

Tighter Nonuniform Trace Reductions from Ring-LWE to MP-LWE

Yihang Cheng¹, Xiangzheng Zhao¹, Hengyi Luo^{2,3},
Yanbin Pan^{2,3}, and Siwei Sun¹

¹ School of Cryptology, University of Chinese Academy of Sciences
chengyihang15@mails.ucas.ac.cn
zhaoxiangzheng24@mails.ucas.ac.cn
sunswei@ucas.ac.cn

² Academy of Mathematics and System Science, Chinese Academy of Sciences
fallingtime2019@outlook.com, panyanbin@amss.ac.cn

³ University of Chinese Academy of Sciences

Abstract. Middle-Product LWE (MP-LWE) can inherit hardness from Ring-LWE over many number fields. Peikert and Pepin developed a direct reduction using field traces. Njah Nchiwo and Pellet-Mary (NP26) proved polynomial loss for the route through Polynomial-LWE under an additional condition on the defining polynomial. Recently, Pellet-Mary and Xia (PX26) unified both routes and proved polynomial loss for defining polynomials with polynomially bounded coefficients under suitable modulus conditions.

We analyze Peikert and Pepin’s trace reduction for monic irreducible integer polynomials f of degree $n \geq 2$ with coefficients of absolute value at most a fixed $B \geq 1$. First, when the prime modulus satisfies $q \nmid [\mathcal{O}_{\mathbb{Q}(\theta)} : \mathbb{Z}[\theta]]$, where θ is a root of f , we prove the existence of multipliers with primal loss $O_B(n\sqrt{\log n / \log \log n})$ and dual loss $O_B(n^2)$. The multipliers have polynomial-size descriptions and serve as advice. Second, whenever the polynomial class is nonempty, every $Q \geq 21$ admits a prime $q \in [Q, 2Q]$ for which the index condition holds for at least a $1 - O_B(n \log n / Q)$ fraction of the polynomials. At the same prime, we prove the same lower bound when counting the distinct number fields defined by these polynomials. Both coverage bounds tend to one as $n \rightarrow \infty$ if $Q/(n \log n) \rightarrow \infty$. Third, we prove that the trace reduction can match the linear noise bounds achieved by NP26’s reduction through Polynomial-LWE. For the same defining polynomial, Ring-LWE variant, and MP-LWE parameters, a suitable trace multiplier exists whose linear noise amplification is no larger than that achieved by NP26’s multipliers.

Keywords: Ring-LWE · Middle-Product LWE · Ideal lattices · Noise loss

1 Introduction

Ring-LWE is a variant of the Learning With Errors (LWE) problem over the ring of integers of a number field. The choice of field is part of the problem, and the hardness of Ring-LWE over different fields is not generally understood. Middle-Product LWE (MP-LWE) offers a way to reduce dependence on the choice of fields. An attack on one MP-LWE problem can yield attacks on Ring-LWE over many fields. If at least one of the Ring-LWE problems covered by the reductions is hard, then the MP-LWE problem is also hard. The argument does not require every covered Ring-LWE problem to be hard.

MP-LWE keeps consecutive middle coefficients of a polynomial product, without reducing the product modulo a defining polynomial. Roşca et al. [1] used this arithmetic to construct quasi-linear public-key encryption. Their reductions connect MP-LWE to Polynomial-LWE (PLWE) over many defining polynomials. PLWE uses noise in polynomial coefficients, whereas Ring-LWE measures noise in the canonical geometry of a number field. Relating the two error distributions is therefore part of obtaining Ring-LWE-based guarantees for MP-LWE.

As MP-LWE constructions developed, their parameter guarantees became more important. Titanium introduced practical encryption and key encapsulation [2]. A rounding variant replaced Gaussian sampling in encryption [3]. Trappeddoors then enabled encryption using identities as public keys [4]. Later work added hierarchical delegation [5], adaptive identity security [6], and smaller master public keys [7]. Further constructions support authorized inner products [8,9]. For signatures, MPSign [10] repaired earlier proof gaps [11]. Other work developed and shortened linkable ring signatures [12,13]. These applications make the same underlying question important: which Ring-LWE problems support the particular MP-LWE parameters used by a construction?

Several works study reductions from Ring-LWE to MP-LWE. Some reductions pass through PLWE [1,14,15]. Peikert and Pepin reduce Ring-LWE directly to MP-LWE [16]. Recently, Pellet-Mary and Xia placed both routes in a unified reduction framework [17]. We first explain the reductions and the restriction in the earlier analysis.

1.1 The reduction framework and its noise loss

Let $R = \mathcal{O}_K$ be the ring of integers of a degree- n number field K , and let q be the integer modulus. In primal Ring-LWE, the secret $s \in R/qR$ is sampled uniformly once. Every sample uses the same secret s and an independently sampled uniform public element $a \in R/qR$. MP-LWE uses the public polynomial $A(x) = \sum_{j=0}^{n-1} A_j x^j$ and secret polynomial $S(x) = \sum_{k=0}^{2n-2} s_k x^k$ in the single variable x . Their scalar coefficients A_j, s_k are independent uniform elements of $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$. We sample a fresh public polynomial $A(x)$ for each sample and keep the same secret polynomial $S(x)$ across samples. We retain n middle coefficients

in each MP-LWE sample. The noisy parts of the samples are

$$\begin{aligned} \text{Ring-LWE: } b &= as + e \pmod{qR}, \\ \text{MP-LWE: } b'_i &= (AS)_{n-1+i} + e'_i \pmod{q}, \quad 0 \leq i < n. \end{aligned}$$

Here $(AS)_k$ is the coefficient of x^k in the ordinary product AS . MP-LWE retains coefficients $n-1, \dots, 2n-2$, without reducing modulo a defining polynomial. The Gaussian errors are measured in K 's canonical geometry for e and in coefficient coordinates for $\mathbf{e}' = (e'_i)_i$. A decision reduction converts Ring-LWE samples to MP-LWE samples and preserves uniform inputs. An MP-LWE distinguisher then distinguishes Ring-LWE.

Where the noise loss comes from. Let T be the linear map on the error. Define

$$\|T\|_{\text{op}} := \sup \left\{ \frac{\|T\mathbf{x}\|_2}{\|\mathbf{x}\|_2} : \mathbf{x} \neq 0 \right\}.$$

For spherical input of standard deviation αq , output standard deviation is at most $\alpha q \|T\|_{\text{op}}$ in any direction. Given $C \geq \|T\|_{\text{op}}$, we can add Gaussian noise to reach spherical MP-LWE noise of standard deviation βq whenever $\beta \geq \alpha C$.

Ring-LWE to MP-LWE via PLWE. Earlier reductions [14,1] follow $\text{RLWE} \rightarrow \text{PLWE} \rightarrow \text{MP-LWE}$. The first step converts Ring-LWE samples to PLWE samples, with errors expressed in polynomial coefficients. The second step applies a PLWE-to-MP-LWE reduction. Both steps can amplify the error. Let f be a monic irreducible defining polynomial for K , and write $K = \mathbb{Q}(\theta)$ with $f(\theta) = 0$. Before changing error coordinates, the reduction converts samples over R to samples over the polynomial order $\mathbb{Z}[\theta]$. A suitable short multiplier μ_2 defines the sample map $(a, b) \mapsto (\mu_2 a, \mu_2^2 b)$, which changes the secret to $\mu_2 s$ and the error to $\mu_2^2 e$. For dual Ring-LWE, another short multiplier μ_1 first converts dual samples to primal samples by $(a, b) \mapsto (a, \mu_1 b)$. After conversion from dual Ring-LWE to samples over $\mathbb{Z}[\theta]$, the error is $\mu_2^2 \mu_1 e$. We write $\mathbf{e}_{\text{alg}}$ for the canonical embedding coordinates of $\mu_2^2 e$ in the primal case and $\mu_2^2 \mu_1 e$ in the dual case. For primes $q \geq 8n$ with $q \nmid f(0) \text{Disc}(f)$, NP26 proves that suitable short multipliers μ_1, μ_2 exist and uses them as advice [15]. Here $\text{Disc}(f)$ is the polynomial discriminant. The multipliers μ_2 and μ_1 implement the conversion from R to $\mathbb{Z}[\theta]$ and the dual-to-primal conversion, respectively. NP26 proves a polynomial bound on the noise amplification caused by multiplying e by μ_2^2 or $\mu_2^2 \mu_1$ when the coefficients of f are polynomially bounded.

Let $\sigma_1, \dots, \sigma_n : K \rightarrow \mathbb{C}$ be the embeddings of K . The values $\sigma_j(\theta)$ are the complex roots of f . We extend each σ_j by real linearity when applying σ_j to the Gaussian error. The evaluation matrix $(V_f)_{j,i} = \sigma_j(\theta)^i$, $1 \leq j \leq n$, $0 \leq i < n$, maps polynomial coefficients to values at these roots. Its inverse converts the embedding-coordinate error to polynomial coefficients. Let M_f map PLWE coefficient error to MP-LWE error before the output coordinates are reversed. The remaining error transformation is

$$\mathbf{e}_{\text{alg}} \xrightarrow{V_f^{-1}} V_f^{-1} \mathbf{e}_{\text{alg}} \xrightarrow{M_f} M_f V_f^{-1} \mathbf{e}_{\text{alg}}.$$

Earlier analyses bound the two maps separately and use

$$\left\| M_f V_f^{-1} \right\|_{\text{op}} \leq \|M_f\|_{\text{op}} \left\| V_f^{-1} \right\|_{\text{op}}.$$

This multiplies the largest stretch of each map. However, a vector stretched most by V_f^{-1} need not be sent to a direction stretched most by M_f . Multiplying the separate bounds can therefore overestimate the total noise loss.

Njah Nchiwo and Pellet-Mary (NP26) [15] instead bound the operator norm of the composed map $M_f V_f^{-1}$ directly. Each MP-LWE sample considered here retains n consecutive middle coefficients, so $M_f V_f^{-1}$ is an $n \times n$ matrix. NP26 denotes the largest absolute entry of this matrix by $D(f)$ and bounds its operator norm as follows:

$$D(f) = \max_{i,j} |(M_f V_f^{-1})_{ij}|, \quad \left\| M_f V_f^{-1} \right\|_{\text{op}} \leq nD(f).$$

For a fixed coefficient bound and input noise rate $0 < \alpha < 1$, NP26 proves primal loss $O(n^{14})$ and dual loss $O(n^{22})$ under the condition $D(f) \leq n$. This condition requires every entry of $M_f V_f^{-1}$ to have absolute value at most n . NP26 observes $D(f) \leq n$ for at least 90% of its random test polynomials. This proportion is experimental. NP26 also exhibits bounded-coefficient families with exponential $D(f)$, so bounded coefficients alone do not imply its entry bound.

Other developments extend the coefficient route. Titanium preserves coefficient noise under restrictions on the polynomials and output length [2]. Universal products enlarge the PLWE polynomial class [18]. They still need the Ring-LWE conversion. Skew middle products extend the approach to noncommutative algebras [19]. A recent survey [20] reviews these developments.

Ring-LWE to MP-LWE using traces. Peikert and Pepin (TCC 2019, journal version PP24) [16] obtain the middle-product relation directly from the Ring-LWE product as . The trace $\text{Tr}(x) = \sum_j \sigma_j(x)$ sums the embedding values. Let $s \in R$ represent the primal Ring-LWE secret above. Let $c \in K \setminus \{0\}$ satisfy $\text{Tr}(cu) \in \mathbb{Z}$ for every $u \in R$. Then $\text{Tr}(cs\theta^k)$ is an integer for every $k \geq 0$. We use these traces modulo q to construct the MP-LWE secret. The residues do not depend on the representative of s . Assume q is prime and $q \nmid [R : \mathbb{Z}[\theta]]$. Every public element $a \in R/qR$ then has a representative $a' = \sum_{j=0}^{n-1} a'_j \theta^j$ with integer a'_j and $a' \equiv a \pmod{qR}$. Define two polynomials over $\mathbb{Z}_q$:

$$A(x) = \sum_{j=0}^{n-1} a'_j x^{n-1-j},$$

$$\tilde{S}(x) = \sum_{k=0}^{2n-2} \tilde{S}_k x^k, \quad \tilde{S}_k = \text{Tr}(cs\theta^k) \bmod q.$$

The public polynomial A lists the coefficients of a' in reverse order. For $0 \leq i < n$, the product of x^{n-1-j} and x^{i+j} is x^{n-1+i} . Therefore the coefficient of x^{n-1+i} in

$A\tilde{S}$ satisfies

$$(A\tilde{S})_{n-1+i} = \sum_{j=0}^{n-1} a'_j \tilde{S}_{i+j} = \text{Tr}(ca's\theta^i) \pmod{q}.$$

Thus the n traces of the product $a's$ are exactly the n middle coefficients required by MP-LWE. To make the secret uniform, draw one uniform polynomial v over $\mathbb{Z}_q$ of degree below $2n-1$ for all samples and set $S = \tilde{S} + v$. Adding the middle coefficients of Av to each transformed sample changes its secret from $\tilde{S}$ to S .

The trace transformation maps the error e to

$$L_{c,n}(e)_i = \text{Tr}(ce\theta^i) = \sum_{j=1}^n \sigma_j(c) \sigma_j(\theta)^i \sigma_j(e), \quad 0 \leq i < n.$$

Section 3.1 states this specialization of PP24's reduction. Our task is to bound $\|L_{c,n}\|_{\text{op}}$.

With $c = 1$, the noise amplification is bounded by $\|V_f\|_{\text{op}}$. Large powers of roots whose absolute values exceed 1 can make this bound exponential. PP24 proves polynomial bounds when the nonleading terms have degree at most $(1 - \delta)n$ for fixed $\delta > 0$ and the coefficients are bounded [16]. Cyclotomic and binomial examples admit smaller bounds [21,22,23].

A unified framework. Recently, Pellet-Mary and Xia (PX26) [17] placed the reduction through PLWE and the direct trace reduction in one framework. Their reduction uses two field elements $y_a, y_s \in K \setminus \{0\}$. The first multiplier converts the public Ring-LWE element to the polynomial order. The second multiplier scales the secret. Their product acts on the error, so the resulting linear error map is

$$e \mapsto (\text{Tr}(y_a y_s e \theta^i))_{i=0}^{n-1} = L_{y_a y_s, n}(e),$$

up to a reversal of the output coordinates, which preserves the norm. Specific choices of the two multipliers recover the earlier reductions. PX26 then uses a weighted lattice argument to prove the existence of multipliers with polynomial loss. The weights compensate for large root powers by requiring smaller multiplier values at the corresponding embeddings. For every monic irreducible degree- n defining polynomial with polynomially bounded coefficients, their general bound applies at primes $q \geq 8n$ satisfying $q \nmid [R : \mathbb{Z}[\theta]]$. PX26 states this arithmetic condition as coprimality with the conductor ideal of $\mathbb{Z}[\theta]$. The two conditions are equivalent, as explained in Section 1.3. The general reduction is nonuniform because the required multipliers are proved to exist but are not found by a general polynomial-time algorithm.

1.2 Our contributions

We developed this work independently of PX26. We analyze Peikert and Pepin's existing trace reduction and do not propose a new reduction framework. Our

analysis shows that the existing trace framework already achieves polynomial loss for every bounded-coefficient defining polynomial under the same polynomial-dependent modulus condition as PX26. Both works use weighted estimates of the products $\sigma_j(c)\sigma_j(\theta)^i$ to prove that a suitable multiplier exists. We obtain smaller general loss bounds by also controlling the interaction between the columns of the trace matrix. Both works therefore extend the polynomial-loss guarantee beyond the polynomial families covered by the earlier bounds. We further quantify coverage at one prime modulus: the proportions of bounded-coefficient polynomials and distinct represented fields covered by the reduction tend to one under our modulus choices. PX26 proves a reduction for every polynomial satisfying its modulus condition, but does not state this counting result.

