




Evasive LWE Assumptions: Definitions, Classes, and Counterexamples

Chris Brzuska¹, Akin Ünal², and Ivy K. Y. Woo¹

¹ Aalto University, Espoo, Finland

² ISTA, Klosterneuburg, Austria

Abstract. The evasive LWE assumption, proposed by Wee [Eurocrypt’22 Wee] for constructing a lattice-based optimal broadcast encryption, has shown to be a powerful assumption, adopted by subsequent works to construct advanced primitives ranging from ABE variants to obfuscation for null circuits. However, a closer look reveals significant differences among the precise assumption statements involved in different works, leading to the fundamental question of how these assumptions compare to each other. In this work, we initiate a more systematic study on evasive LWE assumptions:

- (i) Based on the standard LWE assumption, we construct simple counterexamples against three private-coin evasive LWE variants, used in [Crypto’22 Tsabary, Asiacrypt’22 VWW, Crypto’23 ARYY] respectively, showing that these assumptions are unlikely to hold.
- (ii) Based on existing evasive LWE variants and our counterexamples, we propose and define three classes of plausible evasive LWE assumptions, suitably capturing all existing variants for which we are not aware of non-obfuscation-based counterexamples.
- (iii) We show that under our assumption formulations, the security proofs of [Asiacrypt’22 VWW] and [Crypto’23 ARYY] can be recovered, and we reason why the security proof of [Crypto’22 Tsabary] is also plausibly repairable using an appropriate evasive LWE assumption.

Note. Main changes compared to the proceedings version:

- We strengthen the heuristic obfuscation-based counterexample by [VWW22] against all private-coin evasive LWE variants into a *provable* counterexample, assuming only indistinguishability obfuscation for null circuits and the hardness of LWE with super-polynomial modulus-to-noise ratio. This can be viewed as concrete evidence of the qualitative difference between public- and private-coin evasive LWEs. See Section 7 for details.
- In the proceedings version of this work, we stated all our proposed evasive LWE assumption families with respect to *arbitrary* distributions of matrix \mathbf{B} . Unfortunately, we later noticed that allowing for arbitrarily distributed matrices \mathbf{B} leads to simple counterexamples. See Section 8 for details. Thus, for the time being, we advise to use evasive LWE only with *uniformly* distributed \mathbf{B} . Understanding how the choice of the matrix \mathbf{B} distribution affects hardness of evasive LWE is an interesting open question. We left our definitions of the proposed evasive LWE assumption families in Section 4 unchanged and only added a note referring to Section 8 which discusses the new counterexamples.

1 Introduction

Resolving a decade-long open problem, Wee [Wee22] constructs a lattice-based ciphertext-policy attribute-based encryption (CP-ABE) for NC1 with parameters independent of the circuit size, implying an optimal broadcast encryption, under a new assumption called the evasive LWE assumption. Roughly, the assumption states that, for any PPT Samp outputting an arbitrary matrix \mathbf{P} and auxiliary information aux containing all coin tosses used by Samp ,

$$\begin{array}{ll} \text{if} & (\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top, \mathbf{s}^\top \mathbf{P} + \mathbf{e}_1^\top, \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \$, \text{aux}) \\ \text{then} & (\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}), \end{array} \quad (1)$$

for a uniformly random matrix $\mathbf{B} \leftarrow \$ \mathbb{Z}_q^{n \times m}$, a uniformly random LWE secret $\mathbf{s} \leftarrow \$ \mathbb{Z}_q^n$, Gaussian errors $\mathbf{e}_0, \mathbf{e}_1$ of appropriate dimensions, and where $\mathbf{B}^{-1}(\mathbf{P})$ denotes a short Gaussian preimage of \mathbf{P} with respect to \mathbf{B} , i.e. it holds that $\mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P}) = \mathbf{P} \bmod q$ and $\mathbf{B}^{-1}(\mathbf{P})$ is short in Euclidean norm. Intuitively, the assumption postulates that, to distinguish LWE samples $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top \bmod q$ given a short preimage $\mathbf{B}^{-1}(\mathbf{P})$ as hint, the only thing one can do with this hint is to right-multiply it to $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top \bmod q$ to obtain

another LWE sample $\mathbf{s}^T \mathbf{P} + \mathbf{e}_1^T \bmod q$ and try to distinguish the latter. We shall call this the *public-coin* evasive LWE assumption, to highlight that all random coins of \mathbf{Samp} are given to the distinguisher via \mathbf{aux} . Subsequently, public-coin evasive LWE was adapted by [WWW22,HLL24], to prove security of a multi-authority (MA-)ABE scheme without a random oracle and construct a variant of inner-product functional encryption, respectively.

Concurrently and subsequently, a number of works considered what we shall call *private-coin* variants of evasive LWE, where \mathbf{Samp} does not need to provide all its randomness to the distinguisher. These private-coin variants of evasive LWE (sometimes applied along with other new lattice assumptions) have shown to imply further advanced primitives including multi-input (MI-)ABE [ARYY23], CP-ABE for unbounded depth circuits [AKY24], null-iO [VWW22], and witness encryption [Tsa22,VWW22]. Notably, these primitives tend to be believed to be stronger than what is currently achievable by public-coin evasive LWE.

Underlying the name of evasive LWE, however, are significant differences among the actual assumption statements involved in all aforementioned works, in addition to the distinction between public and private-coin. For example, the public-coin evasive LWE variants in [WWW22,HLL24] involve multiple independent matrices and LWE samples $(\mathbf{B}_i, \mathbf{s}_i^T \mathbf{B}_i + \mathbf{e}_i^T \bmod q)$, where the \mathbf{s}_i 's can be identical or different, and [HLL24] further requires a joint preimage where $\mathbf{B}_i^{-1}(\mathbf{P}) = \mathbf{B}_j^{-1}(\mathbf{P})$ for all i, j . In the private-coin world, [VWW22] formulates an evasive LWE assumption where the matrices \mathbf{B}, \mathbf{P} are not explicitly stated in the *if* and *then* joint distributions, whereas [ARYY23] includes \mathbf{B} into both distributions, but not \mathbf{P} . Unfortunately, up to now, we have little understanding as to how these and other modifications affect the strength and plausibility of the assumptions, with the only cryptanalytic insight being a heuristic obfuscation-based counterexample by [VWW22] against some private-coin variants.

1.1 Our Contributions

We initiate a systematic study on evasive LWE assumptions, summarised by three aspects:

Falsifying Existing Private-coin Variants. We construct counterexamples against subclasses of existing evasive LWE variants. Specifically, we give three examples to show that, each of the three private-coin evasive LWE assumptions stated in [Tsa22,VWW22,ARYY23], respectively, are unlikely to hold. All our counterexamples are conceptually simple, based on the standard LWE assumption and do not rely on obfuscation. Our counterexamples are summarised in Section 2.2 and formalised in Section 5. We sketch additional counterexamples in Section 8 and Appendix C.

Classifications and Definitions. Based on existing variants together with the essences of our counterexamples, we propose a classification of plausible evasive LWE assumptions. Our proposed three families are summarised below:

1. Public-coin evasive LWE, i.e. the PPT sampler \mathbf{Samp} does not hide any of its computation from the distinguisher;
2. Private-coin binding evasive LWE, where (i) \mathbf{Samp} does not input the matrix \mathbf{B} , and (ii) the matrices \mathbf{B}, \mathbf{P} are given to the distinguisher;
3. Private-coin hiding evasive LWE, where (i) \mathbf{Samp} does not input the matrix \mathbf{B} , (ii) \mathbf{B} is not given to the distinguisher, and (iii) the matrix \mathbf{P} is provably sufficiently hidden from the distinguisher.

For each family we formulate a general definition, suitably capturing all variants in existing works that are not subject to simple counterexamples (see Remark 3 for a discussion of obfuscation-based counterexamples). We summarise the rationale in Section 2.3 and give precise definitions in Section 4. We hope that this will serve as a first step to provide the community with a language for communicating about evasive LWE assumptions, leaving its intuition unchanged while simultaneously expressing the qualitative differences between the actual assumptions involved.

Implications on Existing Constructions. We show that the assumption instances in the security proofs of [VWW22,ARYY23] fall under our proposed families of plausible evasive LWE, as such, their related constructions remain secure under our assumptions. More concretely, in Section 6 we prove that the evasive LWE instances in the proof of [VWW22] satisfy our condition of \mathbf{P} being sufficiently hidden, thus falling into the private-coin hiding family. The assumption instances of [ARYY23] fall into the private-coin binding family directly by definition. For the proof of [Tsa22], we discuss in Section 2.4 why we believe it may be repairable with an alternative and plausible evasive LWE assumption.

2 Overview

In Section 2.1 we review a number of evasive LWE variants. In Section 2.2 we sketch our counterexamples against three subclasses, which leads to our proposal of three plausible evasive LWE families summarised in Section 2.3. For existing works that rely on assumptions affected by our counterexamples, we discuss in Section 2.4 to which extent their security proofs may be repairable.

We adopt simplified notation in this overview: Operations are understood to be over \mathbb{Z}_q and we suppress mod q expressions. To denote noise terms, we use curly underline $\underline{\sim}$, e.g. $\underline{\underline{s}}^T \mathbf{B}$ means $\mathbf{s}^T \mathbf{B} + \mathbf{e}^T$ where \mathbf{e} is short relative to q . We use $\$$ to refer to uniformly random values, where multiple $\$$ signs in a joint distribution are understood to be independent uniform samples.

2.1 Existing Evasive LWE Variants and Classifications

To aid understanding the differences among the evasive LWE variants below, we provide two partition grids in Fig. 1 which we will cross-reference with.

Public-coin Evasive LWE. We recall again Wee’s public-coin evasive LWE assumption: For any PPT Samp outputting an arbitrary matrix \mathbf{P} and auxiliary information aux containing all randomness used by Samp ,

$$\begin{array}{ll} \text{if} & (\mathbf{B}, \mathbf{P}, \underline{\underline{s}}^T \mathbf{B}, \underline{\underline{s}}^T \mathbf{P}, \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \$, \text{aux}) \\ \text{then} & (\mathbf{B}, \mathbf{P}, \underline{\underline{s}}^T \mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \end{array} \quad (2)$$

for uniformly random \mathbf{B} and uniformly random LWE secret \mathbf{s} .³ Subsequently, Waters, Wee and Wu [WWW22] define a variant consisting of polynomially many pairs of matrices $(\mathbf{B}_i, \mathbf{P}_i)_i$ and their respective LWE samples $\underline{\underline{s}}_i^T \mathbf{B}_i, \underline{\underline{s}}_i^T \mathbf{P}_i$. More recently, Hsieh, Lin, and Luo [HLL24] define a public-coin variant consisting also of pairs of $(\mathbf{B}_i, \mathbf{P}_i)_i$, but with LWE samples of the form $(\underline{\underline{s}}^T \mathbf{B}_i, \underline{\underline{s}}^T \mathbf{P}_i)_i$ sharing the same secret \mathbf{s} , with a structured error distribution instead of a random Gaussian, and further requiring joint preimages satisfying $\mathbf{B}_i^{-1}(\mathbf{P}) = \mathbf{B}_j^{-1}(\mathbf{P})$ for all i, j .

Private-coin without \mathbf{B}, \mathbf{P} . Vaikunthanathan, Wee and Wichs [VWW22] define the following private-coin evasive LWE assumption: For any PPT Samp which outputs (arbitrary) LWE secret \mathbf{S} , matrix \mathbf{P} , and auxiliary information aux (not necessarily containing all randomness used by Samp),

$$\begin{array}{ll} \text{if} & (\underline{\underline{\mathbf{S}}}\mathbf{B}, \underline{\underline{\mathbf{S}}}\mathbf{P}, \text{aux}) \approx_c (\$, \$, \text{aux}) \\ \text{then} & (\underline{\underline{\mathbf{S}}}\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \end{array} \quad (3)$$

for uniformly random \mathbf{B} . This variant corresponds to the [blue area](#) in Fig. 1b. We observe that, unlike in Eq. (2), in this variant, \mathbf{B} is not included in the if and then distributions, and \mathbf{P} also not necessarily. However, since both aux, \mathbf{P} are generated by Samp , information about \mathbf{P} may or may not be included in the distributions via aux , e.g., Samp could choose $\text{aux} = \mathbf{P}$. Using the above, [VWW22] prove that the well-known GGH15 encodings [GGH15] are secure, implying the existence of null-iO and witness encryption, and recent works also use the above evasive LWE variant to construct SNARG for UP as well as universal computational extractors [MPV24,CM24].

Private-coin without \mathbf{P} . Agrawal, Rossi, Yadav and Yamada (ARYY) [ARYY23] defined the following private-coin variant: For any PPT Samp outputting (arbitrary) LWE secret \mathbf{S} , matrix \mathbf{P} , and auxiliary information aux (not necessarily containing all randomness used by Samp),

$$\begin{array}{ll} \text{if} & (\mathbf{B}, \underline{\underline{\mathbf{S}}}\mathbf{B}, \underline{\underline{\mathbf{S}}}\mathbf{P}, \text{aux}) \approx_c (\mathbf{B}, \$, \$, \text{aux}) \\ \text{then} & (\mathbf{B}, \underline{\underline{\mathbf{S}}}\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{B}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \end{array} \quad (4)$$

³ [Wee22] includes an additional matrix \mathbf{A} and LWE samples $\underline{\underline{s}}^T \mathbf{A}$ in the joint distributions for further expressiveness, which we omit in this overview.

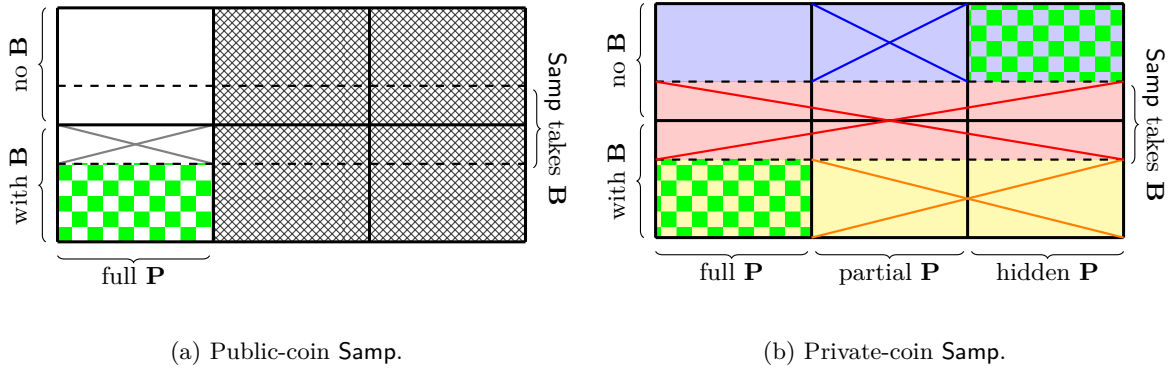


Fig. 1: Partition of evasive LWE assumptions. Background colour $\{\text{yellow, blue, red}\} = \text{Assumption } \{(4), (3), (5)\}$. Cross in $\{\text{yellow, blue, red}\} = \text{Counterexample } \{1, 2, 3\}$ sketched in Section 2.2 and formalised in Section 5. Cross in gray = heuristic counterexample. Green checkboxes are the three proposed classes in Section 2.3.

for uniformly random \mathbf{B} . This corresponds to the **yellow area** in Fig. 1b. We notice that here the matrix \mathbf{B} is included in the distributions, and \mathbf{P} again may or may not be via aux . ARYY show that Eq. (4) implies another private-coin variant, which is almost identical except that aux can be partitioned as $(\text{aux}_1, \text{aux}_2)$, with aux_1 provably pseudorandom, and \mathbf{P} is efficiently computable from aux_2 . The latter is used for proving security of their MI-ABEs⁴, and subsequently also by [AKY24] for proving their CP-ABE for unbounded depth circuits⁵.

