Decision Attacks on Known-Covariance Ring-LWE via Public Gram Determinants

Anis Bkakria

IRT SystemX, France

Abstract. Known-Covariance Ring-LWE (KCov-RLWE) was introduced by Micciancio and Suhl to support simulation-secure threshold encryption. It strengthens ordinary Ring-LWE by revealing an exact quadratic summary of the hidden secret and errors, and its decisional hardness was left open.

We give explicit polynomial-time decision attacks for broad families of prime moduli at ring rank one. Our key observation is that the public sample and the covariance hint together determine a bordered Gram matrix whose determinant exposes a hidden residual with a different distribution for genuine and uniform samples. After reducing the ring into finite-field components, this residual has one of two forms: a product across reciprocal components, or a Hermitian norm when a component is preserved by conjugation. The latter case yields a particularly strong distinguishing signal. Moreover, reduction modulo any divisor preserves the KCov real and random experiments exactly, so every attackable odd prime divisor of a composite modulus transfers the corresponding distinguisher with no loss of advantage.

The attack covers the one-sample setting globally in the self-reciprocal regime and, under the coefficient-product Gaussian specialization, locally in the completely split regime. More generally, we obtain a complete classification of how this residual-based attack behaves for odd prime moduli when $m \geq 2$, and extend these guarantees to arbitrary moduli possessing a suitable odd prime divisor. The classification does not imply insecurity for every parameter family: when the relevant residue fields are sufficiently large, the guaranteed distinguishing advantage in the reciprocal case can become negligible, and pure power-of-two moduli remain outside the divisor-reduction extension. Accordingly, our results invalidate the KCov pseudorandomness assumption in the quantified families; they do not show that every parameter choice permitted by the motivating Ring-PKE sketch is insecure, nor do they give a direct IND-CPA break of the complete threshold-encryption construction.

Keywords: Ring-LWE · post-quantum cryptanalysis · leakage · threshold encryption · cyclotomic rings · finite fields

1 Introduction

1.1 Motivation

Simulation-secure threshold encryption asks for more than the usual guarantee that a ciphertext hides its message. A simulator may also have to reproduce values that remain correlated with secret randomness when parties are corrupted during an execution. Micciancio and Suhl introduced Known-Covariance Ring-LWE (KCov-RLWE) in their construction of simulation-secure threshold public-key encryption precisely to capture this additional requirement [8].

KCov-RLWE starts from the usual noisy Ring-LWE relations, but reveals one extra piece of information: the exact quadratic "energy" of the hidden secret and of the errors used to generate the samples. This

additional value does not reveal the secret itself, nor any individual error term. It is an aggregate second-order relation among hidden quantities, and for that reason it may initially appear much less informative than the Ring-LWE samples themselves.

Its exactness is nevertheless important. Ordinary Ring-LWE hardness does not immediately imply pseudorandomness once such an algebraic relation is made public. The resulting security question is therefore simple to state: *do Ring-LWE samples still look random when this exact quadratic information is revealed?*

Micciancio and Suhl left the Ring-LWE version of this question open [8]. Subsequent work by Boudgoust, Jeudy, Tairi, and Wen (BJTW) gives search-hardness evidence for broad leakage classes, including the exact quadratic hint underlying Known-Covariance M-LWE [2]. Their main reduction, however, requires module rank $d > 1$, explicitly does not cover the ring setting $d = 1$, and leaves the corresponding decision extension open. This paper addresses precisely that rank-one decisional case.

We give a negative answer for broad and explicitly characterized families of prime moduli, and the same distinguishers extend without loss to composite moduli having an attackable odd prime divisor. The attack does not recover the secret or the errors, and it does not depend on detecting a visibly nonuniform projected error distribution. Instead, the exact covariance information completes a public algebraic relation that is absent from ordinary Ring-LWE. Genuine KCov-RLWE samples force the resulting public invariant to follow a special multiplicative law, whereas uniform random samples do not. This difference is enough to distinguish the two experiments.

1.2 Results and contributions

The attack is easiest to understand in the completely split setting, where the cyclotomic polynomial breaks into linear factors modulo the prime. In this case the public covariance value can be combined with the Ring-LWE sample to form a determinant that behaves differently in the real and random experiments. The important point is that the test does not try to recover the secret or the errors, and it does not rely on the short error becoming visibly nonuniform after reduction. The distinguishing bias comes instead from an exact algebraic relation made public by the covariance hint.

Main result, informally. For Known-Covariance Ring-LWE over the power-of-two cyclotomic rings considered here, the exact covariance information creates a public algebraic test that distinguishes genuine samples from uniform ones for broad families of prime moduli. The form of the test depends on how ring conjugation behaves after reduction modulo the prime. In the usual case, conjugation exchanges two different finite-field components, giving a product-type distinguisher whose strength decreases as those components become larger. In one exceptional prime-modulus class, conjugation remains inside a single component; the resulting structure is Hermitian and the distinguisher is substantially stronger.

The completely split setting is covered by the first case and gives a direct non-negligible attack for polynomial-size moduli under the stated anti-degeneracy conditions. For one sample, the Hermitian branch already gives a global attack, while the completely split reciprocal branch is covered under the coefficient-product discrete-Gaussian specialization; these provide concrete one-sample families relevant to the Ring-PKE discussion of Micciancio and Suhl [8]. More generally, every odd prime modulus falls into one of the two algebraic cases above. Thus the result gives a complete classification of the residual channel exposed by the covariance hint for prime moduli. A further divisor-reduction principle shows that a composite-modulus KCov instance inherits any attack available modulo an odd prime divisor, with identical distinguishing advantage. This does *not* mean that every modulus family, or every parameter choice permitted by the motivating Ring-PKE sketch, is broken: in the reciprocal prime branch the signal

becomes weaker as the associated residue fields grow, and moduli without a suitable odd prime divisor are not covered by this transfer.

Two features of this classification are worth highlighting. First, the self-reciprocal case is more vulnerable than residue degree alone would suggest. There the real residual becomes a single Hermitian norm while the random branch becomes a higher-dimensional Hermitian form, and their zero probabilities differ on the same q^{-1} scale as in the completely split case. Second, the argument does not assume that the short secret or error has independent projections across the different ring factors. We condition on all hidden values and use independence only from the public uniform Ring-LWE multipliers.

The reciprocal one-sample case requires a separate small refinement. In the completely split Gaussian regime, a naive zero test can lose its signal when the projected covariance vanishes. Since that covariance value is public, the distinguisher can simply ignore those components. This public gate recovers the local separation proved in Section 4.4; the Hermitian branch, by contrast, already handles every $m \geq 1$ through its active-factor test.

At the algebraic level, the main observation is a rank-one Gram completion. The exact covariance supplies the missing quadratic term needed to build a square public bordered Gram matrix from the Ring-LWE transcript. For genuine samples, the determinant of this matrix collapses to the norm of a single hidden residual. This collapse is specific to ring rank one. At higher module rank the analogous hidden factor is rectangular, and Cauchy–Binet yields a sum of minors rather than one residual. This rank-one determinant mechanism also explains why our negative result coexists with the positive KCov M-LWE evidence of BJTW [2]. Their framework already permits the exact quadratic hint and the HNF-shaped sample map, so the separation is not merely one of leakage type or matrix form. Rather, their main hardness reduction requires module rank $d > 1$ and establishes search hardness, while the square determinant collapse below occurs at ring rank $d = 1$ and yields a decision distinguisher.

The exact quantitative advantages, the residue-degree notation, and the formal anti-concentration conditions are introduced later, after the relevant algebraic objects have been defined. Sections 4 and 5 state the formal attack theorems, and Section 6 combines them into the prime-modulus classification and its composite-modulus transfer.

1.3 Attack intuition

The attack can be summarized without the cyclotomic details. The covariance hint allows us to build one scalar invariant from public information. Genuine KCov-RLWE samples force that invariant to factor in a special way. Reducing the ring modulo an irreducible factor turns this structure either into a product of two field elements or, in the self-reciprocal case, into a field norm. Uniform random samples do not have the same distribution. We distinguish the two experiments by testing how often the invariant vanishes, and we repeat the test over the available factors to amplify the gap. The short-error distribution matters only through mild anti-concentration conditions; the randomization that drives the attack comes from the public uniformly sampled Ring-LWE multipliers.

1.4 Technical overview

We now give the algebra behind the attack intuition above. From this point on we introduce a small amount of notation solely for the overview; Section 2 gives the formal definitions. Let $n \geq 2$ be a power of two, let q be an odd prime, and write $R_q = (\mathbb{Z}/q\mathbb{Z})[X]/(X^n + 1)$. A genuine KCov-RLWE transcript is generated from a public vector $\mathbf{a} \in R_q^m$ of uniformly sampled ring elements, a secret $s \in R_q$, an error vector $\mathbf{e} \in R_q^m$, and the resulting sample $\mathbf{b} \in R_q^m$. The integer m is the number of Ring-LWE samples.

Ring conjugation, induced by $X \mapsto X^{-1}$, is denoted by a bar. The exact covariance is public; in this overview we write $E \in R_q$ for its reduction modulo q, which is all the attack uses. The decisional problem compares this genuine transcript with a random one in which $\mathbf{b}$ is replaced by an independent uniform vector in R_q^m while the same public covariance E is retained.

The argument has three layers. First, the covariance hint makes a single residual statistic public. Second, reducing this residual to the finite-field components of R_q leaves exactly two possible local laws, depending on how conjugation acts on the corresponding irreducible factor. Finally, we control the exceptional components and amplify the resulting local probability gap across all usable factors. The six steps below follow this dependency chain; the exact finite bounds are proved later.

Step 1: exact covariance exposes one public residual. For a column vector $\mathbf{z}$, write $\mathbf{z}^\star = \bar{\mathbf{z}}^T$. A genuine KCov-RLWE sample satisfies $\mathbf{b} = \mathbf{a}s + \mathbf{e}$ and $E = \bar{s}s + \mathbf{e}^\star\mathbf{e}$. Let I_m denote the $m \times m$ identity matrix and consider the public $(m+1) \times (m+1)$ matrix

$$\Gamma = \begin{pmatrix} I_m + \mathbf{a}\mathbf{a}^\star & \mathbf{b} \\ \mathbf{b}^\star & E \end{pmatrix}. \tag{1}$$

For a genuine sample it admits the hidden factorization

$$\Gamma = \begin{pmatrix} I_m & \mathbf{a} \\ \mathbf{e}^\star & \bar{s} \end{pmatrix} \begin{pmatrix} I_m & \mathbf{e} \\ \mathbf{a}^\star & s \end{pmatrix}. \tag{2}$$

Define the hidden residual $v = s - \mathbf{a}^\star\mathbf{e}$. At ring rank one the hidden factor is square, and taking determinants gives

$$W := \det \Gamma = \bar{v}v. \tag{3}$$

Thus the exact covariance turns a hidden quadratic relation into a public scalar statistic W. The identity is purely algebraic: it does not require Gaussianity, complete splitting, or an inverse. The square factorization is also specific to ring rank one; for higher module rank the analogous hidden factor is rectangular, so Cauchy–Binet produces a sum of minors rather than the single residual above. Section 3 makes this point precise.

Equation (3) is the common starting point for the whole attack, but by itself it does not yet distinguish real samples from random ones. To obtain a probability gap, we next inspect the same public residual in the finite-field components of R_q.

Step 2: factor reduction leaves exactly two conjugation patterns. Let P be an irreducible factor of $X^n + 1$ over $\mathbb{F}_q$. All such factors have the same degree $f = \mathrm{ord}_{2n}(q)$, where $\mathrm{ord}_{2n}(q)$ denotes the multiplicative order of q modulo $2n$. Reduction modulo P gives the finite field $K_P = \mathbb{F}_q[X]/(P) \simeq \mathbb{F}_{q^f}$. For any $z \in R_q$, write z_P for its image in K_P. In particular, W_P and E_P denote the projections of the public residual and covariance.

Let $P^\vee$ denote the reciprocal factor, whose roots are the inverses of the roots of P. Ring conjugation acts on irreducible factors in only two ways: it either exchanges P with a distinct factor $P^\vee$, or it preserves P. For the power-of-two cyclotomics considered here, the second possibility occurs exactly when $q \equiv -1 \pmod{2n}$, in which case the factors are self-reciprocal quadratics.

These two orbit types are the only local environments that the residual from Step 1 can encounter. Figure 1 summarizes the branching point: a two-factor orbit produces a product channel, whereas a fixed factor produces a Hermitian norm channel. Steps 3 and 4 analyze the corresponding local probability gaps. No third prime-modulus case remains to be handled.

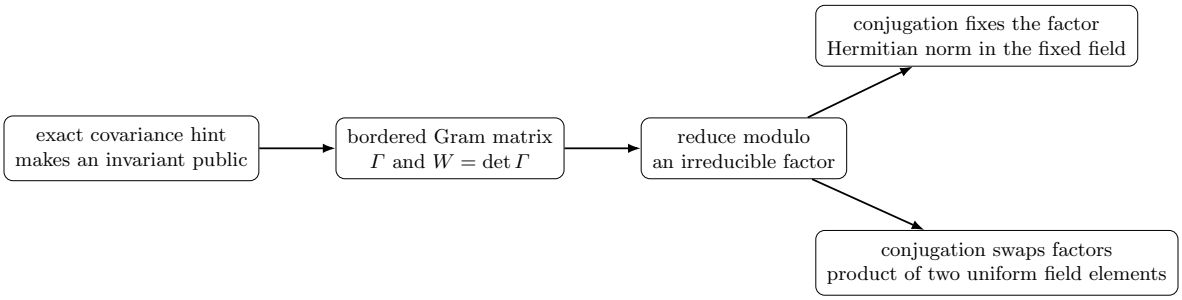


Fig. 1. Attack pipeline. Exact covariance makes one bordered-Gram residual public. Reduction modulo an irreducible factor then leaves two exhaustive conjugation patterns, producing either a product channel or a Hermitian norm channel.

Step 3: reciprocal factor pairs give a product channel. Suppose conjugation exchanges two distinct reciprocal factors P and $P^\vee$. Write $K = \mathbb{F}_{q^f}$ and $Q = q^f$ for the common residue field and its cardinality. Call the pair *useful* when the error projection to neither side of the reciprocal pair vanishes. Conditional on this event, the two public Chinese-remainder (CRT) coordinates supplied by the uniformly sampled Ring-LWE multiplier make the two affine factors of the projected residual independent and uniform in K. Consequently the projected residual has the simple law

$$W_P = UV, \qquad U, V \xleftarrow{\$} K \quad \text{independently}, \tag{4}$$

where $\xleftarrow{\$}$ denotes uniform sampling. This uniformity comes from the public multiplier; it does not require the short secret or errors themselves to become uniform after CRT projection.

For a genuine sample, $\Pr[W_P = 0] = 2/Q - 1/Q^2$. For a random KCov transcript, the same public statistic is a generic bilinear expression and has zero probability $Q^{-1} + O(Q^{-2})$, except for the covariance-zero one-sample branch isolated in Step 5. After the public gate used there, the same Q^{-1} scale is restored. Thus one useful reciprocal pair gives a local distinguishing gap of order $Q^{-1} = q^{-f}$. This is only a local statement; Step 6 will account for how many such pairs are available and amplify the gap. Section 4 derives the exact real and random laws, including the degenerate projected-error cases.

Step 4: fixed factors give a Hermitian norm channel. Suppose instead that conjugation preserves P. By Step 2 this occurs precisely when $q \equiv -1 \pmod{2n}$. The residue field is then $K = \mathbb{F}_{q^2}$, and ring conjugation becomes the Frobenius involution $z \mapsto z^q$. The hidden residual v from Step 1 projects to $v_P \in K$, and the public statistic becomes

$$W_P = v_P^q v_P = \mathrm{N}_{K/\mathbb{F}_q}(v_P) \in \mathbb{F}_q, \tag{5}$$

where $\mathrm{N}_{K/\mathbb{F}_q}$ denotes the field norm from K to its fixed field $\mathbb{F}_q$.

Call a fixed factor *active* when its public covariance component satisfies $E_P \neq 0$. On an active factor, the real zero probability is at most q^{-2}. In the random branch, the corresponding Hermitian form takes values in the smaller field $\mathbb{F}_q$ and hits the relevant target with probability $q^{-1} + O(q^{-2})$. Thus one active fixed factor gives a local gap of order q^{-1}. The direction of the zero test is reversed here: a zero is evidence for the random branch. Again, this is a local statement; Step 6 performs the global amplification. The fixed-field effect also explains why the final theorem cannot be stated solely in terms of residue degree:

whether conjugation swaps or fixes a factor matters just as much. Section 5 gives the exact Hermitian counts.

Step 5: control degenerate components and the one-sample case. The local laws above were stated on components for which the corresponding test is informative. The two branches reach this condition in slightly different ways. In the reciprocal branch, a pair is useful when both projected error vectors are nonzero; anti-concentration bounds ensure that this fails only rarely. When $m = 1$, there is an additional degeneracy: on components with $E_P = 0$, the reciprocal zero test can lose its real/random separation. Since E_P is public, the distinguisher simply discards those components and tests only pairs with $E_P \neq 0$.

In the Hermitian branch we likewise use the public condition $E_P \neq 0$. This gate removes the degenerate real components while preserving the q^{-1}-scale random zero probability from Step 4. Thus the same observable condition provides a clean active set in the fixed-factor case.

