proved in characteristic 2. Indeed, the result holds for prime characteristic q as well. In our work, we require the characteristic q version, stated below; see Appendix A for a proof.

Lemma 16 (Goldreich-Levin Theorem) *Let $q, m \in \mathbb{N}$ such that q is prime, and let $\mathbf{s} \in \mathbb{Z}_q^m$ a secret vector. Suppose there exists a function $\mathsf{Pred} : \mathbb{Z}_q^m \to \mathbb{Z}_q$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}(\mathbf{x}) = \langle\mathbf{s},\mathbf{x}\rangle\big] \geq 1/q + \varepsilon$, for $\varepsilon > 0$. Then, there exists an algorithm Inv, running in time $\mathsf{poly}(m, q, 1/\varepsilon)$, which gets oracle access to Pred and outputs $\mathbf{s}^* \in \mathbb{Z}_q^m$ such that*

$$\Pr_{\mathbf{s}^*\sim\mathsf{Inv}^{\mathsf{Pred}}}\big[\mathbf{s}^* = \mathbf{s}\big] \geq \frac{\varepsilon^2}{8m^3q^5}.$$

4.2.2 Statistical Sampling

Here we show that for a high fraction of $r-$planes $\mathbf{V} \subset \mathbb{Z}_q^m$, choosing a random point $\mathbf{v} \sim \mathbf{V}$ is statistically close to choosing a random vector $\mathbf{v} \sim \mathbb{Z}_q^m$.

Lemma 17 (Plane-vs-Point Sampler) *Let $r, m, q \in \mathbb{N}$ be integers with q prime, and let $\mathcal{V}$ be the set of $r-$dimensional subspaces of $\mathbb{Z}_q^m$. Then, for every subset $S \subset \mathbb{Z}_q^m$ of size $|S| = \lambda q^m$ and $\varepsilon_1 > 0$ such that $\varepsilon_1 < \lambda$,*

$$\Pr_{\mathbf{V}\sim\mathcal{V}}\left[\left|\Pr_{\mathbf{v}\sim\mathbf{V}}\left[\mathbf{v} \in S\right] - \lambda\right| > \varepsilon_1\right] \leq 2\varepsilon_1^{-2} q^{-(r-1)}.$$

Proof. We have

$$\Pr_{\mathbf{V}\sim\mathcal{V}}\left[\left|\Pr_{\mathbf{x}\sim\mathbf{V}}\left[\mathbf{x} \in \mathsf{S}\right] - \lambda\right| \geq \varepsilon_1\right] \leq \frac{1}{\varepsilon_1^2} \cdot \left[\mathbb{E}_{\mathbf{V}\sim\mathcal{V}}\left[\Pr_{\mathbf{x}\sim\mathbf{V}}\left[\mathbf{x} \in \mathsf{S}\right]^2\right] - \lambda^2\right],$$

by Markov's inequality, since $\lambda = \mathbb{E}_{\mathbf{V}\sim\mathcal{V}}\left[\Pr_{\mathbf{x}\sim\mathbf{V}}[\mathbf{x} \in \mathsf{S}]\right]$. We complete the proof by showing that $\mathbb{E} := \mathbb{E}_{\mathbf{V}\sim\mathcal{V}}\left[\Pr_{\mathbf{x}\sim\mathbf{V}}[\mathbf{x} \in \mathbf{V}]^2\right] \leq \lambda^2 + 2q^{-(r-1)}$ using a pairwise independence argument. Let $\mathbb{1}_{\mathsf{S}}$ denote the indicator function for S. We have

$$\mathbb{E} = \mathbb{E}_{\substack{\mathbf{V}\sim\mathcal{V} \\ \mathbf{x}_1,\mathbf{x}_2\sim\mathbf{V}}}\left[\mathbb{1}_{\mathsf{S}}(\mathbf{x}_1) \cdot \mathbb{1}_{\mathsf{S}}(\mathbf{x}_2)\right] \leq \mathbb{E}_{\mathbf{x}_1,\mathbf{x}_2\sim\mathbb{Z}_q^m}\left[\mathbb{1}_{\mathsf{S}}(\mathbf{x}_1) \cdot \mathbb{1}_{\mathsf{S}}(\mathbf{x}_2)\right] + 2q^{-(r-1)} = \lambda^2 + 2q^{-(r-1)},$$

since $2q^{-(r-1)}$ is an upper bound on the probability that two randomly selected vectors from an $r-$plane are linearly dependent. ■

Finally, we show that if $S \subset \mathbb{Z}_q^m$ is a subset of density λ, then when we sample a random basis $\{\mathbf{x}_1, \ldots, \mathbf{x}_r\} \subset \mathbb{Z}_q^m$ for an $r-$plane and a random point $\mathbf{v} \sim \text{Span}(\{\mathbf{x}_1, \ldots, \mathbf{x}_r\})$, then $\{\mathbf{x}_1, \ldots, \mathbf{x}_r, \mathbf{v}\} \subset S$ with probability close to λ^{r+1}.

Corollary 2 *Let $r, m, q \in \mathbb{N}$ such that q is prime, and let $\mathsf{S} \subset \mathbb{Z}_q^m$ be a subset of density λ. Then,*

$$\Pr_{\substack{\mathbf{x}_1,\ldots,\mathbf{x}_r \sim \mathbb{Z}_q^m \\ \alpha_1,\ldots,\alpha_r \sim \mathbb{Z}_q}} \Big[\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i \in S\Big] \geq \frac{\lambda^{r+1}}{2} - \Big(\frac{4}{\lambda q^{r-1}} + 2^{-\Omega(m)}\Big).$$

Proof. We'll say that a subset $\{\mathbf{x}_1, \ldots, \mathbf{x}_r\} \subset \mathbb{Z}_q^m$ is *good* if the following properties hold:

1. $\{\mathbf{x}_1, \ldots, \mathbf{x}_r\}$ is linearly independent.
2. $\{\mathbf{x}_1, \ldots, \mathbf{x}_r\} \subset S$.
3. $\Pr_{\alpha_1,\ldots,\alpha_r \sim \mathbb{Z}_q}\Big[\sum\limits_i \alpha_i \mathbf{x}_i \in S\Big] \geq \lambda/2$.

By Lemma 17, it follows that

$$\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_r \sim \mathbb{Z}_q^m}\big[\{\mathbf{x}_1, \ldots, \mathbf{x}_r\} \text{ good}\big] \geq \lambda^r - \Big(\frac{4}{\lambda q^{r-1}} + 2^{-\Omega(m)}\Big).$$

Thus

$$\Pr_{\substack{\mathbf{x}_1,\ldots,\mathbf{x}_r \sim \mathbb{Z}_q^m \\ \alpha_1,\ldots,\alpha_r \sim \mathbb{Z}_q}} \Big[\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i \in S\Big] \geq \Big(\lambda^r - \Big(\frac{4}{\lambda q^{r-1}} + 2^{-\Omega(m)}\Big)\Big) \cdot \frac{\lambda}{2}$$
$$\geq \frac{\lambda^{r+1}}{2} - \Big(\frac{4}{\lambda q^{r-1}} + 2^{-\Omega(m)}\Big),$$

as desired. ■

4.2.3 Computing a Linearly Independent Subset with a Property

Here we show how to algorithmically compute a linearly independent subset of $\mathbb{Z}_q^m$ ($m, q \in \mathbb{N}$ such that q is prime) with high probability.

Lemma 18 *Let $m, q \in \mathbb{N}$ such that q is prime, $S \subset \mathbb{Z}_q^m$ an efficiently recognizable*[1] *subset of density $\lambda := |S|q^{-m} \geq q^{-c}$ ($c \in \mathbb{N}$), and $k \in \mathbb{N}$ such that $k \leq m - (2c+1)$. Then, there exists an algorithm running in time $\mathsf{poly}(m, q)$ which with probability $1 - 2^{-\Omega(m)}$ over its randomness outputs a linearly independent subset $B \subset \mathbb{Z}_q^m$ of size $|B| = k$. Furthermore, if $T \subset S$ is a subset of density $1 - \varepsilon$, for $\varepsilon > 0$, then $B \subset T$ with probability $1 - k(2\varepsilon + \lambda/2)$.*

Proof. The algorithm $\mathcal{A}_S$ of the lemma works as follows:

1. Initialize $B := \emptyset$.
2. While $|B| < k$, choose $\mathbf{x} \sim \mathbb{Z}_q^m$; if $\mathbf{x} \in S$ and $\mathbf{x} \notin \mathrm{Span}(B)$, then update $B := B \bigcup \{\mathbf{x}\}$; otherwise continue.
3. Output B.

We'll show that $\mathcal{A}_S$ outputs linearly independent $B \subset S$ with probability $1 - 2^{-\Omega(m)}$. First, note that $\mathrm{P}_{S,B} := \Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\left[\mathbf{x} \in S \text{ and } \mathbf{x} \notin \mathrm{Span}(B)\right] \geq \lambda - q^{-(m-k)} \geq \lambda/2$. Let $N = 108m/\lambda^2$. By the Chernoff Bound, we have that with probability $1 - 2^{-\Omega(m)}$ over $\mathbf{x}_1, \ldots, \mathbf{x}_N \sim \mathbb{Z}_q^m$, $\#\{i \in [N] : \mathbf{x}_i \in S \backslash \mathrm{Span}(B)\} \geq (2N/3)\mathrm{P}_{S,B} \geq 1$.

[1] By efficiently recognizable, we mean membership to the subset can be tested in time $\mathsf{poly}(m, q)$.

Let B be the set output by $\mathcal{A}_S$. We have

$$\begin{aligned}
&\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_k \sim \mathbb{Z}_q^m}\left[\{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset T \;\middle|\; \{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset S \text{ L.I.}\right] \\
&\geq \frac{\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_k \sim \mathbb{Z}_q^m}\left[\{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset T \text{ L.I.} \;\middle|\; \{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset S\right]}{\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_k \sim \mathbb{Z}_q^m}\left[\{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \text{ L.I.} \;\middle|\; \{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset S\right]} \\
&\geq (1-k\varepsilon)\cdot \frac{\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_k \sim \mathbb{Z}_q^m}\left[\{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \text{ L.I.} \;\middle|\; \{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset T\right]}{\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_k \sim \mathbb{Z}_q^m}\left[\{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \text{ L.I.} \;\middle|\; \{\mathbf{x}_1,\ldots,\mathbf{x}_k\} \subset S\right]} \\
&\geq (1-k\varepsilon)\cdot \frac{\prod_{i=1}^{k}\left(\lambda(1-\varepsilon) - q^{-(m-(i-1))}\right)}{\prod_{i=1}^{k}\left(\lambda - q^{-(m-(i-1))}\right)} \\
&= (1-k\varepsilon)\cdot \frac{\prod_i (1-\varepsilon - q^{-(m-(i-1))}/\lambda)}{\prod_i (1 - q^{-(m-(i-1))}/\lambda)} \\
&\geq (1-k\varepsilon)\cdot(1-(\varepsilon+\lambda/2))^k \geq (1-k\varepsilon)(1-k(\varepsilon+\lambda/2)) \geq 1-k(2\varepsilon+\lambda/2),
\end{aligned}$$

where the first inequality on the final line follows from $q^{-(m-(i-1))}/\lambda \leq q^{-(m-k)}/\lambda \leq \lambda/2 \leq 1$. ■

4.3 Main Theorem and Proof Overview

Here we outline our main theorem, which provides a high dimensional approximate version of the Goldreich-Levin Theorem with low agreement.