Let $B \geq 1$ be fixed. We write $\mathcal{I}_{n,B}$ for the monic irreducible degree- n integer polynomials with coefficient magnitudes at most B . For $f \in \mathcal{I}_{n,B}$, set $K_f = \mathbb{Q}(\theta)$ with $f(\theta) = 0$. We analyze the error map in PP24's trace reduction and prove the following three results. Section 1.3 then compares these results with NP26 and PX26.

1. Polynomial loss for every bounded-coefficient defining polynomial. For every $f \in \mathcal{I}_{n,B}$ and every prime q with

$$q \nmid j_f, \quad j_f = [\mathcal{O}_{K_f} : \mathbb{Z}[\theta]],$$

we prove that legal multipliers exist with primal loss $O_B(n\sqrt{\log n / \log \log n})$ and dual loss $O_B(n^2)$ (Theorem 1). The input and output noise rates are matched using these bounds. After this choice, $q \nmid j_f$ is the remaining polynomial-dependent admissibility condition. It makes the polynomial order and the ring of integers isomorphic modulo q .

The trace-map entries $\sigma_j(c)\sigma_j(\theta)^i$ have the same product form as the entries in NP26's composed error map. This similarity motivates us to apply the idea of analyzing the complete error map directly to PP24's trace reduction. A large power $\sigma_j(\theta)^i$ need not cause large noise if $|\sigma_j(c)|$ is correspondingly small. We therefore seek a legal multiplier whose embedding values are small where the root powers are large. These values must come from one field element, so we cannot choose them independently at different roots. We use different bounds on $|\sigma_j(c)|$ at different embeddings. One part of each bound compensates for the size of the root powers. A second part accounts for the interaction between different columns of the trace matrix. We control these interactions with a positive semidefinite matrix whose entries are positive. The row sums of this matrix determine the second set of weights. Minkowski's convex body theorem then proves that one legal multiplier satisfies all the weighted bounds. The resulting volume depends on geometric means of the weights, rather than their largest values.

To bound the row sums, we replace each root of absolute value greater than one by the reciprocal of its complex conjugate. All the resulting points lie in the closed unit disk. We relate the squared sums of powers of these points to the discriminant of f . The primal argument compares positive Poisson kernels at different radii. The comparison prevents the power sums from growing too

quickly near the unit circle and yields the factor $\sqrt{\log n / \log \log n}$. For the dual argument, the field discriminant appears both in the multiplier lattice volume and in the bound on the power sums. We keep these two occurrences together and optimize the resulting expression. This proves the dual $O_B(n^2)$ bound without a logarithmic factor. Section 3 proves each of these steps.

We prove that the multiplier has polynomial bit length (Lemma 5). Its description depends on f and the Ring-LWE variant, independently of q and the samples. Given this advice, the reduction runs in polynomial time. The general multiplier search remains open. We also derive the cryptographic effect of the smaller loss (Proposition 5 and Corollary 1): at fixed MP-LWE noise rate and sample count, the bound allows more Ring-LWE noise and a smaller sufficient modulus in the composed ideal-lattice reduction.

2. A coverage bound for polynomials and number fields. The first result applies to every $f \in \mathcal{I}_{n,B}$ once the modulus avoids j_f . We next bound the proportion of polynomials satisfying this condition at one prime modulus. For every degree n with $\mathcal{I}_{n,B} \neq \emptyset$ and every real $Q \geq 21$, we prove that there is a prime $q \in [Q, 2Q]$ such that

$$\frac{|\{f \in \mathcal{I}_{n,B} : q \nmid j_f\}|}{|\mathcal{I}_{n,B}|} \geq 1 - O_B\left(\frac{n \log n}{q}\right)$$

(Theorem 2). The size parameter Q can be chosen before applying the theorem. The theorem guarantees a suitable prime in $[Q, 2Q]$, rather than the same bound for every prime in that interval. For every $f \in \mathcal{I}_{n,B}$ with $q \nmid j_f$, the reduction uses the same target modulus, dimensions, noise rate, sample bound, and encoding at degree n . The proof uses the fact that each index j_f has few large prime divisors. Averaging over primes in $[Q, 2Q]$ bounds the excluded proportion.

For the asymptotic result, choose Q_n with $Q_n/(n \log n) \rightarrow \infty$ and choose a prime $q_n \in [Q_n, 2Q_n]$ that satisfies the displayed bound. Then $q_n/(n \log n) \rightarrow \infty$, so the excluded proportion tends to zero. For the connection to ideal-lattice hardness, we choose polynomial-size Q_n that also satisfy the sufficient conditions in Section 4.2. Corollaries 1 and 3 establish coverage and the ideal-lattice reductions at these parameters. The limits are taken along degrees with $\mathcal{I}_{n,B} \neq \emptyset$. Exceptions may remain at every finite degree.

We prove the same lower bound for the proportion of distinct fields represented by $\mathcal{I}_{n,B}$, counting isomorphic fields once. Different generators can define the same field, so a polynomial count may count one field many times. We apply the prime-divisor argument also to one chosen polynomial per field. A single prime makes both excluded proportions small. For any sequence $f_n \in \mathcal{I}_{n,B}$ with $q_n \nmid j_{f_n}$, Corollary 1 reduces ideal-SIVP over K_{f_n} to MP-LWE. If ideal-SIVP at the approximation factors in that corollary is hard for nonuniform quantum algorithms over at least one such sequence of fields, then MP-LWE is hard at the shared parameters.

3. Retaining smaller loss bounds in the trace framework. For the reduction through PLWE described above, we prove that whenever multipliers satisfying

that reduction's conditions make the complete linear error map have norm at most C , a legal trace multiplier exists whose error map also has norm at most C (Proposition 2). The complete map includes multiplication of the input Ring-LWE error during conversion to samples over the polynomial order. The finite trace reduction retains the loss guarantee up to a constant factor.

We show that NP26's linear error transformation can be written as $e'_i = \text{Tr}_{K/\mathbb{Q}}(ce\theta^i)$, after placing the output coordinates in the same order. Here e is the original Ring-LWE error, and the same legal trace multiplier c is used for every coordinate. This proves equality of the linear error maps after reversing the order of the output coordinates. The identity follows from the algebraic form of the maps and does not require the weighted norm estimate. PX26 also recovers NP26's complete error map within its unified framework [17]. Here we state the transfer within PP24's trace framework and additionally establish the modulus-independent choice below.

For fixed f , the Ring-LWE variant, and n output coefficients, we also prove that a legal multiplier minimizing the operator norm of the trace map exists and has a polynomial-size description independent of the modulus (Corollary 2). Finding this multiplier efficiently remains open.

1.3 Comparison with NP26 and PX26

We compare reductions for fixed B , with n output coefficients and spherical input noise of rate $0 < \alpha < 1$. For PX26, we also set the degree of the defining polynomial and the length of the public MP-LWE polynomial equal to n . NP26, PX26, and our general result use nonuniform multiplier advice. The table lists published general loss bounds for NP26 and PX26.

	NP26 [15]	PX26 [17]	This work
Primal loss	$O_B(n^{14})$	$O_B(n^8)$	$O_B(n\sqrt{\frac{\log n}{\log \log n}})$
Dual loss	$O_B(n^{22})$	$O_B(n^8)$	$O_B(n^2)$
Entry condition	$D(f) \leq n$	None	None
Prime modulus	$q \nmid f(0) \text{ Disc}(f)$ $q \geq 8n$	$q \nmid j_f$ $q \geq 8n$	$q \nmid j_f$

PX26's explicit general factor is $16n^8B^5$. Their more precise bound is

$$8n^3\Delta_f^{2/n}\mathcal{M}(f),$$

where Δ_f is the absolute polynomial discriminant and $\mathcal{M}(f)$ is the Mahler measure. These are upper bounds, not lower bounds on the loss attainable in their framework. Our bounds improve the published general estimates.

The polynomial-dependent modulus conditions in PX26 and this work are equivalent. Write $\mathfrak{c}_f = \{a \in R : aR \subseteq \mathbb{Z}[\theta]\}$ for the conductor ideal. The inclusions $j_f R \subseteq \mathfrak{c}_f \subseteq \mathbb{Z}[\theta] \subseteq R$ imply $j_f \mid [R : \mathfrak{c}_f] \mid j_f^n$. Thus, for a prime q , the condition $\mathfrak{c}_f + qR = R$ used by PX26 is equivalent to $q \nmid j_f$. PX26 also requires $q \geq 8n$ in its general multiplier-existence bound. Both works establish

polynomial loss without NP26's entry condition. Our additional counting theorem quantifies how many polynomials and distinct fields meet the arithmetic condition at one selected prime.

The improvement in the loss has two parts. PP24's order transformation allows us to work directly with all legal trace multipliers. For dual Ring-LWE, these are the nonzero elements of R . PX26's dual reduction requires $y_a \in \mathfrak{c}_f$ and $y_s \in R$. Their product therefore belongs to $\mathfrak{c}_f$, which may be a proper subideal of R . Within PP24's legal multiplier space, using only weights determined by root magnitudes already proves primal $O_B(n^{3/2})$ and dual $O_B(n^{5/2})$ bounds. These follow from Lemma 2, the estimates in (4), and the bound of the matrix norm by n times its largest absolute entry. We derive these intermediate bounds to show how much the trace-matrix analysis improves the loss. PX26 does not state these bounds. Our analysis of interactions between matrix columns then improves these bounds to the two loss estimates in Theorem 1. The latter improvement uses adaptive weights, Poisson kernels, and a joint estimate involving the discriminant.

Under NP26's reduction conditions, we prove that a legal trace multiplier exists whose error map has operator norm no larger than the operator norm of NP26's complete linear error map (Proposition 2). PX26 also includes NP26's error map as a special case. Our modulus-independent minimizing multiplier is stated in Corollary 2.

2 Preliminaries

We compute exactly in an algebraic basis and measure Gaussian error in the canonical embedding. Conversion can amplify noise, so we specify both the basis and the norm. We define the algebra and geometry first, then the decision problems and reduction guarantees.

All logarithms are natural. We write $\log^+ x = \max(0, \log x)$. As $n \rightarrow \infty$, O, Ω, Θ denote upper, lower, and two-sided bounds up to positive constants. Subscripts indicate dependence on fixed parameters. For $u, v > 0$, $u = o(v)$ means $u/v \rightarrow 0$ and $u = \omega(v)$ means $u/v \rightarrow \infty$.

2.1 Number fields and trace maps

Definition 1 (Defining polynomials and number fields). *For integers $n \geq 2$ and $B \geq 1$, let $\mathcal{I}_{n,B}$ consist of the monic $f = x^n + \sum_{i=0}^{n-1} a_i x^i \in \mathbb{Z}[x]$ irreducible over $\mathbb{Q}$ with $|a_i| \leq B$ for $0 \leq i < n$. Put $K_f = \mathbb{Q}[x]/(f)$, abbreviate $K = K_f$, and set $\theta = x \bmod f$. The degree- n field K has power basis $(1, \theta, \dots, \theta^{n-1})$.*

A $\mathbb{Q}$ -basis determines unique rational coordinates in K . A $\mathbb{Z}$ -basis determines unique integer coordinates in the additive subgroup generated by the basis.

Definition 2 (Orders). *The ring of integers $R = \mathcal{O}_K$ consists of elements of K that satisfy a monic integer polynomial. An order is a subring of R of $\mathbb{Z}$ -rank n containing 1. We use $\mathcal{O}_f = \mathbb{Z}[\theta]$ and index $j_f = [R : \mathcal{O}_f] = |R/\mathcal{O}_f|$.*

Definition 3 (Embeddings). The embeddings $\sigma_j : K \rightarrow \mathbb{C}$, $1 \leq j \leq n$, are the field homomorphisms that fix $\mathbb{Q}$. The values $\sigma_1(\theta), \dots, \sigma_n(\theta)$ are the complex roots of f . There are r_1 real embeddings and r_2 complex pairs, so $n = r_1 + 2r_2$. Totally imaginary means $r_1 = 0$.

The embeddings allow us to define the trace, norm, and discriminants.

Definition 4 (Trace and field norm). For $x \in K$, trace and field norm are $\text{Tr}(x) = \sum_j \sigma_j(x)$ and $N_{K/\mathbb{Q}}(x) = \prod_j \sigma_j(x)$.

Definition 5 (Discriminants). For a $\mathbb{Q}$ -basis (u_i) of K , the discriminant is

$$\det(\text{Tr}(u_i u_j)).$$

The power basis has absolute discriminant $\Delta_f = |\text{Disc}(f)|$. A $\mathbb{Z}$ -basis of R has absolute discriminant $\Delta_K = |\text{Disc } K|$. The two discriminants satisfy $\Delta_f = \Delta_K j_f^2$.

Different irreducible polynomials can define the same number field. For example, $\mathbb{Q}[x]/(x^2 - 2)$ and $\mathbb{Q}[x]/(x^2 - 8)$ are both isomorphic to $\mathbb{Q}(\sqrt{2})$. Their integer quotients define the different orders $\mathbb{Z}[\sqrt{2}]$ and $\mathbb{Z}[2\sqrt{2}]$ in that field. Thus the order $\mathcal{O}_f$ and index j_f may change with the defining polynomial even when the field K is unchanged. We therefore distinguish polynomial counts from field counts.

Coordinate notation. Coordinate vectors are columns. For a basis $\mathcal{B} = (u_i)_{i=1}^n$ of K , define $[x]_{\mathcal{B}}$ by $x = \sum_i ([x]_{\mathcal{B}})_i u_i$. Changing the basis leaves x fixed but may change its coordinate column $[x]_{\mathcal{B}}$.

Allowing real basis coordinates defines the real vector space $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$. We use $[x]_{\mathcal{B}}$ for these coordinates as well. We extend each embedding and the trace by real linearity.

Basis matrices place basis vectors in columns in specified coordinates. T_{ij} indexes row i , column j of a matrix T . Coordinate and Gaussian vectors $\mathbf{y}, \mathbf{G}, \mathbf{Z}$ have scalar entries y_i, G_i, Z_i .

For integer $m \geq 1$ and $p(x) = \sum_k p_k x^k$ of degree below m , coefficient coordinates are $(p_0, \dots, p_{m-1})^T$, including zeros. We write $(p)_k = p_k$ for the coefficient of x^k . In vector expressions, a polynomial denotes its coefficient column of the stated length.

The Ring-LWE secret uses a fixed ideal. The worst-case lattice problems allow arbitrary ideals in the same field. Trace duality makes the trace map compatible with integer coefficients.

Definition 6 (Fractional ideals). A fractional R -ideal is a nonzero additive subgroup $I \subseteq K$ closed under multiplication by R . We also require $dI \subseteq R$ for some integer $d \geq 1$. The ideal I is integral if $I \subseteq R$.

For fractional ideals I, J , the product IJ consists of finite sums of xy with $x \in I, y \in J$. The inverse ideal is $I^{-1} = \{z \in K : zI \subseteq R\}$. Ideal norm extends $N(J) = [R : J]$ multiplicatively from integral ideals.

Definition 7 (Trace dual and codifferent). A full $\mathbb{Z}$ -lattice $M \subseteq K$ is the integer span of a $\mathbb{Q}$ -basis. The trace dual of M is $M^\vee = \{z \in K : \text{Tr}(zM) \subseteq \mathbb{Z}\}$. The codifferent is the trace dual $R^\vee$, with $N(R^\vee) = \Delta_K^{-1}$.

Definition 8 (Conductor). The conductor ideal of $\mathcal{O}_f$ in R is $\mathfrak{c}_f = \{a \in R : aR \subseteq \mathcal{O}_f\}$.

Trace duality specifies which multipliers produce integer trace coefficients.