What remains is to show that enough useful or active components survive these restrictions. We express this through anti-concentration conditions on the relevant projections of the secret and errors. These conditions are distribution-generic; coefficient-product discrete Gaussians satisfy them by elementary one-coordinate conditioning arguments. This step is especially important for the motivating application because Micciancio and Suhl explicitly identify one-sample KCov-RLWE as sufficient for their Ring-PKE sketch [8]. The exact anti-concentration parameters and finite bounds are introduced only where they are needed in Sections 4 and 5.

At this point the proof roadmap is complete locally: Steps 3 and 4 provide a probability gap on each informative component, while Step 5 guarantees that sufficiently many such components remain. Only amplification is left.

Step 6: amplify the local gaps and assemble the theorem. The projections of a short secret or error polynomial may be strongly correlated across factors, so amplification cannot assume independent hidden CRT coordinates. We instead condition on the entire hidden tuple $(s, \mathbf{e})$. Conditioning fixes which factors are useful or active. Given this hidden state, the remaining randomness in the local tests comes from CRT coordinates of uniformly sampled public ring elements, and those coordinates are independent across distinct factors. We can therefore amplify conditionally and then average over the hidden state; a concavity argument converts the expected number of informative components into the required global lower bound.

Writing Adv for the absolute difference between the acceptance probabilities in the real and random experiments, the counting now makes the two final scales transparent. In the reciprocal branch there are $n/(2f)$ factor pairs. Under the anti-concentration conditions from Step 5, all but the corresponding small exceptional fraction are useful, and each useful pair has a local gap of order q^{-f}. The exact amplification argument gives

$$\mathrm{Adv} = \Omega\left(\frac{n}{fq^f}\right). \tag{6}$$

In the self-reciprocal Hermitian branch there are $n/2$ factors. Step 5 controls the fraction for which the public gate is inactive, while every active factor has a local gap of order q^{-1}; the corresponding argument gives

$$\mathrm{Adv} = \Omega(n/q). \tag{7}$$

Complete splitting is the special case $f = 1$ of the first branch. Since Step 2 shows that every odd prime modulus belongs to exactly one of the two conjugation patterns, Steps 3–5 establish the local hypotheses in both cases and this final amplification step yields the unified prime-modulus classification. Section 6.3

then observes that reduction modulo a divisor maps the KCov real and random experiments exactly to the corresponding lower-modulus experiments. Consequently, one attackable odd prime divisor is enough to transfer the relevant prime-modulus distinguisher to a composite modulus without loss of advantage.

The prime-modulus classification is exhaustive, but the guaranteed advantage need not always be non-negligible. In the reciprocal branch, families for which q^f grows superpolynomially remain outside the reach of the present distinguisher; pure power-of-two moduli are also outside the divisor-reduction extension. These are limitations of the attack, not evidence that the remaining families are secure.

1.5 Scope and limitations

Our claims are deliberately assumption-level. We show that the decisional KCov-RLWE game is distinguishable in the parameter regimes quantified by our theorems, including composite moduli that admit reduction to one of the vulnerable odd-prime regimes. We do not prove that all KCov-RLWE modulus families are insecure, that every parameter choice permitted by the motivating Ring-PKE sketch is vulnerable, or that the complete construction is IND-CPA insecure; nor do we establish search recovery. In particular, a reciprocal prime factor for which q^f is superpolynomial can make our guaranteed advantage negligible, while pure powers of two are not covered by the present CRT analysis. These are limitations of the attack, not evidence that such families are secure.

1.6 Organization

Section 2 fixes the KCov-RLWE game and the cyclotomic factorization notation. Section 3 proves the bordered-Gram identity. Sections 4 and 5 analyze the reciprocal and Hermitian residual channels. Section 6 combines them into the prime-modulus classification and shows how any vulnerable odd prime divisor transfers the attack to a composite modulus. Section 7 gives the assumption- and construction-level consequences. Section 8 positions the result against leakage-hardness, leftover-hash, threshold-encryption, and algebraic-RLWE literature, and Section 9 concludes.

2 Preliminaries

Status and provenance of this section. This section fixes notation and recalls background used by the attack. We claim no novelty for the statements in this section. The Ring-LWE setting and its cyclotomic notation are standard [7]; the Known-Covariance Ring-LWE experiment is due to Micciancio and Suhl [8]; and the finite-field factorization and reciprocal-factor facts recalled below are standard consequences of the theory of cyclotomic polynomials over finite fields [6, 12]. We nevertheless state the exact forms needed later and give a self-contained proof of Lemma 1 in Appendix A. The attack-specific algebraic mechanism developed in this paper begins in Section 3.

2.1 Ring notation and adjoints

Let $R = \mathbb{Z}[X]/(X^n + 1)$ with $n = 2^k \geq 2$, and for an integer modulus $q \geq 2$ let $R_q = R/qR \simeq (\mathbb{Z}/q\mathbb{Z})[X]/(X^n + 1)$. Whenever an element sampled in R occurs in an expression over R_q, we implicitly apply the canonical reduction map $R \to R_q$. We write $\xleftarrow{\$}$ for uniform sampling from a finite set. This is the standard power-of-two cyclotomic Ring-LWE setting; see, e.g., Lyubashevsky, Peikert, and Regev [7] for general cyclotomic Ring-LWE background.

The standard involution on R and R_q is $\bar{a}(X) = a(X^{-1}) \pmod{X^n+1}$. It is well defined because X is a unit and $X^{-1} = -X^{n-1}$ in R, and it satisfies $\overline{ab} = \bar{a}\,\bar{b}$, $\overline{a+b} = \bar{a} + \bar{b}$, and $\bar{\bar{a}} = a$.

Vectors are columns unless stated otherwise. For $\mathbf{z} = (z_1, \ldots, z_m)^T \in R_q^m$, define $\mathbf{z}^\star = \bar{\mathbf{z}}^T = (\bar{z}_1, \ldots, \bar{z}_m)$, and for a matrix M over R_q set $M^\star = \bar{M}^T$. Hence $\mathbf{x}^\star\mathbf{y} = \sum_{i=1}^m \bar{x}_i y_i \in R_q$ and $\overline{\mathbf{x}^\star\mathbf{y}} = \mathbf{y}^\star\mathbf{x}$. The $\star$ notation is a convention of this paper; no additional algebraic assumption is introduced by it.

2.2 The decisional KCov-RLWE game

We restate the rank-one Known-Covariance Ring-LWE experiment of Micciancio and Suhl [8, Definition 12] in the notation used throughout this paper.

Let χ be any distribution supported on R, and let $m \geq 1$ denote the number of Ring-LWE samples. We write χ^m for the product distribution of m independent samples from χ.

Definition 1 (Decisional KCov-RLWE (Micciancio–Suhl, restated)). *For parameters (R, q, m, χ), define the following two games. Corresponding sampling steps in the two columns use independent randomness from the same distributions.*

$\mathbf{Real}_{R,q,m,\chi}$	$\mathbf{Rand}_{R,q,m,\chi}$
1. $s \leftarrow \chi$, $\mathbf{e} \leftarrow \chi^m$, $\mathbf{a} \xleftarrow{\$} R_q^m$.	*1.* $s \leftarrow \chi$, $\mathbf{e} \leftarrow \chi^m$, $\mathbf{a} \xleftarrow{\$} R_q^m$.
2. Set $\mathbf{b} \leftarrow \mathbf{a}s + \mathbf{e}$.	*2. Sample* $\mathbf{u} \xleftarrow{\$} R_q^m$.
3. Set $E^{\mathrm{int}} \leftarrow s\bar{s} + \sum_{i=1}^m e_i\bar{e}_i$.	*3. Set* $E^{\mathrm{int}} \leftarrow s\bar{s} + \sum_{i=1}^m e_i\bar{e}_i$.
4. Return $(\mathbf{a}, \mathbf{b}, E^{\mathrm{int}})$.	*4. Return* $(\mathbf{a}, \mathbf{u}, E^{\mathrm{int}})$.

For a distinguisher $\mathcal{D}$, define

$$\mathrm{Adv}_{\mathsf{KCov\text{-}RLWE}}(\mathcal{D}) := \left|\Pr[\mathcal{D}(\mathsf{Real}_{R,q,m,\chi}) = 1] - \Pr[\mathcal{D}(\mathsf{Rand}_{R,q,m,\chi}) = 1]\right|. \tag{8}$$

We use this convention without the usual factor $1/2$.

The two games expose exactly the same genuine covariance. The only difference is that the Ring-LWE vector $\mathbf{b} = \mathbf{a}s + \mathbf{e}$ in the real game is replaced by an independent uniform vector $\mathbf{u}$ in the random game. This shared-hint coupling is part of the Micciancio–Suhl experiment, rather than a feature introduced by our attack.

2.3 Prime-modulus factor notation

The factor-based attacks require q to be an odd prime. In this subsection and in Sections 4–6, set $N = 2n$ and let $f = \mathrm{ord}_N(q)$ denote the multiplicative order of q modulo N, i.e., the smallest positive integer f such that $q^f \equiv 1 \pmod{N}$. Because n is a power of two, $X^n + 1 = \Phi_N(X)$, and $q \nmid N$, so the reduction of $X^n + 1$ modulo q is square-free.

The degree and number of irreducible factors below are the standard cyclotomic-factorization theorem over finite fields; see Lidl and Niederreiter [6, Theorem 2.47]. The reciprocal/self-reciprocal dichotomy is likewise standard; an explicit formulation in terms of the condition $N \mid (q^j + 1)$ appears, for example, in Sangwisut–Jitman–Ling–Udomkavanich [12, Lemma 3]. We package these facts into one lemma only to fix the exact notation used by the attack.

Lemma 1 (Standard cyclotomic factorization facts). *The polynomial $X^n + 1$ factors over $\mathbb{F}_q$ into exactly n/f distinct monic irreducible factors, each of degree f. For such a factor P, define its monic reciprocal by $P^\vee(X) = P(0)^{-1}X^f P(X^{-1})$. Then $P^\vee$ is again an irreducible factor of $X^n + 1$, and ring conjugation maps the CRT component indexed by P to the component indexed by $P^\vee$. Moreover, $P = P^\vee$ if and only if $q^j \equiv -1 \pmod N$ for some $0 \le j < f$.*

Proof overview. The irreducible factors are the Frobenius orbits of primitive N-th roots of unity over $\mathbb{F}_q$. Every such orbit has size $f = \mathrm{ord}_N(q)$, which gives the common factor degree and the number of factors. Taking reciprocals inverts the roots, while ring conjugation is induced by the same inversion operation. A factor is therefore self-reciprocal exactly when inversion keeps its Frobenius orbit fixed, which is equivalent to the congruence above. These are standard finite-field facts [6, 12]; a complete proof specialized to our notation is given in Appendix A.

For an irreducible factor P, write $K_P = \mathbb{F}_q[X]/(P)$ and $\pi_P : R_q \to K_P$ for the corresponding CRT projection. Lemma 1 implies $|K_P| = q^f$. Conjugation therefore acts on the irreducible factors in orbits of size one or two: either $P = P^\vee$ or $P \neq P^\vee$ and the two factors are paired. Section 3.3 explains how this standard dichotomy determines the two local forms of our public residual.

3 From Exact Covariance to a Public Residual

Definition 1 reveals the exact covariance $E^{\mathrm{int}} \in R$, but every construction in this section uses it only through the public residue

$$E := E^{\mathrm{int}} \bmod qR \in R_q.$$

Consequently, all distinguishers below remain valid in the weaker leakage experiment that reveals only E rather than the integral covariance.

Mechanism at a glance. All of our distinguishers share the same front end. The exact covariance first completes a public Gram matrix; at ring rank one, its determinant then collapses on a genuine transcript to the norm-like value of a single hidden residual:

$$(\mathbf{a}, \mathbf{b}, E) \;\longrightarrow\; \Gamma \;\longrightarrow\; W = \det \Gamma \overset{\mathsf{Real}}{=} \bar{v}v, \qquad v = s - \mathbf{a}^\star \mathbf{e}. \tag{9}$$

Equation (9) is a roadmap, not an additional assumption. The distinguisher never computes the hidden residual v: it computes W entirely from public data, while $W = \bar{v}v$ describes only the real-branch law.

How to read the roadmap. Sections 3.1–3.2 establish and compute the common public residual; Section 3.3 localizes it through the CRT dichotomy of Lemma 1, leading respectively to Sections 4 and 5. Section 3.4 then records why the determinant collapse is specific to ring rank one. We begin with the Gram completion.

3.1 Exact covariance completes a public Gram matrix

The first step explains precisely what the covariance hint contributes. Without E, the transcript contains the linear relation $\mathbf{b} = \mathbf{a}s + \mathbf{e}$ but not the quadratic quantity needed to close the bottom-right entry of a Gram matrix. The disclosed covariance supplies exactly that term.

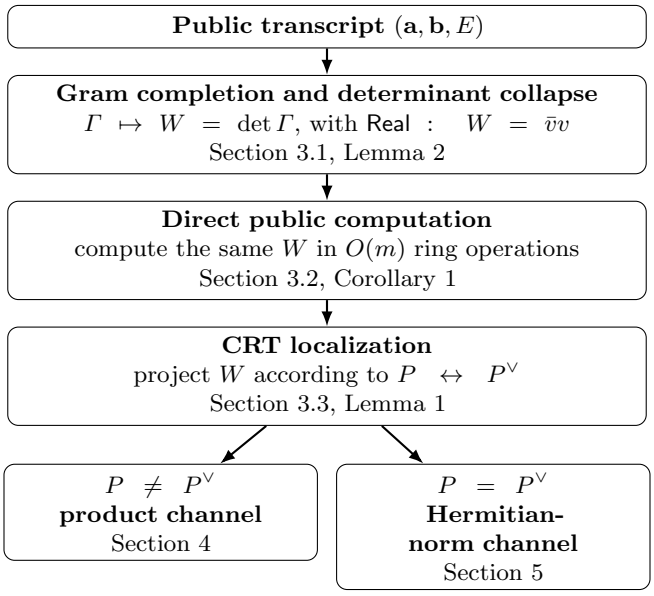


Fig. 2. Reader's map for Section 3. Each stage points to the subsection and formal result that establishes it. After the common front end, Sections 4 and 5 quantify the two local residual laws.

Lemma 2 (Bordered-Gram identity). *Let S be a commutative ring equipped with an involution $x \mapsto \bar{x}$. Let $s \in S$ and $\mathbf{a}, \mathbf{e} \in S^m$, and set $\mathbf{b} = \mathbf{a}s + \mathbf{e}$ and $E = \bar{s}s + \mathbf{e}^\star\mathbf{e}$. Define the public bordered matrix*

$$\Gamma(\mathbf{a}, \mathbf{b}, E) = \begin{pmatrix} I_m + \mathbf{a}\mathbf{a}^\star & \mathbf{b} \\ \mathbf{b}^\star & E \end{pmatrix}. \tag{10}$$

Then, with $v = s - \mathbf{a}^\star\mathbf{e}$, one has the exact identity

$$\det \Gamma = \bar{v}v. \tag{11}$$

Interpretation. Lemma 2 is the structural core of the attack. The matrix Γ is public, whereas v depends on the hidden secret and errors. On a genuine transcript, however, the public determinant is forced to lie in the special norm-like form $\bar{v}v$. Thus the covariance hint does not reveal v itself; it reveals enough exact quadratic information to make a structured function of v publicly computable.

Proof overview. Introduce the hidden square matrix

$$L = \begin{pmatrix} I_m & \mathbf{a} \\ \mathbf{e}^\star & \bar{s} \end{pmatrix}.$$

The relations $\mathbf{b} = \mathbf{a}s + \mathbf{e}$ and $E = \bar{s}s + \mathbf{e}^\star\mathbf{e}$ give the exact factorization $\Gamma = LL^\star$. Because L is square at ring rank one, $\det \Gamma = \det(L)\,\overline{\det(L)}$. A determinant-one block column operation reduces L to a triangular matrix with final diagonal entry $\bar{s} - \mathbf{e}^\star\mathbf{a} = \bar{v}$, giving $\det \Gamma = \bar{v}v$. The complete algebraic proof is in Appendix B.

Applying the lemma with $S = R_q$ therefore defines the public residual $W := \det \Gamma \in R_q$. The argument uses no division and no invertibility assumption, so zero divisors in R_q cause no difficulty.

3.2 A directly computable public residual

The determinant identity explains the structure, but an implementation need not compute an $(m+1)\times(m+1)$ determinant. The rank-one upper-left block gives a closed expression directly in the public transcript.

Corollary 1 (Public residual formula). *Let $h = 1 + \mathbf{a}^\star\mathbf{a} \in S$. Under the hypotheses of Lemma 2,*

$$W = hE - h\,\mathbf{b}^\star\mathbf{b} + (\mathbf{b}^\star\mathbf{a})(\mathbf{a}^\star\mathbf{b}). \tag{12}$$

Hence W is computable using $O(m)$ scalar products/ring operations and no matrix inversion. Moreover, $\bar{W} = W$.

Interpretation. Equation (12) is the form used by the distinguisher. It uses only $(\mathbf{a}, \mathbf{b}, E)$: no secret recovery, error recovery, linear solve, or invertibility test is required. The hidden expression $W = \bar{v}v$ is only the real-branch distributional explanation for this public statistic.