Theorem 4 *Let $m, n, q \in \mathbb{N}$, such that $q = \mathsf{poly}(m, n)$ is prime, and let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix. Suppose there exists a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ such that*

$$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon,$$

where $\varepsilon > 0$ such that $q^{-c} \leq \varepsilon \leq q^{-1}$, for a constant $c \in \mathbb{N}$. Then, there exists an algorithm $\mathcal{A}$ running in time $\mathsf{poly}(m, n, q, 1/\varepsilon, \mathsf{T}_f)$, where T_f is the running time of f, which gets oracle access to f and outputs with probability $\mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$ over its randomness either:

- *a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n \times m}$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\mathbf{A}'\mathbf{x} = \mathbf{A}\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$;*
- *or (the description of) a subset $S \subset \mathbb{Z}_q^m$ of density $|S|q^{-m} \geq \varepsilon$ and a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ of constant dimension $s \in \mathbb{N}$ for which there exists a superspace $\mathbf{W}^* \supset \mathbf{W}$ of constant dimension $d \in \mathbb{N}$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}^* \;\big|\; \mathbf{x} \in S\big] > 1/q$ and*

  $$\Pr_{\substack{\mathbf{x}_1, \ldots, \mathbf{x}_r \sim \mathbb{Z}_q^m \\ \alpha_1, \ldots, \alpha_r \sim \mathbb{Z}_q}}\left[\phi\Big(\sum_i \alpha_i \mathbf{x}_i\Big) = \sum_i \alpha_i \phi(\mathbf{x}_i) \;\middle|\; \mathbf{x}_1, \ldots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i \in T\right] \geq 1 - \gamma,$$

  where $T \subset S$ is the set of $\mathbf{x} \in S$ such that $f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}^*$, $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^{d-s}$ is a function which, when restricted to T, is the projection of $f(T)$ onto $\mathbf{W}^*/\mathbf{W} \cong \mathbb{Z}_q^{d-s}$, $r \in \mathbb{N}$ is a constant, and $\gamma = \mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$.

Remark. Theorem 4 states that if a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ has low agreement with a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, then there exists an algorithm which recovers an approximation of $\mathbf{A}$, unless f admits a function $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^{d-s}$ which satisfies the conditional affine linearity test of Theorem 8 in Chapter 2 Section 2.3. Let us further decompose this conclusion. First, let $\{\mathbf{w}_1, \ldots, \mathbf{w}_s\} \subset \mathbf{W}$ be a basis for $\mathbf{W}$, and extend this to a basis $\{\mathbf{w}_1, \ldots, \mathbf{w}_s, \mathbf{w}_{s+1}, \ldots, \mathbf{w}_d\} \subset \mathbf{W}^*$ for $\mathbf{W}^*$, where $\mathbf{w}_{s+1}, \ldots, \mathbf{w}_d \in \mathbf{W}^* \backslash \mathbf{W}$. Reindex $\{\mathbf{w}_{s+1}, \ldots, \mathbf{w}_d\}$ as $\{\mathbf{w}_1^*, \ldots, \mathbf{w}_{d-s}^*\}$. Clearly, $\{\mathbf{w}_1^* + \mathbf{W}, \ldots, \mathbf{w}_{d-s}^* + \mathbf{W}\} \subset \mathbf{W}^*/\mathbf{W}$ is a basis for $\mathbf{W}^*/\mathbf{W}$. Define $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ as a function which, when restricted to $T = \{\mathbf{x} \in S : f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}^*\}$, is the projection of $f(T)$ onto $\mathbf{W}^*/\mathbf{W}$. Let $\mathbf{M} \in \mathbb{Z}_q^{n \times (d-s)}$ be the matrix whose i^{th} column is $\mathbf{w}_i^*$. Observe that if $\mathbf{x} \in T$, then we can write $f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{M} \cdot \phi(\mathbf{x}) + \mathbf{W}$. If ϕ agreed with an affine map $(\mathbf{A}', \mathbf{b}') \in \mathbb{Z}_q^{(d-s) \times m} \times \mathbb{Z}_q^{d-s}$, then we could write $f(\mathbf{x}) \in (\mathbf{A} - \mathbf{M}\mathbf{A}')\mathbf{x} + \mathbf{b}' + \mathbf{W}$. Moreover if this agreement held with probability at least $1/q$ conditioned on $\mathbf{x} \in S$, then we claim this would effectively eliminate the second point of the conclusion of Theorem 4 (see Lemma 19 below). However, Theorem 8 concludes that this agreement is at least q^{-d}, which is too small in our case. If we could reduce our argument to the case in which $d = 1$, though, then this would suffice. The main barrier in currently achieving this is a restriction imposed on the parameters by our statistical sampling argument, which is further detailed in the proof of Lemma 19. An interesting follow-up question would be if there exists a method to circumvent this parameter restriction, or the sampling argument altogether, and allow us to reduce our argument to the case in which $d = 1$.

Our proof of Theorem 4 is divided into two main parts. First, suppose that $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ and $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ are such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon$. The first part constructs a subset $S \subset \mathbb{Z}_q^m$ and a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in$

$\mathbf{W} \mid \mathbf{x} \in S\big] \geq 1/q + \varepsilon'$, for some $\varepsilon' > 0$. Then, the second part utilizes Goldreich-Levin prediction-implies-inversion arguments together with standard linear algebraic techniques to recover an approximation of the matrix $\mathbf{A}$ (*i.e.*, a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ which has good agreement with $\mathbf{A}$). More concretely, the proof of Theorem 4 follows immediately from Lemmas 19 and 20 below by simply composing the algorithms $\mathcal{A}_1$ and $\mathcal{A}_2$ of the respective lemmas.

Lemma 19 *Let $m, n, q \in \mathbb{N}$, such that $q = \mathsf{poly}(m, n)$ is prime, and let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be a matrix. Suppose there exists a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ such that*

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon,$$

where $\varepsilon > 0$ such that $q^{-c} \leq \varepsilon \leq q^{-1}$, for a constant $c \in \mathbb{N}$. Then, there exists an algorithm $\mathcal{A}_1$ running in time $\mathsf{poly}(m, n, q, 1/\varepsilon, \mathsf{T}_f)$, where T_f is the running time of f, which gets oracle access to f and outputs with probability $\mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$ over its randomness matrices $(\mathbf{Z}, \mathbf{Y}) \in \mathbb{Z}_q^{n\times\ell} \times \mathbb{Z}_q^{m\times\ell}$ and a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ of dimension d, for constants $\ell, d, \in \mathbb{N}$, such that either:

- $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W} \mid \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] \geq q^{-1/3}$*;*
- *or there exists a superspace $\mathbf{W}^* \supset \mathbf{W}$ of constant dimension $d' \in \mathbb{N}$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}^* \mid \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] > 1/q$ and*

  $$\Pr_{\substack{\mathbf{x}_1,\ldots,\mathbf{x}_r\sim\mathbb{Z}_q^m \\ \alpha_1,\ldots,\alpha_r\sim\mathbb{Z}_q}}\left[\phi\Big(\sum_i \alpha_i\mathbf{x}_i\Big) = \sum_i \alpha_i\phi(\mathbf{x}_i) \;\middle|\; \mathbf{x}_1,\ldots,\mathbf{x}_r, \sum_i \alpha_i\mathbf{x}_i \in T\right] \geq 1-\gamma,$$

  where $T \subset \mathbb{Z}_q^m$ is the set of $\mathbf{x} \in \mathbb{Z}_q^m$ such that $\mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}$ and $f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}^*$, $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^{d'-d}$ is a function which, when restricted to T, is the projection of $f(T)$ onto $\mathbf{W}^*/\mathbf{W} \cong \mathbb{Z}_q^{d'-d}$, $r \in \mathbb{N}$ is a constant, and $\gamma = \mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$.

Lemma 20 *Let $m, n, q \in \mathbb{N}$, such that $q = \mathsf{poly}(m, n)$ is prime, and let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ be a matrix. Let $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ be a function such that*

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon,$$

where $\varepsilon > 0$ such that $q^{-c} \leq \varepsilon \leq q^{-1}$, for a constant $c \in \mathbb{N}$. Suppose there exist matrices $(\mathbf{Z}, \mathbf{Y}) \in \mathbb{Z}_q^{n\times\ell} \times \mathbb{Z}_q^{m\times\ell}$ and a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ of dimension d, for constants $\ell, d, \in \mathbb{N}$, such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W} \bigm| \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] \geq q^{-1/3}$. Then, there exists an algorithm $\mathcal{A}_2$ running in time $\mathsf{poly}(m, n, q, 1/\varepsilon, \mathsf{T}_f)$, where T_f is the running time of f, which takes as input $(\mathbf{Z}, \mathbf{Y}, \mathbf{W})$, gets oracle access to f, and outputs with probability $\mathsf{poly}(1/m, 1/n, 1/q, \varepsilon)$ over its randomness a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{A}'\mathbf{x} = \mathbf{A}\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$.

Note that Theorem 4 follows immediately from Lemmas 19 and 20. The remainder of this work is therefore devoted to proving Lemmas 19 and 20. First, we outline their proofs in the next section.

4.3.1 Proof Overview

Here we give an overview of the proofs of Lemmas 19 and 20. Let $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ be a function and $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ a matrix such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon$, for $\varepsilon > 0$. The goal is to use f to recover an approximation $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ of $\mathbf{A}$. Initialize empty matrices $(\mathbf{Z}, \mathbf{Y}) \in \mathbb{Z}_q^{0\times n} \times \mathbb{Z}_q^{0\times m}$ and the trivial subspace $\mathbf{W} := \{\mathbf{0}\} \subset \mathbb{Z}_q^n$. We define the following quantities:

- $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} := \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \mathbf{W} \bigm| \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big]$.

- $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) := \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\langle f(\mathbf{x}), \mathbf{z}\rangle = \langle \mathbf{A}^t\mathbf{z}, \mathbf{x}\rangle \bigm| \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big]$, for $\mathbf{z} \in \mathbb{Z}_q^n$.

Ensuring Progress by Conditional Agreement with A. Observe that initially, $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \varepsilon$, while

$$\begin{aligned}&\mathbb{E}_{\mathbf{z}\sim\mathbf{W}^\perp}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})\big] = \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} + (1-\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}})\cdot \\ &\qquad \mathbb{E}_{\mathbf{x}\sim\mathbb{Z}_q^m}\Big[\Pr_{\mathbf{z}\sim\mathbf{W}^\perp}\big[\langle f(\mathbf{x})-\mathbf{A}\mathbf{x},\mathbf{z}\rangle = 0\big] \;\Big|\; \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}, f(\mathbf{x})-\mathbf{A}\mathbf{x}\notin\mathbf{W}\Big] \\ &\quad = \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} + (1-\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}})/q \geq 1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/2.\end{aligned}$$

Thus, with probability $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4$ over $\mathbf{z}\sim\mathbf{W}^\perp$, we have that $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \geq 1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4$. We can then use the Goldreich-Levin Theorem to recover $\mathbf{y} = \mathbf{A}^t\mathbf{z} \in \mathbb{Z}_q^m$ with good probability as follows. Construct the predictor $\mathsf{Pred}_{\mathbf{z}}(\mathbf{x})$ which checks if $\mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}$, and if so outputs $\langle f(\mathbf{x}),\mathbf{z}\rangle \in \mathbb{Z}_q$, and otherwise outputs a random guess $\alpha\sim\mathbb{Z}_q$. If $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \geq 1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4$, then $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}_{\mathbf{z}}(\mathbf{x}) = \langle\mathbf{A}^t\mathbf{z},\mathbf{x}\rangle\big]$ is at least

$$\begin{aligned}&\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big](1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4) + (1-\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big])/q \\ &= 1/q + \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big]\cdot(\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4),\end{aligned}$$

and so the Goldreich-Levin Theorem recovers $\mathbf{y} = \mathbf{A}^t\mathbf{z}\in\mathbb{Z}_q^m$ with good probability. Suppose it additionally held for our choice of $\mathbf{z}$ that $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \leq \varDelta := 4q^{-1/3}$. Then, we would have

$$\begin{aligned}&\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x})-\mathbf{A}\mathbf{x}\in\mathbf{W} \;\big|\; \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x} \;\&\&\; \langle f(\mathbf{x}),\mathbf{z}\rangle = \langle\mathbf{y},\mathbf{x}\rangle\big] \\ &= \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x})-\mathbf{A}\mathbf{x}\in\mathbf{W} \;\&\&\; \langle f(\mathbf{x}),\mathbf{z}\rangle = \langle\mathbf{y},\mathbf{x}\rangle \;\big|\; \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big]/\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \\ &= \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}\cdot\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\langle f(\mathbf{x}),\mathbf{z}\rangle = \langle\mathbf{y},\mathbf{x}\rangle \;\big|\; \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x} \;\&\&\; f(\mathbf{x})-\mathbf{A}\mathbf{x}\in\mathbf{W}\big]/\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \\ &= \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \geq \varepsilon/\varDelta.\end{aligned}$$

Thus, after updating $\mathbf{Z}$ and $\mathbf{Y}$ by adding to them the rows $\mathbf{z}$ and $\mathbf{y}$, respectively, we have that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \varepsilon/\varDelta$ and $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] \geq \varepsilon$. We can easily see that after continuing this process k times, it holds that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \varepsilon/\varDelta^k \geq \varDelta/4 > 1/q$, for a constant $k\in\mathbb{N}$.