Samp generates $\mathbf{B}^{-1}(\mathbf{P})$. Concurrent to [VWW22], Tsabary [Tsa22] proposes a similar flavour of assumption for constructing witness encryption. Putting formulation differences aside, Tsabary’s assumption can be summarised as follows: For any PPT **Samp** which inputs a matrix \mathbf{B} with its trapdoor $\text{td}_{\mathbf{B}}$, and outputs LWE secret \mathbf{S} , target matrix \mathbf{P} , preimages $\mathbf{B}^{-1}(\mathbf{P})$ sampled using $\text{td}_{\mathbf{B}}$, and auxiliary information aux (not necessarily containing all randomness of **Samp**),

$$\begin{array}{llll}
 \text{if} & (\underline{\mathbf{S}}\mathbf{B}, \underline{\mathbf{S}}\mathbf{P}, \text{aux}) & \approx_c & (\$, \$, \text{aux}) \\
 \text{then} & (\underline{\mathbf{S}}\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) & \approx_c & (\$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}).
 \end{array} \tag{5}$$

This variant covers the **red area** in Fig. 1b. The if and then distributions are identical to Eq. (3). The crucial difference is, however, that **Samp** takes the matrix \mathbf{B} as input and samples the preimages $\mathbf{B}^{-1}(\mathbf{P})$ itself, in contrast to other variants where the preimages are provided by the challenger. This formulation necessitates **Samp** knowing \mathbf{B} and $\text{td}_{\mathbf{B}}$, and as a consequence \mathbf{S}, \mathbf{P} and aux can be arbitrarily correlated with \mathbf{B} .

Summarising from above, we obtain the following factors:

- Public-coin vs. private-coin **Samp** (Fig. 1a vs. Fig. 1b).
- Whether \mathbf{B} is included in the if and then distributions (left axis of Fig. 1b).
- Whether \mathbf{P} is included in the if and then distributions. We separately consider, informally for now, \mathbf{P} is fully available, partially available, or fully hidden from the distinguisher (bottom axis of Fig. 1b), the last to be expanded later.
- Whether **Samp** inputs the matrix \mathbf{B} (right axis of Fig. 1b).

The above jointly form a partition of Fig. 1, as marked by the figure labels.⁶

Remark 1 (On Assumption Strength). The classes of samplers **Samp** in evasive LWE form natural inclusions corresponding to the strength of the assumptions: The largest class contains the most **Samp**, which

⁴ For their MI-ABE for \mathbf{P} , along with a new assumption called extended tensor LWE.

⁵ Along with a new assumption called circular tensor LWE.

⁶ The partition has not taken into account whether the LWE secret \mathbf{S} is generated by **Samp**. This is to be discussed at the end of Section 2.2.

can take the most inputs, e.g. \mathbf{B} , and output arbitrary matrices, representing the strongest assumptions. Underlying are various subclasses, e.g. one containing Samp which does not input \mathbf{B} , and one containing Samp which must output \mathbf{P} in plain, each forming weaker assumptions. Similarly, public-coin samplers are a subclass of private-coin ones.

In contrast, including \mathbf{B} in the distributions or not rather makes the assumptions incomparable: Including \mathbf{B} leads to a stronger if condition to be satisfied, but the assumption simultaneously asserts a stronger then statement with \mathbf{B} .

2.2 Counterexamples

In order to invalidate an evasive LWE assumption, our goal is to construct a PPT Samp such that, with respect to Samp , (1) the if statement holds (assuming plausible assumptions), but (2) the then statement does not.

Counterexample 1: private-coin with \mathbf{B} and partial/hidden \mathbf{P} . We give a counterexample for the case where the distinguisher of if receives the matrix \mathbf{B} , but not \mathbf{P} , corresponding to a subclass of Eq. (4) with a yellow cross in Fig. 1b. The idea is simple: Given \mathbf{B} and $\mathbf{B}^{-1}(\mathbf{P})$, the distinguisher of then can recover \mathbf{P} , so that if \mathbf{P} contains useful information for distinguishing, this can be used by then but not if. Concretely, let Samp return the following:

$$\mathbf{P} = \left(\mathbf{P}_1, \mathbf{P}_2 = \begin{bmatrix} \mathbf{u}^\top \\ \mathbf{R} \end{bmatrix} \right), \quad \text{aux} = ()$$

where \mathbf{u} is a short vector such that $\mathbf{P}_1 \mathbf{u} = \mathbf{0}$, i.e. Samp samples \mathbf{P}_1 together with a trapdoor to generate \mathbf{u} , and \mathbf{R} is uniformly random. To see that the if statement holds, observe that

$$(\mathbf{B}, \underbrace{\mathbf{s}^\top \mathbf{B}}, \underbrace{\mathbf{s}^\top \mathbf{P}_1}, \underbrace{\mathbf{s}^\top \mathbf{P}_2}) \stackrel{\text{stat.}}{\approx} (\mathbf{B}, \underbrace{\mathbf{s}^\top \mathbf{B}}, \$, \$) \stackrel{\text{LWE}}{\approx} (\mathbf{B}, \$, \$, \$)$$

where the first statistical statement holds since \mathbf{R} and \mathbf{P}_1 are (close to) uniformly random (and unknown), and the second follows by LWE, as \mathbf{B} is uniformly random. For the distinguisher of then, observe that, when given

$$(\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}_1, \mathbf{P}_2), \underbrace{\mathbf{s}^\top \mathbf{P}_1}, \underbrace{\mathbf{s}^\top \mathbf{P}_2}),$$

it can compute

$$\mathbf{P}_2 \leftarrow \mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P}_2); \quad \begin{bmatrix} \mathbf{u}^\top \\ \mathbf{R} \end{bmatrix} \leftarrow \mathbf{P}_2; \quad \text{Test } \underbrace{\mathbf{s}^\top \mathbf{P}_1} \cdot \mathbf{u} \approx? \mathbf{0}, \quad (6)$$

where $\underbrace{\mathbf{s}^\top \mathbf{P}_1} \cdot \mathbf{u}$ would likely not be short when replacing $\underbrace{\mathbf{s}^\top \mathbf{P}_1}$ with a random vector. Note that the example works also when \mathbf{P}_1 is given in aux .

Counterexample 2: private-coin without \mathbf{B} and partial \mathbf{P} . Next we turn to the case where \mathbf{B} is not available to the distinguisher of if but \mathbf{P} is partially available, corresponding to a subclass of Eq. (3) with a blue cross in Fig. 1b. Here, our Samp is almost identical to the previous counterexample, except that now it lets $\text{aux} = \mathbf{P}_1$. Similar to above, for the if condition, we have

$$(\underbrace{\mathbf{s}^\top \mathbf{B}}, \underbrace{\mathbf{s}^\top \mathbf{P}_1}, \underbrace{\mathbf{s}^\top \mathbf{P}_2}, \mathbf{P}_1) \stackrel{\text{stat.}}{\approx} (\$, \underbrace{\mathbf{s}^\top \mathbf{P}_1}, \$, \mathbf{P}_1) \stackrel{\text{LWE}}{\approx} (\$, \$, \$, \mathbf{P}_1)$$

where the first relation is due to \mathbf{B}, \mathbf{R} being uniform and unknown to the distinguisher. For distinguishing the then distribution, since \mathbf{B} is not contained in the joint distribution any longer, the distinguisher takes an extra step to recover it. Namely, our crucial observation is that when \mathbf{P}_1 has (at least) as many columns as that of \mathbf{B} and \mathbb{Z}_q is a field⁷ then with high probability, \mathbf{B} is fully determined given $(\mathbf{B}^{-1}(\mathbf{P}_1), \mathbf{P}_1)$, and is efficiently recoverable by a system of linear equations from the relation $\mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P}_1) = \mathbf{P}_1$. This follows since each entry of $\mathbf{B}^{-1}(\mathbf{P}_1)$ is Gaussian-distributed so that $\mathbf{B}^{-1}(\mathbf{P}_1)$ has full rank over \mathbb{Z}_q if the variance of its entries is large enough. After recovering \mathbf{B} , the distinguisher performs the same steps as Eq. (6).

⁷ The condition of \mathbb{Z}_q being a field can be naturally extended to the ring setting, where \mathcal{R}_q splits into subfields of super-polynomial size, so that a random element is invertible with high probability. For simplicity we only consider \mathbb{Z}_q in the rest.

Remark 2 (When aux contains s). Counterexample 2 relies on that \mathbf{aux} does not contain information about the LWE secret \mathbf{s} , so that $(\mathbf{s}^\top \mathbf{P}_1, \mathbf{P}_1) \approx (\$, \mathbf{P}_1)$ holds by LWE. Against settings where \mathbf{aux} is required to contain \mathbf{s} (e.g. that in [VWW22]), one can modify the above slightly to yield another counterexample: Write $\mathbf{s}^\top = (\mathbf{s}_1^\top, \mathbf{s}_2^\top)$, $\mathbf{B} = \begin{pmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \end{pmatrix}$, $\mathbf{P}_1 = \begin{pmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} \end{pmatrix}$. Let $\mathbf{aux} = (\mathbf{s}, \mathbf{Q}_{11}, \mathbf{Q}_{22})$. Now, $(\mathbf{s}^\top \mathbf{P}_1, \mathbf{aux}) \approx (\$, \mathbf{aux})$ still holds since $\mathbf{s}_2^\top \mathbf{Q}_{21}$ and $\mathbf{s}_1^\top \mathbf{Q}_{12}$ are uniform without information about $\mathbf{Q}_{21}, \mathbf{Q}_{12}$. To distinguish then, use the relations $\mathbf{B}_1 \cdot \mathbf{B}^{-1} \begin{pmatrix} \mathbf{Q}_{11} \\ \mathbf{Q}_{21} \end{pmatrix} = \mathbf{Q}_{11}$ and $\mathbf{B}_2 \cdot \mathbf{B}^{-1} \begin{pmatrix} \mathbf{Q}_{12} \\ \mathbf{Q}_{22} \end{pmatrix} = \mathbf{Q}_{22}$ to recover \mathbf{B} , the rest is the same.

Counterexample 3: Private-coin where Samp inputs B. Our last counterexample applies whenever a private-coin \mathbf{Samp} inputs the matrix \mathbf{B} , spanning the whole area in red in Fig. 1b, marked with a red cross.⁸ The idea is again simple: If \mathbf{Samp} knows the matrices $\mathbf{B}, \mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2)$ and wishes to make a secret available to the distinguisher of then who additionally receives a short preimage $\mathbf{B}^{-1}(\mathbf{p}_1)$, then an immediate strategy is to encrypt, specifically with the dual-Regev encryption. As is usual in the lattice-setting, for one to distinguish LWE, an appropriate short preimage suffices. Therefore, let \mathbf{Samp} on input \mathbf{B} output

$$\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2), \quad \mathbf{aux} = \text{ctxt}_{\mathbf{u}_2} = (\mathbf{RB}, \mathbf{Rp}_1 + \lfloor q/2 \rfloor \mathbf{u}_2),$$

where $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2) = \mathbf{B} \cdot (\mathbf{u}_1, \mathbf{u}_2)$ are the images under random short (e.g. Gaussian or binary) preimages $(\mathbf{u}_1, \mathbf{u}_2)$ sampled by \mathbf{Samp} itself, and $\text{ctxt}_{\mathbf{u}_2}$ the dual-Regev encryption of \mathbf{u}_2 under the public key $(\mathbf{B}, \mathbf{p}_1)$. To see that the if statement holds, observe

$$(\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B}, \mathbf{s}^\top \mathbf{P}, \text{ctxt}_{\mathbf{u}_2}) \stackrel{c}{\approx} (\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B}, \mathbf{s}^\top \mathbf{P}, \$) \stackrel{c}{\approx} (\mathbf{B}, \mathbf{P}, \$, \$, \$)$$

where the first \approx_c follows from security of dual-Regev encryption, the second by LWE. Finally, to distinguish the then distributions, we can use $\mathbf{B}^{-1}(\mathbf{p}_1)$ to decrypt by observing that $\mathbf{RB} \cdot \mathbf{B}^{-1}(\mathbf{p}_1) \approx \mathbf{Rp}_1$,

so we can obtain \mathbf{u}_2 . Next, we observe that $\mathbf{B}^{-1}(\mathbf{p}_2) - \mathbf{u}_2$ is a short preimage of $\mathbf{0}$ for the image \mathbf{p}_2 , i.e. $\mathbf{B} \cdot (\mathbf{B}^{-1}(\mathbf{p}_2) - \mathbf{u}_2) \approx \mathbf{0}$. We highlight that in the above we can prove the hardest if condition where \mathbf{B}, \mathbf{P} are in the joint distributions, and we can distinguish the weakest then distributions without \mathbf{B} and \mathbf{P} , since the distinguisher does not require them. As such, the above counterexample covers the whole red area in Fig. 1b spanning all settings with/without \mathbf{B}, \mathbf{P} .

On LWE secret s. We are not aware of counterexamples which specifically require \mathbf{Samp} to know/generate the LWE secret \mathbf{s} . More concretely, for any attack which targets a subclass where \mathbf{Samp} needs to know/generate \mathbf{s} , we realise a more general attack which works over the corresponding superclass where \mathbf{Samp} does not know \mathbf{s} . To illustrate, when \mathbf{Samp} is allowed to generate the LWE secret \mathbf{s} , we can simplify counterexample 1 against Eq. (4) above, by letting \mathbf{Samp} embed \mathbf{s} in e.g. the first row \mathbf{P} , thus skipping generating \mathbf{P}_1 and \mathbf{u} . However, once such “embed-secret-in- \mathbf{P} ” mechanism is possible, the attack generalises to one where \mathbf{Samp} does not need to know \mathbf{s} . Indeed, all attacks that we are aware of work by embedding some secret chosen by \mathbf{Samp} , which can be recovered given $\mathbf{B}^{-1}(\mathbf{P})$. As long as this can be achieved, one can naturally pick a secret that breaks LWE by interacting with (parts of) \mathbf{P} , without interacting with \mathbf{s} .