Proof overview. For $A = I_m + \mathbf{a}\mathbf{a}^\star$, the rank-one identities $\det(A) = h$ and $\mathrm{adj}(A) = hI_m - \mathbf{a}\mathbf{a}^\star$ turn the bordered determinant $\det \Gamma$ into exactly (12). Self-conjugacy follows from $W = \bar{v}v$. Full details are given in Appendix B.

Schur-complement interpretation. If h happens to be a unit, then

$$(I_m + \mathbf{a}\mathbf{a}^\star)^{-1} = I_m - \frac{\mathbf{a}\mathbf{a}^\star}{h},$$

and (12) can equivalently be written as

$$W = h\left(E - \mathbf{b}^\star(I_m + \mathbf{a}\mathbf{a}^\star)^{-1}\mathbf{b}\right). \tag{13}$$

This optional interpretation motivates the term *bordered-Gram residual*. The attack itself always uses the inverse-free formula (12), so no condition on h is imposed.

3.3 From the residual to two CRT attack channels

The global identity becomes useful after projection to the CRT factors from Section 2. Since $\bar{W} = W$, the behavior of one coordinate is controlled by how conjugation acts on its irreducible factor. There are exactly two cases.

Distinct reciprocal factors: product channel. If $P \neq P^\vee$, conjugation exchanges the two components indexed by P and $P^\vee$. After transporting the reciprocal coordinate into the same residue field, the real residual has the form

$$W_P = UV$$

for two coupled coordinates of the hidden residual. Section 4 shows that this multiplicative structure creates a detectable difference from the random branch and derives its exact finite-field advantage.

Self-reciprocal factor: Hermitian-norm channel. If $P = P^\vee$, conjugation acts inside a single residue field. The projected real residual becomes

$$W_P = v_P \overline{v_P},$$

a Hermitian norm into the fixed subfield of the induced involution. Section 5 shows that the real norm law has a q^{-2} zero atom while the corresponding random Hermitian form has a q^{-1}-scale zero atom.

By Lemma 1, these are the only two conjugation patterns for the CRT factors considered here. Sections 4 and 5 analyze the corresponding local laws of the same public residual W.

3.4 Why the collapse is specific to rank one

The essential distinction is dimensional. At ring rank one, the hidden factor L in the proof of Lemma 2 has $(m+1)$ rows and $(m+1)$ columns, so the Gram determinant is the norm of one determinant. At module rank $d > 1$, the analogous hidden factor is rectangular; the determinant of the public Gram matrix therefore expands into a sum of minor norms rather than one residual.

Proposition 1 (Rank-one specificity of the determinant collapse). *Consider the natural module-rank d analogue over a commutative involutive ring S. Let $A \in S^{m\times d}$, $\mathbf{s} \in S^d$, and $\mathbf{e} \in S^m$, and set $\mathbf{b} = A\mathbf{s} + \mathbf{e}$ and $E = \mathbf{s}^\star\mathbf{s} + \mathbf{e}^\star\mathbf{e}$. Then*

$$\begin{pmatrix} I_m + AA^\star & \mathbf{b} \\ \mathbf{b}^\star & E \end{pmatrix} = L_d L_d^\star, \qquad L_d = \begin{pmatrix} I_m & A \\ \mathbf{e}^\star & \mathbf{s}^\star \end{pmatrix}, \tag{14}$$

and Cauchy–Binet gives

$$\det(L_d L_d^\star) = \sum_{\substack{J\subseteq[m+d] \\ |J|=m+1}} \overline{\det((L_d)_J)}\, \det((L_d)_J). \tag{15}$$

For $d = 1$ the sum has a single term and reduces to Lemma 2; for $d > 1$ it contains multiple minors, so this factorization alone does not yield an analogous single affine residual.

Proof overview. Equation (14) follows by direct block multiplication. The matrix L_d has $(m+1)$ rows and $(m+d)$ columns, so applying Cauchy–Binet to $L_d L_d^\star$ produces the minor sum (15). Only when $d = 1$ is L_d square and only one minor survives. The complete proof is in Appendix B.

Proposition 1 is a structural limitation of this particular determinant mechanism, not a hardness statement for module rank greater than one and not an impossibility result for other attacks. The boundary aligns closely with the scope stated by Boudgoust, Jeudy, Tairi, and Wen (BJTW) [2]: their general-distribution hardness theorem requires module rank $d > 1$, and their Known-Covariance application therefore treats the module version rather than the ring case. Moreover, their resulting KCov guarantee is for the search problem, while extending their main framework to decision M-LWE is explicitly left open. Our attack lies precisely in this excluded $d = 1$ decisional regime. Importantly, BJTW already handle the exact quadratic hint and an HNF-shaped map $[A \mid I]$, so neither feature alone separates the two results; the decisive distinction here is the rank-one square Gram collapse exploited by the decision test.

Section takeaway. The common Gram mechanism is now complete; Sections 4 and 5 analyze the two CRT-local laws identified in Section 3.3.

4 The Reciprocal-Factor Product Attack

Section 3 produced the public residual W and showed that $W = \bar{v}v$ on genuine transcripts (Lemma 2). We now follow the $P \neq P^\vee$ branch of Section 3.3 and Figure 2. Projecting to such a reciprocal pair turns the global identity into a finite-field product whose real and random laws differ.

Attack at a glance. The logical path is

$$\boxed{\begin{array}{c} W = \bar{v}v \quad \text{(Lemma 2)} \\ \Downarrow\ P \neq P^\vee \quad \text{(Lemma 1)} \\ W_P = UV \quad \text{in the real branch} \quad \text{(Lemma 4)} \\ \Downarrow \\ \Pr[W_P = 0] = 2/Q - 1/Q^2 \\ \Downarrow \\ \text{compare with the random rank law} \quad \text{(Lemma 5)} \\ \Downarrow \\ \text{local distinguisher} \quad \text{(Theorem 1)} \\ \Downarrow \\ \text{amplify across reciprocal pairs} \quad \text{(Theorem 3).} \end{array}} \tag{16}$$

Section 4 therefore starts from the residual of Section 3, localizes it to one reciprocal pair, and turns its multiplicative structure into a statistical test. We follow the steps in (16) in order.

4.1 From the global residual to one reciprocal pair

Let

$$K_P = \mathbb{F}_q[X]/(P), \qquad K_{P^\vee} = \mathbb{F}_q[X]/(P^\vee).$$

Both fields have cardinality

$$Q = q^f, \qquad f = \mathrm{ord}_{2n}(q).$$

Write π_P and $\pi_{P^\vee}$ for the corresponding CRT projections. Conjugation induces the canonical field isomorphism

$$\tau_P : K_{P^\vee} \longrightarrow K_P, \qquad \tau_P(\pi_{P^\vee}(r)) = \pi_P(\bar{r}). \tag{17}$$

For readability, we transport the $P^\vee$ coordinate through τ_P and write, for $r \in R_q$,

$$r_P := \pi_P(r), \qquad r_{\bar{P}} := \tau_P(\pi_{P^\vee}(r)) \quad \in K_P.$$

With this convention, conjugation simply swaps the two coordinates: $(\bar{r})_P = r_{\bar{P}}$ and $(\bar{r})_{\bar{P}} = r_P$. The transport is only a notational convenience; it places the two exchanged components in one common copy $K \simeq \mathbb{F}_Q$ and adds no algebraic assumption.

Lemma 3 (Public CRT independence on a reciprocal pair). *Let $r \leftarrow R_q$ be uniform. Then $(r_P, r_{\bar{P}})$ is uniform on $K \times K$. Consequently, for a uniform public vector $\mathbf{r} \leftarrow R_q^m$, its two transported coordinate vectors are independent uniform elements of K^m.*

Interpretation. This is the only independence property used by the attack. We do *not* assume that a short secret or error has independent P and $P^\vee$ projections. Whenever independence is invoked below, it comes from a public uniform Ring-LWE coordinate after conditioning on the complete hidden state.

Proof overview. The CRT decomposition is a ring isomorphism onto the product of the distinct field factors. A uniform ring element therefore has independent uniform CRT coordinates, and applying the bijection τ_P preserves uniformity. The full argument is recorded in Appendix C.

4.2 Real branch: the residual becomes a product

We now project the real-branch identity from Lemma 2. For the public multiplier and the errors, write

$$\mathbf{x} = (a_{1,P}, \dots, a_{m,P}), \quad \mathbf{y} = (a_{1,\bar{P}}, \dots, a_{m,\bar{P}}),$$
$$\mathbf{u} = (e_{1,P}, \dots, e_{m,P}), \quad \mathbf{z} = (e_{1,\bar{P}}, \dots, e_{m,\bar{P}}),$$

and put $s_1 = s_P$, $s_2 = s_{\bar{P}}$. The two coordinates of the hidden residual $v = s - \mathbf{a}^\star \mathbf{e}$ are

$$v_P = s_1 - \mathbf{y}^T \mathbf{u}, \qquad v_{\bar{P}} = s_2 - \mathbf{x}^T \mathbf{z}.$$

Lemma 4 (Reciprocal product law). *On a genuine KCov-RLWE transcript,*

$$W_P = v_{\bar{P}} v_P = \big(s_2 - \mathbf{x}^T \mathbf{z}\big)\big(s_1 - \mathbf{y}^T \mathbf{u}\big). \tag{18}$$

If $\mathbf{u} \neq 0$ and $\mathbf{z} \neq 0$, then conditional on the complete hidden tuple $(s, \mathbf{e})$ the two factors in (18) are independent uniform elements of K. Consequently,

$$\Pr[W_P = 0 \mid s, \mathbf{e}, \mathbf{u} \neq 0, \mathbf{z} \neq 0] = \alpha_1 := \frac{2}{Q} - \frac{1}{Q^2}. \tag{19}$$

Why this is the signal. The distinguisher never learns either hidden factor in (18); it only computes their public product W_P. The useful feature is the zero atom: a product of two independent uniform field elements vanishes whenever either factor is zero, giving the probability $2/Q - 1/Q^2$, roughly twice the $1/Q$ scale of a generic field equation.

To account for degenerate error projections without imposing a Gaussian assumption, define for one error draw $e \leftarrow \chi$

$$\rho_P = \Pr[e_P = 0], \qquad \rho_{\bar{P}} = \Pr[e_{\bar{P}} = 0], \qquad \theta_P = \Pr[e_P = 0 \ \wedge \ e_{\bar{P}} = 0], \tag{20}$$

and set $\rho = \max\{\rho_P, \rho_{\bar{P}}\}$. Since the m error samples are independent, $\Pr[\mathbf{u} = 0] = \rho_P^m$ and $\Pr[\mathbf{z} = 0] = \rho_{\bar{P}}^m$. The exact lower bound used later is

$$\Pr[W_P = 0 \mid \mathsf{Real}] \geq \alpha_1 - \beta_1(\rho_P^m + \rho_{\bar{P}}^m) - \frac{\theta_P^m}{Q^2}, \qquad \beta_1 := \frac{1}{Q} - \frac{1}{Q^2}. \tag{21}$$

The derivation is elementary and is deferred to Appendix C; the hidden projections may remain correlated, as emphasized in Section 4.1.

4.3 Random branch: the product bias disappears

A product law alone is not yet a distinguisher. We must evaluate the *same public statistic* when the KCov game replaces $\mathbf{b}$ by an independent uniform vector. After the transport from Section 4.1, write

$$\mathbf{x}, \mathbf{y}, \mathbf{p}, \mathbf{q} \in K^m$$

for the two coordinates of $\mathbf{a}$ and the two coordinates of the random $\mathbf{b}$; these four vectors are mutually independent and uniform. Put

$$c = \mathbf{x}^T\mathbf{y}, \qquad \varepsilon = E_P.$$

The explicit public formula of Corollary 1 becomes

$$W_P = (1+c)\varepsilon - \mathbf{p}^T M_{\mathbf{x},\mathbf{y}}\mathbf{q}, \qquad M_{\mathbf{x},\mathbf{y}} = (1+c)I_m - \mathbf{y}\mathbf{x}^T. \tag{22}$$

Thus the random branch is controlled by the rank of one public matrix.

For later use, if $U, V \leftarrow K^d$ are independent uniform vectors, define

$$\alpha_d := \Pr[U^T V = 0] = \frac{1}{Q} + \frac{Q-1}{Q^{d+1}}, \tag{23}$$

$$\beta_d := \Pr[U^T V = z] = \frac{1}{Q} - \frac{1}{Q^{d+1}} \qquad (z \neq 0). \tag{24}$$

Lemma 5 (Random-branch rank collapse). *Assume $m \geq 2$. For the matrix in* (22)*,*

$$\det M_{\mathbf{x},\mathbf{y}} = (1+c)^{m-1}.$$

Hence

$$\operatorname{rank} M_{\mathbf{x},\mathbf{y}} = \begin{cases} m, & c \neq -1, \\ 1, & c = -1, \end{cases} \qquad \Pr[c = -1] = \beta_m.$$

Conditioned on the public covariance coordinate ε, the exact random zero probabilities are therefore

$$P_0^{\mathsf{Rand}} := \Pr[W_P = 0 \mid \varepsilon = 0] = \beta_m\alpha_1 + (1-\beta_m)\alpha_m, \tag{25}$$

$$P_1^{\mathsf{Rand}} := \Pr[W_P = 0 \mid \varepsilon \neq 0] = \beta_m\alpha_1 + (1-\beta_m)\beta_m. \tag{26}$$

Moreover $P_0^{\mathsf{Rand}} > P_1^{\mathsf{Rand}}$, so P_0^{Rand} is a pointwise upper bound for the random zero probability for every possible leakage value.

Interpretation. The real branch has rank one because exact covariance completes the hidden product. In the random branch the same statistic is generically a rank-m bilinear form. It collapses to rank one only on the explicit public-random event $c = -1$, whose probability is $\beta_m \approx 1/Q$. This *distributional rank collapse*, rather than membership in a special norm image, is the source of the distinguishing gap.

Subtracting the worst random probability from the nondegenerate real atom gives

$$\Delta_0 := \alpha_1 - P_0^{\mathsf{Rand}} = \frac{1}{Q}\left(1 - \frac{1}{Q} + \frac{1}{Q^{m+1}}\right)\left(1 - \frac{1}{Q}\right)\left(1 - \frac{1}{Q^{m-1}}\right). \tag{27}$$

For fixed $m \geq 2$, $\Delta_0 = Q^{-1}(1 + O(Q^{-1}))$.

Theorem 1 (Local reciprocal-factor distinguisher). *Assume $m \geq 2$ and an efficiently accessible reciprocal pair $P \neq P^{\vee}$. The polynomial-time zero test*

$$\mathcal{D}_P(\mathbf{a}, \mathbf{b}, E) = \mathbf{1}[W_P = 0]$$

has advantage at least

$$\mathrm{Adv}(\mathcal{D}_P) \geq \Delta_0 - \beta_1(\rho_P^m + \rho_{\bar{P}}^m) - \frac{\theta_P^m}{Q^2}. \tag{28}$$

In particular,

$$\mathrm{Adv}(\mathcal{D}_P) \geq \Delta_0 - \alpha_1 \rho^m. \tag{29}$$

Thus, if $\rho^m \leq 1/2 - \delta$ for a constant $\delta > 0$, then for sufficiently large Q,

$$\mathrm{Adv}(\mathcal{D}_P) = \Omega_\delta(Q^{-1});$$

if $\rho^m = o(1)$, the advantage is $(1 - o(1))/Q$.

Proof overview. Equation (21) lower-bounds the real zero probability pointwise in the hidden data. Lemma 5 upper-bounds the random zero probability by P_0^{Rand} for every covariance value, so no independence between E and the hidden errors is needed. Subtracting gives (28); the simplified form follows from $2\beta_1 + Q^{-2} = \alpha_1$. Appendix C contains the complete finite-field calculation.

4.4 One sample: recovering the signal with a public gate

The rank argument above genuinely uses $m \geq 2$. For $m = 1$ the matrix $M_{x,y}$ in (22) is always rank one, so the leakage-independent rank gap disappears. The product signal nevertheless survives after removing one publicly visible degeneracy.

Write $x = a_P$, $y = a_{\bar{P}}$, $w = b_P$, $z = b_{\bar{P}}$ and $\varepsilon = E_P$. The public residual is

$$W_P = (1 + xy)\varepsilon - wz. \tag{30}$$

On a genuine sample it also factors as

$$W_P = (s_1 - yu)(s_2 - xv), \tag{31}$$

where $u = e_P$ and $v = e_{\bar{P}}$. If $\varepsilon = 0$, the random branch itself becomes a zero-product test and acquires the same enlarged zero mass as the real branch. Since ε is public, the distinguisher simply discards this branch:

$$\mathcal{D}_P^{(1)} = \mathbf{1}[\varepsilon \neq 0 \ \wedge \ W_P = 0]. \tag{32}$$

Conditional on fixed hidden data with $u, v, \varepsilon \neq 0$, the real acceptance probability is

$$\alpha_1 = \frac{2}{Q} - \frac{1}{Q^2},$$

whereas in the random experiment

$$\Pr[W_P = 0 \mid \varepsilon \neq 0, \mathsf{Rand}] = \gamma_1 := \frac{1}{Q} - \frac{1}{Q^3}. \tag{33}$$

The raw gated gap is therefore

$$\delta_1 := \alpha_1 - \gamma_1 = \frac{1}{Q} - \frac{1}{Q^2} + \frac{1}{Q^3} = Q^{-1}(1 + O(Q^{-1})). \tag{34}$$

Why the gate is essential. For $\varepsilon \neq 0$, the random equation asks an unrelated product wz to hit the nonzero target $(1+xy)\varepsilon$, producing the $1/Q$-scale law in (33). At $\varepsilon = 0$ that target vanishes and the random branch reverts to a zero-product law. The gate removes exactly this publicly recognizable branch and nothing else.