However, note that in the above analysis we have assumed that for each choice of $\mathbf{z} \sim \mathbf{W}^{\perp}$, it holds that $1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4 \leq \mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \leq \mathit{\Delta}$. We saw that with probability $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4$ over $\mathbf{z} \sim \mathbf{W}^{\perp}$, $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \geq 1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4$. So, for each choice of $\mathbf{z} \sim \mathbf{W}^{\perp}$, exactly one of the following conditions must hold:

- $\Pr_{\mathbf{z} \sim \mathbf{W}^{\perp}}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) > \mathit{\Delta}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/8$,
- or $\Pr_{\mathbf{z} \sim \mathbf{W}^{\perp}}\big[1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4 \leq \mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \leq \mathit{\Delta}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/8$.

The above analysis handles the case in which the second condition always holds. However, if the second condition ceases to hold before the end of k consecutive iterations of our procedure, then we must ensure progress under the first condition.

Ensuring Progress by Trading Agreement with A for a Dimension of W. Suppose that $\Pr_{\mathbf{z} \sim \mathbf{W}^{\perp}}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) > \mathit{\Delta}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/8$. In this case, it follows that the function $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}} : \mathbb{Z}_q^m \to \mathbb{C}$ defined above has high variance. We'll perform Fourier analysis on $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}$ to extract a heavy Fourier coefficient $\mathbf{w} \in \mathbb{Z}_q^n \backslash \{\mathbf{0}\}$ from the variance of $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}$. Finally, we'll show that such a heavy Fourier coefficient $\hat{\mathrm{Q}}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(\mathbf{w})$ is equivalent to
$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \big(\mathbf{W} + \mathrm{Span}(\mathbf{w})\big)\backslash \mathbf{W} \;\big|\; \mathbf{Z} \cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big]$. In all, this argument implies that there exists some $\mathbf{w} \in \mathbb{Z}_q^n \backslash \{\mathbf{0}\}$ such that

$$\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \big(\mathbf{W} + \mathrm{Span}(\mathbf{w})\big)\backslash \mathbf{W} \;\big|\; \mathbf{Z} \cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/\mathit{\Delta}.$$

We'll show that iterating this argument constructs a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ of constant dimension such that eventually $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \mathit{\Delta}/4 > 1/q$.

In slightly more detail, first let us denote $\dim(\mathbf{W})$ by d. Initially, $d = 0$. Define the function $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y}} : \mathbb{Z}_q^m \to \mathbb{C}$ by $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) = \mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})$ if $\mathbf{z} \in \mathbf{W}^{\perp}$, and $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) = \mathbf{0}$ otherwise.

It can be shown that $\forall \mathbf{w} \in \mathbb{Z}_q^n$, $\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(\mathbf{w}) = q^{-d} \cdot \mathbb{E}_{\mathbf{z} \sim \mathbf{W}^\perp}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})\big]$, if $\mathbf{w} \in \mathbf{W}$, and $\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(\mathbf{w}) = q^{-(d+1)} \cdot \Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \big(\mathbf{W} + \mathrm{Span}(\mathbf{w})\big) \backslash \mathbf{W} \bigm| \mathbf{Z} \cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big]$ otherwise. We can employ a Fourier analytic argument to bound the variance of $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})$ on a uniformly random $\mathbf{z} \sim \mathbf{W}^\perp$ by a quantity proportional to $\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(\mathbf{w}^*) := \max_{\mathbf{w} \in \mathbb{Z}_q^n \backslash \mathbf{W}}\{\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(\mathbf{w})\}$. By applying the hypothesis that $\Pr_{\mathbf{z} \sim \mathbf{W}^\perp}\big[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) > \mathit{\Delta}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/8$ and Chebyshev's inequality, one can show that $\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(\mathbf{w}^*) \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/\mathit{\Delta}$. By updating $\mathbf{W} := \mathbf{W} + \mathrm{Span}(\mathbf{w}^*)$ and iterating this procedure a constant number of times, it follows that eventually $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \mathit{\Delta}/4 > 1/q$.

The previous paragraph outlines a method for proving the existence of a constant dimensional subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ such that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} > 1/q$. However, in order to use this in any meaningful way, we must provide a method to algorithmically compute $\mathbf{W}$. At each step in the previous procedure, we have computed an explicit subspace $\mathbf{W} \subset \mathbb{Z}_q^n$, and have the guarantee that $\exists \mathbf{w}^* \in \mathbb{Z}_q^n \backslash \{\mathbf{W}\}$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \big(\mathbf{W} + \mathrm{Span}(\mathbf{w}^*)\big) \backslash \mathbf{W} \bigm| \mathbf{Z} \cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/\mathit{\Delta}$. Now, let $r \in \mathbb{N}$ be a constant and define the set $S = \{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{Z} \cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\}$. Consider the random process which chooses $\mathbf{x}_1, \ldots, \mathbf{x}_r \sim \mathbb{Z}_q^m$, $\alpha_1, \ldots, \alpha_r \sim \mathbb{Z}_q$ such that $\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i \in S$, chooses $\mathbf{w}' \sim \mathbf{W}$, and outputs $f(\sum_i \alpha_i \mathbf{x}_i) - \sum_i \alpha_i f(\mathbf{x}_i) - \mathbf{w}' \in \mathbb{Z}_q^n$. It can be shown by a statistical sampling argument that $f(\mathbf{u}) - \mathbf{A}\mathbf{u} \in \big(\mathbf{W} + \mathrm{Span}(\mathbf{w}^*)\big) \backslash \mathbf{W}, \forall \mathbf{u} \in \{\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i\}$, with good probability. Moreover, for each $\mathbf{x}_i$, we can write $f(\mathbf{x}_i) = \mathbf{A}\mathbf{x}_i + \mathbf{w}_i + \beta_i \mathbf{w}^*$, for $\mathbf{w}_i \in \mathbf{W}, \beta_i \in \mathbb{Z}_q \backslash \{0\}$, and we can write $f(\sum_i \alpha_i \mathbf{x}_i) = \mathbf{A}(\sum_i \alpha_i \mathbf{x}_i) + \mathbf{w} + \beta \mathbf{w}^*$, for $\mathbf{w} \in \mathbf{W}, \beta \in \mathbb{Z}_q \backslash \{0\}$. It thus follows that $f(\sum_i \alpha_i \mathbf{x}_i) - \sum_i \alpha_i f(\mathbf{x}_i) = (\mathbf{w} - \sum_i \alpha_i \mathbf{w}_i) + (\beta - \sum_i \alpha_i \beta_i)\mathbf{w}^*$. Assuming our guess $\mathbf{w}'$ for $\mathbf{w} - \sum_i \alpha_i \mathbf{w}_i \in \mathbf{W}$ is correct, then our random process outputs $\gamma \mathbf{w}^*$ ($\gamma \in \mathbb{Z}_q$). If

$\gamma \in \mathbb{Z}_q^*$, then $\mathrm{Span}(\gamma \mathbf{w}^*) = \mathrm{Span}(\mathbf{w}^*)$ and we can update $\mathbf{W} := \mathbf{W} + \mathrm{Span}(\gamma \mathbf{w}^*)$. However, if $\gamma \mathbf{w}^* = \mathbf{0}$, then this procedure fails. In particular, consider the example function f which is completely linear (*i.e.*, $f(\mathbf{x}) = \mathbf{A}\mathbf{x}, \forall \mathbf{x} \in \mathbb{Z}_q^m$). In this case, our random process will always output $\mathbf{0}$. However, it turns out that if our random process outputs $\mathbf{0}$, then f passes type of conditional linearity test. In this case, we'll apply Theorem 8 from Chapter 2 to conclude that f agrees conditionally with an affine map.

Putting it all Together. Now that we have computed $(\mathbf{Z}, \mathbf{Y}, \mathbf{W})$ such that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} > 1/q$, we can apply standard Goldreich-Levin "prediction-implies-inversion" and linear algebraic techniques to efficiently recover an approximation of the matrix $\mathbf{A}$. This process is quite straightforward and uses techniques described above; see the proof of Lemma 20.

4.4 Proofs of Lemmas 19 and 20

In this section, we prove Lemmas 19 and 20.

4.4.1 Proof of Lemma 19

In the course of proving Lemma 19, we continue with the notation established in our main theorem's proof overview (Section 4.3.1).

Proof of Lemma 19. The algorithm $\mathcal{A}_1$ of Lemma 19 is described in Figure 4.1. $\mathcal{A}_1$ first initializes empty matrices $(\mathbf{Z}, \mathbf{Y}) \in \mathbb{Z}_q^{0\times n} \times \mathbb{Z}_q^{0\times n}$ and sets $\mathbf{W} \subset \mathbb{Z}_q^n$ to be the trivial subspace ($\mathbf{W} = \{\mathbf{0}\}$). Let $\varDelta = 4q^{-1/3}$, $k \in \mathbb{N}$ a constant such that $\varDelta^{k+1} \leq \varepsilon$, and $\eta = \varepsilon/(8q^k)$. Since $\mathbb{E}_{\mathbf{z}\sim\mathbb{Z}_q^n}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})\big] \geq 1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/2$ implies that $\Pr_{\mathbf{z}\sim\mathbb{Z}_q^n}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \geq 1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4\big] \geq$

Parameters and Subroutines: Let $r \in \mathbb{N}$ be a constant. $\mathcal{A}_1$ will call the Goldreich-Levin algorithm $\mathcal{A}_{\mathsf{GL}}$ of Lemma 16 as a subroutine.

Inputs and Oracle Access: $\mathcal{A}_1$ takes no input, and gets oracle access to $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$.

1. The First Loop: Initialize $\mathbf{Z} \in \mathbb{Z}_q^{0\times n}$ and $\mathbf{Y} \in \mathbb{Z}_q^{0\times m}$ to empty matrices, and $\mathbf{W} \subset \mathbb{Z}_q^n$ to the trivial subspace $\mathbf{W} = \{\mathbf{0}\}$. Draw $\ell \sim \{0, \dots, n\}$, and do the following ℓ times.

- Draw $\mathbf{z} \sim \mathbb{Z}_q^n$ and $\mathbf{y} \sim \mathcal{A}_{\mathsf{GL}}^{\mathsf{Pred}}$ where $\mathsf{Pred} : \mathbb{Z}_q^m \to \mathbb{Z}_q$ is the function

$$\mathsf{Pred}(\mathbf{x}) = \begin{cases} \langle \mathbf{z}, f(\mathbf{x})\rangle, & \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x} \\ \alpha \sim \mathbb{Z}_q, & \text{otherwise} \end{cases}$$

- Update $\mathbf{Z}$ and $\mathbf{Y}$ by adding the new rows $\mathbf{z}$ and $\mathbf{y}$, respectively.

