NAIBI: Binding Reconciliation KEMs and Ephemeral Key Agreement over Non-Split Commutative Algebras
(Full Version)

Sidoine Djimnaibeye[1], Djiby Sow[2], Mahamat Borgou Hassan[3], Daniel Tieudjo[4], and Ganga Tchawa[5]

[1] ChadUP, N'Djamena, Chad
d.sidoine@chad-up.com
[2] Department of Mathematics and Computer Science, Université Cheikh Anta Diop (UCAD), Dakar, Senegal
dsow@ucad.edu.sn
[3] Department of Mathematics and Computer Science, ENASTIC, N'Djamena, Chad
mahamatborgou@gmail.com
[4] Department of Mathematics and Computer Science, University of Ngaoundéré, Ngaoundéré, Cameroon
tieudjo@yahoo.com
[5] National Agency for Secure Documents (ANATS), N'Djamena, Chad
tchawa8ganga@gmail.com

Abstract. We propose **NAIBI-Full**, a lattice-based key encapsulation mechanism (KEM) together with its forward-secure ephemeral key-agreement protocols, built on the regular representation ρ of the non-split commutative algebra $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ over $R_q = \mathbb{Z}_q[x]/(x^n + 1)$, $k \in \{2, 3\}$, α a non-k-th power. Each party publishes the full matrix $\mathbf{t} = A\rho(\mathbf{s}) + \mathbf{e} \in R_q^{k\times k}$; because $\rho(\mathcal{A}_\alpha)$ is commutative, the cross-product collapses to small noise and a Peikert hint closes the gap to exact agreement, even though the public matrix A is fully generic in $M_k(R_q)$.
Hardness rests on a single, well-localised assumption: structured-secret Module-LWE MLWE^ρ, which we identify exactly with a $\rho(y)$-linked k-sample MLWE problem via column decomposition, placing it inside the well-cryptanalysed MLWE landscape of ML-KEM. NAIBI-Full is the *conservative* member of the family: a clean account in terms of a standard lattice assumption, at the cost of k^2-element public keys and ciphertexts. Three primitives follow from this single core: an IND-CCA2 KEM ($\mathrm{FO}^\not\perp$, i.e. with implicit rejection, in the ROM and QROM) and two forward-secure ephemeral protocols, ephemeral-static and ephemeral-ephemeral. For the KEM we prove in addition a statistical, decapsulation-level *binding correctness* guarantee, with collision probability $\leq (2/3 + \frac{1}{3q})^{\lceil n/2 \rceil} + (8/q)^{n/2} + 2^{-256}$, below 2^{-148} at every parameter set. That binding survives in the malicious-key model on the public-key axis ($\mathsf{MAL\text{-}BIND\text{-}K\text{-}PK}$, with a q_H factor on the mechanism term), with no distributional assumption on the adversarial keys: the property ML-KEM is known to lack, and which the seed key format of FIPS 203 — which

does restore the ciphertext axis — still leaves unattained. The statistical modality attaches to the accept branch; on the rejection branch, where no KEM admits a statistical guarantee, the rejection key hashes the public key, without which the notion falls to a one-line attack reusing a single z across two malicious keys. We deliberately offer no static-static mode, since it would inherit the active key-mismatch attacks of the Ding/Peikert/NewHope family; NAIBI-Full is confined to its key-mismatch-resistant deployments. Parameter sets cover NIST security Categories 1, 3 and 5, all with $\delta \leq 2^{-128}$.

Keywords: Post-quantum cryptography · Module-LWE · Reconciliation-based key exchange · Key encapsulation · Non-split algebra · Binding correctness · MAL-BIND-K-PK

This is the full version. The conference versions defer to this document, which contains in particular: *(i)* the complete correctness proof (Appendix A), including the reconciliation-threshold argument (Lemma 2) and the exact-uniformity lemma (Lemma 1); *(ii)* the full game sequences for IND-CPA of NAIBI-PKE-OTP (Theorem 6), the FO$^{\not\perp}$ transform to IND-CCA2 in the ROM and QROM (Theorems 7, 8), and the key-indistinguishability proofs of the ephemeral–static and ephemeral–ephemeral agreement modes (Theorems 4, 5); *(iii)* the syntactic equivalence between NAIBI-PKE-OTP and one leg of NAIBI-DH (Appendix B); *(iv)* the detailed noise-budget and per-mode parameter tables and the uncompressed hybrid deployment options NAIBI-{I,III,V}-pc (Section 7); *(v)* the per-mode bandwidth, message-by-message key-exchange tables and TLS-integration figures (Section 9); *(vi)* the empirical validation of the slot heuristic (Appendix D), of which the conference versions retain only a three-line summary; *(vii)* the numerical sensitivity analysis of δ with respect to the fresh reconciliation error and the threshold shift (Appendix C); and *(viii)* the complete study of alternative reconciliation mechanisms (Section 6), of which the conference versions retain only the classification table and the degradation theorem. Results are stated once here in full.

1 Introduction

Reconciliation-based lattice key exchange. The starting point of this work is the family of *reconciliation-based lattice key exchange* protocols initiated by Ding, Xie and Lin in 2012 [3], refined by Peikert in 2014 [29], and deployed in BCNS [30] and NewHope [2]. In these schemes the two parties exchange MLWE/RLWE samples and reach approximate equality on a shared value; an additional *reconciliation hint* (a few bits per coefficient, sent from one side to the other) closes the gap to exact agreement. After the one-time publication of static public keys, the per-session cost is a single message, typically the initiator's sample plus reconciliation data. In the strict cryptographic sense these are therefore *one-pass* schemes, though not *non-interactive* ones: the 0-pass primitive of classical Diffie–Hellman [1], where the shared secret is recovered from public keys alone, is a different object.

Truly non-interactive lattice key agreement remained elusive until very recently. CSIDH [5] achieves 0-pass post-quantum key agreement from isogeny commutative group actions but has suffered substantial classical attacks [6] and its concrete security is unsettled. SWOOSH [31], presented at USENIX Security 2024, is the first lattice-based protocol that achieves *true* non-interactivity over MLWE, at the cost of high-dimensional lattices and large key sizes. The present work does *not* attempt 0-pass non-interactivity. It stays firmly inside the one-pass reconciliation-based family; what it contributes is algebraic and structural, not the removal of the reconciliation message.

This paper. We introduce **NAIBI-Full**, an algebraic recasting of reconciliation-based lattice key agreement built on the non-split commutative algebra $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ with $\alpha \in \mathbb{Z}_q^*$ a non-k-th-power. Each party publishes the *full matrix* $\mathbf{t} = A\,\rho(\mathbf{s}) + \mathbf{e} \in R_q^{k\times k}$, where the public matrix $A \in M_k(R_q)$ is generic and the secret $\rho(\mathbf{s})$ lies in the *commutative* subalgebra $\rho(\mathcal{A}_\alpha) \subset M_k(R_q)$ given by the regular representation ρ. The crucial algebraic identity is

$$\mathbf{t}_B\,\rho(\mathbf{s}_A) - \mathbf{t}_A\,\rho(\mathbf{s}_B) = \mathbf{e}_B\,\rho(\mathbf{s}_A) - \mathbf{e}_A\,\rho(\mathbf{s}_B),$$

where the cancellation of the A-terms uses $\rho(\mathbf{s}_B)\rho(\mathbf{s}_A) = \rho(\mathbf{s}_A)\rho(\mathbf{s}_B)$ in $\rho(\mathcal{A}_\alpha)$, even though A is generic in $M_k(R_q)$. A 4-quarter Peikert hint $h \in \{0,1\}^{k^2 n}$ then yields exact agreement with negligible failure probability.

What is new, what is not. Reconciliation-based one-pass key exchange is not new: Ding KE, Peikert KE, BCNS and NewHope are all one-pass after public-key publication. What is new here is the *algebraic packaging*, and the properties it enables:

(1) A *commutative-algebra* formulation in which the public matrix A is fully generic in $M_k(R_q)$ rather than constrained to be symmetric, circulant or otherwise structured. Only the *secrets* are structured (living in $\rho(\mathcal{A}_\alpha)$); errors remain generic.
(2) A unified design in which the ephemeral key agreement (two flavours: ES, EE), the underlying CPA-secure PKE and the IND-CCA2 KEM all share the *same* algebraic core.
(3) A formal *KEM binding correctness* property (Theorem 10): no PPT adversary can craft a single ciphertext that *decapsulates* to the same shared secret under two distinct honest KEM key pairs, with collision probability $\leq (2/3 + \frac{1}{3q})^{\lceil n/2\rceil} + (8/q)^{n/2} + 2^{-256}$ (below 2^{-148} everywhere). We lift it, with no distributional assumption on the adversarial keys, to the malicious-key model on the public-key axis (MAL-BIND-K-PK, Theorem 11) — the notion ML-KEM is known to lack [44] and the seed key format does not restore — at the cost of one specification detail: the rejection key hashes the public key (Remark 20).
(4) A *non-splitness* barrier ($\mathcal{A}_\alpha \cong \prod_i \mathbb{F}_{q^k}$ when α is a non-k-th-power) that systematically obstructs the slot-wise dimension-reduction attacks available in the split case.

The conservative choice. NAIBI-Full is the conservative member of the construction: the public key is the full $k \times k$ matrix $\mathbf{t}$, each column of which is a bona fide rank-k MLWE sample with linked secrets. Hardness rests on structured-secret MLWE MLWE^{ρ}, which we relate as closely as currently possible to *plain MLWE* (Section 8.1); the assumption therefore sits inside the well-cryptanalysed MLWE landscape of Kyber/ML-KEM. Bandwidth is the price of that footing. The matrix-valued public key and ciphertext cost k^2 ring elements where ML-KEM pays k, so NAIBI-Full runs at 2.3–4.3× the ML-KEM bandwidth, bought against the binding guarantees and the clean MLWE reduction.

Caveats. Four limitations should be stated up front. *(a)* A naive static-static reconciliation deployment of this algebraic core would, like every member of the Ding/Peikert/NewHope family, be exposed to active key-mismatch attacks under long-term key reuse: Fluhrer [33], Ding–Alsayigh–Saraswathy–Fluhrer–Lin [34], Ding–Fluhrer–Saraswathy [35] and Qin–Cheng–Ding [36] recover the secret in a few sessions whenever the responder behaviour leaks bits via signal, plaintext-check or final-key-match oracles. **We therefore do not offer a static-static mode at all**; NAIBI-Full is deployed exclusively through the FO-transformed KEM and the ephemeral agreement variants, in which these attacks provably do not apply. Section 8.4 sets out this rationale in full. *(b)* The lightest parameter set (NAIBI-I) is set at $k = 2$, $n = 512$ ($kn = 1024$), the same module-lattice dimension as NAIBI-III; under the full `lattice-estimator` it reaches ≈ 192 classical core-SVP bits with a dual-hybrid cost of 2^{203} (the single authoritative Table 8), clearing the Category 1 floor with wide margin. We flag only that, like every Category-1 lattice set, its *absolute* margin over the floor is the smallest in the family. *(c)* The hardness of the scheme rests on the *standing assumption* that the public $\rho(y)$-link between secret columns confers no attacker advantage: Theorem 13 is an exact identification with a linked k-sample MLWE problem, *not* a reduction to independent-sample MLWE, and our concrete-security estimates (Section 8.2) assume the link is not exploitable. *(d)* Binding correctness is a property of the *KEM*, not of the bare PKE one-time pad (which a single null ciphertext defeats); it is delivered by the FO re-encryption check and implicit rejection (Remark 22).

Relation to truly non-interactive lattice DH. SWOOSH [31] is the relevant point of comparison for actual non-interactivity (0-pass). It achieves true NIKE from M-LWE using a lattice dimension of 8192 and a 214-bit modulus, with public keys and shared secrets in the megabyte range. NAIBI-Full does *not* remove the per-session reconciliation message and is not a NIKE: it sits in the much more bandwidth-efficient one-pass KEM family (plus one-pass ES-DH and two-pass EE-DH ephemeral key agreement), paying the cost of one per-session message rather than removing it.

The algebraic correspondence we exploit is summarised below, with the appropriate caveats:

Classical DH (group $G = \langle g \rangle$, 0-pass)	*NAIBI (algebra $\mathcal{A}_\alpha$ with rep. ρ, 1-pass)*
$\mathsf{pk}_P = g^{s_P}$ $\quad (s_P \in \mathbb{Z}$ secret)	$\mathsf{pk}_P = (\, A,\ \mathbf{t}_P = A\,\rho(\mathbf{s}_P) + \mathbf{e}_P \in R_q^{k\times k}\,)$
$\mathsf{sk}_P = s_P$	$\mathsf{sk}_P = \mathbf{s}_P$ $\quad$ (small / ternary, in $\mathcal{A}_\alpha$)
shared $= g^{s_A s_B} = (g^{s_A})^{s_B} = (g^{s_B})^{s_A}$	shared $\approx \mathbf{t}_B\,\rho(\mathbf{s}_A) = \mathbf{t}_A\,\rho(\mathbf{s}_B)$ (up to small noise)
deterministic (group equality)	**deterministic after reconciliation hint**
0 extra messages (truly non-interactive)	**1 extra message** ($k^2 n$-bit hint from one party)

Why "non-split" commutative algebra. The construction lives over the non-split commutative algebra $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ with $\alpha \in \mathbb{Z}_q^*$ a non-k-th-power; this guarantees $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$, ruling out the slot-wise NTRU-style attacks that would plague the split case. The regular representation $\rho : \mathcal{A}_\alpha \to M_k(R_q)$ produces a *commutative* family of matrices; commutativity is what makes the A-terms cancel (yielding correctness of the DH agreement) and what ensures that *the same* algebraic structure simultaneously supports a KEM with strong binding correctness.

KEM as a by-product. Once the agreement core is in place, a KEM follows almost for free. In the spirit of hashed ElGamal [20], encapsulation under a long-term public key $\mathsf{pk}_B = (A, \mathbf{t}_B)$ amounts to running "one leg" of the DH protocol with an ephemeral $(\mathbf{s}', \mathbf{e}')$, sending the corresponding $\mathbf{u} = A\,\rho(\mathbf{s}') + \mathbf{e}' \in R_q^{k\times k}$, deriving a *one-time pad* $K = \mathsf{H}(\mathsf{PExt}(\mathbf{t}_B\,\rho(\mathbf{s}'))[k])$ from the Peikert-extracted shared key bit, and **XOR ing** the chosen plaintext against K to form the second component of the ciphertext. We apply the standard FO$\not\perp$ transformation [11] on top of this CPA-secure "OTP-PKE" to obtain *IND-CCA2-secure NAIBI-KEM*. Correctness, parameter sizes and hardness all transfer one-to-one from NAIBI-DH; in particular the negligible-δ guarantee is inherited.

Main contributions.

(1) **An algebraic framework for lattice key encapsulation and ephemeral key agreement.** We recast reconciliation-based lattice key exchange in the language of non-split commutative algebras. The construction places the public matrix A fully generic in $M_k(R_q)$ rather than constraining it to the structure of the algebra: all structure is confined to the secret. This yields a uniform algebraic skeleton in which ephemeral key agreement (ES, EE flavours), CPA-secure PKE and IND-CCA2 KEM all share the same 4-quarter Peikert reconciliation kernel applied to $\mathbf{t}\rho(\mathbf{s})$.

(2) **Algebraic noise analysis with generic errors.** We prove (Lemma 3) that the worst-entry ℓ_∞-norm of $\mathbf{e}_B\,\rho(\mathbf{s}_A) - \mathbf{e}_A\,\rho(\mathbf{s}_B)$ is bounded by $2\alpha k n \eta^2$, with an exact characterisation of the $\mathbf{u}$-compression amplification. The non-trivial element is that the errors $\mathbf{e}_A, \mathbf{e}_B$ may be generic in $R_q^{k\times k}$ (no structure required), which simplifies the security reductions and design.

(3) **Structured-secret MLWE as an explicit assumption.** We introduce *structured-secret MLWE* $\mathsf{MLWE}^{\rho}_{n,q,\chi}$, in which A is uniform in $M_k(R_q)$ and $\mathbf{e}$ is generic, but the secret $\rho(\mathbf{s}) \in \rho(\mathcal{A}_\alpha)$ is constrained to the commutative subalgebra. We exhibit a column-decomposition that rewrites a structured-secret sample as a pair of standard rank-k MLWE samples on the same A with *algebraically linked* secret columns ($\mathbf{s}^{(1)} = \rho(y)\mathbf{s}^{(0)}$ for $k = 2$, etc.) and independent generic errors. We are explicit that this is *not* an equivalence with independent-secret 2-sample MLWE, and that $\mathsf{MLWE}^{\rho}_{n,q,\chi}$ must be treated as a structured variant whose hardness rests on the conjecture that the public algebraic link between secret columns confers no advantage on a lattice attacker beyond the same-A multi-sample reuse already absorbed by standard MLWE analyses (Theorem 13).

(4) **IND-CCA2 NAIBI-KEM with binding correctness.** Apply FO$\not\perp$ [11] on top of a CPA-secure OTP-PKE that is exactly "one leg of the agreement protocol plus a one-time pad" (Section 5) to obtain an IND-CCA2 KEM. Optimised parameter sets at NIST security Categories 1, 3 and 5 (served by NAIBI-I/III/V), together with the optional degree-3 Category-5 set NAIBI-V′, achieve negligible δ by a statistical noise argument. Bandwidth at Category 1: $|\mathsf{pk}| = 2\,848$ B, $|\mathsf{ct}| = 3\,104$ B.

(5) **KEM binding correctness, up to the malicious-key model.** We formalise and prove (Theorem 10, with Corollary 4) a *ε-binding correctness* property *for the KEM* with $\varepsilon \le (2/3 + \frac{1}{3q})^{\lceil n/2 \rceil} + (8/q)^{n/2} + 2^{-256}$: no PPT adversary can craft a single ciphertext that *decapsulates* to the same shared secret under two distinct honestly-generated KEM key pairs. It is a statistical, *decapsulation-level* collision-freeness in the sense of robust PKE [41] (HON-BIND-K-PK in the X-BIND taxonomy of [43]), following from the slot-wise non-degeneracy of $\mathcal{A}_\alpha$ on honest secrets together with the FO re-encryption check and implicit rejection (the bare one-time pad does not itself bind, Remark 22). On the K-PK axis we then lift this to the *malicious-key* model. Theorem 11 establishes MAL-BIND-K-PK (Section 5.5) with no distributional assumption on the adversarial keys: the honest-distribution heuristic gives way to two assumption-free lemmas, a minimum-NTT-weight bound (Lemma 10) and a worst-case per-slot collision bound (Lemma 11), the latter established by exhaustive evaluation at the deployed moduli. This is the notion Schmieg [44] proves ML-KEM *lacks* and that the FIPS 203 seed key format does not restore, and the sharpest separation NAIBI-Full offers from the standard. It comes with one design obligation, on the branch where geometry has no purchase: the implicit-rejection key must hash the public key, or a single z reused across two malicious keys wins the game outright (Remark 20).

Why pay for binding. NAIBI-Full's bandwidth premium over ML-KEM (2.3–4.3×) buys a property with documented protocol-level value. Cremers–Dax–Medinger [43] exhibit, by automated analysis, attacks on KEM-based protocols enabled precisely by the absence of binding (ciphertext re-use across recipients, key-confusion across sessions); Grubbs–Maram–Paterson [41] show that robustness/collision-freeness is the enabling property for anonymous

and auditable encryption, which implicit-rejection lattice KEMs do not provide off the shelf. Collision-freeness of FO-based lattice KEMs is not in itself new: [41] and Xagawa [42] establish it generically for NTRU, BIKE, FrodoKEM and HQC, but *computationally, in the QROM*, and on the robustness (CT-PK) axis. What NAIBI-Full adds is different in kind. On the K-PK axis its collision-freeness is *statistical, and holds in the malicious-key model with no distributional assumption on the keys* (Theorem 11): on the accept branch it withstands adversaries bounded only in random-oracle queries, and does not rest on hashing the transcript. That is the regime Schmieg [44] shows ML-KEM does not reach. Deployments that bind keys to transcripts anyway may well prefer ML-KEM, and we say so; our target is the case where the transcript hash is unavailable, unauthenticated, or must survive compromise of the hash inputs.

(6) **A principled exclusion of the static-static mode.** We show (Section 8.4) that any static-static reconciliation deployment of this core inherits the active key-mismatch attacks of [33,34,35,36,38,37] — recovering the long-term secret in only a handful of sessions — and that the non-split barrier, which protects the NTT domain, gives no protection in the coefficient domain where these attacks operate. Rather than patch the mode with key-rotation bounds or NIZK proofs of well-formedness (which would negate its bandwidth advantage), we adopt the pragmatic stance of *not offering it*: NAIBI-Full is deployed only through the IND-CCA2 KEM (immune by the FO re-encryption check) and the ephemeral ES/EE agreement modes (immune because each session uses a fresh secret).

2 Preliminaries

Notation. $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ with representatives in $(-q/2, q/2]$. $\lfloor \cdot \rceil$ denotes rounding to the nearest integer. For $x \in \mathbb{Z}_q$ and integer $d \geq 1$:

$$\mathsf{Compress}(x,d) = \lfloor 2^d x/q \rceil \bmod 2^d, \qquad \mathsf{Decompress}(c,d) = \lfloor q\, c/2^d \rceil.$$

The rounding error satisfies $|x - \mathsf{Decompress}(\mathsf{Compress}(x,d),d)| \leq \lfloor q/2^{d+1} \rfloor =: E_d$. Both operations extend coefficient-wise to polynomials and vectors. Throughout, $\mathsf{Compress}(\cdot, d)$ is the single deterministic rounding primitive used by NAIBI-DH (Section 4); we do not introduce any separate notation for it.

For $f, g \in R_q = \mathbb{Z}_q[x]/(x^n+1)$ with $\|f\|_\infty, \|g\|_\infty \leq B$: $\|fg\|_\infty \leq n\, B^2$. More precisely, for ternary g ($B = 1$): $\|fg\|_\infty \leq n\, \|f\|_\infty$.

Module-LWE. Let n be a power of two, q a prime, $k \geq 1$, and χ an error distribution over R_q. Given $A \leftarrow R_q^{k \times k}$ uniform and $\mathbf{s} \leftarrow \chi^k$, the MLWE distribution is $(A, A\mathbf{s} + \mathbf{e})$ where $\mathbf{e} \leftarrow \chi^k$. The *decisional* MLWE problem $\mathsf{MLWE}_{n,k,q,\chi}$ asks to distinguish this from $(A, \mathbf{u})$ with $\mathbf{u} \leftarrow R_q^k$.

Centred binomial distribution. χ_η samples each coefficient as $\sum_{i=1}^{\eta}(b_i - b_i')$ with $b_i, b_i' \leftarrow \{0,1\}$ i.i.d. Coefficients lie in $\{-\eta, \ldots, \eta\}$ with $\|f\|_\infty \leq \eta$ and $\sigma^2 = \eta/2$. For $\eta = 1$ (ternary): coefficients in $\{-1, 0, 1\}$.

Module Learning With Rounding (MLWR). For integer $p < q$, define the rounding function $\lfloor\cdot\rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$ by $\lfloor x \rceil_p = \lfloor p\,x/q \rceil \bmod p$, extended coefficient-wise. The decisional $\mathsf{MLWR}_{n,k,q,p,\chi}$ problem asks to distinguish $\big(A,\ \lfloor A\mathbf{s}\rceil_p\big)$ from $\big(A,\ \mathbf{u}\big)$ with $\mathbf{u} \leftarrow R_q^k$ projected to $\mathbb{Z}_p^{kn}$, where $A \leftarrow R_q^{k\times k}$ and $\mathbf{s} \leftarrow \chi^k$. Banerjee–Peikert–Rosen [18] prove hardness from (M)LWE only when q/p is *super-polynomial* in the security parameter; the refinements of Alwen et al. [19] and of Bogdanov et al. [47] apply to polynomial moduli but only for an a-priori *bounded* number of samples and with parameter conditions that do not cover the constant-ratio regime $q/p \in [2^3, 2^5]$ relevant to compressed lattice schemes. *No reduction from MLWE to constant-ratio MLWR is known.* Where a rounded object appears in this paper we therefore proceed in one of two ways. *(a)* Whenever the rounded object is a public deterministic function of a sample that is itself pseudorandom, we invoke only MLWE (Remark 1). *(b)* The single remaining rounding-flavoured step — pseudorandomness of products against a *decompressed-uniform* public matrix induced by public-key compression — is stated as an explicit standing assumption (Definition 1), in line with the methodology of Saber [48]; it disappears entirely in the uncompressed parameter option of Remark 2.

Remark 1 (Public post-processing). Let f be any efficiently computable function (e.g. $\mathsf{Compress}(\cdot, d)$). If $(A, \mathbf{t}) \approx_c (A, \mathbf{U})$ under MLWE^ρ, then $(A, f(\mathbf{t})) \approx_c (A, f(\mathbf{U}))$, since any distinguisher of the latter yields one of the former by applying f. Consequently every game hop of the form "replace the published $\mathsf{Compress}(A\rho(s) + \mathbf{e}, d)$ by $\mathsf{Compress}(\mathbf{U}, d)$ with $\mathbf{U}$ uniform" costs exactly $\mathsf{Adv}_{\mathsf{MLWE}^\rho}$ and *requires no rounding assumption.* The simulator publishes $\mathsf{Compress}(\mathbf{U}, d)$, not a uniform string in $\{0, \ldots, 2^d - 1\}^{k^2 n}$: for constant-ratio compression these two distributions are statistically far apart, and our proofs (Theorems 6, 4, 5) never conflate them.

Definition 1 (Coarse-matrix structured MLWE). *For $d \le \lceil \log_2 q \rceil$ let $\mathcal{D}_d$ denote the distribution $\mathsf{Decompress}(\mathsf{Compress}(\mathbf{U}, d), d)$ for $\mathbf{U} \leftarrow M_k(R_q)$ uniform. The problem $\mathsf{MLWE}^{\rho,\mathcal{D}_d}_{n,q,\chi}$ is: given $\mathbf{B} \leftarrow \mathcal{D}_d$, $s \leftarrow \chi(\mathcal{A}_\alpha)$, $\mathbf{e} \leftarrow \chi^{k\times k}$, distinguish $(\mathbf{B}, \mathbf{B}\rho(s) + \mathbf{e})$ from $(\mathbf{B}, \mathbf{U}')$ with $\mathbf{U}'$ uniform. For $d = \lceil \log_2 q \rceil$ one has $\mathcal{D}_d = \mathrm{U}(M_k(R_q))$ and the problem coincides with MLWE^ρ. For $d < \lceil \log_2 q \rceil$ we treat $\mathsf{MLWE}^{\rho,\mathcal{D}_d}$ as a standing assumption: $\mathbf{B}$ retains d bits of entropy per coefficient and the known lattice attacks are insensitive to the coarse support of the public matrix (the attack lattice profile (kn, q, σ) is unchanged), but we know of no reduction from MLWE^ρ, for the same structural reason that public-key compression was dropped from Kyber in Round 2 of the NIST process [49].*

Remark 2 (Provable variant without public-key compression). Setting $d_t = \lceil \log_2 q \rceil$ (no compression of the long-term key $\mathbf{t}$; the ephemeral/ciphertext compression d_u is kept) eliminates $\mathsf{MLWE}^{\rho,\mathcal{D}_{d_t}}$ from all reductions, which then rest on MLWE^ρ alone (Theorems 6, 4, 5). The cost is $k^2 n(\lceil \log_2 q \rceil - d_t)/8$ bytes of public key, i.e. +1024 B at Level I, +512 B at Level III, +2048 B at Level V. We denote these options NAIBI-I/III/V-pc and recommend them when provable

security from MLWE^{ρ} is preferred over bandwidth, mirroring the rationale for which the Kyber team dropped public-key compression in Round 2 [49].

The reconciliation functions (randomised doubling). NAIBI extracts shared key bits by Peikert's *randomised doubling* over $\mathbb{Z}_{2q}$, which makes the extracted key bit *exactly* uniform given the transmitted hint even though q is odd (the naive "4-quarter" partition of $\mathbb{Z}_q$ would leave a $\Theta(1/q)$ per-bit bias that is fatal to the 2^{-256} targets in the security proofs). The two subroutines below are used *identically* by all NAIBI schemes (ES-DH, EE-DH and the KEM) as their shared key-extraction primitives.

Definition 2 (Randomised doubling and reconciliation over $\mathbb{Z}_{2q}$). *Let q be an odd prime. Define $\mathsf{dbl} : \mathbb{Z}_q \to \mathbb{Z}_{2q}$ by $\mathsf{dbl}(v) := 2v - \bar{e}$, where $\bar{e} \in \{-1, 0, 1\}$ is sampled with $\Pr[\bar{e} = 0] = 1/2$ and $\Pr[\bar{e} = \pm 1] = 1/4$. For $\bar{v} \in \mathbb{Z}_{2q}$ (representatives in $[0, 2q)$) define the* key bit *and the* cross-rounding (hint) bit

$$b(\bar{v}) := \lfloor \bar{v}/q \rceil \bmod 2 \in \{0, 1\}, \qquad c(\bar{v}) := \lfloor 2\bar{v}/q \rfloor \bmod 2 \in \{0, 1\}.$$

The reconciliation function $\mathsf{rec} : \mathbb{Z}_{2q} \times \{0, 1\} \to \{0, 1\}$ on input $(\bar{w}, c)$ outputs $b(\bar{v}^*)$, where $\bar{v}^*$ is the point of $\{\bar{x} \in \mathbb{Z}_{2q} : c(\bar{x}) = c\}$ closest to $\bar{w}$ (ties toward 0). The matrix-level primitives are:

- $\mathsf{PExt}(\mathbf{r}; \bar{\mathbf{e}}) \to (h, k) \in \{0, 1\}^{k^2 n} \times \{0, 1\}^{k^2 n}$*: for each of the $k^2 n$ coefficients v of $\mathbf{r} \in R_q^{k \times k}$, set $\bar{v} \leftarrow 2v - \bar{e}_v$ with the dedicated doubling bits $\bar{e}_v$, and output $h^{(v)} := c(\bar{v})$ (the* hint *bit) and $k^{(v)} := b(\bar{v})$ (the* raw key *bit).*
- $\mathsf{PRec}(\mathbf{w}, h) \to k' \in \{0, 1\}^{k^2 n}$*: for each coefficient w with matching hint bit h, output $k'^{(w)} := \mathsf{rec}(2w, h)$.*

In interactive modes the doubling bits $\bar{\mathbf{e}}$ are sampled fresh per session; inside the FO-transformed PKE they are derived deterministically from the encryption coins r (Section 5.1), so encryption stays a deterministic function of $(m, \mathsf{pk}; r)$ and the FO re-encryption check is well defined.

Lemma 1 (Exact uniformity of the key bit given the hint). *Let q be odd, $v \leftarrow \mathrm{U}(\mathbb{Z}_q)$ and $\bar{v} = \mathsf{dbl}(v)$. Then $\bar{v}$ is exactly uniform on $\mathbb{Z}_{2q}$, and $b(\bar{v})$ is exactly uniform on $\{0, 1\}$ and independent of $c(\bar{v})$.*

Proof. For an even target $x \in \mathbb{Z}_{2q}$, $\bar{v} = x$ requires $\bar{e} = 0$ and $v = x/2$: probability $\frac{1}{q} \cdot \frac{1}{2} = \frac{1}{2q}$. For odd x, $\bar{v} = x$ requires $\bar{e} = \pm 1$ with $v = (x \pm 1)/2$: probability $\frac{1}{q} \cdot \frac{1}{4} + \frac{1}{q} \cdot \frac{1}{4} = \frac{1}{2q}$. Hence $\bar{v} \sim \mathrm{U}(\mathbb{Z}_{2q})$. Partition $[0, 2q)$ into the four arcs $Q_0 = [0, q/2)$, $Q_1 = [q/2, q)$, $Q_2 = [q, 3q/2)$, $Q_3 = [3q/2, 2q)$, of integer cardinalities $\lceil q/2 \rceil, \lfloor q/2 \rfloor, \lceil q/2 \rceil, \lfloor q/2 \rfloor$ (equal in antipodal pairs since $2q$ is even and q odd). Here $c(\bar{v})$ is the parity of the arc index and $b(\bar{v})$ its top bit. Conditioned on $c = 0$, $\bar{v}$ is uniform on $Q_0 \cup Q_2$ with $|Q_0| = |Q_2|$ and $b = 0$ on Q_0, $b = 1$ on Q_2; similarly for $c = 1$. Hence $\Pr[b = 0 \mid c] = 1/2$ exactly, for both values of c.

Lemma 2 (Agreement). *Let $v, w \in \mathbb{Z}_q$ with $|v - w \bmod^{\pm} q| \le q/8 - 1$, and $\bar{v} = \mathsf{dbl}(v)$. Then $\mathsf{rec}(2w, c(\bar{v})) = b(\bar{v})$.*

Proof. $|2w - \bar{v} \bmod^{\pm} 2q| = |2(w-v)+\bar{e}| \le 2(q/8-1)+1 < q/4$. Given $c := c(\bar{v})$, the set $\{\bar{x} : c(\bar{x}) = c\}$ is the union of two antipodal arcs of length $q/2$ at mutual distance $q/2$; b is constant on each arc and differs between them. A point within distance $< q/4$ of $\bar{v}$ is strictly closer to the arc containing $\bar{v}$ than to the antipodal one, so rec returns $b(\bar{v})$.

The hint $h \in \{0,1\}^{4n}$ is the per-session reconciliation message of the ephemeral agreement modes. The *agreement guarantee* of these functions — that $\mathsf{PRec}(\mathbf{w}, \mathsf{PExt}(\mathbf{v})[h]) = \mathsf{PExt}(\mathbf{v})[k]$ entry-wise whenever $\|\mathbf{v} - \mathbf{w}\|_\infty \le q/8 - 1$ (Lemma 2), together with exact uniformity of the key bit given the hint (Lemma 1) — is established as Theorem 2.

The multi-sample variant (used only to delimit the excluded mode). A static-static reconciliation deployment of the kernel would leak *multiple* MLWR samples sharing the same secret $\mathbf{s}_B$ across the lifetime of a long-term key. We will appeal to the *multi-sample* hardness assumption $\mathsf{MLWR}^{(T)}$ which is believed to be polynomially equivalent to single-sample MLWR [18,19], with a concrete tightness loss factor of T, only to characterise precisely the regime NAIBI excludes (Section 8.4). In the offered ephemeral-static and ephemeral-ephemeral variants a tight reduction to single-sample MLWR holds.

FO transformation. Given a δ-correct IND-CPA PKE with plaintext space $\mathcal{M}$, the FO$\not\perp$ transformation [11] produces an IND-CCA2-secure KEM in the (Q)ROM with security loss depending on δ. When δ is negligible the loss is cleanest; this is the regime in which all NAIBI parameter sets operate.

3 Algebraic Setting: the Non-Split Algebra $\mathcal{A}_\alpha$

Fix n a power of two, q a prime with $q \equiv 1 \pmod{2n}$, a *degree parameter* $k \in \{2,3\}$, and $\alpha \in \mathbb{Z}_q^*$. The algebra has degree k over R_q and the regular representation lands in $M_k(R_q)$. We require α to be a *non-k-th-power* modulo q so that $y^k - \alpha$ is irreducible:

- $k = 2$: α a *non-square*, $\alpha^{(q-1)/2} \equiv -1 \pmod{q}$;
- $k = 3$: α a *non-cube*, requiring $q \equiv 1 \pmod 3$ and $\alpha^{(q-1)/3} \not\equiv 1 \pmod{q}$.

All parameter sets use $\alpha = 3$ and degree $k = 2$: every NAIBI instance (Levels I, III, V) is realised over $\mathcal{A}_\alpha = R_q[y]/(y^2 - 3)$ and differs only in the ring degree n and the modulus q (see Section 8.2). The degree-3 case is retained in the general exposition below for completeness and as an optional high-security variant (NAIBI-V′, Section 7.2), but is not used by any of the three NGCC parameter sets.

Definition 3 (Non-split commutative algebra). *$\mathcal{A}_\alpha := R_q[y]/(y^k - \alpha)$. Elements are $u = \sum_{i=0}^{k-1} u_i y^i$ with $u_i \in R_q$, and multiplication is reduction modulo $y^k - \alpha$ (so $y^k = \alpha$). For $k = 2$ this is*

$$(u_0 + u_1 y)(v_0 + v_1 y) = (u_0 v_0 + \alpha\, u_1 v_1) + (u_0 v_1 + u_1 v_0)\, y. \tag{1}$$

Since α is a non-k-th-power, $y^k - \alpha$ is irreducible over $\mathbb{F}_q$, so $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$ (a product of degree-k field extensions, non-split*).*

Remark 3 (Why non-split matters). The split algebra $R_q[y]/(y^k - \beta)$ with β a k-th power factors as a product $\prod_j R_q/(y - \zeta_j \sqrt[k]{\beta})$ over the k-th roots, enabling slot-wise attacks that reduce the dimension of any LWE instance by a factor of k. The non-split condition prevents this and is what gives NAIBI-DH its asymptotic security guarantees.