Remark 3 (Obfuscation-based Counterexample). [Wee22, VWW22] suggested heuristic obfuscation-based counterexamples which apply to all private-coin variants discussed above. In Section 7, we show that the adapted versions of their counterexamples can be proven based on LWE and the existence of null-iO.

Briefly summarised, the sampler \mathbf{Samp} in the adapted counterexample outputs a tall $\mathbf{S} \in \mathbb{Z}_q^{m_P \times n}$ and a wide $\mathbf{P} \in \mathbb{Z}_q^{n \times m_P}$, both uniformly random. It also provides via \mathbf{aux} an obfuscation of a circuit $C_{\mathbf{SP} + \mathbf{E}''}$ with $\mathbf{SP} + \mathbf{E}''$ hardwired, where \mathbf{E}'' is short error sampled by \mathbf{Samp} itself. The circuit $C_{\mathbf{SP} + \mathbf{E}''}$ takes as inputs a tall matrix $\mathbf{M}_1 \in \mathbb{Z}_q^{m_P \times m}$ and a wide matrix $\mathbf{M}_2 \in \mathbb{Z}_q^{m_P \times m}$. If $\mathbf{M}_1 \mathbf{M}_2$ is close to $\mathbf{SP} + \mathbf{E}''$ it outputs 1, otherwise 0. In other words, $C_{\mathbf{SP} + \mathbf{E}''}$ checks if the low-rank matrix $\mathbf{M}_1 \mathbf{M}_2$ sufficiently approximates $\mathbf{SP} + \mathbf{E}''$. In the then challenge, the adversary uses the obfuscation of $C_{\mathbf{SP} + \mathbf{E}''}$ to distinguish $\mathbf{SB} + \mathbf{E}$

⁸ Note that the counterexample also applies to Eq. (5) which considers an even larger sampler class that also knows the trapdoor $\mathbf{td}_{\mathbf{B}}$.

from uniform randomness, since $C_{\mathbf{SP}+\mathbf{E}'}$ outputs 1 on input $(\mathbf{SB} + \mathbf{E}, \mathbf{U})$ with overwhelming probability. In the if challenge, the obfuscation of $C_{\mathbf{SP}+\mathbf{E}'}$ provably does not help in distinguishing $\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}'$ from uniform randomness, as the hardcoded $\mathbf{SP} + \mathbf{E}'$ in $C_{\mathbf{SP}+\mathbf{E}'}$ is indistinguishable from a uniformly random matrix $\mathbf{R} \in \mathbb{Z}_q^{m_P \times m_P}$. For $m_P \gg m$, the circuit $C_{\mathbf{R}}$ is zero everywhere. At this point, we can rely on the null-iO security to replace the obfuscation of $C_{\mathbf{R}}$ with one of a padded constant zero circuit which is clearly useless for solving the if challenge.

2.3 Plausible Classes of Assumptions

Recall again the intuition of evasive LWE from [Wee22]: To distinguish $\underline{\mathbf{s}}^T \underline{\mathbf{B}}$ given $\mathbf{B}^{-1}(\mathbf{P})$, the seemingly only meaningful way to use $\mathbf{B}^{-1}(\mathbf{P})$ is to right-multiply it to $\underline{\mathbf{s}}^T \underline{\mathbf{B}}$, obtaining new samples $\underline{\mathbf{s}}^T \underline{\mathbf{P}}$ and distinguishing with the latter. Underlying all counterexamples above are simple alternative uses of $\mathbf{B}^{-1}(\mathbf{P})$ crafted according to the corresponding assumption setting, which we summarise:

Mapping between \mathbf{B}, \mathbf{P} . Both counterexamples 1 and 2 target the setting where in the if distribution only one of the two matrices \mathbf{B} and \mathbf{P} are fully known to the distinguisher, but in the then distribution the additional information of $\mathbf{B}^{-1}(\mathbf{P})$ allows to recover the other hidden matrix via the relation

$$\mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P}) = \mathbf{P}, \quad (7)$$

the latter containing sufficient information for distinguishing LWE.

Encrypt w.r.t. short vector $\mathbf{B}^{-1}(\mathbf{P})$. For counterexample 3, $\mathbf{B}^{-1}(\mathbf{P})$ acts as the secret key of an encryption scheme. In fact, in this scenario, during decryption $\mathbf{B}^{-1}(\mathbf{P})$ is still being multiplied to some LWE samples $\underline{\mathbf{r}}^T \underline{\mathbf{B}}$ of \mathbf{B} to obtain new samples $\underline{\mathbf{r}}^T \underline{\mathbf{P}}$ w.r.t. \mathbf{P} . The crucial difference, however, is that such pair of LWE samples is prepared by \mathbf{Samp} but not the challenger, allowing its embedding of secret in the form of $\underline{\mathbf{r}}^T \underline{\mathbf{P}} + \text{secret}$.

With these in mind, we propose three main families of evasive LWE assumptions where these alternative uses cannot apply. These families are marked with green checkboxes in Fig. 1. See Section 4 for formal definitions.

Public-coin Evasive LWE. The first family is when \mathbf{Samp} is public-coin, meaning that it outputs all randomness used. This captures for example the assumptions in [Wee22, WWW22, HLL24, Wee24, CLW24].

We require that \mathbf{B} is contained in the joint distribution, see Remark 4 for a discussion. We also require that \mathbf{Samp} does not input \mathbf{B} , which is supported by a heuristic counterexample sketched in Appendix C.3.⁹ Note that under this family, \mathbf{P} is always available to the distinguisher (c.f. Fig. 1a), since it is efficiently recoverable from the randomness used by \mathbf{Samp} .

Private-coin Binding Evasive LWE. The second family is when \mathbf{Samp} is private-coin, and \mathbf{B}, \mathbf{P} are explicitly included in the joint distributions. Assumptions of this family take the following form: For any PPT \mathbf{Samp} inputting the security parameter λ and outputting $(\mathbf{S}, \mathbf{P}, \text{aux})$, where aux need not include all randomness used,

$$\begin{aligned} \text{if} \quad & (\mathbf{B}, \mathbf{P}, \underline{\mathbf{SB}}, \underline{\mathbf{SP}}, \text{aux}) && \approx_c && (\mathbf{B}, \mathbf{P}, \$, \$, \text{aux}) \\ \text{then} \quad & (\mathbf{B}, \mathbf{P}, \underline{\mathbf{SB}}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) && \approx_c && (\mathbf{B}, \mathbf{P}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}). \end{aligned}$$

This corresponds to the bottom-left green checkbox area of Fig. 1b and captures the variant Eq. (8) used by [ARYY23, AKY24], to be discussed in Section 2.4.

Private-coin Hiding Evasive LWE. The third and most subtle family is when \mathbf{Samp} is private-coin and \mathbf{B}, \mathbf{P} are hidden from the joint distribution. Our proposed family takes the following form: For any PPT \mathbf{Samp} inputting the security parameter λ and outputting $(\mathbf{S}, \mathbf{P}, \text{aux})$, where aux need not include all randomness used, and it holds $(\mathbf{P}, \text{aux}) \approx_c (\mathbf{P} + \mathbf{R}, \text{aux})$ for a bounded-norm random \mathbf{R} ,

$$\begin{aligned} \text{if} \quad & (\underline{\mathbf{SB}}, \underline{\mathbf{SP}}, \text{aux}) && \approx_c && (\$, \$, \text{aux}) \\ \text{then} \quad & (\underline{\mathbf{SB}}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) && \approx_c && (\$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}). \end{aligned}$$

⁹ We thank a reviewer of AsiaCrypt 2024 for pointing out this counterexample to us!

The bounded-norm \mathbf{R} can be interpreted as noise, and $\mathbf{P} + \mathbf{R}$ some approximation of \mathbf{P} . The indistinguishability between \mathbf{P} and $\mathbf{P} + \mathbf{R}$ serves as a proof of “ \mathbf{P} cannot be approximated given \mathbf{aux} ” and resists algebraic attacks which would, e.g., recover a column of \mathbf{P} , or a sum of two columns, since either allows to distinguish by cross-checking with \mathbf{aux} . See Section 4.3 for a more thorough discussion. This family corresponds to the top-right green checkbox area of Fig. 1b, and seeks to capture existing assumptions taking the form of Eq. (3) while plausibly resisting attacks exploiting the alternative use of $\mathbf{B}^{-1}(\mathbf{P})$ via Eq. (7): Without knowing neither \mathbf{B} nor approximation of \mathbf{P} , it is unclear how Eq. (7) may still be exploited.

Remark 4. Observant readers might have noticed that we have not included the top-left areas of Figs. 1a and 1b, the families where \mathbf{B} is hidden but \mathbf{P} is known. While we have not found counterexamples on these families, we believe they are of relatively low utility, based on the following informal argument: Since \mathbf{P} is fully known, the indistinguishability $\underline{\mathbf{S}}\mathbf{P} \approx_c \$$ implies \mathbf{S} having high entropy given the “if” distribution, in which case an appropriate entropic LWE assumption [BD20] implies $\underline{\mathbf{S}}\mathbf{B} \approx_c \$$ even when \mathbf{B} is known.¹⁰ We therefore expect that, instead of resorting to an assumption in this family, one can utilise the families where \mathbf{B} is known, yielding also a stronger *then* relation where \mathbf{B} is also known.

2.4 Assumption Instances in Existing Works

[ARYY23,AKY24]: Counterexample 1 applies to the evasive LWE variant in Eq. (4), which is the one of the two involved in [ARYY23,AKY24]. Fortunately, both works’ proofs are modular, in particular, ARYY show that Eq. (4) implies a second evasive LWE assumption, the latter used by both works to prove security of their constructions. The second assumption is as follows: For any PPT \mathbf{Samp} outputting an (arbitrary) LWE secret \mathbf{S} , matrix \mathbf{P} , and auxiliary information $\mathbf{aux}_1, \mathbf{aux}_2$ (not necessarily containing all randomness used by \mathbf{Samp}), where \mathbf{P} is efficiently computable given \mathbf{aux}_2 ,

$$\begin{aligned} \text{if} \quad & (\mathbf{B}, \underline{\mathbf{S}}\mathbf{B}, \underline{\mathbf{S}}\mathbf{P}, \mathbf{aux}_1, \mathbf{aux}_2) && \approx_c && (\mathbf{B}, \$, \$, \$, \mathbf{aux}_2) && (8) \\ \text{then} \quad & (\mathbf{B}, \underline{\mathbf{S}}\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \mathbf{aux}_1, \mathbf{aux}_2) && \approx_c && (\mathbf{B}, \$, \mathbf{B}^{-1}(\mathbf{P}), \$, \mathbf{aux}_2) \end{aligned}$$

for uniformly random \mathbf{B} . Since \mathbf{P} is efficiently computable from \mathbf{aux}_2 , one can also phrase Eq. (8) as an instance of our proposed private-coin binding evasive LWE, where \mathbf{B} and \mathbf{P} are both fully available in the joint distribution. Thus, assuming the private-coin binding evasive LWE, the MI-ABE of [ARYY23] and CP-ABE of [AKY24] remain secure.

[VWW22]: Counterexample 2 applies to the evasive LWE variant in Eq. (3), which first appears in [VWW22] and is recently involved in [CM24,MPV24]. Nevertheless, we have not found attacks on the assumption instance (i.e. the specific \mathbf{Samp}) used by [VWW22]. Indeed, in Section 6 we show that, the assumption instance used by [VWW22] falls under our proposed private-coin hiding evasive LWE, where \mathbf{B} is not given in the joint distribution and \mathbf{P} is hidden. The main reason is that each entry of \mathbf{P} contains Gaussian noise that is *independent* of \mathbf{aux} , and when the Gaussian parameter is sufficiently large, the noise term, consequently also \mathbf{P} , is statistically irrecoverable. Thus, assuming the private-coin hiding evasive LWE, the constructions of [VWW22] remain secure.

[Tsa22]: Counterexample 3 applies to the evasive-type assumption by [Tsa22], where a private-coin \mathbf{Samp} inputs \mathbf{B} and its outputs can be arbitrarily correlated with \mathbf{B} . Nevertheless, we observe that the condition of \mathbf{Samp} knowing \mathbf{B} seems to be an artifact of the definitional style used in [Tsa22]. In Tsabary’s evasive-type assumption, \mathbf{Samp} also outputs the pre-image $\mathbf{B}^{-1}(\mathbf{P})$ for the *then* distribution, which necessitates \mathbf{Samp} to input \mathbf{B} and its trapdoor. In contrast, in other evasive LWE assumptions, such $\mathbf{B}^{-1}(\mathbf{P})$ is generated by the challenger but not \mathbf{Samp} . Upon closer inspection, the \mathbf{Samp} instance in the security proof of [Tsa22] indeed *only* uses \mathbf{B} and $\mathbf{td}_{\mathbf{B}}$ for sampling the preimages $\mathbf{B}^{-1}(\mathbf{P})$ but nothing else. In particular \mathbf{aux} contains no further components correlated to \mathbf{B} and $\mathbf{td}_{\mathbf{B}}$. Therefore, it seems plausible to us that Tsabary’s construction can be proved under a variant of private-coin evasive LWE that is not subject to counterexamples. However, verifying this claim is not straightforward, since Tsabary’s assumption and

¹⁰ The argument is informal, because $\underline{\mathbf{S}}\mathbf{P} \approx_c \$$ might not necessarily imply sufficiently high entropy for entropic LWE to apply.

proof comes with many subtle differences. In more detail, Tsabary considers exponential-size distributions over matrices and LWE samples, and the distinguisher accesses them via querying an oracle with an index for the sample. In turn, evasive LWE provides polynomial-size samples directly to the distinguisher, so the adaptation of Tsabary's proof to evasive LWE requires syntactic changes and possibly re-structuring the proof into more game hops. We leave verifying the security of Tsabary's construction to future works.

3 Preliminaries

Denote by $\mathbb{N} = \{1, 2, \dots\}$ the set of natural numbers, by \mathbb{Z} the set of integers and by \mathbb{R} the set of real numbers. For $n \in \mathbb{N}$, we set $[n] = \{1, \dots, n\}$. For a ring extension $\mathcal{R} \supseteq \mathbb{Z}$, we set $\mathcal{R}_q := \mathcal{R}/(q \cdot \mathcal{R})$.

Denote by λ the security parameter. Write $\text{poly}(\lambda) = \bigcup_{d \in \mathbb{N}} O(\lambda^d)$, $\text{negl}(\lambda) = \bigcap_{d \in \mathbb{N}} o(\lambda^{-d})$. A **PPT** algorithm is a probabilistic algorithm whose time complexity lies in $\text{poly}(\lambda)$. For a PPT algorithm \mathcal{A} , we write $\mathcal{A}(\cdot; \text{rand})$ for running which on randomness rand , where rand is understood to be uniformly random over its randomness space. Denote by $\mathcal{U}(S)$ the uniform distribution over a finite set S . The **statistical distance** between two discrete distributions $\mathcal{S}_1, \mathcal{S}_2$ over a set X is $\Delta(\mathcal{S}_1, \mathcal{S}_2) = \frac{1}{2} \sum_{x \in X} |\mathcal{S}_1(x) - \mathcal{S}_2(x)|$. For a vector $\mathbf{x} = (x_1, \dots, x_m)^\top \in \mathbb{R}^m$, its ℓ_2 -norm is $\|\mathbf{x}\| := \|\mathbf{x}\|_2 = \sqrt{\sum_{i \in [m]} x_i^2}$.