2. The Second Loop: Draw $d \sim \{0, \dots, n\}$. For all $j \in [d]$, choose linearly independent $\mathbf{x}_1, \dots, \mathbf{x}_r \sim \mathbb{Z}_q^m$ and $\alpha_1, \dots, \alpha_r \sim \mathbb{Z}_q$ such that $\mathbf{Z}\cdot f(\mathbf{u}) = \mathbf{Y}\mathbf{u}$ holds for all $\mathbf{u} \in \big\{\mathbf{x}_1, \dots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i\big\}$, and compute
$\mathbf{w}_j := f\big(\sum_i \alpha_i \mathbf{x}_i\big) - \sum_i \alpha_i f(\mathbf{x}_i) \in \mathbb{Z}_q^n$. Let $\mathbf{W} := \mathrm{Span}(\mathbf{w}_1, \dots, \mathbf{w}_d)$.

Output: $(\mathbf{Z}, \mathbf{Y}, \mathbf{W})$.

Figure 4.1: The Matrix Recovery Algorithm $\mathcal{A}_1$

$\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4$, then it follows that exactly one of the following conditions must hold:

The Goldreich-Levin Condition (GLC): $\Pr_{\mathbf{z}\sim\mathbb{Z}_q^n}\big[1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4 \leq \mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \leq \mathit{\Delta}\big] \geq \eta$.

The High Variance Condition: $\Pr_{\mathbf{z}\sim\mathbb{Z}_q^n}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) > \mathit{\Delta}\big] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} - \eta$.

Consider the first loop of $\mathcal{A}_1$. Either the GLC condition will hold for at least k iterations of the first loop, in which case we'll show the algorithm makes excellent progress in amplifying the conditional agreement of f and $\mathbf{A}$, or at some point before k iterations complete the GLC will cease to hold. $\mathcal{A}_1$ chooses $\ell \sim \{0, 1, \ldots, n\}$ as a guess for the number t of consecutive iterations for which the GLC will hold (this guess is correct with probability $1/(n+1)$). For each iteration in $[\ell]$, the algorithm chooses a random $\mathbf{z} \sim \mathbb{Z}_q^n$ (note that with probability η of this random choice of $\mathbf{z}$, it holds that $1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}/4 \leq \mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \leq \mathit{\Delta}$). Then, $\mathcal{A}_1$ constructs the predictor as detailed in Figure 4.1. Note that

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}(\mathbf{x}) = \langle \mathbf{x}, \mathbf{A}^t\mathbf{z}\rangle\big] = \mathbb{E}_{\mathbf{z}\sim\mathbb{Z}_q^n}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \bigm| \mathsf{GLC}\text{ holds}\big] \geq 1/q + \varepsilon^2/4.$$

Then, by Lemma 16, $\mathcal{A}_{\mathsf{GL}}^{\mathsf{Pred}}$ outputs $\mathbf{y} = \mathbf{A}^t\mathbf{z} \in \mathbb{Z}_q^m$ with probability $\varepsilon^4/\mathcal{O}(m^3q^5)$. Again since the GLC holds, it follows that after updating $\mathbf{Z}, \mathbf{Y}$ with rows $\mathbf{z}, \mathbf{y}$, respectively, we have that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \varepsilon/\mathit{\Delta}$ with probability $\eta\varepsilon^4/\mathcal{O}(m^3q^5)$ over $(\mathbf{z}, \mathbf{y})$. In total, after ℓ iterations of the first loop, it follows that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq \varepsilon/\mathit{\Delta}^t =: \varepsilon'$ with probability at least $(\eta\varepsilon^4/\mathcal{O}(m^3q^5))^t/(n+1)$.

Now, let $S = \{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{Z} \cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\}$, and note that $|S|q^{-m} \geq \varepsilon$. Before entering the second loop, $\mathcal{A}_1$ chooses $d \sim \{0, 1, \ldots, n\}$ as a guess for $k - t$ (this guess is correct with probability $1/(n+1)$). Since at this point we must have $\varepsilon' \leq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} < \mathit{\Delta}/4$ and that the High Variance Condition holds, by Claim 18 there exists a subspace $\mathbf{W}^* \subset \mathbb{Z}_q^n$

of dimension $k-t$ such that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} \geq \Delta/4 = q^{-1/3}$. Our goal in the second loop is then to algorithmically recover $\mathbf{W}^*$.

Let $T = \{\mathbf{x} \in S : f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W}^*\}$, and note that the density of T in S is $\lambda := |T|/|S| \geq q^{-1/3}$. Let $r \in \mathbb{N}$ be a constant such that $4/(\lambda q^{r-1}) \leq \lambda^{r+1}/8$. For each iteration j of the second loop, $\mathcal{A}_1$ chooses $\mathbf{x}_1, \ldots, \mathbf{x}_r \sim \mathbb{Z}_q^m$, $\alpha_1, \ldots, \alpha_r \sim \mathbb{Z}_q$ such that $\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum\limits_i \alpha_i \mathbf{x}_i \in S$. By Corollary 2, we have that $\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum\limits_i \alpha_i \mathbf{x}_i \in T$ with probability at least $\lambda^{r+1}/4 - 2^{-\Omega(m)}$. So, we must have that $\forall \mathbf{u} \in \{\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum\limits_i \alpha_i \mathbf{x}_i\}$, $\mathbf{w}_j = f(\mathbf{u}) - \mathbf{A}\mathbf{u} \in \mathbf{W}^*$, and thus $f(\sum\limits_i \alpha_i \mathbf{x}_i) - \sum\limits_i \alpha_i f(\mathbf{x}_i) \in \mathbf{W}^*$. Let $\gamma = (q-1)\varepsilon^{r+1}\lambda^{r+1}/(8q^{(r+1)k-r})$. Either $\{\mathbf{w}_1, \ldots, \mathbf{w}_d\} \subset \mathbf{W}^*$ are linearly independent with probability γ, which implies that we have recovered $\mathbf{W} := \mathrm{Span}(\mathbf{w}_1, \ldots, \mathbf{w}_d) = \mathbf{W}^*$ with probability $(\lambda^{r+1}/4 - 2^{-\Omega(m)})\gamma$, or $\mathbf{W} \subset \mathbf{W}^*$ is a proper subspace such that

$$\Pr_{\substack{\mathbf{x}_1, \ldots, \mathbf{x}_r \sim \mathbb{Z}_q^m \\ \alpha_1, \ldots, \alpha_r \sim \mathbb{Z}_q}} \Big[f\big(\sum_i \alpha_i \mathbf{x}_i\big) - \sum_i \alpha_i f(\mathbf{x}_i) \in \mathbf{W} \;\Big|\; \mathbf{x}_1, \ldots, \mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i \in T \Big] \geq 1 - \gamma.$$

In the latter case, let $s = \dim(\mathbf{W})$ and let us extend a basis $\{\hat{\mathbf{w}}_1, \ldots, \hat{\mathbf{w}}_s\} \subset \mathbf{W}$ for $\mathbf{W}$ to a basis $\{\hat{\mathbf{w}}_1, \ldots, \hat{\mathbf{w}}_s, \hat{\mathbf{w}}_{s+1}, \ldots, \hat{\mathbf{w}}_d\} \subset \mathbf{W}^*$ for $\mathbf{W}^*$, and reindex $\{\hat{\mathbf{w}}_{s+1}, \ldots, \hat{\mathbf{w}}_d\}$ as $\{\mathbf{w}_1^*, \ldots, \mathbf{w}_{d-s}^*\} \subset \mathbf{W}^* \backslash \mathbf{W}$. Define the function $\phi : \mathbb{Z}_q^m \to \mathbb{Z}_q^{d-s}$ by $\phi(\mathbf{x}) = (\beta_1, \ldots, \beta_{d-s})$ such that $f(\mathbf{x}) - \mathbf{A}\mathbf{x} - \sum\limits_{j=1}^{d-s} \beta_j \mathbf{w}_j^* \in \mathbf{W}$, if $\mathbf{x} \in T$, and $\phi(\mathbf{x}) = 0$ otherwise. For each $j \in [d-s]$, we denote the j^{th} projection of ϕ by $\phi_j : \mathbb{Z}_q^m \to \mathbb{Z}_q$. Note that if $\mathbf{x}_1, \ldots, \mathbf{x}_r \in \mathbb{Z}_q^m$, $\alpha_1, \ldots, \alpha_r \in \mathbb{Z}_q$ such that $\mathbf{x}_1, \ldots, \mathbf{x}_r, \sum\limits_i \alpha_i \mathbf{x}_i \in T$ then we can write $f(\mathbf{u}) = \mathbf{A}\mathbf{u} + \sum\limits_{j=1}^{d-s} \phi_j(\mathbf{u})\mathbf{w}_j^* + \mathbf{w}_i' \in \mathbb{Z}_q^n, \forall \mathbf{u} \in \{\mathbf{x}_1, \ldots, \mathbf{x}_r\}$ (each $\mathbf{w}_i' \in \mathbf{W}$), and $f\big(\sum\limits_{i=1}^{r} \alpha_i \mathbf{x}_r\big) = \mathbf{A}\big(\sum\limits_{i=1}^{r} \alpha_i \mathbf{x}_r\big) + \sum\limits_{j=1}^{d-s} \phi_j\big(\sum\limits_{i=1}^{r} \alpha_i \mathbf{x}_r\big)\mathbf{w}_i^* + \mathbf{w}' \in \mathbb{Z}_q^n$ ($\mathbf{w}' \in \mathbf{W}$). If additionally it held that $f\big(\sum\limits_{i=1}^{r} \alpha_i \mathbf{x}_r\big) -$

$\sum_{i=1}^{r} \alpha_i f(\mathbf{x}_r) \in \mathbf{W}$, then we'd have

$$\begin{aligned}\mathbf{W} &\ni \mathbf{A}\Big(\sum_i \alpha_i \mathbf{x}_i\Big) + \sum_{j=1}^{d-s} \phi_j\Big(\sum_i \alpha_i \mathbf{x}_i\Big)\mathbf{w}_i^* + \mathbf{w}' - \bigg(\sum_i \alpha_i\Big(\mathbf{A}\mathbf{x}_i + \sum_{j=1}^{d-s}\phi_j(\mathbf{x}_i)\mathbf{w}_j^* + \mathbf{w}_j'\Big)\bigg) \\ &= \sum_{j=1}^{d-s}\bigg(\phi_j\Big(\sum_i \alpha_i \mathbf{x}_i\Big) - \sum_{i=1}^{r}\alpha_i \phi_j(\mathbf{x}_i)\bigg)\mathbf{w}_j^* + \Big(\mathbf{w}' - \sum_{i=1}^{r}\alpha_i \mathbf{w}_i'\Big).\end{aligned}$$

Hence $\sum_{j=1}^{d-s}\bigg(\phi_j\Big(\sum_i \alpha_i \mathbf{x}_i\Big) - \sum_{i=1}^{r}\alpha_i\phi_j(\mathbf{x}_i)\bigg)\mathbf{w}_j^* \in \mathbf{W}$, which implies that $\phi\Big(\sum_i \alpha_i \mathbf{x}_i\Big) = \sum_i \alpha_i \phi(\mathbf{x}_i)$. So, we have

$$\Pr_{\substack{\mathbf{x}_1,\ldots,\mathbf{x}_r \sim \mathbb{Z}_q^m \\ \alpha_1,\ldots,\alpha_r \sim \mathbb{Z}_q}}\bigg[\phi\Big(\sum_i \alpha_i \mathbf{x}_i\Big) = \sum_i \alpha_i \phi(\mathbf{x}_i) \;\bigg|\; \mathbf{x}_1,\ldots,\mathbf{x}_r, \sum_i \alpha_i \mathbf{x}_i \in T\bigg] \geq 1 - \gamma.$$

■

Claim 18 *Continuing the notation established in Lemma 19, suppose that after the execution of the first loop of* $\mathcal{A}_1$, $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\{\mathbf{0}\}} < \mathit{\Delta}/4$ *and the High Variance Condition holds. Then, there exists a subspace* $\mathbf{W}^* \subset \mathbb{Z}_q^n$ *of dimension* $k - t$ *such that* $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} \geq \mathit{\Delta}/4$.