Regular representation. Define $\rho : \mathcal{A}_\alpha \to M_k(R_q)$ as the matrix of multiplication by u in the basis $1, y, \ldots, y^{k-1}$. It is the α-circulant

$$\rho\Big(\textstyle\sum_i u_i y^i\Big)_{r,c} = \begin{cases} u_{r-c} & r \geq c, \\ \alpha\, u_{k+r-c} & r < c, \end{cases} \qquad 0 \leq r, c \leq k-1. \tag{2}$$

Explicitly, for $k = 2$ and $k = 3$ (with $\alpha = 3$ in the right-hand forms):

$$\rho_{k=2}(u_0 + u_1 y) = \begin{pmatrix} u_0 & \alpha\, u_1 \\ u_1 & u_0 \end{pmatrix}, \qquad \rho_{k=3}(u_0 + u_1 y + u_2 y^2) = \begin{pmatrix} u_0 & \alpha u_2 & \alpha u_1 \\ u_1 & u_0 & \alpha u_2 \\ u_2 & u_1 & u_0 \end{pmatrix}.$$

ρ is an injective ring homomorphism. Crucially, since $\mathcal{A}_\alpha$ is commutative, $\rho(u)\,\rho(v) = \rho(v)\,\rho(u)$ for all $u, v \in \mathcal{A}_\alpha$.

Public matrix. $A \leftarrow M_k(R_q)$, sampled uniformly from a 32-byte seed via XOF: $A = \mathsf{XOF}(\mathtt{seed}_A) \in M_k(R_q)$. We impose *no* structural constraint on A: it ranges over the full $k^2 n$-dimensional matrix algebra, *not* over the kn-dimensional subring $\rho(\mathcal{A}_\alpha)$. All NAIBI parties share the same A (it is a system parameter, like the curve generator in ECDH).

The Diffie–Hellman commutativity identity. The single algebraic fact that powers NAIBI-DH is the commutativity of $\rho(\mathcal{A}_\alpha) \subset M_k(R_q)$: for $\mathbf{s}, \mathbf{s}' \in \mathcal{A}_\alpha$,

$$\rho(\mathbf{s})\,\rho(\mathbf{s}') \;=\; \rho(\mathbf{s}\,\mathbf{s}') \;=\; \rho(\mathbf{s}'\,\mathbf{s}) \;=\; \rho(\mathbf{s}')\,\rho(\mathbf{s}) \qquad (\text{in } M_k(R_q)). \tag{3}$$

Note that A itself need *not* commute with anything — the identity acts only on the secret factors $\rho(\mathbf{s}), \rho(\mathbf{s}')$, which live in $\rho(\mathcal{A}_\alpha)$. The error matrices play no role in the cancellation and are left generic in $R_q^{k\times k}$. This is the *exact analogue of $g^{ab} = g^{ba}$ in classical DH*: the operation by which Alice and Bob "raise" the partner's public key to their own secret commutes, even though the base A has no commutativity of its own.

Lemma 3 (Algebraic noise bound). *Let $\mathbf{s}, \mathbf{s}' \in \mathcal{A}_\alpha$ have all coefficients bounded by η, and let $\mathbf{e}, \mathbf{e}' \in R_q^{k\times k}$ be* generic *matrices with all entry coefficients bounded by η. In $M_k(R_q)$ with $\|\cdot\|_\infty$ the entry-wise ℓ_∞ norm:*

$$\big\|\mathbf{e}\,\rho(\mathbf{s}')\big\|_\infty \leq \alpha\, k\, n\eta^2, \qquad \big\|\mathbf{e}\,\rho(\mathbf{s}') - \mathbf{e}'\,\rho(\mathbf{s})\big\|_\infty \leq 2\alpha\, k\, n\eta^2 \;=:\; C_{\mathrm{alg}}(k)\, n\eta^2. \tag{4}$$

Write $C_{\mathrm{alg}}(k) := 2\alpha k$ *for the worst-entry coefficient. For* $\alpha = 3$, $\eta = 1$*:* $C_{\mathrm{alg}}(2) = \mathbf{12}$ *(the* $12n$ *bound for the* $k = 2$ *instances, i.e. all three NGCC sets) and* $C_{\mathrm{alg}}(3) = \mathbf{18}$ *(the* $18n$ *bound for the optional degree-3 variant). Correctness is governed by this worst entry.*

Proof. Each entry of $\mathbf{e}\,\rho(\mathbf{s}')$ is the inner product of a row of $\mathbf{e}$ (k entries, each of ℓ_∞-norm $\le \eta$) and a column of $\rho(\mathbf{s}')$. By (2) every column of $\rho(\mathbf{s}')$ has k entries, each of the form u_i or αu_i with $\|u_i\|_\infty \le \eta$, hence of norm $\le \alpha\eta$. Using $\|fg\|_\infty \le n\|f\|_\infty\|g\|_\infty$ in R_q, each of the k products has norm $\le \alpha n\eta^2$, so the entry has norm $\le k\alpha n\eta^2$. For the difference $\mathbf{e}\,\rho(\mathbf{s}') - \mathbf{e}'\,\rho(\mathbf{s})$ the triangle inequality doubles the bound, giving $2\alpha k n\eta^2$. For $\alpha = 3$, $\eta = 1$: $12n$ ($k = 2$) and $18n$ ($k = 3$).

Remark 4 (Where the structure lives). The public matrix A is genuinely uniform in the full k^2n-dimensional matrix algebra $M_k(R_q)$; the algebraic structure that drives the DH commutativity is confined to the *secrets*, which lie in the abelian image $\rho(\mathcal{A}_\alpha) \subset M_k(R_q)$. The errors $\mathbf{e}_P \in R_q^{k\times k}$ are generic small matrices, carrying no algebraic structure at all. The underlying hardness assumption is therefore plain rank-k MLWE over R_q with a *structured secret* $\rho(\mathbf{s}_P) \in \rho(\mathcal{A}_\alpha)$ and an ordinary generic error (Section 8).

3.1 The Compact Construction

Since the DH product contracts row-wise ($\big(\mathbf{t}_B\rho(\mathbf{s}_A)\big)_{i,:} = (\mathbf{t}_B)_{i,:}\rho(\mathbf{s}_A)$) and the KDF consumes only 256 bits out of the $kn \ge 1024$ row coefficients, we publish only row 0 of $\mathbf{t}$: with $\mathbf{a}^\top = A_{0,:}$, $\hat{\tau}_P = \mathsf{Compress}\big((A\rho(\mathbf{s}_P) + \mathbf{e}_P)_{0,:},\, d_t\big) \in \{0,\dots,2^{d_t}-1\}^{kn}$. Compact DH computes row 0 of the product plus a fresh one-row error (compact form of Remark 7): $\hat{\mathbf{r}}_A = \hat{\tau}'_B\,\rho(\mathbf{s}_A) + \hat{\mathbf{e}}'' \in R_q^{1\times k}$ on the hint-emitting side, $\hat{\mathbf{r}}_B = \hat{\tau}'_A\,\rho(\mathbf{s}_B)$ on the recovering side. Key/hint bits are extracted by $\mathsf{PExt}/\mathsf{PRec}$ (Definition 2) on a fixed subset of $N_{\mathrm{key}} = 256$ row coefficients, so the hint is $\hat{h} \in \{0,1\}^{256}$ (32 B); the KEM uses a one-row ephemeral $\hat{\mathbf{u}} = \mathsf{Compress}\big((A\rho(\mathbf{s}') + \mathbf{e}')_{0,:},\, d_u\big)$ and derandomises $(\hat{\mathbf{e}}'', \bar{\mathbf{e}})$ from the FO coins. KeyGen, $\mathsf{PExt}/\mathsf{PRec}$, FO and KDF are unchanged: the four schemes (kernel, ES-DH, EE-DH, KEM) still share all code paths.

Lemma 4 (Row identity). *With* $\boldsymbol{\varepsilon}_{t,P}$ *the* d_t*-decompression error* ($\|\cdot\|_\infty \le E_{d_t}$ *entry-wise), in* $R_q^{1\times k}$*:* $\hat{\tau}'_B\rho(\mathbf{s}_A) - \hat{\tau}'_A\rho(\mathbf{s}_B) = \big(\mathbf{e}_B\rho(\mathbf{s}_A) - \mathbf{e}_A\rho(\mathbf{s}_B)\big)_{0,:} + \big(\boldsymbol{\varepsilon}_{t,B}\rho(\mathbf{s}_A) - \boldsymbol{\varepsilon}_{t,A}\rho(\mathbf{s}_B)\big)_{0,:}$*. By commutativity of the secrets, the* A*-terms cancel, leaving row* 0 *of the noise cross-product; adding the compression errors gives the stated form. Full proof in [?].*

Theorem 1 (Compact correctness). *Recovery noise has* kn *coefficients of the* same per-coefficient law *as in the full scheme, so a union bound over* $N_{\mathrm{key}} = 256$ *extracted coefficients gives* $\delta_{\mathrm{cpt}} \le 256 \cdot \mathrm{erfc}((q/8-1)/(\sigma\sqrt{2})) \le \frac{1}{k}\delta_{\mathrm{full}}$ *— correctness only* improves*, and is far below* 2^{-128} *at every set of Table* **??***.*

Lemma 5 (Reversal identity: the one-row pairing is the $\mathcal{A}_\alpha$ product). *Define the coordinate reversal* $\mathrm{rev} : \mathcal{A}_\alpha \to \mathcal{A}_\alpha$, $\mathrm{rev}\big(\sum_{m=0}^{k-1} a_m y^m\big) := \sum_{m=0}^{k-1} a_{k-1-m}\, y^m$, *extended to row vectors in* $R_q^{1\times k}$ *coordinate-wise (*$\mathrm{rev}(\mathbf{v})_j := \mathbf{v}_{k-1-j}$*). Then for all* $a, \mathbf{s} \in \mathcal{A}_\alpha$ *(identifying* a *with the row* $\mathbf{a}^\top = (a_0, \ldots, a_{k-1})$*),*

$$\mathrm{rev}\big(\mathbf{a}^\top \rho(\mathbf{s})\big) \;=\; \mathrm{rev}(a) \cdot \mathbf{s} \qquad \text{in } \mathcal{A}_\alpha,$$

where $\cdot$ *is the ring product of* $\mathcal{A}_\alpha$*. Consequently the map* $(\mathbf{a}, \mathbf{b}) \mapsto (\mathrm{rev}(\mathbf{a}), \mathrm{rev}(\mathbf{b}))$ *is an involutive bijection of* $R_q^{1\times k} \times R_q^{1\times k}$ *that sends the* $\mathsf{sRLWE}_{\mathcal{A}_\alpha}$ *distribution* $(\mathbf{a},\, \mathbf{a}^\top\rho(\mathbf{s}) + \mathbf{e}^\top)$ exactly *to the plain single-sample Ring-LWE distribution* $(a',\, a'\cdot \mathbf{s} + e')$ *over* $\mathcal{A}_\alpha$ *(with* a' *uniform, the* same *secret* $\mathbf{s}$*, and* $e' = \mathrm{rev}(\mathbf{e}) \sim \chi^k$*), and uniform pairs to uniform pairs.*

Corollary 1 (Exact equivalence with Ring-LWE over $\mathcal{A}_\alpha$; implication from the full-matrix assumption). (i) $\mathsf{sRLWE}_{\mathcal{A}_\alpha, q, \chi}$ *is tightly equivalent (identical advantages, in both directions) to the single-sample decisional Ring-LWE problem over* $\mathcal{A}_\alpha$ *with secret* $\mathbf{s} \leftarrow \chi(\mathcal{A}_\alpha)$ *and coordinate-wise error* χ*; the same holds between* sRLWR *and single-sample Ring-LWR over* $\mathcal{A}_\alpha$*, and between the* T*-sample / two-block variants on each side.* (ii) MLWE^ρ *is, row by row, exactly the* k*-sample form of the same problem: each row* $\mathbf{t}_{i,:} = A_{i,:}\rho(\mathbf{s}) + \mathbf{e}_{i,:}$ *is an independent-*$\mathbf{a}$ *sample with the same secret. Hence any adversary against* $\mathsf{sRLWE}_{\mathcal{A}_\alpha}$ *yields an adversary against* MLWE^ρ *with the* same *advantage, by restriction to row* 0*: hardness of* MLWE^ρ *implies hardness of* $\mathsf{sRLWE}_{\mathcal{A}_\alpha}$*. The one-row scheme therefore rests on an assumption* implied by *that of a full-matrix instantiation of the same core; the dial of this paper trades bandwidth against the number of published samples of a single underlying problem — Ring-LWE over* $\mathcal{A}_\alpha$ *— not against a structurally new pairing.*

Remark 5 (What is genuinely new: the ring, not the pairing). By Corollary 1 the object of study is single-sample decisional Ring-LWE over $\mathcal{A}_\alpha$, the mod-q reduction of the order $\mathcal{O} = \mathbb{Z}[\zeta_{2n}, \sqrt[k]{\alpha}]$ of $K = \mathbb{Q}(\zeta_{2n}, \sqrt[k]{\alpha})$, of degree kn. For the deployed case ($k = 2$, $\alpha = 3$), $K = \mathbb{Q}(\zeta_{2n}, \sqrt{3})$ is *abelian* and contained in $\mathbb{Q}(\zeta_{6n})$, since $\sqrt{3} = \zeta_{12} + \zeta_{12}^{-1}$. *(a)* The assumption is thus a standard-form Ring-LWE instance over an order in an abelian subfield of a cyclotomic — the natural target for the Order-LWE worst-case theory of [45]; whether $\mathcal{O}$ is maximal away from 2 and 3, and how the canonical embedding of $(1, \sqrt{3})$ distorts the error, are the concrete open questions. *(b)* K has proper subfields (e.g. $\mathbb{Q}(\zeta_{2n})$, index 2), so norm/trace maps exist over the number field even though the mod-q slots $\mathbb{F}_{q^k}$ admit none (Remark **??**); the trace maps one sample over $\mathcal{A}_\alpha$ to one over R_q with doubled error and halved information, and we have found no way to exploit this against ternary errors — an avenue future cryptanalysis should address.

One problem, two instantiations. The identities above organise the whole design:

$$\underbrace{\text{Ring-LWE over } \mathcal{A}_\alpha}_{1 \text{ sample}} \overset{\text{Lem. 5}}{\equiv} \mathsf{sRLWE}_{\mathcal{A}_\alpha} \overset{\text{Cor. 1}}{\Longleftarrow} \underbrace{\mathsf{MLWE}^\rho}_{k \text{ samples}} \Longrightarrow \text{IND-CPA} \overset{\text{FO}}{\Longrightarrow} \text{IND-CCA2},$$

while the binding guarantee of Section 5.3 is statistical and assumption-free. All schemes below are parametrised by the number $\ell \in \{1, k\}$ of published rows of $\mathbf{t} = A\rho(\mathbf{s}) + \mathbf{e}$ (and of the ephemeral value): $\ell = k$ is the *full-matrix* profile, whose security rests on MLWE^{ρ}; $\ell = 1$ is the *compact* profile, resting on $\mathsf{sRLWE}_{\mathcal{A}_\alpha}$. Slot-wise correctness, the reconciliation kernel and the binding theory are identical for the two profiles; we present the algorithms in compact form ($\ell = 1$) and flag full-matrix differences where they occur.

4 Ephemeral Lattice Key Agreement and the Reconciliation Kernel

This section defines the ephemeral key-agreement layer of NAIBI and its shared *reconciliation kernel*, and proves correctness and security. Despite the algebraic analogy with classical Diffie–Hellman (cf. the table in Section 1), NAIBI is *not* a 0-pass non-interactive key exchange (NIKE) in the sense of that protocol, of CSIDH or of SWOOSH [31]. Just as importantly, **NAIBI does not offer a static-static mode**: as Section 8.4 explains in detail, reusing a long-term secret in a reconciliation protocol opens the active key-mismatch attack surface of the Ding/Peikert/NewHope family, and we exclude that deployment outright rather than patch it. The kernel below is therefore always run with at least one *fresh* per-session secret.

We offer two forward-secure flavours — one-pass ES-DH and two-pass EE-DH:

- **ES-DH (ephemeral-static).** Alice generates a fresh ephemeral key per session; Bob reuses a long-term key, but only *inside the FO-protected / authenticated* workflow that denies the adversary a key-mismatch oracle. Tightest reduction to single-sample structured MLWE; weak FS.
- **EE-DH (ephemeral-ephemeral).** Both parties generate fresh keys per session. Perfect forward secrecy (PFS) by construction, and immune to key-mismatch by design.

Both rely on the same algebraic core (Equation 3) and the same rounding mechanism. We present the reconciliation kernel first, then the two ephemeral protocols built on it.

4.1 Common Setup and Key Generation

Public parameters. (n, q, η, d_t), a fixed seed $\mathtt{seed}_A$, and the matrix $A = \mathsf{XOF}(\mathtt{seed}_A) \in M_2(R_q)$ (uniformly distributed in the full matrix algebra; *no* ρ-wrap). These are global, shared by every participant (analogous to a curve in ECDH). The parameter d_t controls public-key compression (Section 7); key extraction uses $\mathsf{PExt}/\mathsf{PRec}$ with the fixed modulus q (no additional rounding parameter needed).

Algorithm 1 NI-DH.KeyGen() (each party, executed once)

1: $\mathbf{s}_P \leftarrow \chi_\eta^2$ ▷ $\mathbf{s}_P \in \mathcal{A}_\alpha$, ternary; resample if non-invertible in $\mathcal{A}_\alpha$
2: $\mathbf{e}_P \leftarrow \chi_\eta^{2\times 2}$ ▷ $\mathbf{e}_P \in R_q^{2\times 2}$, generic ternary matrix
3: $\mathbf{t}_P \leftarrow A\,\rho(\mathbf{s}_P) + \mathbf{e}_P \in R_q^{2\times 2}$
4: $\bar{\mathbf{t}}_P \leftarrow \mathsf{Compress}(\mathbf{t}_P, d_t)$ ▷ public-key compression to d_t bits/coef
5: **return** $\big(\mathsf{pk}_P = \bar{\mathbf{t}}_P,\ \mathsf{sk}_P = \mathbf{s}_P\big)$

Symmetric primitives. Throughout this paper, KDF denotes SHAKE-256 (FIPS 202) used in a variable-output-length mode: $\mathsf{KDF}_N(x) := \text{SHAKE-256}(x, N)$, where the output length N is an instance parameter rather than a fixed constant, since SHAKE-256 is natively an extendable-output function. We write N_{ss} for the length used by the final KEM.Encaps/Decaps output $\mathsf{ss} \leftarrow \mathsf{KDF}_{N_{\text{ss}}}(\cdot)$: $N_{\text{ss}} = 256$ at every named set (NAIBI-I/III/V/V′), matching the classical security strength targeted at each level. The parameter is kept explicit so that a wider shared secret can be selected without touching any other part of the specification. Every other use of KDF in this paper — deriving the per-session key K in the ephemeral agreement modes, and the $\bar{K}$ half consumed internally by the FO transform — keeps the default $N = 256$, since those are intermediate values, not the final shared-secret output.

Widening N_{ss} costs nothing beyond re-parametrising the same SHAKE-256 call. Every security and binding-correctness argument in this paper treats KDF as a random oracle, and the resulting advantage and collision bounds (Theorems 7–8, 10–11) depend on the entropy of KDF's *input* — the message space $\mathcal{M} = \{0,1\}^{256}$ and the implicit-rejection seed $z \in \{0,1\}^{256}$ — not on its output length. The 2^{-256} random-oracle-collision term in Theorems 10–11 is in fact $2^{-N_{\text{ss}}}$, so widening N_{ss} can only tighten it: every bound stated with the literal constant 2^{-256} remains valid. The hash function $\mathsf{H} : \{0,1\}^* \to \{0,1\}^{256}$ is instantiated as SHA3-256, and the random oracle $\mathsf{G} : \{0,1\}^* \to \{0,1\}^{512}$ as SHA3-512 (split into two 256-bit halves). The extendable output function XOF used for matrix expansion is SHAKE-128. These choices align with the NIST PQC framework conventions and with ML-KEM's symmetric-primitive suite.

Long-term key generation. Each party P runs KeyGen once and publishes pk_P (Algorithm 1):

Public-key directory. We assume a standard mechanism (PKI / authenticated directory) by which pk_A and pk_B are reliably known to the corresponding counterparts. Authentication of pk is *out of scope* for the key-agreement layer itself, exactly as it is for classical DH.

Remark 6 (Why the secret is resampled to be invertible). The resampling in line 1 of Algorithm 1 excludes the non-invertible elements of $\mathcal{A}_\alpha$. Since $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$, a uniform element is non-invertible with probability $1 - (1 - q^{-k})^n \le n\,q^{-k}$; for a ternary secret the probability is no larger, so resampling fires with probability $\le n/q^k$ (e.g. $\le 2^{-23}$ for NAIBI-III), once per lifetime of a key. We

Algorithm 2 NI-DH.Derive$_A(\mathsf{sk}_A, \mathsf{pk}_A, \mathsf{pk}_B)$ (initiator, emits hint h_A)

Parse $\mathsf{sk}_A = \mathbf{s}_A$, $\mathsf{pk}_B = \bar{\mathbf{t}}_B \in \{0,\dots,2^{d_t}-1\}^{k^2 n}$
$\mathbf{t}'_B \leftarrow \mathsf{Decompress}(\bar{\mathbf{t}}_B, d_t) \in R_q^{k\times k}$
$\mathbf{e}''_A \leftarrow \chi_\eta^{k\times k}$; $\mathbf{r}_A \leftarrow \mathbf{t}'_B\,\rho(\mathbf{s}_A) + \mathbf{e}''_A \in R_q^{k\times k}$ ▷ "DH product" + fresh error (Remark 7)
Sample doubling bits $\bar{\mathbf{e}}$; $(h_A, k_A) \leftarrow \mathsf{PExt}(\mathbf{r}_A; \bar{\mathbf{e}})$ ▷ Def. 2: $h_A^{(v)} = c(2v - \bar{e}_v)$, $k_A^{(v)} = b(2v - \bar{e}_v)$; $h_A, k_A \in \{0,1\}^{k^2 n}$
$K \leftarrow \mathsf{KDF}\big(k_A \,\|\, \mathsf{H}(\mathsf{pk}_{\min}) \,\|\, \mathsf{H}(\mathsf{pk}_{\max})\big)$ ▷ $\mathsf{H}(\mathsf{pk}_{\min}) \le \mathsf{H}(\mathsf{pk}_{\max})$, lex. on the two pks
return $(K,\ h_A)$; transmit h_A to Bob

stress that invertibility is a *correctness/entropy hygiene* condition, *not* a security requirement, for three independent reasons. *(i) The reduction does not need it.* The hardness reduction of Section 8 (Theorem 13) embeds the MLWE challenge in $\mathbf{t}$ and never inverts $\mathbf{s}$; MLWE is no easier when the secret is drawn from the ternary distribution conditioned on invertibility, because excluding a $\le nq^{-k}$ fraction of secrets changes the secret distribution — and hence the advantage of any distinguisher — by at most that statistically negligible amount. *(ii) An attacker cannot detect or exploit it.* The bit "$\mathbf{s}$ is invertible" is not revealed by $\mathsf{pk} = \mathsf{Compress}(A\rho(\mathbf{s}) + \mathbf{e})$, which is pseudorandom under MLWE; and even a hypothetical non-invertible $\mathbf{s}$ would yield an equally hard MLWE/lattice recovery instance, so there is no class of weak keys to target. *(iii) The only effect is on raw-key entropy.* If $\mathbf{s}$ vanished in some CRT slot, the raw DH value $\mathbf{t}_B\rho(\mathbf{s}_A)$ would lose the entropy of that slot; requiring invertibility keeps the slot-wise map a bijection, which is what guarantees the full-entropy input to the KDF and underlies the binding argument of Theorem 10 (which assumes $\mathbf{s}_0, \mathbf{s}_1 \in \mathcal{A}_\alpha^\times$). Even without resampling, the raw key has $kn \gg 256$ coefficients and the KDF extracts only 256 bits, so the entropy margin is enormous; resampling is therefore a cheap belt-and-suspenders measure rather than a necessity.

4.2 The Reconciliation Kernel (shared by ES-DH, EE-DH and the KEM)

We describe the kernel in its symmetric two-key form; the ephemeral protocols of Sections 4.5–4.6 and the KEM of Section 5 are obtained by substituting a fresh ephemeral key for one (ES) or both (EE) of the two keys below. The emitter runs Derive_A to produce a $4n$-bit Peikert hint; the recoverer runs Derive_B. Both parties agree on identical raw key bits with negligible-failure ($\delta \le 2^{-128}$) correctness. We emphasise that the symmetric two-static-key presentation is expository only: **NAIBI never deploys the kernel with two reused long-term keys** (Section 8.4); at least one of the two keys is always a per-session ephemeral.

Properties.

Algorithm 3 $\mathsf{NI\text{-}DH.Derive}_B(\mathsf{sk}_B, \mathsf{pk}_A, \mathsf{pk}_B, h_A)$ (responder, recovers K from the hint)

1: Parse $\mathsf{sk}_B = \mathbf{s}_B$, $\mathsf{pk}_A = \bar{\mathbf{t}}_A$
2: $\mathbf{t}'_A \leftarrow \mathsf{Decompress}(\bar{\mathbf{t}}_A, d_t)$; $\mathbf{r}_B \leftarrow \mathbf{t}'_A\,\rho(\mathbf{s}_B) \in R_q^{k\times k}$
3: $k_B \leftarrow \mathsf{PRec}(\mathbf{r}_B, h_A)$ ▷ Def. 2: per coefficient w with hint bit h, $k_B^{(w)} = \mathsf{rec}(2w, h)$ over $\mathbb{Z}_{2q}$
4: $K \leftarrow \mathsf{KDF}\big(k_B \parallel \mathsf{H}(\mathsf{pk}_{\min}) \parallel \mathsf{H}(\mathsf{pk}_{\max})\big)$
5: **return** K

- **One k^2n-bit hint.** The emitter transmits a single k^2n-bit reconciliation hint alongside its (ephemeral) sample; the recoverer reads the hint and locally derives K. No raw key bits are ever transmitted.
- **Both sides agree.** Each party uses the decompressed $\mathbf{t}' = \mathbf{t} + \boldsymbol{\varepsilon}_t$, where $\boldsymbol{\varepsilon}_t \in R_q^{k\times k}$ is the public-key-compression error with entry-wise $\|\boldsymbol{\varepsilon}_t\|_\infty \le E_{d_t} = \lfloor q/2^{d_t+1} \rfloor$. By the commutativity identity (3), and writing $\mathbf{e}''_A$ for the emitter's fresh reconciliation error (Remark 7):

$$\begin{aligned}\mathbf{r}_A - \mathbf{r}_B &= \big(\mathbf{t}_B + \boldsymbol{\varepsilon}_{t,B}\big)\rho(\mathbf{s}_A) + \mathbf{e}''_A - \big(\mathbf{t}_A + \boldsymbol{\varepsilon}_{t,A}\big)\rho(\mathbf{s}_B)\\ &= \mathbf{e}_B\,\rho(\mathbf{s}_A) - \mathbf{e}_A\,\rho(\mathbf{s}_B) + \boldsymbol{\varepsilon}_{t,B}\,\rho(\mathbf{s}_A) - \boldsymbol{\varepsilon}_{t,A}\,\rho(\mathbf{s}_B) + \mathbf{e}''_A,\end{aligned}$$

  where the cancellation uses $A\rho(\mathbf{s}_B)\rho(\mathbf{s}_A) = A\rho(\mathbf{s}_A)\rho(\mathbf{s}_B)$ (commutativity of the *secrets* in $\rho(\mathcal{A}_\alpha)$). By Lemma 3 the noise part $\mathbf{e}_B\rho(\mathbf{s}_A) - \mathbf{e}_A\rho(\mathbf{s}_B)$ has worst-entry norm $\le 2\alpha kn\eta^2$ ($= 12n$ for $k = 2$, $\alpha = 3$, $\eta = 1$); the $\mathbf{t}$-compression parts each contribute $\le (\alpha+1)knE_{d_t}\eta/2$ entry-wise per term, and the fresh error adds $\le \eta$. For $k = 2$, $\alpha = 3$, $\eta = 1$:

$$\big\|\mathbf{r}_A - \mathbf{r}_B\big\|_\infty \;\le\; 12n + 8nE_{d_t} + 1.$$

  Note that the error matrices $\mathbf{e}_A, \mathbf{e}_B, \mathbf{e}''_A$ are generic in $R_q^{k\times k}$ and need carry no algebraic structure; only the secrets $\rho(\mathbf{s}_A), \rho(\mathbf{s}_B)$ live in the abelian image $\rho(\mathcal{A}_\alpha)$, which is all the cancellation requires.
- **Symmetric KDF input.** Hashes of public keys are concatenated in a canonical order so that both parties feed the same string into KDF, regardless of who is "A" or "B".

Theorem 2 (Reconciliation-kernel correctness). *Let $\mathbf{r}_A = \mathbf{t}'_B\,\rho(\mathbf{s}_A) + \mathbf{e}''_A$ and $\mathbf{r}_B = \mathbf{t}'_A\,\rho(\mathbf{s}_B) + \mathbf{e}''_B$ in $R_q^{k\times k}$ (with $\mathbf{e}''_\bullet$ the fresh reconciliation errors of Section 4.2). If a given coefficient of $\mathbf{r}_A - \mathbf{r}_B$ has magnitude $\le q/8 - 1$, then the recovery procedure of Algorithm 3, applied to Bob's $\mathbf{r}_B$ together with Alice's hint h_A from Algorithm 2, returns the matching bit $k_B = k_A$ (Lemma 2); if this holds for all k^2n coefficients then $K_A = K_B$. Moreover, by Lemma 1, when the reconciliation input is uniform the extracted key bits are exactly uniform given the transmitted hint — the property used in the security proofs of Theorems 6, 4 and 5.*

Consequently the agreement-failure probability is bounded by the per-coefficient Gaussian tail (10) with the agreement variance of (9):

$$\delta_{\mathrm{DH}} \;\le\; k^2 n \cdot \operatorname{erfc}\!\left(\tfrac{q/8-1}{\sigma_{\mathrm{DH}}\sqrt{2}}\right) \;\le\; 2^{-128} \tag{5}$$

for the parameters of Table 4 (see Table 6 for the per-instance σ_{DH} and δ_{DH}).

Proof. The doubling scheme of Definition 2 sends a uniform $v \in \mathbb{Z}_q$ to a uniform $\bar{v} \in \mathbb{Z}_{2q}$, on which the cross-rounding $c(\bar{v})$ and key bit $b(\bar{v})$ are the parity and top bit of the arc index. By Lemma 2, a recovering coefficient w with $|v-w|_q \le q/8-1$ satisfies $\mathsf{rec}(2w, c(\bar{v})) = b(\bar{v})$, so the bit is recovered correctly; a union bound over the $k^2 n$ coefficients against the Gaussian tail (10) gives the stated δ_{DH}. Exact uniformity of $b(\bar{v})$ given $c(\bar{v})$ is Lemma 1.

Remark 7 (Fresh reconciliation error and FO derandomisation). Following Peikert [29] and BCNS [30], the hint-*emitting* party adds a fresh error to its reconciliation input: wherever an algorithm of this paper computes $\mathbf{r} \leftarrow \mathbf{t}'\rho(\mathbf{s})$ on the emitting side, read instead

$$\mathbf{e}'' \leftarrow \chi_\eta^{k\times k}, \qquad \mathbf{r} \leftarrow \mathbf{t}'\rho(\mathbf{s}) + \mathbf{e}'' \in R_q^{k\times k},$$

followed by $(h, k) \leftarrow \mathsf{PExt}(\mathbf{r}; \bar{\mathbf{e}})$ with fresh doubling bits $\bar{\mathbf{e}}$ (Definition 2). This is what makes $\mathbf{r}$ a genuine MLWE^ρ sample and lets the security proofs avoid any learning-with-rounding assumption. Inside the FO-transformed PKE the coins are expanded as $(\mathbf{s}', \mathbf{e}', \mathbf{e}'', \bar{\mathbf{e}}) \leftarrow \mathsf{XOF}(r)$, so encryption is a deterministic function of $(m, \mathsf{pk}; r)$ and the re-encryption check recomputes the hint bit-for-bit. The recovering side is unchanged except that PRec is the $\mathbb{Z}_{2q}$ version $\mathsf{rec}(2\cdot, h)$. The per-coefficient recovery difference thereby gains the single additive coefficient of $\mathbf{e}''$ (variance $v = \eta/2$); the doubling offset is absorbed by the threshold shift $q/8 \to q/8-1$ of Lemma 2. Both effects are numerically invisible at the parameters of Table 4: every failure probability in Tables 5–6 is unchanged at the displayed precision (Appendix C). We stress that $\mathbf{e}''$ provides *passive* pseudorandomness only and does *not* make the static-static mode safe against active key-mismatch attacks; see Remark 42.

4.3 Security Model: the Key-Indistinguishability Experiment

The three security theorems of this section (Theorems 3, 4, 5) are proved in the standard execution model of Bellare–Rogaway / Canetti–Krawczyk [21], specialised to the message flows of Algorithms 4 and 5. All hash functions $\mathsf{H}, \mathsf{G}, \mathsf{KDF}$ are modelled as programmable random oracles; $N_{\mathrm{ss}} \ge 256$ is the shared-secret length; $\Delta(\cdot,\cdot)$ denotes statistical distance; and for a game G we write $\Pr[G \Rightarrow 1]$ for the probability that the experiment outputs 1.

Definition 4 (Sessions, partnering, transcript). *A* session *is an instance π_i^t (the t-th instance at party P_i) holding a role $\in \{\mathsf{init}, \mathsf{resp}\}$, an intended peer, ephemeral state, an acceptance flag, and (on acceptance) a key $K \in \{0,1\}^{N_{\mathrm{ss}}}$.*

Its transcript *is the ordered tuple of messages it sent and received. Two accepted sessions are* partners *if they have opposite roles, name each other as peer, and share matching transcripts (for ES-DH the single message* $(\bar{\mathbf{u}}_A, h_A)$*; for EE-DH the pair* $\big(\mathbf{u}_A, (\mathbf{u}_B, h_B)\big)$*).*

Definition 5 (KI experiment $\mathrm{Exp}_{\Pi}^{\text{KI-}b}$**).** *Fix* $\Pi \in \{\textit{EE-DH}, \textit{ES-DH}\}$ *and a bit* $b \in \{0,1\}$*. The challenger generates the long-term keys that* Π *uses (none for EE-DH; one pair* $(\mathsf{pk}_B, \mathbf{s}_B)$ *per server for ES-DH), draws the random-oracle tables lazily, and gives* $\mathcal{A}$ *the public parameters.* $\mathcal{A}$ *may adaptively issue:*

- $\mathsf{Send}(\pi_i^t, \mathsf{msg})$*: deliver* msg *(possibly* $\mathcal{A}$*-chosen) to the session and return its outgoing message, if any. Honest delivery realises a* passive *adversary;* $\mathcal{A}$*-chosen messages realise an* active *one.*
- $\mathsf{RevealKey}(\pi_i^t)$*: if accepted, return* K*.*
- $\mathsf{Corrupt}(P_i)$*: return the long-term secret* $\mathbf{s}_i$ *(meaningful for ES-DH; enables the KCI setting).*
- $\mathsf{RO}(\cdot)$*: query* $\mathsf{H}, \mathsf{G}, \mathsf{KDF}$*.*
- $\mathsf{Test}(\pi_{i^*}^{t^*})$*, asked once on a* fresh *(Definition 6) accepted session: return* K *if* $b = 0$ *and a uniform* $K^* \leftarrow \{0,1\}^{N_{\text{ss}}}$ *if* $b = 1$*.*

$\mathcal{A}$ *finally outputs a bit, which the experiment returns. Set*

$$\mathsf{Adv}_{\Pi}^{\text{KI}}(\mathcal{A}) := \big|\Pr[\mathrm{Exp}_{\Pi}^{\text{KI-0}}(\mathcal{A}) \Rightarrow 1] - \Pr[\mathrm{Exp}_{\Pi}^{\text{KI-1}}(\mathcal{A}) \Rightarrow 1]\big|.$$

Definition 6 (Freshness). *The tested session* $\pi_{i^*}^{t^*}$ *is* fresh *if: (F1) neither it nor its partner (if any) was a* RevealKey *target; and (F2) it has a partner,* or *its intended peer was never* Corrupt*ed before acceptance. For ES-DH we also use the* KCI-passive *relaxation: the tested session is an honest* initiator *session whose flow was delivered unmodified, and* Corrupt *of the peer is allowed (the adversary may hold* $\mathbf{s}_B$*), but no active modification of the tested flow is permitted. "Active in the initiator role" means* $\mathcal{A}$ *may use* $\mathcal{A}$*-chosen* Send *messages to* responder *sessions other than the test partner.*

Remark 8 (Canonical key = emitter key). In both modes the *hint-emitting* party derives its raw key *directly* via PExt, never via PRec: in EE-DH this is Bob (Algorithm 5, line 4), in ES-DH it is the ephemeral initiator Alice (Algorithm 4, line 4). We take the session key of a partnered pair to be that emitter key. By the agreement guarantee (Lemma 2, Theorem 2) the partner's PRec output equals it except with probability $\leq \delta_{\text{DH}}$; since the reductions below always compute and output the *emitter* key, which is exact, no δ term enters the KI advantage, and a RevealKey on a partnered pair is answered by this single canonical value.

A multi-instance form of the two-block assumption. The per-session reductions below replace, in one step, a public-key sample and the hint input that shares its secret. Across many sessions with *independent* secrets this is the following multi-instance generalisation of Definition 12.