For the motivating fully split family, the gate and the two error-zero events can be controlled explicitly for coefficient-product discrete Gaussians.

Theorem 2 (One-sample fully split distinguisher). *Let $n = 2^r \geq 2$, let q be an odd prime with $q \equiv 1 \pmod{2n}$, and let $s, e \leftarrow D_{R,\sigma}$ independently with independent coefficient-product discrete-Gaussian coordinates. For any reciprocal linear pair, the test* (32) *has signed real-minus-random gap at least*

$$\begin{aligned} \left(q^{-1} - q^{-2} + q^{-3}\right)(1 - 2h_\sigma) - \left(2q^{-1} - q^{-2}\right) h_\sigma, \\ h_\sigma := \min\{1, \sigma^{-1} + q^{-1}\}. \end{aligned} \tag{35}$$

Its gap is at most $q^{-1} - q^{-2} + q^{-3}$. Consequently, if $q = \mathrm{poly}(n)$ with $q \to \infty$ and $\sigma \to \infty$, then

$$\mathrm{Adv}(\mathcal{D}_P^{(1)}) = q^{-1}(1 + o(1)),$$

and in particular the one-sample setting remains distinguishable with non-negligible advantage.

Proof overview. The finite-field contribution is exactly $\delta_1 = \alpha_1 - \gamma_1$. The remaining loss comes from the gated and zero-projection exceptional events; a one-coefficient anti-concentration argument bounds both by the Gaussian residue atom h_σ. Appendix C gives the complete calculation.

4.5 From a local gap to the global reciprocal distinguisher

We now return from one CRT pair to the original ring transcript. When no factor is self-reciprocal, the n/f irreducible factors form

$$L = \frac{n}{2f} \tag{36}$$

reciprocal pairs. The short-error projections across these pairs may be strongly correlated; the amplification argument therefore never treats them as independent. Instead, we condition on the complete hidden tuple $(s, \mathbf{e})$. Under this conditioning the hidden degeneracy pattern is fixed, while the public CRT coordinates of $\mathbf{a}$ (and of the random $\mathbf{b}$) remain independent across distinct factors by Lemma 3.

The following theorem makes the amplification explicit. The distinguisher randomly selects a bounded number of reciprocal pairs and outputs "real" if any selected local residual is zero. Random subsampling keeps the accumulated zero probability in the linear regime and avoids any assumption on how the hidden bad pairs are arranged.

Theorem 3 (Global reciprocal-factor distinguisher). *Assume $m \geq 2$, that all factors occur in distinct reciprocal pairs, and define*

$$\rho := \max_P \Pr_{e \leftarrow \chi}[e_P = 0].$$

Suppose $\rho^m \leq 1/2 - \delta$ for some constant $\delta > 0$. Put

$$Q = q^f, \qquad L = \frac{n}{2f}, \qquad J = \min\{L, \lfloor \delta Q/16 \rfloor\}.$$

Then the polynomial-time subsampled OR distinguisher using only the public residual coordinates has signed real-minus-random gap at least

$$\begin{aligned}\mathrm{Adv} \geq (1-\rho^m)\big(1-(1-Q^{-1})^{2J}\big) \\ -\big(1-(1-P_0^{\mathsf{Rand}})^J\big),\end{aligned} \tag{37}$$

where P_0^{Rand} *is the exact quantity in* (25)*. Moreover, for* $Q \geq 16/\delta$*,*

$$\mathrm{Adv} \geq \delta\,\frac{J}{Q}. \tag{38}$$

Consequently,

$$\mathrm{Adv} = \Omega_\delta\left(\frac{n}{fq^f}\right). \tag{39}$$

The hidden CRT projections need not be independent.

Proof overview. Condition on all hidden errors. If a selected pair has two nonzero error vectors, its real zero probability is $1-(1-1/Q)^2$; if exactly one side is nonzero, it is at least $1/Q$; and if both vanish we use the conservative lower bound zero. Thus, if T is the number of nonzero error sides among the selected pairs, the conditional real acceptance probability is at least $1-(1-1/Q)^T$. Concavity and uniform subsampling give the first term of (37) because $\mathbb{E}[T] \geq 2J(1-\rho^m)$.

In the random branch, Lemma 5 bounds every selected local zero probability by P_0^{Rand}, and conditional CRT independence gives the second term of (37). The finite-field formulas also give the uniform bound $P_0^{\mathsf{Rand}} \leq Q^{-1}+2Q^{-2}$ for every $m \geq 2$. With $x = J/Q \leq \delta/16$, the inequalities $1-(1-Q^{-1})^{2J} \geq 2x-2x^2$ and $\rho^m \leq 1/2-\delta$ yield (38) for $Q \geq 16/\delta$. The two possible values of J then imply (39). Appendix C contains the complete finite derivation and constants.

Section takeaway. Distinct reciprocal factors turn the public residual into a product channel, yielding the global bound of Theorem 3 for $m \geq 2$; the public gate additionally handles the fully split one-sample case of Theorem 2. Section 5 now follows the other branch of Figure 2, where the same residual becomes a Hermitian norm.

5 The Self-Reciprocal Hermitian Attack

Section 4 followed the branch in which conjugation exchanges two CRT factors. We now follow the remaining branch of Figure 2, where an irreducible factor is fixed by conjugation. The public residual W is unchanged; only its local law changes. This change reverses the direction of the zero test and gives a q^{-1}-scale signal even though the residue field has q^2 elements.

Attack at a glance. The argument is the fixed-factor analogue of Section 4:

$$\mathsf{Real}: \quad W = \bar{v}v \qquad \text{(Lemma 2)}$$

$$P = P^\vee \quad \Longrightarrow \quad q \equiv -1 \pmod{2n}, \qquad K_P \simeq \mathbb{F}_{q^2}, \qquad \bar{z} = z^q$$

$$\Downarrow$$

$$\mathsf{Real}: \quad W_P = v_P^{q+1} \in \mathbb{F}_q$$

$$\text{on an active factor } E_P \neq 0:$$

$$\Pr_{\mathsf{Real}}[W_P = 0 \mid E_P \neq 0] \leq q^{-2}, \qquad \Pr_{\mathsf{Rand}}[W_P = 0 \mid E_P \neq 0] = q^{-1} + O(q^{-2})$$

$$\Downarrow$$

$$\text{zero favors } \mathsf{Rand}, \qquad \text{local gap } = \Theta(q^{-1})$$

$$\text{constant active fraction} \quad \Longrightarrow \quad \mathrm{Adv} = \Omega(n/q).$$

The final line is conditional on a constant fraction of factors being active; Theorem 5 states the exact activity hypothesis and finite bound. Sections 5.1–5.5 establish the remaining steps in order.

5.1 From a fixed CRT factor to $\mathbb{F}_{q^2}/\mathbb{F}_q$

We first specialize the factor criterion of Lemma 1. For the power-of-two cyclotomics used here, a factor can be fixed by inversion only in one prime-modulus class.

Corollary 2 (The fixed-factor regime). *Let $N = 2n$ with n a power of two, let q be an odd prime, and let P be an irreducible factor of $X^n + 1$ over $\mathbb{F}_q$. Then*

$$P = P^\vee \qquad \Longleftrightarrow \qquad q \equiv -1 \pmod{2n}.$$

In this case $f = \mathrm{ord}_{2n}(q) = 2$, every irreducible factor is self-reciprocal of degree two, there are $n/2$ such factors, and for $K_P \simeq \mathbb{F}_{q^2}$ the induced involution is

$$\bar{z} = z^q. \tag{40}$$

Proof overview. Lemma 1 says that self reciprocity is equivalent to $q^j \equiv -1 \pmod{2n}$ for some j. Modulo 4, this forces j to be odd; for odd j, the quotient $(q^j + 1)/(q + 1)$ is odd. Since $2n$ is a power of two, divisibility of $q^j + 1$ by $2n$ therefore already forces $q \equiv -1 \pmod{2n}$. The converse is immediate. The remaining statements follow from $f = 2$ and from the unique nontrivial automorphism of $\mathbb{F}_{q^2}/\mathbb{F}_q$. Appendix D gives the complete proof.

Thus Section 5 concerns a single, explicit arithmetic regime. In particular, unlike the reciprocal branch, there is no residue-degree parameter left to vary: the local field is always the quadratic extension $\mathbb{F}_{q^2}/\mathbb{F}_q$.

5.2 Real branch: the residual becomes a field norm

Fix one self-reciprocal factor P and write

$$K := K_P \simeq \mathbb{F}_{q^2}, \qquad \varepsilon := E_P \in \mathbb{F}_q.$$

The inclusion $\varepsilon \in \mathbb{F}_q$ follows from $\bar{E} = E$ and (40). Projecting the real-branch identity from Lemma 2 gives the promised norm law.

Lemma 6 (Hermitian norm law). *On a genuine KCov-RLWE transcript,*

$$W_P = \bar{v}_P v_P = v_P^{q+1} = \mathrm{N}_{K/\mathbb{F}_q}(v_P), \qquad v_P = s_P - \mathbf{a}_P^\star \mathbf{e}_P. \tag{41}$$

Call P active *when its public covariance coordinate satisfies $\varepsilon \neq 0$. Conditional on the complete hidden tuple $(s, \mathbf{e})$, every active factor satisfies*

$$\Pr[W_P = 0 \mid s, \mathbf{e}, \varepsilon \neq 0, \mathsf{Real}] \leq A_1 := q^{-2}. \tag{42}$$

If $\mathbf{e}_P \neq 0$, equality holds in (42); if $\mathbf{e}_P = 0$, the probability is zero.

Why the active gate works. A field norm vanishes only at zero. If $\mathbf{e}_P \neq 0$, the uniformly sampled public multiplier makes the affine residual v_P uniform in K, so it vanishes with probability $|K|^{-1} = q^{-2}$. If $\mathbf{e}_P = 0$, then $v_P = s_P$ and $\varepsilon = \mathrm{N}(s_P)$; activity therefore implies $s_P \neq 0$ and the real zero event is impossible. Thus the public condition $E_P \neq 0$ removes the degenerate real components without any independence or anti-concentration assumption on the hidden error projections. The full proof is in Appendix D.

5.3 Random branch: a Hermitian form has a larger zero atom

We now evaluate the same public statistic in the random KCov experiment. The necessary finite-field counts are most conveniently stated for the standard Hermitian form

$$H_d(\mathbf{x}) := \mathbf{x}^\star \mathbf{x} = \sum_{i=1}^{d} x_i^{q+1} \in \mathbb{F}_q, \qquad \mathbf{x} \in K^d.$$

For $d \geq 1$, define

$$A_d := \Pr_{\mathbf{x} \leftarrow K^d}[H_d(\mathbf{x}) = 0], \tag{43}$$

$$B_d := \Pr_{\mathbf{x} \leftarrow K^d}[H_d(\mathbf{x}) = t] \quad (t \in \mathbb{F}_q^*). \tag{44}$$

Lemma 7 (Exact Hermitian-form counts). *For every $d \geq 1$,*

$$A_d = \frac{1}{q} + (-1)^d \frac{q-1}{q^{d+1}}, \tag{45}$$

$$B_d = \frac{1}{q} - (-1)^d \frac{1}{q^{d+1}}. \tag{46}$$

In particular, $A_1 = q^{-2}$, and $B_d = q^{-1} + O(q^{-d-1})$.

Proof overview. For a nontrivial additive character ψ of $\mathbb{F}_q$, the one-coordinate norm sum is

$$\sum_{z \in K} \psi(t z^{q+1}) = \begin{cases} q^2, & t = 0, \\ -q, & t \neq 0. \end{cases}$$

The d-coordinate sum factors, and Fourier inversion over $\mathbb{F}_q$ yields (45)–(46). The complete calculation is in Appendix D.

In the random branch, write the projected public vectors again as $\mathbf{a}, \mathbf{b} \in K^m$ and set

$$h := 1 + \mathbf{a}^\star \mathbf{a} \in \mathbb{F}_q, \qquad M_{\mathbf{a}} := hI_m - \mathbf{a}\mathbf{a}^\star. \tag{47}$$

Corollary 1 becomes

$$W_P = h\varepsilon - \mathbf{b}^\star M_{\mathbf{a}} \mathbf{b}. \tag{48}$$

Lemma 8 (Random Hermitian residual law). *For uniform independent* $\mathbf{a}, \mathbf{b} \leftarrow K^m$*,*

$$\operatorname{rank}(M_{\mathbf{a}}) = \begin{cases} m, & h \neq 0, \\ 1, & h = 0, \end{cases} \qquad \Pr[h = 0] = B_m. \tag{49}$$

For every active covariance value $\varepsilon \neq 0$*,*

$$P^{\mathrm{H}}_{\mathsf{Rand}} := \Pr[W_P = 0 \mid \varepsilon \neq 0, \mathsf{Rand}] = B_m A_1 + (1 - B_m) B_m. \tag{50}$$

Consequently,

$$\Delta^{\mathrm{H}}_m := P^{\mathrm{H}}_{\mathsf{Rand}} - A_1 = (1 - B_m)(B_m - A_1) = q^{-1}(1 + O(q^{-1})). \tag{51}$$

The bound is uniform for every $m \geq 1$*.*

Interpretation. The contrast with Section 4 is now explicit. In the real branch, zero requires one element of the q^2-element field to vanish, giving at most q^{-2}. In the random branch, the public statistic is a Hermitian form whose target lies in the q-element fixed field; its zero probability is therefore on the q^{-1} scale. Hence a zero is evidence for *Rand*, the opposite direction from the reciprocal product test. This is also why residue degree alone does not predict the attack strength: a self-reciprocal quadratic factor gives a q^{-1} signal rather than a q^{-2} one.

5.4 The active-factor distinguisher

Define the public test

$$\mathcal{D}^{\mathrm{H}}_P(\mathbf{a}, \mathbf{b}, E) := \mathbf{1}[E_P \neq 0 \ \wedge \ W_P = 0], \tag{52}$$

and interpret output 1 as evidence for the random experiment. Let

$$\kappa_P := \Pr[E_P = 0], \tag{53}$$

where the probability is over the hidden secret and errors, which are sampled identically in the two KCov branches.

Theorem 4 (Local self-reciprocal distinguisher). *For every* $m \geq 1$ *and every self-reciprocal factor* P*,*

$$\Pr[\mathcal{D}^{\mathrm{H}}_P = 1 \mid \mathsf{Rand}] - \Pr[\mathcal{D}^{\mathrm{H}}_P = 1 \mid \mathsf{Real}] \geq (1 - \kappa_P)\Delta^{\mathrm{H}}_m. \tag{54}$$

In particular, if $\kappa_P \leq 1 - \delta$ *for a constant* $\delta > 0$*, then*

$$\operatorname{Adv}(\mathcal{D}^{\mathrm{H}}_P) = \Omega_\delta(q^{-1}).$$

Unlike the rank-separation argument of Section 4, this statement already includes the one-sample case $m = 1$*.*

Proof overview. Condition on the hidden tuple. On every active factor, Lemma 6 upper-bounds the real zero probability by $A_1 = q^{-2}$, while Lemma 8 gives the exact random probability $P^{\mathrm{H}}_{\mathsf{Rand}}$. The gate is inactive with probability κ_P in both experiments. Subtracting gives (54). Appendix D gives the complete argument.

For coefficient-product discrete Gaussians, activity follows from the same one-coordinate anti-concentration principle used in Section 4.4.

Lemma 9 (Gaussian activity bound). *Let the coefficients of $s, e_1, \ldots, e_m$ be independent samples from $D_{\mathbb{Z},\sigma}$. For every self-reciprocal quadratic factor P,*

$$\kappa_P \le 2h_\sigma, \qquad h_\sigma := \min\{1, \sigma^{-1} + q^{-1}\}. \tag{55}$$

Proof overview. Condition on all hidden coefficients except the constant coefficient Z of the secret. Writing $s_P = Z + c$, the equation $E_P = 0$ becomes a monic quadratic in $Z \in \mathbb{F}_q$, hence has at most two residue-class solutions. Each residue class has discrete-Gaussian mass at most h_σ. Full details are in Appendix D.

5.5 From active factors to the global Hermitian attack

In the fixed-factor regime, Corollary 2 gives exactly

$$L = \frac{n}{2} \tag{56}$$

self-reciprocal factors. The active pattern may be arbitrarily correlated across these factors. As in Section 4.5, we condition on the complete hidden tuple. This fixes the active set, while the CRT coordinates of the public uniform vectors remain independent across distinct factors.

The global distinguisher computes W, keeps only active factors, and outputs "random" if $W_P = 0$ for at least one of them.

Theorem 5 (Global self-reciprocal Hermitian distinguisher). *Assume $q \equiv -1 \pmod{2n}$ and $m \ge 1$, and define*

$$\kappa := \max_P \Pr[E_P = 0]$$

over the $n/2$ self-reciprocal factors. Then the polynomial-time global test above satisfies the explicit lower bound

$$\mathrm{Adv} \ge \frac{n}{4}(1-\kappa)\Delta^{\mathrm{H}}_m. \tag{57}$$

Consequently, if $\kappa \le 1 - \delta$ for a constant $\delta > 0$, then

$$\mathrm{Adv} = \Omega_\delta(n/q). \tag{58}$$

No independence assumption is made on the hidden CRT projections.