3.1 Lattices and Gaussian Distributions

Definition 1 (Lattice and Dual Lattice). A *lattice* is a discrete additive subgroup $\Lambda \subset \mathbb{R}^m$. The *rank* of a lattice Λ is defined to be the vector space dimension by the space spanned by the elements of Λ . The *dual lattice* of a lattice Λ is given by

$$\Lambda^* = \{\mathbf{y} \in \mathbb{R}^m \mid \forall \mathbf{x} \in \Lambda : \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}\}.$$

Definition 2 (q -ary Lattices and Cosets). For a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, define its orthogonal lattice by

$$\Lambda^\perp(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{e} = \mathbf{0} \pmod{q}\} \subseteq \mathbb{Z}^m$$

and the lattice spanned by its rows by

$$\Lambda(\mathbf{A}) := \{\mathbf{y} \in \mathbb{Z}^m \mid \exists \mathbf{x} \in \mathbb{Z}^n, \mathbf{y} = \mathbf{x}^\top \mathbf{A} \pmod{q}\} \subseteq \mathbb{Z}^m.$$

Additionally, for a syndrome $\mathbf{u} \in \mathbb{Z}_q^n$, define the following coset of $\Lambda^\perp(\mathbf{A})$

$$\Lambda_{\mathbf{u}}^\perp(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A}\mathbf{e} = \mathbf{u} \pmod{q}\} \subseteq \mathbb{Z}^m.$$

For $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_k) \in \mathbb{Z}_q^{n \times k}$, the notation extends naturally to $\Lambda_{\mathbf{U}}^\perp(\mathbf{A}) = \Lambda_{\mathbf{u}_1}^\perp(\mathbf{A}) \times \dots \times \Lambda_{\mathbf{u}_k}^\perp(\mathbf{A}) \subseteq \mathbb{Z}^{m \times k}$.

Note that we have $q \cdot \Lambda(\mathbf{A}) = (\Lambda^\perp(\mathbf{A}))^*$ and $\Lambda_{\mathbf{u}}^\perp(\mathbf{A}) = \Lambda^\perp(\mathbf{A}) + \mathbf{t}$, if there exists a $\mathbf{t} \in \mathbb{Z}^m$ s.t. $\mathbf{A}\mathbf{t} = \mathbf{u} \pmod{q}$.

Definition 3 (Gaussian measure). For $\mathbf{x} \in \mathbb{R}^m$, the *Gaussian measure* with parameter $\Sigma \in \mathbb{R}^{m \times m}$ and center $\mathbf{c} \in \mathbb{R}^m$, where Σ is positive semi-definite, is $\rho_{\Sigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \cdot (\mathbf{x} - \mathbf{c})^\top \Sigma^{-1} (\mathbf{x} - \mathbf{c}))$. If $\Sigma = \chi^2 \mathbf{I}$, we write $\rho_{\chi, \mathbf{c}}(\mathbf{x})$. If $\mathbf{c} = \mathbf{0}$, we simply write $\rho_\Sigma(\mathbf{x})$ (resp. $\rho_\chi(\mathbf{x})$). For a discrete set $\Lambda \subset \mathbb{R}^m$, we set $\rho_{\Sigma, \mathbf{c}}(\Lambda) := \sum_{\mathbf{x} \in \Lambda} \rho_{\Sigma, \mathbf{c}}(\mathbf{x})$.

Definition 4 (Discrete Gaussian Distribution). For a non-empty discrete set $\Lambda \subset \mathbb{R}^m$, define the *discrete Gaussian distribution* over Λ with parameter $\Sigma \in \mathbb{R}^{m \times m}$ and center $\mathbf{c} \in \mathbb{R}^m$, where Σ is positive definite, as

$$D_{\Lambda, \Sigma, \mathbf{c}}(\mathbf{x}) = \frac{\rho_{\Sigma, \mathbf{c}}(\mathbf{x})}{\rho_{\Sigma, \mathbf{c}}(\Lambda)} \quad \text{if } \mathbf{x} \in \Lambda,$$

and $D_{\Lambda, \Sigma, \mathbf{c}}(\mathbf{x}) = 0$ otherwise. If $\Sigma = \chi^2 \mathbf{I}$, we write $D_{\Lambda, \chi, \mathbf{c}}$ for $D_{\Lambda, \Sigma, \mathbf{c}}$. If $\mathbf{c} = \mathbf{0}$, we simply write $D_{\Lambda, \Sigma}$ (resp. $D_{\Lambda, \chi}$).

For $\Lambda_{\mathbf{U}}^\perp(\mathbf{A}) = \Lambda_{\mathbf{u}_1}^\perp(\mathbf{A}) \times \dots \times \Lambda_{\mathbf{u}_k}^\perp(\mathbf{A})$ where each $\Lambda_{\mathbf{u}_i}^\perp(\mathbf{A})$ is non-empty, we write $D_{\Lambda_{\mathbf{U}}^\perp(\mathbf{A}), \Sigma, \mathbf{c}}$ for the (horizontal) concatenation of the discrete Gaussian distributions over each $\Lambda_{\mathbf{u}_i}^\perp(\mathbf{A})$.

Definition 5 (Smoothing Parameter [MR04]). For a full-rank lattice $\Lambda \subset \mathbb{R}^m$ and $\epsilon > 0$, its smoothing parameter is

$$\eta_\epsilon(\Lambda) := \inf\{\chi > 0 \mid \rho_{1/\chi}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon\}.$$

The following lemma allows to bound the smoothing parameter for orthogonal lattices of random matrices:

Lemma 1 ([Pei07,GPV08]). Let q be prime, $m \geq 2n \cdot \log q$. We have

$$\Pr_{\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}} \left[\eta_{2^{-n}}(\Lambda^\perp(\mathbf{A})) \leq \frac{4}{\sqrt{\pi}} \cdot \sqrt{\log(2m) + \log(1 + 2^n)} \right] \geq 1 - q^{-n}.$$

Lemma 1 is an implication of results of [Pei07,GPV08]. For details, we refer the reader to Lemmas 13 and 15 in Appendix A.1.

Lemma 2 ([MR04,PR06]). Let $\Lambda \subset \mathbb{R}^m$ be an m -dimensional full-rank lattice. For any $\mathbf{c} \in \mathbb{R}^m$, $\epsilon > 0$, $\chi \geq 2\eta_\epsilon(\Lambda)$, and $\mathbf{y} \in \Lambda$, it holds that

$$D_{\Lambda, \chi, \mathbf{c}}(\mathbf{y}) \leq 2^{-m} \cdot \frac{1 + \epsilon}{1 - \epsilon}.$$

Lemma 3 (Leftover Hash Lemma [HILL99,BDK⁺11]). Let $m, n, q \in \mathbb{N}$. If we draw $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{x} \leftarrow \{0, 1\}^m$, we have

$$\Delta((\mathbf{A}, \mathbf{A}\mathbf{x} \bmod q), (\mathbf{A}, \mathbf{y})) \leq \frac{1}{2} \cdot \frac{q^{n/2}}{2^{m/2}}.$$

Lemma 4 (Noise Flooding). Let $\chi = D_{\mathbb{Z}, \sigma}$. For all $t \in \mathbb{Z}$, we have $\Delta(\chi, \chi + t) \leq \sqrt{\frac{\pi}{2}} \cdot \frac{\|t\|}{\sigma}$. In particular, if $\chi \in \lambda^{\omega(1)}\|t\|$, then $\Delta(\chi, \chi + t) \in \text{negl}(\lambda)$.

For a proof of Lemma 4, see [BDE⁺18, Appendix A.2].

3.2 Lattice Assumptions and Lattice Trapdoors

Definition 6 (Learning with Errors). Let q, n, m, χ be parametrised by λ , where $n, m, q \in \mathbb{N}$ with $n, m \in \text{poly}(\lambda)$. The (decisional) **Learning with Errors** $\text{LWE}_{q, n, m, \chi}$ assumption states that for every PPT distinguisher \mathcal{D} , it holds that

$$\left| \Pr \left[b = 0 \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi} \\ \mathbf{b}^\top := \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \bmod q \\ b \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{b}) \end{array} \right] - \Pr \left[b = 0 \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{b} \leftarrow \mathbb{Z}_q^m \\ b \leftarrow \mathcal{D}(\mathbf{A}, \mathbf{b}) \end{array} \right] \right| \in \text{negl}(\lambda).$$

In their seminal work, [MP12] gave algorithms for sampling (almost) uniformly random matrices together with trapdoors that allow for LWE inversion and preimage lattice sampling. Currently, their algorithms are the status quo for these tasks, and we will usually assume that the challengers for the evasive LWE assumptions will resort to them. Hence, we will give here an overview of the guarantees given in [MP12] for these algorithms:

Theorem 1 ([MP12]). Let $m \geq 3n \cdot \log(q)$ and set $w = n \cdot \lceil \log q \rceil$. Let $\chi \in \Omega(n \cdot \sqrt{\log q})$. There exist algorithms ($\text{TrapGen}, \text{SampPre}$) with the following properties:

1. $\text{TrapGen}(1^n, 1^m)$ outputs two matrices $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, $\text{td} \in \{-1, 0, 1\}^{w \times (m-w)}$ s.t. the statistical distance between \mathbf{B} and $\mathcal{U}(\mathbb{Z}_q^{n \times m})$ is upper-bounded by 2^{-n} .
2. Draw $(\mathbf{B}, \text{td}) \leftarrow \text{TrapGen}(1^n, 1^m)$, let $\mathbf{u} \in \mathbb{Z}_q^n$. For $\mathbf{e}' \leftarrow \text{SampPre}(\mathbf{B}, \text{td}, \mathbf{u}, \chi)$ and $\mathbf{e} \leftarrow D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{B}), \chi}$, we have $\Delta((\mathbf{B}, \mathbf{e}), (\mathbf{B}, \mathbf{e}')) \in O(2^{-n})$.

$\text{Pre}_A^b(1^\lambda)$	$\text{Post}_B^b(1^\lambda)$
$\mathbf{B} \leftarrow \mathcal{D}$	$\mathbf{B} \leftarrow \mathcal{D}$
$(\mathbf{A}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda; \text{rand})$	$(\mathbf{A}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda; \text{rand})$
assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$	assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$
$(\mathbf{S}, \mathbf{S}_A) \leftarrow \mathcal{S}$	$(\mathbf{S}, \mathbf{S}_A) \leftarrow \mathcal{S}$
$\text{hint} := (f(\mathbf{S}; \text{rand}_f), f_A(\mathbf{S}_A; \text{rand}_{f_A}))$	$\text{hint} := (f(\mathbf{S}; \text{rand}_f), f_A(\mathbf{S}_A; \text{rand}_{f_A}))$
if $b = 0$ then	if $b = 0$ then
$\mathbf{E}_A \leftarrow \mathcal{X}_A, \mathbf{E}_B \leftarrow \mathcal{X}_B, \mathbf{E}_P \leftarrow \mathcal{X}_P$	$\mathbf{E}_A \leftarrow \mathcal{X}_A, \mathbf{E}_B \leftarrow \mathcal{X}_B$
$\mathbf{C}_A := \mathbf{S}_A \mathbf{A} + \mathbf{E}_A \text{ mod } q$	$\mathbf{C}_A := \mathbf{S}_A \mathbf{A} + \mathbf{E}_A \text{ mod } q$
$\mathbf{C}_B := \mathbf{S} \mathbf{B} + \mathbf{E}_B \text{ mod } q$	$\mathbf{C}_B := \mathbf{S} \mathbf{B} + \mathbf{E}_B \text{ mod } q$
$\mathbf{C}_P := \mathbf{S} \mathbf{P} + \mathbf{E}_P \text{ mod } q$	$\mathbf{U} \leftarrow D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$
if $b = 1$ then	if $b = 1$ then
$\mathbf{C}_A \leftarrow \mathcal{R}_q^{t_A \times m_A}$	$\mathbf{C}_A \leftarrow \mathcal{R}_q^{t_A \times m_A}$
$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t \times m}$	$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t \times m}$
$\mathbf{C}_P \leftarrow \mathcal{R}_q^{t \times m_P}$	$\mathbf{U} \leftarrow D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$
return $\mathcal{A} \left(\begin{array}{l} \mathbf{A}, \mathbf{B}, \mathbf{P}, \text{hint}, \text{aux}, \\ \mathbf{C}_A, \mathbf{C}_B, \mathbf{C}_P, \\ \text{rand}, \text{rand}_{f_A}, \text{rand}_f \end{array} \right)$	return $\mathcal{B} \left(\begin{array}{l} \mathbf{A}, \mathbf{B}, \mathbf{P}, \text{hint}, \text{aux}, \\ \mathbf{C}_A, \mathbf{C}_B, \mathbf{U}, \\ \text{rand}, \text{rand}_{f_A}, \text{rand}_f \end{array} \right)$

Fig. 2: Experiments Pre and Post for public-coin evasive LWE.

4 Evasive LWE: Definitions, Classes

We formally define the three classes of evasive LWEs outlined in Section 2.3.

4.1 Public-coin Evasive LWE

Definition 7 (Public-coin Evasive LWE). *Let the parameters*

$$\text{param} = (\mathcal{R}, q, n, n_A, m, m_P, m_A, t, t_A, \mathcal{D}, \mathcal{S}, \mathcal{X}_B, \mathcal{X}_P, \mathcal{X}_A, f, f_A, \Sigma)$$

be parametrised by λ , where \mathcal{R} is a ring admitting an embedding as a lattice in \mathbb{R}^φ for some $\varphi \in \mathbb{N}$, $\mathcal{D} \sim \mathcal{R}_q^{n \times m}$, $\mathcal{S} \sim \mathcal{R}_q^{t \times n} \times \mathcal{R}_q^{t_A \times n_A}$, $\mathcal{X}_B \sim \mathcal{R}^{t \times m}$, $\mathcal{X}_P \sim \mathcal{R}^{t \times m_P}$, and $\mathcal{X}_A \sim \mathcal{R}^{t_A \times m_A}$ are distributions, $\Sigma \in \mathbb{R}^{\varphi m \times \varphi m}$ is positive definite, and f, f_A are PPT algorithms. Let Samp be a PPT algorithm which, on input 1^λ , outputs

$$(\mathbf{A} \in \mathcal{R}_q^{n_A \times m_A}, \mathbf{P} \in \mathcal{R}_q^{n \times m_P}, \text{aux} \in \{0, 1\}^*).$$

Denote

$$\begin{aligned} \text{Adv}_A^{\text{Pre}}(\lambda) &:= |\Pr[\text{Pre}_A^0(1^\lambda) = 1] - \Pr[\text{Pre}_A^1(1^\lambda) = 1]|, \\ \text{Adv}_B^{\text{Post}}(\lambda) &:= |\Pr[\text{Post}_B^0(1^\lambda) = 1] - \Pr[\text{Post}_B^1(1^\lambda) = 1]|, \end{aligned}$$

where the experiments Pre_A^b and Post_B^b are defined in Fig. 2. The $\text{PublicEvLWE}_{\text{param}}$ assumption states that for any PPT Samp and \mathcal{B} there exists a PPT \mathcal{A} such that $\text{Adv}_A^{\text{Pre}}(\lambda) \geq \text{Adv}_B^{\text{Post}}(\lambda) / \text{poly}(\lambda) - \text{negl}(\lambda)$.