Definition 7 (Multi-instance two-block assumption). *For $N \geq 1$ and $d \leq \lceil \log_2 q \rceil$, the problem $\mathsf{MLWE}^{\rho,(N;\mathcal{D}_d)}_{n,q,\chi}$ is: given $A \leftarrow \mathrm{U}(M_k(R_q))$, $\mathbf{B} \leftarrow \mathcal{D}_d$, and, for* independent *secrets $\mathbf{s}_1, \ldots, \mathbf{s}_N \leftarrow \chi(\mathcal{A}_\alpha)$ and independent generic errors $\mathbf{e}'_j, \mathbf{e}''_j \leftarrow \chi^{k\times k}$, the family $\left\{\left(A\rho(\mathbf{s}_j) + \mathbf{e}'_j,\ \mathbf{B}\rho(\mathbf{s}_j) + \mathbf{e}''_j\right)\right\}_{j=1}^{N}$, distinguish it from N independent uniform pairs. For $N = 1$ this is Definition 12(Definition 12); for $d = \lceil \log_2 q \rceil$ both block matrices are uniform and the problem is MLWE^{ρ} with stacked $2k \times k$ matrices.*

Lemma 6 (Independent-secret amplification). $\mathsf{Adv}^{\mathsf{MLWE}^{\rho,(N;\mathcal{D}_d)}}(\mathcal{B}) \leq N \cdot \mathsf{Adv}^{\mathsf{MLWE}^{\rho,(1;\mathcal{D}_d)}}(\mathcal{B}')$.

Proof. A standard hybrid $H_0, \ldots, H_N$, where H_ℓ presents the first ℓ pairs uniform and the remaining $N-\ell$ real under the *same* sampled $A, \mathbf{B}$. A distinguisher of $H_{\ell-1}$ from H_ℓ yields a single-instance distinguisher $\mathcal{B}'$: it plants its challenge pair in position ℓ, samples the first $\ell-1$ pairs uniform and the last $N-\ell$ as honest real samples (it knows their secrets, having drawn them) reusing the given $A, \mathbf{B}$, and runs $\mathcal{B}$. Telescoping over ℓ gives the bound. Independence of the secrets is precisely what makes the off-target pairs simulable; the *same* device fails for the same-secret reuse of Theorem 3, which is why that mode carries a genuine T-fold loss.

4.4 Security of the Kernel: Why Reuse Is Excluded

Run with a fresh per-session secret on at least one side (ES-DH, EE-DH, and the KEM), the kernel reduces *tightly* to structured-secret MLWE — via the fresh reconciliation error of Remark 7 and the exact key-bit uniformity of Lemma 1 — as established in Sections 4.5–4.6 and 5. For completeness, and to make the exclusion of the static-static mode quantitative, we record what *would* happen under long-term reuse of both keys: the kernel would leak, over the lifetime of a key, multiple rounded samples sharing one secret, so its (passive) security would rest on the *multi-sample* version of constant-ratio MLWR — itself only a heuristic assumption (cf. Definition 1) — with a T-fold tightness loss. This tightness loss is the milder of the two problems; the decisive one is that the same reuse opens the active key-mismatch attack surface of Section 8.4. We state the multi-sample reduction only to delimit precisely the regime NAIBI does not use.

Theorem 3 (Multi-sample security of the kernel under (excluded) reuse). *Were the kernel deployed static-static, let $\mathcal{A}$ be a passive adversary in the key-indistinguishability game against a target session (P, Q), with T total active sessions involving either pk_P or pk_Q, and making q_H random-oracle queries. Then there exist PPT adversaries $\mathcal{B}_1, \mathcal{B}_2$ such that*

$$\mathsf{Adv}^{\mathrm{KI}}(\mathcal{A}) \ \leq\ \mathsf{Adv}^{\mathsf{MLWR}}(\mathcal{B}_1) + \mathsf{Adv}^{\mathsf{MLWR}^{(T)}}(\mathcal{B}_2) + \frac{q_H}{2^{256}},$$

where $\mathsf{MLWR}^{(T)}$ denotes the T-sample variant (itself only a heuristic assumption at constant rounding ratio, cf. the discussion after Definition 1 — one more

reason this regime is excluded). This bounds only passive security; it offers no *protection against the active key-mismatch adversary of Section 8.4, which is why the static-static mode is excluded.*

Proof. This theorem characterises the regime NAIBI *does not* deploy; its proof legitimately uses the (heuristic) rounding assumptions MLWR and $\mathsf{MLWR}^{(T)}$, because static-static reuse leaks rounded products of one reused secret across T sessions, with no fresh ephemeral on the reused side to invoke the fresh-error patch (Remark 7). Let the tested session be (P, Q), let T be the number of active sessions touching pk_P or pk_Q, and q_H the number of RO queries; we use $k = 2$ for concreteness. We exhibit a chain $G_0, \ldots, G_3$ with G_0 the real-key world ($b = 0$) and G_3 the random-key world ($b = 1$), and write $S_i := \Pr[G_i \Rightarrow 1]$, so that $\mathsf{Adv}^{\mathrm{KI}}(\mathcal{A}) = |S_0 - S_3| \le \sum_i |S_i - S_{i+1}|$.

Game G_0. Real KI game on the static-static kernel, $b = 0$.

Game G_1 (uniformise both long-term keys). Replace each published $\bar{\mathbf{t}}_P = \mathsf{Compress}(A\rho(\mathbf{s}_P) + \mathbf{e}_P, d_t)$ by uniform $\bar{\mathbf{t}}_P \leftarrow \{0, \ldots, 2^{d_t} - 1\}^{4n}$, and likewise $\bar{\mathbf{t}}_Q$.

Reduction $\mathcal{B}_1$ to MLWR. By Lemma 13 each column of $A\rho(\mathbf{s}_P) + \mathbf{e}_P$ is a rank-2 MLWE sample with shared matrix A; compressing at d_t turns the two columns into two structured-secret MLWR samples per party (the $\rho(y)$-link between the secret columns is the exact identification of Theorem 13). $\mathcal{B}_1$ embeds its MLWR challenge into these rounded columns for both P and Q, simulating the rest of the game honestly. A real challenge yields G_0, a uniform one G_1; thus $|S_0 - S_1| \le \mathsf{Adv}^{\mathsf{MLWR}}(\mathcal{B}_1)$. After this hop $\mathbf{t}'_P = \mathsf{Decompress}(\bar{\mathbf{t}}_P, d_t)$ is uniform on a coarse sublattice of $R_q^{2\times 2}$, indistinguishable from uniform on $R_q^{2\times 2}$ for the next hop.

Game G_2 (uniformise the reconciled products across all T sessions). In each of the T sessions involving pk_Q, the reconciliation input is $\mathbf{t}_P^{(j)}\rho(\mathbf{s}_Q)$, a right-action of the reused structured factor $\rho(\mathbf{s}_Q) \in \rho(\mathcal{A}_\alpha)$ on the now-uniform $\mathbf{t}_P^{(j)}$; the transmitted pair $(h_A^{(j)}, k_A^{(j)}) = \mathsf{PExt}\big(\mathbf{t}_P^{(j)}\rho(\mathbf{s}_Q)\big)$ is a deterministic 2-bit-per-entry function of it. Replace every $(h_A^{(j)}, k_A^{(j)})$, $j = 1, \ldots, T$, by fresh uniform bits.

Reduction $\mathcal{B}_2$ to $\mathsf{MLWR}^{(T)}$. The T products $\{\mathbf{t}_P^{(j)}\rho(\mathbf{s}_Q)\}_j$ are T MLWR samples sharing the *single* reused secret $\mathbf{s}_Q$ against uniform matrices $\mathbf{t}_P^{(j)}$ — precisely a T-sample MLWR instance (over R_q, via Theorem 13). $\mathcal{B}_2$ plants its T-sample challenge as these products, applies the public PExt map entry-wise to obtain the $(h_A^{(j)}, k_A^{(j)})$, and runs $\mathcal{A}$; a distinguisher of the bits is a distinguisher of the samples, so $|S_1 - S_2| \le \mathsf{Adv}^{\mathsf{MLWR}^{(T)}}(\mathcal{B}_2)$.

Why the T-fold loss is intrinsic. The T samples share the *same* secret $\mathbf{s}_Q$, so the independent-secret amplification of Lemma 6 does *not* apply: the off-target samples cannot be simulated without $\mathbf{s}_Q$. The reduction is therefore to the genuine multi-sample assumption $\mathsf{MLWR}^{(T)}$, believed only polynomially equivalent to single-sample MLWR with a concrete factor T (cf. the discussion after Definition 1). This is the milder of the two reasons the mode is excluded; the decisive one is the active key-mismatch attack of Section 8.4, against which this passive bound offers no protection.

Game G_3 (idealise the tested key). With the tested session's (h_A, k_A) now uniform and independent of all public keys, $K = \mathsf{KDF}(k_A \| \cdots)$ is replaced by uniform; in the ROM this costs $\leq q_H/2^{256}$. G_3 is the $b = 1$ world.

Summation.

$$\mathsf{Adv}^{\mathrm{KI}}(\mathcal{A}) \;\leq\; \mathsf{Adv}^{\mathsf{MLWR}}(\mathcal{B}_1) + \mathsf{Adv}^{\mathsf{MLWR}^{(T)}}(\mathcal{B}_2) + \frac{q_H}{2^{256}},$$

the bound of the statement. This delimits the *excluded* static-static regime only; the offered ES-DH and EE-DH modes avoid both the T-factor and the active attack.

Remark 9 (Why this regime is excluded). The T-factor in the multi-sample MLWR reduction is the passive price of a static-static scheme: an adversary who corrupts T sessions sees T rounded samples of the same secret. That loss alone would be tolerable for moderate T. The reason NAIBI nonetheless excludes the static-static mode is *active*: under reuse, the transmitted hint plus a key-match/abort oracle recover the long-term secret in a handful of sessions (Section 8.4). ES-DH and EE-DH (next subsections) avoid both problems and are the offered schemes.

4.5 The Ephemeral-Static Variant (NAIBI-ES-DH)

The ephemeral-static variant (Algorithm 4) targets server-authenticated, TLS-like deployments: Alice generates a fresh ephemeral $(\mathbf{s}_A^{\mathrm{eph}}, \mathbf{e}_A^{\mathrm{eph}})$ per session, while Bob reuses his long-term key $(\mathsf{pk}_B, \mathsf{sk}_B)$.

Remark 10 (Hint direction and active key-mismatch resistance). In a reconciliation protocol the party that *emits* the Peikert hint exposes, through that hint, the secret factor of the product it reconciles. In the ephemeral-static setting only one secret is reused across sessions — Bob's long-term $\mathbf{s}_B$ — so soundness against an active initiator requires that the emitted hint reconcile a product involving the *fresh* secret, never the reused one. We therefore have the hint emitted by the *ephemeral* party (Alice), on $\mathbf{r}_A = \mathbf{t}'_B\,\rho(\mathbf{s}_A^{\mathrm{eph}})$, and never by the long-term party. Bob recovers locally with PRec, emitting no second message. This is exactly the hint direction of the KEM (Algorithm 7, where the encapsulator emits the hint), so ES-DH is structurally "one leg of the KEM" (Appendix B). The opposite convention, with Bob emitting a hint on $\mathbf{u}'_A\,\rho(\mathbf{s}_B)$, is the tempting one — it saves nothing, and it is broken by an active initiator, as Remark 11 shows.

Remark 11 (Why the long-term party must not emit the hint). Were Bob (the long-term party) to emit the hint $h_B = \mathsf{PExt}(\mathbf{u}'_A\rho(\mathbf{s}_B))$, an *active* initiator would recover $\mathbf{s}_B$ in a handful of sessions, exactly the key-mismatch attack template of Section 8.4. Concretely, the attacker sends a malformed $\bar{\mathbf{u}}_A = \mathsf{Compress}(c\,E_{ij}, d_u)$ for a scalar $c \in \mathbb{Z}_q$ and an elementary matrix E_{ij}, so that $\mathbf{u}'_A\rho(\mathbf{s}_B)$ has a row equal to $c\cdot$(a row of $\rho(\mathbf{s}_B)$). The emitted cross-rounding bits $c(\overline{c\,(\mathbf{s}_B)_\ell})$ are then

Algorithm 4 NI-DH-ES: ephemeral-static, one-pass, *ephemeral party emits the hint*

Setup: $A = \mathsf{XOF}(\mathsf{seed}) \in M_k(R_q)$ (common); $\mathsf{pk}_B = \bar{\mathbf{t}}_B$ long-term (compressed at d_t).

Pass 1 (Alice $\to$ Bob): ephemeral sample *and* hint

1: Alice: $\mathbf{s}_A^{\mathrm{eph}} \leftarrow \chi_\eta^k$ $(\in \mathcal{A}_\alpha)$, $\mathbf{e}_A^{\mathrm{eph}} \leftarrow \chi_\eta^{k\times k}$, $\mathbf{e}''_A \leftarrow \chi_\eta^{k\times k}$ ▷ fresh errors, Remark 7

2: Alice: $\mathbf{u}_A \leftarrow A\,\rho(\mathbf{s}_A^{\mathrm{eph}}) + \mathbf{e}_A^{\mathrm{eph}}$; $\bar{\mathbf{u}}_A \leftarrow \mathsf{Compress}(\mathbf{u}_A, d_u)$

3: Alice: $\mathbf{t}'_B \leftarrow \mathsf{Decompress}(\bar{\mathbf{t}}_B, d_t)$; $\mathbf{r}_A \leftarrow \mathbf{t}'_B\,\rho(\mathbf{s}_A^{\mathrm{eph}}) + \mathbf{e}''_A$ ▷ hint reconciles the *ephemeral* secret

4: Alice: sample doubling bits $\bar{\mathbf{e}}_A$; $(h_A, k_A) \leftarrow \mathsf{PExt}(\mathbf{r}_A; \bar{\mathbf{e}}_A)$ ▷ $h_A \in \{0,1\}^{k^2 n}$ hint, k_A raw key

5: Alice: $K_A \leftarrow \mathsf{KDF}(k_A \parallel \bar{\mathbf{u}}_A \parallel \mathsf{pk}_B)$; output K_A

6: Alice $\to$ Bob: $(\bar{\mathbf{u}}_A,\ h_A)$

Local (Bob): recover shared secret

7: Bob: $\mathbf{u}'_A \leftarrow \mathsf{Decompress}(\bar{\mathbf{u}}_A, d_u)$; $\mathbf{r}_B \leftarrow \mathbf{u}'_A\,\rho(\mathbf{s}_B) \in R_q^{k\times k}$ ▷ Bob emits *no* hint

8: Bob: $k_B \leftarrow \mathsf{PRec}(\mathbf{r}_B,\ h_A)$; $K_B \leftarrow \mathsf{KDF}(k_B \parallel \bar{\mathbf{u}}_A \parallel \mathsf{pk}_B)$; output K_B

threshold indicators on each ternary coefficient $(\mathbf{s}_B)_\ell$; the fresh error $\mathbf{e}''_B$ (amplitude $\le \eta = 1$) is far too small to blur a threshold of spacing $\Theta(q)$, and being resampled while $\mathbf{s}_B$ is reused it averages out under majority voting. Sweeping c over an $O(\log q)$ grid and reading the toggle directions recovers the whole of $\mathbf{s}_B$ in $\approx$ 15–256 sessions, independent of n (each session's $k^2 n$-bit hint exposes all coefficients in parallel). Algorithm 4 closes this at the root: the only hint ever transmitted, h_A, reconciles $\mathbf{t}'_B\rho(\mathbf{s}_A^{\mathrm{eph}})$, whose secret factor $\mathbf{s}_A^{\mathrm{eph}}$ is *fresh per session*; a malformed long-term $\mathbf{t}_B$ chosen by a malicious responder can probe only that ephemeral secret, which is erased afterward, so no per-session leakage accumulates. Bob transmits no hint, and $\mathbf{s}_B$ never appears in any product whose reconciliation data is revealed.

Lemma 7 (The ES-DH transcript is $\mathbf{s}_B$-independent given pk_B). *In Algorithm 4, every message an active-initiator adversary observes or injects — the long-term $\mathsf{pk}_B = \bar{\mathbf{t}}_B$ and, for each honest initiator session j, the pair $(\bar{\mathbf{u}}_A^{(j)}, h_A^{(j)})$ — is a deterministic function of pk_B and of per-session* ephemeral *randomness $(\mathbf{s}_A^{\mathrm{eph},(j)}, \mathbf{e}_A^{\mathrm{eph},(j)}, \mathbf{e}''^{(j)}_A, \bar{\mathbf{e}}_A^{(j)})$* alone. *In particular, conditioned on pk_B, the joint distribution of the entire transcript is independent of $\mathbf{s}_B$.*

Proof. Read off Algorithm 4. Alice computes $\mathbf{u}_A = A\rho(\mathbf{s}_A^{\mathrm{eph}}) + \mathbf{e}_A^{\mathrm{eph}}$, $\bar{\mathbf{u}}_A = \mathsf{Compress}(\mathbf{u}_A, d_u)$, $\mathbf{t}'_B = \mathsf{Decompress}(\bar{\mathbf{t}}_B, d_t)$, $\mathbf{r}_A = \mathbf{t}'_B\rho(\mathbf{s}_A^{\mathrm{eph}}) + \mathbf{e}''_A$, and $(h_A, k_A) = \mathsf{PExt}(\mathbf{r}_A; \bar{\mathbf{e}}_A)$: each right-hand side depends only on the public A, the public $\mathbf{t}'_B$ (a function of pk_B), and the ephemeral coins; $\mathbf{s}_B$ appears nowhere. Bob's only use of $\mathbf{s}_B$, $\mathbf{r}_B = \mathbf{u}'_A\rho(\mathbf{s}_B)$ followed by $k_B = \mathsf{PRec}(\mathbf{r}_B, h_A)$, produces *no outgoing message* (Algorithm 4, lines 7–8: Bob emits no hint). Hence no transmitted bit is a function of $\mathbf{s}_B$ except through pk_B, and the conditional independence follows.

Corollary 2 (Unconditional active-initiator resistance). *No active-initiator adversary, however adaptive, obtains any information about* $\mathbf{s}_B$ *beyond* pk_B*. Consequently the key-mismatch oracle of Remark 11 has no reused secret to probe: the only secret ever reconciled by a* transmitted *hint is the per-session ephemeral* $\mathbf{s}_A^{\mathrm{eph}}$*, erased after one use.*

Theorem 4 (ES-DH security, single sample). *Run as in Algorithm 4 (ephemeral party emits the hint), NAIBI-ES-DH is secure against an* active *adversary in the initiator role, as well as against passive and KCI-passive adversaries. For any such adversary making* q_{sess} *session queries and* q_H *RO queries, there exist PPT* $\mathcal{B}_0, \mathcal{B}_1, \mathcal{B}_2$ *with*

$$\begin{aligned}\mathsf{Adv}^{\mathrm{KI}}_{\mathrm{ES}}(\mathcal{A}) \;&\le \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_0) + q_{\mathrm{sess}}\,\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_1)\\ &\quad + \mathsf{Adv}^{\mathsf{MLWE}^{\rho,(q_{\mathrm{sess}};\mathcal{D}_{d_t})}}(\mathcal{B}_2) + \frac{q_H}{2^{256}}.\end{aligned}$$

The active-initiator resistance is unconditional *in the sense that no hint reconciles the reused secret: every transmitted* $h_A^{(j)}$ *is a function of a fresh ephemeral* $\mathbf{s}_A^{\mathrm{eph},(j)}$*, so the active key-mismatch oracle of Remark 11 has no reused target to probe. The third term is the* q_{sess}*-sample structured-MLWE form over Bob's long-term* $\mathbf{s}_B$ *reconstructed by* PRec *from the honest* $\mathbf{t}_B$ *— standard multi-sample MLWE, no T-fold loss. The key-bit hop is free by Lemma 1.*

Proof. The statement has two independent halves: an *information-theoretic* clause (active-initiator resistance: no transmitted hint reconciles the reused secret), which is exactly Corollary 2; and a *computational* clause (key indistinguishability of the tested session), which we prove now. We treat the non-KCI active-initiator setting (the peer is not corrupted before the tested session accepts), where the setup hop is sound; Remark 13 discharges the KCI-passive case.

Let $\mathcal{A}$ make q_{sess} session and q_H RO queries. By Remark 8 the tested key is the initiator key $K_A = \mathsf{KDF}(k_A \,\|\, \bar{\mathbf{u}}_A \,\|\, \mathsf{pk}_B)$ with $(h_A, k_A) = \mathsf{PExt}(\mathbf{r}_A; \bar{\mathbf{e}}_A)$, $\mathbf{r}_A = \mathbf{t}'_B \rho(\mathbf{s}_A^{\mathrm{eph}}) + \mathbf{e}''_A$ and $\mathbf{t}'_B = \mathsf{Decompress}(\bar{\mathbf{t}}_B, d_t)$. Games $G_0, \ldots, G_4$, $S_i := \Pr[G_i \Rightarrow 1]$:

Game G_0**.** Real KI game, $b = 0$.

Game G_1 **(setup hop: uniformise the long-term key).** Replace $\mathbf{t}_B = A\rho(\mathbf{s}_B) + \mathbf{e}_B$ by $\mathbf{U} \leftarrow \mathrm{U}(M_k(R_q))$ and publish $\bar{\mathbf{t}}_B = \mathsf{Compress}(\mathbf{U}, d_t)$; the initiator-local matrix $\mathbf{t}'_B = \mathsf{Decompress}(\bar{\mathbf{t}}_B, d_t)$ is then distributed *exactly* as $\mathcal{D}_{d_t}$ (Definition 1).

Reduction $\mathcal{B}_0$ *to* MLWE^{ρ}*.* On challenge (A, T), set the common matrix to A and publish $\bar{\mathbf{t}}_B := \mathsf{Compress}(T, d_t)$. Since $\mathsf{Compress}$ is a public function of an MLWE^{ρ} sample, Remark 1 gives $|S_0 - S_1| \le \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_0)$. (If KeyGen conditions on $\mathbf{s}_B$ invertible, add the statistical $\varepsilon_{\mathrm{inv}} \le n/q^k$ of Theorem 6, omitted in the displayed bound.)

Answering oracles after G_1 *without* $\mathbf{s}_B$*.* The simulator no longer holds $\mathbf{s}_B$; it still answers every query. *(i) Honest initiator sessions* (including the

tested one): the emitter key K_A is computed from the ephemeral secret and the public $\mathbf{t}'_B$ only (Lemma 7). *(ii)* $\mathsf{RevealKey}$ *on a responder session partnered with an honest initiator*: answered by copying that initiator's canonical key (Remark 8). *(iii) Unmatched (forged) responder sessions*, where $\mathcal{A}$ injected $(\bar{\mathbf{u}}_A, h_A)$ produced by no honest initiator: the responder key would be $\mathsf{KDF}\big(\mathsf{PRec}(\mathbf{u}'_A\rho(\mathbf{s}_B), h_A) \,\|\, \bar{\mathbf{u}}_A \,\|\, \mathsf{pk}_B\big)$; the simulator answers with a fresh uniform K^*, recorded per transcript. This is detectable only if $\mathcal{A}$ queries KDF at the matching point, which requires $k_B = \mathsf{PRec}(\mathbf{u}'_A\rho(\mathbf{s}_B), h_A)$; by Lemma 7 the whole view is independent of $\mathbf{s}_B$ given pk_B, so $\mathcal{A}$ cannot predict the $k^2n \geq 256$-bit k_B better than guessing, and the total simulation error over q_H queries is $\leq q_H/2^{256}$, absorbed into the final RO term.

Game G_2 (per-session hop: uniformise the ephemeral samples). In every honest initiator session $j = 1, \ldots, q_{\mathrm{sess}}$ replace, *jointly*,

$$\mathbf{u}_A^{(j)} = A\rho(\mathbf{s}_A^{\mathrm{eph},(j)}) + \mathbf{e}_A^{\mathrm{eph},(j)} \quad \text{and} \quad \mathbf{r}_A^{(j)} = \mathbf{t}'_B\rho(\mathbf{s}_A^{\mathrm{eph},(j)}) + \mathbf{e}_A''^{(j)}$$

by independent uniforms $(\mathbf{U}_1^{(j)}, \mathbf{U}_2^{(j)})$, then set $\bar{\mathbf{u}}_A^{(j)} = \mathsf{Compress}(\mathbf{U}_1^{(j)}, d_u)$ and $(h_A^{(j)}, k_A^{(j)}) = \mathsf{PExt}(\mathbf{U}_2^{(j)}; \bar{\mathbf{e}}_A^{(j)})$. This is exactly the multi-instance two-block assumption of Definition 7 over the *ephemeral* secrets, with first block matrix $A \leftarrow \mathrm{U}(M_k(R_q))$ and second block matrix $\mathbf{t}'_B \sim \mathcal{D}_{d_t}$ (made coarse-uniform by G_1), independent secrets across sessions and independent fresh errors.

Reduction $\mathcal{B}_2$. Receive $\big(A, \mathbf{B}, \{(\mathbf{X}_j, \mathbf{Y}_j)\}_j\big)$; set the common matrix to A and publish $\bar{\mathbf{t}}_B$ with $\mathsf{Decompress}(\bar{\mathbf{t}}_B, d_t) = \mathbf{B}$ (a coarse-uniform handle is part of the $\mathcal{D}_{d_t}$ instance). For session j set $\mathbf{u}_A^{(j)} := \mathbf{X}_j$, $\mathbf{r}_A^{(j)} := \mathbf{Y}_j$, derive $\bar{\mathbf{u}}_A^{(j)}, h_A^{(j)}, k_A^{(j)}, K_A^{(j)}$ as above with fresh $\bar{\mathbf{e}}_A^{(j)}$; forged-responder and reveal queries are answered as in G_1. Real instance $\Rightarrow G_1$, uniform $\Rightarrow G_2$, so $|S_1 - S_2| \leq \mathsf{Adv}^{\mathsf{MLWE}^{\rho,(q_{\mathrm{sess}};\mathcal{D}_{d_t})}}(\mathcal{B}_2)$.

Soundness of the joint move and active immunity. The two objects share $\mathbf{s}_A^{\mathrm{eph},(j)}$, so they must move in one hop (a single-sample split would be unsound, exactly as in Theorem 6, Game 2). The replaced quantities are tied to the *ephemeral* secret; an active choice of $\bar{\mathbf{u}}_A$ or a malformed $\mathbf{t}_B$ by $\mathcal{A}$ therefore cannot accumulate any information on $\mathbf{s}_B$, which by Lemma 7 never enters a transmitted value.

Game G_3 (key bits exact). With every $\mathbf{r}_A^{(j)}$ uniform, Lemma 1 makes each $(h_A^{(j)}, k_A^{(j)})$ have $k_A^{(j)}$ *exactly* uniform given $h_A^{(j)}$: $\Delta(G_2, G_3) = 0$, so $S_2 = S_3$.

Game G_4 (idealise the tested key). Replace the tested K_A by uniform K^*. In the ROM this costs $\leq q_H/2^{256}$ as before; together with the forged-responder simulation error of G_1 the RO term remains $q_H/2^{256}$. G_4 is the $b = 1$ world.

Summation.

$$\mathsf{Adv}^{\mathrm{KI}}_{\mathrm{ES}}(\mathcal{A}) \;\leq\; \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_0) + \mathsf{Adv}^{\mathsf{MLWE}^{\rho,(q_{\mathrm{sess}};\mathcal{D}_{d_t})}}(\mathcal{B}_2) + \frac{q_H}{2^{256}}.$$

Forward secrecy is *weak*: the tested key depends on Alice's erased ephemeral secret, so corrupting $\mathbf{s}_B$ afterwards reveals nothing about it. The active-initiator clause is Corollary 2.

Remark 12 (Reconciliation with the displayed bound). The bound displayed in Theorem 4 carries an additional $q_{\text{sess}}\,\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_1)$ term for the "$\mathbf{u}_A$ replacements". The proof's G_2 moves $\mathbf{u}_A^{(j)}$ and $\mathbf{r}_A^{(j)}$ *together* inside the single multi-instance two-block call, so that term is subsumed; the displayed three-term form is the looser accounting in which the two blocks are separated and the $\mathbf{u}_A$ block is counted as a per-session single-sample MLWE^{ρ}. By Lemma 6 the proven multi-instance term is in turn $\le q_{\text{sess}}\,\mathsf{Adv}^{\mathsf{MLWE}^{\rho,(1;\mathcal{D}_{d_t})}}$, the single-instance two-block advantage of Definition 12, so both forms upper-bound the same quantity. In the uncompressed option $d_t = \lceil \log_2 q \rceil$ (Remark 2) the coarse matrix $\mathbf{t}'_B$ is uniform and the whole bound rests on MLWE^{ρ} alone.

Remark 13 (The KCI-passive case). Under $\mathsf{Corrupt}(B)$ the value $\mathbf{t}_B$ must stay consistent with the revealed $\mathbf{s}_B$, so the setup hop G_1 is unavailable. Two routes recover the claim. *(i) Information-theoretic*: by Lemma 7 the tested transcript is a function of the ephemeral secret and pk_B only, so an adversary *given* $\mathbf{s}_B$ learns nothing extra about the tested key beyond what the per-session hop G_2 already idealises; the proof goes through verbatim once one treats the honest published $\mathbf{t}'_B$ as the second-block matrix. *(ii) Assumption*: invoke the natural standing assumption that structured-secret MLWE with a fresh reconciliation error is hard against the *published* (non-adversarial) coarse matrix $\mathbf{t}'_B$ — the same posture as Definition 1. Either route leaves the displayed bound unchanged; neither lets the active key-mismatch oracle of Remark 11 act, as no transmitted hint reconciles $\mathbf{s}_B$.

Remark 14 (Relation to the KEM and deployment guidance). Algorithm 4 is the agreement-mode reading of the KEM: Alice runs "one leg" of the protocol with a fresh $(\mathbf{s}_A^{\text{eph}}, \mathbf{e}_A^{\text{eph}})$, emits $(\bar{\mathbf{u}}_A, h_A)$, and the long-term party recovers the shared secret locally (cf. Appendix B). No second message is sent. Deployments that already use the IND-CCA2 KEM obtain the same forward-secure ephemeral-static guarantee through encapsulation to pk_B: KEM encapsulation *is* NI-DH-ES in one pass, with the FO$^{\perp}$ transform providing IND-CCA2 security on top. In both cases the reused secret $\mathbf{s}_B$ is never the secret factor of a transmitted hint, which is what makes the mode immune to the active key-mismatch attacks excluded in Section 8.4.

4.6 The Ephemeral-Ephemeral Variant (NAIBI-EE-DH)

For settings demanding perfect forward secrecy without static keys (e.g. Noise's NN pattern), both parties run KeyGen per session. The protocol is given in Algorithm 5.

Theorem 5 (EE-DH security, single sample). *For any passive adversary against NAIBI-EE-DH:*

$$\mathsf{Adv}^{\text{KI}}_{\text{EE}}(\mathcal{A}) \;\le\; 3\,\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}) + \frac{q_H}{2^{256}}.$$

Algorithm 5 NI-DH-EE: ephemeral-ephemeral, two-pass, PFS, with Peikert hint

Setup: $A = \mathsf{XOF}(\mathtt{seed}) \in M_k(R_q)$ (common).

Pass 1 (Alice $\to$ Bob)

1: Alice: $\mathbf{s}_A \leftarrow \chi_\eta^k$ $(\in \mathcal{A}_\alpha)$, $\mathbf{e}_A \leftarrow \chi_\eta^{k\times k}$ $(\in R_q^{k\times k})$; $\mathbf{u}_A \leftarrow A\,\rho(\mathbf{s}_A) + \mathbf{e}_A \in R_q^{k\times k}$; send $\mathbf{u}_A$

Pass 2 (Bob $\to$ Alice)

2: Bob: $\mathbf{s}_B \leftarrow \chi_\eta^k$ $(\in \mathcal{A}_\alpha)$, $\mathbf{e}_B \leftarrow \chi_\eta^{k\times k}$ $(\in R_q^{k\times k})$; $\mathbf{u}_B \leftarrow A\,\rho(\mathbf{s}_B) + \mathbf{e}_B \in R_q^{k\times k}$

3: Bob: $\mathbf{e}''_B \leftarrow \chi_\eta^{k\times k}$; $\mathbf{r}_B \leftarrow \mathbf{u}_A\,\rho(\mathbf{s}_B) + \mathbf{e}''_B \in R_q^{k\times k}$; sample doubling bits $\bar{\mathbf{e}}_B$; $(h_B,\, k_B) \leftarrow \mathsf{PExt}(\mathbf{r}_B; \bar{\mathbf{e}}_B)$ ▷ hint from Bob's local DH product + fresh error (Remark 7)

4: Bob $\to$ Alice: $(\mathbf{u}_B,\ h_B)$ ▷ ephemeral public key + $k^2 n$-bit reconciliation hint

Local key derivation (both)

5: Bob: $K_B \leftarrow \mathsf{KDF}(k_B \,\|\, \mathbf{u}_A \,\|\, \mathbf{u}_B)$

6: Alice: $\mathbf{r}_A \leftarrow \mathbf{u}_B\,\rho(\mathbf{s}_A) \in R_q^{k\times k}$; $k_A \leftarrow \mathsf{PRec}(\mathbf{r}_A,\ h_B)$; $K_A \leftarrow \mathsf{KDF}(k_A \,\|\, \mathbf{u}_A \,\|\, \mathbf{u}_B)$

7: All ephemeral material is erased.

Hops: (1) $\mathbf{u}_A$ uniform; (2) $\mathbf{u}_B$ uniform; (3) $\mathbf{r}_B = \mathbf{u}_A\rho(\mathbf{s}_B)+\mathbf{e}''_B$ uniform — since $\mathbf{u}_A$ is now uniform (EE-DH transmits it uncompressed) and $\mathbf{e}''_B$ is fresh, this is a plain MLWE^ρ sample; (4) key bits exactly uniform given the hint (Lemma 1); (5) KDF in the ROM. Forward secrecy follows from the ephemerality of both keys. No rounding assumption appears.

Proof. Passive sessions of EE-DH are mutually independent (both keys ephemeral and single-use), so it suffices to bound the advantage on a single tested session (Alice, Bob). We use a chain $G_0, \ldots, G_4$ with G_0 the real-key world ($b = 0$) and G_4 the random-key world ($b = 1$), $S_i := \Pr[G_i \Rightarrow 1]$, so $\mathsf{Adv}^{\mathrm{KI}}_{\mathrm{EE}}(\mathcal{A}) = |S_0 - S_4| \le \sum_i |S_i - S_{i+1}|$. By Remark 8 the tested key is the emitter key $K_B = \mathsf{KDF}(k_B \,\|\, \mathbf{u}_A \,\|\, \mathbf{u}_B)$ with $(h_B, k_B) = \mathsf{PExt}(\mathbf{r}_B; \bar{\mathbf{e}}_B)$, $\mathbf{r}_B = \mathbf{u}_A\rho(\mathbf{s}_B) + \mathbf{e}''_B$.

Game G_0. Real KI game, $b = 0$.

Game G_1 (replace $\mathbf{u}_A$ by uniform). Sample $\mathbf{u}_A \leftarrow \mathrm{U}(M_k(R_q))$ instead of $A\rho(\mathbf{s}_A)+\mathbf{e}_A$; everything else is unchanged, so Bob still forms $\mathbf{r}_B = \mathbf{u}_A\rho(\mathbf{s}_B)+\mathbf{e}''_B$ from the now-uniform $\mathbf{u}_A$.

Reduction $\mathcal{B}_1$ to MLWE^ρ. On challenge (A, T) set the common matrix to A and the first flow to $\mathbf{u}_A := T$. $\mathcal{B}_1$ draws Bob's $\mathbf{s}_B, \mathbf{e}_B, \mathbf{e}''_B, \bar{\mathbf{e}}_B$ itself, computes $\mathbf{u}_B = A\rho(\mathbf{s}_B)+\mathbf{e}_B$, $\mathbf{r}_B = \mathbf{u}_A\rho(\mathbf{s}_B)+\mathbf{e}''_B$, $(h_B, k_B) = \mathsf{PExt}(\mathbf{r}_B; \bar{\mathbf{e}}_B)$ and the tested key $K_B = \mathsf{KDF}(k_B \| \mathbf{u}_A \| \mathbf{u}_B)$ exactly as the protocol prescribes. Real $T \Rightarrow G_0$, uniform $T \Rightarrow G_1$, so $|S_0 - S_1| \le \mathsf{Adv}^{\mathsf{MLWE}^\rho}(\mathcal{B}_1)$.

Game G_2 (replace $(\mathbf{u}_B, \mathbf{r}_B)$ jointly by uniform). Sample $\mathbf{u}_B \leftarrow \mathrm{U}(M_k(R_q))$ *and* $\mathbf{r}_B \leftarrow \mathrm{U}(M_k(R_q))$ independently, then proceed $(h_B, k_B) = \mathsf{PExt}(\mathbf{r}_B; \bar{\mathbf{e}}_B)$, $K_B = \mathsf{KDF}(k_B \| \mathbf{u}_A \| \mathbf{u}_B)$. In G_1 both $\mathbf{u}_B = A\rho(\mathbf{s}_B)+\mathbf{e}_B$ and $\mathbf{r}_B = \mathbf{u}_A\rho(\mathbf{s}_B)+\mathbf{e}''_B$ are structured-MLWE samples of the *same* secret $\mathbf{s}_B$ against the two independent uniform matrices A and $\mathbf{u}_A$ (the latter uniform after G_1, and uncompressed in EE-DH); no single-sample hybrid can replace one while simulating the other, since the off-sample needs $\mathbf{s}_B$. This is the two-block instance of Definition 12 at

$d = \lceil \log_2 q \rceil$ (both block matrices uniform), i.e. MLWE^{ρ} with the stacked $2k \times k$ matrix $\left[\begin{smallmatrix} A \\ \mathbf{u}_A \end{smallmatrix}\right]$.