Definition 9 (Legal multipliers and trace maps). For $I = R$ (primal) or $I = R^\vee$ (dual), a legal multiplier is $0 \neq c \in R^\vee I^{-1}$. Thus $c \in R^\vee$ in the primal case and $c \in R$ in the dual case.

For $1 \leq h \leq n$, define the real-linear trace map on $K_\mathbb{R}$ by $L_{c,h}(e) = (\text{Tr}(ce\theta^i))_{0 \leq i < h}$. The condition $c \in R^\vee I^{-1}$ ensures $L_{c,h}(I) \subseteq \mathbb{Z}^h$.

For reductions at the same integer modulus $q \geq 2$, define admissible polynomials by $\mathcal{A}_{n,B}(q) = \{f \in \mathcal{I}_{n,B} : \gcd(q, j_f) = 1\}$. Polynomial coverage is $|\mathcal{A}_{n,B}(q)|/|\mathcal{I}_{n,B}|$.

Let $\mathcal{K}_{n,B}$ contain one representative of each $\mathbb{Q}$ -isomorphism class of fields defined by $\mathcal{I}_{n,B}$. For field coverage, we count a field once if at least one defining polynomial belongs to $\mathcal{A}_{n,B}(q)$. We divide this count by $|\mathcal{K}_{n,B}|$. A field with many defining polynomials contributes many times to a polynomial count but only once to a field count. Both fractions require nonempty denominators.

Our examples include cyclotomic fields, which are generated by roots of unity.

Definition 10 (Cyclotomic fields). For integer $m \geq 1$, let $\omega_m = e^{2\pi i/m}$ be a primitive m th root of unity. Write Φ_m for the minimal polynomial of ω_m over $\mathbb{Q}$. The field $K_m = \mathbb{Q}(\omega_m)$ has embeddings $\sigma_u(\omega_m) = \omega_m^u$ for invertible residues $u \in (\mathbb{Z}/m\mathbb{Z})^\times$. The degree of K_m is $\varphi(m) = |(\mathbb{Z}/m\mathbb{Z})^\times|$. The integer m is the multiplicative order of ω_m . A prime-power cyclotomic polynomial is Φ_m with m a prime power.

2.2 Geometry and noise

We now define the geometry used to measure noise. Even when the input error is spherical, the trace map can produce correlated output coordinates. We use the operator norm to bound the increase in standard deviation. We use covariance to determine how much independent Gaussian noise to add to obtain a spherical output.

Definition 11 (Canonical geometry). The canonical embedding $\sigma(x) = (\sigma_j(x))_{j=1}^n$ identifies $K_\mathbb{R}$ with the subspace of $\mathbb{C}^n$ having real coordinates at real embeddings and conjugate coordinates at complex pairs. The inner product is $\langle x, y \rangle = \sum_j \sigma_j(x) \overline{\sigma_j(y)}$, so $\|x\|^2 = \sum_j |\sigma_j(x)|^2$ counts each complex pair twice. In orthonormal real coordinates, for invertible $E \in \mathbb{R}^{n \times n}$, the lattice $E\mathbb{Z}^n$ has covolume $|\det E|$.

We measure the trace map with matrix norms in these coordinates.

Definition 12 (Matrix norms and adjoints). For a matrix T , the operator norm is $\|T\|_{\text{op}} = \sup_{\|x\|=1} \|Tx\|$. We write T^* for the adjoint and T^*T for the Gram matrix. The notation $\text{diag}(a_j)$ denotes a diagonal matrix with entries a_j . The Frobenius norm satisfies $\|T\|_F^2 = \sum_{ij} |T_{ij}|^2$. We also use $\|T\|_\infty = \max_i \sum_j |T_{ij}|$ and $\|x\|_\infty = \max_i |x_i|$.

Definition 13 (Positive semidefinite order). A complex square matrix T is Hermitian if $T = T^*$. For Hermitian matrices T_1, T_2 of equal size, write $T_1 \preceq T_2$ if $\mathbf{z}^*(T_2 - T_1)\mathbf{z} \geq 0$ for every complex vector $\mathbf{z}$. A matrix T is positive semidefinite if $0 \preceq T$. For real symmetric matrices, it suffices to test real vectors.

Definition 14 (Covariance). Writing $\mathbb{E}$ for expectation, the covariance of a real random vector X is $\text{Cov}(X) = \mathbb{E}[(X - \mathbb{E}X)(X - \mathbb{E}X)^T]$. A linear map satisfies $\text{Cov}(TX) = T \text{Cov}(X) T^T$.

Definition 15 (Continuous Gaussians). For $0 \preceq W$, $N(0, W)$ is Gaussian with mean zero and covariance W . For the $m \times m$ identity $\mathbf{I}_m$, $N(0, s^2 \mathbf{I}_m)$ is spherical, with standard deviation $s > 0$ in every orthonormal coordinate. For $r > 0$, a density proportional to $e^{-\pi \|x\|^2 / r^2}$ has standard deviation $r / \sqrt{2\pi}$.

The trace map depends on the roots of the defining polynomial. We first record the matrices that convert coordinates. We then define the root weights and sums used to bound their norms.

Definition 16 (Coordinate matrices). For the roots of f , write $V_f = (\sigma_j(\theta)^i)_{1 \leq j \leq n, 0 \leq i < n}$ for the Vandermonde matrix, with embeddings as rows and powers as columns. For the route through PLWE, define $M_f \in \mathbb{Z}^{n \times n}$ by

$$(M_f)_{i+1, k+1} = (x^{i+k} \bmod f)_0, \quad 0 \leq i, k < n.$$

The remainder $x^{i+k} \bmod f$ has degree below n . Before output reversal, M_f maps PLWE error coefficients to MP-LWE error coefficients [15, Definition 2.3]. We use NP26's notation $D(f) = \max_{i,j} |(M_f V_f^{-1})_{ij}|$.

Definition 17 (Polynomial norms and root weights). The coefficient norms are $\|f\|_p = (1 + \sum_{i < n} |a_i|^p)^{1/p}$ for $p = 1, 2$. The Mahler measure and weights determined by root magnitudes are

$$\begin{aligned} \mathcal{M}(f) &= \prod_j \max(1, |\sigma_j(\theta)|), & w_j &= \max(1, |\sigma_j(\theta)|^{n-1}), \\ \prod_j w_j &= \mathcal{M}(f)^{n-1}. \end{aligned} \tag{1}$$

The weights w_j account for the magnitudes of powers of the roots. To bound the trace coordinates together, we also need the relative positions of the roots. Reflecting roots outside the unit circle into the disk allows us to use one convergent series for all roots.

Definition 18 (Reflected roots and Poisson kernel). For $1 \leq j \leq n$, define

$$\zeta_j = \frac{\sigma_j(\theta)}{\max(1, |\sigma_j(\theta)|)^2}.$$

Thus $\zeta_j = \sigma_j(\theta)$ when $|\sigma_j(\theta)| \leq 1$, and $\zeta_j = 1/\overline{\sigma_j(\theta)}$ when $|\sigma_j(\theta)| > 1$. Reflection preserves the angle of each root and respects complex conjugate pairs. Every ζ_j lies in the closed unit disk. For $w \in \mathbb{C}$ with $|w| < 1$, the Poisson kernel is

$$P(w) = \frac{1 - |w|^2}{|1 - w|^2} = 1 + 2 \operatorname{Re} \sum_{k \geq 1} w^k.$$

In particular, $P(w) > 0$.

Definition 19 (Sums of reflected root powers). For integers $k \geq 1$ and $0 < \rho < 1$, define

$$S_k = \sum_{j=1}^n \zeta_j^k, \quad F_f(\rho) = \sum_{k \geq 1} \rho^k |S_k|^2, \quad E_f(\rho) = \sum_{k \geq 1} \frac{\rho^k |S_k|^2}{k}.$$

Both series converge because $|S_k| \leq n$.

The parameter ρ keeps the series convergent even for roots on the unit circle. The sum $F_f(\rho)$ will bound the combined contribution of the roots to the trace norm. The factor $1/k$ in $E_f(\rho)$ allows us to relate the sum $E_f(\rho)$ to the discriminant.

2.3 Decision problems and reduction guarantees

We compare two decision problems. Ring-LWE measures error in canonical coordinates, whereas MP-LWE measures error in coefficient coordinates. In both definitions, $\mathbf{U}$ denotes uniform sampling.

Definition 20 (Decision Ring-LWE). Fix $f \in \mathcal{I}_{n,B}$, integers $q \geq 2, t \geq 1$, and noise rate $\alpha > 0$. Set $I = R$ (primal) or $I = R^\vee$ (dual). One secret $s \leftarrow \mathbf{U}(I/qI)$ is shared by at most t samples. For each sample, draw fresh independent a, e and compute b as follows:

$$a \leftarrow \mathbf{U}(R/qR), \quad b = as + e \pmod{qI}, \quad e \leftarrow N(0, \alpha^2 q^2 \mathbf{I}_n) \quad (2)$$

The Gaussian error uses canonical coordinates and has standard deviation αq . The problem $\text{RLWE}_{K,I,q,\alpha}^{[t]}$ distinguishes the samples from independent uniform pairs on $(R/qR) \times (K_{\mathbb{R}}/qI)$. Haar measure on $K_{\mathbb{R}}/qI$ uses uniform coordinates in $[0, q)^n$ in any $\mathbb{Z}$ -basis of I .

Definition 21 (Middle product). For integers $q \geq 2$ and $m \geq 1$, write $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ and let $\mathbb{Z}_q^{<m}[x]$ contain polynomials of degree below m , stored as m coefficients. For integers $1 \leq h \leq n$, $a \in \mathbb{Z}_q^{<n}[x]$, and $s \in \mathbb{Z}_q^{<n+h-1}[x]$, the middle product is $a \odot_h s = \sum_{i=0}^{h-1} (as)_{n-1+i} x^i$. The output has h coefficients. For $h = n$, the retained indices of as are $n-1, \dots, 2n-2$.

Definition 22 (Decision MP-LWE). Fix integers $q \geq 2$, $1 \leq h \leq n$, noise rate $\beta > 0$, and integer $t \geq 1$. One secret $s \leftarrow \mathcal{U}(\mathbb{Z}_q^{<n+h-1}[x])$ is shared by at most t samples. For each sample, draw fresh independent a, e and compute b as follows:

$$a \leftarrow \mathcal{U}(\mathbb{Z}_q^{<n}[x]), \quad b = a \odot_h s + e \pmod{q}, \quad e \leftarrow N(0, \beta^2 q^2 \mathbf{I}_h), \quad (3)$$

Here $b \in (\mathbb{R}/q\mathbb{Z})^h$ and the error has standard deviation βq . Distinguishing the samples from independent uniform pairs on $\mathbb{Z}_q^n \times (\mathbb{R}/q\mathbb{Z})^h$ defines $\text{MP-LWE}_{q,n,h,\beta}^{[t]}$.

In MP-LWE, n counts coefficients of the public polynomial a , h counts noisy output coefficients, and the secret has $n + h - 1$ coefficients. The integer t limits sample pairs sharing one secret, so at most th noisy scalar outputs are available. In our degree-preserving reductions, the Ring-LWE field degree is also n . The integer modulus q determines coefficient residues and the quotients R/qR and I/qI . The modulus q is distinct from the defining polynomial f .

The rates α, β are dimensionless. Before modular reduction, the standard deviations are αq in Ring-LWE's orthonormal canonical coordinates and βq in MP-LWE's coefficient coordinates. Consequently, changing q while fixing a noise rate changes the absolute error magnitude.

Our decision reductions use an MP-LWE distinguisher to distinguish Ring-LWE samples from uniform samples. The distinguishing advantage is the absolute difference between the acceptance probabilities in the two experiments.

Suppose we have a bound $C \geq \|L_{c,n}\|_{\text{op}}$, and put $\Gamma = 2C$. Under the conditions of Lemma 1, Ring-LWE with rate β/Γ then reduces to MP-LWE with rate β . We call Γ the noise-loss factor. The factor two allows us to add independent Gaussian noise and obtain the required spherical covariance. It also leaves room for numerical approximation. At fixed MP-LWE parameters, a smaller Γ therefore permits a larger Ring-LWE rate $\alpha = \beta/\Gamma$.

The sample bound $[t]$ applies to each distinguisher call. We omit $[t]$ when the sample bound is unspecified.

Continuous Gaussian samples require finite encodings for computation. We define the encoding and then bound simulation error on encoded samples.

We represent R and the Ring-LWE ideal I by $\mathbb{Z}$ -bases of polynomial bit length in rational power-basis coordinates.

Definition 23 (Finite representation of samples). Encode $a \in R/qR$ exactly. In the basis $\mathcal{B}_I = (u_i)$ of I , write $b = \sum_i y_i u_i \pmod{qI}$ with $\mathbf{y} \in [0, q]^n$. Thus $\mathbf{y} = \tilde{b}|_{\mathcal{B}_I} \bmod q\mathbb{Z}^n$ for any representative $\tilde{b}$ of b . For integer precision $r_{\text{in}} \geq 0$, store $\lfloor 2^{r_{\text{in}}} \mathbf{y} \rfloor$. Decoding returns $2^{-r_{\text{in}}} \lfloor 2^{r_{\text{in}}} \mathbf{y} \rfloor$. Each coordinate error is below $2^{-r_{\text{in}}}$.

For an MP-LWE coordinate $z \in [0, q]$, use $\text{Enc}_{r_{\text{out}}}(z) = \lfloor 2^{r_{\text{out}}} z \rfloor$, $r_{\text{out}} \in \mathbb{Z}_{\geq 0}$, or nearest-integer rounding modulo q . For each problem, encode both the LWE and uniform experiments identically.

Parameters and bases. Reductions to one fixed MP-LWE problem use identical (q, n, h, β, t) and sample encoding, even when the Ring-LWE fields differ.

To bound basis descriptions across $\mathcal{L}_{n,B}$, we use Hermite normal form (HNF): an upper-triangular integer column basis H with $H_{ii} > 0$ and $0 \leq H_{ij} < H_{ii}$ for $i < j$. Lemma 5 bounds the descriptions of HNF-based bases. Proposition 3 also allows other bases of polynomial bit length.

Simulation error. In each noisy or uniform experiment, we compare the joint simulation with exact sampling followed by encoding.

For distributions P, Q , total variation distance is $\text{TV}(P, Q) = \sup_E |P(E) - Q(E)|$, over measurable events E . A coupling chooses $X \sim P$ and $Y \sim Q$ jointly. For any such coupling, $\text{TV}(P, Q) \leq \Pr[X \neq Y]$.

An error $\epsilon(n)$ is negligible if $n^c \epsilon(n) \rightarrow 0$ for every fixed $c > 0$. Summing negligible errors over polynomially many reduction calls leaves a negligible total. Negligible total error preserves inverse-polynomial distinguishing advantage.

Computation and advice. We measure computation in bit operations. The notation poly denotes an unspecified polynomial bound. Public rational Ring-LWE and MP-LWE parameters have $\text{poly}(n)$ bit length. A positive real parameter u of either problem has $|\log u| \leq \text{poly}(n)$ and can be approximated to k bits in time $\text{poly}(n, k)$.

Uniform reductions use one algorithm for all inputs, without advice. Our general reductions instead use nonuniform advice of polynomial size. We fix the advice separately for each f, I . The advice is independent of q , secrets, samples, and every worst-case fractional input ideal $\mathfrak{a}$.

We prove that suitable advice exists and has polynomial bit length. We do not provide an efficient algorithm to compute the general multiplier. Lemma 5 bounds the advice size. The input ideal $\mathfrak{a}$ can vary while the field and the Ring-LWE secret ideal I remain fixed.

2.4 Ideal-lattice problems

We finally define the lattice problems used to state the cryptographic consequences. The field stays fixed, but the input ideal may vary.