Proof overview. Condition on all hidden values and let T be the number of active factors. In the random branch, each active local zero event has probability $P^{\mathrm{H}}_{\mathsf{Rand}}$; in the real branch its probability is at most A_1. Conditional CRT independence therefore makes the difference between the two OR tests at least

$$(1 - A_1)^T - (1 - P^{\mathrm{H}}_{\mathsf{Rand}})^T \geq \frac{T}{2}\Delta^{\mathrm{H}}_m.$$

The last inequality is fully finite: the exact Hermitian counts give $P^{\mathrm{H}}_{\mathsf{Rand}} \leq q^{-1} + 2q^{-2}$, while $q \equiv -1 \pmod{2n}$ implies $n/2 \leq (q+1)/4$; hence $(1 - P^{\mathrm{H}}_{\mathsf{Rand}})^{T-1} \geq 1/2$ (with $q = 3$ handled separately). Averaging gives $\mathbb{E}[T] \geq (n/2)(1-\kappa)$ and hence (57). Thus no asymptotic estimate is used in the amplification itself. Appendix D gives the complete finite inequality, including the smallest modulus.

Combining Theorem 5 with Lemma 9 gives, whenever $2h_\sigma < 1$,

$$\mathrm{Adv} \geq \frac{n}{4}(1 - 2h_\sigma)\Delta^{\mathrm{H}}_m. \tag{59}$$

Thus if $q = \mathrm{poly}(n)$, $q \to \infty$, and $\sigma \to \infty$, the self-reciprocal family is distinguishable with advantage $\Omega(n/q)$ for every $m \geq 1$.

Section takeaway. Sections 4 and 5 have now exhausted the two conjugation patterns of Lemma 1. The reciprocal branch gives a product channel, while the fixed branch gives a Hermitian channel whose zero test runs in the opposite direction and already covers $m = 1$. Section 6 now assembles these local results into the prime-modulus classification.

6 Prime-Modulus Classification and Composite Transfer

Sections 4 and 5 have exhausted the two ways in which ring conjugation can act on a prime-modulus CRT factor. We first combine those local analyses into one prime-modulus classification theorem, then show that an exact divisor reduction transfers every such attack unchanged to larger composite moduli. The point of this section is not to introduce a third attack, but to make explicit which local channel applies, what advantage is proved, and how far those guarantees propagate beyond prime moduli.

6.1 Prime-modulus synthesis and phase diagram

Write $N = 2n$ and

$$f = \mathrm{ord}_N(q).$$

For $m \geq 2$, the two global theorems from Sections 4 and 5 give an exhaustive prime-modulus statement.

Theorem 6 (Prime-modulus attack classification). *Let n be a power of two, let q be an odd prime, let $R = \mathbb{Z}[X]/(X^n + 1)$, and assume $m \geq 2$. Put $f = \mathrm{ord}_{2n}(q)$. Exactly one of the following two cases occurs.*

1. *If $q \equiv -1 \pmod{2n}$, then every irreducible factor of $X^n + 1$ modulo q is self-reciprocal quadratic. With*

$$\kappa := \max_P \Pr[E_P = 0],$$

if $\kappa \leq 1 - \delta$ for a constant $\delta > 0$, the global Hermitian distinguisher of Theorem 5 satisfies

$$\mathrm{Adv} = \Omega_\delta(n/q). \tag{60}$$

Table 1. Prime-modulus attack phase diagram for $R = \mathbb{Z}[X]/(X^n + 1)$, with n a power of two.

Regime	CRT orbit / channel	Proven guarantee
$q \equiv 1 \pmod{2n}$ $m \geq 2$	reciprocal linear product; zero $\Rightarrow$ Real	global $\Omega_\delta(n/q)$
$q \equiv 1 \pmod{2n}$ $m = 1$, Gaussian	reciprocal linear gated product	local bound (35); $q^{-1}(1 + o(1))$ if $q, \sigma \to \infty$
$q \equiv -1 \pmod{2n}$ $m \geq 1$	fixed quadratic Hermitian; zero $\Rightarrow$ Rand	global $\Omega_\delta(n/q)$
$q \not\equiv \pm 1 \pmod{2n}$ $m \geq 2$	reciprocal degree f product; zero $\Rightarrow$ Real	global $\Omega_\delta(n/(fq^f))$

2. *Otherwise no factor is self-reciprocal, so all irreducible factors occur in distinct reciprocal pairs. With*

$$\rho := \max_P \Pr_{e \leftarrow \chi}[e_P = 0],$$

if $\rho^m \leq 1/2 - \delta$ for a constant $\delta > 0$, the global product distinguisher of Theorem 3 satisfies

$$\mathrm{Adv} = \Omega_\delta\left(\frac{n}{fq^f}\right). \tag{61}$$

In particular, $q \equiv 1 \pmod{2n}$ gives $f = 1$ and therefore $\mathrm{Adv} = \Omega_\delta(n/q)$.

The two cases are exhaustive for every odd prime modulus. The theorem classifies the guarantee of the present distinguisher; it is not a security characterization of parameter families for which the displayed lower bound is negligible.

Proof overview. Corollary 2 shows that an irreducible factor is fixed by conjugation if and only if $q \equiv -1 \pmod{2n}$. In that case Theorem 5 applies. In every other prime-modulus class all factors occur in reciprocal pairs, so Theorem 3 applies. The completely split case is the reciprocal case with $f = 1$. Appendix E records the complete synthesis argument and the elementary order calculation used below.

Corollary 3 (One-sample coverage). *For $m = 1$, the following two guarantees are already proved.*

1. *If $q \equiv -1 \pmod{2n}$ and $\max_P \Pr[E_P = 0] \leq 1 - \delta$, then the global Hermitian distinguisher has advantage $\Omega_\delta(n/q)$.*
2. *If $q \equiv 1 \pmod{2n}$ and the secret and error have independent coefficient-product discrete-Gaussian coefficients, then for any reciprocal linear pair Theorem 2 gives the explicit finite bound (35); in particular, when $q = \mathrm{poly}(n)$, $q \to \infty$, and $\sigma \to \infty$, its local advantage is $q^{-1}(1 + o(1))$.*

No global $m = 1$ reciprocal-factor claim beyond the proved cases above is used in this paper.

Table 1 summarizes both the exhaustive $m \geq 2$ classification and the one-sample statements above. The distinction in the last column is deliberate: the split $m = 1$ result is currently local, whereas the other rows state global amplified guarantees.

Why residue degree alone is not enough. The phase diagram contains a useful structural asymmetry. A reciprocal factor of degree two gives a product-channel signal on the q^{-2} scale, while a self-reciprocal quadratic factor gives a Hermitian signal on the q^{-1} scale. Thus two factors with the same residue-field size q^2 can differ by a factor q in local distinguishing strength. The orbit of the involution, not merely the residue degree, determines the relevant statistical channel. This is why the self-reciprocal row must remain separate from the generic $f = 2$ reciprocal row.

Reading f from the 2-adic position of q. Because $N = 2n$ is a power of two, the residue degree $f = \mathrm{ord}_N(q)$ is itself a power of two and is determined by how close q is 2-adically to $+1$ or -1. In particular, $q \equiv 1 \pmod{2n}$ gives $f = 1$, whereas $q \equiv -1 \pmod{2n}$ gives $f = 2$. The explicit closed form and its short derivation are recorded in Appendix E; they are included there only to support the phase diagram, not as a new arithmetic result.

6.2 Non-negligible regimes and the remaining boundary

The reciprocal theorem gives the scale

$$\frac{n}{fq^f}.$$

For a constant anti-degeneracy margin δ, this guaranteed advantage is non-negligible whenever

$$fq^f = \mathrm{poly}(n). \tag{62}$$

Since $q \geq 3$, condition (62) is equivalently captured, up to constants in the polynomial exponent, by

$$f \log q = O(\log n). \tag{63}$$

The criterion makes the reach of the classification transparent.

Completely split moduli. If $q \equiv 1 \pmod{2n}$, then $f = 1$. For $m \geq 2$ the global product theorem gives $\Omega_\delta(n/q)$, which is non-negligible whenever $q = \mathrm{poly}(n)$. For $m = 1$, Theorem 2 supplies the gated Gaussian local distinguisher summarized in Corollary 3.

Self-reciprocal moduli. If $q \equiv -1 \pmod{2n}$, the generic reciprocal criterion is not the right measure: the Hermitian channel is stronger. Under the activity hypothesis of Theorem 5, the advantage is $\Omega_\delta(n/q)$ for every $m \geq 1$. Thus polynomial-size q again gives a non-negligible distinguisher, even though the residue fields have size q^2.

Fixed small residue degree. More generally, in the reciprocal branch, if $f = O(1)$ and $q = \mathrm{poly}(n)$ then $q^f = \mathrm{poly}(n)$ and (62) holds. The product attack therefore covers prime-modulus families well beyond complete splitting; what matters quantitatively is the size of the residue field q^f together with the number $n/(2f)$ of reciprocal pairs.

Large residue degree. At the opposite extreme, the present guarantee can become negligible. For example, for the fixed prime $q = 3$ and $n \geq 4$ a power of two, the order formula in Appendix E gives $f = n/2$, so $q^f = 3^{n/2}$ and the reciprocal lower bound is exponentially small. More generally, any family with $f \log q = \omega(\log n)$ lies outside the non-negligible regime certified by Theorem 6. We deliberately say that such families are *not ruled out by the present distinguisher*; the theorem does not imply that they are secure.

6.3 Composite moduli by divisor reduction

The prime-modulus analysis above also yields attacks on a broader class of moduli essentially for free. Let $Q \geq 2$ be an arbitrary modulus and let $d \mid Q$. Write

$$\pi_{Q\to d} : R_Q \longrightarrow R_d$$

for coefficientwise reduction, and apply it componentwise to public vectors. The exact integral covariance $E^{\mathrm{int}} \in R$ is left unchanged.

Lemma 10 (Divisor reduction preserves KCov exactly). *For every $m \geq 1$ and every hidden distribution χ on R, the map*

$$(\mathbf{a}, \mathbf{b}, E^{\mathrm{int}}) \longmapsto (\pi_{Q\to d}(\mathbf{a}), \pi_{Q\to d}(\mathbf{b}), E^{\mathrm{int}})$$

has the exact push-forward identities

$$\pi_{Q\to d}(\mathsf{Real}_{R,Q,m,\chi}) = \mathsf{Real}_{R,d,m,\chi}, \qquad \pi_{Q\to d}(\mathsf{Rand}_{R,Q,m,\chi}) = \mathsf{Rand}_{R,d,m,\chi}.$$

Consequently, every polynomial-time distinguisher modulo d lifts by public reduction to a distinguisher modulo Q with identical signed gap and identical advantage.

Proof. Reduction modulo d is a surjective ring homomorphism because $d \mid Q$. Hence a uniform element of R_Q maps to a uniform element of R_d: every fiber has cardinality $(Q/d)^n$. In the real game, $\mathbf{b} = \mathbf{a}s + \mathbf{e} \pmod{Q}$ therefore reduces to $\pi(\mathbf{b}) = \pi(\mathbf{a})[s]_d + [\mathbf{e}]_d \pmod{d}$ with the same hidden $s, \mathbf{e}$ and the same exact covariance $E^{\mathrm{int}} = s\bar{s} + \sum_i e_i\bar{e}_i$. In the random game, the two independent uniform public vectors remain independent and uniform after reduction. The two push-forward distributions are thus exactly the lower-modulus KCov games, so composing any lower-modulus distinguisher with $\pi_{Q\to d}$ preserves both acceptance probabilities.

Corollary 4 (Composite-modulus transfer). *Let Q be an integer modulus and let $p \mid Q$ be an efficiently available odd prime divisor. Every attack guarantee proved in Theorem 6 or Corollary 3 for the parameters (R, p, m, χ) yields an attack with identical advantage against (R, Q, m, χ). In particular, for $m \geq 2$:*

1. *if $p \equiv -1 \pmod{2n}$ and the activity hypothesis modulo p holds, then the original modulus-Q instance inherits advantage $\Omega_\delta(n/p)$;*
2. *otherwise, if $f_p = \mathrm{ord}_{2n}(p)$ and the reciprocal anti-degeneracy hypothesis modulo p holds, it inherits advantage $\Omega_\delta(n/(f_p p^{f_p}))$.*

The corresponding one-sample statements transfer in exactly the same way.

Thus a composite modulus is not an escape from the attack merely because it is not prime: one vulnerable odd prime divisor suffices. For asymptotic families with $Q = \mathrm{poly}(n)$, all prime divisors can in particular be found by trial division in polynomial time in n. The transfer does not cover pure power-of-two moduli, and it gives no non-negligible guarantee when every available odd prime divisor falls outside the non-negligible regimes identified above.

Section takeaway. For $m \geq 2$, every odd prime modulus is assigned to one of two exhaustive attack channels, with an explicit guaranteed scale, and Lemma 10 propagates any such attack unchanged through larger composite moduli. The completely split and self-reciprocal prime classes both give $\Omega(n/q)$ behavior for different algebraic reasons, while generic reciprocal families weaken as q^f grows. The remaining question is no longer which local channel applies, but what these assumption-level attacks imply for the motivating cryptographic construction; Section 7 turns to that question.

7 Cryptographic Consequences and Limitations

Sections 4–6 establish an assumption-level attack. We now separate three logically distinct consequences: failure of the corresponding decisional KCov-RLWE instances, distinguishability of the covariance-revealing public-key hybrid suggested by the motivating Ring-PKE sketch, and what does *not* follow about the full encryption scheme. This separation is important: the first two statements are formal consequences of the preceding sections, whereas an IND-CPA or simulation-security attack would require an additional argument that we do not give.

7.1 From the classification to assumption failure

The hardness statement proposed by Micciancio and Suhl is decisional: a genuine Ring-LWE output should remain pseudorandom even when the same hidden secret and errors determine the public covariance hint [8]. The distinguishers constructed above therefore refute that assumption exactly in the asymptotic families where their proven advantage is non-negligible.

Corollary 5 (Assumption-level consequence). *Consider any asymptotic family of rank-one KCov-RLWE instances covered directly by Theorem 6 or Corollary 3, or indirectly through the divisor reduction of Corollary 4. If the corresponding lower bound on* Adv *is non-negligible, then decisional KCov-RLWE does not satisfy its stated pseudorandomness requirement for that family. In particular, for $m \geq 2$ this gives a non-negligible break whenever some relevant odd prime modulus or prime divisor satisfies the stated anti-degeneracy hypothesis and either the Hermitian regime has $p = \mathrm{poly}(n)$ or the reciprocal regime has $f_p p^{f_p} = \mathrm{poly}(n)$. It also applies to the $m = 1$ families covered by Corollary 3 and transferred by Corollary 4 whenever the resulting advantage is non-negligible.*

Proof overview. Theorems 3 and 5, together with the one-sample results summarized in Corollary 3, give explicit polynomial-time distinguishers for the very real/random experiments in the KCov-RLWE definition. Lemma 10 preserves these experiments and their distinguishing gaps exactly when passing from a larger modulus to a divisor. A non-negligible lower bound on any resulting advantage is therefore incompatible with decisional hardness. Appendix F records the remaining formal one-line reduction.

This statement is deliberately parameterized. It does not say that every modulus instance is distinguishable with non-negligible advantage. In the reciprocal branch, Section 6.2 identifies the large-q^f prime factors for which the present lower bound can become negligible, while pure power-of-two moduli are not reached by the divisor reduction of Section 6.3.

7.2 From the public padding descriptor to a distinguishable key hybrid

The preceding attack is defined from the integral covariance E^{int}, but in fact uses only its public residue $E = [E^{\mathrm{int}}]_q$; see Section 3. It is therefore natural to ask whether the motivating Ring-PKE sketch exposes

enough public information to reconstruct this input. Micciancio and Suhl's Section 5.3.3 explicitly presents a Ring-PKE *sketch* whose public key contains the RLWE samples together with a padding descriptor (F, λ) derived from the same hidden covariance [8]. To avoid overloading a ring element with its matrix representation, let $M_h \in \mathbb{Z}^{n\times n}$ denote negacyclic multiplication by $h \in R = \mathbb{Z}[X]/(X^n+1)$ in the power basis. In this notation the source descriptor has the form

$$E^{\mathrm{int}} = \bar{s}s + \sum_{i=1}^{m} e_i\bar{e}_i, \qquad F = \lambda I_n - \sigma_r^2 M_{E^{\mathrm{int}}}, \tag{64}$$

where the Gaussian width σ_r and λ are public.

Proposition 2 (Recovering covariance from the public padding descriptor). *Assume the descriptor in* (64) *is available in an exact, efficiently manipulable representation, with public λ and $\sigma_r^2 > 0$. Then E^{int}, and hence $E = [E^{\mathrm{int}}]_q$, is recoverable in polynomial time in the public encoding length.*

More generally, suppose an approximate matrix $\widehat{F}$ is accompanied by a certified error bound satisfying

$$\max_{0\le j<n} |(\widehat{F} - F)_{j,0}| < \frac{\sigma_r^2}{2}. \tag{65}$$

Then the integral coefficients of E^{int} are again recovered uniquely by scaling the first column and rounding.

Proof overview. Multiplication by E^{int} sends the power-basis vector representing 1 to the coefficient vector of E^{int}. Hence the first column of $\lambda I_n - F = \sigma_r^2 M_{E^{\mathrm{int}}}$ is exactly σ_r^2 times that coefficient vector. Division by the public nonzero scale recovers E^{int}; under (65), every resulting coefficient is within strictly less than $1/2$ of its unique integer value. The full encoding-level argument is in Appendix F.