Proof. Initialize $\mathbf{W}^* := \{\mathbf{0}\} \subset \mathbb{Z}_q^n$ to be the trivial subspace. Define $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} : \mathbb{Z}_q^m \to [0,1]$ by

$$\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) = \begin{cases} \mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}), & \mathbf{z} \perp \mathbf{W}^* \\ 0, & \text{otherwise} \end{cases}$$

Note that $\mathbb{E}_{\mathbf{z}\perp\mathbf{W}^*}\big[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z})\big] = 1/q + \big(1 - 1/q\big)\cdot \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}$. The following claim characterizes the Fourier coefficients of $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}$; we prove it below, outside the current proof.

Claim 19 *Let* $d = \dim(\mathbf{W}^*)$. *For all* $\mathbf{w} \in \mathbb{Z}_q^n$, $q^d \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}) = \mathbb{E}_{\mathbf{z}\perp\mathbf{W}^*}\big[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z})\big]$, *if* $\mathbf{w} \in \mathbf{W}^*$, *and* $q^d \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}) = q^{-1} \cdot \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\Big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \big(\mathbf{W}^* + \mathrm{Span}(\mathbf{w})\big) \setminus \mathbf{W}^* \Big| \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\Big]$, *otherwise.*

In order to prove the claim, we must show that whenever $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} \leq \mathit{\Delta}/4$ and the High Variance Condition holds, letting d denote the current value of $\dim(\mathbf{W}^*)$, $\exists \mathbf{w}^* \notin \mathbf{W}^*$ such that $q^{d+1} \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}^*) \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/\mathit{\Delta}$. It then follows that $\mathrm{P}_{\mathbf{Z},\mathbf{Y},(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))} \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/\mathit{\Delta}$, and so our procedure updates $\mathbf{W}^* := \mathbf{W}^* + \mathrm{Span}(\mathbf{w}^*)$. Note that if the High Variance Condition holds, then it will also hold for the next iteration. To see this, observe that if $\mathrm{Pr}_{\mathbf{z}\perp\mathbf{W}^*}\left[1/q + \varepsilon/4 \leq \mathrm{Q}^*(\mathbf{z}) \leq \mathit{\Delta}\right] < \nu$, then

$$\mathrm{Pr}_{\mathbf{z}\perp(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))}\left[1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))}/4 \leq \mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))}(\mathbf{z}) \leq \mathit{\Delta}\right]$$
$$\leq q \cdot \mathrm{Pr}_{\mathbf{z}\perp\mathbf{W}^*}\left[1/q + \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/4 \leq \mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) \leq \mathit{\Delta}\right] < q \cdot \nu,$$

since $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))}(\mathbf{z}) = \mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z})$ for all $\mathbf{z} \perp \mathbf{W}^*_{i+1}$ and $\mathrm{P}_{\mathbf{Z},\mathbf{Y},(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))} \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}$. Hence $\mathrm{Pr}_{\mathbf{z}\perp(\mathbf{W}^*+\mathrm{Span}(\mathbf{w}^*))}\left[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) > \mathit{\Delta}\right] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/4 - q\nu$. So, initially when the High Variance Condition holds we have $\mathrm{Pr}_{\mathbf{z}\perp\mathbf{W}^*}\left[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) > \mathit{\Delta}\right] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/4 - \eta$, and at the beginning of each iteration $i \in [k-t]$ of our procedure we'll have $\mathrm{Pr}_{\mathbf{z}\perp\mathbf{W}^*}\left[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) > \mathit{\Delta}\right] \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/4 - q^{i-1}\eta \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/4 - q^k\eta = \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/4 - \varepsilon/8 \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/8$, since $\eta = \varepsilon/(8q^k)$. Therefore, after $k-t$ consecutive iterations of our procedure we have $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} \geq \mathit{\Delta}/4$.

Now, consider an iteration i of our procedure. We'll show $\exists \mathbf{w}^* \notin \mathbf{W}^*$ such that $q^{d+1} \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}^*) \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/\varDelta$. We have

$$\begin{aligned}
&\frac{\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}}{8} \leq \mathrm{Pr}_{\mathbf{z}\perp\mathbf{W}^*}\left[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) > \varDelta\right] \\
&\leq \mathrm{Pr}_{\mathbf{z}\perp\mathbf{W}^*}\left[\left(\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) - q^d \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{0})\right)^2 \geq \varDelta^2/4\right] \\
&\leq \frac{4}{\varDelta^2} \cdot \left[\mathbb{E}_{\mathbf{z}\perp\mathbf{W}^*}\left[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z})^2\right] - q^{2d} \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{0})^2\right] \\
&= \frac{4q^d}{\varDelta^2} \cdot \left[\mathbb{E}_{\mathbf{z}\sim\mathbb{Z}_q^n}\left[\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z})^2\right] - \sum_{\mathbf{w}\in\mathbf{W}^*} \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w})^2\right] \\
&= \frac{4q^d}{\varDelta^2} \cdot \sum_{\mathbf{w}\notin\mathbf{W}^*} \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w})^2 \leq \frac{4q^d}{\varDelta^2} \cdot \max_{\mathbf{w}^*\notin\mathbf{W}^*} \left\{\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}^*)\right\} \cdot \sum_{\mathbf{w}\in\mathbb{Z}_q^n} \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}) \\
&= \frac{4q^d}{\varDelta^2} \cdot \max_{\mathbf{w}^*\notin\mathbf{W}^*} \left\{\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}^*)\right\}.
\end{aligned}$$

Thus, $\exists \mathbf{w}^* \notin \mathbf{W}^*$ such that $q^{d+1} \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}^*) \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} \cdot \left(q\varDelta^2/32\right) \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}/\varDelta$, since $\varDelta = 4q^{-1/3}$. The computation above has used $q^d \cdot \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{0}) = \mathbb{E}_{\mathbf{z}\perp\mathbf{W}^*}\left[\mathrm{Q}^*(\mathbf{z})\right]$ and $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*} \leq \varDelta/4$ for the second inequality; Markov's inequality for the third; $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{z}) = 0$ for $\mathbf{z} \not\perp \mathbf{W}^*$, and $\hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}) = \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w})$ for $\mathbf{w} \in \mathbf{W}^*$ for the equality on the fourth line; Parseval's identity for the next equality; and $\sum_{\mathbf{w}\in\mathbb{Z}_q^n} \hat{\mathrm{Q}}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{w}) = \mathrm{Q}_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}(\mathbf{0}) = 1$ for the final equality. ■

Proof of Claim 19. For convenience we'll denote the function $\mathrm{Q}^*_{\mathbf{Z},\mathbf{Y},\mathbf{W}^*}$ by Q^*. Let $\mathbf{w} \in \mathbb{Z}_q^n$. We have

$$\begin{aligned} q^d \cdot \hat{\mathrm{Q}}^*(\mathbf{w}) &= q^d \cdot \mathbb{E}_{\mathbf{z}\sim\mathbb{Z}_q^n}\left[\mathrm{Q}^*(\mathbf{z}) \cdot \omega^{-\langle \mathbf{z},\mathbf{w}\rangle}\right] = \mathbb{E}_{\mathbf{z}\perp\mathbf{W}}\Big[\mathrm{Pr}_{\mathbf{x}}\big[f(\mathbf{x}) - \mathbf{A}\mathbf{x} \perp \mathbf{z} \big| \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}\big] \cdot \\ \omega^{-\langle \mathbf{z},\mathbf{w}\rangle}\Big] &= \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \cdot \mathbb{E}_{\mathbf{z}\perp\mathbf{W}^*}\left[\omega^{-\langle \mathbf{z},\mathbf{w}\rangle}\right] + \frac{1}{q} \cdot (1 - \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}) \cdot \mathbb{E}_{\substack{\mathbf{z}\perp\mathbf{W}^* \\ f(\mathbf{x})-\mathbf{A}\mathbf{x}\perp\mathbf{z}}}\left[\omega^{-\langle \mathbf{z},\mathbf{w}\rangle}\right], \end{aligned}$$

where $\omega = e^{2\pi i/q}$ is a complex $q-$th root of unity. Thus $q^d \cdot \hat{\mathrm{Q}}^*(\mathbf{w}) = 1/q + (1 - 1/q) \cdot \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}$ whenever $\mathbf{w} \in \mathbf{W}^*$, since the $\omega^{-\langle \mathbf{z},\mathbf{w}\rangle}$ quantities are all equal to 1 in this case. When $\mathbf{w} \notin \mathbf{W}$, the first term is always zero, and the second term is zero unless $f(\mathbf{x}) - \mathbf{A}\mathbf{x} \in \mathbf{W} + \mathrm{Span}(\mathbf{w})$, in which case the $\omega^{-\langle \mathbf{z},\mathbf{w}\rangle}$ quantities in the second term are all equal to 1. The claim follows.

■

4.4.2 Proof of Lemma 20

We first prove the following claim, which shows if $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} > 1/q$, then we can repeat a modified version of the first loop of $\mathcal{A}_1$ to construct matrices $(\mathbf{Z}', \mathbf{Y}')$ such that $\mathrm{P}_{\mathbf{Z}',\mathbf{Y}',\mathbf{W}} \geq 1 - 1/q$.

Claim 20 *If* $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq q^{-1/3}$*, then there exists an efficient algorithm running in time* $\mathsf{poly}(m, n, q)$ *which with probability* $\mathsf{poly}(1/m, 1/n, 1/q)$ *over its randomness outputs matrices* $(\mathbf{Z}', \mathbf{Y}') \in \mathbb{Z}_q^{\ell+t\times n} \times \mathbb{Z}_q^{\ell+t\times m}$*, for a constant* $t \in \mathbb{N}$*, such that* $\mathrm{P}_{\mathbf{Z}',\mathbf{Y}',\mathbf{W}} \geq 1 - 1/q$*.*

Proof. Let $t \in \mathbb{N}$ be a constant such that $(1 - 1/(4q))^{t+1} \leq q^{-1/3}$. Suppose $\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} < 1 - 1/q$. Then, $\mathrm{R}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} := 1 - \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq 1/q$. Note that $\forall \mathbf{z} \in \mathbf{W}^\perp$, $\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \geq \mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}} \geq q^{-1/3}$, while $\mathbb{E}_{\mathbf{z}\perp\mathbf{W}}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})\big] = 1 - \mathrm{R}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(1 - 1/q)$. Then,

$$\begin{aligned}\Pr_{\mathbf{z}\perp\mathbf{W}}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) > 1 - 1/(4q)\big] &\leq \frac{\mathbb{E}_{\mathbf{z}\perp\mathbf{W}}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z})\big]}{1 - 1/(4q)} \\ &= \frac{1 - \mathrm{R}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(1 - 1/q)}{1 - 1/(4q)} \leq \frac{1 - 1/(2q)}{1 - 1/(4q)},\end{aligned}$$

since $\mathrm{R}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}(1 - 1/q) \geq (1/q)(1/2)$. Thus

$$\Pr_{\mathbf{z}\perp\mathbf{W}}\big[\mathrm{Q}_{\mathbf{Z},\mathbf{Y}}(\mathbf{z}) \leq 1 - 1/(4q)\big] \geq 1 - \frac{1 - 1/(2q)}{1 - 1/(4q)} \geq \frac{1/(4q)}{1 - 1/(4q)} \geq \frac{1}{4q},$$

Then, following the same analysis as the first loop of $\mathcal{A}_1$, we can choose a random $\mathbf{z} \perp \mathbf{W}$, construct the predictor algorithm detailed in Figure 4.1, and apply the Goldreich Levin Theorem (Lemma 16) to compute $\mathbf{y} = \mathbf{A}^t\mathbf{z} \in \mathbb{Z}_q^m$, all with probability $(q^{-1/3} - q^{-1})^2/\mathcal{O}(m^3q^6)$. We then have

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W} \;\big|\; \mathbf{Z}\cdot f(\mathbf{x}) = \mathbf{Y}\mathbf{x}, \langle \mathbf{z}, f(\mathbf{x})\rangle = \langle \mathbf{x}, \mathbf{y}\rangle\big] \geq \frac{\mathrm{P}_{\mathbf{Z},\mathbf{Y},\mathbf{W}}}{1 - 1/(4q)},$$

and so we update $(\mathbf{Z}, \mathbf{Y})$ by adding the rows $(\mathbf{z}, \mathbf{y})$, respectively. After continuing this process t times, the resulting matrices $(\mathbf{Z}', \mathbf{Y}')$ are then such that

$$\mathrm{P}_{\mathbf{Z}',\mathbf{Y}',\mathbf{W}} \geq \frac{q^{-1/3}}{(1 - 1/(4q))^t} \geq 1 - \frac{1}{4q} \geq 1 - \frac{1}{q}$$

with probability $((q^{-1/3} - q^{-1})^2/\mathcal{O}(m^3q^6))^t$. ■

Now, we may prove Lemma 20.