Reduction $\mathcal{B}_2$. On a two-block challenge $(A, \mathbf{B}, \mathbf{X}, \mathbf{Y})$ set the common matrix to A, the first flow to $\mathbf{u}_A := \mathbf{B}$ (uniform, as in G_1), $\mathbf{u}_B := \mathbf{X}$, $\mathbf{r}_B := \mathbf{Y}$; draw fresh $\bar{\mathbf{e}}_B$, compute $k_B = \mathsf{PExt}(\mathbf{r}_B; \bar{\mathbf{e}}_B)$ and $K_B = \mathsf{KDF}(k_B \| \mathbf{u}_A \| \mathbf{u}_B)$ — no knowledge of $\mathbf{s}_B$ is required. Real $\Rightarrow G_1$, uniform $\Rightarrow G_2$, so $|S_1 - S_2| \leq \mathsf{Adv}^{\mathsf{MLWE}^{\rho,(2)}}(\mathcal{B}_2) = \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_2)$ (stacked).

Game G_3 (key bits exactly uniform given the hint). In G_2 the input $\mathbf{r}_B$ is uniform and the doubling bits $\bar{\mathbf{e}}_B$ are honest fresh randomness, so by Lemma 1, applied to the $k^2 n$ entries, each key bit is exactly uniform and independent of the emitted hint bit; thus (h_B, k_B) has the same distribution as $(h_B, \text{uniform } k_B)$, $\Delta(G_2, G_3) = 0$, $S_2 = S_3$. After this hop Alice's recovered $k_A = \mathsf{PRec}(\mathbf{u}_B \rho(\mathbf{s}_A), h_B)$ receives a hint carrying no information about k_B.

Game G_4 (idealise the key). Replace K_B by uniform $K^* \leftarrow \{0,1\}^{N_{\mathrm{ss}}}$; in the ROM, with the now-uniform $k^2 n \geq 256$-bit k_B, distinguishing requires a KDF query at $k_B \| \mathbf{u}_A \| \mathbf{u}_B$, probability $\leq q_H / 2^{256}$. So $|S_3 - S_4| \leq q_H / 2^{256}$, and in G_4 the tested key is uniform and independent of the transcript ($b = 1$ world).

Summation.

$$\mathsf{Adv}^{\mathrm{KI}}_{\mathrm{EE}}(\mathcal{A}) \;\leq\; \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_1) + \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_2) + \frac{q_H}{2^{256}}.$$

Forward secrecy (PFS) is immediate: both ephemeral secrets are erased after the session (Algorithm 5, line 7), so a later $\mathsf{Corrupt}$ recovers no value entering K_B. No rounding assumption appears.

Remark 15 (Relation to the displayed "3 $\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}$"). The proof gives the *tighter* bound $2\,\mathsf{Adv}^{\mathsf{MLWE}^{\rho}} + q_H/2^{256}$: the second and third sketched hops form a single two-block instance, which Definition 12 identifies with one MLWE^{ρ} on a stacked matrix. The displayed $3\,\mathsf{Adv}^{\mathsf{MLWE}^{\rho}} + q_H/2^{256}$ is recovered as a valid (looser) upper bound by counting that two-block instance as two single-sample MLWE^{ρ} instances. Either way the reduction is tight (no T-factor), since the EE secrets are single-use.

4.7 Summary of NAIBI-DH Variants

Table 1 distills the trade-off space. Both offered variants share the same algebraic core — computing $\mathbf{t}'\rho(\mathbf{s})$ and applying the 4-quarter Peikert reconciliation — but differ in how keys are managed and what forward-secrecy guarantee follows. The static-static mode is *not* listed because it is excluded (Section 8.4).

NAIBI-ES-DH uses an ephemeral key on the initiator's side, yielding a tight (single-sample) reduction and weak forward secrecy: compromise of the server's long-term key does not reveal past session keys. This mirrors the TLS 1.3 static-server model.

NAIBI-EE-DH uses fresh ephemeral keys on both sides, achieving perfect forward secrecy at the cost of a second pass. The security reduction is again tight and single-sample.

Table 1. The two offered flavours of NAIBI ephemeral key agreement. "# msg." counts *key-agreement* messages *after* any long-term public key has been distributed. The static-static mode is excluded (Section 8.4).

Property	ES-DH (§4.5)	EE-DH (§4.6)
# msg. for KA	1 ($\bar{\mathbf{u}}_A + h_A$)	2 ($\mathbf{u}_A + (\mathbf{u}_B, h_B)$)
Hint size per session	$k^2 n$ bits (Alice)	$k^2 n$ bits (Bob)
Key model	ephemeral–static	ephemeral–ephemeral
Forward secrecy	weak (wFS)	perfect (PFS)
Reduction	MLWE^{ρ} (single)	MLWE^{ρ} (single)
Tightness	1 per session	1 per session
Use case	TLS-like static server	Noise NN

4.8 The Common Algebraic Skeleton

Both variants execute the *same* local computation: one party calls PExt on its DH product plus a fresh error and transmits the hint; the other uses PRec to recover the same raw key bits:

$$\begin{aligned} (h, k) &\leftarrow \mathsf{PExt}\big(\mathbf{t}\,\rho(\cdot)\mathbf{s} + \mathbf{e}''\!;\ \bar{\mathbf{e}}\big), \\ k' &\leftarrow \mathsf{PRec}\big(\mathbf{t}'\,\rho(\cdot)\mathbf{s}'\!,\ h\big), \\ K &\leftarrow \mathsf{KDF}(k \parallel \text{transcript}). \end{aligned} \tag{6}$$

What changes between variants is only *which* $\mathbf{t}$ stands in for the DH product: the ephemeral $\mathbf{u}_A$ in ES-DH (Bob's side) or $\mathbf{u}_B$ in EE-DH (Alice's side), and the long-term $\mathbf{t}'$ in KEM encryption. Implementations therefore share essentially all code paths.

5 NAIBI-KEM: Key Encapsulation as a By-Product

A key encapsulation mechanism follows from NAIBI-DH the same way hashed ElGamal follows from classical Diffie–Hellman: the receiver keeps a long-term DH key, the sender runs "one leg" of DH ephemerally, and both sides extract a key by KDF. We then apply the Fujisaki–Okamoto transform [11] on top of an underlying CPA-secure PKE to obtain IND-CCA2 security. NAIBI-KEM is therefore NAIBI-DH plus a thin wrapper.

5.1 The Underlying CPA-Secure PKE (NAIBI-PKE-OTP)

The PKE encrypts a 256-bit message $m \in \{0,1\}^{256}$ under a long-term DH-style public key by **XOR ing m against a hash of the rounded DH share**. This is the lattice analogue of hashed ElGamal [20]: the role of "g^{ab}" is played by $\mathsf{PExt}(\mathbf{t}\,\rho(\mathbf{s}'))[k] \in \{0,1\}^{4n}$, the role of the OTP key by its hash, and the ciphertext consists of the ephemeral $\bar{\mathbf{u}} = \mathsf{Compress}(\mathbf{u}, d_u)$, the 32-byte XOR mask $c = m \oplus \mathsf{H}(k)$, and the $4n$-bit reconciliation hint h.

Three properties of the construction matter below:

Algorithm 6 NAIBI-PKE-OTP.KeyGen() (same as NI-DH.KeyGen)

1: $\mathtt{seed}_A \leftarrow \{0,1\}^{256}$; $A \leftarrow \mathsf{XOF}(\mathtt{seed}_A) \in M_2(R_q)$
2: $\mathbf{s} \leftarrow \chi_\eta^2$ ($\in \mathcal{A}_\alpha$, resample if non-invertible); $\mathbf{e} \leftarrow \chi_\eta^{2\times 2}$ ($\in R_q^{2\times 2}$)
3: $\mathbf{t} \leftarrow A\,\rho(\mathbf{s}) + \mathbf{e} \in R_q^{2\times 2}$
4: $\bar{\mathbf{t}} \leftarrow \mathsf{Compress}(\mathbf{t}, d_t)$ ▷ public-key compression, d_t bits/coef
5: **return** $\big(\mathsf{pk} = (\mathtt{seed}_A, \bar{\mathbf{t}}),\ \mathsf{sk} = \mathbf{s}\big)$

Algorithm 7 NAIBI-PKE-OTP.Enc$\big(\mathsf{pk},\ m \in \{0,1\}^{256};\ r\big)$ (≡ "one leg" of NAIBI-DH plus a one-time pad)

1: Parse $\mathsf{pk} = (\mathtt{seed}_A, \bar{\mathbf{t}})$; $A \leftarrow \mathsf{XOF}(\mathtt{seed}_A) \in M_k(R_q)$; $\mathbf{t}' \leftarrow \mathsf{Decompress}(\bar{\mathbf{t}}, d_t) \in R_q^{k\times k}$
2: $(\mathbf{s}',\ \mathbf{e}',\ \mathbf{e}'',\ \bar{\mathbf{e}}) \leftarrow \mathsf{XOF}(r)$ ▷ coin expansion (Remark 7): ephemeral secret $\mathbf{s}' \in \mathcal{A}_\alpha$, errors $\mathbf{e}', \mathbf{e}'' \in R_q^{k\times k}$, doubling bits $\bar{\mathbf{e}}$ — all deterministic in r
3: $\mathbf{u} \leftarrow A\,\rho(\mathbf{s}') + \mathbf{e}' \in R_q^{k\times k}$ ▷ Alice's "ephemeral DH share"
4: $(h,\ k) \leftarrow \mathsf{PExt}\big(\mathbf{t}'\,\rho(\mathbf{s}') + \mathbf{e}'';\ \bar{\mathbf{e}}\big)$ ▷ $h \in \{0,1\}^{k^2 n}$ reconciliation hint; $k \in \{0,1\}^{k^2 n}$ raw key bits
5: $K \leftarrow \mathsf{H}(k) \in \{0,1\}^{256}$ ▷ the one-time pad
6: $c \leftarrow m \oplus K \in \{0,1\}^{256}$ ▷ OTP encryption
7: **return** $\mathsf{ct} \leftarrow \big(\mathsf{Compress}(\mathbf{u}, d_u),\ c,\ h\big)$

- the ciphertext is $\big(\mathsf{Compress}(\mathbf{u}, d_u),\ c,\ h\big)$ where $(h, k) = \mathsf{PExt}(\mathbf{t}'\,\rho(\mathbf{s}'))$, $c = m \oplus \mathsf{H}(k)$, and $h \in \{0,1\}^{4n}$ is the reconciliation hint;
- correctness reduces to *Peikert agreement* ($\mathsf{PRec}(\mathbf{u}'\,\rho(\mathbf{s}), h) = k$), the very same condition that governs NAIBI-DH (Theorem 2);
- the IND-CPA proof is a hashed-ElGamal-style argument (Theorem 6).

Binding correctness is established at the *KEM* level (Section 5.3), after the FO transform: the bare one-time pad does not itself satisfy a ciphertext-binding property (Remark 22).

Public hash. $\mathsf{H} : \{0,1\}^{4n} \to \{0,1\}^{256}$ is the hash function specified in Section 4.1 (SHA3-256), modelled as a random oracle. It maps the Peikert raw key bits $k \in \{0,1\}^{4n}$ (one bit per coefficient of the 2×2 matrix) into a 256-bit one-time pad. The exponent $4n$ reflects the 2×2 matrix structure: $4n$ coefficients in $R_q^{2\times 2}$.

Remark 16 (Direct correspondence with NAIBI-DH). Dropping the OTP step (line 6 of Enc and line 4 of Dec) reduces the construction to *exactly* NAIBI-ES-DH (Algorithm 4) under the identification $\mathbf{s}_A^{\mathrm{eph}} \leftrightarrow \mathbf{s}'$, $\mathsf{pk}_B \leftrightarrow \mathbf{t}$. Encryption computes $(h, k) = \mathsf{PExt}(\mathbf{t}'\,\rho(\mathbf{s}'))$, transmits $(\bar{\mathbf{u}}, c, h)$, and uses $K = \mathsf{H}(k)$ as the OTP; decryption recomputes $k' = \mathsf{PRec}(\mathbf{u}'\,\rho(\mathbf{s}), h)$ and hashes to K'; the XOR step $c = m \oplus K$ is the only thing that turns this shared key into a chosen-plaintext encryption. The PKE is therefore *literally* "NAIBI-ES-DH followed by a one-time pad", with the Peikert hint h playing the same role in both: bridging the noise gap so that the sender's k and the receiver's k' agree exactly.

Algorithm 8 NAIBI-PKE-OTP.Dec($\mathsf{sk} = \mathbf{s}$, ct)

1: Parse $\mathsf{ct} = (\bar{\mathbf{u}}, c, h)$; $\mathbf{u}' \leftarrow \mathsf{Decompress}(\bar{\mathbf{u}}, d_u) \in R_q^{k\times k}$
2: $k' \leftarrow \mathsf{PRec}\big(\mathbf{u}'\,\rho(\mathbf{s}),\ h\big)$ ▷ by (3), the A-terms cancel; h bridges the gap
3: $K' \leftarrow \mathsf{H}(k')$
4: **return** $m' \leftarrow c \oplus K'$

Tightened Decryption Noise

Lemma 8 (Rounding-agreement noise). *Let*

$$\mathbf{t}' = \mathsf{Decompress}(\mathsf{Compress}(\mathbf{t}, d_t), d_t), \quad \mathbf{u}' = \mathsf{Decompress}(\mathsf{Compress}(\mathbf{u}, d_u), d_u),$$

with compression errors $\boldsymbol{\varepsilon}_t := \mathbf{t}' - \mathbf{t}$ and $\boldsymbol{\varepsilon}_u := \mathbf{u}' - \mathbf{u}$ bounded entry-wise by E_{d_t} and E_{d_u} respectively, and let $\mathbf{e}'' \leftarrow \chi_\eta^{k\times k}$ be the encryptor's fresh reconciliation error (Remark 7). Write $W(k) := (k-1)\alpha + 1$ for the worst-column weight of $\rho(\mathbf{s})$ — the largest column sum of coefficient magnitudes in the α-circulant ρ of (2), attained at the last column (so $W(2) = \alpha + 1$ and $W(3) = 2\alpha + 1$; for $\alpha = 3$, $W(2) = 4$ and $W(3) = 7$). Then

$$\big\|\big(\mathbf{t}'\,\rho(\mathbf{s}') + \mathbf{e}''\big) \;-\; \mathbf{u}'\,\rho(\mathbf{s})\big\|_\infty \;\le\; W(k)\,n\eta\,E_{d_t} \;+\; W(k)\,n\eta\,E_{d_u} \;+\; 2\alpha k\,n\eta^2 \;+\; \eta, \tag{7}$$

where $E_d = \lfloor q/2^{d+1} \rfloor$. For $\alpha = 3$, $\eta = 1$ this gives $4nE_{d_t} + 4nE_{d_u} + 12n + 1$ when $k = 2$, and $7nE_{d_t} + 7nE_{d_u} + 18n + 1$ when $k = 3$. Consequently, whenever the right-hand side of (7) is $\le q/8 - 1$, the Peikert recovery yields $k = k'$ in

$$\mathsf{PRec}(\mathbf{u}'\,\rho(\mathbf{s}), h) = \mathsf{PExt}(\mathbf{t}'\,\rho(\mathbf{s}') + \mathbf{e}''; \bar{\mathbf{e}})[k],$$

hence $K = K'$ and decryption succeeds. (As elsewhere, correctness is in practice governed by the statistical noise model of Section 7, not by this worst case; the worst-case constants are recorded here only for completeness and the negligible-δ calibration uses (9).)

Proof. Expand the difference:

$$\begin{aligned}\big(\mathbf{t}'\rho(\mathbf{s}') + \mathbf{e}''\big) - \mathbf{u}'\rho(\mathbf{s}) &= (\mathbf{t} + \boldsymbol{\varepsilon}_t)\rho(\mathbf{s}') + \mathbf{e}'' - (\mathbf{u} + \boldsymbol{\varepsilon}_u)\rho(\mathbf{s}) \\ &= \big[\mathbf{t}\,\rho(\mathbf{s}') - \mathbf{u}\,\rho(\mathbf{s})\big] + \boldsymbol{\varepsilon}_t\,\rho(\mathbf{s}') - \boldsymbol{\varepsilon}_u\,\rho(\mathbf{s}) + \mathbf{e}''.\end{aligned}$$

By the commutativity identity (3) (applied as in Theorem 2), the bracket equals $\mathbf{e}\,\rho(\mathbf{s}') - \mathbf{e}'\,\rho(\mathbf{s})$, of entry-wise ℓ_∞-norm $\le 2\alpha k\,n\eta^2$ by Lemma 3. Each compression term contributes $\le W(k)\,n\eta\,E_{d_*}$ entry-wise: each entry is the sum of the k products of a row of $\boldsymbol{\varepsilon}$ (factors of norm $\le E_{d_*}$) against a column of $\rho(\mathbf{s})$, whose entries are u_i or $\alpha^j u_i$ (norm $\le \eta$ or $\le \alpha\eta$); using $\|fg\|_\infty \le n\|f\|_\infty\|g\|_\infty$ and summing the worst column's coefficient weights $W(k) = (k-1)\alpha + 1$ gives $n\eta\,E_{d_*}\,W(k)$ (which is $4n\eta\,E_{d_*}$ for $k = 2$ and $7n\eta\,E_{d_*}$ for $k = 3$, with $\alpha = 3$); the fresh error adds $\le \eta$. Triangle inequality gives the bound. Two $\mathbb{Z}_q$-values within $q/8 - 1$ of each other satisfy the Peikert recovery condition of Lemma 2.

Corollary 3 (Negligible decryption failure). *For $\eta = 1$, $\alpha \in \{2, 3\}$ and the parameters of Table 4, the per-coefficient recovery noise has standard deviation σ given by (9), and by (10)*

$$\delta \;=\; \Pr\big[\mathsf{PRec}(\mathbf{u}'\,\rho(\mathbf{s}), h) \neq k\big] \;\le\; k^2 n \cdot \operatorname{erfc}\!\Big(\tfrac{q/8-1}{\sigma\sqrt{2}}\Big) \;\le\; 2^{-128}.$$

Hence $K = K'$ and $m' = m$ except with negligible probability.

Remark 17 (KEM correctness ≡ agreement correctness — unified condition). The KEM and ephemeral-agreement correctness conditions are *both* governed by the same per-coefficient Gaussian tail (10), differing only in the variance (9): the EE-DH reconciliation between two ephemeral samples carries two compression terms of the same kind, while KEM/ES-DH carries one $\mathbf{t}$- and one $\mathbf{u}$-compression term. All three offered schemes (ES-DH, EE-DH, KEM) rely on the same Peikert recovery guarantee and the same threshold $q/8 - 1$; correctness is a single statistical statement, with no decoding step.

Remark 18 (Why the statistical bound is safe). The worst-case sup-norm bound (Lemma 3) aligns all n ternary products and grows as $O(n)$; it is essentially never approached by random χ_1-secrets, whose per-coefficient noise concentrates as $O(\sqrt{n})$. Equation (9) captures this $O(\sqrt{n})$ behaviour exactly, and the Gaussian tail (10) is the standard ML-KEM-style estimate (a central-limit approximation that is conservative in the tail for bounded ternary summands). With the moduli of Table 4 every instance satisfies $\delta \le 2^{-128}$, so a decapsulation/agreement failure is cryptographically negligible without requiring the worst case.

IND-CPA Security

Theorem 6 (CPA security of NAIBI-PKE-OTP). *For any PPT adversary $\mathcal{A}$ in the IND-CPA game making q_H queries to the random oracle H, there exist PPT $\mathcal{B}_1, \mathcal{B}_2$ such that*

$$\mathsf{Adv}^{\mathrm{CPA}}_{\mathrm{PKE\text{-}OTP}}(\mathcal{A}) \;\le\; \mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_1) + \mathsf{Adv}^{\mathsf{MLWE}^{\rho,(2;\mathcal{D}_{d_t})}}(\mathcal{B}_2) + \frac{q_H}{2^{256}} + \varepsilon_{\mathrm{inv}},$$

where $\varepsilon_{\mathrm{inv}} = \Pr[\mathbf{s}' \in \mathcal{A}_\alpha$ non-invertible$] \le n/q^k$ and $\mathsf{MLWE}^{\rho,(2;\mathcal{D}_{d_t})}$ is the two-block structured assumption of Definition 12. In the uncompressed option ($d_t = \lceil \log_2 q \rceil$, Remark 2) the second term is $\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}$ with a $2k \times k$ stacked matrix, so the whole bound rests on MLWE^{ρ} alone. No rounding assumption appears.

Proof. A five-game hybrid; the lattice analogue of the hashed-ElGamal CPA proof of Abdalla–Bellare–Rogaway [20].

Game 0. The IND-CPA game (real Enc, real $\bar{\mathbf{t}}$).

Game 1. Replace the key generator's $\mathbf{t} = A\rho(\mathbf{s}) + \mathbf{e}$ by uniform $\mathbf{U} \leftarrow M_k(R_q)$ and publish $\bar{\mathbf{t}} = \mathsf{Compress}(\mathbf{U}, d_t)$. By Remark 1 (compression is public post-processing of an MLWE^{ρ} sample), Games 0 and 1 are indistinguishable up to

Algorithm 9 NAIBI-KEM.KeyGen()

1: $(\mathsf{pk}, \mathbf{s}) \leftarrow$ NAIBI-PKE-OTP.KeyGen()
2: $z \leftarrow \{0,1\}^{256}$
3: **return** $\big(\mathsf{pk},\ \mathsf{sk} = (\mathbf{s}, \mathsf{pk}, \mathsf{H}(\mathsf{pk}), z)\big)$

$\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}(\mathcal{B}_1)$; the invertibility conditioning of the key generator costs an additional statistical ε_{inv}. The encryptor's local matrix $\mathbf{t}' = \mathsf{Decompress}(\bar{\mathbf{t}}, d_t)$ is now distributed exactly as $\mathcal{D}_{d_t}$.

Game 2. In the challenge, replace *jointly*

$$\mathbf{u} = A\rho(\mathbf{s}') + \mathbf{e}' \quad \text{and} \quad \mathbf{r} = \mathbf{t}'\rho(\mathbf{s}') + \mathbf{e}''$$

by independent uniforms $(\mathbf{U}_1, \mathbf{U}_2)$. Since A is uniform, $\mathbf{t}' \sim \mathcal{D}_{d_t}$, and $\mathbf{e}', \mathbf{e}''$ are independent fresh errors over the *same* secret $\mathbf{s}'$, this is precisely an instance of Definition 12: cost $\mathsf{Adv}^{\mathsf{MLWE}^{\rho,(2;\mathcal{D}_{d_t})}}(\mathcal{B}_2)$. (The two replacements must be done in one hop because they share $\mathbf{s}'$; performing them separately would be unsound.)

Game 3. With $\mathbf{r}$ uniform, Lemma 1 gives that the k^2n key bits k are *exactly* uniform and independent of the transmitted hint h (the doubling bits $\bar{\mathbf{e}}$ are honest randomness derived from the uniform challenge coins). Replacing k by uniform given h is therefore *free* (statistical distance 0) — no MLWR assumption is needed here, in contrast to the earlier version of this proof.

Game 4. Replace $K = \mathsf{H}(k)$ by uniform. In the ROM the adversary distinguishes only by querying H at the uniform $k \in \{0,1\}^{k^2n}$, probability $\leq q_H 2^{-k^2n} \leq q_H/2^{256}$ ($k^2n \geq 256$).

Game 5. K uniform, so $c = m \oplus K$ is uniform and independent of m; the advantage is 0. Summing the hops yields the bound.

Remark 19 (OTP simplicity). The proof is one structured-MLWE hop on the joint $(\mathbf{u}, \mathbf{r})$ followed by a *statistically free* key-bit hop (Lemma 1) and a random-oracle step on $\mathsf{H}(k)$, structurally a lattice analogue of the classical hashed-ElGamal proof [20]. The hint h in the ciphertext does not weaken the argument: by exact uniformity it carries no information about k.

5.2 The KEM via FO

NAIBI-KEM applies the FO$\not\perp$ transformation [11] (with implicit rejection) to NAIBI-PKE-OTP. Let $\mathsf{H} : \{0,1\}^* \to \{0,1\}^{256}$ (the same RO used inside the OTP-PKE, but applied here to ciphertexts and public keys), $\mathsf{G} : \{0,1\}^* \to \{0,1\}^{256} \times \{0,1\}^{256}$, $\mathsf{KDF} : \{0,1\}^* \to \{0,1\}^{256}$ be the hash functions specified in Section 4.1 (SHA3-256, SHA3-512, and SHAKE-256 respectively), modelled as (Q)ROs.

Remark 20 (The rejection key must be bound to pk*).* Two features of Algorithm 11 are load-bearing for Section 5.3 and are easy to lose in an implementation.

Algorithm 10 NAIBI-KEM.Encaps(pk)

$m \leftarrow \{0,1\}^{256}$
$(\bar{K}, r) \leftarrow \mathsf{G}(m \,\|\, \mathsf{H}(\mathsf{pk}))$
$\mathsf{ct} \leftarrow \mathsf{NAIBI\text{-}PKE\text{-}OTP.Enc}(\mathsf{pk}, m; r)$
$\mathsf{ss} \leftarrow \mathsf{KDF}_{N_{\mathrm{ss}}}(\mathtt{0x01} \,\|\, \bar{K} \,\|\, \mathsf{H}(\mathsf{ct}))$ ▷ $N_{\mathrm{ss}} = 256$; 0x01: accept-branch separator
return $(\mathsf{ct}, \mathsf{ss})$

Algorithm 11 NAIBI-KEM.Decaps(sk, ct)

Parse $\mathsf{sk} = (\mathbf{s}, \mathsf{pk}, h, z)$
$m' \leftarrow \mathsf{NAIBI\text{-}PKE\text{-}OTP.Dec}(\mathbf{s}, \mathsf{ct})$
$(\bar{K}', r') \leftarrow \mathsf{G}(m' \,\|\, h)$
$\mathsf{ct}' \leftarrow \mathsf{NAIBI\text{-}PKE\text{-}OTP.Enc}(\mathsf{pk}, m'; r')$
if $\mathsf{ct}' = \mathsf{ct}$ **then return** $\mathsf{KDF}_{N_{\mathrm{ss}}}(\mathtt{0x01} \,\|\, \bar{K}' \,\|\, \mathsf{H}(\mathsf{ct}))$
else return $\mathsf{KDF}_{N_{\mathrm{ss}}}(\mathtt{0x02} \,\|\, z \,\|\, \mathsf{H}(\mathsf{pk}) \,\|\, \mathsf{H}(\mathsf{ct}))$ ▷ implicit rejection; $\mathsf{H}(\mathsf{pk})$ *recomputed* from the pk field of sk, never the cached h
end if

(i) $\mathsf{H}(\mathsf{pk})$ *enters the rejection key, and is recomputed.* A rejection key of the form $\mathsf{KDF}(z \,\|\, \mathsf{H}(\mathsf{ct}))$ — the ML-KEM shape, and the one used in earlier versions of this work — is a deterministic function of (z, ct) alone. In the malicious-key model the adversary supplies z, so it may place the *same* z in two secret keys with $\mathsf{pk}_0 \neq \mathsf{pk}_1$ and submit any ciphertext that fails re-encryption under both — for instance the null ciphertext of Remark 22, which is guaranteed to fail under every key. Both decapsulations then return $\mathsf{KDF}(z \,\|\, \mathsf{H}(\mathsf{ct}^*))$: the malicious-key binding game is lost with probability 1, with no random-oracle query and independently of every lattice parameter. This is exactly why ML-KEM misses MAL-BIND-K-PK even in the seed key format [44], and it has nothing to do with the reconciliation mechanism. Feeding $\mathsf{H}(\mathsf{pk})$ into the rejection branch makes the two KDF inputs differ whenever $\mathsf{pk}_0 \neq \mathsf{pk}_1$, which is what Case (B) of Theorem 11 needs. The value must be recomputed from the pk field of sk: the cached h is unverified, and an adversary would simply plant $h_0 = h_1$. The cost is one H call on the rejection path only; implementations may equivalently verify $h = \mathsf{H}(\mathsf{pk})$ once at key import and then reuse h.

(ii) Domain separation. The separators 0x01/0x02 keep the accept- and reject-branch KDF inputs in disjoint prefixes, so a cross-branch agreement is a random-oracle collision and nothing else.

What is *not* changed: KeyGen still caches $h = \mathsf{H}(\mathsf{pk})$ and the FO coins $\mathsf{G}(m' \,\|\, h)$ still consume it unverified, so the adversary retains the freedom Schmieg exploits, and Case (A) — the statistical statement about the reconciliation mechanism, which is the subject of this paper — remains reachable and is proved under that freedom. The patch concerns the rejection branch only, where no statistical guarantee is available to anyone (Remark 24).

Theorem 7 (IND-CCA2, ROM). *For an adversary making q_H hash and q_D decapsulation queries:*

$$\mathsf{Adv}^{\mathrm{CCA2}}_{\mathrm{KEM}}(\mathcal{A}) \;\le\; \frac{q_H}{|\mathcal{M}|} + \sqrt{q_H \cdot \mathsf{Adv}^{\mathrm{CPA}}_{\text{PKE-OTP}}} + q_D\,\delta.$$

The additive correctness term $q_D\,\delta$ is bounded by $q_D \cdot 2^{-128}$ (Cor. 3) and is therefore negligible for any polynomially-bounded q_D; the FO⊥ transform of [11] is correctness-preserving up to exactly this term.

Theorem 8 (IND-CCA2, QROM).

$$\mathsf{Adv}^{\mathrm{CCA2}}_{\mathrm{KEM}}(\mathcal{A}) \;\le\; \frac{q_H^2}{|\mathcal{M}|} + q_H\sqrt{\mathsf{Adv}^{\mathrm{CPA}}_{\text{PKE-OTP}}},$$

following [12,13].

5.3 Binding Correctness

Binding correctness is a statistical, *decapsulation-level* property that complements the hash-based binding profile of ML-KEM (Section 5.4): no adversary can construct one ciphertext that decapsulates to the same shared secret under two distinct KEM key pairs. In the X-BIND taxonomy of Cremers–Dax–Medinger [43], where X-BIND-P-Q asks that the components in P determine those in Q, one ciphertext reaching the same shared secret under two *distinct public keys* is a break of the K-PK axis; the K-CT axis asks instead that two *distinct ciphertexts* never reach the same shared secret, and is settled for NAIBI-KEM by the presence of $\mathsf{H}(\mathsf{ct})$ in both KDF inputs (Remark 23). The substantive axis here is therefore K-PK, and it is the one on which the reconciliation geometry does the work. We state and prove it directly for the *KEM*, where the FO re-encryption check and implicit rejection are available, rather than for the bare PKE one-time pad. The reason for that scoping is a concrete attack on any PKE-level formulation; we isolate the attack first (Remark 22), then give the KEM statement it motivates.

Definition 8 (KEM binding correctness game, honest keys — HON-BIND-K-PK). *Game* $\mathsf{Bind}(\lambda)$ *for the KEM: (1) Challenger generates* $(\mathsf{pk}_0,\mathsf{sk}_0),(\mathsf{pk}_1,\mathsf{sk}_1) \leftarrow \mathsf{KEM.KeyGen}$ *($\mathsf{pk}_0 \neq \mathsf{pk}_1$) and gives* $(\mathsf{pk}_0,\mathsf{pk}_1)$ *to $\mathcal{A}$. (2) $\mathcal{A}$ outputs* ct^**. (3) Challenger computes* $\mathsf{ss}_b \leftarrow \mathsf{KEM.Decaps}(\mathsf{sk}_b,\mathsf{ct}^*)$*, $b \in \{0,1\}$ (under implicit rejection ss_b is always defined, never $\perp$). (4) $\mathcal{A}$ wins if* $\mathsf{ss}_0 = \mathsf{ss}_1$*. The KEM has* ε-binding correctness *if* $\Pr[\mathcal{A}\ \textit{wins}] \le \varepsilon$ *for all PPT $\mathcal{A}$.*

Remark 21 (Why ct^* is left unrestricted). One might be tempted to require in addition that ct^* be "not honestly generated under either key". That restriction would remove the substantive case: Case (A) of Theorem 10 and clause (a) of Corollary 4 treat precisely the ciphertexts that *are* honest under one key, which is the only way both decapsulations reach the accept branch. The definition therefore places no condition on ct^*, and the game is the honest-key specialisation of Definition 9.

Remark 22 (Why binding is a KEM property here, not a bare-PKE property). A PKE-level formulation of binding — one ciphertext decrypting to the same *message* under two honest PKE secret keys — is *not* achievable for NAIBI-PKE-OTP, and it is instructive to see why, because the same gap is exactly what the FO transform closes. Decryption (Algorithm 8) never outputs $\perp$ and computes $m'_b = c \oplus \mathsf{H}\big(\mathsf{PRec}(\mathbf{u}'\,\rho(\mathbf{s}_b), h)\big)$ with $\mathbf{u}' = \mathsf{Decompress}(\bar{\mathbf{u}}, d_u)$. An adversary may choose $\bar{\mathbf{u}} = \mathsf{Compress}(\mathbf{0}, d_u)$; since $\mathsf{Compress}_q(0, d) = 0$ and $\mathsf{Decompress}_q(0, d) = 0$, this yields $\mathbf{u}' = \mathbf{0}$, hence $\mathsf{PRec}(\mathbf{0} \cdot \rho(\mathbf{s}_b), h) = \mathsf{PRec}(\mathbf{0}, h)$ *independently of* $\mathbf{s}_b$. Then $m'_0 = m'_1$ with probability 1 under *any* two keys: a bare-PKE binding statement fails outright, and no slot-counting bound over honest secrets can rescue it, because $\mathbf{u}'$ is adversarial, not drawn from any distribution.

This is precisely why we formulate binding at the KEM level. Under the FO transform (Section 5.2), decapsulation re-encrypts the recovered message and accepts only if the re-encryption reproduces ct^* exactly; on failure it returns the rejection key of Algorithm 11. The null ciphertext above re-encrypts, under the honest encryption coins, to some $\bar{\mathbf{u}} = \mathsf{Compress}(A\rho(\mathbf{s}') + \mathbf{e}', d_u)$ with $\mathbf{u} \neq \mathbf{0}$, so it *fails* the re-encryption test under both keys and is mapped to $\mathsf{KDF}(\mathsf{0x02} \,\|\, z_b \,\|\, \mathsf{H}(\mathsf{pk}_b) \,\|\, \mathsf{H}(\mathsf{ct}^*))$, $b \in \{0, 1\}$, whose KDF inputs differ because $\mathsf{pk}_0 \neq \mathsf{pk}_1$. The attack therefore does not break KEM binding; it merely shows that the property must be stated for the KEM, which is the object of interest in the X-BIND taxonomy of [43] in any case.

The second half of that sentence is not free, and it is worth being explicit about what it costs. Under honest key generation the two rejection keys differ because z_0, z_1 are independent and uniform, and $\mathsf{H}(\mathsf{pk})$ plays no role. Under *malicious* keys that argument evaporates: the adversary chooses z_b, and with the ML-KEM rejection shape $\mathsf{KDF}(z \,\|\, \mathsf{H}(\mathsf{ct}))$ it need only set $z_0 = z_1$ and submit this very null ciphertext to win with probability 1, with no random-oracle query. The null ciphertext thus attacks not only the PKE-level statement but any KEM-level statement whose rejection branch is pk-oblivious. Algorithm 11 is specified accordingly (Remark 20), and this is what makes Theorem 11 true as stated.

Remark 23 (The $\mathsf{K\text{-}CT}$ *axis is settled by the* KDF *interface).* Both branches of Algorithm 11 feed $\mathsf{H}(\mathsf{ct})$ into KDF. Two decapsulations with $\mathsf{ct}_0 \neq \mathsf{ct}_1$ therefore query the random oracle on distinct points, and agreement is a KDF/H collision: NAIBI-KEM is $\mathsf{MAL\text{-}BIND\text{-}K\text{-}CT}$ with advantage $\leq q_H^2\, 2^{-257}$, by the same generic argument that gives seed-format ML-KEM the property [44]. This is a computational statement in the ROM, and nothing stronger is available: for any KEM with 256-bit shared secrets the map $\mathsf{ct} \mapsto \mathsf{Decaps}(\mathsf{sk}, \mathsf{ct})$ is a deterministic function into $\{0, 1\}^{256}$ on a domain far larger than its range, so an adversary unbounded in random-oracle queries always finds a colliding pair. *No* KEM has a statistical $\mathsf{K\text{-}CT}$ guarantee, and the interesting question on that axis is closed before the mechanism is reached. What the reconciliation geometry buys, and what the rest of this section is about, is the $\mathsf{K\text{-}PK}$ axis, where the two decapsulations may share a ciphertext and the KDF interface alone decides nothing.

Remark 24 (No statistical guarantee on the rejection branch, for any KEM). On the rejection path the shared secret is a deterministic function $F(\mathsf{sk}, \mathsf{ct})$ of

adversarially chosen material, with 256-bit output. An adversary unbounded in time but bounded in random-oracle queries can enumerate offline only through KDF, so the best one can hope for is a collision term in q_H; an adversary unbounded in *oracle* queries wins outright, since $F(\cdot, \mathsf{ct}^*)$ maps a domain of size $\gg 2^{256}$ into $\{0,1\}^{256}$. The statistical modality claimed in this paper is therefore necessarily restricted to the accept branch, where the collision event is decided by the geometry of PRec and not by the output length of a hash — and where, correspondingly, the bound $\beta_q^{\lceil n/2 \rceil}$ has nothing to do with 2^{-256} (Remark 26). Theorem 11 is stated as a hybrid on purpose: statistical in Case (A), random-oracle in Case (B). Claiming an unconditional 2^{-256} in Case (B) would amount to reinstating honest key generation inside the malicious game.