Remark 5. We parametrise the assumption by the modulus, matrix dimensions, noise parameters, etc.. This is analogous to how the LWE assumption is defined when done precisely, which formally is parametrised by $(\mathcal{R}, q, n, m, \chi)$. We emphasise that we believe the plausibility of an evasive LWE assumption should depend on these parameters, analogous to that the plausibility of LWE depends on e.g. the dimension n and the ratio q/χ . As we will see, existing public-coin evasive LWE all fall under Definition 7 with specially chosen parameters.

Definition 7 is a versatile definition designed to capture both the public-coin flavour and many different (sometimes implicit) features involved in existing definitions. We elaborate on some key aspects.

Randomness of Samp. The randomness `rand` used by `Samp` is made explicit syntactically, to highlight the public-coin nature. The same convention was used in [HLL24]. Given `rand`, the inputs $(\mathbf{A}, \mathbf{P}, \text{aux})$ to \mathcal{A} and \mathcal{B} may be omitted as they can be derived from `rand`; we include them for clarity.

Matrix \mathbf{A} . The matrix \mathbf{A} serves to represent the LWE matrix not involved in the preimage-image relation which the evasive LWE assumption concerns, i.e. in the `Post` experiment, the acquired preimages do not involve \mathbf{A} . This serves to increase expressiveness of the assumption while leaving the intuition of evasive LWE unchanged, and is required in a number of works [Wee22, WWW22, HLL24, CLW24]. By setting \mathbf{A} the empty matrix, we recover the simpler cases outlined in Section 2.

Check $\mathbf{P} = \mathbf{B}\mathcal{R}^{m_P}$. This check makes the evasive LWE assumption formally well-defined: Without this check, in case $\mathbf{P} = (\mathbf{p}_i)_i$ is not in the \mathcal{R} -span of \mathbf{B} , the distribution $D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$, and consequently `Post`^{*b*}, would be ill-defined, since some $A_{\mathbf{p}_i}^\perp(\mathbf{B})$ would be empty. In existing works, this check is implicit since in those settings w.h.p. \mathbf{B} is primitive and all $A_{\mathbf{p}_i}^\perp(\mathbf{B}) \neq \emptyset$.

Distribution of \mathbf{B} . We let \mathbf{B} be sampled from a distribution \mathcal{D} , the latter being itself a parameter of the assumption. In the simple setting of [Wee22], \mathcal{D} is the uniform distribution over $\mathbb{Z}_q^{n \times m}$. As discussed below, by setting \mathcal{D} appropriately, other existing variants, e.g. where an evasive LWE is defined over multiple \mathbf{B}_i 's, can also be naturally captured. To understand and gain confidence in this generalisation, we discuss some special cases of \mathcal{D} . (1) Low-entropy \mathcal{D} (in the extreme case \mathbf{B} is deterministic) and $n < m$: In this case LWE w.r.t. \mathbf{B} is likely easy, so that both `Pre`^{*b*} and `Post`^{*b*} can be efficiently distinguished, and the assumption is vacuously true. (2) \mathcal{D} is such that \mathbf{B} is (likely) not primitive: If \mathbf{P} is in the column span of \mathbf{B} then the rationale of the assumption stays unchanged; otherwise, the check $\mathbf{P} = \mathbf{B}\mathcal{R}^{m_P}$ fails, so that the winning probability in both `Pre`^{*b*}, `Post`^{*b*} are zero and the assumption is again true. (3) \mathcal{D} is such that the lattice $A^\perp(\mathbf{B})$ (likely) contains no short vector: The assumption, in particular the distribution $D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$, is still well-defined, although with the unusual scenario that \mathbf{U} is (likely) not short, which intuitively only makes distinguishing `Post`^{*b*} harder.

Note on December 10, 2024. The generalisation to arbitrary distributions of \mathbf{B} introduces more weaknesses than we initially realised. See Section 8 for a simple counterexample and related discussions. As such, at the moment, we advise to restrict to uniform \mathbf{B} only.

LWE Secret Distributions. The LWE secret distribution \mathcal{S} , determining (correlation between) the LWE secret for samples w.r.t. to (\mathbf{B}, \mathbf{P}) and \mathbf{A} respectively, is parametrised by the assumption (instead of e.g. fixed to be uniform). This is in line with the intuition of evasive LWE (and already implicit in existing private-coin variants), which says that secret distributions should not matter to the if-then relation that the assumption postulates, since intuitively $\mathbf{B}^{-1}(\mathbf{P})$ should not be able to interact with \mathbf{S} and \mathbf{S}_A meaningfully. Note that a poor secret distribution would allow to distinguish `Pre`^{*b*} in the first place so that the assumption is vacuously true. As to be discussed below, this treatment further allows us to interpret some existing variants in an arguably more intuitive way.

Public-coin Hint. The hint on the LWE secrets, which are outputs of the functions f, f_A parametrised by the assumption, manifests the intuition of “secret distributions should not matter to the if-then relation”, in that leakages on the secret may be given to the distinguishers. We view this as an evasive analogue of entropic LWE [BD20]. (In private-coin variants, such hint is implicit in `aux` since `Samp` can generate it itself.) To be consistent with the spirit of “public-coin”, we require the randomness `rand`_{*f_A*}, `rand`_{*f*} used in generating hint (if any) to be provided to the distinguisher. We note again that a poor leakage would allow a distinguisher to distinguish `Pre`^{*b*} so that the assumption is vacuously true.

\mathbf{B} not known to Samp. Similar to the private-coin setting to be discussed up next, in Definition 7 we forbid the sampler `Samp` to input \mathbf{B} . Although we are unable to provide a provable counterexample against this case, in Appendix C we provide a heuristic counterexample to support this restriction.

Relating to existing definitions. All existing public-coin evasive LWE can be viewed as special cases of Definition 7, which we briefly summarise. Most works define an assumption over the integers $\mathcal{R} = \mathbb{Z}$.¹¹

¹¹ On the other hand, suppose one is to optimise any involved constructions for efficiency, then it is easy to see that an analogous evasive LWE instance over other rings \mathcal{R} , e.g. the ring of integers of some cyclotomic field, is to be involved.

The original evasive LWE of [Wee22], later also used in [Wee24], has secret distribution \mathcal{S} such that $\mathbf{S}_A = \mathbf{S} = \mathbf{s}^\top$ where \mathbf{s} is uniform, and f_A, f the empty function, i.e. no hint involved. [WWW22] defined a public-coin evasive LWE with multiple independent and uniform $\mathbf{B}_i \in \mathbb{Z}_q^{n \times m}$, with LWE samples $\mathbf{s}_i^\top \mathbf{B}_i + \mathbf{e}_i^\top \bmod q$ under independent uniform secrets \mathbf{s}_i , and additionally with samples w.r.t. \mathbf{A} under the

concatenated secret $(\mathbf{s}_1^\top, \dots, \mathbf{s}_k^\top)$. This can be expressed by our definition as having $\mathbf{B} = \begin{pmatrix} \mathbf{B}_1 & & \\ & \ddots & \\ & & \mathbf{B}_k \end{pmatrix}$

and secret distribution \mathcal{S} such that $\mathbf{S} = \mathbf{S}_A = (\mathbf{s}_1^\top, \dots, \mathbf{s}_k^\top)$. [CLW24] used a public-coin evasive LWE that is similar to that of [WWW22], but let \mathcal{S} be such that $\mathbf{S} = (\mathbf{s}_1^\top, \mathbf{s}^\top, \dots, \mathbf{s}_k^\top, \mathbf{s}^\top)$, $\mathbf{S}_A = (\mathbf{s}_1^\top, \dots, \mathbf{s}_k^\top, \mathbf{s}^\top)$, where \mathbf{s}_i, \mathbf{s} are all uniform. [HLL24] defined a public-coin evasive LWE with matrices and LWE samples of the (arguably ad-hoc) forms

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}_0 \\ \vdots \\ \mathbf{B}_k \end{pmatrix}, \quad \mathbf{s}_i^\top (\mathbf{B}_0, \dots, \mathbf{B}_k) + (\mathbf{0}^\top, \mathbf{g}^\top \otimes \mathbf{e}_{i,1}^\top) + \mathbf{e}_{i,2}^\top \bmod q \quad \text{for all } i \in [I]$$

(with preimages $\mathbf{U} = \mathbf{B}^{-1}(\mathbf{P})$ w.r.t. \mathbf{B}), where $\mathbf{g}^\top = (2^0, \dots, 2^{k-1})$, and with further samples of the form $(\mathbf{s}_i^\top, \mathbf{s}_0^\top) \mathbf{A} + \mathbf{e}_{i,0}^\top \bmod q$ for all $i \in [I]$; This can be summarised by our definition as letting $\mathbf{B} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{nk \times m}$ be simply uniformly random, and with structured secret and error distributions $\mathcal{S}, \chi_B, \chi_A$ such that

$$\mathbf{S} = \begin{pmatrix} \hat{\mathbf{S}} & & \\ & \ddots & \\ & & \hat{\mathbf{S}} \end{pmatrix} \quad \text{where } \hat{\mathbf{S}} = \begin{pmatrix} \mathbf{s}_1^\top \\ \vdots \\ \mathbf{s}_I^\top \end{pmatrix}, \quad \mathbf{S}_A = \begin{pmatrix} \mathbf{s}_1^\top, \mathbf{s}_0^\top \\ \vdots \\ \mathbf{s}_I^\top, \mathbf{s}_0^\top \end{pmatrix},$$

$$\mathbf{E}_B = \begin{pmatrix} (\mathbf{0}^\top, \mathbf{g}^\top \otimes \mathbf{e}_{1,1}^\top) + \mathbf{e}_{1,2}^\top \\ \vdots \\ (\mathbf{0}^\top, \mathbf{g}^\top \otimes \mathbf{e}_{I,1}^\top) + \mathbf{e}_{I,2}^\top \end{pmatrix}, \quad \mathbf{E}_A = \begin{pmatrix} \mathbf{e}_{1,0}^\top \\ \vdots \\ \mathbf{e}_{I,0}^\top \end{pmatrix}.$$

Remark 6 (Relation to evasive circular LWE of [HLL23]). In [HLL23] an “evasive circular LWE assumption” is proposed, which can be viewed as involving a non-trivial hint function. However, despite being regarded as a public-coin assumption by [HLL23], this does not fall into the public-coin family under our characterisation (also when factoring out its circular nature), due to its hint function being not public-coin. Further discussion on this in Appendix D.

4.2 Private-coin Binding Evasive LWE

Definition 8 (Private-coin Binding Evasive LWE). *Let the parameters*

$$\text{param} = (\mathcal{R}, q, n, m, m_P, t, \mathcal{D}, \chi_B, \chi_P, \Sigma)$$

be parametrised by λ , where \mathcal{R} is a ring admitting an embedding as a lattice in \mathbb{R}^φ for some $\varphi \in \mathbb{N}$, $\mathcal{D} \sim \mathcal{R}_q^{n \times m}$, $\chi_B \sim \mathcal{R}^{t \times m}$, and $\chi_P \sim \mathcal{R}^{t \times m_P}$ are distributions, and $\Sigma \in \mathbb{R}^{\varphi m \times \varphi m}$ is positive definite. Let Samp be a PPT algorithm which, on input 1^λ , outputs

$$(\mathbf{S} \in \mathcal{R}_q^{t \times n}, \mathbf{P} \in \mathcal{R}_q^{n \times m_P}, \text{aux} \in \{0, 1\}^*).$$

Let Pre_A^b and Post_B^b be the experiments defined in Fig. 3 and denote

$$\text{Adv}_A^{\text{Pre}}(\lambda) := |\Pr[\text{Pre}_A^0(1^\lambda) = 1] - \Pr[\text{Pre}_A^1(1^\lambda) = 1]|,$$

$$\text{Adv}_B^{\text{Post}}(\lambda) := |\Pr[\text{Post}_B^0(1^\lambda) = 1] - \Pr[\text{Post}_B^1(1^\lambda) = 1]|.$$

The $\text{PrivateBindEvLWE}_{\text{param}}$ assumption states that for any PPT Samp and \mathcal{B} there exists a PPT \mathcal{A} such that $\text{Adv}_A^{\text{Pre}}(\lambda) \geq \text{Adv}_B^{\text{Post}}(\lambda) / \text{poly}(\lambda) - \text{negl}(\lambda)$.

We explain some key aspects of Definition 8.

$\text{Pre}_A^b(1^\lambda)$	$\text{Post}_B^b(1^\lambda)$
$\mathbf{B} \leftarrow \mathcal{D}$	$\mathbf{B} \leftarrow \mathcal{D}$
$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$
assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$	assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$
if $b = 0$ then	if $b = 0$ then
$\mathbf{E}_B \leftarrow \chi_B, \mathbf{E}_P \leftarrow \chi_P$	$\mathbf{E}_B \leftarrow \chi_B$
$\mathbf{C}_B := \mathbf{S}\mathbf{B} + \mathbf{E}_B \bmod q$	$\mathbf{C}_B := \mathbf{S}\mathbf{B} + \mathbf{E}_B \bmod q$
$\mathbf{C}_P := \mathbf{S}\mathbf{P} + \mathbf{E}_P \bmod q$	$\mathbf{U} \leftarrow D_{A_{\mathbf{P}}(\mathbf{B}), \Sigma}$
if $b = 1$ then	if $b = 1$ then
$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t \times m}$	$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t \times m}$
$\mathbf{C}_P \leftarrow \mathcal{R}_q^{t \times m_P}$	$\mathbf{U} \leftarrow D_{A_{\mathbf{P}}(\mathbf{B}), \Sigma}$
return $\mathcal{A} \left(\begin{array}{c} \mathbf{B}, \mathbf{P}, \text{aux} \\ \mathbf{C}_B, \mathbf{C}_P \end{array} \right)$	return $\mathcal{B} \left(\begin{array}{c} \mathbf{B}, \mathbf{P}, \text{aux} \\ \mathbf{C}_B, \mathbf{U} \end{array} \right)$

Fig. 3: Experiments Pre and Post for private-coin binding evasive LWE.

Randomness of Samp. As its name suggests, in Definition 8 the randomness of **Samp** need not be given to the distinguishers. This makes the assumption more susceptible to future attacks, as **Samp** can embed secret information in **aux**. For example, as discussed in [VWW22], one may potentially include in **aux** a carefully crafted obfuscation containing a trapdoor of **P** (cf. Remark 3).