Definition 24 (Ideal lattices and Euclidean duals). *For a nonzero fractional R -ideal $\mathfrak{a}$, $\Lambda_{\mathfrak{a}} = \sigma(\mathfrak{a}) \subset K_{\mathbb{R}}$ is the rank- n ideal lattice spanned over $\mathbb{Z}$ by any embedded $\mathbb{Z}$ -basis of $\mathfrak{a}$. For a full lattice Λ in canonical space, the Euclidean dual is $\Lambda^* = \{y \in K_{\mathbb{R}} : \langle y, \Lambda \rangle \subseteq \mathbb{Z}\}$. The two duals satisfy $\Lambda_{\mathfrak{a}}^* = \overline{\sigma(\mathfrak{a}^{\vee})}$. Complex conjugation preserves lengths.*

Definition 25 (Successive minima). *For a full lattice Λ in canonical space, the i th successive minimum $\lambda_i(\Lambda)$, $1 \leq i \leq n$, is the least radius containing i linearly independent lattice vectors. Write $\lambda_i(\mathfrak{a}) = \lambda_i(\Lambda_{\mathfrak{a}})$.*

Definition 26 (Smoothing parameter). For a full lattice Λ in canonical space and $0 < \epsilon < 1$, the smoothing parameter of Λ is

$$\eta_\epsilon(\Lambda) = \inf \left\{ r > 0 : \sum_{0 \neq y \in \Lambda^*} e^{-\pi r^2 \|y\|^2} \leq \epsilon \right\}.$$

Write $\eta_\epsilon(\mathfrak{a}) = \eta_\epsilon(\Lambda_{\mathfrak{a}})$.

The tolerance ϵ bounds the summed contribution of nonzero dual vectors. The width $\eta_\epsilon(\Lambda)$ uses the same length units as lattice vectors. Increasing r decreases the sum, so a smaller tolerance can only increase the smoothing threshold.

At widths at least the threshold, the tolerance controls the closeness of a continuous Gaussian modulo Λ to uniform. The smoothing tolerance and the error allowed in finite sample simulation are chosen separately.

Definition 27 (Discrete Gaussian sampling). For a nonzero fractional R -ideal $\mathfrak{a}$ and $r > 0$, the discrete Gaussian $D_{\mathfrak{a},r}$ assigns $x \in \mathfrak{a}$ a probability proportional to $e^{-\pi \|\sigma(x)\|^2 / r^2}$.

Ideal-DGS (discrete Gaussian sampling) uses a positive real-valued threshold function g_{DGS} on ideals. Given a basis of $\mathfrak{a}$ and $r \geq g_{\text{DGS}}(\mathfrak{a})$, the task is to sample within negligible total variation distance of $D_{\mathfrak{a},r}$.

The function g_{DGS} specifies the required sampling widths. A lower g_{DGS} requires sampling narrower Gaussians. The Gaussian width r is a real number. In contrast, the encoding precisions $r_{\text{in}}, r_{\text{out}}$ count fractional bits.

Definition 28 (Shortest independent vectors). For real $\gamma \geq 1$, *Ideal-SIVP $_{K,\gamma}$ (shortest independent vectors)* takes a basis of a nonzero fractional R -ideal $\mathfrak{a}$ and asks for n linearly independent elements of $\mathfrak{a}$, each with embedded norm at most $\gamma \lambda_n(\mathfrak{a})$. The vectors must be independent, but need not form a $\mathbb{Z}$ -basis of the ideal.

The radius $\lambda_n(\mathfrak{a})$ is the best possible maximum length of n independent vectors, without a computational restriction. The approximation factor γ permits output vectors up to γ times longer. A smaller γ requires shorter output vectors. A reduction with smaller γ therefore turns a distinguisher into a stronger ideal-lattice algorithm.

2.5 Road Map

In Section 3, we first express the noise loss as the norm of a legal trace map. We bound its Gram matrix using the Poisson kernel of the reflected roots. This comparison determines weights for the multiplier's embeddings, and a weighted lattice argument proves that a legal multiplier whose trace map has small norm exists. The discriminant bounds $E_f(\rho)$. Comparing the sums at different values of ρ sharpens the primal estimate, while retaining the discriminant in the calculation sharpens the dual estimate. Integer trace coordinates then prove that the

multipliers have polynomial-size descriptions. None of these arguments requires NP26's condition $D(f) \leq n$.

In Section 4, we choose one prime for almost all polynomials in $\mathcal{I}_{n,B}$ and distinct represented fields. We then choose sufficient parameters for the lattice-hardness reduction and show that the corresponding modulus size makes both coverage proportions tend to one.

In Section 5, we express NP26's complete linear error map as a legal trace map. For n outputs, we prove that one short multiplier independent of the modulus meets both our general norm bound and the norm bounds achieved by NP26's multipliers. Section 6 discusses the remaining limits.

3 Low-Loss Trace Maps

We prove lower-loss multipliers with polynomial-size descriptions for every $f \in \mathcal{I}_{n,B}$, without NP26's condition $D(f) \leq n$. PP24's framework then implements the reduction given the multiplier, under $q \nmid j_f$ and the stated noise conditions.

3.1 From a Norm Bound to a Reduction

We state the specialization of PP24's reduction used in our norm analysis.

Lemma 1 (Trace reduction, following PP24). *Let $f \in \mathcal{I}_{n,B}$, and let q be prime with $q \nmid j_f$. Fix $I = R$ or $R^\vee$, $1 \leq h \leq n$, and $\alpha, \beta > 0$. Fix $0 \neq c \in R^\vee I^{-1}$ of polynomial bit length and a proved bound $C \geq \|L_{c,n}\|_{\text{op}}$, fixed by f, I . If $\beta \geq \alpha C$, there is an exact continuous decision reduction from $\text{RLWE}_{K,I,q,\alpha}$ to $\text{MP-LWE}_{q,n,h,\beta}$, preserving the number of samples and the decision advantage before numerical approximation. For the finite implementation, assume polynomial-size descriptions of the public parameters, C , and the ring and ideal bases. Each positive real parameter $u \in \{\alpha, \beta, C\}$ must satisfy $|\log u| \leq \text{poly}(n)$ and be approximable to k bits in time $\text{poly}(n, k)$. If $\beta \geq 2\alpha C$, the reduction has a polynomial-time finite-string implementation using polynomially many precision bits, uniform given c and C (Proposition 3). If c is supplied as advice, the resulting reduction is nonuniform.*

Proof. Multiplication by c sends (s, b, I) to (cs, cb, cI) bijectively between the corresponding quotients, with error ce . Since $j_f R \subseteq \mathcal{O}_f$ and $q \nmid j_f$, the inverse transform in PP24's Theorem 4.5 replaces the public ring R by $\mathcal{O}_f$ without changing the error [16, Theorem 4.5]. The inclusion $cI \subseteq R^\vee \subseteq \mathcal{O}_f^\vee$ then permits expanding the secret ideal to $\mathcal{O}_f^\vee$ [16, Theorem 4.4]. PP24's Theorem 5.3 applies the power-basis traces and produces MP-LWE samples. Retaining h coefficients makes the error map $L_{c,h}$ [16, Section 5.1 and Theorem 5.3]. As in PP24 [16, proof of Theorem 4.8], one shared uniform mask $v \in \mathbb{Z}_q^{\leq n+h-1}[x]$ makes the MP-LWE secret uniform. These operations preserve uniform inputs and the sample count. No inverse of c modulo q is needed.

By PP24 [16, Corollary 5.5], we add an independent Gaussian $\mathbf{G} \sim N(0, W)$ with

$$W = (\beta q)^2 \mathbf{I}_h - (\alpha q)^2 L_{c,h} L_{c,h}^* \succeq 0.$$

The inequality follows from $\beta \geq \alpha C$. The resulting error is spherical with standard deviation βq , which proves the exact decision reduction. The finite implementation under $\beta \geq 2\alpha C$ follows from Proposition 3. $\square$

3.2 General Noise Bounds

We prove that a legal multiplier with a small complete trace norm exists. The proof adapts the multiplier weights to a positive Gram matrix. A discriminant estimate controls the energy of the reflected roots. Comparing that energy at different radii yields the primal bound. Retaining the field discriminant in the optimization yields the dual bound.

Throughout this subsection, let $f \in \mathcal{I}_{n,B}$ and abbreviate $M = \mathcal{M}(f)$, $\eta = M^{1-1/n}$, and $y = \Delta_K^{1/n}$. The reflected roots ζ_j , their power sums S_k , the functions F_f, E_f , and the Poisson kernel P are as defined in the preliminaries. Jensen's formula, concavity of the logarithm, and Parseval's identity imply

$$\begin{aligned} M &\leq \|f\|_2 \leq B\sqrt{n+1}, & \Delta_f &\leq n^n M^{2n-2}, \\ \Delta_K^{1/n} &\leq \Delta_f^{1/n} \leq H_n := n(n+1)B^2. \end{aligned} \quad (4)$$

Indeed, $\log M = (2\pi)^{-1} \int_0^{2\pi} \log |f(e^{iu})| du$. Each row of V_f has norm at most $\sqrt{n} w_j$. Hadamard's inequality, $|\det V_f|^2 = \Delta_f$, and $\prod_j w_j = M^{n-1}$ prove the discriminant estimate.

A weighted lattice bound.

Lemma 2. *Let J be a nonzero fractional ideal of R , and let positive weights v_j agree on conjugate embeddings. There is $0 \neq c \in J$ such that*

$$\max_j v_j |\sigma_j(c)| \leq \kappa_K \left(\sqrt{\Delta_K} N(J) \prod_j v_j \right)^{1/n}, \quad \kappa_K = (2/\pi)^{r_2/n} \leq 1. \quad (5)$$

The proof applies Minkowski's convex body theorem to the body $v_j |\sigma_j(x)| \leq r$. Appendix E records the volume computation, including the factors from complex embeddings.

The legal multiplier ideal is

$$J = R^\vee I^{-1}; \quad J = R^\vee \quad (I = R), \quad J = R \quad (I = R^\vee). \quad (6)$$

Its norm is respectively Δ_K^{-1} or 1.

Discriminant and reflected-root energy.

Lemma 3. *For every $0 < \rho < 1$,*

$$E_f(\rho) \leq 2(n-1) \log M - \log \Delta_f + n \log \frac{1}{1-\rho} - \frac{n(n-1)}{2} \log \rho. \quad (7)$$

Proof. Put $m_j = \max(1, |\sigma_j(\theta)|)$ and $d_{jk} = |\sigma_j(\theta) - \sigma_k(\theta)| / (m_j m_k)$ for $j \neq k$. If both roots lie inside the closed unit disk, or both lie outside, then $d_{jk} = |\zeta_j - \zeta_k|$. If one root lies inside and the other outside, then $d_{jk} = |1 - \zeta_j \bar{\zeta}_k|$. The identity

$$|1 - u\bar{v}|^2 - |u - v|^2 = (1 - |u|^2)(1 - |v|^2) \geq 0 \quad (|u|, |v| \leq 1)$$

therefore proves $d_{jk} \leq |1 - \zeta_j \bar{\zeta}_k|$ in every case. Also, for $|w| \leq 1$,

$$|1 - \rho w|^2 - \rho |1 - w|^2 = (1 - \rho)(1 - \rho |w|^2) \geq 0.$$

Consequently $|1 - \rho \zeta_j \bar{\zeta}_k| \geq \sqrt{\rho} d_{jk}$ for $j \neq k$. The discriminant product satisfies

$$\prod_{j < k} d_{jk}^2 = \Delta_f M^{-2(n-1)}.$$

Expanding $-\log(1 - w)$ in its absolutely convergent power series yields

$$E_f(\rho) = - \sum_{j,k} \log |1 - \rho \zeta_j \bar{\zeta}_k|.$$

The off-diagonal terms are bounded using the preceding product. For the diagonal terms, $1 - \rho |\zeta_j|^2 \geq 1 - \rho$. These estimates prove (7). $\square$

A positive Gram matrix and adaptive weights.

Proposition 1 (General operator bound). *For every $0 < \rho < 1$ and each $I = R, R^\vee$, there is a legal multiplier c , depending only on f, I, ρ , such that*

$$\|L_{c,n}\|_{\text{op}}^2 \leq 4\rho^{-(n-1)} \kappa_K^2 n \eta^2 \left(1 + \frac{F_f(\rho)}{n^2}\right) \begin{cases} y^{-1}, & I = R, \\ y, & I = R^\vee. \end{cases} \quad (8)$$

Proof. Let $\mathbf{U} = V_f^\top \text{diag}(w_j^{-1})$, so $\mathbf{U}_{i+1,j} = \sigma_j(\theta)^i / w_j$. Partition the columns into $\mathcal{J}_{\text{in}} = \{j : |\sigma_j(\theta)| \leq 1\}$ and $\mathcal{J}_{\text{out}} = \{j : |\sigma_j(\theta)| > 1\}$. Define three matrices on the full index set by

$$\mathbf{Q}_{i+1,j} = \zeta_j^i \quad (0 \leq i < n), \quad (\mathbf{H}_\rho)_{jk} = \frac{1}{1 - \rho \zeta_j \bar{\zeta}_k}, \quad \mathbf{A}_\rho = \mathbf{H}_\rho + \overline{\mathbf{H}_\rho}.$$

The geometric-series expansion shows that $\mathbf{H}_\rho$ is positive semidefinite. Comparing its first n summands with $\mathbf{Q}^* \mathbf{Q}$ proves

$$\mathbf{Q}^* \mathbf{Q} \preceq \rho^{-(n-1)} \mathbf{H}_\rho \preceq \rho^{-(n-1)} \mathbf{A}_\rho.$$

The second inequality holds because $\overline{H_\rho}$ is also positive semidefinite. Moreover, $(A_\rho)_{jk} = 1 + P(\rho \overline{\zeta_j} \zeta_k) > 0$ is real and symmetric. Let $b_j(\rho)$ be the row sum of A_ρ within the group containing j , and put $D = \text{diag}(\sqrt{b_j(\rho)})$ on all columns. For a group $\mathcal{J}$, let $A_{\rho, \mathcal{J}}$ denote the principal submatrix indexed by $\mathcal{J}$. For any complex column $\mathbf{v} = (v_j)_{j \in \mathcal{J}}$, the identity

$$\mathbf{v}^* (\text{diag}_{j \in \mathcal{J}}(b_j(\rho)) - A_{\rho, \mathcal{J}}) \mathbf{v} = \frac{1}{2} \sum_{j, k \in \mathcal{J}} (A_\rho)_{jk} |v_j - v_k|^2 \geq 0$$

proves that each normalized group of $\mathbf{Q}$ has squared operator norm at most $\rho^{-(n-1)}$.

The inside group of $\mathbf{U}$ equals the corresponding group of $\mathbf{Q}$. Writing $\sigma_j(\theta) = |\sigma_j(\theta)| e^{i\phi_j}$ for an outside root yields

$$U_{i+1, j} = e^{i(n-1)\phi_j} \overline{\zeta_j}^{n-1-i}.$$

Thus the outside group is obtained from $\mathbf{Q}$ by conjugation, row reversal, and multiplication of columns by numbers of modulus one. The real diagonal matrix D commutes with the column factors, and conjugation preserves the norm after normalization. The squared norm of a horizontal concatenation is at most the sum of the squared norms of its blocks. Applying this fact to the two normalized groups, with any empty group omitted, proves

$$\|\mathbf{U}D^{-1}\|_{\text{op}}^2 \leq 2\rho^{-(n-1)}, \quad \sum_j b_j(\rho) \leq 2n^2 + 2F_f(\rho). \quad (9)$$

For the second inequality, the within-group sums are at most the sum over all pairs because $(A_\rho)_{jk} > 0$. The identity $\sum_{j, k} P(\rho \overline{\zeta_j} \zeta_k) = n^2 + 2F_f(\rho)$ then applies.