Heuristic 9 (Slot-wise model for honest ternary secrets). *Let $\mathbf{s}_0, \mathbf{s}_1 \leftarrow \chi_\eta(\mathcal{A}_\alpha)$ be independent honestly generated secrets conditioned on $\mathbf{s}_0 \neq \mathbf{s}_1$, and $\Delta\mathbf{s} = \mathbf{s}_0 - \mathbf{s}_1$. Under the NTT decomposition $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$ we model: (i) the events $\{\Delta\mathbf{s}^{(i)} = 0\}$ as independent across slots, each of probability at most $2/q$; (ii) at every slot i with $\Delta\mathbf{s}^{(i)} \neq 0$, the pair $(\rho(\mathbf{s}_0)^{(i)}, \rho(\mathbf{s}_1)^{(i)})$ as uniform on its support and independent across slots. This is the standard NTT-domain heuristic used in ML-KEM-style correctness analyses [7,10]; we validate it empirically in Appendix D (10^7 sampled pairs per parameter set; per-slot vanishing frequencies and pairwise slot correlations within sampling error).* This heuristic governs honest secrets only; it is invoked solely in Case (A) of Theorem 10, where the ciphertext is constrained to be honest under one of the two keys, and plays no role in the adversarial Case (B), which rests on implicit rejection alone.

Lemma 9 (Hint-independent per-coefficient collision bound). *Fix any hint bit $h \in \{0,1\}$. For $X, Y \leftarrow \mathrm{U}(\mathbb{Z}_q)$ independent, $\Pr[\mathsf{rec}(2X, h) = \mathsf{rec}(2Y, h)] \leq \frac{1}{2} + \frac{8}{q^2} \leq \frac{3}{4}$.*

Proof. $x \mapsto \mathsf{rec}(2x, h)$ is a two-valued function on $\mathbb{Z}_q$ whose level sets are unions of arcs of total cardinalities qp_h and $q(1-p_h)$ with $|p_h - 1/2| \leq 2/q$ (the arcs of $\mathbb{Z}_{2q}$ have cardinalities $\lfloor q/2 \rfloor$ or $\lceil q/2 \rceil$ and $x \mapsto 2x$ is a bijection onto the even residues). The collision probability is $p_h^2 + (1-p_h)^2 = \frac{1}{2} + 2(p_h - \frac{1}{2})^2 \leq \frac{1}{2} + 8/q^2$, uniformly in h, hence for an adversarially chosen hint.

Theorem 10 (KEM binding correctness of NAIBI-KEM). *Under Heuristic 9(i), NAIBI-KEM has ε-binding correctness (Definition 8) with*

$$\varepsilon \;\leq\; \beta_q^{\lceil n/2 \rceil} + (8/q)^{n/2} + 2^{-256}, \qquad \beta_q = \tfrac{2}{3} + \tfrac{1}{3q} \text{ (Lemma 11).}$$

Concretely, for any ciphertext $\mathsf{ct}^* = (\bar{\mathbf{u}}, c, h)$ (the hint h included) and any two honestly-generated KEM key pairs with secrets $\mathbf{s}_0 \neq \mathbf{s}_1$ in $\mathcal{A}_\alpha^{\times}$,

$$\Pr_{(\mathbf{s}_0,\mathbf{s}_1)\leftarrow\mathsf{KeyGen}}\big[\mathsf{Decaps}(\mathsf{sk}_0, \mathsf{ct}^*) = \mathsf{Decaps}(\mathsf{sk}_1, \mathsf{ct}^*)\big] \;\leq\; \beta_q^{\lceil n/2 \rceil} + (8/q)^{n/2} + 2^{-256},$$

where the first term is the raw Peikert-bit agreement on honest re-encryptions, the second the slot-deficiency tail, and the third a random-oracle (KDF/H) collision. For $n \geq 512$ (all NAIBI parameter sets) this is below 2^{-148}, hence well below 2^{-128}. The adversary may be unbounded except for the random-oracle term.

Proof. We do *not* route through any PKE-level decryption statement (which fails, Remark 22). Decapsulation (Algorithm 11) under each key either accepts (the FO re-encryption reproduces ct^*) or rejects. Fix ct^* and honest $\mathsf{sk}_0, \mathsf{sk}_1$, and split exhaustively on the joint outcome.

Case (B): ct^* is rejected under at least one key. Say it is rejected under sk_b. Implicit rejection returns $\mathsf{ss}_b = \mathsf{KDF}(\mathtt{0x02} \,\|\, z_b \,\|\, \mathsf{H}(\mathsf{pk}_b) \,\|\, \mathsf{H}(\mathsf{ct}^*))$ with z_b uniform and independent of everything in the adversary's view (Algorithm 9, line 2), while ss_{1-b} is some value (accept or reject) determined without access to z_b. Then $\mathsf{ss}_0 = \mathsf{ss}_1$ forces a collision on the 256-bit KDF/H output involving the independent uniform z_b, of probability $\leq 2^{-256}$. In particular the null-ciphertext attack of Remark 22 lands here, since its re-encryption produces $\mathbf{u} \neq \mathbf{0}$ and so is rejected under both keys. Note that this step, and this step only, is where honest key generation is used: it is the uniformity of z_b that yields the 2^{-256}, not the structure of the rejection key. In the malicious model z_b is adversarial and the argument runs instead through $\mathsf{H}(\mathsf{pk}_b)$ (Theorem 11, Case (B)).

Case (A): ct^* is accepted under *both* keys. Acceptance under sk_0 means $\mathsf{ct}^* = \mathsf{Enc}(\mathsf{pk}_0, m_0'; r_0')$ is *honestly formed* under pk_0, with $m_0' = \mathsf{Dec}(\mathsf{sk}_0, \mathsf{ct}^*)$ and FO-derived coins r_0'. Both decapsulations are then on the accept branch, so $\mathsf{ss}_0 = \mathsf{ss}_1$ requires equal pre-KDF keys, which (by correctness, up to the δ-failure already counted in Theorem 7) forces the underlying raw Peikert keys to agree: $k_0' = k_1'$, where $k_b' = \mathsf{PRec}(\mathbf{u}'\, \rho(\mathbf{s}_b), h)$ and $\mathbf{u}' = \mathsf{Decompress}(\bar{\mathbf{u}}, d_u)$ is the *honest* ephemeral share $\mathbf{u} = A\rho(\mathbf{s}') + \mathbf{e}'$ fixed by the accepted ciphertext. Because the ciphertext is honest under sk_0, its share $\mathbf{u}'$ is the image of an honestly sampled $(\mathbf{s}', \mathbf{e}')$; in particular $\mathbf{u}'$ is non-zero on all but a negligible fraction of slots, and we may invoke Heuristic 9 on the honest secrets $(\mathbf{s}_0, \mathbf{s}_1)$.

Write $\mathbf{d} := \mathbf{u}'\, \rho(\Delta\mathbf{s}) \in R_q^{k \times k}$ with $\Delta\mathbf{s} := \mathbf{s}_0 - \mathbf{s}_1 \neq \mathbf{0}$, and work slot-by-slot in $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$ (Definition 3).

Slot-count tail. By Heuristic 9(i) and a binomial tail, the number of slots with $\Delta\mathbf{s}^{(i)} \neq 0$ is at least $n_0 := \lceil n/2 \rceil$ except with probability $\leq \binom{n}{n/2}(2/q)^{n/2} \leq (8/q)^{n/2}$ (using $\binom{n}{n/2} \leq 2^n$), $\leq 2^{-217}$ already at NAIBI-I; we absorb this into the second additive term. By non-splitness of α, $\mathbb{F}_{q^k}$ is a field, so each non-zero $\Delta\mathbf{s}^{(i)}$ is invertible and $\mathbf{u}'^{(i)} \mapsto \mathbf{u}'^{(i)} \rho(\Delta\mathbf{s})^{(i)}$ is a bijection of $M_k(\mathbb{F}_q)$.

Per-slot collision bound (single coordinate). At each of the $\geq n_0$ non-vanishing slots i, agreement of the raw keys requires in particular agreement on *coordinate* $(0,0)$ of the slot: $f_h(w_0) = f_h(w_1)$ with $w_b = \left(\mathbf{u}'^{(i)} \rho(\mathbf{s}_b)^{(i)}\right)_{0,0} = \langle (\mathbf{u}'^{(i)})_{0,:},\, \mathrm{col}_0(\rho(\mathbf{s}_b)^{(i)}) \rangle$. Both columns are nonzero ($\mathbf{s}_b \in \mathcal{A}_\alpha^\times$, so the slot values are units of $\mathbb{F}_{q^k}$) and distinct ($\Delta\mathbf{s}^{(i)} \neq 0$). Over the uniform $(\mathbf{u}'^{(i)})_{0,:}$: if the two columns are linearly independent, (w_0, w_1) is uniform on $\mathbb{F}_q^2$ and Lemma 9 gives collision probability $\leq \frac{1}{2} + 8/q^2$; if dependent, $w_1 = \lambda w_0$ with $\lambda \neq 1$ and Lemma 11 gives $\leq \beta_q = \frac{2}{3} + \frac{1}{3q}$. In both cases the slot collides with probability $\leq \beta_q$, *uniformly over the slot values and the hint* — so no independence or uniformity of the k^2 intra-slot entries is invoked (Heuristic 9(ii) is not needed), and the bound holds for the worst-case slot values, only the slot count of part (i) being heuristic here.

Union over slots. The slot values $\mathbf{u}'^{(i)}$ of the honest ephemeral share are independent uniform across slots (bijective NTT change of basis). Across the $\geq n_0 = \lceil n/2 \rceil$ non-vanishing slots,

$$\Pr[k'_0 = k'_1 \mid \mathsf{ct}^* \text{ honest under } \mathsf{sk}_0,\ \mathbf{s}_0 \neq \mathbf{s}_1] \ \leq\ \beta_q^{n_0} \ \leq\ \beta_q^{\lceil n/2 \rceil}.$$

Combining. A binding win requires either Case (B) (prob. $\leq 2^{-256}$) or Case (A) with $k'_0 = k'_1$ (prob. $\leq \beta_q^{\lceil n/2 \rceil}$), plus the slot-deficiency tail $(8/q)^{n/2}$, giving the three-term bound. For all NAIBI parameter sets ($\lceil n/2 \rceil \in \{256, 512, 1024\}$), $\beta_q^{\lceil n/2 \rceil} \leq (2/3)^{256} \approx 2^{-149.7}$, well below 2^{-128}. We stress what is and is not heuristic: Case (B), the per-coordinate bounds (Lemmas 9 and 11, the latter established by exhaustive evaluation at the deployed moduli) and slot-wise invertibility are assumption-free on the keys; only the slot-count modelling of honest secrets in Case (A) — Heuristic 9(i), replaceable by the unconditional Lemma 10 of Section 5.5 — is heuristic, exactly as in the correctness analyses of [7,10].

Remark 25 (Status of the heuristic). After the single-coordinate repair above, the only ingredient of Theorem 10 that models the honest keys is Heuristic 9(i) (the slot count), and it enters *only* in Case (A); part (ii) is no longer used, the per-slot bound of Lemmas 9/11 holding for worst-case slot values. Case (B) (implicit rejection) is unconditional, and part (i) itself can be replaced by the assumption-free Lemma 10 (Section 5.5). We validate the heuristic empirically in Appendix D and conjecture that the true Case (A) collision probability is at most a small constant factor larger than the bound above.

Remark 26 (Why the $\beta_q^{\lceil n/2 \rceil}$ term is present and not replaced by 2^{-256}). In Case (A) it is tempting to argue that, since the KDF/hash output is 256 bits, agreement should hold up to 2^{-256}. This conflates two *disjoint* events. Agreement of the two decapsulations needs the pre-KDF keys to coincide, which happens either because (I) the *raw* Peikert keys already coincide, $k'_0 = k'_1$, or (II) they differ but the hash collides. The random oracle bounds only event (II), at 2^{-256}. In event (I) the hash agrees *trivially*, so idealising the hash buys nothing: what stands between the adversary and a Case (A) break is that two distinct secret keys are unlikely to recover the *same* raw value from an honest re-encryption. That is exactly the $\beta_q^{\lceil n/2 \rceil}$ term, and it dominates at $n = 512$. The relevant quantity is the *agreement probability of two decoders on the accepted honest share*, bounded per slot by $\beta_q = \frac{2}{3} + \frac{1}{3q}$ (not $\frac{1}{2}$: the two slot inputs are correlated) via Lemma 11. Dropping the $\beta_q^{\lceil n/2 \rceil}$ term would be *unsound*. At NAIBI-I, $(2/3)^{256} \approx 2^{-149.7}$ dominates 2^{-256}, so we report the honest (larger) figure 2^{-149} as the operative bound.

Corollary 4 (Binding from the PKE viewpoint). *Equivalently, the FO transform* lifts *the (otherwise unachievable) PKE collision-freeness to a KEM guarantee: any ciphertext that decapsulates to the same shared secret under two distinct honest KEM key pairs must (a) be honest under one key and have its honest re-encryption recovered identically under the other (prob. $\leq \beta_q^{\lceil n/2 \rceil} + (8/q)^{n/2}$*

by Theorem 10, Case (A), under Heuristic 9(i)), or (b) be implicitly rejected under some key and collide on the KDF/H *output (prob.* $\leq 2^{-256}$*, Case (B)). Hence KEM binding holds with bound* $\leq \beta_q^{\lceil n/2 \rceil} + (8/q)^{n/2} + 2^{-256}$*, with no appeal to a bare-PKE binding property.*

5.4 Relation to Robustness and Key-Binding Notions

The property of Theorem 10 should be situated in the recent literature on KEM binding rather than presented as simply "absent elsewhere". In the terminology of Grubbs–Maram–Paterson [41], our notion is a statistical *collision-freeness* guarantee for honestly generated keys (SCFR-style): no single ciphertext *decapsulates* to the same shared secret under two distinct honest secret keys. Like ML-KEM, NAIBI-KEM uses implicit rejection, so decapsulation never outputs $\perp$ and $\perp$-robustness notions are vacuous; collision-freeness is the meaningful target for implicit-rejection KEMs, as [41] observe for Kyber-style schemes. We emphasise (Remark 22) that the guarantee is genuinely a KEM property: at the level of the bare PKE one-time pad a single adversarial null ciphertext decrypts identically under all keys, and it is the FO re-encryption check that lifts the property to the KEM. In the X-BIND taxonomy of Cremers–Dax–Medinger [43], Corollary 4 yields HON-BIND-K-PK and HON-BIND-K,CT-PK (honest-key models): the derived shared secret binds the key pair that produced it. The $\mathsf{K\text{-}CT}$ axis — two *distinct* ciphertexts reaching one shared secret — is a separate and, for this scheme, generic matter (Remark 23).

On the $\mathsf{K\text{-}PK}$ axis we go further: Section 5.5 removes the honest-distribution hypothesis of Heuristic 9 entirely and establishes $\mathsf{MAL\text{-}BIND\text{-}K\text{-}PK}$ with no distributional assumption on the adversarial keys (Theorem 11), the very notion [44] proves ML-KEM does *not* satisfy — and which the seed key format of FIPS 203, which does repair the $\mathsf{K\text{-}CT}$ axis, still leaves unattained. The two heuristic ingredients are replaced by assumption-free lemmas: a minimum-NTT-weight bound (Lemma 10, a finite first-moment condition verified computationally with exponential margin) shows that an adversarially chosen $\Delta\mathbf{s} \neq 0$ — intuitively, coefficient-domain-sparse but NTT-domain dense — cannot vanish on a majority of slots, so the slot-counting argument retains its $\lceil n/2 \rceil$ nonzero slots; and a correlated per-slot collision bound (Lemma 11, established by exhaustive evaluation at the deployed moduli) replaces Lemma 9 for adversarial slot values. The statistical modality attaches to the accept branch; the rejection branch is a random-oracle statement and cannot be otherwise (Remark 24). The LEAK model is covered a fortiori, MAL being the strongest of the three adversary classes of [43].

One misreading is worth heading off: binding is not simply "absent" in ML-KEM. As mapped by [43] and Schmieg [44], ML-KEM (FIPS 203) achieves several binding properties *generically*, by hashing the public key into the key schedule ($\bar{K} = \mathsf{G}(m \,\|\, \mathsf{H}(\mathrm{ek}))$), while failing others in the MAL model [44]. Collision-freeness of FO-based lattice KEMs is, moreover, already known: Grubbs–Maram–Paterson [41] and Xagawa [42] establish it for NTRU, BIKE, FrodoKEM and HQC. Two axes separate that line of work from ours. First, their guarantee is

computational, proved in the QROM, whereas ours is statistical. Second, in the X-BIND taxonomy of [43] their collision-freeness sits on the robustness (CT-PK) axis — binding two *keys* to one ciphertext — while our guarantee is on the K-PK axis, binding the *shared secret* to the public key, and is lifted to the malicious-key model (Section 5.5), the regime Schmieg [44] shows ML-KEM does not reach. What NAIBI adds is thus different in kind: a *statistical, decapsulation-level* K-PK collision-freeness that does not rest on hashing the transcript, holds against unbounded adversaries (up to the RO term and Heuristic 9), and therefore survives where hash-based transcript binding is unavailable or insufficient — e.g. anonymity- and auditability-sensitive deployments [41], and the protocol-level attacks from non-binding KEMs automatically found in [43], which motivate binding as a first-class design goal.

5.5 From Honest to Malicious Keys: MAL-BIND-K-PK

Section 5.4 deliberately confined the guarantee of Theorem 10 to the honest-key models HON-BIND-K-PK and HON-BIND-K, CT-PK, for two reasons: the Case (A) slot-counting argument invoked Heuristic 9, which models the *honestly sampled* secrets, and the Case (B) argument used the uniformity of z_b, which honest key generation supplies and a malicious adversary does not. We now remove both hypotheses and obtain a MAL-BIND-K-PK guarantee with no distributional assumption on the adversarial keys — precisely where Schmieg [44] shows ML-KEM falls short. The second hypothesis is removed by the specification of the rejection key (Remark 20), not by an argument: no argument is possible there (Remark 24). For the first, two lemmas replace the two roles of Heuristic 9: Lemma 10 replaces part (i) (slot-deficiency under an adversarial $\Delta\mathbf{s}$), and Lemma 11 handles the per-slot collision for correlated, adversarially chosen slot values (a single coordinate per slot suffices, so part (ii) is not needed at all). Neither lemma assumes anything about the key distribution — their own epistemic status (finite computational verification) is stated in Remark 31 — and only an NTT-orthogonality fact about the *honest* ephemeral share survives from the honest analysis, being a property of the accepted ciphertext, not of the adversarial keys.

We work throughout in $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ with $R_q = \mathbb{Z}_q[x]/(x^n + 1)$, n a power of two, $q \equiv 1 \pmod{2n}$. Writing w for a primitive $2n$-th root of unity in $\mathbb{F}_q$, the NTT evaluation points are $\zeta_i = w^{2i+1}$, $i = 0, \ldots, n-1$ (the roots of $x^n + 1$). For $\Delta\mathbf{s} = \sum_{j=0}^{k-1} P_j(x) y^j \in \mathcal{A}_\alpha$, slot i vanishes iff $P_j(\zeta_i) = 0$ for all j; let $Z(\Delta\mathbf{s})$ count the vanishing slots. A difference of two ternary secrets has all coefficients in $\{-2, \ldots, 2\}$; we call any $\Delta\mathbf{s}$ with coefficient vector in $\{-2, \ldots, 2\}^{kn}$ *ternary-realizable*.

Minimum NTT-weight of ternary differences.

Lemma 10 (Minimum NTT-weight). *Let $B = 5$. If*

$$\binom{n}{m} B^{kn} < q^{km}, \tag{8}$$

then no nonzero ternary-realizable $\Delta\mathbf{s} \in \mathcal{A}_\alpha$ *has* $Z(\Delta\mathbf{s}) \geq m$*. In particular, if* (8) *holds at* $m = \lceil n/2 \rceil$*, then every nonzero ternary-realizable* $\Delta\mathbf{s}$ *satisfies* $Z(\Delta\mathbf{s}) < n/2$*.*

Proof. A first-moment (union-bound) argument.

(i) Vanishing on a fixed slot-set is a codimension-km linear condition. Fix $S \subseteq \{0, \ldots, n-1\}$, $|S| = m$. For $\Delta\mathbf{s}$ with coefficient vector $c = (c_{j,e}) \in \mathbb{Z}_q^{kn}$, vanishing on S reads $\sum_{e=0}^{n-1} c_{j,e}\zeta_i^e = 0$ in $\mathbb{F}_q$ for all $i \in S$, $j \in \{0, \ldots, k-1\}$ — that is km linear forms. For each fixed j the m rows $(\zeta_i^e)_{e<n}$, $i \in S$, form a Vandermonde block in the *distinct* nodes $\{\zeta_i\}_{i\in S}$ (distinct because w is primitive), hence have rank m; blocks for different j occupy disjoint coordinate ranges. The $\mathbb{F}_q$-solution space V_S therefore has dimension $kn - km$.

(ii) Counting ternary-realizable points in V_S*.* Let $\mathcal{B} = \{-2, \ldots, 2\}^{kn} \subset \mathbb{Z}_q^{kn}$ (the entries are distinct residues since $q \geq 5$). The surjective linear map $L_S : \mathbb{F}_q^{kn} \to \mathbb{F}_q^{km}$ given by the km forms has fibres of equal size q^{kn-km}; $V_S = L_S^{-1}(0)$ is the fibre over 0. By the character-sum identity

$$|\mathcal{B} \cap V_S| = \frac{1}{q^{km}} \sum_{\chi \in (\mathbb{F}_q^{km})^\vee} \prod_{(j,e)} \Big(\sum_{c \in \{-2,\ldots,2\}} \psi\big(\chi \cdot \ell^{(j,e)}\, c\big) \Big),$$

where ψ is a fixed nontrivial additive character of $\mathbb{F}_q$ and $\ell^{(j,e)}$ the (j,e)-column of L_S. The $\chi = 0$ term equals $B^{kn}/q^{km} = 5^{kn}/q^{km}$ exactly; each $\chi \neq 0$ term is bounded in modulus, via $|\sum_{c\in\{-2,\ldots,2\}} \psi(\theta c)| \leq 5$ and the linear independence of the forms, by a product that sums (geometrically) to $o(5^{kn}/q^{km})$. Hence $|\mathcal{B} \cap V_S| \leq (1 + o(1))\, 5^{kn}/q^{km}$.

(iii) Union over slot-sets. A nonzero ternary-realizable $\Delta\mathbf{s}$ with $Z(\Delta\mathbf{s}) \geq m$ vanishes on some m-subset of its vanishing slots, so

$$\#\{\Delta\mathbf{s} \in \mathcal{B} \setminus \{0\} : Z(\Delta\mathbf{s}) \geq m\} \;\leq\; \sum_{|S|=m} |\mathcal{B} \cap V_S| \;\leq\; (1 + o(1)) \binom{n}{m} \frac{5^{kn}}{q^{km}}.$$

Under (8) the right-hand side is below 1; being a non-negative integer it is 0. Taking $m = \lceil n/2 \rceil$ gives $Z(\Delta\mathbf{s}) < n/2$.

Remark 27 (Numerical margin). Taking logarithms, (8) reads $\log_2 \binom{n}{m} + kn \log_2 5 < km \log_2 q$, whose left side peaks near $m = n/2$ at $\leq n + kn \log_2 5$. For the NAIBI sets the bound is below 1 already for all $m \geq m_0$ with m_0 far below $n/2$: $m_0 = 96$ (NAIBI-I, k=2, n=512, q=19457), $m_0 = 100$ (NAIBI-III), $m_0 = 178$ (NAIBI-V). At the majority threshold $m = \lceil n/2 \rceil$ the right-hand side of the displayed bound is 2^{-4410}, 2^{-4130}, 2^{-9916} respectively. Thus not merely $Z(\Delta\mathbf{s}) < n/2$, but $Z(\Delta\mathbf{s}) < 0.2\, n$ for every nonzero ternary-realizable $\Delta\mathbf{s}$ — in quantitative agreement with the independent lattice-reduction estimate of the critical zeroable fraction $f^* = \ln(2\sqrt{2\pi e})/\ln q \approx 0.20$–$0.23$ (verified by BKZ on the slot-vanishing lattice L_S). Exhaustive enumeration in small dimension (n=4, k=2, q=17) confirms the count: for every m at which the first-moment bound drops below 1, the exact number of nonzero ternary-realizable $\Delta\mathbf{s}$ vanishing on m slots is 0.

Per-slot collision under malicious correlated values. We strengthen Lemma 9. In the malicious model the adversary chooses, at each slot, the two slot values $g_0 \neq g_1 \in \mathbb{F}_q$ of $\rho(\mathbf{s}_0), \rho(\mathbf{s}_1)$ *and* the hint h; the honest ephemeral slot value $\mathbf{u}'^{(i)} =: u'$ is uniform and (by non-splitness) invertible. The recoverer sees the value $z \in \mathbb{Z}_q$ as $2z \in \mathbb{Z}_{2q}$; write $f_h(z) := \mathsf{rec}(2z, h) \in \{0,1\}$. The per-slot collision is $\Pr_{u' \sim \mathrm{U}(\mathbb{F}_q)}[f_h(u'g_0) = f_h(u'g_1)]$, a function of the two *correlated* inputs $u'g_0, u'g_1$, unlike the independent pair of Lemma 9.

Lemma 11 (Malicious per-slot collision). *Let q be one of the deployed moduli of Table 4 ($q \in \{13\,313,\ 19\,457,\ 40\,961\}$). For every $g_0 \neq g_1$ in $\mathbb{F}_q^\times$ and every $h \in \{0,1\}$,*

$$\Pr_{u' \sim \mathrm{U}(\mathbb{F}_q)}\big[f_h(u'g_0) = f_h(u'g_1)\big] \;\leq\; \beta_q := \tfrac{2}{3} + \tfrac{1}{3q}.$$

The maximum over (g_0, g_1, h) equals $\frac{2}{3} - \frac{1}{3q} < \frac{2}{3}$ at these moduli (all $\equiv 2 \pmod 3$) and is attained at $g_1 \equiv -\frac{1}{3}\, g_0$ (equivalently $g_0 = -3g_1$). We conjecture $\max_{g_0 \neq g_1, h} \leq \frac{2}{3} + \frac{1}{3q}$ for every odd prime $q \geq 5$; it is verified exhaustively for $q \in \{97, 257, 1009, 13\,313, 19\,457, 40\,961\}$ (Remark 28), which covers every modulus at which the lemma is invoked. If $g_0 = 0$ or $g_1 = 0$ the collision probability can be 1, but invertible secrets ($\mathbf{s}_b \in \mathcal{A}_\alpha^\times$, enforced by KeyGen) have all slot values nonzero, so that case does not arise at the $\geq \lceil n/2 \rceil$ nonzero slots of Lemma 10.

Proof. Since $g_0 \in \mathbb{F}_q^\times$, $X := u'g_0$ is uniform on $\mathbb{F}_q$; with $r := g_1 g_0^{-1} \neq 1$ we bound $p(r,h) = \Pr_X[f_h(X) = f_h(rX)]$.

(i) Structure of f_h. By the exact-uniformity property of randomised doubling (Lemma 9 and its proof), f_h is the indicator of a *single* cyclic arc $I_h \subset \mathbb{F}_q$ of cardinality $(q \pm 1)/2$ (verified along with the scan below); with $\varphi_h := 1 - 2f_h \in \{\pm 1\}$ one has $|\frac{1}{q}\sum_X \varphi_h(X)| \leq 1/q$.

(ii) Collision is a dilation correlation. From $\mathbf{1}[f_h(X) = f_h(rX)] = \frac{1}{2}(1 + \varphi_h(X)\varphi_h(rX))$,

$$\begin{aligned} p(r,h) &= \tfrac{1}{2} + \tfrac{1}{2}\, C(r,h), \\ C(r,h) &= \tfrac{1}{q} \sum_{X \in \mathbb{F}_q} \varphi_h(X)\varphi_h(rX) = \tfrac{4}{q}\,|I_h \cap r^{-1} I_h| - 1 + O(1/q), \end{aligned}$$

so bounding p is bounding the overlap of the arc I_h with its dilate $r^{-1}I_h$ (an arithmetic progression of step r^{-1}).

(iii) Exhaustive evaluation. $p(r,h)$ is a finite function of (q,r,h); we evaluate it *exactly*, for all $r \in \mathbb{F}_q \setminus \{0,1\}$ and both h, at each of the six moduli listed in the statement (an $O(q^2)$ computation included in the artifact scripts). The maximum is $\frac{2}{3} + \varepsilon_q$ with $|\varepsilon_q| = \frac{1}{3q}$ (sign + iff $q \equiv 1 \bmod 3$), attained at $r \equiv -\frac{1}{3} \bmod q$ in every case (Remark 28); all deployed moduli have $q \equiv 2 \bmod 3$, whence the strict bound $\frac{2}{3} - \frac{1}{3q}$. Since the lemma is only ever applied at the deployed moduli, this finite verification establishes it where used; a general proof of the $\frac{2}{3}$ extremal-overlap claim (e.g. via the three-distance theorem) remains open. A natural guess

— that the extremum sits at the dilations $r = 2^{\pm 1}$, with value $\frac{3}{4} + \Theta(1/q)$ — is ruled out by the same evaluation, which gives $p(2^{\pm 1}, h) = \frac{1}{2} + O(1/q)$ there.

Remark 28 (Numerical confirmation). Exhaustive maximisation of $p(r, h)$ over $r \in \mathbb{F}_q \setminus \{0, 1\}$, $h \in \{0, 1\}$, with f_h computed exactly from Definition 2, gives 0.670103 (q=97), 0.665370 (q=257), 0.667000 (q=1009), 0.666642 (q=13 313), 0.666650 (q=19 457), 0.666659 (q=40 961), i.e. $\frac{2}{3} \pm \frac{1}{3q}$, with the maximizer at $r \equiv -3^{-1} \bmod q$ in every case (e.g. $r = 32$ at $q = 97$: $32 \cdot 3 \equiv -1$; $r = 12\,971$ at $q = 19\,457$; $r = 27\,307$ at $q = 40\,961$). For comparison, the dilation $r = 2^{-1}$ — the natural first guess — yields only $\frac{1}{2} + O(1/q)$ (0.50515 at q=97, 0.50005 at q=19 457). The independent-pair value of Lemma 9 is likewise confirmed at $\frac{1}{2} + O(1/q^2)$.

Malicious-key binding.

Definition 9 (Malicious-key binding — MAL-BIND-K-PK, after [43]). *In game* $\mathsf{MalBind}(\lambda)$ *the adversary outputs two key pairs* $(\mathsf{pk}_0, \mathsf{sk}_0)$, $(\mathsf{pk}_1, \mathsf{sk}_1)$ *with* $\mathsf{pk}_0 \neq \mathsf{pk}_1$ *— generated by any means, subject only to* sk_b *parsing as a well-formed secret key (for NAIBI:* $\mathsf{sk}_b = (\mathbf{s}_b, \mathsf{pk}_b, h_b, z_b)$ *with* $\mathbf{s}_b \in \mathcal{A}_\alpha^\times$*, no consistency required between* h_b *and* $\mathsf{H}(\mathsf{pk}_b)$*, nor between* pk_b *and* $\mathbf{s}_b$*, and* z_b *chosen freely, including* $z_0 = z_1$*) — together with two ciphertexts* $\mathsf{ct}_0, \mathsf{ct}_1$*, not necessarily distinct. It wins if* $\mathsf{Decaps}(\mathsf{sk}_0, \mathsf{ct}_0) = \mathsf{Decaps}(\mathsf{sk}_1, \mathsf{ct}_1)$*. The KEM is* ε-MAL-BIND-K-PK *if every adversary wins with probability at most* ε*. The companion notion* MAL-BIND-K-CT *is the same game with the winning condition* $\mathsf{ct}_0 \neq \mathsf{ct}_1$ *in place of* $\mathsf{pk}_0 \neq \mathsf{pk}_1$ *(Remark 23).*

Remark 29 (Reading the definition). Three points, since the axis labels of [43] are easy to transpose. *(i)* X-BIND-P-Q reads "P binds Q": the adversary must make the components of P agree while some component of Q differs. With $P = \{\mathsf{K}\}$ and $Q = \{\mathsf{PK}\}$ the win is one shared secret reached under two distinct public keys, which is what Definitions 8 and 9 formalise. *(ii)* The single-ciphertext case $\mathsf{ct}_0 = \mathsf{ct}_1 = \mathsf{ct}^*$ is the substantive one: as soon as $\mathsf{ct}_0 \neq \mathsf{ct}_1$ the $\mathsf{H}(\mathsf{ct})$ in the KDF input reduces the game to a random-oracle collision (Remark 23), which is also the whole of the K-CT axis. *(iii)* z_b is adversarial, exactly like h_b; the reader used to honest key generation should resist substituting a uniform z_b anywhere in the analysis, and Remark 24 says what that costs.

Theorem 11 (MAL-BIND-K-PK of NAIBI-KEM). *Suppose* (8) *holds at* $m = \lceil n/2 \rceil$ *(Lemma 10), as it does with exponential margin at every NAIBI parameter set (Remark 27), and let* Decaps *be Algorithm 11 (rejection key bound to* $\mathsf{H}(\mathsf{pk})$*, Remark 20). Then NAIBI-KEM is* MAL-BIND-K-PK *in the sense of Definition 9: for any adversary making at most* q_H *random-oracle queries and otherwise computationally unbounded, that outputs two key pairs with* $\mathsf{pk}_0 \neq \mathsf{pk}_1$ *and secrets* $\mathbf{s}_0 \neq \mathbf{s}_1$ *in* $\mathcal{A}_\alpha^\times$ *together with* $\mathsf{ct}_0, \mathsf{ct}_1$*,*

$$\Pr\big[\mathsf{Decaps}(\mathsf{sk}_0, \mathsf{ct}_0) = \mathsf{Decaps}(\mathsf{sk}_1, \mathsf{ct}_1)\big] \;\le\; \underbrace{q_H\, \beta_q^{\lceil n/2 \rceil}}_{\text{Case (A), statistical}} \;+\; \underbrace{q_H^2\, 2^{-256}}_{\text{Case (B), ROM}},$$

$$\beta_q = \tfrac{2}{3} + \tfrac{1}{3q} \text{ (Lemma 11).}$$

The same bound with the first term deleted holds for MAL-BIND-K-CT *(Remark 23). For NAIBI-I/III ($n = 512$) the mechanism term is $\beta_q^{256} \approx 2^{-149.7}$, meeting a 2^{-128} target for $q_H \le 2^{21}$; for NAIBI-V ($n = 1024$) it is $\beta_q^{512} \approx 2^{-299}$ and is negligible at any q_H for which the generic collision term — which every KEM pays, and which caps the statement at $q_H \lesssim 2^{64}$ for a 2^{-128} target — is itself meaningful. Remark 30 discusses sizing the mechanism term at Category 1.*

Two scoping remarks. The statistical modality applies to Case (A), and only there: Case (B) is a random-oracle statement and cannot be anything else (Remark 24). The result holds with the cached $\mathsf{H}(\mathsf{pk})$ *retained in* sk *and consumed unverified by the FO coins — the format that breaks ML-KEM [44] — and it is on the axis that the seed key format of FIPS 203 still misses.*

Remark 30 (Sizing the mechanism term). The q_H factor is not an artefact: having fixed $(\mathbf{s}_0, \mathbf{s}_1)$, the adversary knows both secrets and can therefore *test* candidate messages offline, each random-oracle query $\mathsf{G}(m \,\|\, h)$ yielding fresh coins, a fresh honest $\mathbf{u}'$ and an independent trial of the slot-wise collision. (No such union bound is needed in Theorem 10: there the secrets are the challenger's, and the adversary cannot evaluate the event it is trying to cause.) A 2^{-128} target against q_H offline trials requires $\lceil n/2 \rceil \log_2(1/\beta_q) \ge 128 + \log_2 q_H$ with $\log_2(1/\beta_q) \approx 0.585$: satisfied with room to spare at $n = 1024$, and at $n = 512$ for $q_H \le 2^{21}$. Deployments wanting the mechanism term to dominate at Category 1 for larger q_H should raise the number of extracted coefficients rather than the ring degree; the compact profile states this trade-off in the coefficient domain, where the bound reads $q_H \beta_q^{N_{key}}$ and $N_{key} = 640$ costs 48 B of ciphertext.

Proof. We reuse the Case (A)/Case (B) split of Theorem 10, replacing the two heuristic ingredients by Lemmas 10 and 11.

Reduction to a common ciphertext. If $\mathsf{ct}_0 \ne \mathsf{ct}_1$ then the two KDF inputs contain $\mathsf{H}(\mathsf{ct}_0) \ne \mathsf{H}(\mathsf{ct}_1)$ unless H collides, and agreement is then a KDF collision on distinct points: $\le q_H^2\, 2^{-257}$ for the two events together. Assume henceforth $\mathsf{ct}_0 = \mathsf{ct}_1 =: \mathsf{ct}^*$.