Proof of Lemma 20. Let $t \in \mathbb{N}$ be a constant such that $q^{-t} \leq \varepsilon \leq q^{-(t-1)}$, where $\varepsilon > 0$ is the initial agreement of f with $\mathbf{A}$, and $c = 2t + 1$. Since $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W} \;\big|\; \mathbf{x} \in S\big] \geq 1 - 1/q$, where $\dim(\mathbf{W}) = d = \mathcal{O}(1)$, then by Lemma 18, we can compute a linearly independent subset $\{\mathbf{x}_1, \ldots, \mathbf{x}_{m-c}\} \subset T$ with probability $(1 - 2^{-\Omega(m)})(1 - m(2/q + \varepsilon/2)) \geq (1 - 2^{-\Omega(m)})(1 - 5m/(2q)) \geq 1 - 3m/q$. For each $i \in [m - c]$, write $f(\mathbf{x}_i) = \mathbf{A}\mathbf{x} + \mathbf{w}_i$

($\mathbf{w}_i \in \mathbf{W}$). Also, since $\mathbf{W}^{\perp}$ is efficiently recognizable and has dimension $n-d$, then again by Lemma 18 we can compute a linearly independent subset $\{\mathbf{z}_1, \ldots, \mathbf{z}_{n-d}\} \subset \mathbf{W}^{\perp}$ with probability $1 - 2^{-\Omega(m)}$.

Extend $\{\mathbf{z}_1, \ldots, \mathbf{z}_{n-d}\}$ to a basis $\{\mathbf{z}_1, \ldots, \mathbf{z}_n\}$ for $\mathbb{Z}_q^n$. Let $\mathbf{B} \in \mathbb{Z}_q^{n\times n}$ be the matrix whose i^{th} row is $\mathbf{z}_i$. Note that $\forall i \in [m-c]$, the first $n-d$ components of $\mathbf{B}\cdot f(\mathbf{x}_i) - \mathbf{B}\mathbf{A}\mathbf{x}_i \in \mathbb{Z}_q^n$ are zero, and so we can use a routine linear algebraic argument to recover $\mathbf{B}' \in \mathbb{Z}_q^{(n-d)\times m}$ such that $\operatorname{rank}\big(\mathbf{B}' - (\mathbf{B}\mathbf{A})^{(n-d)}\big) = \mathcal{O}(1)$, where $(\mathbf{B}\mathbf{A})^{(n-d)}$ denotes the $(n-d) \times m$ matrix over $\mathbb{Z}_q$ consisting of the first $n-d$ rows of $\mathbf{B}\mathbf{A}$. Extend $\mathbf{B}'$ to a matrix $\bar{\mathbf{B}} \in \mathbb{Z}_q^{n\times m}$ where the last d rows are chosen uniformly at random. It follows that $\operatorname{rank}(\bar{\mathbf{B}} - \mathbf{B}\mathbf{A}) \leq \operatorname{rank}(\mathbf{B}' - (\mathbf{B}\mathbf{A})^{(n-d)}) + d \leq \mathcal{O}(1)$, which implies that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\bar{\mathbf{B}}\mathbf{x} = \mathbf{B}\mathbf{A}\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$. Our algorithm then simply outputs $\mathbf{B}^{-1}\mathbf{B}' \in \mathbb{Z}_q^{n\times m}$. ■

Chapter 5

Conclusions

In this dissertation, we have proven two novel high dimensional linearity testing theorems. Both of our high dimensional linearity tests work in characteristic q, and can be thought of as a "hybrid" between the high and low test acceptance parameter regimes. What we mean by this is that while these linearity tests both have a high test acceptance, they each have some property which allows us to view them through a certain lens as a low test acceptance test (although these low test acceptance conditions are not sufficient for proving the theorems). We demonstrate a concrete application of both linearity tests to lattice-based cryptography and the Goldreich-Levin Theorem in high dimension, respectively.

Our first novel linearity tests shows that if there exists a function $h : \mathbb{Z}_q^n \to \mathbb{Z}_q^n$ ($n, q \in \mathbb{N}$ such that q is prime) such that $\forall \alpha_1, \alpha_2 \in \mathbb{Z}_q^2$, $\Pr_{\mathbf{a}_1, \mathbf{a}_2 \sim \mathbb{Z}_q^n}\big[h(\alpha_1 \mathbf{a}_1 + \alpha_2 \mathbf{a}_2) \in \mathrm{Span}(h(\mathbf{a}_1), h(\mathbf{a}_2))\big] \geq 1 - \varepsilon$, for some $\varepsilon > 0$, then there exists a linear map $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^n}\big[h(\mathbf{a}) \in \mathrm{Span}(\mathbf{H}\mathbf{a})\big] \geq 1 - \mathcal{O}(n^2 q \sqrt{q\varepsilon})$. We remark that our linearity test at first glance does not appear to fit into the general template of linearity tests. Note, however, that by averaging, our conclusion can be re-written as follows: $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^n}\big[h(\mathbf{a}) = \mathbf{H}\mathbf{a}\big] \geq (1 - \mathcal{O}(n^2 q\sqrt{q\varepsilon}))/q$. Although this appears to indicate a traditional linearity test with high test acceptance and low agreement, note that again by averaging our hypothesis, we similarly obtain a low test acceptance. To be clear, this low test acceptance condition *is not* sufficient for our conclusion, but rather supplies a data point on how we can view our linearity test through a certain lens as a "hybrid" between the high and low test acceptance parameter regimes. Finally, we emphasize that while all previous linearity testing theorems we are aware of in the literature use combinatorial and Fourier analytic techniques, our proof is purely algebraic, and follows the proof of the Fundamental Theorem of Projective Geometry ([Art57]).

We demonstrate an application of our first linearity testing theorem to lattice-based cryptography. Specifically we show that there cannot exist a certain type of reduction from Learning with Errors (LWE) to Learning with Rounding (LWR), unless LWE is computationally tractable. Reductions from LWE to LWR are known when the modulus q is super-polynomial in the lattice dimension n, and when q is polynomial in n with an *a priori* bound on the number m of input samples to the reduction. Each prior reduction from LWE to LWR in the literature conforms to a specific template, which we characterize by a function $f : \mathbb{Z}_q^n \times \mathbb{Z}_q \to \mathbb{Z}_q^n \times \mathbb{Z}_p$ ($n, q, p \in \mathbb{N}$ such that q is prime and $p < q$) that operates point-wise on the input samples to the reduction. Our result addresses the case of polynomial q and unbounded m; we show that if there exists a reduction f from LWE to LWR, then f satisfies our high dimensional linearity test, hence agrees with a linear map with good probability. Then, we give an efficient algorithm which directly solves LWE given such a reduction function f which has good agreement with a linear map. Overall, our result *does not* suggest that LWR is computationally tractable, but rather that LWE cannot be reduced to LWR, when q is polynomial in n and m is unbounded, using techniques from the prior work.

Our second novel linearity test shows that if there exists a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^d$ ($m, d, q \in \mathbb{N}$ such that q is prime) and $S \subset \mathbb{Z}_q^m$ is a subset of sufficiently large density $\lambda := |S|q^{-m}$ such that $\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y}) \bigm| \mathbf{x}, \mathbf{y} \in S\big] \geq 1 - \gamma$, for $\gamma > 0$, then there exists an affine map $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{d\times m} \times \mathbb{Z}_q^d$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} \bigm| \mathbf{x} \in S\big] \geq (1 + \mathsf{poly}(\lambda, q^{-d}))^d$. Again, we remark that our result is not a traditional linearity test; indeed, it can be viewed as a high dimensional conditional version of [BLR93] in characteristic q. Also while our linearity test appears to be a high test acceptance and low agreement test, by removing the condition in the hypothesis, we obtain a low test acceptance. We again remark that this low test acceptance condition is not sufficient for our conclusion, but instead allows us to view our result as a "hybrid" between the high and low test acceptance parameter regimes.

We apply our second novel linearity test to proving a high dimensional approximation version of the celebrated Goldreich-Levin Theorem ([GL89]) with low agreement. Let $m, n, q \in \mathbb{N}$ such that $q = \mathsf{poly}(m, n)$ is prime, and suppose there exists a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ and a function $f : \mathbb{Z}_q^m \to \mathbb{Z}_q^n$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[f(\mathbf{x}) = \mathbf{A}\mathbf{x}\big] \geq \varepsilon$, where $q^{-n} < \varepsilon \leq q^{-1}$. Indeed, in this parameter regime, it turns out that one cannot hope to exactly recover $\mathbf{A}$, but instead we consider the problem of recovering an approximation of $\mathbf{A}$ (*i.e.*, a matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ such that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathbf{A}'\mathbf{x} = \mathbf{A}\mathbf{x}\big] \geq q^{-\mathcal{O}(1)}$). Such a high dimensional approximate

Goldreich-Levin Theorem continues a line of work relating to list-decoding group homomorphism codes, and has applications to approximate list-decodable codes and effective algebraic property testing. In our work, we construct an algorithm which uses f to recover an approximation of $\mathbf{A}$, and characterize the conditions under which our algorithm fails to do so. Specifically, we first show that we can compute (the description of) a subset $S \subset \mathbb{Z}_q^m$ and a subspace $\mathbf{W} \subset \mathbb{Z}_q^n$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[f(\mathbf{x}) \in \mathbf{A}\mathbf{x} + \mathbf{W} \bigm| \mathbf{x} \in S\big] > 1/q$, and then use standard Goldreich-Levin "prediction-implies-inversion" and linear algebraic techniques to recover an approximation of $\mathbf{A}$. Finally, we show that the only way in which our algorithm fails to recover an approximation of $\mathbf{A}$ is if f satisfies our high dimensional conditional linearity test.

Bibliography

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

[ABN$^+$92] Noga Alon, Jehoshua Bruck, Joseph Naor, Moni Naor, and Ron M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Trans. Inf. Theory*, 38(2):509–516, 1992.

[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.

[AGHP93] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Addendum to "simple construction of almost k-wise independent random variables". *Random Struct. Algorithms*, 4(1):119–120, 1993.

[AIK11] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. How to garble arithmetic circuits. In Rafail Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 120–129. IEEE Computer Society, 2011.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996.

[Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In Jirí Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *Automata, Languages and Programming, 26th International Colloquium, ICALP'99, Prague, Czech Republic, July 11-15, 1999, Proceedings*, volume 1644 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 1999.

[Ajt10] Miklós Ajtai. Oblivious rams without cryptographic assumptions. In *STOC*, pages 181–190. ACM, 2010.

[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.