Complex conjugation permutes each group and preserves $b_j(\rho)$. We may therefore apply Lemma 2 to J in (6) with weights $w_j \sqrt{b_j(\rho)}$. It supplies a nonzero legal c with

$$\max_j w_j \sqrt{b_j(\rho)} |\sigma_j(c)| \leq \kappa_K \eta \left(\prod_j b_j(\rho) \right)^{1/(2n)} \begin{cases} y^{-1/2}, & I = R, \\ y^{1/2}, & I = R^\vee. \end{cases}$$

In canonical coordinates the trace map has matrix

$$\mathbf{U}D^{-1} \text{diag}(w_j \sqrt{b_j(\rho)} \sigma_j(c))$$

restricted to the canonical real subspace of $\mathbb{C}^n$. The trace coordinates are real on this subspace. Its inclusion is isometric, so the trace operator norm is at most the complex matrix norm. Finally, $(\prod_j b_j(\rho))^{1/n} \leq n^{-1} \sum_j b_j(\rho)$ and (9) prove (8). $\square$

Comparing Poisson scales.

Lemma 4. *Let $\zeta_1, \dots, \zeta_n$ be any points in the closed unit disk. Put $S_k = \sum_j \zeta_j^k$, $F(\rho) = \sum_{k \geq 1} \rho^k |S_k|^2$, and $E(\rho) = \sum_{k \geq 1} \rho^k |S_k|^2 / k$. Then every $0 < \rho < 1$ satisfies*

$$E(\rho) \geq \frac{1-\rho}{4} F(\rho) \log \left(1 + \frac{F(\rho)}{n^2} \right). \quad (10)$$

Proof. For $0 \leq r \leq \rho < 1$ and $|w| \leq 1$, the Poisson kernel satisfies

$$P(rw) \geq \frac{(1-\rho)(1+r)}{(1+\rho)(1-r)} P(\rho w). \quad (11)$$

The ratio $P(rw)/P(\rho w)$ is minimized at $w = 1$. Appendix E verifies this claim by differentiating first in the angle and then in the magnitude of w .

Put $\delta = 1 - \rho$, $\varepsilon = \rho - r$, and $x = F(\rho)/n^2$. Since $|S_k| \leq n$, we have $x \leq \rho/\delta$. The factor in (11) is at least $\delta/(\delta + 2\varepsilon)$, because

$$(1+r)(\delta + 2\varepsilon) - (1+\rho)(\delta + \varepsilon) = 2r\varepsilon \geq 0.$$

Sum (11) over all pairs of points and use $\sum_{j,k} P(r\zeta_j \overline{\zeta_k}) = n^2 + 2F(r)$. The result is

$$F(r) \geq n^2 \frac{\delta x - \varepsilon}{\delta + 2\varepsilon}.$$

The case $x = 0$ is immediate. Otherwise, on $0 \leq \varepsilon \leq \delta x/2$ the last expression is at least $\delta F(\rho)/(2(\delta + 2\varepsilon))$. The bound $\delta x \leq \rho$ ensures $r \geq \rho/2 > 0$ on this interval. Termwise integration of the nonnegative series yields $E(\rho) = \int_0^\rho F(r) dr/r$. Since $1/r \geq 1$, integration over the indicated interval yields

$$E(\rho) \geq \int_0^{\delta x/2} \frac{\delta F(\rho)}{2(\delta + 2\varepsilon)} d\varepsilon = \frac{\delta F(\rho)}{4} \log(1 + x).$$

□

Explicit factors for the reductions. The following integer factors include the factor two for finite implementation. All square roots are rounded upward after the indicated rational expression is formed:

$$\ell_{n,B} = \lceil \log_2(B^2 n(n+1)) \rceil + 1, \quad v_n = \lfloor \log_2(1 + 4 \lfloor \log_2 n \rfloor) \rfloor, \quad (12)$$

$$\Gamma_p(n, B) = 4B \left\lceil \sqrt{3n(n+1) \left(1 + \frac{16\ell_{n,B}}{v_n} \right)} \right\rceil, \quad \Gamma_d(n, B) = 10B^2 n(n+1). \quad (13)$$

Here $v_n \geq 1$. Both factors are nondecreasing in B and can be computed by integer comparisons and integer square roots.

Theorem 1 (General noise-loss bounds). *Let $f \in \mathcal{I}_{n,B}$, and let q be prime with $q \nmid j_f$. Fix $1 \leq h \leq n$, $\beta > 0$, and an integer sample bound $1 \leq t \leq \text{poly}(n)$. Put $\Gamma_R = \Gamma_p(n, B)$ and $\Gamma_{R^\vee} = \Gamma_d(n, B)$ as in (13). For each $I = R, R^\vee$, use public parameters and bases of R and I with polynomial-size descriptions. Each positive real parameter u must satisfy $|\log u| \leq \text{poly}(n)$ and be approximable to k bits in time $\text{poly}(n, k)$. For each (f, I) , there exists a legal multiplier c of polynomial bit length, fixed independently of q , the secret, and the samples, with $\|L_{c,n}\|_{\text{op}} \leq \Gamma_I/2$. It induces a nonuniform polynomial-time decision reduction $\text{RLWE}_{K,I,q,\beta/\Gamma_I}^{[t]} \rightarrow \text{MP-LWE}_{q,n,h,\beta}^{[t]}$. For fixed B , the factors satisfy, as $n \rightarrow \infty$,*

$$\Gamma_p = O_B\left(n\sqrt{\frac{\log n}{\log \log n}}\right), \quad \Gamma_d = O_B(n^2). \quad (14)$$

These are upper bounds. No optimality assertion is made. Given c , a uniform algorithm implements the reduction. For MP-LWE encodings with polynomially many fractional bits, Proposition 3 proves negligible total variation error in each encoded experiment. Finding a general multiplier in polynomial time remains open.

Proof. Primal bound. Take $\rho = 1 - 1/n$ and put $A_B = \log(B^2 n(n+1)) + 1/2$. Since $\Delta_f \geq 1$, $M \leq B\sqrt{n+1}$, and $-(n-1)\log(1-1/n) \leq 1$, Lemma 3 implies $E_f(\rho) \leq nA_B$. For $x = F_f(\rho)/n^2$, Lemma 4 therefore proves

$$x \log(1+x) \leq 4A_B, \quad x \leq \frac{8A_B}{\log(1+4A_B)}. \quad (15)$$

The second inequality follows by inverting the increasing function $x \log(1+x)$; Appendix E verifies the displayed upper bound and the rounding estimate $x \leq 16\ell_{n,B}/v_n$. Now $y \geq 1$, $\eta^2 \leq B^2(n+1)$, and $\rho^{-(n-1)} \leq e < 3$. Apply Proposition 1 with $I = R$:

$$\|L_{c,n}\|_{\text{op}}^2 \leq 4eB^2 n(n+1) \left(1 + \frac{16\ell_{n,B}}{v_n}\right) \leq (\Gamma_p(n, B)/2)^2.$$

Dual bound. Take $\rho = e^{-5/(6n)}$ and $u = e^{-1/(3n)}$. The inequality $ke^{-k/(2n)} \leq 2n/e$ for $k > 0$ proves $F_f(\rho) \leq (2n/e)E_f(u)$. By (4) and $\Delta_f = j_f^2 \Delta_K$,

$$0 < y \leq Y := n\eta^2 j_f^{-2/n}.$$

For $s > 0$, $1 - e^{-s} \geq se^{-s/2}$. Apply this inequality with $s = 1/(3n)$ to (7). The regularization terms have sum at most $\log(3n) + 1/6$, so

$$\frac{E_f(u)}{n} \leq 2\log \eta - \frac{2\log j_f}{n} - \log y + \log(3n) + \frac{1}{6} = \log(Y/y) + \lambda, \quad \lambda = \log 3 + \frac{1}{6}.$$

Write $x = y/Y \in (0, 1]$. Proposition 1, with $I = R^\vee$ and $\rho^{-(n-1)} \leq e^{5/6}$, supplies a legal multiplier satisfying

$$\|L_{c,n}\|_{\text{op}}^2 \leq 4e^{5/6} \kappa_K^2 n^2 \eta^4 j_f^{-2/n} x \left(1 + \frac{2}{e}(\lambda - \log x)\right).$$

The last factor increases on $(0, 1]$: its derivative is $1 + (2/e)(\lambda - 1 - \log x) > 0$. Consequently

$$\|L_{c,n}\|_{\text{op}} \leq C_* \kappa_K n \eta^2 j_f^{-1/n}, \quad C_* = 2e^{5/12} \sqrt{1 + \frac{2}{e} \left(\log 3 + \frac{1}{6} \right)} < 5. \quad (16)$$

For the last inequality, $5/2 < e < 3$ and $\log 3 < 7/6$ imply $C_*^2 < 124/5 < 25$. Since $\kappa_K \leq 1$, $j_f \geq 1$, and $\eta^2 \leq B^2(n+1)$, the resulting bound is at most $\Gamma_d(n, B)/2$.

Advice and reduction. In either case the chosen multiplier depends only on f, I . The bound $\|L_{c,n}\|_{\text{op}} \leq \Gamma_I/2 = \text{poly}(n, B)$ and Lemma 5 establish the polynomial bit length of the multiplier and the ideal bases. Deleting trace coordinates cannot increase the operator norm, so this multiplier works for every $h \leq n$. Apply Lemma 1 with $C = \Gamma_I/2$ and use Proposition 3 for the encoding error. Finally, for fixed B , $\ell_{n,B} = \Theta_B(\log n)$ and $v_n = \Theta(\log \log n)$. The integer formulas prove (14). $\square$

3.3 Finite Implementation

The norm bound also controls the multiplier's description length. For a legal c , the coordinates $\text{Tr}(c\theta^i)$ of $L_{c,n}(1)$ are integers of magnitude at most $\sqrt{n} \|L_{c,n}\|_{\text{op}}$. An invertible integer trace-pairing matrix recovers the power-basis coordinates of c . Lemma 5, proved in Appendix A, establishes polynomial bit length for the multiplier and the ring and ideal bases.

Given this advice, the condition $\beta \geq 2\alpha C$ leaves a positive margin in the added Gaussian covariance. Polynomially many precision bits then suffice for sampling and encoding. Proposition 3 states and proves the full guarantee in Appendix A: each encoded noisy or uniform experiment has negligible total variation error over polynomially many samples. Input precision may depend on the field and basis, while the MP-LWE output encoding remains shared.

4 Coverage at Shared MP-LWE Parameters

The general loss bound applies when $q \nmid j_f$. We first bound the proportions of polynomials and represented number fields satisfying this condition at one selected prime modulus. We then choose sufficient parameters for the reduction from ideal-SIVP to Ring-LWE. The resulting modulus grows fast enough for both coverage proportions to tend to one.

4.1 Coverage of Polynomials and Number Fields

Our noise-loss bound holds for every $f \in \mathcal{I}_{n,B}$. Each index has few large prime divisors, so averaging over primes will meet $q \nmid j_f$ for most polynomials.

Polynomial coverage can overcount fields with many defining polynomials. We count each field once by preselecting $f_K \in \mathcal{I}_{n,B}$ with $\mathbb{Q}[x]/(f_K) \simeq K$ for

each $K \in \mathcal{K}_{n,B}$. The prime-divisor argument needs no enumeration algorithm and can also retain a prescribed polynomial list.

Theorem 2. *Fix integers $n \geq 2$ and $B \geq 1$ with $\mathcal{I}_{n,B} \neq \emptyset$. For each $K \in \mathcal{K}_{n,B}$, choose $f_K \in \mathcal{I}_{n,B}$ with $\mathbb{Q}[x]/(f_K) \simeq K$. Let $Q \geq 41/2$ and let $\mathcal{S} \subseteq \mathcal{I}_{n,B}$ be a prescribed set of size k . Put $a_Q = 5n \log H_n/(6Q)$. If $ka_Q < 1$, there is a prime $q \in [Q, 2Q]$ with $q \nmid j_f$ for every $f \in \mathcal{S}$ and*

$$\max \left\{ \frac{|\{f \in \mathcal{I}_{n,B} : q \mid j_f\}|}{|\mathcal{I}_{n,B}|}, \frac{|\{K \in \mathcal{K}_{n,B} : q \mid j_{f_K}\}|}{|\mathcal{K}_{n,B}|} \right\} \leq \min \left\{ 1, \frac{2a_Q}{1 - ka_Q} \right\}. \quad (17)$$

Proof. Let m_Q count primes in $[Q, 2Q]$ and put $\nu_Q = \lfloor n \log H_n/(2 \log Q) \rfloor$. Since $j_f^2 \leq \Delta_f \leq H_n^n$, each index excludes at most ν_Q primes in the interval. Thus the set $\mathcal{G}_Q$ of primes avoiding all prescribed indices has size at least $m_Q - k\nu_Q$. Rosser and Schoenfeld [23, Corollary 3, (3.8)] prove $m_Q > 3Q/(5 \log Q)$, hence $\nu_Q/m_Q \leq a_Q$ and $m_Q - k\nu_Q \geq m_Q(1 - ka_Q) > 0$.

We now average over primes that retain the prescribed list. Let $b_{\text{pol}}(q), b_{\text{fld}}(q)$ be the excluded fractions in (17), counting polynomials and field representatives. Counting excluded primes for each polynomial or representative shows $\sum_{q \in \mathcal{G}_Q} b_{\text{pol}}(q) \leq \nu_Q$ and $\sum_{q \in \mathcal{G}_Q} b_{\text{fld}}(q) \leq \nu_Q$. Therefore

$$\frac{1}{|\mathcal{G}_Q|} \sum_{q \in \mathcal{G}_Q} (b_{\text{pol}}(q) + b_{\text{fld}}(q)) \leq \frac{2\nu_Q}{m_Q - k\nu_Q} \leq \frac{2a_Q}{1 - ka_Q}.$$

Some prime in $\mathcal{G}_Q$ meets this bound on the sum. Each excluded fraction is nonnegative, at most the sum, and at most one. $\square$

The loss factors depend only on n, B and the Ring-LWE variant. Consequently, all covered polynomials use the same input rate β/Γ_I for fixed MP-LWE rate β . Appendix F records the common finite target with the computational assumptions from Theorem 1.

Fix $B \geq 1$ and choose $Q_n \geq 21$ with $Q_n/(n \log n) \rightarrow \infty$. For every sufficiently large n with $\mathcal{I}_{n,B} \neq \emptyset$, Theorem 2, applied with $\mathcal{S} = \emptyset$, proves that a prime $q_n \in [Q_n, 2Q_n]$ satisfies

$$\frac{|\{f \in \mathcal{I}_{n,B} : q_n \nmid j_f\}|}{|\mathcal{I}_{n,B}|}, \quad \frac{|\{K \in \mathcal{K}_{n,B} : q_n \nmid j_{f_K}\}|}{|\mathcal{K}_{n,B}|} \geq 1 - O_B(n \log n/Q_n) \rightarrow 1.$$

Here f_K is the defining polynomial chosen for K before selecting q_n . Thus the fraction of polynomials satisfying the index condition and the fraction of fields whose chosen defining polynomials satisfy the index condition both tend to one. The limit is taken over degrees n for which $\mathcal{I}_{n,B}$ is nonempty.

Taking $\mathcal{S} = \emptyset$ imposes no requirement on any individual polynomial. Corollary 3 in Appendix F also allows us to choose a set $\mathcal{S}_n \subseteq \mathcal{I}_{n,B}$ before selecting q_n , with $|\mathcal{S}_n| \leq \lfloor 3Q_n/(5n \log H_n) \rfloor$. The selected prime then satisfies $q_n \nmid j_f$ for every $f \in \mathcal{S}_n$, while the two proportions above still tend to one.