Case (B): ct^* *rejected under at least one key.* This is *not* identical to Theorem 10, and the difference is the whole content of the malicious model on this branch: z_0, z_1 are chosen by the adversary and may coincide, so no entropy is available in z_b and what separates the two decapsulations is the public key. If ct^* is rejected under both keys the outputs are $\mathsf{KDF}(\mathtt{0x02} \,\|\, z_b \,\|\, \mathsf{H}(\mathsf{pk}_b) \,\|\, \mathsf{H}(\mathsf{ct}^*))$, and since $\mathsf{pk}_0 \ne \mathsf{pk}_1$ the two KDF inputs are distinct unless H collides on $\mathsf{pk}_0, \mathsf{pk}_1$; if ct^* is rejected under exactly one key, the two inputs carry distinct domain separators and are distinct unconditionally. In both situations agreement is a random-oracle collision on two distinct points, and over the q_H queries available to the adversary

$$\Pr[\text{Case (B) agreement}] \;\le\; q_H^2\, 2^{-257}.$$

With the pk-oblivious rejection key $\mathsf{KDF}(z \,\|\, \mathsf{H}(\mathsf{ct}^*))$ the same event would have probability 1 — set $z_0 = z_1$ and take for ct^* the null ciphertext of Remark 22 — so this case is settled by the specification of Algorithm 11 and not by an argument; none is possible (Remark 24).

Case (A): ct^* *accepted under both keys.* Acceptance forces ct^* to be honestly formed under one key, fixing an honest ephemeral share $\mathbf{u}' = A\rho(\mathbf{s}') + \mathbf{e}'$ that is uniform and, by non-splitness, invertible on all but a negligible fraction of slots. On the accept branch the output is $\mathsf{KDF}(\mathtt{0x01} \,\|\, \bar{K}'_b \,\|\, \mathsf{H}(\mathsf{ct}^*))$ with $(\bar{K}'_b, r'_b) = \mathsf{G}(m'_b \,\|\, h_b)$, so agreement forces $\bar{K}'_0 = \bar{K}'_1$ and hence $(m'_0, h_0) = (m'_1, h_1)$ up to a random-oracle collision already charged above; in particular $h_0 = h_1$ — the planted cached hash, which the adversary is free to choose — and $m'_0 = m'_1$. Since $m'_b = c \oplus \mathsf{H}(k'_b)$ with the same payload c, agreement of the two decapsulations requires the raw Peikert keys to coincide at every slot, in particular on *coordinate* 0 of every slot. By Lemma 10, the adversarial $\Delta\mathbf{s} = \mathbf{s}_0 - \mathbf{s}_1 \neq 0$ vanishes on fewer than $n/2$ slots, so at least $\lceil n/2 \rceil$ slots have $\Delta\mathbf{s}^{(i)} \neq 0$. Fix such a slot i and consider coordinate 0 of the row product: $w_b := \big(\mathbf{u}'^{(i)}\rho(\mathbf{s}_b)^{(i)}\big)_0 = \langle \mathbf{u}'^{(i)}, \mathbf{c}_b \rangle$ with $\mathbf{c}_b := \mathrm{col}_0\big(\rho(\mathbf{s}_b)^{(i)}\big) \in \mathbb{F}_q^k$. Since $\mathbf{s}_b \in \mathcal{A}_\alpha^\times$ and $\mathbb{F}_{q^k}$ is a field, both $\mathbf{c}_0, \mathbf{c}_1$ are nonzero, and $\Delta\mathbf{s}^{(i)} \neq 0$ gives $\mathbf{c}_0 \neq \mathbf{c}_1$. Two exhaustive subcases, over the honest uniform $\mathbf{u}'^{(i)} \in \mathbb{F}_q^k$: if $\mathbf{c}_0, \mathbf{c}_1$ are $\mathbb{F}_q$-linearly *independent*, (w_0, w_1) is uniform on $\mathbb{F}_q^2$ and the collision probability is $\leq \frac{1}{2} + 8/q^2$ by Lemma 9; if they are *dependent*, $\mathbf{c}_1 = \lambda\mathbf{c}_0$ with $\lambda \in \mathbb{F}_q^\times \setminus \{1\}$, so w_0 is uniform on $\mathbb{F}_q$, $w_1 = \lambda w_0$, and the collision probability is $\leq \beta_q$ by Lemma 11, *uniformly* over the adversarial slot values and hint. In both subcases the coordinate-0 collision at slot i has probability $\leq \beta_q$, and the all-coordinate collision at that slot is bounded by it a fortiori (no intra-slot independence is assumed). Across distinct NTT slots the values $\mathbf{u}'^{(i)}$ are independent uniform (the NTT is a bijective change of basis and the honest $\mathbf{u}'$ is uniform up to the pseudorandomness already accounted for); this is a property of the *honest* ephemeral share fixed by the accepted ciphertext, not of the adversarial $(\mathbf{s}_0, \mathbf{s}_1)$, and therefore needs no distributional assumption on the keys. Taking the product over the $\geq \lceil n/2 \rceil$ nonzero slots,

$$\Pr[\text{Case (A) collision}] \;\leq\; \beta_q^{\lceil n/2 \rceil}.$$

The adversary may repeat this experiment: having fixed $(\mathbf{s}_0, \mathbf{s}_1)$ it can test candidate messages offline, each random-oracle query yielding a fresh honest $\mathbf{u}'$ and an independent trial, so a union bound over the at most q_H trials gives $q_H\, \beta_q^{\lceil n/2 \rceil}$ (Remark 30). Adding the random-oracle terms of the reduction step and of Case (B) gives the bound. At the deployed moduli β_q may even be replaced by the strict $\frac{2}{3} - \frac{1}{3q}$ (Lemma 11); numerically $\beta_q^{\lceil n/2 \rceil} \leq (2/3)^{256} \approx 2^{-149.7}$ at $n = 512$ and $\leq (2/3)^{512} \approx 2^{-299}$ at $n = 1024$. One might be tempted to raise this bound to the k-th power per slot by treating the k coordinate collisions as independent. That step is unsound: the $2k$ coordinate values are functions of only k uniform coordinates of $\mathbf{u}'^{(i)}$. We therefore keep the single-coordinate bound above.

Remark 31 (What the upgrade does and does not assume). Neither lemma makes a distributional assumption on the adversarial keys, but their epistemic statuses should be stated precisely. Lemma 10 is a finite first-moment inequality whose hypothesis (8) is a closed-form numerical condition met with exponential slack;

its equidistribution step is verified computationally (Remark 27), by exact enumeration in small dimension and by BKZ on L_S. Lemma 11 is a statement about the reconciliation geometry established by *exhaustive exact evaluation* at every modulus at which it is invoked (Remark 28), with the general-q form stated as a conjecture. The sole ingredient inherited from the honest analysis is the slot-independence of the honest ephemeral $\mathbf{u}'$ used in Case (A); it is a property of the NTT and of $\mathbf{u}'$ — which is honest even in the malicious-*key* model, the accepted ciphertext being well-formed under one key — and not of the adversarial secrets. A third ingredient is not inherited but *specified*: the rejection key of Algorithm 11 hashes pk, without which Case (B) collapses (Remarks 20, 24). Consequently Theorem 11 upgrades the binding guarantee from HON to MAL on the $\mathsf{K\text{-}PK}$ axis, the property that [44] shows ML-KEM lacks and that the seed key format does not restore, with no assumption on the key distribution; the $\mathsf{K\text{-}CT}$ axis is settled generically by the KDF interface (Remark 23). Heuristic 9 is thereby *no longer needed* for $\mathsf{K\text{-}PK}$ binding: part (i) is superseded by Lemma 10 and part (ii) by the worst-case per-slot bound of Lemma 11, which also tightens the honest-key constants of Theorem 10.

Remark 32 (The cached $\mathsf{H}(\mathsf{pk})$: where it is tolerated and where it is not). KeyGen (Algorithm 9) stores $h = \mathsf{H}(\mathsf{pk})$ inside sk and the FO coins $\mathsf{G}(m' \,\|\, h)$ consume the cached value without recomputation — the very design choice through which Schmieg [44] constructs his counterexample for ML-KEM (a malicious key pair whose cached hash is inconsistent with the embedded public key). NAIBI-KEM keeps it, deliberately: it is what makes Case (A) reachable, and Theorem 11 is proved with $h_0 = h_1$ planted by the adversary. On the accept branch the mechanism, not the key schedule, carries the guarantee, and that is the claim of this paper.

The rejection branch is the exception, and the asymmetry is not cosmetic. There the secret key contributes only z_b, which the adversary also chooses; no geometric argument is available, and $\mathsf{H}(\mathsf{pk})$ must enter the rejection key, recomputed from the pk embedded in sk rather than read from the cache (Remark 20). Implementations MAY instead verify $h = \mathsf{H}(\mathsf{pk})$ once at key import and reuse h thereafter; what they MUST NOT do is revert the rejection branch to $\mathsf{KDF}(z \,\|\, \mathsf{H}(\mathsf{ct}))$, which loses $\mathsf{MAL\text{-}BIND\text{-}K\text{-}PK}$ completely — not by a degraded bound but by a probability-1 attack requiring no random-oracle query — while leaving every other property of the scheme intact. It is the one place where NAIBI-KEM's binding profile is a serialisation matter, and we make it a normative requirement of Algorithm 11 rather than an implementation note.

6 Reconciliation Asymmetry and the Threshold–Binding Trade-Off

The malicious-key binding guarantee of Theorem 11 rests on a *single* property of the reconciliation map: the per-slot collision constant $\beta_q \le \frac{2}{3} + \frac{1}{3q}$ of Lemma 11. This section develops the full study behind the one-line design choice "one-bit

Peikert reconciliation": it shows that the $\frac{2}{3}$ ceiling is *fragile* — destroyed by every natural attempt to raise the correctness threshold of the mechanism, and absent from every other standard reconciliation or encoding mechanism we have examined — and that the obstruction is, in all cases, the same *reflection attack*. The conference versions retain only the classification of Table 3 and the degradation theorem (Theorem 12); the evaluation framework, the generalised Peikert-t theory, the reflection-attack proofs and the per-mechanism analysis appear here.

6.1 Evaluation framework

We rank candidate reconciliation mechanisms against six requirements, in decreasing order of non-negotiability (Table 2). The discriminating requirement is **P3**, the per-slot malicious collision ceiling, which is new: it is exactly the quantity β_q that Lemma 11 bounds by $\frac{2}{3} + \frac{1}{3q}$ for one-bit Peikert reconciliation, and it is the property on which MAL-BIND-K-PK (Theorem 11) stands or falls.

Table 2. Evaluation requirements for a reconciliation mechanism inside NAIBI, in decreasing order of non-negotiability. P3 is the new, discriminating requirement.

#	Requirement	Stake for NAIBI
P1	Exact key-bit uniformity under randomisation (Lemma 1)	Deterministic IND-CCA2 correctness; rounding-assumption-free proofs
P2	FO-compatible derandomisation (the hint reproduces the bit)	FO re-encryption check (Section 5.2)
P3	Malicious per-slot collision ceiling β (Lemma 11)	MAL-BIND-K-PK (Theorem 11)
P4	Threshold/σ_{KEM} ratio and hint cost	Bandwidth and failure probability δ
P5	Realistic constant-time on a 16-bit datapath	Safe implementation
P6	Support for the contributory ES-/EE-DH modes	Ephemeral key agreement (Section 4)

One-bit Peikert reconciliation satisfies P1–P6. The rest of this section shows that the mechanisms competing on P4 (a higher correctness threshold, hence potentially smaller q or larger noise budget) forfeit P3, and that the mechanisms sharing an $x \mapsto -x$ symmetry forfeit P3 outright.

6.2 Generalised Peikert reconciliation with t hint bits

The natural way to raise the correctness threshold of randomised doubling is to transmit a finer cross-rounding hint.

Definition 10 (Peikert-t). *Let $t \geq 1$ and q odd. Randomised doubling is unchanged, $\bar{v} = 2v - \bar{e}_v$ with $\bar{e}_v \in \{-1, 0, 1\}$ balanced. The t-*bit cross-rounding *and the key bit are*

$$c_t(\bar{x}) = \left\lfloor \tfrac{2^t \bar{x}}{q} \right\rceil \bmod 2^t, \qquad b(\bar{x}) = \left\lfloor \tfrac{\bar{x}}{q} \right\rceil \bmod 2, \qquad \bar{x} \in \mathbb{Z}_{2q},$$

and $\mathsf{rec}_t(\bar{w}, h)$ outputs $b(\bar{x}^*)$ for $\bar{x}^*$ the point of $\{\bar{x} : c_t(\bar{x}) = h\}$ closest to $\bar{w}$ (ties toward 0). The case $t = 1$ is Definition 2.

Lemma 12 (Threshold and exact uniformity of Peikert-t). *For every $t \geq 1$ and odd q:* (i) *if $|w - v| \leq \theta_t - 1$ with $\theta_t = \frac{q}{4}\big(1 - 2^{-t}\big)$, then $\mathsf{rec}_t\big(2w,\, c_t(\mathrm{dbl}(v))\big) = b(\mathrm{dbl}(v))$;* (ii) *for v uniform on $\mathbb{Z}_q$, the key bit $b(\mathrm{dbl}(v))$ is* exactly *uniform given $c_t(\mathrm{dbl}(v))$.*

Proof. (i) The class $\{c_t = h\}$ consists of exactly two antipodal arcs of length $q/2^t$ in $\mathbb{Z}_{2q}$; the nearest adverse point of the same class lies at distance $\geq q(1 - 2^{-t})/2$ from the interior worst case, and doubling halves distances in $\mathbb{Z}_q$. The worst case is $2w = 2v - 1$ (antipodal to the arc edge), which leaves distance $q(1 - 2^{-t})/2$ to the first adverse point. For $t = 1$ this recovers the familiar $q/8$; as $t \to \infty$, $\theta_t \to q/4$, the additive-encoding limit. We verified the thresholds exactly at $q = 1009$: $\theta_1 = 126$, $\theta_2 = 189$, $\theta_3 = 221$ against $q(1 - 2^{-t})/4 = 126.1,\ 189.2,\ 220.7$. (ii) The map $\bar{x} \mapsto \bar{x} + q$ is a bijection of $\mathbb{Z}_{2q}$ that preserves each class $\{c_t = h\}$ setwise (it exchanges its two arcs, integer translates by q, hence of *equal* cardinality) and flips b; since $\mathrm{dbl}(v)$ is exactly uniform on $\mathbb{Z}_{2q}$ (Lemma 1), both values of b are equiprobable within each class.

Peikert-t therefore looks strictly better than Peikert-1 on P1 and P4: the correctness threshold grows from $q/8$ towards $q/4$, at a hint cost of t bits per extracted coefficient, and exact uniformity (P1) is retained. The binding constant P3, however, moves in the opposite direction.

6.3 The degradation law and the reflection attack

Theorem 12 (Binding degradation of Peikert-t). *Let $f_h^{(t)}(w) = \mathsf{rec}_t(2w, h)$ and, for $r \in \mathbb{F}_q^\times \setminus \{1\}$, let $p_t(r, h) = \Pr_{X \sim \mathrm{U}(\mathbb{F}_q)}[f_h^{(t)}(X) = f_h^{(t)}(rX)]$. Then*

$$\beta_t \;:=\; \max_{r \neq 1,\, h} p_t(r, h) \;\geq\; 1 - 2^{-t} - O(1/q),$$

the maximum being attained at the reflection $r = -1$ for $t \geq 2$. Consequently a MAL-BIND-K-PK *bound of the form $q_H\, \beta_t^{\lceil n/2 \rceil} + q_H^2\, 2^{-256} \leq 2^{-128}$ requires $\lceil n/2 \rceil \log_2 \frac{1}{1 - 2^{-t}} \geq 128$: at $n = 512$ this holds only for $t = 1$.*

Proof. The set $(f_h^{(t)})^{-1}(1)$ is a single arc I of length $(q \pm 1)/2$ whose position, as h ranges over $\{0, \ldots, 2^t - 1\}$, moves in steps of $\approx q/2^{t+1}$. Under the reflection $r = -1$ the collision probability is $1 - 4\delta/q + O(1/q)$, where δ is the offset of the centre of I from a fixed point of $x \mapsto -x$; choosing the best h gives

$\delta \le q/2^{t+2}$, whence $p_t(-1,h) \ge 1 - 2^{-t} - O(1/q)$. The adversary realises $r = -1$ by taking $\mathbf{s}_1 = -\mathbf{s}_0$ (both ternary, invertible iff $\mathbf{s}_0$ is), a globally admissible choice: every slot then satisfies $g_1 = -g_0$. Exhaustive evaluation at $q = 1009$ confirms sharpness:

t	1	2	3
$\max_{r,h} p_t$	0.6670	0.7502	0.8751
attained at	$r \equiv -\frac{1}{3}$	$r = -1$	$r = -1$
$1 - 2^{-t}$	0.5	0.75	0.875

At $t = 1$ the arc positions ($h \in \{0, 1\}$) are the two *asymmetric* arcs produced by the half-integer randomisation, neither of which can be centred on a reflection fixed point: $\Pr[f_h(X) = f_h(-X)] = \frac{1}{2} + O(1/q)$ (verified at $q \in \{97, 257, 1009, 13313\}$), and the true extremum is the milder dilation $r \equiv -\frac{1}{3}$ of Lemma 11.

6.4 The additive-encoding limit and the classification

As $t \to \infty$, Peikert-t converges to *additive encoding* à la LPR [7]: the ciphertext carries $v = w + \lfloor q/2 \rfloor \mu$ (+ noise) and decryption computes $\hat{\mu} = \lfloor 2(v - w')/q \rceil \bmod 2$, i.e. the key map is the indicator of an arc of length $q/2$ whose *position is chosen by the adversary* through v. The reflection attack then achieves collision *exactly* 1: setting $\mathbf{s}_1 = -\mathbf{s}_0$ and centring the arc on a fixed point of $x \mapsto -x$,

$$\max_{r \ne 1,\ \text{arc position}} \Pr[\mathbf{1}_J(X) = \mathbf{1}_J(rX)] \;=\; 1 \quad \text{at } r = -1,$$

verified exhaustively at $q \in \{97, 257, 1009\}$. Statistical binding collapses entirely; what remains is the computational, hash-based binding of the FO layer — and Schmieg's counterexamples [44] show how far that carries for ML-KEM.

The same symmetry argument disposes of the remaining mechanisms in the literature (Table 3). In each case the adversary reaches an $x \mapsto -x$ symmetry of the composed key map, either exactly (a lattice symmetry of a Voronoi decoder) or by adversarial positioning (an LPR-style arc), and the reflection $r = -1$ drives the per-slot collision to $1 - O(1/q)$.

Remark 33 (The structural criterion). The pattern is uniform: *a reconciliation or encoding mechanism supports statistical malicious-key public-key binding if and only if its composed key map admits no isometry of $\mathbb{Z}_q$ (reflection or translation) that the adversary can reach through its choice of keys and hint.* One-bit Peikert reconciliation passes because randomised doubling forces the key map to be a single arc of *odd-adjusted* length $(q \pm 1)/2$ at one of two fixed, asymmetric positions; every scheme that lets the sender's value shift the decision regions, or whose decoder inherits a lattice symmetry, fails. The correctness threshold and the binding constant are therefore in genuine tension — $\theta_t \uparrow q/4$ forces $\beta_t \uparrow 1$ — and $t = 1$ is the unique point where a 2^{-128} statistical binding target is met at practical dimensions.

Table 3. Classification of reconciliation and encoding mechanisms by the symmetry criterion. Only one-bit Peikert reconciliation supports a statistical MAL-BIND-K-PK guarantee at the deployed dimensions.

Mechanism	**Symmetry available to Per-slot collision adversary**	
Peikert-1 [29]	none (single asymmetric arc, 2 fixed positions)	$\leq \frac{2}{3} + \frac{1}{3q}$ (Lemma 11)
Peikert-t, $t \geq 2$ (Def. 10)	near-centring via hint granularity $q/2^{t+1}$	$\geq 1 - 2^{-t}$ (Thm. 12)
Ding signal [3]; OKCN/AKCN	centred/positionable arc from LPR-type decoding	$1 - O(1/q)$
LPR / NewHope-Simple / ML-KEM encoding [7,2,22]	arc freely positioned by v	$= 1$ (verified)
Voronoi $\widetilde{D}_4$, E_8 [2]	decoder invariant under $x \mapsto -x$ (lattice symmetry)	$= 1$
Code-offset / ECC layer (LAC, Round5 style)	underlying arc geometry unchanged by the code	$1 - O(1/q)$

Remark 34 (Design rationale: Peikert-1 vs. reconciliation alternatives). The one-bit Peikert reconciliation of Definition 2 was chosen for NAIBI not only for its bandwidth/correctness trade-off (P4) but because it is, among the reconciliation mechanisms in the lattice-KEM literature, the only one whose per-slot collision constant β_q (Lemma 11) is bounded strictly below $\frac{3}{4}$ under adversarial slot values (P3). Every alternative surveyed above admits the $x \mapsto -x$ auto-symmetry either exactly or by adversarial choice and is broken by a reflection attack with per-slot collision $1 - O(1/q)$. The choice of one-bit Peikert reconciliation is therefore *structurally required* by the MAL-BIND-K-PK property of Theorem 11, not merely convenient: among standard lattice-KEM reconciliation mechanisms, none is a drop-in replacement that preserves the guarantee at the deployed dimensions, and none beats the $\frac{2}{3}$ ceiling or saves enough bandwidth to make the loss of statistical binding worth accepting.

7 Parameter Sets

The parameter sets in this section instantiate the *full-matrix* profile ($\ell = k$); the compact one-row profile ($\ell = 1$) reuses the same (n, k, q, η) values. NAIBI instances use $\alpha = 3$ and a ternary structured secret/error ($\eta = 1$).

The NIST security categories. Parameter sets are indexed by the NIST post-quantum security categories, each pegged to the cost of exhaustive key search on a block cipher:

NIST category	AES analogue	classical / quantum floor (bits)
Category 1	AES-128	128 / 64
Category 3	AES-192	192 / 96
Category 5	AES-256	256 / 128

Floors are stated in *core-SVP bits*: in that cost model the classical strength is $0.292\,\beta^*$ and the quantum strength $0.265\,\beta^*$ (Section 8.2), and the quantum figure above is the Grover-halved classical exponent. Because the quantum floor is half the classical one at every category (a ratio of 0.5, well under the cost-model ratio $0.265/0.292 = 0.91$), the *classical* floor is the tighter constraint: an instance meeting it meets the quantum floor automatically. We report both figures for every set in Table 8. (Where a *gate-count* comparison is needed — the dual-hybrid column of Table 8 is a gate count, not a core-SVP figure — the corresponding AES thresholds are $2^{143}/2^{207}/2^{272}$; Section 8.2 uses those and only those for that column.)

Category→instance mapping. We serve the three categories with three instances of non-decreasing module-lattice dimension kn:

- **Category 1** $\triangleq$ **NAIBI-I**: $k = 2$, $n = 512$, $kn = 1024$ ($\approx$ 192 classical / 174 quantum bits, core-SVP primal-uSVP $\beta^* = 656$; clears Category 1 with wide margin — its classical figure alone lands within half a bit of the Category 3 floor, 191.6 against 192).
- **Category 3** $\triangleq$ **NAIBI-III**: $k = 2$, $n = 512$, $kn = 1024$ ($\approx$ 199 classical / 181 quantum bits, $\beta^* = 683$).
- **Category 5** $\triangleq$ **NAIBI-V**: $k = 2$, $n = 1024$, $kn = 2048$ ($\approx$ 406 classical / 368 quantum bits).
- **Category 5, degree-3 option** $\triangleq$ **NAIBI-V′**: $k = 3$, $n = 512$, $kn = 1536$ ($\approx$ 312 classical / 283 quantum bits, $\beta^* = 1069$), the maximal-non-split-barrier alternative at the same category (Section 7.2).

The three default instances use the same degree-2 regular representation ρ: $\mathcal{A}_\alpha = R_q[y]/(y^2 - 3)$, $\rho \to M_2(R_q)$, $\mathbf{e} \in R_q^{2\times 2}$, with the security level set entirely by the ring degree n and the modulus q. NAIBI-I uses $n = 512$ ($kn = 1024$), the same module-lattice dimension as NAIBI-III, so it clears the Category 1 floor with substantial margin; the Category 5 set simply increases n. The degree-3 regular representation $\mathcal{A}_\alpha = R_q[y]/(y^3 - 3)$ with 3 a non-cube, $\rho \to M_3(R_q)$, $\mathbf{e} \in R_q^{3\times 3}$, is used by the optional Category-5 set NAIBI-V′ alone; every other quantity in this section is unchanged by the switch, since k enters the correctness and security analyses only through the weights $\{w_j^2\}$ and the dimension kn.

Intermediate strengths. Because security in this family is tuned almost entirely through the ring degree n at fixed $k = 2$ (the mapping above is monotone in kn), any strength between two named sets is reachable without structural change: raising the ring degree above the Category 5 set ($n = 1024$, $\approx$ 406 classical / 368 quantum bits) lands an instance at any higher strength, inheriting the same ρ,

kernels, and correctness analysis. We name no such set only because the NIST ladder stops at Category 5; the family extends beyond it by a single parameter choice.

Within an instance the same (n, q, η, k) values serve all four schemes. The KEM ciphertext consists of $\mathsf{Compress}(\mathbf{u}, d_u)$, a 32-byte OTP mask c, and the k^2n-bit Peikert hint h (one bit per coefficient of the $k \times k$ matrix $\mathbf{r}$). The public-key compression parameter d_t governs $\bar{\mathbf{t}}$. Key extraction in all schemes uses $\mathsf{PExt}/\mathsf{PRec}$, so no separate rounding parameter is needed.

From worst-case to a statistical noise budget. The earlier worst-case bound (algebraic coefficient $C_{\mathrm{alg}}(k) = 2\alpha k$, compression weight $W(k)$ — the sum of α-powers in the worst column of $\rho(\mathbf{s})$, $W(2) = 4$, $W(3) = 7$) aligns all n ternary products in the same direction and is therefore extremely pessimistic: the true per-coefficient noise is a sum of many independent centred products and concentrates as $O(\sqrt{n})$, not $O(n)$. Following the ML-KEM/Saber methodology we calibrate parameters to a *negligible* failure probability $\delta \le 2^{-128}$ rather than to the worst case, which permits substantially smaller moduli (and hence smaller keys and ciphertexts) at the same security level.

Statistical model. Fix one of the k^2n output coefficients of the recovery difference $\mathbf{r}_{\mathrm{enc}} - \mathbf{r}_{\mathrm{dec}}$. It is the sum of: (i) the algebraic part $\mathbf{e}_B\rho(\mathbf{s}_A) - \mathbf{e}_A\rho(\mathbf{s}_B)$ (resp. $\mathbf{e}\rho(\mathbf{s}') - \mathbf{e}'\rho(\mathbf{s})$ for the KEM), an inner product of k ring-products, each a centred sum of n coefficient products; and (ii) the compression terms $\boldsymbol{\varepsilon}_t\rho(\mathbf{s})$, $\boldsymbol{\varepsilon}_u\rho(\mathbf{s})$, with $\boldsymbol{\varepsilon}$ entries modelled as uniform on $[-E_d, E_d]$ (variance $E_d^2/3$). Writing $v = \eta/2$ for the variance of a χ_η coefficient and $w_j \in \{1, \alpha\}$ for the α-weights of a column of $\rho(\mathbf{s})$, the variance of one coefficient is

$$\sigma^2_{\mathrm{KEM}} = \underbrace{2\sum_{j=1}^{k} n\, v^2 w_j^2}_{\text{algebraic}} + \underbrace{\sum_{j=1}^{k} n\, \tfrac{E_{d_t}^2}{3}\, w_j^2 v}_{\mathbf{t}\text{-comp}} + \underbrace{\sum_{j=1}^{k} n\, \tfrac{E_{d_u}^2}{3}\, w_j^2 v}_{\mathbf{u}\text{-comp}} + \underbrace{v}_{\mathbf{e}''}, \tag{9}$$

where the last term is the variance of the single fresh reconciliation error $\mathbf{e}''$ added on the hint-emitting side (Remark 7); numerically $v = 1/2$ is invisible against $\sigma^2 \approx 10^4$. The worst-column squared weights are $\{w_j^2\} = \{1, \alpha^2\}$ for $k = 2$ and $\{1, \alpha^2, \alpha^2\}$ for $k = 3$ ($v = \eta/2$). By a Gaussian tail bound and a union bound over all k^2n coefficients, Peikert recovery fails with probability

$$\delta \;\le\; k^2 n \cdot \mathrm{erfc}\Big(\tfrac{q/8-1}{\sigma\sqrt{2}}\Big), \tag{10}$$

since a single coefficient is recovered correctly whenever its magnitude stays $\le q/8 - 1$ (Theorem 2). All parameter sets in Table 4 satisfy $\delta \le 2^{-128}$.

7.1 Core Parameter Values

Modulus verification.

Table 4. NAIBI parameter sets (*bandwidth-optimised*). NIST Categories 1, 3 and 5 are served by NAIBI-I, NAIBI-III and NAIBI-V respectively; NAIBI-V′ is the optional degree-3 alternative at Category 5 (Section 7.2). All sets are calibrated to $\delta \le 2^{-128}$. **u**-compression d_u and public-key compression d_t; key extraction in all schemes uses PExt/PRec. kn is the module-lattice dimension; the hint length is $k^2 n$ bits.

NIST category	Instance	k	n	q	η	$\lceil \log_2 q \rceil$	d_u	d_t	kn
Category 1	NAIBI-I	2	512	19 457	1	15	11	11	1024
Category 3	NAIBI-III	2	512	13 313	1	14	11	12	1024
Category 5	NAIBI-V	2	1024	40 961	1	16	12	12	2048
Cat. 5 (opt.)	NAIBI-V′	3	512	18 433	1	15	12	12	1536

- $q = 19\,457 = 1 + 19 \cdot 1024 \equiv 1 \pmod{1024}$: prime, $q \equiv 1 \pmod{2n}$ for $n = 512$; $\alpha = 3$ is non-square ($k = 2$).
- $q = 13\,313 = 1 + 13 \cdot 1024 \equiv 1 \pmod{1024}$: prime, $q \equiv 1 \pmod{2n}$ for $n = 512$; $\alpha = 3$ is non-square ($k = 2$).
- $q = 40\,961 = 1 + 20 \cdot 2048 \equiv 1 \pmod{2048}$: prime, $q \equiv 1 \pmod{2n}$ for $n = 1024$ (NAIBI-V); $\alpha = 3$ is non-square ($k = 2$).
- $q = 18\,433 = 1 + 18 \cdot 1024 \equiv 1 \pmod{1024}$: prime, $q \equiv 1 \pmod{2n}$ for $n = 512$ (NAIBI-V′); $q \equiv 1 \pmod 3$ and $\alpha = 3$ is a *non-cube* modulo q, so $y^3 - 3$ is irreducible ($k = 3$).

7.2 The Degree-3 Option at Category 5 (NAIBI-V′)

NAIBI-V′ is the one instance in this paper built on the degree-3 regular representation: $\mathcal{A}_\alpha = R_q[y]/(y^3 - 3)$ with 3 a non-cube modulo q, $\rho \to M_3(R_q)$, $\mathbf{e} \in R_q^{3\times 3}$. It targets the same NIST Category 5 as NAIBI-V and is offered alongside it, not in place of it: the two sets trade bandwidth against the strength of the non-split barrier.

Choice of (n, k, q). At $k = 3$ the non-split condition is $q \equiv 1 \pmod 3$ together with $3^{(q-1)/3} \not\equiv 1 \pmod q$, which rules out the classical NTT primes (3 *is* a cube modulo 12 289 and 15 361). Under the NTT, non-split and 16-bit-datapath constraints the admissible moduli at $n = 512$ are $\{18\,433, 39\,937, 58\,369, 61\,441, 64\,513\}$; we take the smallest, $q = 18\,433 = 1 + 18 \cdot 1024$, since a smaller q/σ ratio makes the uSVP embedding harder at fixed kn (the same argument as for the $k = 2$ sets). This gives module-lattice dimension $kn = 1536$ and $\lceil \log_2 q \rceil = 15$, so the uniform 16-bit datapath is preserved.

Noise budget. The worst-column squared weights are $\{w_j^2\} = \{1, \alpha^2, \alpha^2\}$ for $k = 3$, i.e. a weight sum of 19 against 10 at $k = 2$. With $\eta = 1$ and $d_t = d_u = 12$ we get $E_d = \lfloor q/2^{13} \rfloor = 2$ and, from (9), $\sigma_{\mathrm{KEM}} = 133.5$ against the threshold $q/8 = 2\,304$. The envelope (10) over the $k^2 n = 4\,608$ coefficients gives $\delta \le 2^{-207}$, comfortably below the 2^{-128} target (Table 5). The heavier weight sum is more

than offset by the smaller compression error: the recovery-failure ratio $\frac{q/8}{\sigma} = 17.3$ is the largest of all four sets.

Security. The calibrated core-SVP estimate is $\beta^* = 1069$, i.e. ≈ 312 classical / 283 quantum bits at $kn = 1536$, clearing the Category 5 floor of 256 classical core-SVP bits. We state plainly that this margin is thinner than NAIBI-V's (406 classical bits at $kn = 2048$): NAIBI-V′ buys its per-slot barrier at a smaller lattice dimension, and a deployment that values raw core-SVP margin above the barrier should keep NAIBI-V. The non-split barrier itself is the strongest available here — each slot is the full $\mathbb{F}_{q^3}$, with no intermediate subfield modulo q for a trace- or norm-based attack to descend through.

Bandwidth: an honest accounting. In the full-matrix profile of this paper the public matrix carries $k^2 = 9$ ring elements rather than 4, so NAIBI-V′ is *more* expensive than NAIBI-V, not less: $|\mathsf{pk}| + |\mathsf{ct}| = 14\,464$ B against $12\,864$ B, a 12.4% premium (Table 7). The bandwidth advantage sometimes attributed to the degree-3 choice belongs to the *one-row* profile, in which $|\mathsf{pk}|$ and $|\mathsf{ct}|$ scale as kn rather than k^2n: there NAIBI-V′ is the optimum at the Category-5 floor, $4\,704$ B against $6\,240$ B for NAIBI-V (-24.6%), precisely because $kn = 1536 < 2048$. The two statements are consistent and must not be conflated: $k = 3$ helps when the transmitted object is a row and hurts when it is a matrix.

Arithmetic cost. A degree-3 matrix needs 9 length-512 NTTs where NAIBI-V needs 4 length-1024; the butterfly counts are $9 \cdot \frac{512}{2} \log_2 512$ versus $4 \cdot \frac{1024}{2} \log_2 1024$, within 1.3% of each other. The switch therefore costs essentially nothing in transform work, and all protocol logic, the reconciliation kernel, the FO transform and every proof in this paper apply verbatim: k enters only through $\{w_j^2\}$, the dimension kn, and the hint length k^2n.

7.3 Correctness: Noise Budget

Table 5. KEM failure probability per instance under the statistical model (9)–(10). σ is the per-coefficient standard deviation; recovery needs $|\cdot| \le q/8 - 1$. All instances reach $\delta \le 2^{-128}$.

Instance	k	E_{d_u}	E_{d_t}	σ_{KEM}	$q/8$	δ_{KEM}
NAIBI-I	2	4	4	172.8	2 432	2^{-136}
NAIBI-III	2	3	1	105.3	1 664	2^{-173}
NAIBI-V	2	5	5	300.8	5 120	2^{-201}
NAIBI-V′	3	2	2	133.5	2 304	2^{-207}

$E_d = \lfloor q/2^{d+1} \rceil$ from Table 4; the recovery-failure ratio $\frac{q/8}{\sigma}$ ranges from 14.1 (NAIBI-I, the tightest, giving $\delta \approx 2^{-136}$) to 17.3 (NAIBI-V′, the widest; 17.0

for NAIBI-V). The worst-case bound of Lemma 3 is no longer required to hold below $q/8 - 1$; it is superseded by the much tighter $O(\sqrt{n})$ statistical bound.

Why the compression is asymmetric ($d_t \neq d_u$ for NAIBI-III). The public key $\mathbf{t}$ and the ciphertext $\mathbf{u}$ enter the noise budget (9) *symmetrically*: their compression errors contribute the identical per-coefficient variance $n\,(E_{d_t}^2/3)\,w_j^2 v$ and $n\,(E_{d_u}^2/3)\,w_j^2 v$. The asymmetry is therefore *not* forced by correctness; it is the outcome of minimising bandwidth subject to $\delta \leq 2^{-128}$. Both compressions cost the same per bit ($k^2 n/8$ bytes per unit of d), so we minimise the integer sum $d_t + d_u$ over the feasibility region $\delta(d_t, d_u) \leq 2^{-128}$. Because $E_d = \lfloor q/2^{d+1} \rfloor$ jumps in integer steps, that region has a staircase boundary and its cheapest lattice point need not be symmetric. Concretely, for NAIBI-III ($q = 13\,313$): the symmetric choice $(d_u, d_t) = (11, 11)$ gives $E_{d_t} = E_{d_u} = 3$, hence $\sigma \approx 134$ and $\delta \approx 2^{-104}$ — *above* the 2^{-128} target, so it is infeasible. The next-cheapest options share $d_t + d_u = 23$, namely $(11, 12)$ and $(12, 11)$; both give the identical $\sigma = 105.3$ and $\delta = 2^{-173}$ (the budget is symmetric), at the identical bandwidth. We break this tie toward $(d_u, d_t) = (11, 12)$ — the *finer* public key and *coarser* ciphertext — because in the ES-DH and KEM settings $\mathbf{t}$ is published once and amortised, whereas the per-session ephemeral $\mathbf{u}$ is transmitted every session; shifting one bit of precision from the recurring $\mathbf{u}$ to the one-time $\mathbf{t}$ lowers the amortised cost at no change in δ. For NAIBI-I and NAIBI-V the discrete optimum is already symmetric ($(11, 11)$ and $(12, 12)$ respectively), because at those moduli reducing either parameter by one step crosses the 2^{-128} boundary; there is no slack to exploit, hence no asymmetry.