B not known to Samp. The assumption is restricted to the class of **Samp** which does not input the matrix **B**. This avoids counterexamples such as ours, sketched in Section 2.2 and detailed in Section 5.3.

Correlations among S, P and aux. **Samp** outputs also the LWE secret **S**, implying that **S** can be correlated to **P** and **aux** secretly and arbitrarily. This leads to another potential attack angle of exploiting such correlations.

Other outputs of Samp. Relative to Definition 7, other components such as **A**, \mathbf{S}_A and hint functions are omitted, since these can now be generated by **Samp** and contained in **aux**.

Relating to existing definitions. The private-coin variant of [ARYY23, Lemma 3.4] falls under Definition 8, where they formulated **aux** as $(\text{aux}_1, \text{aux}_2)$, where aux_1 can be proven pseudorandom and **P** is efficiently computable from aux_2 . As reference, the security proofs of [ARYY23] involve calling private-coin evasive LWE iteratively, some instances due to that **aux** contains preimages w.r.t. components of **P** sampled using a trapdoor which cannot be leaked to the distinguisher, and some others due to that **aux** involves secret correlation with **S**; all instances do not involve correlation between **S** and **P**. The same variant is later used by [AKY24]. Moreover, the evasive circular LWE assumption of [HLL23] is also a member of Definition 8, for which we discuss in Appendix D.

4.3 Private-coin Hiding Evasive LWE

Definition 9 (Private-coin Hiding Evasive LWE). *Let the parameters*

$$\text{param} = (\mathcal{R}, q, n, m, m_P, t, \mathcal{D}, \chi_B, \chi_P, \ell, \Sigma)$$

*be parametrised by λ , where \mathcal{R} is a ring admitting an embedding as a lattice in \mathbb{R}^φ for some $\varphi \in \mathbb{N}$, $\mathcal{D} \sim \mathcal{R}_q^{n \times m}$, $\chi_B \sim \mathcal{R}^{t \times m}$, and $\chi_P \sim \mathcal{R}^{t \times m_P}$ are distributions, $\Sigma \in \mathbb{R}^{\varphi m \times \varphi m}$ is positive definite, and $\ell \in \{1, 2, \dots, q\}$. Let **Samp** be a PPT algorithm which, on input 1^λ , outputs*

$$(\mathbf{S} \in \mathcal{R}_q^{t \times n}, \mathbf{P} \in \mathcal{R}_q^{n \times m_P}, \text{aux} \in \{0, 1\}^*).$$

Let Pre_A^b , Pre_A^b and Post_B^b be the experiments defined in Fig. 4 and denote

$$\begin{aligned} \text{Adv}_A^{\text{Pre}^1}(\lambda) &:= |\Pr[\text{Pre}_A^0(1^\lambda) = 1] - \Pr[\text{Pre}_A^1(1^\lambda) = 1]|, \\ \text{Adv}_A^{\text{Pre}^2}(\lambda) &:= |\Pr[\text{Pre}_A^0(1^\lambda) = 1] - \Pr[\text{Pre}_A^1(1^\lambda) = 1]|, \end{aligned}$$

Pre1 $_A^b(1^\lambda)$	Post $_B^b(1^\lambda)$
$\mathbf{B} \leftarrow \mathcal{D}$	$\mathbf{B} \leftarrow \mathcal{D}$
$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$
assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$	assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$
if $b = 0$ then	if $b = 0$ then
$\mathbf{E}_B \leftarrow \mathcal{X}_B, \mathbf{E}_P \leftarrow \mathcal{X}_P$	$\mathbf{E}_B \leftarrow \mathcal{X}_B$
$\mathbf{C}_B := \mathbf{SB} + \mathbf{E}_B \bmod q$	$\mathbf{C}_B := \mathbf{SB} + \mathbf{E}_B \bmod q$
$\mathbf{C}_P := \mathbf{SP} + \mathbf{E}_P \bmod q$	$\mathbf{U} \leftarrow D_{\mathcal{A}_B^\perp(\mathbf{B}), \Sigma}$
if $b = 1$ then	if $b = 1$ then
$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t \times m}$	$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t \times m}$
$\mathbf{C}_P \leftarrow \mathcal{R}_q^{t \times m_P}$	$\mathbf{U} \leftarrow D_{\mathcal{A}_B^\perp(\mathbf{B}), \Sigma}$
return $\mathcal{A}(\mathbf{C}_B, \mathbf{C}_P, \text{aux})$	return $\mathcal{B}(\mathbf{C}_B, \mathbf{U}, \text{aux})$
<hr/>	
Pre2 $_A^b(1^\lambda)$	
$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	
if $b = 0$ then	
return $\mathcal{A}(\mathbf{P}, \text{aux})$	
if $b = 1$ then	
$\mathbf{R} \leftarrow \mathcal{U}(\{0, 1, \dots, \ell\}^{n \times m_P})$	
return $\mathcal{A}(\mathbf{P} + \mathbf{R} \bmod q, \text{aux})$	

Fig. 4: Experiments Pre1, Pre2 and Post for private-coin hiding evasive LWE.

$$\text{Adv}_B^{\text{Post}}(\lambda) := |\Pr[\text{Post}_B^0(1^\lambda) = 1] - \Pr[\text{Post}_B^1(1^\lambda) = 1]|.$$

The PrivateHideEvLWE $_{\text{param}}$ assumption states that for any PPT Samp and \mathcal{B} there exists a PPT \mathcal{A} such that $\text{Adv}_A^{\text{Pre1}}(\lambda) + \text{Adv}_A^{\text{Pre2}}(\lambda) \geq \text{Adv}_B^{\text{Post}}(\lambda)/\text{poly}(\lambda) - \text{negl}(\lambda)$.

The experiments Pre1, Post in Definition 9 are almost identical to the experiments Pre, Post in Definition 8, except that the matrix \mathbf{B} is not given to the distinguishers \mathcal{A} and \mathcal{B} , and \mathbf{P} also not necessarily. The obvious distinction is the additional experiment Pre2, which we define below.

Experiment Pre2. This experiment seeks to ensure that “both \mathbf{B}, \mathbf{P} are sufficiently hidden from the distinguishers”. First, observe that the indistinguishability of Pre1^0 and Pre1^1 , i.e., $(\underline{\mathbf{SB}}, \underline{\mathbf{SP}}, \text{aux}) \approx_c (\mathcal{S}, \mathcal{S}, \text{aux})$, implies that $(\underline{\mathbf{SB}}, \underline{\mathbf{SP}}, \text{aux})$ does not leak information about \mathbf{B} (computationally), since aux is independent of \mathbf{B} . Moreover, the only possible way for a PPT adversary \mathcal{A} to obtain information about \mathbf{P} is via aux. Experiment Pre2^b then ensures that the latter is also impossible, in that (\mathbf{P}, aux) and $(\mathbf{P} + \text{noise}, \text{aux})$ are indistinguishable. In words, given aux, no PPT \mathcal{A} can learn sufficiently about (an approximation of) \mathbf{P} , in that \mathbf{P} and $\mathbf{P} + \text{noise}$ look the same to \mathcal{A} .

On ℓ and $\mathcal{U}(\{0, 1, \dots, \ell\})$. The parameter ℓ parametrises the strength of Pre2 and the overall evasive LWE assumption. For example, if $\ell = q$, then Pre2 can be seen as requiring \mathbf{P} to be pseudorandom, and since this is the hardest case to achieve, the resulting Hiding Evasive LWE assumption is the weakest. As ℓ decreases, i.e. less noise is added to \mathbf{P} , Pre2 becomes easier to be satisfied, and the evasive LWE assumption grows stronger. The noise distribution $\mathcal{U}(\{0, 1, \dots, \ell\})$ may alternatively be replaced by other natural distributions, e.g. discrete Gaussian over \mathcal{R} , with the Gaussian parameter being a suitable ℓ parametrising the hardness of Pre2.

Relating to existing definitions. Suppose $\ell = 0$ (which is disallowed in Definition 9), then the experiment Pre2 is trivial and Definition 9 collapses to the private-coin variant in [VWW22]. Letting $m_P = 2m$, our counterexamples prove that the assumption for this setting is false (conditioned on other appropriate parameters). On the other hand, even just by setting $\ell = 1$, we are unaware of counterexamples (except the obfuscation-based one by [VWW22], c.f. Remark 3). Moreover, in Section 6 we will see that Definition 9 with a large ℓ can be applied to the security proof of [VWW22], without altering their parameters.

$\text{Pre}_{\mathcal{A}}^{b,\beta,\gamma}(1^\lambda)$	$\text{Post}_{\mathcal{B}}^{b,\beta,\gamma}(1^\lambda)$
$\mathbf{B} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$	$\mathbf{B} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$
if $\beta = 0$ then	if $\beta = 0$ then
$(\mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	$(\mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$
if $\beta = 1$ then	if $\beta = 1$ then
$(\mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda, \mathbf{B})$	$(\mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda, \mathbf{B})$
if $b = 0$ then	if $b = 0$ then
$\mathbf{e}_0 \leftarrow_{\$} D_{\mathbb{Z}^m, \chi}, \mathbf{e}_0 \leftarrow_{\$} D_{\mathbb{Z}, \chi}, \mathbf{s} \leftarrow_{\$} \mathbb{Z}_q^n$	$\mathbf{e}_0 \leftarrow_{\$} D_{\mathbb{Z}, \chi}, \mathbf{s} \leftarrow_{\$} \mathbb{Z}_q^n$
$\mathbf{c} := \mathbf{s}^\top \mathbf{B} + \mathbf{e}_B^\top \pmod q$	$\mathbf{c} := \mathbf{s}^\top \mathbf{B} + \mathbf{e}_B^\top \pmod q$
$\mathbf{c}_P := \mathbf{s}^\top \mathbf{P} + \mathbf{e}_P^\top \pmod q$	$\mathbf{D} \leftarrow_{\$} D_{\Lambda_{\mathbb{F}}^\perp(\mathbf{B}), \chi}$
if $b = 1$ then	if $b = 1$ then
$\mathbf{c} \leftarrow_{\$} \mathbb{Z}_q^m$	$\mathbf{c} \leftarrow_{\$} \mathbb{Z}_q^m$
$\mathbf{c}_P \leftarrow_{\$} \mathbb{Z}_q^{m_P}$	$\mathbf{D} \leftarrow_{\$} D_{\Lambda_{\mathbb{F}}^\perp(\mathbf{B}), \chi}$
if $\gamma = 0$ then	if $\gamma = 0$ then
return $\mathcal{A}(\mathbf{c}, \mathbf{c}_P, \text{aux})$	return $\mathcal{B}(\mathbf{c}, \mathbf{D}, \text{aux})$
if $\gamma = 1$ then	if $\gamma \in \{1, 2\}$ then
return $\mathcal{A}(\mathbf{B}, \mathbf{c}, \mathbf{c}_P, \text{aux})$	return $\mathcal{B}(\mathbf{B}, \mathbf{c}, \mathbf{D}, \text{aux})$
if $\gamma = 2$ then	
return $\mathcal{A}(\mathbf{B}, \mathbf{c}, \mathbf{P}, \mathbf{c}_P, \text{aux})$	

Fig. 5: Experiments Pre and Post for private-coin evasive LWE variants, to which our counterexamples apply.

Remark 7 (What if Pre2 restricts some but not all entries of \mathbf{P}). In Appendix C.1, we sketch an alternative counterexample against the assumption in [VWW22], which demonstrates that it is necessary for most entries of \mathbf{P} to be irrecoverable, as even leaking only m entries of \mathbf{P} (for m the number of columns of \mathbf{B}) would lead to a successful distinguisher for Post^b .

Remark 8 (Alternative Pre2 candidate). Another way to potentially capture the intuition of “cannot approximate \mathbf{P} given aux ” might be to ask for indistinguishability of $\mathbf{P} \pmod \ell$ and $\mathbf{R} \pmod \ell$, for a uniform \mathbf{R} over \mathbb{Z}_q .¹² Other than the complication in formalisation and potential issues due to number-theoretic relations between q and ℓ , this alternative is also intuitively a weaker pre-condition (hence leading to stronger evasive assumption). For more context, in Appendix C.2 we provide a heuristic counterexample against such alternative Pre2 for $\ell = 2$ and $m_P = O(m^2)$, which may be further generalised to be against constant ℓ and $m_P = O(m^\ell)$ [AG11, NMSÜ24].

5 Counterexamples to Existing Variants

We present our counterexamples against a number of existing evasive LWE variants sketched in Section 2.2. These variants are formally defined in Fig. 5 and parametrised by β, γ , each controlling if Samp receives \mathbf{B} as input or not, and if the distinguishers receive \mathbf{B} , (\mathbf{B}, \mathbf{P}) , or none. We remark that this definition is a special case of the private-coin variants in existing works, where we consider the restricted setting of the LWE sample \mathbf{s} being sampled honestly by the experiments, not available to Samp .

Denote the advantages of distinguishers \mathcal{A}, \mathcal{B} and $\beta \in \{0, 1\}, \gamma \in \{0, 1, 2\}$ by

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{Pre}, \beta, \gamma}(\lambda) &= \left| \Pr[\text{Pre}_{\mathcal{A}}^{0, \beta, \gamma}(1^\lambda) = 0] - \Pr[\text{Pre}_{\mathcal{A}}^{1, \beta, \gamma}(1^\lambda) = 0] \right| \\ \text{Adv}_{\mathcal{A}}^{\text{Post}, \beta, \gamma}(\lambda) &= \left| \Pr[\text{Post}_{\mathcal{A}}^{0, \beta, \gamma}(1^\lambda) = 0] - \Pr[\text{Post}_{\mathcal{A}}^{1, \beta, \gamma}(1^\lambda) = 0] \right| \end{aligned}$$

¹² These mean, fix the representatives of $\mathbb{Z}_q = \{0, 1, \dots, q-1\}$ and map each entry in \mathbb{Z}_q to their representative over \mathbb{Z} , then further apply $\pmod \ell$ operation. For an intuition, if $\ell = 2^k$ is a power of 2, then $\mathbf{P} \pmod \ell$ may be interpreted as the last k bits of (each entry of) \mathbf{P} .

In all counterexamples, we consider the usual case of the ring of integers $\mathcal{R} = \mathbb{Z}$ and assume that error noise is distributed according to a discrete Gaussian distribution $D_{\mathbb{Z}^m, \chi}$ for a parameter $\chi > 0$ to be specified. We require q to be prime, and abide to the following parameter restrictions for n, m, χ, q

$$m \geq 3n \log(q), \quad \chi \geq \lambda m, \quad q \geq \lambda^3 m^2 \chi^2.$$

Remark 9 (On Parameters). The above parameters are polynomial (in particular the modulus q can be chosen to lie in $O(n^4 \lambda^6)$, for example) and the attacks we propose in the following have a high advantage of at least $1 - O(1/\lambda)$ to win the then challenge. It is possible to increase the advantage to be overwhelming in λ , by choosing $q \gg \chi \gg m$ such that q is super-polynomially larger than χ , and χ is super-polynomially larger than m .