The nonempty degrees are unbounded: $x^n + 1 \in \mathcal{I}_{n,B}$ when n is a power of two. Counting one chosen representative per field lower-bounds coverage of distinct fields. Polynomial coverage and field coverage use the same selected prime, but the two proportions need not be equal. Finite-degree exceptions may remain.

Section 4.2 now chooses polynomial-size Q_n that meet the hypotheses of the ideal-lattice reduction and satisfy $Q_n/(n \log n) \rightarrow \infty$.

4.2 Modulus Choices for Worst-Case Hardness Reductions

We now choose a modulus large enough for the ideal-lattice reduction and apply the coverage bound. At MP-LWE noise rate β , loss Γ allows dual Ring-LWE noise rate $\alpha_d = \beta/\Gamma$. The lattice-to-Ring-LWE step requires sufficiently large $\alpha_d q$. A smaller Γ therefore lowers the sufficient modulus threshold. The trace step itself only requires $q \nmid j_f$ and the noise-rate inequality.

Appendix B states the hypotheses of the lattice-to-Ring-LWE reduction [24, Corollary 7.3] and checks every parameter in the following corollary. The comparison in Appendix D also records the effect of changing the loss factor and an illustrative numerical example.

The next corollary uses all n middle coefficients and logarithmically many samples. Its noise rates and modulus interval are shared across $\mathcal{I}_{n,B}$.

Corollary 1. *Fix $B \geq 1$. For every $n \geq 2$, define*

$$\begin{aligned} t_n &= \lceil \log_2 n \rceil, & h_n &= n, & \beta_n &= \frac{1}{4nt_n(16t_n(n+1)+1)}, \\ s_n &= 2\lceil n^{1/4} \rceil, & u_n &= \lceil \sqrt{t_n} \rceil \lceil \log_2(t_n+1) \rceil, \\ \alpha_{d,n} &= \beta_n/\Gamma_d(n, B), & Q_n &= \Gamma_d(n, B)s_n u_n/\beta_n. \end{aligned} \quad (18)$$

Let q_n be any prime in $[Q_n, 2Q_n]$. Then

$$\begin{aligned} q_n &= \Theta_B(n^{17/4}(\log n)^{5/2} \log \log n), \\ \beta_n &= \Theta(n^{-2}(\log n)^{-2}), & \alpha_{d,n} &= \Theta_B(n^{-4}(\log n)^{-2}). \end{aligned} \quad (19)$$

Along any sequence $f_n \in \mathcal{I}_{n,B}$ satisfying $q_n \nmid j_{f_n}$, there is a nonuniform polynomial-time quantum reduction

$$\text{IdealSIVP}_{K_{f_n}, \gamma_n} \leq \text{MP-LWE}_{q_n, n, \beta_n}^{[t_n]}, \quad \gamma_n = c_B n^{19/4}(\log n)^{7/2} \log \log n \quad (20)$$

for sufficiently large n and a suitable constant c_B . The reduction supports any prescribed polynomial number of binary fractional bits or nearest-integer rounding, with the error bound in Proposition 3. All polynomials with $q_n \nmid j_f$ share MP-LWE parameters and encoding. Advice depends on f_n , independently of the input ideal $\mathfrak{a}$.

Substituting the modulus size into the coverage bound. The modulus interval in (18) satisfies

$$\frac{n \log n}{Q_n} = \Theta_B \left(\frac{1}{n^{13/4} (\log n)^{3/2} \log \log n} \right) \rightarrow 0.$$

Thus Corollary 3 selects a prime $q_n \in [Q_n, 2Q_n]$ for which both excluded proportions are $O_B(n \log n / Q_n)$ and tend to zero along degrees with $\mathcal{I}_{n,B} \neq \emptyset$. At this prime, Corollary 1 applies to every $f \in \mathcal{I}_{n,B}$ with $q_n \nmid j_f$ at shared MP-LWE parameters. The scale $q_n = n^{17/4+o(1)}$ is sufficient for our noise and sample choices, not required for all Ring-LWE parameters.

At the parameters in (18) and the prime selected by Corollary 3, MP-LWE hardness follows if ideal-SIVP at factor γ_n is hard for nonuniform quantum polynomial-time algorithms over one sequence of fields K_{f_n} with $f_n \in \mathcal{I}_{n,B}$ and $q_n \nmid j_{f_n}$. The reductions do not prove independence of these field-specific hardness assumptions. We next compare the NP26 and trace error maps directly.

5 Retaining NP26's Loss Guarantees

We express NP26's complete error map as a legal trace map with the same operator norm (Proposition 2). PX26 also identifies this map within its unified framework [17]. For n outputs, norm minimization then retains both our general bound and every complete NP26 bound (Corollary 2).

NP26 uses $0 \neq \mu_2 \in \mathfrak{c}_f$ to change samples from R to $\mathcal{O}_f$. For dual Ring-LWE, $0 \neq \mu_1 \in (R^\vee)^{-1}$ also converts dual samples to primal samples [15, Lemmas B.1–B.2]. The error becomes ue , where

$$u = \begin{cases} \mu_2^2, & I = R, \\ \mu_1 \mu_2^2, & I = R^\vee. \end{cases} \quad (21)$$

Both choices satisfy $uI \subseteq \mathcal{O}_f$. Let $a_0 = f(0)$, and let $M_f^{(h)}$ contain the first h rows of M_f . Before output reversal and added noise, NP26's complete error map is

$$\mathsf{T}_{\text{NP},h}(e) = M_f^{(h)} V_f^{-1} \sigma(ue), \quad e \in K_{\mathbb{R}}.$$

The following identity only needs the inclusion $uI \subseteq \mathcal{O}_f$. Both operator norms use canonical input and Euclidean output coordinates.

Proposition 2 (Transferring NP26's multipliers). *Let $f \in \mathcal{I}_{n,B}$, $I = R$ or $R^\vee$, and $0 \neq u \in K$ satisfy $uI \subseteq \mathcal{O}_f$. Then*

$$c_{\text{NP}} = -\frac{a_0 u}{\theta f'(\theta)} \in R^\vee I^{-1} \setminus \{0\}. \quad (22)$$

For every $1 \leq h \leq n$,

$$\mathsf{T}_{\text{NP},h} = L_{c_{\text{NP}},h}, \quad \|\mathsf{T}_{\text{NP},h}\|_{\text{op}} = \|L_{c_{\text{NP}},h}\|_{\text{op}}. \quad (23)$$

If u has polynomial bit length, so does c_{NP} . Given u , the multiplier c_{NP} is computable in polynomial time.

Proof. Set

$$c_0 = -a_0/(\theta f'(\theta)). \quad (24)$$

The denominator is nonzero since f is irreducible and $n \geq 2$. For $y \in K_{\mathbb{R}}$, let y_0 be its constant power-basis coordinate. Lagrange interpolation at zero proves $\text{Tr}(c_0 y) = y_0$, hence $c_0 \in \mathcal{O}_f^\vee$. For $x \in I$ and $r \in R$, the inclusion $uI \subseteq \mathcal{O}_f$ implies $urx \in \mathcal{O}_f$ and $\text{Tr}(c_0 urx) \in \mathbb{Z}$. Thus $c_0 u \in R^\vee I^{-1}$.

By the definition of M_f in Definition 16, the i th output coordinate, indexed from zero, is the constant coefficient of $\theta^i u e$. Hence

$$(M_f^{(h)} V_f^{-1} \sigma(u e))_i = (\theta^i u e)_0 = \text{Tr}(c_0 u e \theta^i) = L_{c_0 u, h}(e)_i.$$

Here $(\theta^i u e)_0$ means the constant power-basis coordinate. This proves the map and norm identities. Output reversal preserves the norm. Multiplication and inversion in $\mathbb{Q}[x]/(f)$ are polynomial-time rational operations with polynomial output bit length. $\square$

The identical linear maps have identical Gaussian output covariance. Independent Gaussian noise added by NP26 can only increase that covariance, so any bound on NP26's complete loss also bounds the trace norm. Moreover, $j_f \mid N(\mathbf{c}_f)$, and NP26's prime condition $q \nmid a_0 N(\mathbf{c}_f)$ implies $q \nmid j_f$ [15, Theorem 5.3].

The transferred multiplier may depend on q if NP26's multipliers do. Appendix C removes this dependence by minimizing the norm over all legal trace multipliers. The minimum exists by lattice discreteness, and Lemma 5 bounds its description length. Corollary 2 proves that one multiplier meets both the general bound and every complete NP26 bound for n outputs. The same appendix checks the common $\sqrt{n}$ norm for $f = x^n + 1$ when n is a power of two.

6 Conclusion and Open Problems

We prove lower-loss multipliers without NP26's entry condition and asymptotic coverage for $\mathcal{I}_{n,B}$ and the represented fields at shared sufficient lattice-hardness parameters. We also express NP26's complete linear error map as a legal trace map. For n outputs, we prove the existence of one short multiplier, independent of the modulus, whose norm meets both our general bound and the bounds achieved by NP26's multipliers. Finite implementation preserves the smaller loss guarantee up to a constant factor.

A limit of integral dual trace multipliers. For $0 \neq c \in R$ and $1 \leq h \leq n$, we have $\|L_{c,h}\|_{\text{op}} \geq \sqrt{n}$ (Proposition 4, proved in Appendix C). Independent added noise cannot cancel the trace error, so $o(\sqrt{n})$ loss is impossible within these trace maps. Our larger general exponents need not be optimal.

Finding general multipliers efficiently without advice remains open. We also leave open sharper general bounds and hardness guarantees at concrete MP-LWE scheme parameters.

Use of generative AI. The authors conceived and developed the central theme, research questions, and core ideas of this paper. OpenAI’s ChatGPT and Codex were used as supporting tools for language polishing, formatting, mathematical derivations, and the refinement and verification of proofs. The authors have thoroughly reviewed, rigorously verified, and approved all content generated or refined with AI assistance. The authors take full responsibility for the accuracy, correctness, and integrity of all proofs and other content in this paper.

References

1. Roşca, M., Sakzad, A., Stehlé, D., Steinfeld, R.: Middle-product learning with errors. In: *Advances in Cryptology – CRYPTO 2017*. LNCS, vol. 10403, pp. 283–297. Springer (2017). https://doi.org/10.1007/978-3-319-63697-9_10
2. Steinfeld, R., Sakzad, A., Zhao, R.K.: Practical MP-LWE-based encryption balancing security-risk versus efficiency. *Designs, Codes and Cryptography* **87**(12), 2847–2884 (2019). <https://doi.org/10.1007/s10623-019-00654-5>, <https://eprint.iacr.org/2019/1179>
3. Bai, S., Boudgoust, K., Das, D., Roux-Langlois, A., Wen, W., Zhang, Z.: Middle-product learning with rounding problem and its applications. In: *Advances in Cryptology – ASIACRYPT 2019*. LNCS, vol. 11921, pp. 55–81. Springer (2019). https://doi.org/10.1007/978-3-030-34578-5_3
4. Lombardi, A., Vaikuntanathan, V., Vuong, T.D.: Lattice trapdoors and IBE from middle-product LWE. In: *Theory of Cryptography – TCC 2019*. LNCS, vol. 11891, pp. 24–54. Springer (2019). https://doi.org/10.1007/978-3-030-36030-6_2
5. Le, H.Q., Duong, D.H., Susilo, W., Pieprzyk, J.: Trapdoor delegation and HIBE from middle-product LWE in standard model. In: *Applied Cryptography and Network Security – ACNS 2020*. LNCS, vol. 12146, pp. 130–149. Springer (2020). https://doi.org/10.1007/978-3-030-57808-4_7
6. Fan, J., Lu, X., Au, M.H.: Adaptively secure identity-based encryption from middle-product learning with errors. In: *Information Security and Privacy – ACISP 2023*. LNCS, vol. 13915, pp. 320–340. Springer (2023). https://doi.org/10.1007/978-3-031-35486-1_15
7. Fan, J., Lu, X., Au, M.H., Yiu, S.M.: Compact adaptively secure identity-based encryption from middle-product learning with errors. In: *Information and Communications Security – ICICS 2025*. LNCS, vol. 16217, pp. 218–236. Springer (2026). https://doi.org/10.1007/978-981-95-3540-8_12
8. Yang, N., Yang, S., Zhao, Y., Wu, W.: Inner product encryption from middle-product learning with errors. In: *Security and Privacy in Social Networks and Big Data – SocialSec 2022*. CCIS, vol. 1663, pp. 94–113. Springer (2022). https://doi.org/10.1007/978-981-19-7242-3_7
9. Yang, N., Yang, S., Zhao, Y., Wu, W., Wang, X.: Inner product encryption from middle-product learning with errors. *Computer Standards & Interfaces* **87**, 103755 (2024). <https://doi.org/10.1016/j.csi.2023.103755>
10. Bai, S., Das, D., Hiromasa, R., Roşca, M., Sakzad, A., Stehlé, D., Steinfeld, R., Zhang, Z.: MPSign: A signature from small-secret middle-product learning with errors. In: *Public-Key Cryptography – PKC 2020*. LNCS, vol. 12111, pp. 66–93. Springer (2020). https://doi.org/10.1007/978-3-030-45388-6_3
11. Hiromasa, R.: Digital signatures from the middle-product LWE. In: *Provable Security – ProvSec 2018*. LNCS, vol. 11192, pp. 239–257. Springer (2018). https://doi.org/10.1007/978-3-030-01446-9_14
12. Das, D., Au, M.H., Zhang, Z.: Ring signatures based on middle-product learning with errors problems. In: *Progress in Cryptology – AFRICACRYPT 2019*. LNCS, vol. 11627, pp. 139–156. Springer (2019). https://doi.org/10.1007/978-3-030-23696-0_8
13. Lin, H., Sun, S.F., Wang, M., Liu, J.K., Wang, W.: Shorter linkable ring signature based on middle-product learning with errors problem. *The Computer Journal* **66**(12), 2974–2989 (2023). <https://doi.org/10.1093/comjnl/bxac141>

14. Roşca, M., Stehlé, D., Wallet, A.: On the Ring-LWE and Polynomial-LWE problems. In: *Advances in Cryptology – EUROCRYPT 2018*. LNCS, vol. 10820, pp. 146–173. Springer (2018). https://doi.org/10.1007/978-3-319-78381-9_6
15. Njah Nchiwo, R.Y., Pellet-Mary, A.: Improved reduction from RLWE to MP-LWE. In: *Public-Key Cryptography – PKC 2026*. LNCS, vol. 16551, pp. 167–196. Springer (2026). https://doi.org/10.1007/978-3-032-26731-3_6, <https://eprint.iacr.org/2026/342>, full version cited
16. Peikert, C., Pepin, Z.: Algebraically structured LWE, revisited. *Journal of Cryptology* **37**, 28 (2024). <https://doi.org/10.1007/s00145-024-09508-3>, conference version: TCC 2019
17. Pellet-Mary, A., Xia, W.: A unified reduction from RLWE to MP-LWE. *Cryptology ePrint Archive*, Paper 2026/2026 (2026), <https://eprint.iacr.org/2026/2026>, version dated September 14, 2026
18. Yang, S., Huang, X.: Universal product learning with errors: A new variant of LWE for lattice-based cryptography. *Theoretical Computer Science* **915**, 90–100 (2022). <https://doi.org/10.1016/j.tcs.2022.02.032>
19. Ling, C., Mendelsohn, A.: Middle-products of skew polynomials and learning with errors. In: *Cryptography and Coding – IMACC 2023*. LNCS, vol. 14421, pp. 199–219. Springer (2024). https://doi.org/10.1007/978-3-031-47818-5_11
20. Tezuka, M., Tanaka, K.: A survey on middle-product learning with errors cryptography. In: *Mathematical Foundations for Post-Quantum Cryptography, Mathematics for Industry*, vol. 40, pp. 251–271. Springer (2026). https://doi.org/10.1007/978-981-96-1218-5_13
21. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: *Advances in Cryptology – EUROCRYPT 2013*. LNCS, vol. 7881, pp. 35–54. Springer (2013). https://doi.org/10.1007/978-3-642-38348-9_3, lemma 3.2 in the conference version
22. Attema, T., Cramer, R., Xing, C.: A note on short invertible ring elements and applications to cyclotomic and trinomials number fields. *Cryptology ePrint Archive*, Paper 2019/1200 (2019), <https://eprint.iacr.org/2019/1200>, revised March 14, 2021
23. Rosser, J.B., Schoenfeld, L.: Approximate formulas for some functions of prime numbers. *Illinois Journal of Mathematics* **6**(1), 64–94 (1962). <https://doi.org/10.1215/ijm/1255631807>
24. Peikert, C., Regev, O., Stephens-Davidowitz, N.: Pseudorandomness of Ring-LWE for any ring and modulus. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 461–473. ACM (2017). <https://doi.org/10.1145/3055399.3055489>, <https://eprint.iacr.org/2017/258>, theorem numbering refers to the full version dated June 4, 2020
25. Mehlhorn, K., Sagraloff, M., Wang, P.: From approximate factorization to root isolation with application to cylindrical algebraic decomposition. *Journal of Symbolic Computation* **66**, 34–69 (2015). <https://doi.org/10.1016/j.jsc.2014.02.001>, <https://arxiv.org/abs/1301.4870>
26. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. *Journal of the ACM* **60**(6), 43:1–43:35 (2013). <https://doi.org/10.1145/2535925>, <https://eprint.iacr.org/2012/230>, theorem numbering refers to the full version dated June 25, 2013

A Finite-Precision Implementation

We prove the finite implementation guarantee stated in Proposition 3. The multiplier description is supplied as advice.