Ephemeral-agreement reconciliation. In EE-DH both reconciled samples are ephemeral (two $\mathbf{u}$-compression terms); in ES-DH one is the long-term $\mathbf{t}$ and one the ephemeral $\mathbf{u}$, identical to the KEM. The per-coefficient σ_{DH} and failure probability below are representative upper bounds for these modes (computed for two compression terms); the hint h_B is one bit per coefficient of $\mathbf{r} \in R_q^{k \times k}$ ($k^2 n$ bits), and the raw key has the same length. For the offered ephemeral modes the two reconciled samples carry either two $\mathbf{u}$-compression terms (EE-DH) or one $\mathbf{t}$- and one $\mathbf{u}$-compression term (ES-DH, identical to the KEM); we report the ES-DH/KEM figures below, since ES-DH is the larger of the two offered modes and hence the conservative envelope for both. (The excluded static-static mode would carry two $\mathbf{t}$-compression terms, of the same order.)

All four instances achieve $\delta \leq 2^{-128}$ for both KEM and agreement reconciliation, so decapsulation/agreement failure is cryptographically negligible. The KDF compresses the $k^2 n$ raw bits down to a 256-bit session key, ensuring ample min-entropy throughout.

7.4 Key, Ciphertext and Message Sizes

Encoding. $|\mathtt{seed}_A| = 32\,\mathrm{B}$; $|\bar{\mathbf{t}}| = k^2\, n\, d_t/8\,\mathrm{B}$ for $\mathbf{t} \in R_q^{k \times k}$ compressed at d_t bits/coef; $|\mathbf{s}| = kn/4\,\mathrm{B}$ (ternary, 2 bits/coef, $\mathbf{s} \in \mathcal{A}_\alpha$ of dimension kn);

Table 6. Per-session reconciliation for the ephemeral modes: per-coefficient σ, failure probability, hint length and raw key bits (conservative two-compression-term envelope).

Instance	k	n	q	σ_{DH}	$q/8$	δ_{DH}	$\lvert h_B\rvert$ (B)	Raw bits
NAIBI-I	2	512	19 457	172.8	2 432	2^{-136}	256	2 048
NAIBI-III	2	512	13 313	105.3	1 664	2^{-173}	256	2 048
NAIBI-V	2	1024	40 961	300.8	5 120	2^{-201}	512	4 096
NAIBI-V′	3	512	18 433	133.5	2 304	2^{-207}	576	4 608

$|\bar{\mathbf{u}}| = k^2 n\, d_u/8\,\mathrm{B}$ for $\mathbf{u} \in R_q^{k\times k}$; $|c| = 32\,\mathrm{B}$ (the OTP-encrypted FO message); $|\mathsf{H}(\mathsf{pk})| = |z| = 32\,\mathrm{B}$. The hint is $k^2 n$ bits.

Table 7. NAIBI-KEM and NAIBI-DH sizes (bytes), with public-key compression $\bar{\mathbf{t}} = \mathsf{Compress}(\mathbf{t}, d_t)$. $|\mathsf{pk}| = 32 + k^2 n\, d_t/8$ (seed + compressed matrix). $|\mathsf{sk}| = kn/4 + |\mathsf{pk}| + 64$. KEM ciphertext $|\mathsf{ct}_{\mathrm{KEM}}| = k^2 n\, d_u/8 + 32 + k^2 n/8$ (compressed matrix $\bar{\mathbf{u}}$ + 32-byte XOR mask + $k^2 n$-bit reconciliation hint h). The offered agreement modes include a $k^2 n$-bit reconciliation hint in their per-session message: ES-DH hint h_A emitted by Alice (alongside $\bar{\mathbf{u}}_A$, one pass); EE-DH hint h_B emitted by Bob (alongside Bob's ephemeral $\mathbf{u}_B$).

Instance	k	$\lvert\mathsf{pk}\rvert$	$\lvert\mathsf{sk}\rvert$	$\lvert\mathsf{ct}_{\mathrm{KEM}}\rvert$	$\lvert\mathbf{u}_A\rvert_{\mathrm{ES/EE}}$	$\lvert\mathsf{ss}\rvert$	$\lvert h\rvert$ (all variants)
NAIBI-I	2	2 848	3 168	**3 104**	2 816	32	**256 B**
NAIBI-III	2	3 104	3 424	**3 104**	2 816	32	**256 B**
NAIBI-V	2	6 176	6 752	**6 688**	6 144	32	**512 B**
NAIBI-V′	3	6 944	7 392	**7 520**	6 912	32	**576 B**

Reading the table. For *NAIBI-ES-DH*, Alice transmits $(\bar{\mathbf{u}}_A, h_A)$ (compressed ephemeral plus the $k^2 n$-bit hint); Bob recovers the shared secret locally, sending nothing back. For *NAIBI-EE-DH*, Bob transmits his ephemeral $\mathbf{u}_B$ plus the $k^2 n$-bit hint h_B alongside Alice's $\mathbf{u}_A$. For *NAIBI-KEM*, the encapsulator transmits ct (3 104 B at Level I, including the 256-byte hint h). The bandwidth-optimised moduli (negligible $\delta \le 2^{-128}$) cut $|\mathsf{pk}| + |\mathsf{ct}|$ relative to the worst-case $\delta = 0$ sets, the saving coming directly from the smaller q.

The one-row profile (NAIBI-compact). Every set above publishes the full $k \times k$ matrix $\bar{\mathbf{t}}$. A *one-row* deployment instead publishes only row 0, $\hat{\tau} = \mathbf{a}^\top \rho(\mathbf{s}) + \hat{\mathbf{e}}$ with $\mathbf{a}^\top = A_{0,:}$, and reconciles on the fixed 256-coefficient subset S, which divides both transmitted objects by k:

$$|\mathsf{pk}| = 32 + kn\, d_t/8, \qquad |\mathsf{ct}| = kn\, d_u/8 + 32 + 32,$$

the two 32-byte terms being the OTP mask and the 256-bit hint. The same $(n, k, q, \eta, d_t, d_u)$ of Table 4 then give $|\mathsf{pk}| + |\mathsf{ct}| = 2\,912/3\,040/6\,240/4\,704\,\mathrm{B}$ for

NAIBI-compact-I/III/V/V′. We list these figures in Table 9 for size comparison only: publishing one row releases a *single* sample, so the profile rests on a single-sample assumption over $\mathcal{A}_\alpha$ rather than on the k-sample MLWE^ρ analysed here, and it is *not* covered by the reductions of Section 8. This is also where the degree-3 set inverts the ordering: at $kn = 1536 < 2048$, NAIBI-compact-V′ is the smaller of the two Category-5 sets.

8 Security Analysis

8.1 Structured-Secret MLWE: Statement and Limits of Reduction

A central question is whether using a secret of the structured form $\rho(\mathbf{s}) \in \rho(\mathcal{A}_\alpha)$ (with $\mathbf{s} \in \mathcal{A}_\alpha$ ternary) weakens the underlying MLWE assumption. The structured-secret form is *not* equivalent to independent two-sample MLWE, but it admits a column decomposition that exposes its hardness in terms of a clearly identified, mildly non-standard variant. We treat MLWE^ρ as an explicit standing assumption and discuss which standard MLWE results transfer to it and which require additional care. The error is taken *generic*, $\mathbf{e} \in R_q^{2\times 2}$, which makes the column decomposition clean.

Definition 11 (Structured-secret MLWE). *The* structured-secret MLWE problem $\mathsf{MLWE}^\rho_{n,q,\chi}$ *is: given* $A \leftarrow M_2(R_q)$ *uniform, a structured secret* $\mathbf{s} \leftarrow \chi(\mathcal{A}_\alpha)$ *and a* generic *error matrix* $\mathbf{e} \leftarrow \chi^{2\times 2}$ *(each of the four entries an independent* χ*-sample in* R_q*), distinguish* $\big(A,\ A\,\rho(\mathbf{s}) + \mathbf{e}\big) \in M_2(R_q) \times M_2(R_q)$ *from* $(A, \mathbf{U})$ *with* $\mathbf{U} \leftarrow M_2(R_q)$ *uniform.*

Definition 12 (Two-block structured MLWE). $\mathsf{MLWE}^{\rho,(2;\mathcal{D}_d)}_{n,q,\chi}$*: given* $A \leftarrow \mathrm{U}(M_k(R_q))$, $\mathbf{B} \leftarrow \mathcal{D}_d$, $\mathbf{s}' \leftarrow \chi(\mathcal{A}_\alpha)$ *and independent* $\mathbf{e}', \mathbf{e}'' \leftarrow \chi^{k\times k}$*, distinguish* $\big(A, \mathbf{B},\ A\rho(\mathbf{s}') + \mathbf{e}',\ \mathbf{B}\rho(\mathbf{s}') + \mathbf{e}''\big)$ *from* $(A, \mathbf{B}, \mathbf{U}_1, \mathbf{U}_2)$ *uniform. For* $d = \lceil \log_2 q \rceil$, $\mathbf{B}$ *is uniform and the problem is exactly* MLWE^ρ *with the stacked public matrix* $\left[\begin{smallmatrix} A \\ \mathbf{B} \end{smallmatrix}\right] \in R_q^{2k\times k}$*, i.e. standard structured-secret MLWE with* $2k$ *samples per secret — the multi-sample form already absorbed by standard MLWE analyses. For* $d < \lceil \log_2 q \rceil$ *it is the two-block form of the coarse-matrix assumption of Definition 1.*

The public matrix A is generic in the full $4n$-dimensional algebra $M_2(R_q)$, the error $\mathbf{e}$ is generic in $M_2(R_q)$, and the only structure sits in the secret: $\rho(\mathbf{s}) \in \rho(\mathcal{A}_\alpha)$ has just $2n$ free R_q-coefficients.

Lemma 13 (Column decomposition). *For* $\mathbf{s} = s_0 + s_1 y \in \mathcal{A}_\alpha$*, write the two columns of* $\rho(\mathbf{s})$ *as* $\mathbf{s}^{(0)} = (s_0, s_1)^\top$ *and* $\mathbf{s}^{(1)} = (\alpha s_1, s_0)^\top = \rho(y)\mathbf{s}^{(0)}$*, and let* $\mathbf{e} = [\,\mathbf{e}^{(0)} \,|\, \mathbf{e}^{(1)}\,]$ *be the column split of the generic error, with* $\mathbf{e}^{(0)}, \mathbf{e}^{(1)} \in R_q^2$ independent *generic ternary columns. Then*

$$A\,\rho(\mathbf{s}) + \mathbf{e} \;=\; \big[\,A\,\mathbf{s}^{(0)} + \mathbf{e}^{(0)},\ \ A\,\mathbf{s}^{(1)} + \mathbf{e}^{(1)}\,\big] \;\in\; R_q^{2\times 2}.$$

Each column is a standard rank-2 MLWE sample over R_q *with the shared matrix* A *and an* independent *generic error, but the two secret columns satisfy the deterministic public relation* $\mathbf{s}^{(1)} = \rho(y)\mathbf{s}^{(0)}$.

The decomposition *identifies* structured-secret MLWE with a multi-sample MLWE problem whose secrets are *algebraically linked*; we name that problem explicitly.

Definition 13 (σ-linked 2-sample MLWE). *For a fixed public $\mathbb{F}_q$-linear bijection $\sigma : R_q^k \to R_q^k$ (in our case $\sigma = \rho(y)$), σ-linked 2-sample MLWE $\mathsf{MLWE}^{(2,\sigma)}_{n,k,q,\chi}$ asks to distinguish $(A,\ A\mathbf{s} + \mathbf{e}^{(0)},\ A\sigma(\mathbf{s}) + \mathbf{e}^{(1)})$ from $(A, \mathbf{u}_0, \mathbf{u}_1)$, with $A \leftarrow M_k(R_q)$ uniform, $\mathbf{s} \leftarrow \chi^k$ ternary, and $\mathbf{e}^{(0)}, \mathbf{e}^{(1)} \leftarrow \chi^k$ independent generic errors.*

Theorem 13 (Structured-secret MLWE $\equiv$ linked-secret 2-sample MLWE). *There is a polynomial-time, advantage-preserving equivalence between $\mathsf{MLWE}^{\rho}_{n,q,\chi}$ (Definition 11) and $\rho(y)$-linked 2-sample MLWE $\mathsf{MLWE}^{(2,\rho(y))}_{n,2,q,\chi}$:*

$$\mathsf{Adv}^{\mathsf{MLWE}^{\rho}}_{n,q,\chi}(\mathcal{A}) \;=\; \mathsf{Adv}^{\mathsf{MLWE}^{(2,\rho(y))}}_{n,2,q,\chi}(\mathcal{A}).$$

The same equivalence holds for the rounded variants MLWR^{ρ} and $\mathsf{MLWR}^{(2,\rho(y))}$ (and with T-fold loss for multi-sample variants).

Proof. The equivalence is by direct rewriting via Lemma 13. Given a structured-secret sample $(A, T = A\rho(\mathbf{s}) + \mathbf{e})$, output $(A, T_{:,0}, T_{:,1}) = (A, A\mathbf{s}^{(0)} + \mathbf{e}^{(0)}, A\mathbf{s}^{(1)} + \mathbf{e}^{(1)})$ with $\mathbf{s}^{(1)} = \rho(y)\mathbf{s}^{(0)}$; this is exactly a $\rho(y)$-linked 2-sample MLWE instance. Conversely, given a linked 2-sample instance with $\sigma = \rho(y)$, reassemble the columns into $T \in R_q^{2\times 2}$ and output (A, T). Uniformity is preserved on both sides since uniform columns reassemble to a uniform matrix. The extension to the rounded variants is term-by-term.

Remark 35 (Epistemic status of Theorem 13). Theorem 13 is an exact *syntactic identification*, not a reduction to standard independent-sample MLWE. The hardness of NAIBI-full rests on the standing assumption that the public $\rho(y)$-link between secret columns confers no advantage beyond same-matrix multi-sample reuse; Section 8.2 estimates concrete security on that basis, and Section 10 lists the reduction (or a link-exploiting attack) as an open problem. The abstract and introduction are worded accordingly: we claim an exact identification with a linked k-sample variant, not equivalence with plain MLWE.

Remark 36 (Relation to Order-LWE and algebraically structured LWE). The ring $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ is (the mod-q reduction of) an order in a degree-k extension of the $2n$-th cyclotomic field, so LWE problems over $\mathcal{A}_\alpha$ with the *ring* product fall within the Order-LWE framework of Bolboceanu–Brakerski–Perlman–Sharma [45] and the unified treatment of algebraically structured LWE of Peikert–Pepin [46], which supply worst-case ideal-lattice foundations for such rings. For the deployed case $k = 2$, $\alpha = 3$ the field $K = \mathbb{Q}(\zeta_{2n}, \sqrt{3})$ is moreover *abelian* and contained in the cyclotomic field $\mathbb{Q}(\zeta_{6n})$ (since $\sqrt{3} = \zeta_{12} + \zeta_{12}^{-1}$), placing $\mathcal{A}_\alpha$ in an abelian subfield of a cyclotomic — much closer to studied territory than an ad-hoc order, and the natural target for the worst-case theory

above (the concrete questions being whether the order $\mathbb{Z}[\zeta_{2n}, \sqrt{3}]$ is maximal away from 2 and 3, and how the canonical embedding of the basis $(1, \sqrt{3})$ distorts the error). We also flag that K has proper subfields — e.g. $\mathbb{Q}(\zeta_{2n})$ of index 2 — so norm/trace maps to subfields exist *over the number field*; the non-split condition guarantees only that the mod-q slots $\mathbb{F}_{q^k}$ admit no intermediate subfield, and a characteristic-zero subfield attack against ternary-error instances is an avenue future cryptanalysis should address. A caveat keeps the Order-LWE results from applying directly to NAIBI: the full-matrix scheme uses a *generic* public matrix $A \in M_k(R_q)$ acting on $\rho(\mathcal{A}_\alpha)$, not the ring product, so the full-matrix object is a clean MLWE sample over $M_k(R_q)$ rather than an LWE sample under the natural $\mathcal{A}_\alpha$-multiplication. Connecting either problem to the Order-LWE worst-case theory — or showing a separation — is an open problem that we add to Section 10.

Remark 37 (Linked 2-sample MLWE vs. independent 2-sample MLWE). Linked 2-sample MLWE is *not* reducible from independent 2-sample MLWE in a black-box way: a reducer producing independent $(\mathbf{s}^{(0)}, \mathbf{s}^{(1)})$ from a single MLWE secret $\mathbf{s}$ would have to know $\mathbf{s}$, which is precisely what the assumption hides. The relation $\mathbf{s}^{(1)} = \rho(y)\mathbf{s}^{(0)}$ is a fixed public bijection of $\mathcal{A}_\alpha \cong R_q^k$ that the attacker can apply at will to one sample's secret; this gives the attacker more algebraic structure than in the independent case, and could in principle confer additional leverage.

We do not exclude that linked 2-sample MLWE is strictly easier than independent 2-sample MLWE. Three observations partially mitigate this concern.

(a) *Standard MLWE is invariant under invertible $\mathbb{F}_q$-linear preprocessing of the secret.* For any fixed invertible $\mathbb{F}_q$-linear $\sigma : R_q^k \to R_q^k$, $\mathsf{MLWE}(A, A\mathbf{s} + \mathbf{e})$ has the same hardness as $\mathsf{MLWE}(A', A'\sigma(\mathbf{s}) + \mathbf{e})$ with $A' = A\sigma^{-1}$ uniform. Thus rewriting the second column as a plain MLWE sample with secret $\sigma(\mathbf{s})$ *and rerandomised matrix $A\sigma^{-1}$* is benign. The actual structural constraint imposed by the link is the *shared* A between the two samples, exactly the constraint already absorbed by standard same-A multi-sample MLWE analyses [8,9].

(b) *The link is publicly invertible and does not aid lattice attacks at known dimensions.* In all of the standard lattice attacks (primal uSVP, dual, dual-hybrid) the cost depends only on the lattice profile (kn, q, σ) and the number of available samples per secret [9,17]. The deterministic link does not introduce a smaller-dimension target lattice; an attacker can compute both columns from a candidate $\mathbf{s}$ but gains no shorter vector thereby.

(c) *No published attack exploits such an algebraic link.* In the considerable body of work on attacking structured MLWE/RLWE, algebraic links between secrets of distinct samples (rather than within a single Ring-LWE / Module-LWE secret) have not, to the best of our knowledge, been used.

Nevertheless, we treat MLWE^{ρ} as a *new structured assumption* distinct from standard MLWE. Theorem 13 identifies it cleanly with a linked 2-sample variant; our concrete security analysis (Section 8.2) conservatively estimates the cost using the lattice-estimator on the underlying (kn, q, σ) profile, with the explicit caveat that no published attack currently exploits the $\rho(y)$-link.

Remark 38 (Generic error simplifies the proof). A structured error $\rho(\mathbf{e})$ would force the two error columns to satisfy $\mathbf{e}^{(1)} = \rho(y)\mathbf{e}^{(0)}$, an additional algebraic constraint. Taking $\mathbf{e} \in R_q^{2\times 2}$ generic removes this constraint: the two error columns are independent.

Remark 39 (Why the non-split condition is essential). If α were a k-th power (split case), $\mathcal{A}_\alpha \cong R_q^k$, and each slot of $\rho(\mathbf{s})$ would be *diagonalisable* over R_q: an attacker could decouple the secret coordinates and solve k independent rank-k LWE problems with effective dimension reduced by a factor of k. Non-splitness (α a non-k-th-power: non-square for $k = 2$, non-cube for $k = 3$) ensures the secret matrix $\rho(\mathbf{s})$ has no $\mathbb{F}_q$-rational eigenspace at any NTT slot, so the k columns of $\mathbf{t}$ cannot be split apart. The barrier is therefore structural, not a parameter choice one could relax.

Remark 40 (Degree-k reduction). For $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ the column-decomposition argument of Theorem 13 gives a reduction from k-sample rank-k MLWE: the k columns of $A\rho(\mathbf{s}) + \mathbf{e} \in R_q^{k\times k}$ are k rank-k MLWE samples on the shared A, with independent generic error columns and secret columns related to $\mathbf{s}^{(0)}$ by the fixed public maps $\mathbf{s}^{(j)} = \rho(y^j)\mathbf{s}^{(0)}$. Hence $\mathsf{Adv}^{\mathsf{MLWE}^\rho} \le \mathsf{Adv}^{\mathsf{MLWE}^{(k)}}_{n,k,q,\chi}$ for both $k \in \{2, 3\}$. The default sets use $k = 2$, reducing to rank-2 MLWE in dimension kn (1024 for Level I, 2048 for Level V); the optional NAIBI-V′ reduces to rank-3 MLWE in dimension $kn = 1536$.

Remark 41 (Standing assumption for NAIBI). The hardness assumption underlying NAIBI's full-matrix scheme is therefore $\rho(y)$-linked k-sample MLWE on a uniform $A \in M_k(R_q)$. We conjecture that this is essentially as hard as independent k-sample MLWE in the same lattice dimension kn, but, as discussed above, we do *not* have a black-box reduction to that effect, and the precise quantitative relationship is left open. Concrete security estimates (Section 8.2) use the standard lattice-estimator on the (kn, q, σ) profile, on the conservative basis that no known attack exploits the $\rho(y)$-link.

8.2 Concrete Security Estimates

We use the `lattice-estimator` [16] on each MLWE instance $(n, k, q, \sigma = \sqrt{1/2})$, taking the minimum cost across primal uSVP, dual, and dual-hybrid attacks under the core-SVP cost model [14] and the refined dual estimates [17]. The structured-secret form $\rho(\mathbf{s}_P) \in \rho(\mathcal{A}_\alpha)$ is conservatively *not* exploited in attacks, justified by Theorem 13: NAIBI's hardness is no easier than k-sample rank-k MLWE on a uniform $A \in M_k(R_q)$ with a ternary secret of the natural dimension kn.

Table 8. Single authoritative concrete-security table for NAIBI-Full. All classical/quantum strength and attack-cost estimates quoted anywhere in this paper derive from this table; the bandwidth comparison (Table 9) carries no independent security figures and refers here. Concrete security for the bandwidth-optimised parameter sets, confirmed with the full `lattice-estimator` [16] (run on each MLWE instance $(kn, q, \sigma = \sqrt{1/2})$, ternary secret). Columns β^*_{primal} and the core-SVP bits (classical $\approx 2^{0.292\beta^*}$, quantum $\approx 2^{0.265\beta^*}$) report primal-uSVP; the full run agrees with the calibrated self-contained estimator on β^* to within ± 1. The cheapest attack in every case is the dual-hybrid; we list its cost in the refined gate-count model of MATZOV [17] ($\log_2$ ROP). The last column gives the NIST security category served (Categories 1, 3 and 5 by NAIBI-I, NAIBI-III and NAIBI-V; NAIBI-V′ is the optional degree-3 alternative at Category 5). The dual-hybrid figure for NAIBI-V′ is not reported: the $k = 3$ profile was not covered by the same `lattice-estimator` run, and we prefer a blank to an extrapolation. All figures are computed on the (kn, q, σ) profile under the standing assumption that the $\rho(y)$-link is not exploitable (Remark 41); they apply to both the compressed default parameters (whose CPA proof additionally invokes the coarse-matrix variant of Definition 1) and the uncompressed `-pc` option (provable from MLWE^{ρ} alone, Remark 2), which share the same lattice profile.

Instance	k	kn	q	β^*_{primal}	core-SVP (uSVP)		dual-hybrid	NIST
					Classical	Quantum	$\log_2$ ROP	category
NAIBI-I	2	1024	19 457	656	192	174	203	Category 1
NAIBI-III	2	1024	13 313	683	199	181	211	Category 3
NAIBI-V	2	2048	40 961	1389	406	368	396	Category 5
NAIBI-V′	3	1536	18 433	1069	312	283	—	Cat. 5 (opt.)

NAIBI-I is set at $k = 2$, $n = 512$ ($kn = 1024$, $q = 19457$), matching the reference implementation and Table 4; its figures above are computed with the `lattice-estimator` at this dimension. Since $kn = 1024$ equals the NAIBI-III dimension (with a slightly larger modulus, hence $\beta^* = 656$ vs. 683), NAIBI-I clears Category 1 with substantial margin — its 191.6 classical core-SVP bits fall only 0.45 bit short of the Category 3 floor of 192, so it is a Category-1 set with a Category-3-sized margin rather than a Category-3 set.

Why Level I uses $k = 2$, $n = 512$. NAIBI-I is a $k = 2$ instance at $n = 512$, giving module-lattice dimension $kn = 1024$ — the *same* dimension as the NAIBI-III set, differing only in the modulus ($q = 19457$ vs. 13313). This over-provisions the Category 1 floor (a $k = 2$, $n = 256$ instance at $kn = 512$ cannot reach 128-bit security for any modulus, so $n = 512$ is the smallest admissible power-of-two ring degree at $k = 2$), trading minimal bandwidth for a comfortable security margin. The `lattice-estimator` gives primal-uSVP $\beta^* = 656$, i.e. ≈ 192 classical / 174 quantum core-SVP bits at $kn = 1024$, $q = 19457$ (Table 8) — only marginally below NAIBI-III ($\beta^* = 683$) because of the slightly larger modulus, and well

above the Category 1 floor (indeed just short of the Category 3 classical floor of 192). The non-split barrier is $\mathbb{F}_{q^2}$ per slot, as at Categories 3 and 5.

Effect of the smaller modulus. Reducing q at fixed dimension kn and fixed noise $\sigma = \sqrt{1/2}$ *increases* the BKZ block size required (a smaller q/σ ratio makes the unique-SVP embedding harder), so the bandwidth-optimised sets of Table 4 are *no weaker* than larger-q sets at the same kn. NAIBI-III ($kn = 1024$) reaches ≈ 199 classical bits and NAIBI-V ($kn = 2048$) ≈ 406, comfortably meeting Levels III and V respectively.

Cheapest attack (dual-hybrid). Across the full `lattice-estimator` run (MATZOV gate-count model), the cheapest attack on every instance is the dual-hybrid, not primal-uSVP. Its cost is 2^{203} (NAIBI-I), 2^{211} (NAIBI-III) and 2^{396} (NAIBI-V). These are gate counts, so they are compared against the *gate-count* counterparts of the category floors, i.e. the AES key-search costs $2^{143}/2^{207}/2^{272}$ rather than the 128/192/256 core-SVP bits used elsewhere; all three clear them. NAIBI-III has the slimmest margin (≈ 4 bits over Category 3), while NAIBI-I clears the Category-1 threshold by ≈ 60 bits; a $kn = 768$ instance would clear it by only ≈ 18.

Honest assessment. NAIBI-I sits at $k = 2$, $n = 512$ ($kn = 1024$), sharing the module-lattice dimension of NAIBI-III rather than the marginal $kn = 768$ one might pick to shave bandwidth. That choice keeps it clear of the borderline regime in which a Level-I claim depends on the cost model: the `lattice-estimator` gives 191.6 classical core-SVP bits and a dual-hybrid cost of 2^{203} in the MATZOV gate-count model, clearing the Category-1 floor with wide margin in both accountings (128 core-SVP bits, 2^{143} gates). At $kn = 768$ the core-SVP dual-hybrid would dip to $\approx 2^{137}$, below the AES-128 line — the same cost-model sensitivity ML-KEM-512 exhibits. We prefer the extra dimension.

Cross-check vs. ML-KEM. The estimator reproduces Kyber-512/768/1024 core-SVP block sizes ($\approx 406/623/873$) to within one unit, which is why the NAIBI figures can be read on the same scale. NAIBI-I ($k = 2$, $kn = 1024$) shares the dimension of NAIBI-III, so both exceed Kyber-768; NAIBI-V ($kn = 2048$) exceeds Kyber-1024.

8.3 Algebraic and Structural Attacks

Non-split barrier. $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$ rules out slot-wise NTRU/meet-in-the-middle attacks that exploit a split product. Each slot lives in $\mathbb{F}_{q^k}$ ($\mathbb{F}_{q^2}$ for the $k = 2$ sets NAIBI-I/III/V; $\mathbb{F}_{q^3}$ for NAIBI-V′), not $\mathbb{F}_q$.

Secret invertibility. $\mathbf{s}$ is conditioned to be invertible in $\mathcal{A}_\alpha$. $\Pr[\mathbf{s} \text{ non-invertible}] \leq n/q^2 \leq 2^{-25}$ for all instances; resampling terminates quickly.

Multi-key and key-reuse (KEM). Independent seeds $\mathtt{seed}_A$ give independent MLWE instances. Under a fixed A shared by Q users, multi-target attacks gain $\leq \log_2 Q$ bits, absorbed by the security margins.

Multi-session and key-reuse (ephemeral modes). The offered ES-DH and EE-DH modes use a fresh ephemeral secret each session, so no long-term secret accumulates MLWR samples: the single-sample reduction applies and there is no T-tightness loss. The static-static regime that *would* leak T multi-sample MLWR observations (Theorem 3) is excluded outright (Section 8.4).

Decryption / agreement failure attacks. Since $\delta \leq 2^{-128}$ for both KEM and DH (statistical argument), no failure-boosting attack applies.

8.4 Why NAIBI Excludes the Static-Static Mode

A static-static reconciliation deployment would reuse a long-term secret across many sessions, placing it squarely within the threat model of *key-mismatch* (a.k.a. *plaintext-checking*) attacks against reconciliation-based lattice key exchange — a substantial body of work matured since Fluhrer's original observation [33]. This subsection explains why we exclude that mode rather than ship it with countermeasures.

The key-mismatch attack template. The adversary impersonates one party and submits *malformed* public values $\mathbf{t}^*$ that probe the counterparty's long-term secret. After each session it either (i) observes the responder's signal function [33,34], (ii) observes the reconciliation hint [39], or (iii) learns only whether the final session keys match (a plaintext-check oracle [35,36,37,38]). Because the reconciliation product against a reused secret is *deterministic and noise-free* (the secret error sits in the attacker's own key, not in the product it observes), a small constant number of malformed keys of the form $\mathbf{t}^* = c \cdot I$ isolates the ternary secret coefficient by coefficient via scalar multiplication.

The non-split barrier is in the wrong domain. One might hope the non-split structure $\mathcal{A}_\alpha \cong \prod_i \mathbb{F}_{q^k}$ helps. It does not: non-splitness obstructs slot-wise dimension reduction in the *NTT* (frequency) domain, against *passive lattice* attacks, whereas key-mismatch operates in the *coefficient* domain on an adversarially chosen $\mathbf{t}^*$. The two are orthogonal; the central algebraic guarantee of NAIBI gives no protection here. A static-static NAIBI would therefore fall to the same attacks as Ding KE and NewHope, and just as quickly.

Remark 42 (The fresh reconciliation error does not rescue static-static). A natural objection, given the fresh error $\mathbf{e}''$ added on the hint-emitting side (Remark 7), is whether $\mathbf{e}''$ might itself defeat the key-mismatch adversary. It does not, and it is worth being explicit about why, since $\mathbf{e}''$ now appears throughout the protocol. The role of $\mathbf{e}''$ is to make the reconciliation input $\mathbf{r} = \mathbf{t}'\rho(\mathbf{s}) + \mathbf{e}''$ *pseudorandom under* MLWE^ρ for a *passive* observer; this is what lets the security proofs avoid any rounding assumption (Theorems 6, 4, 5). The active key-mismatch adversary lives in a different world: it never observes $\mathbf{r}$. It submits a malformed public value $\mathbf{t}^* = c \cdot I$ (not a well-formed MLWE sample) and reads a single match/abort bit per session, so the victim computes $\mathbf{r}_B = c\,\rho(\mathbf{s}_B) + \mathbf{e}''_B$ and the adversary

sweeps c to locate the per-coefficient toggle boundaries of $c\,\rho(\mathbf{s}_B)$, recovering $\mathbf{s}_B$ coordinate by coordinate. Three facts make $\mathbf{e}''_B$ inoperative against this. (i) *Amplitude.* $\mathbf{e}''_B$ is ternary, $\|\mathbf{e}''_B\|_\infty \le \eta = 1$, far too small to blur a toggle boundary whose spacing is $\Theta(q)$; widening it would wreck correctness (δ blows up), the same security-vs-correctness tension that already forces the dimension choice of Section 8.2. (ii) *Freshness against a fixed target.* $\mathbf{e}''_B$ is resampled every session while $\mathbf{s}_B$ is reused; over the adversary's many queries the fresh independent noise averages out and only the persistent secret survives the aggregation — precisely the regime of Qin–Cheng–Ding [36] and Ding–Fluhrer–Saraswathy [35], which already handle reconciliation noise of this kind. (iii) *Out-of-distribution input.* Passive pseudorandomness of $\mathbf{r}$ is established only for well-formed keys; the adversary deliberately violates that hypothesis with $\mathbf{t}^* = c \cdot I$, so the MLWE^ρ guarantee says nothing about the victim's behaviour here. The static-static exclusion below therefore stands unchanged: $\mathbf{e}''$ buys passive pseudorandomness, not resistance to an active chosen-$\mathbf{t}^*$ oracle.

Why we do not patch it. The only sound repairs of the static-static mode are (M1) public-key-validation/NIZK proofs of well-formedness attached to every static key (the SWOOSH route, in the adversarial-key NIKE model of Freire–Hofheinz–Kiltz–Paterson [50]), which denies the adversary a chosen $\mathbf{t}^*$ but costs several kilobytes per key and so *negates the bandwidth advantage* that was the only reason to want a static-static mode; or (M2) key-rotation bounds, which merely cap the damage. Neither yields an attractive object. We therefore make the pragmatic choice of **not offering the static-static mode at all**. NAIBI is deployed exclusively as:

(a) The ephemeral agreement modes (ES-DH, EE-DH). A fresh ephemeral secret on at least one side means each session probes a different secret, so per-session leakage cannot accumulate: the attack template does not apply. EE-DH is immune on both sides; ES-DH protects Alice's ephemeral perfectly, and Bob's long-term key is used only inside the FO/authenticated KEM workflow of (b).

(b) The IND-CCA2 KEM. The FO transformation applied to NAIBI-PKE-OTP eliminates plaintext-check and key-mismatch oracles by design: an invalid ciphertext is rejected by the re-encryption check (implicit rejection), and a valid ciphertext does not reveal whether intermediate values match. In the language of [38], NAIBI-KEM is a properly-transformed CCA-secure object, whereas a static-static reconciliation agreement is a CPA object under reuse. This is the recommended default deployment.

Consequence. Excluding the static-static mode removes the key-mismatch surface entirely, at no security cost and with the bandwidth profile of a standard one-pass KEM. It does forgo the static-static key model (both parties deriving a shared secret from pre-published keys); that model, delivered safely, is exactly SWOOSH's contribution, against which we do not compete. An actively secure static-static NIKE from $\mathcal{A}_\alpha$ (with the requisite NIZK proofs of knowledge of the secret) is left as an avenue for future work.

8.5 Side-Channel Considerations

Ternary secrets ($\eta = 1$) make polynomial multiplication constant-time via add/subtract, and centred binomial sampling at $\eta = 1$ is inherently constant-time (2 random bits per coefficient); the rounding function and the FO re-encryption `memcmp` are constant-time as well. None of this addresses the generic chosen-ciphertext side-channel attacks that target the FO re-encryption comparison in CCA-secure lattice KEMs [40]; a full side-channel evaluation is recommended before deployment.

9 Bandwidth Comparison

We compare NAIBI (full-matrix scheme) against the principal post-quantum KEM families: lattice-based (ML-KEM [22], NTRU [23], FrodoKEM [24]), code-based (HQC [25], BIKE [26], Classic McEliece [27]). The relevant figure of merit for transmission is the bandwidth $\mathrm{BW} = |\mathsf{pk}| + |\mathsf{ct}|$, which captures the cost of one full key exchange. The NAIBI rows (label *MLWE+*, Sections 4–8) rest on the structured-secret MLWE assumption, with public-key compression handled by the coarse-matrix variant of Definition 1 (no rounding assumption).

Table 9. Bandwidth of post-quantum KEMs (bytes) at matching NIST security categories. BW = $|\mathsf{pk}| + |\mathsf{ct}|$. Entries marked † standardised by NIST (FIPS 203) or selected for standardisation; δ is the (worst-case or estimated) decapsulation failure probability. The BIKE rows (‡) are the *IND-CPA* (ephemeral-key) parameter sets, the only entries here not claiming IND-CCA2; their sizes are therefore not directly comparable with the CCA-secure rows, NAIBI's included. Rows are ordered by increasing BW within each category. The NAIBI *MLWE+* rows rest on the structured-secret MLWE standing assumption (MLWE^{ρ}, identified with linked k-sample MLWE but not reduced to independent-sample MLWE), with compression handled by the coarse-matrix variant. The *1-sample* $\mathcal{A}_\alpha$ rows are the one-row profile of Section 7: same parameters, a single published sample, and therefore a single-sample assumption *not* analysed in this paper — they are listed for size comparison only and should be read with correspondingly lower confidence. NAIBI-KEM-V′ is the optional degree-3 alternative at Category 5 (Section 7.2); in this full-matrix profile it costs 12.4% more than NAIBI-KEM-V, the ordering being reversed in the one-row profile.