The following lemma upper-bounds the probability of a square Gaussian matrix sampled over random cosets being not invertible over \mathbb{Z}_q , which be useful for our counterexamples. Its proof in more generality is given in Appendix A.

Lemma 5. *Let $\mathbf{B}, \mathbf{P}_1 \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$ be uniformly random. Let $\chi \in \omega(\sqrt{n})$, $m \geq 2n \log q$ and let $q > n \cdot \chi$ be prime. We have*

$$\Pr_{\mathbf{D}_1 \leftarrow_{\$} D_{\Lambda_{\mathbf{P}_1}^{\perp}(\mathbf{B}), \chi}} [\det(\mathbf{D}_1) = 0 \pmod q] \leq O(m/\chi).$$

Remark 10 (Alternative to Lemma 5). As an alternative to Lemma 5, one can use a similar result by Regev [Reg05], which states that m^2 many Gaussian vectors $\mathbf{d}_1, \dots, \mathbf{d}_{m^2} \leftarrow_{\$} D_{\mathbb{Z}^m, \chi}$ will contain a basis of \mathbb{Z}^m with probability at least $1 - 2^{-\Omega(n)}$. Correspondingly, the numbers of columns of the matrices \mathbf{P}_1 and \mathbf{P}_3 in Sections 5.1 and 5.2 is required to increase from m to m^2 .

5.1 Counterexample 1

We give a counterexample for the case $\beta = 0, \gamma = 1$, i.e., \mathbf{B} is given to the distinguishers, but $\mathbf{P} = (\mathbf{P}_1, \mathbf{P}_2)$ is not and aux is empty. Our idea is to embed a trapdoor in \mathbf{P}_2 , which we use to distinguish $\mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T$ from uniform randomness in the Post^b challenge. Concretely, let $\text{Samp}_1(1^\lambda)$ output the following:

$$(\mathbf{P} = (\mathbf{P}_1, \mathbf{P}_2), \text{aux} = \perp)$$

where

$$\begin{aligned} \mathbf{u}' &\leftarrow_{\$} \{0, 1\}^{m-1}, & \mathbf{P}'_1 &\leftarrow_{\$} \mathbb{Z}_q^{n \times (m-1)}, & \mathbf{P}'_2 &\leftarrow_{\$} \mathbb{Z}_q^{(n-1) \times m}, \\ \mathbf{u}^T &= ((\mathbf{u}')^T, 1) \in \{0, 1\}^{1 \times m}, & \mathbf{P}_1 &= (\mathbf{P}'_1 \parallel -\mathbf{P}'_1 \mathbf{u}' \pmod q), & \mathbf{P}_2 &= \begin{pmatrix} \mathbf{P}'_2 \\ \mathbf{u}^T \end{pmatrix}. \end{aligned}$$

Proposition 1. *Let \mathcal{A} be a PPT adversary. Under the $\text{LWE}_{\mathbb{Z}, q, n, 2m, \chi}$ assumption, we have for the experiment $\text{Pre}_{\mathcal{A}}^{\beta=0, \gamma=1}(1^\lambda)$ in Fig. 5 instantiated with Samp_1*

$$\text{Adv}_{\mathcal{A}}^{\text{Pre}, 0, 1}(\lambda) \in \text{negl}(\lambda).$$

Proof. The proof proceeds via four hybrid experiments:

\mathcal{D}_0 : The joint distribution of the ifstatement, i.e.

$$(\mathbf{B}, \quad \mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T \pmod q, \quad \mathbf{s}^T \mathbf{P}_1 + \mathbf{e}_1^T \pmod q, \quad \mathbf{s}^T \mathbf{P}_2 + \mathbf{e}_2^T \pmod q)$$

where $\mathbf{B} \leftarrow_{\$} \mathcal{R}_q^{n \times m}$, $\mathbf{s} \leftarrow_{\$} \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow_{\$} \chi^m$, $\mathbf{e}_1 \leftarrow_{\$} \chi^m$, $\mathbf{e}_2 \leftarrow_{\$} \chi^m$.

\mathcal{D}_1 : The last term $\mathbf{s}^T \mathbf{P}_2 + \mathbf{e}_2^T \pmod q$ is replaced by a random vector \mathbf{y}_2 , i.e.

$$(\mathbf{B}, \quad \mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T \pmod q, \quad \mathbf{s}^T \mathbf{P}_1 + \mathbf{e}_1^T \pmod q, \quad \mathbf{y}_2)$$

for $\mathbf{y}_2 \leftarrow_{\$} \mathbb{Z}_q^m$.

\mathcal{D}_2 : As \mathcal{D}_1 , but \mathbf{P}_1 is swapped to random.

$\mathcal{D}_3 : \mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top \bmod q$ and $\mathbf{s}^\top \mathbf{P}_1 + \mathbf{e}_1^\top \bmod q$ are swapped to random, i.e.

$$(\mathbf{B}, \mathbf{y}_0, \mathbf{y}_1, \mathbf{y}_2)$$

for $\mathbf{y}_0, \mathbf{y}_1, \mathbf{y}_2 \leftarrow \mathbb{Z}_q^m$.

The statistical distance between \mathcal{D}_0 and \mathcal{D}_1 is bounded by q^{-n+1} . This is because for $\mathbf{s} = (\mathbf{s}', s_n) \leftarrow \mathbb{Z}_q^n$, $\mathbf{s}^\top \mathbf{P}_2 = \mathbf{s}'^\top \mathbf{P}'_2 + s_n \cdot \mathbf{u}^\top \bmod q$ is uniformly random if \mathbf{s}' is non-zero (even if we know $\mathbf{B}, \mathbf{P}_1, \mathbf{s}$ and \mathbf{u}).

Since $m \geq 2n \log q$, the Leftover-Hash Lemma 3 implies that the statistical distance between $\mathbf{P}_1 = (\mathbf{P}'_1 | - \mathbf{P}'_1 \mathbf{u}' \bmod q)$ and a uniformly random matrix is bounded by 2^{-n} . Hence, the statistical distance between \mathcal{D}_1 and \mathcal{D}_2 is also bounded by 2^{-n} . Finally, $\text{LWE}_{\mathbb{Z}, q, n, 2m, \chi}$ states that \mathcal{D}_2 and \mathcal{D}_3 are computationally indistinguishable.

Therefore, \mathcal{A} 's distinguishing advantage between $\text{Pre}_{\mathcal{A}}^{0,0,1}(1^\lambda) = \mathcal{D}_0$ and $\text{Pre}_{\mathcal{A}}^{1,0,1}(1^\lambda) = \mathcal{D}_3$ is bounded by 2^{-n+1} plus a negligible term stemming from the LWE assumption. \square

We introduce two more lemmas, their proofs are given in Appendix B.

Lemma 6. Let $(\mathbf{P}_1 \in \mathbb{Z}_q^{n \times m}, \mathbf{u} \in \{0, 1\}^m) \leftarrow \text{Samp}_1(1^\lambda)$. For $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{D}_1 \leftarrow D_{\Lambda_{\mathbf{P}_1}^\perp(\mathbf{B}), \chi}$, we have

$$\Pr[\mathbf{D}_1 \cdot \mathbf{u} = 0 \bmod q] \leq O(m/\chi).$$

Lemma 7. Let $(\mathbf{P}_1 \in \mathbb{Z}_q^{n \times m}, \mathbf{u} \in \{0, 1\}^m) \leftarrow \text{Samp}_1(1^\lambda)$. Sample $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \chi}$ and $\mathbf{D}_1 \leftarrow D_{\Lambda_{\mathbf{P}_1}^\perp(\mathbf{B}), \chi}$. We have

$$\Pr[|\mathbf{e}_0^\top \cdot \mathbf{D}_1 \cdot \mathbf{u}| \geq \lambda^2 m^2 \chi^2] \leq 2^{-\lambda}.$$

Proposition 2. There is a PPT adversary \mathcal{B} s.t. we have for the experiment $\text{Post}_{\mathcal{B}}^{\beta=0, \gamma=1}(1^\lambda)$ in Fig. 5 instantiated with Samp_1

$$\text{Adv}_{\mathcal{B}}^{\text{Post}, 0, 1}(\lambda) \geq 1 - O(1/\lambda).$$

Proof. Let $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$ and $(\mathbf{S}, \mathbf{P}, \text{aux} = \perp) \leftarrow \text{Samp}_1(1^\lambda)$. Sample $\mathbf{D} \leftarrow D_{\Lambda_{\mathbf{P}}^\perp(\mathbf{B}), \chi}$ and note that $\mathbf{D} \in \mathbb{Z}^{n \times 2m}$ can be split into two equally large parts $\mathbf{D} = (\mathbf{D}_1 | \mathbf{D}_2)$, which are distributed as

$$\mathbf{D}_1 \leftarrow D_{\Lambda_{\mathbf{P}_1}^\perp(\mathbf{B}), \chi} \quad \text{and} \quad \mathbf{D}_2 \leftarrow D_{\Lambda_{\mathbf{P}_2}^\perp(\mathbf{B}), \chi}.$$

Recall that in the experiment $\text{Post}_{\mathcal{B}}^{\beta=0, \gamma=1}(1^\lambda)$, \mathcal{B} has to decide if \mathbf{c}^\top in

$$(\mathbf{B}, \mathbf{D}_1, \mathbf{D}_2, \mathbf{c}^\top)$$

equals $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top \bmod q$, for $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \chi}$ and $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, or has been sampled uniformly at random from $\mathbb{Z}_q^{1 \times m}$.

Our adversary \mathcal{B} proceeds as follows:

1. It recovers $\mathbf{P}_2 := \mathbf{B} \cdot \mathbf{D}_2 \bmod q$.
2. Since Samp_1 samples \mathbf{P}_2 as $\begin{pmatrix} \mathbf{P}'_2 \\ \mathbf{u}^\top \end{pmatrix}$, \mathcal{B} can extract the binary vector $\mathbf{u} \in \{0, 1\}^m$ from the last row of \mathbf{P}_2 .
3. \mathcal{B} computes $r := \mathbf{c}^\top \cdot \mathbf{D}_1 \cdot \mathbf{u} \bmod q$.
4. If $r \in \{-\lambda^2 \chi^2 m^2, \dots, \lambda^2 \chi^2 m^2\} \subset \mathbb{Z}_q$, then \mathcal{B} outputs 0. Otherwise, it outputs 1.

If \mathbf{c} is drawn uniformly at random from \mathbb{Z}_q^m , then the probability that r lies in $\{-\lambda^2 \chi^2 m^2, \dots, \lambda^2 \chi^2 m^2\}$ is bounded by $O(\lambda^2 \chi^2 m^2 / q) \subseteq O(1/\lambda)$. This is because r is distributed uniformly at random in \mathbb{Z}_q if $\mathbf{D}_1 \cdot \mathbf{u} \bmod q$ is non-zero, which happens with probability at least $1 - O(m/\chi) = 1 - O(1/\lambda)$ by Lemma 6. Hence, if \mathbf{c} is uniformly random, \mathcal{B} outputs 1 with probability $1 - O(1/\lambda)$.

Else, if $\mathbf{c} = \mathbf{B}^\top \mathbf{s} + \mathbf{e}_0 \bmod q$ for $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \chi}$ and $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, then

$$\begin{aligned} r &= \mathbf{c}^\top \cdot \mathbf{D}_1 \mathbf{u} = (\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top) \cdot \mathbf{D}_1 \mathbf{u} = \mathbf{s}^\top \mathbf{B} \cdot \mathbf{D}_1 \mathbf{u} + \mathbf{e}_0^\top \cdot \mathbf{D}_1 \mathbf{u} \\ &= \mathbf{s}^\top \mathbf{P}_1 \mathbf{u} + \mathbf{e}_0^\top \mathbf{D}_1 \mathbf{u} = \mathbf{e}_0^\top \mathbf{D}_1 \mathbf{u} \bmod q. \end{aligned}$$

By Lemma 7, $\|\mathbf{e}_0^\top \mathbf{D}_1 \mathbf{u}\| \leq \lambda^2 m^2 \chi^2$ with probability at least $1 - 2^{-\lambda}$. Hence, \mathcal{B} will output 0 if $\mathbf{c} = \mathbf{B}^\top \mathbf{s} + \mathbf{e}_0 \bmod q$ with overwhelming probability. It follows that the advantage of \mathcal{B} lies in $1 - O(1/\lambda)$. \square

5.2 Counterexample 2

We give a counterexample for the case $\beta = \gamma = 0$, i.e. \mathbf{B} is not given to the distinguisher. The sampler is similar to that in Section 5.1, except that we add a third part $\mathbf{P}_3 \leftarrow \mathbb{Z}_q^{n \times m}$ to \mathbf{P} , which is also included in aux . While \mathbf{P}_3 is harmless on its own, the distinguisher for Post^b can use it to recover \mathbf{B} and continue as in the first counterexample. Concretely, let $\text{Samp}_2(1^\lambda)$ output the following:

$$(\mathbf{P} = (\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3), \quad \text{aux} = \mathbf{P}_3)$$

where $\mathbf{P}_3 \leftarrow \mathbb{Z}_q^{n \times m}$ and

$$\begin{aligned} \mathbf{u}' &\leftarrow \{0, 1\}^{m-1}, & \mathbf{P}'_1 &\leftarrow \mathbb{Z}_q^{n \times (m-1)}, & \mathbf{P}'_2 &\leftarrow \mathbb{Z}_q^{(n-1) \times m}, \\ \mathbf{u}^\top &= ((\mathbf{u}')^\top, 1) \in \{0, 1\}^{1 \times m}, & \mathbf{P}_1 &= (\mathbf{P}'_1 | -\mathbf{P}'_1 \mathbf{u}' \bmod q), & \mathbf{P}_2 &= \begin{pmatrix} \mathbf{P}'_2 \\ \mathbf{u}^\top \end{pmatrix}. \end{aligned}$$

Proposition 3. *Let \mathcal{A} be a PPT adversary. Under the $\text{LWE}_{\mathbb{Z}, q, n, 3m, \chi}$ assumption, we have for the experiment $\text{Pre}_{\mathcal{A}}^{\beta=0, \gamma=0}(1^\lambda)$ in Fig. 5 instantiated with Samp_2*

$$\text{Adv}_{\mathcal{A}}^{\text{Pre}, 0, 0}(\lambda) \in \text{negl}(\lambda).$$

Proof. Let $\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3$ be the output of $\text{Samp}_2(1^\lambda)$. Note that in this case, \mathcal{A} has to distinguish between the distribution

$$(\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top, \quad \mathbf{s}^\top \mathbf{P}_1 + \mathbf{e}_1^\top, \quad \mathbf{s}^\top \mathbf{P}_2 + \mathbf{e}_2^\top, \quad \mathbf{s}^\top \mathbf{P}_3 + \mathbf{e}_3^\top, \quad \mathbf{P}_3) \bmod q \quad (9)$$

for $\mathbf{B} \leftarrow \mathcal{R}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow \chi^m$, $\mathbf{e}_1 \leftarrow \chi^m$, $\mathbf{e}_2 \leftarrow \chi^m$, and the distribution

$$(\mathbf{y}_0, \quad \mathbf{y}_1, \quad \mathbf{y}_2, \quad \mathbf{y}_3, \quad \mathbf{P}_3),$$

for $\mathbf{y}_0, \mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3 \leftarrow \mathbb{Z}_q^m$. By the same argument as in the proof of Proposition 1, the statistical distance between the distribution in Eq. (9) and

$$(\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top \bmod q, \quad \mathbf{s}^\top \mathbf{P}'_1 + \mathbf{e}_1^\top \bmod q, \quad \mathbf{y}_2, \quad \mathbf{s}^\top \mathbf{P}_3 + \mathbf{e}_3^\top \bmod q, \quad \mathbf{P}_3), \quad (10)$$

for $\mathbf{P}'_1 \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{y}_2 \leftarrow \mathbb{Z}_q^m$ is bounded by 2^{-n+1} . Now, the claim follows by invoking $\text{LWE}_{\mathbb{Z}, q, n, 3m, \chi}$. \square

Proposition 4. *There is a PPT adversary \mathcal{B} s.t. we have for the experiment $\text{Post}_{\mathcal{B}}^{\beta=0, \gamma=0}(1^\lambda)$ in Fig. 5 instantiated with Samp_2*

$$\text{Adv}_{\mathcal{B}}^{\text{Post}, 0, 0}(\lambda) \geq 1 - O(1/\lambda).$$

Proof. For $i = 1, 2, 3$, let $\mathbf{D}_i \leftarrow D_{A_{\mathbf{P}_i}^\perp(\mathbf{B}), \chi}$ be the short preimages that \mathcal{B} is given in Post^b . Denote the adversary of Proposition 2 from counterexample 1 by \mathcal{B}' . We want to invoke \mathcal{B}' on Post^b , however, note that \mathbf{B} is missing.