Proposition 2 lets us state the construction-level consequence without conflating it with encryption security. Let $\mathsf{Pad}(E^{\mathrm{int}})$ denote the source rule that generates (F, λ) from the hidden covariance. First sample $s, \mathbf{e}$ exactly as in the KCov game, form $E^{\mathrm{int}} = \bar{s}s + \mathbf{e}^\star\mathbf{e}$, and generate the common descriptor $(F, \lambda) \leftarrow \mathsf{Pad}(E^{\mathrm{int}})$. Conditional on this common hidden sampling, define

$$\begin{aligned} \mathsf{KeyReal}_m: &\quad \mathbf{a} \leftarrow R_q^m, \quad \mathbf{b} = \mathbf{a}s + \mathbf{e}, &&\text{return } (\mathbf{a}, \mathbf{b}, F, \lambda), \\ \mathsf{KeyUnif}_m: &\quad \mathbf{a}, \mathbf{u} \leftarrow R_q^m \text{ independently}, &&\text{return } (\mathbf{a}, \mathbf{u}, F, \lambda). \end{aligned} \tag{66}$$

Thus only the RLWE output vector changes between the two distributions; the hidden tuple and covariance-derived descriptor are coupled identically.

Proposition 3 (Transfer to the covariance-revealing public-key hybrid). *Suppose $E = [E^{\mathrm{int}}]_q$ is recovered exactly from (F, λ). Every KCov-RLWE distinguisher whose computation depends on the covariance only through the reduced value E induces a polynomial-time distinguisher between $\mathsf{KeyReal}_m$ and $\mathsf{KeyUnif}_m$ with identical distinguishing advantage. Consequently, all attack bounds from Sections 4–6 transfer to the matching key hybrid whenever their secret/error distributional hypotheses are met.*

If recovery fails with probabilities at most $\varepsilon_{\mathsf{Real}}$ and $\varepsilon_{\mathsf{Rand}}$ in the two hybrid worlds, respectively, the transferred advantage is at least the KCov advantage minus $\varepsilon_{\mathsf{Real}} + \varepsilon_{\mathsf{Rand}}$.

Proof overview. Recover $E = [E^{\mathrm{int}}]_q$ from the descriptor and run the KCov distinguisher on $(\mathbf{a}, \mathbf{b}, E)$ or $(\mathbf{a}, \mathbf{u}, E)$. With exact recovery these reconstructed triples have exactly the KCov real and random laws, respectively, because the descriptor is generated from the same hidden covariance in both hybrids.

Approximate recovery can change the output only on a recovery failure event. Appendix F gives the formal coupling.

The conclusion of Proposition 3 is the strongest construction-level statement needed here: in the affected parameter regimes, the RLWE public output cannot be replaced by an independent uniform vector while retaining the same covariance-derived padding descriptor and appealing to KCov pseudorandomness. This is a failure of the intended KCov-based hybrid step; it is not, by itself, an IND-CPA attack on the complete encryption algorithm.

7.3 One sample, construction-level meaning, and limits

The one-sample results matter directly to the motivation of the assumption. Micciancio and Suhl remark in their Ring-PKE sketch that $m = 1$ would suffice if KCov-RLWE were hard already for one sample [8]. Our classification shows that reducing the public key to one RLWE sample is not a generic escape from the covariance issue. In the self-reciprocal class $q \equiv -1 \pmod{2n}$, Theorem 5 already applies for every $m \geq 1$ under its activity hypothesis. In the completely split class $q \equiv 1 \pmod{2n}$, Theorem 2 gives the gated one-sample distinguisher when the hidden law is the coefficient-product Gaussian specialization used there. We do not claim a global $m = 1$ theorem for generic nonsplit reciprocal factors. The original Ring-PKE sketch also permits a broader modulus design space than the vulnerable congruence classes analyzed above; our claim is therefore that the quantified prime families, and composite families with a vulnerable odd prime divisor, invalidate the KCov assumption in those regimes, not that every parameter choice permitted by the sketch is insecure.

The resulting hierarchy of claims is summarized below. The distinctions are part of the theorem statement, not merely caveats about implementation.

Level	✓ Established	× Does not follow
KCov-RLWE assumption	Non-negligible distinguisher	All moduli; search/recovery
Public-key hybrid	Real $\not\approx$ Uniform*	Secret-key recovery
Ring-PKE proof step	KCov replacement step fails*	Full IND-CPA; simulation-security break

*When the public descriptor permits recovery of $E = [E^{\mathrm{int}}]_q$ as in Proposition 2.

Section takeaway. The cryptanalytic consequence therefore has a precise three-level form. The preceding sections refute decisional KCov-RLWE in the quantified regimes; the public descriptor of the motivating Ring-PKE sketch can expose the covariance residue needed by the distinguisher; and this transfers the attack to the matching real-versus-uniform public-key hybrid. What remains deliberately open is a direct attack on the full encryption security game.

8 Relation to Prior Work and Contribution Boundary

The closest literature spans positive leakage-hardness results, leftover-hash lemmas, algebraic attacks on ordinary Ring-LWE, and threshold constructions that avoid known-covariance assumptions. Table 2 compares the properties that matter for our result. A checkmark means that the cited work explicitly addresses the property in the relevant setting; △ denotes an adjacent or partial treatment; "–" means that the property is not addressed.

Table 2. Property matrix for the closest prior work. ✓ = explicit; △ = adjacent/partial; – = not addressed.

Work	Focus	L	Q	G	A	O	H
Micciancio–Suhl [8]	open	✓	✓	–	–	–	△
Wang–Lai–Liu [13]	hardness	✓	–	–	–	–	–
BJTW [2]	search	✓	✓	–	–	–	–
Boudgoust–Lapiha [3]	LHL	✓	△	–	–	–	–
ELOS / Peikert / CLS [5, 10, 4]	attack	–	–	–	✓	△	–
Passelègue–Stehlé / Okada–Takagi [11, 9]	construct	△	△	–	–	–	–
This work	**attack**	✓	✓	✓	✓	✓	✓

L: leakage-aware LWE/RLWE; Q: exact KCov hint; G: bordered-Gram residual; A: direct cryptanalytic attack; O: CRT/involution classification; H: public-key-hybrid consequence.

Positive leakage and extraction results. Wang–Lai–Liu prove decision R/MLWE hardness under linear leakage [13]. BJTW go substantially further: their framework covers broad leakage classes and explicitly instantiates the exact quadratic hint $x^\star x$ for Known-Covariance M-LWE [2]. This is the closest positive result to ours. Its formal scope is nevertheless different in two decisive ways: the main hardness reduction requires module rank $d > 1$ (and explicitly excludes the ring setting $d = 1$), and the resulting KCov guarantee is search hardness, while a decision extension is left open. Thus the separation is not "quadratic versus nonquadratic leakage" or "HNF versus non-HNF"—BJTW already cover both the quadratic hint and an HNF-shaped $[A \mid I]$ transcript. Our contribution is a decision attack precisely in the excluded rank-one regime, where the hidden factor becomes square and its Gram determinant collapses to one residual. Boudgoust–Lapiha give cyclotomic leftover-hash lemmas with broad leakage, and Albrecht et al. establish a Gaussian LHL over number-field modules [3, 1]. Those extraction results use different hash/distribution hypotheses and do not analyze the rank-one KCov Gram collapse exploited here.

Closest attack and construction lines. Micciancio–Suhl introduce the KCov question and leave the Ring-LWE case open [8]. Classical finite-field attacks of Elias–Lauter–Ozman–Stange, Peikert, and Chen–Lauter–Stange also project RLWE samples to residue fields, but their signal comes from projected-error geometry or nonuniformity rather than from covariance completion [5, 10, 4]. On the construction side, Passelègue–Stehlé rely on a related known-norm assumption, whereas Okada–Takagi subsequently pursue standard-assumption alternatives; neither work cryptanalyzes KCov itself [11, 9]. Our use of CRT is therefore not new; the new object is the covariance-induced residual that is localized only after it has been made public.

Novelty boundary. We do not claim the first use of CRT projections against Ring-LWE, the first LWE result with leakage, or the first treatment of quadratic leakage. To the best of our knowledge, prior work does not give the chain

$$\begin{aligned}\text{exact KCov hint} &\longrightarrow \text{bordered Gram completion} \longrightarrow \det \Gamma = \bar{v}v,\\ &\longrightarrow \{\text{reciprocal-product law, Hermitian-norm law}\}.\end{aligned}$$

which yields the explicit decision attacks, the prime-modulus classification, the composite-modulus transfer, and the public-key-hybrid consequence proved in Sections 4–7.

9 Conclusion

Exact quadratic covariance is not benign merely because it is aggregate leakage. At ring rank one, the KCov hint completes a square bordered Gram matrix whose public determinant collapses on genuine samples to the norm of one hidden affine residual. Localizing this same residual through the CRT yields exactly two statistical channels: a product across distinct reciprocal factors and a Hermitian norm on self-reciprocal factors. This dichotomy gives an exhaustive prime-modulus attack classification for $m \geq 2$, with global advantage $\Omega_\delta(n/(fq^f))$ in the reciprocal regime and $\Omega_\delta(n/q)$ in the self-reciprocal regime, together with the one-sample guarantees stated in Sections 4–6. A divisor-reduction lemma then lifts every one of these guarantees, with identical advantage, to any larger composite modulus containing the corresponding odd prime divisor. The resulting picture also shows that residue degree alone does not determine attack strength: the orbit of the involution is an essential part of the security geometry.

The same distinction carries to the motivating Ring-PKE sketch only at the level that can be justified formally. When its public padding descriptor permits recovery of the covariance residue, the corresponding KCov distinguisher transfers to the matching real-versus-uniform public-key hybrid, so the intended KCov-based pseudorandomness replacement cannot be invoked unchanged in the affected families. We do not derive secret recovery or a direct IND-CPA or simulation-security attack on the full construction. Reciprocal prime factors with superpolynomial q^f, pure power-of-two moduli, generic nonsplit one-sample instances, and mechanisms beyond the rank-one determinant collapse therefore remain open. More broadly, our results illustrate that exact structured leakage should be analyzed jointly with the public linear map that accompanies it: leakage size or algebraic degree alone can miss identities that become visible only after the two are combined.

References

1. Albrecht, M.R., Felderhoff, J., Lai, R.W.F., Lapiha, O., Woo, I.K.Y.: A gaussian leftover hash lemma for modules over number fields. In: Advances in Cryptology – EUROCRYPT 2026, Part IV. Lecture Notes in Computer Science, vol. 16544, pp. 124–153. Springer (2026). https://doi.org/10.1007/978-3-032-25327-9˙5
2. Boudgoust, K., Jeudy, C., Tairi, E., Wen, W.: Hardness of m-lwe with general distributions and applications to leaky variants. In: Public-Key Cryptography – PKC 2026. Lecture Notes in Computer Science, vol. 16551, pp. 3–37. Springer (2026). https://doi.org/10.1007/978-3-032-26731-3˙1
3. Boudgoust, K., Lapiha, O.: Leftover hash lemma(s) over cyclotomic rings. In: Advances in Cryptology – EUROCRYPT 2026, Part IV. Lecture Notes in Computer Science, vol. 16544, pp. 62–91. Springer (2026). https://doi.org/10.1007/978-3-032-25327-9˙3
4. Chen, H., Lauter, K.E., Stange, K.E.: Attacks on the search rlwe problem with small errors. SIAM Journal on Applied Algebra and Geometry **1**(1), 665–682 (2017). https://doi.org/10.1137/16M1096566
5. Elias, Y., Lauter, K.E., Ozman, E., Stange, K.E.: Provably weak instances of Ring-LWE. In: Advances in Cryptology – CRYPTO 2015, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 63–92. Springer (2015). https://doi.org/10.1007/978-3-662-47989-6˙4
6. Lidl, R., Niederreiter, H.: Finite Fields, Encyclopedia of Mathematics and its Applications, vol. 20. Cambridge University Press, 2 edn. (1997)
7. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Advances in Cryptology – EUROCRYPT 2010. Lecture Notes in Computer Science, vol. 6110, pp. 1–23. Springer (2010). https://doi.org/10.1007/978-3-642-13190-5˙1
8. Micciancio, D., Suhl, A.: Simulation-secure threshold pke from lwe with polynomial modulus. Cryptology ePrint Archive, Paper 2023/1728 (2023), final version in IACR Communications in Cryptology, vol. 1, no. 4

9. Okada, H., Takagi, T.: Low communication threshold FHE from standard (module-)LWE. In: Advances in Cryptology – ASIACRYPT 2025, Part VII. Lecture Notes in Computer Science, vol. 16251, pp. 165–198. Springer (2025). https://doi.org/10.1007/978-981-95-5122-4˙6
10. Peikert, C.: How (not) to instantiate Ring-LWE. In: Security and Cryptography for Networks – SCN 2016. Lecture Notes in Computer Science, vol. 9841, pp. 411–430. Springer (2016). https://doi.org/10.1007/978-3-319-44618-9˙22
11. Passelègue, A., Stehlé, D.: Low communication threshold fully homomorphic encryption. In: Advances in Cryptology – ASIACRYPT 2024, Part I. Lecture Notes in Computer Science, vol. 15484, pp. 297–329. Springer (2024). https://doi.org/10.1007/978-981-96-0875-1˙10
12. Sangwisut, E., Jitman, S., Ling, S., Udomkavanich, P.: Hulls of cyclic and negacyclic codes over finite fields. Finite Fields and Their Applications **33**, 232–257 (2015). https://doi.org/10.1016/j.ffa.2014.12.008
13. Wang, Z., Lai, Q., Liu, F.H.: Ring/module learning with errors under linear leakage – hardness and applications. In: Public-Key Cryptography – PKC 2024. Lecture Notes in Computer Science, vol. 14602, pp. 275–304. Springer (2024). https://doi.org/10.1007/978-3-031-57722-2˙9

A Cyclotomic Factorization Details

This appendix gives the complete proof of Lemma 1.

Proof (Lemma 1). Let α be a primitive N-th root of unity in an algebraic closure of $\mathbb{F}_q$. Because $q \nmid N$, the polynomial $X^N - 1$ is separable over $\mathbb{F}_q$, hence so is its factor $\Phi_N(X) = X^n + 1$.

The Frobenius automorphism sends α to α^q. The size of the Frobenius orbit of α is therefore the smallest positive integer j for which $\alpha^{q^j} = \alpha$, equivalently the smallest j such that $q^j \equiv 1 \pmod{N}$. By definition, this size is $f = \mathrm{ord}_N(q)$. Every root of Φ_N is a primitive N-th root, so every Frobenius orbit has the same size f. Since $\deg \Phi_N = n$, the polynomial $X^n + 1$ therefore factors over $\mathbb{F}_q$ into exactly n/f distinct monic irreducible factors, all of degree f.

Let P be one of these factors. Its roots form a Frobenius orbit $\{\alpha, \alpha^q, \ldots, \alpha^{q^{f-1}}\}$ for some primitive N-th root α. The roots of the monic reciprocal $P^\vee(X) = P(0)^{-1} X^f P(X^{-1})$ are the inverses of these roots, namely $\{\alpha^{-1}, \alpha^{-q}, \ldots, \alpha^{-q^{f-1}}\}$. This set is again one Frobenius orbit of primitive N-th roots, so $P^\vee$ is another irreducible factor of $X^n + 1$ of degree f.

The ring involution $a(X) \mapsto a(X^{-1})$ sends the ideal generated by P to the ideal generated, up to a nonzero scalar and a power of X, by $P^\vee$. Consequently it exchanges the corresponding CRT components, and induces an isomorphism between the quotients indexed by P and $P^\vee$.

Finally, $P = P^\vee$ exactly when inversion preserves the Frobenius orbit of α. This happens precisely when $\alpha^{-1} = \alpha^{q^j}$ for some $0 \le j < f$. Since α has exact order N, the latter equality is equivalent to $q^j \equiv -1 \pmod{N}$, proving the stated self-reciprocity criterion.

B Bordered-Gram Residual Proofs

This appendix contains the complete algebraic proofs for the formal results of Section 3. The main text states the results and gives the proof mechanisms needed to follow the attack without entering these calculations.

Proof (Lemma 2). Consider the square $(m+1) \times (m+1)$ matrix

$$L = \begin{pmatrix} I_m & \mathbf{a} \\ \mathbf{e}^\star & \bar{s} \end{pmatrix}. \tag{67}$$

Its adjoint is

$$L^\star = \begin{pmatrix} I_m & \mathbf{e} \\ \mathbf{a}^\star & s \end{pmatrix}.$$

Using $\mathbf{b} = \mathbf{a}s + \mathbf{e}$ and $E = \bar{s}s + \mathbf{e}^\star\mathbf{e}$, direct multiplication gives

$$LL^\star = \begin{pmatrix} I_m + \mathbf{a}\mathbf{a}^\star & \mathbf{b} \\ \mathbf{b}^\star & E \end{pmatrix} = \Gamma. \tag{68}$$

Because S is commutative,

$$\det \Gamma = \det(L)\det(L^\star) = \det(L)\,\overline{\det(L)}.$$

It remains to compute $\det(L)$. Right-multiply L by the determinant-one block matrix

$$T = \begin{pmatrix} I_m & -\mathbf{a} \\ 0 & 1 \end{pmatrix}.$$

Then

$$LT = \begin{pmatrix} I_m & 0 \\ \mathbf{e}^\star & \bar{s} - \mathbf{e}^\star\mathbf{a} \end{pmatrix},$$

so

$$\det(L) = \bar{s} - \mathbf{e}^\star\mathbf{a} = \overline{s - \mathbf{a}^\star\mathbf{e}} = \bar{v}.$$

Consequently $\det \Gamma = \bar{v}v$, as claimed.