Scheme	Family	$\lvert\mathsf{pk}\rvert$	$\lvert\mathsf{sk}\rvert$	$\lvert\mathsf{ct}\rvert$	BW	δ
NIST Category 1 (AES-128 equivalent)						
NTRU-HPS-2048-509	lattice (NTRU)	699	935	699	1 398	negl.
ML-KEM-512†	lattice (MLWE)	800	1 632	768	1 568	2^{-139}
NAIBI-compact-I	**1-sample $\mathcal{A}_\alpha$**	**1 440**	**1 760**	**1 472**	**2 912**	$\mathbf{2^{-136}}$
BIKE-L1‡	code (QC-MDPC)	1 541	5 223	1 573	3 114	$\approx 2^{-128}$
NAIBI-KEM-I	**MLWE+**	**2 848**	**3 168**	**3 104**	**5 952**	$\mathbf{2^{-136}}$
HQC-128†	code (QC)	2 249	2 305	4 433	6 682	$< 2^{-128}$
FrodoKEM-640	lattice (LWE)	9 616	19 888	9 720	19 336	2^{-138}
McEliece-348864	code (Goppa)	261 120	6 492	96	261 216	**0**
NIST Category 3 (AES-192 equivalent)						
NTRU-HPS-2048-677	lattice (NTRU)	930	1 235	930	1 860	negl.
ML-KEM-768†	lattice (MLWE)	1 184	2 400	1 088	2 272	2^{-164}
NAIBI-compact-III	**1-sample $\mathcal{A}_\alpha$**	**1 568**	**1 888**	**1 472**	**3 040**	$\mathbf{2^{-173}}$
BIKE-L3‡	code (QC-MDPC)	3 083	10 105	3 115	6 198	$\approx 2^{-192}$
NAIBI-KEM-III	**MLWE+**	**3 104**	**3 424**	**3 104**	**6 208**	$\mathbf{2^{-173}}$
HQC-192†	code (QC)	4 522	4 586	8 978	13 500	$< 2^{-192}$
FrodoKEM-976	lattice (LWE)	15 632	31 296	15 744	31 376	2^{-148}
McEliece-460896	code (Goppa)	524 160	13 608	156	524 316	**0**
NIST Category 5 (AES-256 equivalent)						
NTRU-HPS-4096-821	lattice (NTRU)	1 230	1 592	1 230	2 460	negl.
ML-KEM-1024†	lattice (MLWE)	1 568	3 168	1 568	3 136	2^{-174}
NAIBI-compact-V′	**1-sample $\mathcal{A}_\alpha$**	**2 336**	**2 784**	**2 368**	**4 704**	$\mathbf{2^{-207}}$
NAIBI-compact-V	**1-sample $\mathcal{A}_\alpha$**	**3 104**	**3 680**	**3 136**	**6 240**	$\mathbf{2^{-201}}$
BIKE-L5‡	code (QC-MDPC)	5 122	16 494	5 154	10 276	$\approx 2^{-256}$
NAIBI-KEM-V	**MLWE+**	**6 176**	**6 752**	**6 688**	**12 864**	$\mathbf{2^{-201}}$
NAIBI-KEM-V′	**MLWE+**	**6 944**	**7 392**	**7 520**	**14 464**	$\mathbf{2^{-207}}$
HQC-256†	code (QC)	7 245	7 317	14 421	21 666	$< 2^{-256}$
FrodoKEM-1344	lattice (LWE)	21 520	43 088	21 632	43 152	2^{-199}
McEliece-8192128	code (Goppa)	1 357 824	14 120	208	1 358 032	**0**

Reading the BW column top-down within each category, several points stand out.

- **The full-matrix NAIBI-KEM buys a clean assumption with bandwidth.** At 6.0–14.5 kB it lands in the code-based band (between the BIKE and HQC sizes, and smaller than HQC at the Category 3/5 points) and $\approx$ 3–5$\times$ below FrodoKEM, retaining negligible $\delta \leq 2^{-128}$ and binding correctness, while reducing to *plain* rank-k MLWE. It adds a guarantee none of the comparison schemes provide: binding correctness, together with the unified forward-secure ephemeral agreement that shares its core.
- **The degree-3 option is a barrier trade, not a size win here.** NAIBI-KEM-V$'$ sits 1 600 B above NAIBI-KEM-V because the transmitted matrix carries $k^2 = 9$ ring elements instead of 4; what it buys is the maximal per-slot non-split barrier $\mathbb{F}_{q^3}$ at the same category. The comparison inverts once the transmitted object is a row rather than a matrix, where $|\mathsf{pk}|, |\mathsf{ct}|$ scale as $kn = 1536 < 2048$ (Section 7.2).
- **The one-row profile trades assumption confidence for size.** NAIBI-compact sits at 2.9–6.2 kB, immediately behind ML-KEM and NTRU and *below every code-based KEM* in all three categories (2 912 B against 3 114 B for BIKE-L1, 4 704 B against 10 276 B for BIKE-L5), at 6.6–10.3$\times$ under FrodoKEM. At equal parameters it is smaller than the full-matrix set by a factor of essentially k — 2.0$\times$ at $k = 2$, 3.1$\times$ at $k = 3$, since $k^2 n$ becomes kn — but on a single-sample assumption this paper does not analyse; the two families are alternatives, not a ladder.
- **ML-KEM and NTRU remain the smallest** (1.4–3.1 kB) and likewise enjoy negligible-or-zero δ, but forfeit binding correctness; the NAIBI premium is the price of that property and of the shared forward-secure agreement core.
- **FrodoKEM** pays 19–43 kB for unstructured-LWE conservatism, and **Classic McEliece** achieves $\delta = 0$ with 96–208 B ciphertexts but multi-megabyte public keys (> 250 kB), impractical except for servers with persistent state.

In short, NAIBI occupies a niche no existing KEM fills: binding correctness together with a unified forward-secure ephemeral agreement that shares the KEM core. The full-matrix scheme pays for it by sitting in the code-based bandwidth band, under a clean reduction to plain MLWE and with $\delta \leq 2^{-128}$ at every level.

10 Conclusion

We have presented **NAIBI-Full**, an algebraic recasting of reconciliation-based lattice key agreement and derived KEMs, built on the non-split commutative algebra $\mathcal{A}_\alpha = R_q[y]/(y^k - \alpha)$ and its regular representation $\rho : \mathcal{A}_\alpha \to M_k(R_q)$. The public key is the full $k \times k$ matrix $\mathbf{t} = A\rho(\mathbf{s}) + \mathbf{e}$, whose hardness lives in the (structured-secret) MLWE landscape: each column is a bona fide rank-k MLWE sample with linked secrets, and the column-decomposition of Theorem 13 identifies the assumption precisely with a $\rho(y)$-linked k-sample MLWE problem. This is the conservative footing — a clean account in terms of a standard lattice assumption — at the cost of k^2-element public keys and ciphertexts.

We are explicit about what NAIBI-Full is and what it is not. It is *not* a non-interactive (0-pass) primitive comparable to classical Diffie–Hellman, CSIDH or SWOOSH [31]; SWOOSH remains the state of the art for true lattice-based NIKE. NAIBI-Full is a one-pass KEM (in the same sense as ML-KEM) together with two forward-secure ephemeral key-agreement variants sharing its algebraic core. And **NAIBI-Full does not offer a static-static mode**: a static-static reconciliation deployment would inherit the active key-mismatch attack surface of the Ding/Peikert/NewHope family, which recovers the long-term secret after only a few sessions and against which the non-split barrier — an NTT-domain guarantee — gives no protection. The only sound repairs, NIZK proofs of well-formedness, would negate the bandwidth advantage that motivated the mode in the first place. Rather than ship a fragile mode with heavy countermeasures, we exclude it outright (Section 8.4) and confine NAIBI-Full to its key-mismatch-resistant deployments: the FO-transformed KEM (immune by the re-encryption check) and the ephemeral ES/EE agreement modes (immune because each session uses a fresh secret).

Summary of contributions.

- **An algebraic framework** for lattice key encapsulation and forward-secure ephemeral key agreement over non-split commutative algebras, recovering the Ding/Peikert template as a special case while admitting a fully generic public matrix.
- **An algebraic noise bound** $\|\mathbf{e}_B\rho(\mathbf{s}_A) - \mathbf{e}_A\rho(\mathbf{s}_B)\|_\infty \leq 2\alpha kn\eta^2$ with generic errors, and an exact characterisation of the $\mathbf{u}$-compression amplification.
- **Explicit structured-secret MLWE assumption** $\mathsf{MLWE}^{\rho}_{n,q,\chi}$ together with a column decomposition (Theorem 13) that identifies it precisely with a $\rho(y)$-linked k-sample MLWE problem; we are clear that this is not equivalent to standard independent-secret multi-sample MLWE and must be treated as a standing assumption in its own right.
- **Parameter sets** at NIST security Categories 1, 3 and 5 (served by NAIBI-I/III/V, with the optional degree-3 set NAIBI-V′ as a maximal-non-split-barrier alternative at Category 5), all with negligible $\delta \leq 2^{-128}$; all concrete-security estimates are consolidated in the single authoritative Table 8 (NAIBI-I clears Category 1 at $\approx$ 192 classical core-SVP bits, dual-hybrid 2^{203}).
- **IND-CCA2 NAIBI-KEM with KEM binding correctness** (Corollary 4, bound $(2/3 + \frac{1}{3q})^{\lceil n/2 \rceil} + (8/q)^{n/2} + 2^{-256}$, below 2^{-148} everywhere): a statistical, *decapsulation-level* collision-freeness, in the sense that no FO-validated ciphertext decapsulates to the same shared secret under two distinct honest keys. It complements the generic hash-based binding of ML-KEM [43,44] (Section 5.4), and we lift it, with no distributional assumption on the adversarial keys, to $\mathsf{MAL\text{-}BIND\text{-}K\text{-}PK}$ (Theorem 11) — the malicious-key notion ML-KEM is proven to lack [44] and the seed key format does not restore — given the pk-bound rejection key of Algorithm 11 (Remark 20). The guarantee comes from the FO re-encryption test and implicit rejection,

not from the bare PKE one-time pad (Remark 22). The KEM, together with the ephemeral ES/EE modes, is free of the key-mismatch attack surface that a static-static agreement mode would inherit.
– **A principled exclusion of the static-static mode** (Section 8.4): we show why a reused-key reconciliation deployment falls to the key-mismatch attacks of [33,34,35,36,38,37], why the non-split barrier does not help in the coefficient domain, and why the only sound repair (NIZK proofs of well-formedness) negates the bandwidth advantage — justifying the design choice not to offer the mode.

Open problems.

– A rigorous reduction of $\rho(y)$-linked k-sample MLWE to independent k-sample MLWE, or a concrete cryptanalysis exploiting the link if one exists.
– Connecting MLWE^{ρ} to the Order-LWE worst-case theory of [45,46] over the order $\mathcal{A}_\alpha$, or establishing a separation (Remark 36).
– A QROM version of Theorem 11. The K-PK axis is *resolved* in the ROM in Section 5.5 (Theorem 11): NAIBI-Full is MAL-BIND-K-PK with no assumption on the key distribution, the notion [44] shows ML-KEM lacks, via the minimum-NTT-weight bound (Lemma 10) and the correlated per-slot collision bound (Lemma 11) on the accept branch, and via the pk-bound rejection key on the reject branch (Remark 20). Case (B) is where a quantum random-oracle treatment would be needed, and where no statistical statement is available in any model (Remark 24).
– An actively secure static-static NIKE from $\mathcal{A}_\alpha$ in the multi-user, adversarial-public-key model of Freire–Hofheinz–Kiltz–Paterson [50] (requiring NIZK proofs of knowledge of static secrets) — the only sound way to reinstate the static-static mode this work deliberately excludes.
– Authenticated ephemeral NAIBI agreement (lattice analogue of HMQV), along the lines of the ideal-lattice AKE of [4] or the generic QROM compiler of [32], and lattice signatures from $\mathcal{A}_\alpha$ via Fiat–Shamir in the style of [15].
– Implementation and side-channel evaluation, including AVX2 and Cortex-M targets.

References

1. W. Diffie and M. Hellman, "New Directions in Cryptography," *IEEE Trans. Inf. Theory*, IT-22(6):644–654, 1976.
2. E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe, "Post-Quantum Key Exchange — A New Hope," *USENIX Security 2016*, pp. 327–343.
3. J. Ding, X. Xie, and X. Lin, "A Simple Provably Secure Key Exchange Scheme Based on the Learning with Errors Problem," IACR ePrint 2012/688.
4. J. Zhang, Z. Zhang, J. Ding, M. Snook, and Ö. Dagdelen, "Authenticated Key Exchange from Ideal Lattices," *EUROCRYPT 2015*, LNCS 9057, pp. 719–751.

5. W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes, "CSIDH: An Efficient Post-Quantum Commutative Group Action," *ASIACRYPT 2018*, LNCS 11274, pp. 395–427.
6. W. Castryck and T. Decru, "An Efficient Key Recovery Attack on SIDH," *EUROCRYPT 2023*, LNCS 14008, pp. 423–447.
7. V. Lyubashevsky, C. Peikert, and O. Regev, "On Ideal Lattices and Learning with Errors over Rings," *J. ACM*, 60(6):43, 2013.
8. A. Langlois and D. Stehlé, "Worst-Case to Average-Case Reductions for Module Lattices," *Des. Codes Cryptogr.*, 75(3):565–599, 2015.
9. M. Albrecht, R. Player, and S. Scott, "On the Concrete Hardness of Learning with Errors," *J. Math. Cryptol.*, 9(3):169–203, 2015.
10. R. Avanzi et al., "CRYSTALS-Kyber (ML-KEM) Algorithm Specifications," NIST PQC Standard FIPS 203, 2024.
11. D. Hofheinz, K. Hövelmanns, and E. Kiltz, "A Modular Analysis of the Fujisaki–Okamoto Transformation," *TCC 2017*, LNCS 10677, pp. 341–371.
12. J. Don, S. Fehr, C. Majenz, and C. Schaffner, "Security of the Fiat–Shamir Transformation in the Quantum Random Oracle Model," *CRYPTO 2019*.
13. N. Bindel, M. Hamburg, K. Hövelmanns, A. Hülsing, and E. Persichetti, "Tighter Proofs of CCA Security in the Quantum Random Oracle Model," *TCC 2019*.
14. A. Becker, L. Ducas, N. Gama, and T. Laarhoven, "New Directions in Nearest Neighbor Searching with Applications to Lattice Sieving," *SODA 2016*, pp. 10–24.
15. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé, "CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme," *TCHES* 2018(1):238–268.
16. M. Albrecht, B. Curtis, A. Deo, A. Davidson, R. Player, E. Postlethwaite, F. Virdia, and T. Wunderer, "Estimate all the {LWE, NTRU} schemes!" *SCN 2018*; also https://github.com/malb/lattice-estimator.
17. MATZOV, "Report on the Security of LWE: Improved Dual Lattice Attack," Tech. Report, 2022.
18. A. Banerjee, C. Peikert, and A. Rosen, "Pseudorandom Functions and Lattices," *EUROCRYPT 2012*, LNCS 7237, pp. 719–737.
19. J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs, "Learning with Rounding, Revisited," *CRYPTO 2013*, LNCS 8042, pp. 57–74.
20. M. Abdalla, M. Bellare, and P. Rogaway, "The Oracle Diffie–Hellman Assumptions and an Analysis of DHIES," *CT-RSA 2001*, LNCS 2020, pp. 143–158.
21. R. Canetti and H. Krawczyk, "Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels," *EUROCRYPT 2001*, LNCS 2045, pp. 453–474.
22. NIST, "Module-Lattice-Based Key-Encapsulation Mechanism Standard," FIPS 203, August 2024.
23. C. Chen, O. Danba, J. Hoffstein, A. Hülsing, J. Rijneveld, J. M. Schanck, P. Schwabe, W. Whyte, and Z. Zhang, "NTRU: Algorithm Specifications and Supporting Documentation," NIST PQC Round-3 submission, 2020.
24. M. R. Albrecht, E. Alkim, J. W. Bos, L. Ducas, K. Easterbrook, B. LaMacchia, P. Longa, I. Mironov, V. Nikolaenko, C. Peikert, A. Raghunathan, and D. Stebila, "FrodoKEM: Learning With Errors Key Encapsulation," ISO/IEC standard proposal, 2025; https://frodokem.org/.
25. C. Aguilar-Melchor, N. Aragon, S. Bettaieb, L. Bidoux, O. Blazy, J.-C. Deneuville, P. Gaborit, E. Persichetti, G. Zémor, J. Bos, A. Hülsing, A. Karpman, and A. Petzoldt, "Hamming Quasi-Cyclic (HQC)," NIST PQC selection (additional KEM), 2025; https://pqc-hqc.org/.

26. N. Aragon, P. Barreto, S. Bettaieb, L. Bidoux, O. Blazy, J.-C. Deneuville, P. Gaborit, S. Gueron, T. Guneysu, C. Aguilar Melchor, R. Misoczki, E. Persichetti, N. Sendrier, J.-P. Tillich, V. Vasseur, and G. Zémor, "BIKE: Bit Flipping Key Encapsulation," NIST PQC Round-4 submission, 2022.
27. D. J. Bernstein, T. Chou, T. Lange, I. von Maurich, R. Misoczki, R. Niederhagen, E. Persichetti, C. Peters, P. Schwabe, N. Sendrier, J. Szefer, and W. Wang, "Classic McEliece: Conservative Code-Based Cryptography," NIST PQC Round-4 submission, 2022.
28. Open Quantum Safe project, "liboqs: C library for prototyping and experimenting with quantum-resistant cryptography," https://github.com/open-quantum-safe/liboqs, 2016–2026. D. Stebila and M. Mosca, "Post-quantum key exchange for the Internet and the Open Quantum Safe project," *Selected Areas in Cryptography – SAC 2016*, LNCS 10532, pp. 14–37, 2017.
29. C. Peikert, "Lattice Cryptography for the Internet," *PQCrypto 2014*, LNCS 8772, pp. 197–219.
30. J. W. Bos, C. Costello, M. Naehrig, and D. Stebila, "Post-Quantum Key Exchange for the TLS Protocol from the Ring Learning with Errors Problem," *IEEE Symposium on Security and Privacy 2015*, pp. 553–570.
31. P. Gajland, B. de Kock, M. Quaresma, G. Malavolta, and P. Schwabe, "Swoosh: Efficient Lattice-Based Non-Interactive Key Exchange," *USENIX Security 2024*; IACR ePrint 2023/271.
32. K. Hövelmanns, E. Kiltz, S. Schäge, and D. Unruh, "Generic Authenticated Key Exchange in the Quantum Random Oracle Model," *PKC 2020*, LNCS 12111, pp. 389–422. *(For the multi-user NIKE model with adversarial public keys, see Freire–Hofheinz–Kiltz–Paterson [50].)*
33. S. R. Fluhrer, "Cryptanalysis of Ring-LWE-Based Key Exchange with Key Share Reuse," IACR ePrint 2016/085.
34. J. Ding, S. Alsayigh, R. V. Saraswathy, S. Fluhrer, and X. Lin, "Leakage of Signal Function with Reused Keys in RLWE Key Exchange," *IEEE ICC 2017*, pp. 1–6; IACR ePrint 2016/1176.
35. J. Ding, S. Fluhrer, and S. Saraswathy R.V., "Complete Attack on RLWE Key Exchange with Reused Keys, Without Signal Leakage," *ACISP 2018*, LNCS 10946, pp. 467–486.
36. Y. Qin, C. Cheng, and J. Ding, "A Complete and Optimized Key Mismatch Attack on NIST Candidate NewHope," *ESORICS 2019*, LNCS 11736, pp. 504–520.
37. Y. Qin, C. Cheng, X. Zhang, Y. Pan, L. Hu, and J. Ding, "A Systematic Approach and Analysis of Key Mismatch Attacks on Lattice-Based NIST Candidate KEMs," *ASIACRYPT 2021*, LNCS 13093, pp. 92–121.
38. A. Bauer, H. Gilbert, G. Renault, and M. Rossi, "Assessment of the Key-Reuse Resilience of NewHope," *CT-RSA 2019*, LNCS 11405, pp. 272–292.
39. C. Liu, Z. Zheng, and G. Zou, "Key Reuse Attack on NewHope Key Exchange Protocol," *ICISC 2018*, LNCS 11396, pp. 163–176.
40. P. Ravi, S. S. Roy, A. Chattopadhyay, and S. Bhasin, "Generic Side-Channel Attacks on CCA-Secure Lattice-Based PKE and KEMs," *TCHES* 2020(3):307–335.
41. P. Grubbs, V. Maram, and K. G. Paterson, "Anonymous, Robust Post-Quantum Public Key Encryption," EUROCRYPT 2022, LNCS 13277, pp. 402–432.
42. K. Xagawa, "Anonymity of NIST PQC Round 3 KEMs," EUROCRYPT 2022, LNCS 13277, pp. 551–581.
43. C. Cremers, A. Dax, and N. Medinger, "Keeping Up with the KEMs: Stronger Security Notions for KEMs and Automated Analysis of KEM-based Protocols," ACM CCS 2024; IACR ePrint 2023/1933.

44. S. Schmieg, "Unbindable Kemmy Schmidt: ML-KEM is Neither MAL-BIND-K-CT nor MAL-BIND-K-PK," IACR ePrint 2024/523.
45. M. Bolboceanu, Z. Brakerski, R. Perlman, and D. Sharma, "Order-LWE and the Hardness of Ring-LWE with Entropic Secrets," ASIACRYPT 2019, LNCS 11922, pp. 91–120.
46. C. Peikert and Z. Pepin, "Algebraically Structured LWE, Revisited," TCC 2019, LNCS 11891, pp. 1–23.
47. A. Bogdanov, S. Guo, D. Masny, S. Richelson, and A. Rosen, "On the Hardness of Learning with Rounding over Small Modulus," TCC 2016-A, LNCS 9562, pp. 209–224.
48. J.-P. D'Anvers, A. Karmakar, S. Sinha Roy, and F. Vercauteren, "Saber: Module-LWR Based Key Exchange, CPA-Secure Encryption and CCA-Secure KEM," AFRICACRYPT 2018, LNCS 10831, pp. 282–305.
49. R. Avanzi et al., "CRYSTALS-Kyber: Algorithm Specifications and Supporting Documentation (Round 2)," NIST PQC submission, 2019.
50. E. S. V. Freire, D. Hofheinz, E. Kiltz, and K. G. Paterson, "Non-Interactive Key Exchange," PKC 2013, LNCS 7778, pp. 254–271.

A Complete Decryption Correctness Proof

We expand the NAIBI-PKE-OTP decryption fully. Let $\mathbf{t} = A\,\rho(\mathbf{s}) + \mathbf{e}$ with generic error $\mathbf{e} \in R_q^{2\times 2}$, and $\mathbf{t}' = \mathbf{t}+\boldsymbol{\varepsilon}_t$ where $\boldsymbol{\varepsilon}_t \in R_q^{2\times 2}$ is the $\mathbf{t}$-decompression error (entry-wise $\|\cdot\|_\infty \le E_{d_t}$, since $\mathbf{t}' = \mathsf{Decompress}(\bar{\mathbf{t}}, d_t)$ on $\bar{\mathbf{t}} = \mathsf{Compress}(\mathbf{t}, d_t)$). Similarly $\mathbf{u} = A\rho(\mathbf{s}')+\mathbf{e}'$ and $\mathbf{u}' = \mathbf{u}+\boldsymbol{\varepsilon}_u$ with $\|\boldsymbol{\varepsilon}_u\|_\infty \le E_{d_u}$, and let $\mathbf{e}'' \leftarrow \chi_\eta^{k\times k}$ be the encryptor's fresh reconciliation error (Remark 7). Both parties' rounding inputs are

$$\mathbf{r}_{\text{enc}} := \mathbf{t}'\,\rho(\mathbf{s}') + \mathbf{e}'' \in R_q^{k\times k}, \qquad \mathbf{r}_{\text{dec}} := \mathbf{u}'\,\rho(\mathbf{s}) \in R_q^{k\times k}.$$

We bound their difference. First, expand at the un-compressed level:

$$\begin{aligned}\mathbf{t}\,\rho(\mathbf{s}') - \mathbf{u}\,\rho(\mathbf{s}) &= \big(A\rho(\mathbf{s}) + \mathbf{e}\big)\rho(\mathbf{s}') - \big(A\rho(\mathbf{s}') + \mathbf{e}'\big)\rho(\mathbf{s}) \\ &= A\,\rho(\mathbf{s})\rho(\mathbf{s}') + \mathbf{e}\,\rho(\mathbf{s}') - A\,\rho(\mathbf{s}')\rho(\mathbf{s}) - \mathbf{e}'\,\rho(\mathbf{s}) \\ &\overset{(3)}{=} \mathbf{e}\,\rho(\mathbf{s}') - \mathbf{e}'\,\rho(\mathbf{s}).\end{aligned}$$

The A-terms cancel because the *secrets* commute, $\rho(\mathbf{s})\rho(\mathbf{s}') = \rho(\mathbf{s}')\rho(\mathbf{s})$ inside $\rho(\mathcal{A}_\alpha)$, even though A itself does not commute with anything and the errors $\mathbf{e}, \mathbf{e}'$ are generic. Now account for the two compression errors and the fresh error:

$$\begin{aligned}\mathbf{r}_{\text{enc}} - \mathbf{r}_{\text{dec}} &= \big(\mathbf{t}\,\rho(\mathbf{s}') - \mathbf{u}\,\rho(\mathbf{s})\big) + \boldsymbol{\varepsilon}_t\,\rho(\mathbf{s}') - \boldsymbol{\varepsilon}_u\,\rho(\mathbf{s}) + \mathbf{e}'' \\ &= \mathbf{e}\,\rho(\mathbf{s}') - \mathbf{e}'\,\rho(\mathbf{s}) + \boldsymbol{\varepsilon}_t\,\rho(\mathbf{s}') - \boldsymbol{\varepsilon}_u\,\rho(\mathbf{s}) + \mathbf{e}''.\end{aligned}$$

By Lemma 3 the worst-case norm of the noise part is $\le 4\alpha n\eta^2 = 12n$ (plus η from $\mathbf{e}''$), but for correctness we use the statistical model of (9): each of the $k^2 n$ coefficients of $\mathbf{r}_{\text{enc}} - \mathbf{r}_{\text{dec}}$ is a sum of independent centred products with per-coefficient standard deviation

$$\sigma_{\text{KEM}} = \Big(2\sum_{j=1}^{k} n v^2 w_j^2 + \sum_{j=1}^{k} n\tfrac{E_{d_t}^2}{3} w_j^2 v + \sum_{j=1}^{k} n\tfrac{E_{d_u}^2}{3} w_j^2 v + v\Big)^{1/2},$$

with $v = \eta/2$ and $\{w_j^2\} = \{1, \alpha^2\}$ for $k = 2$ ($\{1, \alpha^2, \alpha^2\}$ for $k = 3$); the last term v is the $\mathbf{e}''$ contribution, numerically invisible. Whenever a coefficient's magnitude is $\leq q/8 - 1$, the Peikert recovery guarantee of Lemma 2 applies: $\mathsf{PRec}(\mathbf{r}_{\mathrm{dec}}, \mathsf{PExt}(\mathbf{r}_{\mathrm{enc}}; \bar{\mathbf{e}})[h]) = \mathsf{PExt}(\mathbf{r}_{\mathrm{enc}}; \bar{\mathbf{e}})[k]$ coefficient-wise. By the Gaussian tail (10) this fails for the whole matrix with probability $\leq k^2 n \operatorname{erfc}\big(\frac{q/8-1}{\sigma_{\mathrm{KEM}}\sqrt{2}}\big) \leq 2^{-128}$, so except with negligible probability $k' = k$, $K = \mathsf{H}(k) = \mathsf{H}(k') = K'$, and decryption returns $m' = c \oplus K' = m \oplus K \oplus K = m$. □

B Equivalence between NAIBI-PKE-OTP and "one leg of NAIBI-DH"

To make the "KEM as by-product" framing literal, we exhibit the syntactic correspondence between the OTP-PKE encryption and one leg of the ES-DH protocol.

Encryption (Enc). Take coins r, derive $(\mathbf{s}', \mathbf{e}', \mathbf{e}'', \bar{\mathbf{e}}) \leftarrow \mathsf{XOF}(r)$ with $\mathbf{s}' \in \mathcal{A}_\alpha$ and $\mathbf{e}', \mathbf{e}'' \in R_q^{k\times k}$ (generic). Compute

$$\begin{aligned} \mathbf{u} &= A\,\rho(\mathbf{s}') + \mathbf{e}' \in R_q^{k\times k}, \\ (h, k) &= \mathsf{PExt}\big(\mathbf{t}\,\rho(\mathbf{s}') + \mathbf{e}''; \bar{\mathbf{e}}\big), \\ K &= \mathsf{H}(k), \qquad c = m \oplus K. \end{aligned}$$

Send $(\bar{\mathbf{u}}, c, h) = \big(\mathsf{Compress}(\mathbf{u}, d_u), c, h\big)$.

Alice's local computation in NAIBI-ES-DH. Generate ephemeral $\mathbf{s}_A^{\mathrm{eph}} \in \mathcal{A}_\alpha$ and $\mathbf{e}_A^{\mathrm{eph}} \in R_q^{k\times k}$ (generic). Compute

$$\mathbf{u}_A = A\,\rho(\mathbf{s}_A^{\mathrm{eph}}) + \mathbf{e}_A^{\mathrm{eph}} \in R_q^{k\times k}, \qquad k_A = \mathsf{PRec}\big(\mathbf{t}_B\,\rho(\mathbf{s}_A^{\mathrm{eph}}), h_B\big).$$

Send $\bar{\mathbf{u}}_A$, receive hint h_B from Bob (computed by Bob from $\mathbf{u}'_A\rho(\mathbf{s}_B) + \mathbf{e}''_B$), derive key $K_A = \mathsf{KDF}(k_A \| \cdots)$.

Side-by-side. Identifying $\mathbf{s}' \leftrightarrow \mathbf{s}_A^{\mathrm{eph}}$, $\mathbf{e}' \leftrightarrow \mathbf{e}_A^{\mathrm{eph}}$, $\mathbf{t} \leftrightarrow \mathbf{t}_B$, the formulas are *identical* up to: (i) the hint direction (KEM: encryptor adds $\mathbf{e}''$, extracts with the doubling bits $\bar{\mathbf{e}}$ and sends h; ES-DH: Bob adds $\mathbf{e}''_B$, extracts and sends h_B; the recovery side calls PRec in both); (ii) the encryption uses an additional $\mathsf{Compress}(\mathbf{u}, d_u)$ for ciphertext compression, present in neither ES-DH nor EE-DH; (iii) the encryption appends $c = m \oplus \mathsf{H}(k)$ to publish a ciphertext, while DH treats $\mathsf{H}(k)$ *itself* as the shared key; (iv) in the KEM, $(\mathbf{e}'', \bar{\mathbf{e}})$ are derandomised from the FO coins r, whereas the agreement modes sample them fresh.

Conclusion. NAIBI-PKE-OTP is structurally NAIBI-ES-DH with (i) $\mathbf{u}$ additionally compressed for bandwidth, (ii) the Peikert-extracted raw key k used as a one-time-pad key on a chosen plaintext. Apply FO on top of NAIBI-PKE-OTP and obtain the IND-CCA2 KEM. In particular, the underlying assumption (structured-secret MLWE with a fresh reconciliation error: secret in $\rho(\mathcal{A}_\alpha)$, errors generic in $R_q^{k\times k}$, cf. Theorem 6) is *the same* as the one powering ES-DH, and noise budgets are governed by *the same* commutativity identity.

C Sensitivity of δ to the Fresh Reconciliation Error and the Threshold Shift

Two features of the reconciliation kernel enter the noise model of the main text and are worth isolating numerically. The fresh error $\mathbf{e}''$ of Remark 7 adds a single term $v = \eta/2 = 0.5$ to σ^2 in (9), and the randomised doubling moves the agreement threshold from $q/8$ to $q/8 - 1$ (Lemma 2). Both are cheap; the table below prices them, by the same erfc-plus-union-bound method as the main text.

Instance	σ_{base}	$\sigma_{\text{depl.}}$	thr. $q/8$	thr. $q/8-1$	$\log_2 \delta_{\text{base}}$	$\log_2 \delta_{\text{depl.}}$
NAIBI-I	172.8	172.8	2432	2431	-136.0	-135.9
NAIBI-III	105.3	105.3	1664	1663	-173.2	-173.0
NAIBI-V	300.8	300.8	5120	5119	-201.3	-201.2

Here the base columns drop the v term and use the threshold $q/8$; the depl. columns are the deployed model. The added variance $v = 1/2$ is negligible against $\sigma^2 \approx 10^4$, and the unit threshold shift costs < 0.3 bits of $\log_2 \delta$ (the σ and threshold columns are those of Tables 4 and 5). Every displayed value in Tables 4, 5 and 6 is therefore unchanged at the precision reported, and the parameter sets $(n, k, q, \eta, d_u, d_t)$ need no adjustment. The uncompressed full-matrix options NAIBI-{I,III,V}-pc of Remark 2 share the same sets, with $|\mathsf{pk}| = 3872/3616/8224$ B at Levels I/III/V.

D Empirical Validation of Heuristic 9

Heuristic 9 is the only non-rigorous ingredient of the binding proof (Theorem 10). We summarise the sampling-based evidence; full scripts accompany the reference implementation.

Protocol. For each parameter set we draw $N = 10^7$ independent pairs $(\mathbf{s}_0, \mathbf{s}_1) \leftarrow \chi_1(\mathcal{A}_\alpha)^2$ conditioned on $\mathbf{s}_0 \neq \mathbf{s}_1$, compute $\Delta\mathbf{s} = \mathbf{s}_0 - \mathbf{s}_1$, map to the NTT domain $\mathcal{A}_\alpha \cong \prod_{i=1}^{n} \mathbb{F}_{q^k}$, and record (i) the per-slot vanishing frequency $\Pr[\Delta\mathbf{s}^{(i)} = 0]$ and (ii) the empirical pairwise correlations between the vanishing indicators of distinct slots and between the slot values $\rho(\mathbf{s}_b)^{(i)}$.

Findings. Table 10 summarises the results for all three parameter sets at $N = 10^7$.

– **Per-slot vanishing frequency (Heuristic 9(i), bound $\leq 2/q$).** The worst-case per-slot frequency is 2×10^{-7} (one or two observed events at a single slot), at least two orders of magnitude below $2/q$ for every instance. At all three $k = 2$ instances the measured mean $\hat{p}$ matches q^{-k} to within Poisson sampling noise: NAIBI-I 13 events vs. $\mu \approx 13.5$ expected ($\hat{p} = 2.54 \times 10^{-9}$ vs. $q^{-k} = 2.64 \times 10^{-9}$), NAIBI-III 27 vs. $\mu \approx 28.8$, and NAIBI-V 9 vs. $\mu \approx 6.1$ (all within $\approx 1\,\sigma$ of the Poisson mean). All n per-slot frequencies lie strictly below $2/q$.

Table 10. Empirical validation of Heuristic 9 ($N = 10^7$ honest-secret pairs per instance, seed fixed for reproducibility). $\hat{p}$ is the mean per-slot vanishing frequency of $\Delta\mathbf{s}^{(i)}$ in $\mathbb{F}_{q^k}$; "Events" is the total count of vanishing (slot, pair) observations across all $n \cdot N$ pairs; "χ^2 rej." counts rejections of per-slot $\mathbb{Z}_q$-uniformity tests at $\alpha_0 = 0.01$ (out of $60 \times k$ tests); $|r|_{\max}$ is the maximum Pearson correlation of NTT slot values between distinct slots (over $40 \times k$ tested pairs).

Instance	k	q^{-k}	$2/q$	$\hat{p}$	Events	χ^2 rej.	$\lvert r\rvert_{\max}$
NAIBI-I	2	2.64×10^{-9}	1.03×10^{-4}	2.54×10^{-9}	13	2/120	7.0×10^{-4}
NAIBI-III	2	5.64×10^{-9}	1.50×10^{-4}	5.27×10^{-9}	27	1/120	8.3×10^{-4}
NAIBI-V	2	5.96×10^{-10}	4.88×10^{-5}	8.79×10^{-10}	9	1/120	6.2×10^{-4}

- **Pairwise slot-vanishing correlations (Heuristic 9(i), independence).** No vanishing event is shared by any of the 40 randomly selected cross-half slot pairs; the Phi coefficient $\varphi \equiv 0$ for all pairs, consistent with full independence. (At NAIBI-I the expected number of joint vanishings is $\approx 10^7 \cdot (q^{-k})^2 \approx 7 \times 10^{-11}$, undetectable at any feasible N.)
- **Slot-value uniformity and independence (Heuristic 9(ii)).** Across 360 chi-squared goodness-of-fit tests (60 sampled slots $\times\, k$ components per instance, expected counts per bin 244–751), only 4 are rejected at $\alpha_0 = 0.01$, against an expectation of 3.6 under H_0; median p-values are 0.48–0.53. Cross-slot Pearson correlations between NTT values at distinct slots satisfy $|r| \leq 8.3 \times 10^{-4}$ with no $|z| > 3.5\,\sigma$ outlier in 240 tested pairs, consistent with independence.

Caveat. The sampling validates Heuristic 9(i) for honest secrets; it is not needed in the MAL model, where the slot-count is instead given by the assumption-free Lemma 10. The key-independent ingredients of the binding proof — the per-coordinate bounds (Lemmas 9 and 11) and slot-wise invertibility under non-splitness — do not rely on this appendix.