A.1 Polynomial-size descriptions

Lemma 5. *Let $f \in \mathcal{I}_{n,B}$, $I = R$ or $R^\vee$, and $0 \neq c \in R^\vee I^{-1}$ satisfy $\|L_{c,n}\|_{\text{op}} \leq C$ for some $C > 0$. Then c has a rational power-basis description of $\text{poly}(n, \log(B+1), \log(C+1))$ bits. The HNF-derived bases of R and I , and the element 1, have descriptions of $\text{poly}(n, \log(B+1))$ bits. In particular, a polynomial norm bound implies a polynomial-size multiplier description. For fixed (f, I) , we may choose the multiplier and bases once, independently of q , the secret, and the samples.*

Proof (of Lemma 5). The integral trace-pairing matrix of the power basis has determinant of absolute value Δ_f . The adjugate formula implies

$$\mathcal{O}_f \subseteq R \subseteq R^\vee \subseteq \mathcal{O}_f^\vee \subseteq \Delta_f^{-1} \mathcal{O}_f. \quad (25)$$

In power coordinates, $\Delta_f \mathbb{Z}^n \subseteq \Delta_f I \subseteq \mathbb{Z}^n$. An HNF basis of $\Delta_f I$ therefore has entries bounded by Δ_f . Dividing by Δ_f produces an I -basis with coordinates bounded by one and description length $O(n^2 \log(\Delta_f + 1))$. The same argument applies to R , and (4) bounds the bit length of Δ_f .

In both cases, the legal multiplier c lies in $R^\vee$. Consequently its trace coordinates

$$\boldsymbol{\tau} = (\tau_i)_{0 \leq i < n} = L_{c,n}(1), \quad \tau_i = \text{Tr}(c\theta^i),$$

are integers. Since $\|1\| = \sqrt{n}$, the norm bound implies $\|\boldsymbol{\tau}\| \leq \sqrt{n}C$. Write $c = \sum_{k=0}^{n-1} d_k \theta^k$ and $\mathbf{d} = (d_k)_{0 \leq k < n}$. The integer trace-pairing matrix satisfies

$$\mathbf{T}_f \mathbf{d} = \boldsymbol{\tau}, \quad \mathbf{T}_f = (\text{Tr}(\theta^{i+k}))_{0 \leq i, k < n}, \quad |\det \mathbf{T}_f| = \Delta_f \geq 1.$$

Every root has magnitude at most $1 + B$, so each entry of $\mathbf{T}_f$ has magnitude at most $n(1 + B)^{2n-2}$. Cramer's rule expresses each d_k as an integer determinant divided by $\det \mathbf{T}_f$. The entry bounds for $\mathbf{T}_f$ and $\boldsymbol{\tau}$ bound both numerator and denominator lengths by $\text{poly}(n, \log(B+1), \log(C+1))$. The element 1 has the power-basis coordinates $(1, 0, \dots, 0)$. This argument bounds the description of an existing multiplier. It does not compute a suitable multiplier. $\square$

A.2 Computing and encoding samples

Proposition 3. *Let $f \in \mathcal{I}_{n,B}$, let q be prime with $q \nmid j_f$, and fix $I = R$ or $R^\vee$, $1 \leq h \leq n$, and $\alpha, \beta > 0$. Let $0 \neq c \in R^\vee I^{-1}$ have polynomial bit length. Let $C > 0$ be a represented upper bound satisfying*

$$\|L_{c,n}\|_{\text{op}} \leq C, \quad \beta \geq 2\alpha C.$$

Assume that the public parameters, C , and bases of R and I have polynomial-size descriptions. Each positive real parameter $u \in \{\alpha, \beta, C\}$ must satisfy $|\log u| \leq \text{poly}(n)$ and be approximable to k bits in time $\text{poly}(n, k)$. Given c and C , there is a polynomial-time finite-string reduction from Ring-LWE of rate α to MP-LWE of rate β , for any polynomially bounded numbers of output fractional bits and samples. Each encoded experiment has negligible total variation error. Input precision may depend on the field and basis. Output encoding can remain identical across reductions. Advice makes the reduction nonuniform.

Computing samples. Given the multiplier, we compute a' through the isomorphism $\mathcal{O}_f/q\mathcal{O}_f \simeq R/qR$ using modular linear algebra. For the basis (u_j) of I , we use rational arithmetic modulo f and Newton's identities to compute the integer trace matrix

$$\mathbf{L}_{ij} = \text{Tr}(cu_j\theta^{i-1}) \in \mathbb{Z}, \quad 1 \leq i \leq h, \quad 1 \leq j \leq n. \quad (26)$$

We decode b to coordinates $\mathbf{y}$ and compute the trace image as $\mathbf{L}\mathbf{y}$. Rounding $\mathbf{y}$ down to r_{in} fractional bits changes each output coordinate by at most $\|\mathbf{L}\|_{\infty}2^{-r_{\text{in}}}$. Let E_I contain the u_j in orthonormal real coordinates, using $\sqrt{2}\text{Re}\sigma_k$ and $\sqrt{2}\text{Im}\sigma_k$ for one embedding per complex pair. Using its Gram matrix $\mathbf{G}_I = E_I^{\top}E_I$, we express the added-noise covariance as

$$W = (\beta q)^2 \mathbf{I}_h - (\alpha q)^2 \mathbf{L} \mathbf{G}_I^{-1} \mathbf{L}^{\top}. \quad (27)$$

The factor-two condition $\beta \geq 2\alpha C$ and $\|L_{c,h}\|_{\text{op}} \leq C$ imply

$$\frac{3}{4}(\beta q)^2 \mathbf{I}_h \preceq W \preceq (\beta q)^2 \mathbf{I}_h. \quad (28)$$

Let b_I denote the basis description length. The identity $|\det E_I| = \sqrt{\Delta_K}N(I)$ and the entry bounds imply

$$\lambda_{\min}(\mathbf{G}_I) \geq 2^{-\text{poly}(n, \log(B+1), b_I)}.$$

Here $\lambda_{\min}$ denotes the smallest eigenvalue. Root refinement [25] and Lemma 5 therefore require only polynomially many precision bits. The margin in (28) allows stable inversion and Cholesky factorization of W .

Proof (of Proposition 3). Put $\sigma_{\text{MP}} = \beta q$, $D_{\text{grid}} = 2^{r_{\text{out}}} + 2/\sigma_{\text{MP}}$, $\tau = 2^{-3n}$, and $\delta = \tau/(hD_{\text{grid}})$. We choose r_{in} so $\|\mathbf{L}\|_{\infty}2^{-r_{\text{in}}} \leq \delta$. To sample $\mathbf{G} = F_W \mathbf{Z}$ with $F_W F_W^{\top} = W$, we draw independent $U_j \leftarrow \mathcal{U}((0, 1))$, $1 \leq j \leq h$, and set $Z_j = F_N^{-1}(U_j)$. Here F_N is the standard normal cumulative distribution function. We couple each normal with its approximation through shared uniform bits. We discard tails of total probability $O(\tau)$ by clipping uniforms to $[\tau/h, 1 - \tau/h]$. On this interval the quantile and its derivative have magnitudes bounded by $2^{\text{poly}(n, \log h)}$. We integrate the exponential series and use bisection with the covariance margin. The resulting polynomial-bit approximations have coordinate error at most δ , except with probability $O(\tau)$.

To bound rounding errors, we sum a normal density over grid translates and bound both monotone tails by integrals. The resulting density is at most D_{grid} near a rounding boundary. Conditional on the secret, mask, and first component, exact noisy output coordinates have standard deviation σ_{MP} . Hence the probability of any coordinate lying within 2δ of a boundary is $O(h\delta D_{\text{grid}}) = O(\tau)$. Uniform outputs satisfy a smaller bound. If neither approximation nor boundary failure occurs, the encodings agree. Polynomially capped rejection sampling produces uniform mask coefficients with failure $O(\tau)$. A union bound over polynomially many samples proves negligible error, also for adaptive requests. All approximations use polynomially many precision bits. Shifted grids cover nearest-integer rounding as well. $\square$

B Derivation of the Ideal-Lattice Parameters

We state the lattice-to-Ring-LWE hypotheses explicitly, derive the ideal-SIVP factor, and verify the parameters of Corollary 1.

The lattice-to-Ring-LWE input condition. We use [24, Corollary 7.3] for dual Ring-LWE (our $I = R^\vee$), rate α_d , and integer sample count $1 \leq t = t(n) \leq \text{poly}(n)$. Put $S_{n,t} = (nt/\log(nt))^{1/4}$. Fix $u_n = \omega(\sqrt{\log n})$ and negligible smoothing error ϵ with $\log(1/\epsilon) \leq \text{poly}(n)$. For a public prime q satisfying

$$0 < \sqrt{2\pi}\alpha_d < 1, \quad \sqrt{2\pi}\alpha_d q \geq S_{n,t}u_n, \quad (29)$$

the cited theorem reduces ideal-DGS to spherical t -sample dual Ring-LWE in quantum polynomial time. For nonzero fractional input ideal $\mathfrak{a}$, the DGS threshold is

$$g_{\text{DGS}}(\mathfrak{a}) = \max \left\{ \frac{\eta_\epsilon(\mathfrak{a})S_{n,t}u_n}{\sqrt{\pi}\alpha_d}, \quad \frac{\sqrt{2n}}{\lambda_1(\mathfrak{a}^\vee)} \right\} \quad (30)$$

The DGS input ideal $\mathfrak{a}$ varies, while $I = R^\vee$ stays fixed. We obtain (29)–(30) by substituting $\xi = \sqrt{2\pi}\alpha_d$ for the PRS Gaussian noise rate. No further condition restricts how q factors in R .

From Gaussian sampling to short independent vectors. The conversion from DGS to SIVP [26, Section 4] achieves approximation factor $O(\sqrt{n}g_{\text{DGS}}(\mathfrak{a})/\lambda_n(\mathfrak{a}))$. We remove the dependence on the input ideal using the smoothing bound $\eta_\epsilon(\mathfrak{a}) \leq \sqrt{\log(n/\epsilon)}\lambda_n(\mathfrak{a})$ from [26, Lemma 2.2]. Also $\lambda_1(\mathfrak{a}^\vee)\lambda_n(\mathfrak{a}) \geq 1$: a shortest nonzero vector of $\Lambda_{\mathfrak{a}}^*$ has a nonzero integer pairing with at least one of n independent vectors of $\Lambda_{\mathfrak{a}}$ of length at most $\lambda_n(\mathfrak{a})$. We obtain the following sufficient approximation factor for all input ideals:

$$O \left(n + \frac{\sqrt{n}S_{n,t}u_n\sqrt{\log(n/\epsilon)}}{\alpha_d} \right). \quad (31)$$

A Ring-LWE distinguisher with inverse-polynomial advantage therefore solves ideal-SIVP at this factor in quantum polynomial time. At fixed MP-LWE noise

rate β , substituting $\alpha_d = \beta/\Gamma$ shows that the sufficient modulus and the noise-dependent term in (31) are proportional to Γ . Appendix D states the formal comparison and a numerical example.

Proof (of Corollary 1). For $n \geq 2$, $t_n/\log(nt_n) \leq 2/\log 2 < 16$, so $S_{n,t_n} \leq 2n^{1/4} \leq s_n$. Also $u_n = \Theta(\sqrt{\log n \log \log n})$ is admissible and $\alpha_{d,n}q_n \geq s_n u_n$. Since $\alpha_{d,n} \rightarrow 0$, the rate condition holds eventually. Set $\epsilon = e^{-(\log n)^2}$ and use the multiplier advice of Lemma 5. By Theorem 1, the multiplier satisfies $2\|L_{c,n}\|_{\text{op}} \leq \Gamma_d(n, B)$. Substituting $\alpha_{d,n}$ into (29) and (31), then applying Lemma 1, proves the composed reduction. Here $S_{n,t_n} = \Theta(n^{1/4})$, $\sqrt{\log(n/\epsilon)} = \Theta(\log n)$, and $\Gamma_d(n, B) = 10B^2n(n+1) = \Theta_B(n^2)$. Equations (14) and (18) imply the displayed bounds. Let D be a distinguisher for the encoded MP-LWE problem. Fix the input precision required by Proposition 3, and let $\mathcal{B}$ encode its Ring-LWE input, apply the finite trace transformation, and run D . Proposition 3 shows that the advantage of $\mathcal{B}$ on the Ring-LWE and uniform experiments is at least the advantage of D minus a negligible function. We apply [24, Corollary 7.3] to the complete algorithm $\mathcal{B}$. Its behavior on intermediate input distributions is its actual behavior and requires no comparison with an ideal trace oracle. The chosen precision is independent of the worst-case input ideal. $\square$

C A Short Minimizer and Its Finite Norm Bound

One short multiplier that retains both bounds. For the comparison with advice fixed independently of q , we use all n output coefficients. Define

$$C_I^*(f) = \min_{0 \neq c \in R^{\vee I-1}} \|L_{c,n}\|_{\text{op}}.$$

The next corollary proves that this minimum is attained by a multiplier with a short description.

Corollary 2 (Combining the two loss guarantees). *Let $f \in \mathcal{I}_{n,B}$ and $I = R$ or $R^\vee$. Put $\Gamma_R = \Gamma_p(n, B)$ and $\Gamma_{R^\vee} = \Gamma_d(n, B)$. There exists a minimizing multiplier $c_I^*(f)$, fixed by f, I , with $\text{poly}(n, \log(B+1))$ bits in rational power-basis coordinates. For every $0 \neq \mu_2 \in \mathfrak{c}_f$ and, when $I = R^\vee$, every $0 \neq \mu_1 \in (R^\vee)^{-1}$, put $u = \mu_2^2$ for $I = R$ and $u = \mu_1\mu_2^2$ for $I = R^\vee$. Then*

$$\|L_{c_I^*(f),n}\|_{\text{op}} = C_I^*(f) \leq \min \left\{ \Gamma_I/2, \|\mathbf{T}_{\text{NP},n}\|_{\text{op}} \right\}. \quad (32)$$

In particular, an $O(F(n))$ bound for NP26's full-output noise loss on a polynomial family implies an $O(F(n))$ trace reduction with polynomial-size advice independent of q , at every prime q with $q \nmid j_f$. The finite trace loss is at most the smaller of Γ_I and four times the full-output NP26 bound.