Proof (Corollary 1). Put $A = I_m + \mathbf{a}\mathbf{a}^\star$. The standard rank-one determinant and adjugate identities over a commutative ring are

$$\det(A) = 1 + \mathbf{a}^\star\mathbf{a} = h, \qquad \operatorname{adj}(A) = hI_m - \mathbf{a}\mathbf{a}^\star.$$

The bordered-determinant formula therefore yields

$$\det \Gamma = E\det(A) - \mathbf{b}^\star \operatorname{adj}(A)\mathbf{b},$$

which expands to (12). Self-conjugacy follows from Lemma 2, because $W = \bar{v}v = \overline{\bar{v}v}$.

Proof (Proposition 1). Direct block multiplication gives

$$L_d L_d^\star = \begin{pmatrix} I_m + AA^\star & A\mathbf{s} + \mathbf{e} \\ (A\mathbf{s} + \mathbf{e})^\star & \mathbf{e}^\star\mathbf{e} + \mathbf{s}^\star\mathbf{s} \end{pmatrix},$$

which is exactly (14). Since L_d has $m+1$ rows and $m+d$ columns, the Cauchy–Binet identity gives (15), with one term for every choice of $m+1$ columns. When $d = 1$, $m + d = m + 1$ and there is only one such choice, so the sum is the single norm term from Lemma 2. When $d > 1$ there are multiple choices of J, hence multiple minor norms. Therefore the factorization alone no longer collapses to one affine residual.

C Reciprocal-Factor Probability Laws and Proofs

This appendix gives the complete finite-field and amplification arguments used in Section 4. No Gaussian assumption is used except in the final one-sample specialization.

Proof (Lemma 3). Because $q \nmid 2n$, the polynomial $X^n + 1$ is square-free modulo q. The Chinese remainder theorem therefore gives a ring isomorphism from R_q to the product of its irreducible-factor fields. Under any bijection between finite sets, the uniform distribution maps to the uniform distribution; uniformity on a direct product is the product of the uniform coordinate laws. Hence the P and $P^\vee$ coordinates of a uniform ring element are independent uniform field elements. Applying the field isomorphism τ_P to the second coordinate preserves this law. The vector statement follows coordinate-wise.

Proof (Lemma 4). Project $W = \bar{v}v$ from Lemma 2 to P. By the transport convention, $(\bar{v})_P = v_{\bar{P}}$, hence $W_P = v_{\bar{P}} v_P$. Moreover

$$v_P = s_1 - \mathbf{y}^T \mathbf{u}, \qquad v_{\bar{P}} = s_2 - \mathbf{x}^T \mathbf{z},$$

which proves (18). Conditional on fixed hidden data, $\mathbf{x}$ and $\mathbf{y}$ are independent uniform vectors by Lemma 3. If $\mathbf{u} \neq 0$, the map $\mathbf{y} \mapsto \mathbf{y}^T\mathbf{u}$ is a surjective linear functional on K^m, so v_P is uniform on K; similarly $v_{\bar{P}}$ is uniform when $\mathbf{z} \neq 0$. Their independence follows from that of $\mathbf{x}$ and $\mathbf{y}$. Finally, two independent uniform field elements have nonzero product exactly when both are nonzero, giving

$$1 - (1 - Q^{-1})^2 = 2Q^{-1} - Q^{-2}.$$

For completeness we derive (21). Let $r_P = \rho_P^m$, $r_{\bar{P}} = \rho_{\bar{P}}^m$ and $t_P = \theta_P^m$. On the event that neither projected error vector vanishes the zero probability is α_1. If exactly one vanishes, the remaining affine factor is uniform, so the zero probability is at least $1/Q$; if both vanish we use the lower bound zero. Inclusion–exclusion yields

$$\begin{aligned} \Pr[W_P = 0 \mid \mathsf{Real}] &\geq \alpha_1(1 - r_P - r_{\bar{P}} + t_P) + Q^{-1}(r_P + r_{\bar{P}} - 2t_P) \\ &= \alpha_1 - \beta_1(r_P + r_{\bar{P}}) - Q^{-2}t_P, \end{aligned}$$

which is (21).

Uniform bilinear-form law. Let $U, V \leftarrow K^d$ independently and uniformly. If $U = 0$, then $U^T V = 0$ surely; if $U \neq 0$, the map $V \mapsto U^T V$ is a surjective linear functional and every field value has probability $1/Q$. Hence

$$\Pr[U^T V = 0] = Q^{-d} + (1 - Q^{-d})Q^{-1} = \alpha_d,$$

and for any fixed $z \neq 0$,

$$\Pr[U^T V = z] = (1 - Q^{-d})Q^{-1} = \beta_d.$$

If M has rank d, invertible row and column changes put it in rank-normal form $\mathrm{diag}(I_d, 0)$; uniformity is invariant under those changes, so $P^T M Q'$ has the same law.

Proof (Lemma 5). The rank-one determinant identity gives

$$\det((1 + c)I_m - \mathbf{y}\mathbf{x}^T) = (1 + c)^{m-1}\big((1 + c) - \mathbf{x}^T\mathbf{y}\big) = (1 + c)^{m-1}.$$

Thus $c \neq -1$ implies full rank. If $c = -1$, then $M = -\mathbf{y}\mathbf{x}^T$; the condition $\mathbf{x}^T\mathbf{y} = -1$ forces both vectors to be nonzero, so M has rank one. Since c is the dot product of two independent uniform vectors and $-1 \neq 0$, the bilinear-form law gives $\Pr[c = -1] = \beta_m$.

On $c = -1$, the shift $(1 + c)\varepsilon$ vanishes and the rank-one form is zero with probability α_1. On $c \neq -1$, the form has rank m. If $\varepsilon = 0$, zero requires the form to take its zero value, with probability α_m; if $\varepsilon \neq 0$, it must hit a fixed nonzero target, with probability β_m. Averaging over the event $c = -1$ gives (25)–(26). Their difference is

$$P_0^{\mathsf{Rand}} - P_1^{\mathsf{Rand}} = \frac{Q^m(Q-1)+1}{Q^{2m+1}} > 0,$$

so P_0^{Rand} is the worst leakage value.

A direct simplification of $\alpha_1 - P_0^{\mathsf{Rand}}$ gives (27).

Proof (Theorem 1). Equation (21) gives

$$\Pr[\mathcal{D}_P = 1 \mid \mathsf{Real}] \geq \alpha_1 - \beta_1(\rho_P^m + \rho_{\bar{P}}^m) - \theta_P^m/Q^2.$$

For every possible covariance value, Lemma 5 gives

$$\Pr[\mathcal{D}_P = 1 \mid \mathsf{Rand}, E_P] \leq P_0^{\mathsf{Rand}}.$$

Averaging over E_P preserves this pointwise upper bound. Subtracting and using $\Delta_0 = \alpha_1 - P_0^{\mathsf{Rand}}$ proves (28). Since $\rho_P^m, \rho_{\bar{P}}^m, \theta_P^m \leq \rho^m$ and $2\beta_1 + Q^{-2} = \alpha_1$, we obtain (29). Finally, for fixed $m \geq 2$, $\Delta_0 = Q^{-1}(1 + O(Q^{-1}))$ and $\alpha_1 = 2Q^{-1}(1 + O(Q^{-1}))$, which gives the two asymptotic statements. The test uses $O(m)$ ring operations to compute W and one reduction modulo P, hence is polynomial-time.

Exact one-sample random law. For $m = 1$ and $\varepsilon \neq 0$, condition on x, y. If $1 + xy = 0$, which occurs with probability β_1, the equation $(1 + xy)\varepsilon - wz = 0$ asks for $wz = 0$ and succeeds with probability α_1. Otherwise it asks a product of two uniform field elements to hit a specified nonzero value, which succeeds with probability β_1. Therefore

$$\Pr[W_P = 0 \mid \varepsilon \neq 0, \mathsf{Rand}] = \beta_1\alpha_1 + (1 - \beta_1)\beta_1 = Q^{-1} - Q^{-3},$$

proving (33); subtraction from α_1 gives (34).

One-coefficient anti-concentration. For the fully split setting of Theorem 2, let Z_s and Z_e be the constant coefficients of s and e, and let

$$\eta_s = \max_{r \in \mathbb{F}_q} \Pr[Z_s \equiv r \pmod q], \qquad \eta_e = \max_{r \in \mathbb{F}_q} \Pr[Z_e \equiv r \pmod q].$$

Fix all secret coefficients except Z_s and fix the complete error. At a reciprocal root pair (ζ, ζ^{-1}), the covariance coordinate is a monic quadratic in Z_s, so at most two residue classes make it zero. Hence $\Pr[E_\zeta = 0] \leq 2\eta_s$. Repeating with Z_e gives the alternative bound $2\eta_e$, and therefore

$$\Pr[E_\zeta = 0] \leq 2\min\{\eta_s, \eta_e\}.$$

Similarly, after fixing all nonconstant error coefficients, either equation $e(\zeta) = 0$ or $e(\zeta^{-1}) = 0$ fixes at most one residue class of Z_e, so each has probability at most η_e.

For $Z \leftarrow D_{\mathbb{Z},\sigma}$, the mass of any residue class modulo q is at most the centered-coset mass. Writing $\rho_t(\mathbb{Z}) = \sum_{z\in\mathbb{Z}} e^{-\pi z^2/t^2}$,

$$\max_r \Pr[Z \equiv r \pmod q] \le \frac{\rho_{\sigma/q}(\mathbb{Z})}{\rho_\sigma(\mathbb{Z})} \le \frac{1+\sigma/q}{\sigma} = \sigma^{-1} + q^{-1},$$

where the second inequality uses the elementary upper bound $\rho_{\sigma/q}(\mathbb{Z}) \le 1 + \sigma/q$ and Poisson summation $\rho_\sigma(\mathbb{Z}) = \sigma\rho_{1/\sigma}(\mathbb{Z}) \ge \sigma$. Truncating at one gives h_σ.

Proof (Theorem 2). Let $\kappa = \Pr[E_P = 0]$ and let ρ_e be the maximum of the two projected error-zero probabilities. A direct partition into the events where neither, exactly one, or both projected errors vanish gives the exact signed-gap bound

$$\delta_1(1-\kappa) - \alpha_1\rho_e \le \Pr[\mathcal{D}_P^{(1)} = 1 \mid \mathsf{Real}] - \Pr[\mathcal{D}_P^{(1)} = 1 \mid \mathsf{Rand}] \le \delta_1.$$

The anti-concentration argument above gives $\kappa \le 2h_\sigma$ and $\rho_e \le h_\sigma$. Substitution yields (35). Since $\delta_1 = q^{-1}(1 + O(q^{-1}))$ and $h_\sigma = o(1)$ under the stated asymptotic conditions, the claimed $q^{-1}(1 + o(1))$ advantage follows. The algorithm uses only a constant number of field evaluations and operations after the public residual is computed.

Proof (Theorem 3). Let $a = 1 - Q^{-1}$ and let $\mathcal{P}$ be the $L = n/(2f)$ reciprocal pairs. Choose a uniform subset $S \subseteq \mathcal{P}$ of size $J = \min\{L, \lfloor \delta Q/16 \rfloor\}$; for sufficiently large Q this is nonzero. The distinguisher accepts when $W_P = 0$ for at least one $P \in S$.

First condition on the full hidden tuple in the real experiment. For a selected pair P, let $t_P \in \{0,1,2\}$ count how many of its two projected error vectors are nonzero. The calculation behind (21) gives a conditional zero probability at least $1 - a^{t_P}$. The public multiplier CRT coordinates are independent across distinct selected pairs, so if $T_S = \sum_{P\in S} t_P$ then

$$\Pr[\text{accept} \mid s, \mathbf{e}, S, \mathsf{Real}] \ge 1 - a^{T_S}.$$

The function $F(t) = 1 - a^t$ is concave on $[0, 2J]$ and satisfies $F(0) = 0$; hence it lies above the chord joining 0 and $2J$:

$$F(t) \ge \frac{t}{2J} F(2J).$$

Averaging over the random subset and the hidden errors gives

$$\Pr[\text{accept} \mid \mathsf{Real}] \ge (1 - \rho^m)\big(1 - a^{2J}\big),$$

because each of the $2L$ error-side vectors is zero with probability at most ρ^m and uniform subsampling preserves the average expected count.

In the random experiment, condition on the complete covariance E. By Lemma 5, every selected pair has zero probability at most P_0^{Rand}, and the public $(\mathbf{a}, \mathbf{b})$ CRT coordinates are independent across distinct pairs. Therefore

$$\Pr[\text{accept} \mid \mathsf{Rand}] \le 1 - (1 - P_0^{\mathsf{Rand}})^J \le J P_0^{\mathsf{Rand}}.$$

Together with the real lower bound this proves the finite inequality (37).

We now make the constants explicit. From (25),

$$P_0^{\mathsf{Rand}} = \alpha_m + \beta_m(\alpha_1 - \alpha_m).$$

For every $m \geq 2$, $\alpha_m \leq Q^{-1} + Q^{-2}$, $\beta_m \leq Q^{-1}$, and $\alpha_1 - \alpha_m \leq Q^{-1}$. Hence the uniform finite bound

$$P_0^{\mathsf{Rand}} \leq Q^{-1} + 2Q^{-2} \tag{69}$$

holds without treating m as fixed.

Put $x = J/Q$. By construction $x \leq \delta/16$, and

$$1 - (1 - Q^{-1})^{2J} \geq 1 - e^{-2x} \geq 2x - 2x^2.$$

Since $\rho^m \leq 1/2 - \delta$ (and therefore $\delta \leq 1/2$), (69) gives

$$\begin{aligned} \mathrm{Adv} &\geq (1/2 + \delta)(2x - 2x^2) - x(1 + 2/Q) \\ &= 2\delta x - (1 + 2\delta)x^2 - 2x/Q. \end{aligned}$$

If $Q \geq 16/\delta$, then $(1 + 2\delta)x^2 \leq \delta x/8$ and $2x/Q \leq \delta x/8$; consequently

$$\mathrm{Adv} \geq \delta x = \delta J/Q,$$

which is (38).

If $J = L$, this directly gives $\mathrm{Adv} \geq \delta L/Q$. Otherwise, for $Q \geq 32/\delta$, $J = \lfloor \delta Q/16 \rfloor \geq \delta Q/32$, so $\mathrm{Adv} \geq \delta^2/32$. Moreover $Q = q^f \equiv 1 \pmod{2n}$ implies $Q \geq 2n + 1$ and therefore $L/Q < 1/(4f) \leq 1/4$; hence this constant lower bound is also $\Omega_\delta(L/Q)$. The finitely many smaller Q are irrelevant to the asymptotic claim. Substituting $L = n/(2f)$ and $Q = q^f$ proves (39).

D Self-Reciprocal Hermitian Probability Laws and Proofs

This appendix gives the complete arithmetic, finite-field, anti-concentration, and amplification proofs used in Section 5.

Proof (Corollary 2). Let $N = 2n$, so N is a power of two with $N \geq 4$. By Lemma 1, $P = P^\vee$ if and only if $q^j \equiv -1 \pmod{N}$ for some $0 \leq j < f$. Such a j cannot be zero. Reducing modulo 4, the congruence $q^j \equiv -1 \pmod 4$ implies $q \equiv -1 \pmod 4$ and j odd: if $q \equiv 1 \pmod 4$, all powers are 1 modulo 4, while if $q \equiv -1 \pmod 4$, an even power is 1 modulo 4.

For odd j,

$$q^j + 1 = (q + 1)(q^{j-1} - q^{j-2} + \cdots - q + 1).$$

The second factor is a sum of j odd terms modulo 2 and is therefore odd. Hence

$$v_2(q^j + 1) = v_2(q + 1).$$

Since $N \mid q^j + 1$, it follows that $N \mid q + 1$, i.e. $q \equiv -1 \pmod N$. Conversely, if $q \equiv -1 \pmod N$, the criterion of Lemma 1 holds with $j = 1$, so every irreducible factor is self-reciprocal.

The congruence $q \equiv -1 \pmod N$ gives $q^2 \equiv 1 \pmod N$ while $q \not\equiv 1 \pmod N$, hence $f = \mathrm{ord}_N(q) = 2$. Lemma 1 then gives $n/2$ irreducible factors, each of degree two. Therefore $K_P \simeq \mathbb{F}_{q^2}$. The involution induced by ring conjugation is a nontrivial $\mathbb{F}_q$-automorphism of this quadratic field, so it is the Frobenius automorphism $z \mapsto z^q$.

Proof (Lemma 6). Project the identity $W = \bar{v}v$ from Lemma 2 to the fixed factor P. By Corollary 2, conjugation on $K = K_P$ is $z \mapsto z^q$, and therefore

$$W_P = v_P^q v_P = \mathrm{N}_{K/\mathbb{F}_q}(v_P).$$

Projecting $v = s - \mathbf{a}^\star \mathbf{e}$ gives the displayed formula for v_P.