[ALM$^+$92] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and hardness of approximation problems. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*, pages 14–23. IEEE Computer Society, 1992.

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

[BBL$^+$14] Abhishek Banerjee, Hai Brenner, Gaëtan Leurent, Chris Peikert, and Alon Rosen. SPRING: fast pseudorandom functions from rounded ring products. In Carlos Cid and Christian Rechberger, editors, *Fast Software Encryption - 21st International Workshop, FSE 2014, London, UK, March 3-5, 2014. Revised Selected Papers*, volume 8540 of *Lecture Notes in Computer Science*, pages 38–57. Springer, 2014.

[BBW18] László Babai, Timothy J. F. Black, and Angela Wuu. List-decoding homomorphism codes with arbitrary codomains. In Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, volume 116 of *LIPIcs*, pages 29:1–29:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[BCH$^+$96] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos A. Kiwi, and Madhu Sudan. Linearity testing in characteristic two. *IEEE Trans. Inf. Theory*, 42(6):1781–1795, 1996.

[BCLR04] Michael Ben-Or, Don Coppersmith, Michael Luby, and Ronitt Rubinfeld. Non-abelian homomorphism testing, and distributions close to their self-convolutions. *Electron. Colloquium Comput. Complex.*, (052), 2004.

[BGLR93] Mihir Bellare, Shafi Goldwasser, Carsten Lund, and A. Russeli. Efficient probabilistically checkable proofs and applications to approximations. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 294–304. ACM, 1993.

[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.

[BGS95] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, pcps and nonapproximability - towards tight results. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*, pages 422–431. IEEE Computer Society, 1995.

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

[BI01] Amos Beimel and Yuval Ishai. Information-theoretic private information retrieval: A unified construction. In *ICALP*, volume 2076 of *Lecture Notes in Computer Science*, pages 912–926. Springer, 2001.

[BIM04] Amos Beimel, Yuval Ishai, and Tal Malkin. Reducing the servers' computation in private information retrieval: PIR with preprocessing. *J. Cryptology*, 17(2):125–151, 2004.

[BIPW] Elette Boyle, Yuval Ishai, Rafael Pass, and Mary Wooters. Can a database be accessed privately and locally? *submitted to these proceedings*.

[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.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.

[Blu83] Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *SIGACT News*, 15(1):23–27, 1983.

[BN00] Daniel Bleichenbacher and Phong Q. Nguyen. Noisy polynomial interpolation and noisy chinese remaindering. In *EUROCRYPT*, volume 1807 of *Lecture Notes in Computer Science*, pages 53–69. Springer, 2000.

[BNS13] Dan Boneh, Valeria Nikolaenko, and Gil Segev. Attribute-based encryption for arithmetic circuits. *IACR Cryptol. ePrint Arch.*, 2013:669, 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.

[Bra09] Mark Braverman. Poly-logarithmic independence fools ac^0 circuits. In *IEEE Conference on Computational Complexity*, pages 3–8. IEEE Computer Society, 2009.

[Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical gapsvp. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 868–886. Springer, 2012.

[BS94] Mihir Bellare and Madhu Sudan. Improved non-approximability results. In Frank Thomson Leighton and Michael T. Goodrich, editors, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 184–193. ACM, 1994.

[BTV12] Mihir Bellare, Stefano Tessaro, and Alexander Vardy. Semantic security for the wiretap channel. In *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 294–311. Springer, 2012.

[BV11a] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 97–106. IEEE Computer Society, 2011.

[BV11b] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-lwe and security for key dependent messages. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 505–524. Springer, 2011.

[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.

[CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *J. Cryptol.*, 25(4):601–639, 2012.

[CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, 1998.

[CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In *EUROCRYPT*, volume 1592 of *Lecture Notes in Computer Science*, pages 402–414. Springer, 1999.

[CS03] Don Coppersmith and Madhu Sudan. Reconstructing curves in three (and higher) dimensional space from noisy data. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 136–142, 2003.

[CSP+] Raymond Cheng, William Scott, Bryan Parno, Irene Zhang, Arvind Krishnamurthy, and Thomas Anderson. Talek: a private publish-subscribe protocol.

[DEL+21] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. *CoRR*, abs/2111.04808, 2021.

[DGKS08] Irit Dinur, Elena Grigorescu, Swastik Kopparty, and Madhu Sudan. Decodability of group homomorphisms beyond the johnson bound. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 275–284. ACM, 2008.

[DMN11] Ivan Damgård, Sigurd Meldgaard, and Jesper Buus Nielsen. Perfectly secure oblivious RAM without random oracles. In *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 144–163. Springer, 2011.

[Dod12] Yevgeniy Dodis. Shannon impossibility, revisited. In *ICITS*, volume 7412 of *Lecture Notes in Computer Science*, pages 100–110. Springer, 2012.

[FS95] Katalin Friedl and Madhu Sudan. Some improvements to total degree tests. In *Third Israel Symposium on Theory of Computing and Systems, ISTCS 1995, Tel Aviv, Israel, January 4-6, 1995, Proceedings*, pages 190–198. IEEE Computer Society, 1995.

[FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptol. ePrint Arch.*, 2012:144, 2012.

[GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the cryptographic applications of random functions. In G. R. Blakley and David Chaum, editors, *Advances in Cryptology, Proceedings of CRYPTO '84, Santa Barbara, California, USA, August 19-22, 1984, Proceedings*, volume 196 of *Lecture Notes in Computer Science*, pages 276–288. Springer, 1984.

[GKS06] Elena Grigorescu, Swastik Kopparty, and Madhu Sudan. Local decoding and testing for homomorphisms. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2006 and 10th International Workshop on Randomization and Computation, RANDOM 2006, Barcelona, Spain, August 28-30 2006, Proceedings*, volume 4110 of *Lecture Notes in Computer Science*, pages 375–385. Springer, 2006.

[GKS10] Parikshit Gopalan, Subhash Khot, and Rishi Saket. Hardness of reconstructing multivariate polynomials over finite fields. *SIAM Journal on Computing*, 39(6):2598–2621, 2010.

[GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 660–670. ACM, 2018.

[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.

[GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *J. ACM*, 43(3):431–473, 1996.

[Gol11] Oded Goldreich. A sample of samplers: A computational perspective on sampling. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, pages 302–332. 2011.

[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.

[GR17] Oded Goldreich and Guy N. Rothblum. Simple doubly-efficient interactive proof systems for locally-characterizable sets. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:18, 2017.

[GS99] Venkatesan Guruswami and Madhu Sudan. Improved decoding of reed-solomon and algebraic-geometry codes. *IEEE Trans. Information Theory*, 45(6):1757–1767, 1999.

[GS02] Oded Goldreich and Madhu Sudan. Locally testable codes and pcps of almost-linear length. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 13–22. IEEE Computer Society, 2002.

[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.

[GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 469–477. ACM, 2015.

[Hås96] Johan Håstad. Clique is hard to approximate within $\mathrm{n}^{\text{1-epsilon}}$. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 627–636. IEEE Computer Society, 1996.

[Hås97] Johan Håstad. Some optimal inapproximability results. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 1–10. ACM, 1997.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

[IJK09] Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Approximate list-decoding of direct product codes and uniform hardness amplification. *SIAM J. Comput.*, 39(2):564–605, 2009.

[IJKW10] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. Uniform direct product theorems: Simplified, optimized, and derandomized. *SIAM J. Comput.*, 39(4):1637–1665, 2010.

[KM93] Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the fourier spectrum. *SIAM J. Comput.*, 22(6):1331–1348, 1993.

[KO97] Eyal Kushilevitz and Rafail Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. In *FOCS*, pages 364–373. IEEE Computer Society, 1997.

[KR17] Yael Kalai and Ran Raz. personal communication. 2017.

[KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In F. Frances Yao and Eugene M. Luks, editors, *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 80–86. ACM, 2000.

[Luo09] Bao Luong. *Fourier Analysis on Finite Abelian Groups*. Birkhäuser, 2009.

[Mat93] Mitsuru Matsui. Linear cryptanalysis method for DES cipher. In Tor Helleseth, editor, *Advances in Cryptology - EUROCRYPT '93, Workshop on the Theory and Application of of Cryptographic Techniques, Lofthus, Norway, May 23-27, 1993, Proceedings*, volume 765 of *Lecture Notes in Computer Science*, pages 386–397. Springer, 1993.

[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.

[MP13] Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. *IACR Cryptol. ePrint Arch.*, 2013:69, 2013.

[MRS09] Ben Morris, Phillip Rogaway, and Till Stegers. How to encipher messages on a small domain. In *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 286–302. Springer, 2009.

[NP06] Moni Naor and Benny Pinkas. Oblivious polynomial evaluation. *SIAM J. Comput.*, 35(5):1254–1281, 2006.

[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.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.*, 25(2):252–271, 1996.

[Sam07] Alex Samorodnitsky. Low-degree tests at large distances. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on*

Theory of Computing, San Diego, California, USA, June 11-13, 2007, pages 506–515. ACM, 2007.

[San12] Tom Sanders. On the bogolyubov–ruzsa lemma. *Analysis & PDE*, 5(3):627–655, 2012.

[SCSL11] Elaine Shi, T.-H. Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious RAM with o((logn)3) worst-case cost. In *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, pages 197–214, 2011.

[SS22] Ana Salagean and Pantelimon Stanica. Improving bounds on probabilistic affine tests to estimate the nonlinearity of boolean functions. *Cryptogr. Commun.*, 14(2):459–481, 2022.

[Sud97a] Madhu Sudan. Decoding of reed solomon codes beyond the error-correction bound. *J. Complex.*, 13(1):180–193, 1997.

[Sud97b] Madhu Sudan. Decoding of reed solomon codes beyond the error-correction bound. *J. Complexity*, 13(1):180–193, 1997.

[SW06] Amir Shpilka and Avi Wigderson. Derandomizing homomorphism testing in general groups. *SIAM J. Comput.*, 36(4):1215–1230, 2006.

[Ta-17] Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 238–251. ACM, 2017.

[TW14] Madhur Tulsiani and Julia Wolf. Quadratic goldreich-levin theorems. *SIAM J. Comput.*, 43(2):730–766, 2014.

[Zuc97] David Zuckerman. Randomness-optimal oblivious sampling. *Random Struct. Algorithms*, 11(4):345–367, 1997.

Appendix A

The Goldreich-Levin Theorem in Characteristic q

Here, we prove the Goldreich-Levin Theorem (Chapter 4 Lemma 16) in prime characteristic q. For convenience, we restate the result below as a new theorem.

Theorem 21 (Goldreich-Levin Theorem in Characteristic q) *Let $q, m \in \mathbb{N}$ such that q is prime, and let $\mathbf{s} \in \mathbb{Z}_q^m$ a secret vector. Suppose there exists a function $\mathsf{Pred} : \mathbb{Z}_q^m \to \mathbb{Z}_q$ such that $\Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\mathsf{Pred}(\mathbf{x}) = \langle \mathbf{s}, \mathbf{x} \rangle\big] \geq 1/q + \varepsilon$, for $\varepsilon > 0$. Then, there exists an algorithm Inv, running in time $\mathsf{poly}(m, q, 1/\varepsilon)$, which gets oracle access to Pred and outputs $\mathsf{s}^* \in \mathbb{Z}_q^m$ such that*

$$\Pr_{\mathbf{s}^* \sim \mathsf{Inv}^{\mathsf{Pred}}}\big[\mathbf{s}^* = \mathbf{s}\big] \geq \frac{\varepsilon^2}{8m^3q^5}.$$

Proof. Let $k = \big\lceil \log(8m^2/\varepsilon^2)/\log(q) \big\rceil + 4 \in \mathbb{N}$; observe that $\varepsilon^2/(8m^2q) \leq q^{-(k-4)} \leq \varepsilon^2/(8m^2)$. Let $\{\mathbf{e}_1, \ldots, \mathbf{e}_m\} \subset \mathbb{Z}_q^m$ be the standard basis for $\mathbb{Z}_q^m$.