Proof (of Corollary 2). The first trace coordinate is $e \mapsto \text{Tr}(ce)$. Its Euclidean row norm is $(\sum_j |\sigma_j(c)|^2)^{1/2}$. Hence

$$\|\sigma(c)\|_2 \leq \|L_{c,n}\|_{\text{op}}.$$

The ideal $R^\vee I^{-1}$ is a lattice in canonical space. For every fixed H , it contains only finitely many c with $\|L_{c,n}\|_{\text{op}} \leq H$. The weighted argument in Theorem 1 supplies a nonzero candidate with norm at most $\Gamma_I/2$. Minimizing over this finite nonempty set attains the global minimum. The definition of the set uses f, I and does not use q . Proposition 2 supplies a legal candidate for each NP26 choice, which proves (32).

For the description length, the minimizing multiplier is legal and satisfies $\|L_{c_I^*(f),n}\|_{\text{op}} \leq \Gamma_I/2 = \text{poly}(n, B)$. The integer-trace argument in Lemma 5 therefore proves polynomial bit length in rational power-basis coordinates.

For finite implementation, we also need a represented norm bound. Since $\Delta_f c_I^*(f)$ is a nonzero algebraic integer, the arithmetic-geometric mean inequality and the first-row bound imply $C_I^*(f) \geq \sqrt{n}/\Delta_f$. There is $H_I(f) = 2^k$ with $k \in \mathbb{Z}$ and $C_I^*(f) \leq H_I(f) < 2C_I^*(f)$. Let $B_f = \max(1, |a_0|, \dots, |a_{n-1}|)$ be the coefficient height of f . The general bound also applies with B_f in place of B . Set

$$\overline{C}_I(f) = \min\{\Gamma_I(n, B_f)/2, H_I(f)\}.$$

This is a rational number of polynomial bit length, fixed by f, I . It satisfies $C_I^*(f) \leq \overline{C}_I(f) \leq 2C_I^*(f)$. The formulas for $\Gamma_I(n, B)$ are nondecreasing in B , so $\Gamma_I(n, B_f) \leq \Gamma_I(n, B)$. By Lemma 1 and Proposition 3, the finite trace loss is $2\overline{C}_I(f) \leq \min\{\Gamma_I, 4\|\mathbf{T}_{\text{NP},n}\|_{\text{op}}\}$. The condition $q \nmid j_f$ permits the trace reduction. $\square$

The real number $C_I^*(f)$ itself need not have an exact rational encoding. The proof uses the rational upper bound $\overline{C}_I(f)$ instead. Both $c_I^*(f)$ and $\overline{C}_I(f)$ can be supplied as short advice. Finding a minimizing multiplier efficiently remains open. The minimization is essential: the explicit general formula Γ_I can exceed a bound tailored to a particular polynomial. For $h < n$, deleting rows preserves the full-output bound in (32). This does not compare against a bound optimized specifically for h outputs.

The positive cases in NP26. NP26's positive results concern the class $D(f) \leq n$ and the random bounded-coefficient polynomials studied in its experiments [15]. Proposition 2 and Corollary 2 apply to every polynomial for which the complete NP26 chain has a proved small loss.

As a concrete check, take $f = x^n + 1$ with n a power of two and dual Ring-LWE. Here $R = \mathbb{Z}[\theta]$, $R^\vee = n^{-1}R$, and $c_0 = 1/n$. The explicit choices $\mu_1 = n$ and $\mu_2 = 1$ satisfy NP26's ideal conditions and produce $c_{\text{NP}} = 1$. Fourier orthogonality proves $V_f^* V_f = n\mathbf{I}_n$, so both linear error maps have norm $\sqrt{n}$. This check recovers the familiar square-root loss by direct substitution in NP26's maps and the trace map.

C.1 A lower bound for integral dual multipliers

The norm cannot decrease below square-root order, even for a defining polynomial with particularly favorable geometry.

Proposition 4. *For every nonzero multiplier $c \in R$ for dual Ring-LWE and $1 \leq h \leq n$,*

$$\|L_{c,h}\|_{\text{op}} \geq \sqrt{n} |N_{K/\mathbb{Q}}(c)|^{1/n} \geq \sqrt{n}. \quad (33)$$

Proof. The θ^0 row of $L_{c,h}$ is $e \mapsto \text{Tr}(ce)$. On the canonical real subspace, this row has Euclidean norm $(\sum_j |\sigma_j(c)|^2)^{1/2}$, counting both conjugate embeddings. The operator norm is at least the norm of this row. The arithmetic–geometric mean inequality proves the first bound. The second follows because $N_{K/\mathbb{Q}}(c)$ is a nonzero integer. $\square$

D Comparisons of Sufficient Hardness Parameters

Section 4.2 proves the ideal-SIVP reduction at the parameters used for coverage. Here we record how changing the trace loss affects the sufficient parameters, and compare the bounds numerically.

D.1 Comparing two trace-loss bounds

Proposition 5. *Fix $f \in \mathcal{I}_{n,B}$ and dual Ring-LWE over K_f . Choose a prime $q \nmid j_f$, $1 \leq h \leq n$ output coefficients, an integer sample bound $1 \leq t \leq \text{poly}(n)$, and MP-LWE noise rate $\beta > 0$. Set $S_{n,t} = (nt/\log(nt))^{1/4}$. Fix a positive sequence $u_n = \omega(\sqrt{\log n})$ and a negligible $0 < \epsilon < 1$ with $\log(1/\epsilon) \leq \text{poly}(n)$. Let $c \in R \setminus \{0\}$ have polynomial bit length. Let $\Gamma \in \mathbb{Q}_{>0}$ have polynomial bit length and satisfy $\Gamma \geq 2 \|L_{c,n}\|_{\text{op}}$. Set $\alpha_d = \beta/\Gamma$. If $\sqrt{2\pi}\beta/\Gamma < 1$ and*

$$q \geq \frac{\Gamma S_{n,t} u_n}{\sqrt{2\pi}\beta}, \quad (34)$$

then there is a reduction from ideal-SIVP to MP-LWE with sufficient approximation factor

$$\gamma_\Gamma = C_{\text{lat}} \left(n + \frac{\sqrt{n} S_{n,t} u_n \sqrt{\log(n/\epsilon)} \Gamma}{\beta} \right), \quad (35)$$

where C_{lat} is a suitable lattice-reduction constant shared by all noise-loss bounds being compared. The reductions are quantum and nonuniform when c is advice.

For $i = 0, 1$, let $c_i \in R \setminus \{0\}$ and $\Gamma_i \in \mathbb{Q}_{>0}$ have polynomial bit length, with $\Gamma_i \geq 2 \|L_{c_i,n}\|_{\text{op}}$ and $\sqrt{2\pi}\beta/\Gamma_i < 1$. Suppose $\Gamma_1 \leq \Gamma_0$. If $q \geq \Gamma_0 S_{n,t} u_n / (\sqrt{2\pi}\beta)$, both reductions use identical q, n, h, β, t and encoding. Replacing Γ_0 by Γ_1 multiplies the Ring-LWE rate by Γ_0/Γ_1 and the term $\gamma_\Gamma - C_{\text{lat}}n$ by Γ_1/Γ_0 .

Proof. We apply Lemma 1 at $\alpha_d = \beta/\Gamma$ and substitute into eqs. (29) to (31). We use Proposition 3 to compute the input precision for each multiplier. The parameters q, n, h, β, t and the output encoding determine the MP-LWE distribution. $\square$

D.2 A numerical comparison

Titanium-CCA Std128 uses $n = 1024$, $h = 511 + 256 = 767$, $t = 10$, and $q_0 = 118273$, with binomial-difference errors of standard deviation $\sqrt{2}$ ([2, Table 2 and Algorithm 5]). We compare Gaussian MP-LWE at the same n, h, t and fixed rate $\beta_0 = \sqrt{2}/q_0$. The standard deviation $\beta_0 q$ varies with q . For $B = 2$, $s_{1024} = 12 > S_{1024,10}$, and illustrative margin $u_{1024} = 16$, Table 1 evaluates $Q_\Gamma = \lceil \Gamma s_{1024} u_{1024} / \beta_0 \rceil$. Each loss Γ includes the finite factor of two. NP26’s factor follows its Theorem 5.3 and proof of Corollary 5.5 [15].

Table 1. Illustrative modulus scales from the asymptotic bounds.

Reduction	Γ	$\log_2 Q_\Gamma \approx$
NP26, general	$1024 \lceil \sqrt{nh} \rceil n^{21} B^{14}$	267.73
This work, general	41 984 000	49.26
PP24/ours, $x^n + 1$	64	29.94

NP26’s row requires $D(f) \leq n$, $q \geq 8n$, and $q \nmid f(0) \text{Disc}(f)$. Our rows require $q \nmid j_f$. The $x^n + 1$ row uses a power-of-two degree and was already supported by PP24. For the general bound, $\Gamma = 10B^2n(n+1) = 41\,984\,000$ yields $Q_\Gamma = 674\,148\,631\,232\,317$ and $\log_2 Q_\Gamma \approx 49.26006$. The general sufficient threshold falls from about 268 to 50 modulus bits, but still exceeds Titanium’s $\log_2 q_0 \approx 16.85$. The table compares sufficient bounds, not minimum moduli.

If the standard deviation σ_{MP} is fixed instead, then $\alpha_d = \sigma_{\text{MP}}/(\Gamma q)$ and (29) requires $\sqrt{2\pi}\sigma_{\text{MP}}/\Gamma \geq S_{n,t}u_n$. Increasing q cannot help this inequality. The Gaussian comparison does not establish security for Titanium’s different error distribution.

E Auxiliary Details for the General Norm Bounds

The main proof retains the discriminant estimate, the positive Gram matrix argument, the Poisson-scale integration, and the final primal and dual bounds. We record three elementary supporting calculations here.

E.1 The weighted Minkowski volume

Proof (of Lemma 2). In r_1 real and r_2 unscaled complex coordinates, the co-volume of J is $2^{-r_2} \sqrt{\Delta_K} N(J)$. The symmetric convex body $v_j |\sigma_j(x)| \leq r$ has volume

$$\frac{2^{r_1} \pi^{r_2} r^n}{\prod_j v_j}.$$

The product counts each complex pair twice. Minkowski’s convex body theorem applies whenever

$$r^n > (2/\pi)^{r_2} \sqrt{\Delta_K} N(J) \prod_j v_j.$$

Decrease r to this threshold. A bounded body contains only finitely many lattice points, so a nonzero point remains at the threshold. $\square$

E.2 Comparing Poisson kernels at two radii

Proof (of (11)). Assume $0 \leq r \leq \rho < 1$ and $|w| \leq 1$. Write $w = R_0 e^{i\phi}$ with $0 \leq R_0 \leq 1$. For fixed R_0 , put $t = \cos \phi$. The ratio is

$$\frac{P(rw)}{P(\rho w)} = \frac{1 - r^2 R_0^2}{1 - \rho^2 R_0^2} \frac{1 - 2\rho R_0 t + \rho^2 R_0^2}{1 - 2r R_0 t + r^2 R_0^2}.$$

The derivative of the second fraction with respect to t is

$$\frac{-2R_0(\rho - r)(1 - r\rho R_0^2)}{(1 - 2r R_0 t + r^2 R_0^2)^2} \leq 0.$$

The first fraction is positive, so the ratio is minimized at $t = 1$. At that angle the ratio equals $(1 - \rho R_0)(1 + r R_0)/((1 + \rho R_0)(1 - r R_0))$. Its logarithmic derivative with respect to R_0 is $2r/(1 - r^2 R_0^2) - 2\rho/(1 - \rho^2 R_0^2) \leq 0$, since $a/(1 - a^2 R_0^2)$ increases with $a \in [0, 1)$. Its value at $R_0 = 1$ is the factor in (11). $\square$

E.3 Inversion and integer rounding in the primal bound

For $n \geq 2$ and $B \geq 1$, put $A_B = \log(B^2 n(n+1)) + 1/2$. The primal proof establishes $x \log(1+x) \leq 4A_B$ for $x = F_f(1 - 1/n)/n^2$. We verify both estimates used after (15).

Proof. If $H > 0$, set $x_0 = 2H/\log(1+H)$ and $v = \sqrt{1+H} > 1$. The inequality $\log v \leq v - 1$ implies $x_0 \geq v + 1$, whence $\log(1+x_0) \geq \frac{1}{2} \log(1+H)$ and $x_0 \log(1+x_0) \geq H$. Monotonicity of $x \log(1+x)$ proves the second inequality in (15) with $H = 4A_B$. We have $A_B \leq \ell_{n,B}$ and $A_B \geq 2 \log n \geq \lfloor \log_2 n \rfloor$, since $\log 2 > 1/2$. Thus $\log(1 + 4A_B) \geq v_n \log 2 \geq v_n/2$, and $x \leq 16\ell_{n,B}/v_n$. $\square$

F Asymptotic Coverage with Prescribed Polynomials

The finite counting theorem can preserve a prescribed list while the coverage proportions tend to one. The list may grow with the degree.

Corollary 3. *Fix an integer $B \geq 1$ and choose real numbers $Q_n \geq 21$ with $Q_n/(n \log n) \rightarrow \infty$. Prescribe a set $\mathcal{S}_n \subseteq \mathcal{I}_{n,B}$ satisfying $|\mathcal{S}_n| \leq \lfloor \frac{3Q_n}{5n \log H_n} \rfloor$. For each $K \in \mathcal{K}_{n,B}$, fix $f_K \in \mathcal{I}_{n,B}$ with $\mathbb{Q}[x]/(f_K) \simeq K$. For sufficiently large n with $\mathcal{I}_{n,B} \neq \emptyset$, there is a prime $q_n \in [Q_n, 2Q_n]$ with $q_n \nmid j_f$ for every $f \in \mathcal{S}_n$, and both proportions satisfy*

$$\frac{|\mathcal{A}_{n,B}(q_n)|}{|\mathcal{I}_{n,B}|}, \quad \frac{|\{K \in \mathcal{K}_{n,B} : q_n \nmid j_{f_K}\}|}{|\mathcal{K}_{n,B}|} \geq 1 - O_B(n \log n / Q_n) \rightarrow 1. \quad (36)$$

Proof. The list bound implies $|\mathcal{S}_n| a_{Q_n} \leq 1/2$, so both excluded fractions are at most $4a_{Q_n}$ by Theorem 2. Since $\log H_n = O_B(\log n)$ and $Q_n/(n \log n) \rightarrow \infty$, both excluded fractions tend to zero. $\square$

A common finite target. Select a prime q as in Theorem 2. For the computational conclusion, suppose the selected prime q , the public parameters, and the ring and ideal bases have polynomial-size descriptions. Suppose also that each positive real parameter u satisfies $|\log u| \leq \text{poly}(n)$ and, for every precision $p \geq 1$, is approximable to p bits in time $\text{poly}(n, p)$. Fix $1 \leq h \leq n$, $\beta > 0$, integer sample bound $1 \leq t \leq \text{poly}(n)$, and an encoding with polynomially many fractional bits. Then, for $I = R$ or $R^\vee$ and every $f \in \mathcal{I}_{n,B}$ with $q \nmid j_f$, Ring-LWE at noise rate β/Γ_I reduces to the same $\text{MP-LWE}_{q,n,h,\beta}^{[t]}$ with multiplier advice. Here $\Gamma_R = \Gamma_p(n, B)$ and $\Gamma_{R^\vee} = \Gamma_d(n, B)$.

Applying Theorem 1 to each polynomial proves this claim: n, B, h, β, t, q and the output encoding are fixed, and so is Γ_I . Only the input field, its bases, and its multiplier advice vary.