Now condition on the complete hidden tuple $(s, \mathbf{e})$ and assume $\varepsilon = E_P \neq 0$. If $\mathbf{e}_P \neq 0$, then $\overline{\mathbf{a}_P}$ is uniform on K^m because Frobenius is a bijection, and the map

$$\overline{\mathbf{a}_P} \longmapsto \overline{\mathbf{a}_P}^T \mathbf{e}_P$$

is a surjective K-linear functional. Hence $v_P = s_P - \mathbf{a}_P^\star \mathbf{e}_P$ is uniform on K, and

$$\Pr[W_P = 0 \mid s, \mathbf{e}, \varepsilon \neq 0, \mathsf{Real}] = \Pr[v_P = 0] = q^{-2},$$

because the field norm vanishes exactly at zero.

If instead $\mathbf{e}_P = 0$, then $v_P = s_P$. The projected covariance is

$$\varepsilon = \mathrm{N}(s_P) + \sum_{i=1}^{m} \mathrm{N}(e_{i,P}) = \mathrm{N}(s_P).$$

Activity implies $\mathrm{N}(s_P) \neq 0$, hence $s_P \neq 0$ and $W_P \neq 0$. This proves all claims.

Proof (Lemma 7). Fix a nontrivial additive character $\psi : \mathbb{F}_q \to \mathbb{C}^*$. For $t \in \mathbb{F}_q$, define

$$S(t) := \sum_{z \in K} \psi\big(t\, \mathrm{N}_{K/\mathbb{F}_q}(z)\big)\,.$$

Clearly $S(0) = q^2$. If $t \neq 0$, the norm map sends K^* surjectively onto $\mathbb{F}_q^*$ and every nonzero value has exactly $q+1$ preimages. Therefore

$$S(t) = 1 + (q+1) \sum_{u \in \mathbb{F}_q^*} \psi(tu) = 1 - (q+1) = -q.$$

Let $N_d(c)$ denote the number of $\mathbf{x} \in K^d$ with $H_d(\mathbf{x}) = c$. Additive-character orthogonality gives

$$\begin{aligned} N_d(c) &= \frac{1}{q} \sum_{t \in \mathbb{F}_q} \psi(-tc) \left(\sum_{z \in K} \psi(t\, \mathrm{N}(z)) \right)^d \\ &= q^{2d-1} + \frac{(-q)^d}{q} \sum_{t \in \mathbb{F}_q^*} \psi(-tc). \end{aligned}$$

If $c = 0$, the last sum is $q - 1$, so

$$N_d(0) = q^{2d-1} + (-1)^d (q-1) q^{d-1}.$$

If $c \neq 0$, the last sum is -1, so

$$N_d(c) = q^{2d-1} - (-1)^d q^{d-1}.$$

Dividing by $|K|^d = q^{2d}$ yields (45) and (46). Substituting $d = 1$ gives $A_1 = q^{-2}$.

We use the following elementary Hermitian diagonalization fact inside the random law proof. If $M \in K^{m\times m}$ is Hermitian of rank r, then there exists $C \in \mathrm{GL}_m(K)$ such that

$$C^\star MC = \mathrm{diag}(I_r, 0). \tag{70}$$

Indeed, quotient by the radical and apply Hermitian Gram–Schmidt. A nonzero Hermitian form has a vector of nonzero self-product; otherwise polarization would make the whole form zero. Each nonzero diagonal coefficient lies in $\mathbb{F}_q^*$ and can be normalized to 1 because the norm map $K^* \to \mathbb{F}_q^*$ is surjective. Iteration proves (70). Consequently, for uniform $\mathbf{b}$, the distribution of $\mathbf{b}^\star M\mathbf{b}$ is exactly the distribution of H_r.

Proof (Lemma 8). In the random KCov branch, $\mathbf{a}$ and $\mathbf{b}$ are independent uniform vectors in K^m and independent of ε. Projecting Corollary 1 gives

$$W_P = h\varepsilon - h\mathbf{b}^\star\mathbf{b} + (\mathbf{b}^\star\mathbf{a})(\mathbf{a}^\star\mathbf{b}) = h\varepsilon - \mathbf{b}^\star(hI_m - \mathbf{a}\mathbf{a}^\star)\mathbf{b},$$

which is (48).

Put $c = \mathbf{a}^\star\mathbf{a}$, so $h = 1 + c$. The rank-one determinant identity gives

$$\det(hI_m - \mathbf{a}\mathbf{a}^\star) = h^{m-1}(h - c) = h^{m-1}.$$

Thus $h \neq 0$ implies rank m. If $h = 0$, then $M_{\mathbf{a}} = -\mathbf{a}\mathbf{a}^\star$. The relation $c = -1$ implies $\mathbf{a} \neq 0$, so this matrix has rank one. Since $c = H_m(\mathbf{a})$ and $-1 \in \mathbb{F}_q^*$, Lemma 7 gives

$$\Pr[h = 0] = \Pr[H_m(\mathbf{a}) = -1] = B_m.$$

Now fix $\varepsilon \neq 0$. On $h = 0$, the target $h\varepsilon$ is zero and the rank-one Hermitian form vanishes with probability A_1. On $h \neq 0$, the matrix has full rank and the target $h\varepsilon$ is nonzero; by (70) and Lemma 7, the hit probability is B_m. Averaging over the two cases gives

$$P^{\mathrm{H}}_{\mathsf{Rand}} = B_m A_1 + (1 - B_m)B_m.$$

Subtracting A_1 yields

$$\Delta^{\mathrm{H}}_m = (1 - B_m)(B_m - A_1).$$

For every $m \geq 1$, equation (46) gives $B_m = q^{-1} + O(q^{-2})$, while $A_1 = q^{-2}$, and therefore $\Delta^{\mathrm{H}}_m = q^{-1}(1 + O(q^{-1}))$ uniformly in m.

Proof (Theorem 4). Let $\mathcal{A}_P$ denote the public event $E_P \neq 0$. The hidden variables and hence E_P have the same distribution in the real and random KCov experiments, so $\Pr[\mathcal{A}_P] = 1 - \kappa_P$ in both branches. Conditional on any fixed hidden tuple in $\mathcal{A}_P$, Lemma 6 gives real zero probability at most A_1, whereas Lemma 8 gives the exact random zero probability $P^{\mathrm{H}}_{\mathsf{Rand}}$. Therefore

$$\begin{aligned}\Pr[\mathcal{D}^{\mathrm{H}}_P = 1 \mid \mathsf{Rand}] - \Pr[\mathcal{D}^{\mathrm{H}}_P = 1 \mid \mathsf{Real}] \\ \geq (1 - \kappa_P)(P^{\mathrm{H}}_{\mathsf{Rand}} - A_1) = (1 - \kappa_P)\Delta^{\mathrm{H}}_m.\end{aligned}$$

The asymptotic statement follows from (51). The distinguisher computes W using Corollary 1, projects two public ring elements to one quadratic factor, and performs a constant number of field operations, so it is polynomial-time.

Proof (Lemma 9). Fix a self-reciprocal quadratic factor P and identify its residue field with $K = \mathbb{F}_{q^2}$. Condition on every hidden coefficient except the constant coefficient Z of the secret polynomial s. Then

$$s_P = Z + c$$

for a fixed $c \in K$, while all projected errors are fixed. Because $Z \in \mathbb{F}_q$ and conjugation is Frobenius,

$$\begin{aligned} E_P &= \mathrm{N}(Z + c) + \sum_{i=1}^{m} \mathrm{N}(e_{i,P}) \\ &= Z^2 + Z(c + c^q) + cc^q + C, \end{aligned}$$

where $C \in \mathbb{F}_q$ is fixed by the conditioning. Hence $E_P = 0$ is a monic quadratic equation in the residue class of Z modulo q, and has at most two solutions in $\mathbb{F}_q$.

Let

$$\eta = \max_{r \in \mathbb{F}_q} \Pr[Z \equiv r \pmod{q}].$$

The preceding conditioning gives $\kappa_P \le 2\eta$. For $Z \leftarrow D_{\mathbb{Z},\sigma}$, exactly as in the one-coordinate calculation of Appendix C,

$$\eta \le \frac{\rho_{\sigma/q}(\mathbb{Z})}{\rho_\sigma(\mathbb{Z})} \le \sigma^{-1} + q^{-1}.$$

Since probabilities are at most one, $\eta \le h_\sigma = \min\{1, \sigma^{-1} + q^{-1}\}$, which proves (55).

Proof (Theorem 5). Let $\mathcal{P}$ be the $L = n/2$ self-reciprocal factors and condition on the complete hidden tuple $(s, \mathbf{e})$. This fixes the active subset

$$\mathcal{A} = \{P \in \mathcal{P} : E_P \neq 0\}, \qquad T := |\mathcal{A}|.$$

The CRT coordinates of the public uniform vectors are independent across distinct factors.

In the random experiment, Lemma 8 shows that each active zero event has probability $p := P^{\mathrm{H}}_{\mathsf{Rand}}$, so the conditional probability that at least one active factor vanishes is exactly

$$1 - (1 - p)^T.$$

In the real experiment, Lemma 6 gives a conditional zero probability $r_P \le A_1$ on each active factor. Conditional CRT independence therefore gives

$$\Pr[\text{global test} = 1 \mid s, \mathbf{e}, \mathsf{Real}] = 1 - \prod_{P \in \mathcal{A}} (1 - r_P) \le 1 - (1 - A_1)^T.$$

Hence the conditional signed random-minus-real gap is at least

$$\begin{aligned} (1 - A_1)^T - (1 - p)^T &= (p - A_1) \sum_{j=0}^{T-1} (1 - A_1)^{T-1-j} (1 - p)^j \\ &\ge T \Delta^{\mathrm{H}}_m (1 - p)^{T-1}. \end{aligned} \tag{71}$$

For $T = 0$ the bound is interpreted as zero.

It remains to lower-bound the last factor. First suppose $q = 3$. The congruence $q \equiv -1 \pmod{2n}$ and $n \geq 2$ force $2n = 4$, hence $L = 1$ and $(1-p)^{T-1} = 1$ whenever $T = 1$. Now suppose $q \geq 7$. From (46) and $A_1 = q^{-2}$,

$$p = B_m A_1 + (1 - B_m)B_m \leq B_m + A_1 \leq q^{-1} + 2q^{-2}.$$

Moreover $2n \mid q+1$ gives $T \leq L = n/2 \leq (q+1)/4$. Therefore

$$(T-1)p \leq \frac{(q+1)(q+2)}{4q^2} \leq \frac{1}{2},$$

where the last inequality holds for $q \geq 7$. Bernoulli's inequality yields

$$(1-p)^{T-1} \geq 1 - (T-1)p \geq \frac{1}{2}.$$

Thus (71) is always at least $T\Delta_m^{\mathrm{H}}/2$.

Finally average over the hidden tuple. By definition of κ,

$$\mathbb{E}[T] = \sum_{P \in \mathcal{P}} \Pr[E_P \neq 0] \geq L(1-\kappa) = \frac{n}{2}(1-\kappa).$$

Hence

$$\mathrm{Adv} \geq \frac{1}{2}\mathbb{E}[T]\Delta_m^{\mathrm{H}} \geq \frac{n}{4}(1-\kappa)\Delta_m^{\mathrm{H}},$$

which is (57). If $\kappa \leq 1 - \delta$, then (51) gives $\mathrm{Adv} = \Omega_\delta(n/q)$. No step assumes independence among the hidden projections; all cross-factor independence comes from the public uniform CRT coordinates after conditioning on the hidden tuple.

E Prime-Modulus Classification Details

This appendix records the short synthesis proof for Theorem 6 and the elementary order formula used in Section 6.1. No additional cryptanalytic ingredient is introduced here.

Proof (Theorem 6). By Corollary 2, an irreducible factor of $X^n + 1$ modulo q is self-reciprocal if and only if $q \equiv -1 \pmod{2n}$. In that case every irreducible factor is self-reciprocal quadratic, and Theorem 5 gives $\mathrm{Adv} = \Omega_\delta(n/q)$ under the stated activity hypothesis.

Suppose instead that $q \not\equiv -1 \pmod{2n}$. Lemma 1 then implies that no irreducible factor is fixed by reciprocity. Since reciprocity is an involution on the set of irreducible factors, all factors therefore partition into distinct reciprocal pairs. The hypotheses of Theorem 3 are exactly the stated condition $m \geq 2$ together with $\rho^m \leq 1/2 - \delta$, and that theorem gives

$$\mathrm{Adv} = \Omega_\delta\left(\frac{n}{fq^f}\right).$$

If $q \equiv 1 \pmod{2n}$ then $f = \mathrm{ord}_{2n}(q) = 1$, yielding the displayed $\Omega_\delta(n/q)$ specialization. The two cases are exhaustive because Corollary 2 characterizes the only possible fixed-factor class.

Proof of Corollary 3. Part (1) is Theorem 5 specialized to $m = 1$. Part (2) is Theorem 2. No further amplification claim is used.

Explicit order formula. Write $2n = 2^t$ and let $\nu_2(x)$ denote the 2-adic valuation of x, i.e., the largest r for which $2^r \mid x$. For every odd prime q,

$$f = \mathrm{ord}_{2^t}(q) = \begin{cases} 2^{\max\{0,\, t-\nu_2(q-1)\}}, & q \equiv 1 \pmod 4, \\ 2^{\max\{1,\, t-\nu_2(q+1)\}}, & q \equiv 3 \pmod 4. \end{cases} \tag{72}$$

If $q \equiv 1 \pmod 4$, the lifting-the-exponent identity gives, for every $j \geq 0$,

$$\nu_2(q^{2^j} - 1) = \nu_2(q-1) + j.$$

Hence the least power 2^j for which $q^{2^j} \equiv 1 \pmod{2^t}$ is $2^{\max\{0, t-\nu_2(q-1)\}}$.

If $q \equiv 3 \pmod 4$, then for every $j \geq 1$,

$$\nu_2(q^{2^j} - 1) = \nu_2(q-1) + \nu_2(q+1) + j - 1 = \nu_2(q+1) + j,$$

because $\nu_2(q-1) = 1$. The order is even, and the least 2^j with $j \geq 1$ for which the valuation reaches t is therefore $2^{\max\{1, t-\nu_2(q+1)\}}$. This proves (72).

F Construction-Level Consequence Proofs

This appendix gives the short formal proofs for the consequences stated in Section 7. No additional cryptanalytic assumption is introduced here.

Proof of Corollary 5. The decisional KCov-RLWE hardness requirement quantifies over all probabilistic polynomial-time distinguishers and requires their advantage to be negligible. Theorem 6 and Corollary 3 exhibit explicit polynomial-time distinguishers in their respective parameter regimes. Whenever the corresponding lower bound is non-negligible, that distinguisher violates the decisional hardness requirement. No additional reduction is needed.

Proof (Proposition 2). Let $e_0 = (1, 0, \ldots, 0)^T \in \mathbb{Z}^n$. By definition of the power-basis multiplication matrix,

$$M_{E^{\mathrm{int}}} e_0 = \mathrm{coeff}(E^{\mathrm{int}}),$$

the column vector of the n integral coefficients of E^{int}. From (64),

$$\frac{\lambda e_0 - F e_0}{\sigma_r^2} = M_{E^{\mathrm{int}}} e_0 = \mathrm{coeff}(E^{\mathrm{int}}).$$

Thus an exact efficiently manipulable representation of the public quantities recovers the coefficient vector of E^{int} using $O(n)$ entry operations plus one public scaling operation per coefficient. Coefficientwise reduction then gives the public residue $E = [E^{\mathrm{int}}]_q$.

For the approximate case, define

$$y := \frac{\lambda e_0 - \widehat{F} e_0}{\sigma_r^2}.$$

For every coordinate j, condition (65) gives

$$|y_j - \mathrm{coeff}(E^{\mathrm{int}})_j| < \frac{1}{2}.$$

Since the target coefficient is integral, nearest-integer rounding recovers it uniquely. A certified error bound and polynomial-time arithmetic in the public encoding therefore give polynomial-time recovery in the approximate case as well.

Proof (Proposition 3). Let $\mathcal{D}$ be a KCov-RLWE distinguisher that reads the covariance only through the reduced value E. Construct $\mathcal{D}'$ on input $(\mathbf{a}, \mathbf{c}, F, \lambda)$ by computing $E' \leftarrow \mathsf{Rec}(F, \lambda)$ and returning $\mathcal{D}(\mathbf{a}, \mathbf{c}, E')$.

Under $\mathsf{KeyReal}_m$, exact recovery gives $E' = E = [E^{\mathrm{int}}]_q$ and $\mathbf{c} = \mathbf{a}s + \mathbf{e}$, with the same joint hidden sampling used to generate E. Hence the reconstructed triple has exactly the KCov real law. Under $\mathsf{KeyUnif}_m$, $\mathbf{c} = \mathbf{u}$ is independent uniform while the hidden tuple and covariance descriptor are sampled identically to the real hybrid, so the reconstructed triple has exactly the KCov random law. Therefore the two acceptance probabilities, and hence the distinguishing advantage, are identical to those of $\mathcal{D}$.

If recovery can fail, couple each hybrid to the corresponding exact-recovery experiment. The acceptance bit can differ only on the recovery failure event, so each acceptance probability changes by at most the respective failure probability. The triangle inequality then lowers the transferred absolute gap by at most $\varepsilon_{\mathsf{Real}} + \varepsilon_{\mathsf{Rand}}$.