We now construct the algorithm Inv which gets oracle access to Pred, and works as follows:

1. Choose $\mathbf{x}_1, \ldots, \mathbf{x}_k \sim \mathbb{Z}_q^m$, $\alpha_1, \ldots, \alpha_k \sim \mathbb{Z}_q$. For all $\mathbf{z} \in \mathbb{Z}_q^k$, let $\mathbf{x_z} = \sum\limits_{j=1}^{k} z_j \mathbf{x}_j \in \mathbb{Z}_q^m$ and $\alpha_{\mathbf{z}} = \sum\limits_{j=1}^{k} z_j \alpha_j \in \mathbb{Z}_q$.
2. For all $i \in [m]$, compute $y_i = \text{Plurality}_{\mathbf{z} \in \mathbb{Z}_q^k \setminus \{\mathbf{0}\}} \{\mathsf{Pred}(\mathbf{e}_i + \mathbf{x_z}) - \alpha_{\mathbf{z}}\} \in \mathbb{Z}_q$.
3. Output $\mathbf{y} := (y_i)_{i=1}^m \in \mathbb{Z}_q^m$.

It's clear that the running time of Inv is $\mathsf{poly}(m, q, 1/\varepsilon)$. The high-level idea is that each $\alpha_j \sim \mathbb{Z}_q$ is a guess for $\langle \mathbf{x}_j, \mathbf{s} \rangle \in \mathbb{Z}_q$ that is correct with probability $1/q$. So, with probability q^{-k} it holds that each $\alpha_{\mathbf{z}} = \langle \mathbf{x_z}, \mathbf{s} \rangle \in \mathbb{Z}_q$ ($\mathbf{z} \in \mathbb{Z}_q^k$). Next, if $\forall i \in [m], \gamma \in \mathbb{Z}_q^*$,

$$\#\{\mathbf{z} : \mathsf{Pred}(\mathbf{e}_i + \mathbf{x_z}) = \langle \mathbf{e}_i + \mathbf{x_z}, \mathbf{s} \rangle\} > \#\{\mathbf{z} : \mathsf{Pred}(\mathbf{e}_i + \mathbf{x_z}) = \langle \mathbf{e}_i + \mathbf{x_z}, \mathbf{s} \rangle + \gamma\},$$

then each $y_i = \langle \mathbf{e}_i + \mathbf{x}_{\mathbf{z}^*}, \mathbf{s} \rangle - \alpha_{\mathbf{z}^*} = s_i$, for some $\mathbf{z}^* \in \mathbb{Z}_q^k \setminus \{\mathbf{0}\}$, as desired. We now proceed formally.

Let $\mathcal{D}$ denote the distribution of the randomness of Inv; namely, $\mathcal{D}$ chooses $\mathbf{x}_1, \ldots, \mathbf{x}_k \sim \mathbb{Z}_q^m$, $\alpha_1, \ldots, \alpha_k \sim \mathbb{Z}_q$, and outputs $(\{\mathbf{x}_j\}_j, \{\alpha_j\}_j)$. We begin by establishing the following quantities, the last two of which implicitly are a function of the randomness $(\{\mathbf{x}_j\}_j, \{\alpha_j\}_j) \sim \mathcal{D}$.

- For all $\gamma \in \mathbb{Z}_q$, let $\mathrm{P}(\gamma) := \Pr_{\mathbf{x} \sim \mathbb{Z}_q^m}\big[\mathsf{Pred}(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s} \rangle + \gamma\big]$.
- For all $i \in [m], \gamma \in \mathbb{Z}_q$, $\mathbf{z} \in \mathbb{Z}_q^k$, let $\mathrm{E}_i(\mathbf{z}, \gamma)$ denote the event "$\mathsf{Pred}(\mathbf{e}_i + \mathbf{x_z}) = \langle \mathbf{e}_i + \mathbf{x_z}, \mathbf{s} \rangle + \gamma$".
- For all $i \in [m], \gamma \in \mathbb{Z}_q$, let $\mathrm{Q}_i(\gamma) := \Pr_{\mathbf{z} \sim \mathbb{Z}_q^k \setminus \{\mathbf{0}\}}\big[\mathrm{E}_i(\mathbf{z}, \gamma)\big]$.

It suffices to show that $\forall i \in [m], \gamma \in \mathbb{Z}_q^*$, $\Pr_{\mathcal{D}}\big[\mathrm{Q}_i(0) > \mathrm{Q}_i(\gamma)\big] \geq 1 - 8q^{-(k-3)}/\varepsilon^2$, since then

$$\Pr_{\mathcal{D}}\big[\mathbf{y} = \mathbf{s}\big] \geq \Pr_{\mathcal{D}}\big[\text{Ea. } \alpha_j = \langle \mathbf{x}_j, \mathbf{s}\rangle \;\&\; \mathrm{Q}_i(0) > \mathrm{Q}_i(\gamma), \forall i \in [m], \gamma \in \mathbb{Z}_q^*\big]$$

$$\geq q^{-k}(1 - 8mq^{-(k-4)}/\varepsilon^2) \geq (\varepsilon^2 m^{-2} q^{-5}/8)(1 - 1/m) \geq \varepsilon^2 m^{-3} q^{-5}/8.$$

The following two claims complete the proof.

Claim 21 *For all* $\gamma \in \mathbb{Z}_q^*$, $\mathrm{P}(0) - \mathrm{P}(\gamma) > \varepsilon/q$.

Proof. By hypothesis, we have that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}(\mathbf{x}) = \langle \mathbf{s}, \mathbf{x}\rangle\big] \geq 1/q + \varepsilon$. We'll show there exists another prediction algorithm Pred' which has the guarantee that $\forall \gamma \in \mathbb{Z}_q^*$,

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s}\rangle\big] - \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s}\rangle + \gamma\big] > \varepsilon/q.$$

Hence to complete the proof it suffices to redefine Pred as Pred'.

Now, define $\mathsf{Pred}' : \mathbb{Z}_q^m \to \mathbb{Z}_q$ where $\mathsf{Pred}'(\mathbf{x})$ chooses $\alpha \sim \mathbb{Z}_q^*$, and outputs $\alpha^{-1} \cdot \mathsf{Pred}(\alpha\mathbf{x}) \in \mathbb{Z}_q$. First, it's clear that $\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s}\rangle\big] = \mathrm{P}(0) \geq 1/q + \varepsilon$. Next, let $\gamma \in \mathbb{Z}_q^*$. We have

$$\Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s}\rangle + \gamma\big] = \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\gamma^{-1} \cdot \mathsf{Pred}'(\mathbf{x}) = \langle \gamma^{-1}\mathbf{x}, \mathbf{s}\rangle + 1\big]$$

$$= \Pr_{\bar{\mathbf{x}}\sim\mathbb{Z}_q^m}\big[\gamma^{-1} \cdot \mathsf{Pred}'(\gamma\bar{\mathbf{x}}) = \langle \bar{\mathbf{x}}, \mathbf{s}\rangle + 1\big] = \Pr_{\bar{\mathbf{x}}\sim\mathbb{Z}_q^m, \alpha\sim\mathbb{Z}_q^*}\big[\alpha^{-1}\gamma^{-1} \cdot \mathsf{Pred}(\alpha\gamma\bar{\mathbf{x}}) = \langle \bar{\mathbf{x}}, \mathbf{s}\rangle + 1\big]$$

$$= \Pr_{\bar{\mathbf{x}}\sim\mathbb{Z}_q^m, \bar{\alpha}\sim\mathbb{Z}_q^*}\big[\bar{\alpha}^{-1} \cdot \mathsf{Pred}(\bar{\alpha}\bar{\mathbf{x}}) = \langle \bar{\mathbf{x}}, \mathbf{s}\rangle + 1\big] = \Pr_{\bar{\mathbf{x}}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\bar{\mathbf{x}}) = \langle \bar{\mathbf{x}}, \mathbf{s}\rangle + 1\big].$$

Hence $\mathrm{R} := \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s}\rangle + 1\big] = \Pr_{\mathbf{x}\sim\mathbb{Z}_q^m}\big[\mathsf{Pred}'(\mathbf{x}) = \langle \mathbf{x}, \mathbf{s}\rangle + \gamma\big], \forall \gamma \in \mathbb{Z}_q^*$, and we have $1 \geq 1/q + \varepsilon + (q-1)\mathrm{R}$, which implies that $\mathrm{R} \leq (1 - 1/q - \varepsilon)/(q-1) = 1/q - \varepsilon/(q-1) < \mathrm{P}(0) - \varepsilon/(q-1)$. So $\mathrm{P}(0) - \mathrm{R} > \varepsilon/(q-1) \geq \varepsilon/q$. ■

Claim 22 *For all $i \in [m], \gamma \in \mathbb{Z}_q^*$, $\mathrm{Pr}_{\mathcal{D}}\big[\mathrm{Q}_i(0) > \mathrm{Q}_i(\gamma)\big] \geq 1 - 8q^{-(k-3)}/\varepsilon^2$.*

Proof. Let $i \in [m]$ and $\gamma \in \mathbb{Z}_q$. We have

$$
\begin{aligned}
&\mathrm{Pr}_{\mathcal{D}}\left[\Big|\mathrm{Q}_i(\gamma) - \mathbb{E}_{\mathcal{D}}\big[\mathrm{Q}_i(\gamma)\big]\Big| > \frac{\varepsilon}{2q}\right] \leq \frac{4q^2}{\varepsilon^2}\Big(\mathbb{E}_{\mathcal{D}}\big[\mathrm{Q}_i(\gamma)^2\big] - \mathbb{E}_{\mathcal{D}}\big[\mathrm{Q}_i(\gamma)\big]^2\Big) \\
&= \frac{4q^2}{\varepsilon^2}\left(\mathbb{E}_{\mathcal{D}}\Big[\mathrm{Pr}_{\mathbf{z},\mathbf{z}'\sim\mathbb{Z}_q^k\setminus\{\mathbf{0}\}}\big[\mathrm{E}_i(\mathbf{z},\gamma)\ \&\ \mathrm{E}_i(\mathbf{z}',\gamma)\big]\Big] - \mathrm{P}(\gamma)^2\right) \\
&= \frac{4q^2}{\varepsilon^2}\left(\mathrm{Pr}_{\mathbf{z},\mathbf{z}'\sim\mathbb{Z}_q^k\setminus\{\mathbf{0}\}}\big[\mathbf{z} = \mathbf{z}'\big] \cdot \mathrm{P}(\gamma) + \Big(1 - \mathrm{Pr}_{\mathbf{z},\mathbf{z}'\sim\mathbb{Z}_q^k\setminus\{\mathbf{0}\}}\big[\mathbf{z} = \mathbf{z}'\big]\Big)\mathrm{P}(\gamma)^2 - \mathrm{P}(\gamma)^2\right) \\
&= \frac{4q^2}{\varepsilon^2}\left(\frac{1}{q^k - 1}\Big(\mathrm{P}(\gamma) - \mathrm{P}(\gamma)^2\Big)\right) \leq \frac{4}{\varepsilon^2 q^{k-3}},
\end{aligned}
$$

where the first inequality follows from Chebyshev's inequality; the second line follows from $\mathbb{E}_{\mathcal{D}}\big[\mathrm{Q}_i(\gamma)\big] = \mathrm{P}(\gamma)$; and the third line follows from pairwise independence. Now, let $\gamma \in \mathbb{Z}_q^*$. We have that

$$\mathrm{Q}_i(\gamma) \leq \mathrm{P}(\gamma) + \varepsilon/(2q) < \mathrm{P}(0) - \varepsilon/(2q) \leq \mathrm{Q}_i(0)$$

with probability $1 - 8q^{-(k-3)}/\varepsilon^2$ over $\mathcal{D}$, by the above argument and Claim 21. ■ ■