Since \mathbf{B} and \mathbf{P}_3 have been sampled uniformly at random, Lemma 5 implies that $\mathbf{D}_3 \bmod q$ is invertible with probability $1 - O(1/\lambda)$. Hence, \mathcal{B} on input

$$(\mathbf{D}_1, \quad \mathbf{D}_2, \quad \mathbf{D}_3, \quad \mathbf{c}^\top, \quad \text{aux} = \mathbf{P}_3)$$

proceeds as follows:

1. If $\mathbf{D}_3 \bmod q$ is not invertible, then \mathcal{B} outputs a random bit and stops.
2. If $\mathbf{D}_3 \bmod q$ is invertible, \mathcal{B} computes $\mathbf{B} = \mathbf{D}_3^{-1} \cdot \mathbf{P}_3 \bmod q$.
3. \mathcal{B} runs $\mathcal{B}'(\mathbf{B}, \mathbf{D}_1, \mathbf{D}_2, \mathbf{c}^\top)$ and defers its output to the post challenger.

By Lemma 5 and Proposition 1, the advantage of \mathcal{B} is at least $1 - O(1/\lambda)$. \square

5.3 Counterexample 3

We give a counterexample for the cases $\beta = 1, \gamma \in \{0, 1, 2\}$. Note that $\beta = 1$ implies that **Samp** gets **B** as input. We prove indistinguishability of Pre^b in the strongest case, where the joint distribution contains **B** and **P**, i.e. $\gamma = 2$. We also show that, there exists PPT distinguisher for the weakest Post^b challenge, where the joint distribution will contain neither **P** nor **B**, i.e. $\gamma = 0$.

Let $\text{Samp}_3(1^\lambda, \mathbf{B})$, on input **B**, output a matrix **P** constructed as follows:

$$\mathbf{U} = (\mathbf{u}_1 | \mathbf{u}_2) \leftarrow_{\$} \{0, 1\}^{m \times 2}, \quad \mathbf{P} = (\mathbf{p}_1 | \mathbf{p}_2) := \mathbf{B}\mathbf{U} \bmod q \in \mathbb{Z}_q^{n \times 2}.$$

Additionally, Samp_3 outputs **aux**, which consists of the Dual Regev encryptions (under $(\mathbf{B} | \mathbf{p}_1) \in \mathbb{Z}_q^{n \times (m+1)}$) of the bits $u_{2,1}, \dots, u_{2,m}$ of **u**, that is,

$$\mathbf{aux} = \left(\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i \bmod q, \quad \mathbf{r}_i^T \mathbf{p}_1 + f_{i,m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2,i} \bmod q \right)_{i \in [m]} \in \mathbb{Z}_q^{m \times (m+1)}$$

where Samp_3 samples $\mathbf{r}_1, \dots, \mathbf{r}_m \leftarrow_{\$} \mathbb{Z}_q^n$ and $\mathbf{f}_i = (f'_i, f_{i,m+1}) \leftarrow_{\$} D_{\mathbb{Z}^{m+1}, \chi}$.

Proposition 5. *Let \mathcal{A} be a PPT adversary. Under the $\text{LWE}_{\mathbb{Z}, q, n, m+2, \chi}$ assumption, we have for the experiment $\text{Pre}_{\mathcal{A}}^{\beta=1, \gamma=2}(1^\lambda)$ in Fig. 5 instantiated with Samp_3*

$$\text{Adv}_{\mathcal{A}}^{\text{Pre}, 1, 2}(\lambda) \in \text{negl}(\lambda).$$

Proof. We proceed via the following hybrid experiments:

\mathcal{D}_0 : This distribution corresponds to the view of \mathcal{A} in $\text{Pre}^{0,1,2}(1^\lambda)$, i.e.

$$\left(\mathbf{B}, \quad \mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T \bmod q, \quad \mathbf{p}_1, \quad \mathbf{p}_2, \quad \mathbf{s}^T \mathbf{p}_1 + e_1^T \bmod q, \quad \mathbf{s}^T \mathbf{p}_2 + e_2^T \bmod q, \right. \\ \left. \mathbf{aux} = \left(\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i \bmod q, \quad \mathbf{r}_i^T \mathbf{p}_1 + f_{i,m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2,i} \bmod q \right)_{i=1, \dots, m} \right)$$

where $\mathbf{B} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$, $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{aux}) \leftarrow \text{Samp}_3(\mathbf{B}, 1^\lambda)$, $\mathbf{s} \leftarrow_{\$} \mathbb{Z}_q^n$, $\mathbf{e}_0, \leftarrow_{\$} D_{\mathbb{Z}^m, \chi}$, $e_1, e_2 \leftarrow_{\$} D_{\mathbb{Z}, \chi}$.

\mathcal{D}_1 : \mathcal{D}_1 resembles \mathcal{D}_0 , but we replace \mathbf{p}_1 in the joint distribution and in the auxiliary information **aux** of the sampler with a uniformly random vector $\mathbf{p}'_1 \leftarrow_{\$} \mathbb{Z}_q^m$.

\mathcal{D}_2 : We replace $\mathbf{aux} = \left(\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i \bmod q, \mathbf{r}_i^T \mathbf{p}'_1 + f_{i,m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2,i} \bmod q \right)_{i \in [m]}$ by a uniformly random matrix $\mathbf{aux}' \leftarrow_{\$} \mathbb{Z}_q^{m \times (m+1)}$ in \mathcal{D}_1 .

\mathcal{D}_3 : We replace \mathbf{p}_2 by a uniformly random vector in \mathcal{D}_2 .

\mathcal{D}_4 : We replace $\mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T$, $\mathbf{s}^T \mathbf{p}'_1 + e_1^T$, and $\mathbf{s}^T \mathbf{p}'_2 + e_2^T$ by uniformly random vectors and numbers over \mathbb{Z}_q^n .

\mathcal{D}_5 : We swap \mathbf{p}'_2 back to \mathbf{p}_2 in \mathcal{D}_4 .

\mathcal{D}_6 : We revert the changes on \mathbf{aux}' and put again

$$\mathbf{aux} = \left(\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i \bmod q, \mathbf{r}_i^T \mathbf{p}'_1 + f_{i,m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2,i} \bmod q \right)_{i=1, \dots, m}.$$

\mathcal{D}_7 : Finally, we replace the uniformly random vector \mathbf{p}'_1 in \mathcal{D}_6 by the vector $\mathbf{p}_1 = \mathbf{B}\mathbf{u}_1$ outputted by the sampler. Note that \mathcal{D}_7 equals now

$$\left(\mathbf{B}, \quad \mathbf{c}_0, \quad \mathbf{p}_1, \quad \mathbf{p}_2, \quad c'_1, \quad c'_2, \right. \\ \left. \mathbf{aux} = \left(\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i \bmod q, \quad \mathbf{r}_i^T \mathbf{p}_1 + f_{i,m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2,i} \bmod q \right)_{i=1, \dots, m} \right)$$

for $\mathbf{c}_0 \leftarrow_{\$} \mathbb{Z}_q^m$, $c'_1, c'_2 \leftarrow_{\$} \mathbb{Z}_q$. Hence, \mathcal{D}_7 is the view of \mathcal{A} in $\text{Pre}^{1,1,2}(1^\lambda)$.

We claim that the advantage of \mathcal{A} to distinguish between the hybrids is negligible: By the Leftover-Hash Lemma 3, the statistical distance between \mathcal{D}_0 and \mathcal{D}_1 , between \mathcal{D}_2 and \mathcal{D}_3 , between \mathcal{D}_4 and \mathcal{D}_5 and between \mathcal{D}_6 and \mathcal{D}_7 is bounded by 2^{-n} . Going from \mathcal{D}_1 to \mathcal{D}_2 , we replace the m -fold dual-Regev encryption **aux** by a uniformly random matrix $\mathbf{aux}' \leftarrow_{\$} \mathbb{Z}_q^{m \times (m+1)}$, which we revert again when we go from \mathcal{D}_5 to \mathcal{D}_6 . In total, we invoke the $\text{LWE}_{\mathbb{Z}, q, n, m+1, \chi}$ assumption in both directions m times. Finally, the $\text{LWE}_{\mathbb{Z}, q, n, m+2, \chi}$ assumption stipulates that it is hard for \mathcal{A} to distinguish between \mathcal{D}_3 and \mathcal{D}_4 .

Concluding, the advantage of \mathcal{A} to win $\text{Pre}^{\beta=1, \gamma=2}(1^\lambda)$ is bounded by $O(2^{-n})$ plus a negligible term stemming from $(2m+1)$ -times of invoking LWE. \square

We require two more lemmas, their proofs are given in Appendix B.

Lemma 8. For $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{u}_2 \leftarrow \{0, 1\}^m$, $\mathbf{p}_2 = \mathbf{B} \cdot \mathbf{u}_2$ and $\mathbf{d}_2 \leftarrow D_{\Lambda_{\mathbf{p}_2}^\perp(\mathbf{B}), \chi}$, we have $\Pr[\mathbf{d}_2 = \mathbf{u}_2] \leq 2^{-m+1} + q^{-n}$.

Lemma 9. Draw $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{u}_2 \leftarrow \{0, 1\}^m$ and set $\mathbf{p}_2 = \mathbf{B}\mathbf{u}_2 \bmod q$. Draw $\mathbf{d}_2 \leftarrow D_{\Lambda_{\mathbf{p}_2}^\perp(\mathbf{B}), \chi}$ and $\mathbf{f}'_1, \dots, \mathbf{f}'_m \leftarrow D_{\mathbb{Z}^m, \chi}$.

We have

$$\Pr[\exists i \in [m]: |\mathbf{f}'_i \cdot \mathbf{d}_2| \geq \lambda^2 m \chi^2] \leq 2^{-\lambda}.$$

Additionally, we have for $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}_q^m, \chi}$,

$$\Pr[|\mathbf{e}_0 \cdot (\mathbf{d}_2 - \mathbf{u}_2)| \geq \lambda \chi \cdot (\lambda \chi + 1) \cdot m] \leq 2^{-\lambda}.$$

Proposition 6. There is a PPT adversary \mathcal{B} s.t. we have for the experiment $\text{Post}_{\mathcal{B}}^{\beta=1, \gamma=0}(1^\lambda)$ in Fig. 5 instantiated with Samp_3

$$\text{Adv}_{\mathcal{B}}^{\text{Post}, 1, 0}(\lambda) \geq 1 - O(1/(\lambda m)).$$

Proof. Recall that the post challenge consists of

$$\left(\begin{array}{l} \mathbf{c}, \mathbf{d}_1^T, \mathbf{d}_2^T, \\ \text{aux} = \left(\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i \bmod q, \mathbf{r}_i^T \mathbf{p}_1 + f_{i, m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2, i} \bmod q \right)_{i=1, \dots, m} \end{array} \right)$$

for $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $(\mathbf{p}_1, \mathbf{p}_2, \text{aux}) \leftarrow \text{Samp}_3(\mathbf{B}, 1^\lambda)$, $\mathbf{d}_1 \leftarrow D_{\Lambda_{\mathbf{p}_1}^\perp(\mathbf{B}), \chi}$, $\mathbf{d}_2 \leftarrow D_{\Lambda_{\mathbf{p}_2}^\perp(\mathbf{B}), \chi}$. \mathcal{B} has to decide if \mathbf{c} has been sampled uniformly at random from \mathbb{Z}_q^m or is of shape $\mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T \bmod q$ for $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \chi}$.

The distinguisher \mathcal{B} , we propose, proceeds as follows:

1. For $i \in [m]$, it computes

$$\begin{aligned} g_i &:= \mathbf{r}_i^T \mathbf{p}_1 + f_{i, m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2, i} - (\mathbf{r}_i^T \mathbf{B} + \mathbf{f}'_i) \cdot \mathbf{d}_1 && \bmod q \\ &= \mathbf{r}_i^T \mathbf{p}_1 + f_{i, m+1} + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2, i} - \mathbf{r}_i^T \cdot \mathbf{p}_1 - \mathbf{f}'_i \cdot \mathbf{d}_1 && \bmod q \\ &= f_{i, m+1} - \mathbf{f}'_i \cdot \mathbf{d}_1 + \left\lfloor \frac{q}{2} \right\rfloor \cdot u_{2, i} && \bmod q. \end{aligned}$$

Further, it sets $u'_i = 1$ if $|g_i| \geq \frac{q}{4}$ and else 0, and $\mathbf{u}' = (u'_1, \dots, u'_m) \in \{0, 1\}^m$.

2. It computes $r := \mathbf{c}^T \cdot (\mathbf{d}_2 - \mathbf{u}') \bmod q$.
3. If $r \in \{-\lambda \chi (\lambda \chi + 1) \cdot m, \dots, \lambda \chi (\lambda \chi + 1) \cdot m\} \subset \mathbb{Z}_q$, output 0. Else, output 1.

We claim that the dual-Regev decryption of $u_{2,1}, \dots, u_{2,m}$ in step 1 succeeds with overwhelming probability. Indeed, we have $u'_i = u_{2,i}$ whenever $|f_{i, m+1} - \mathbf{f}'_i \cdot \mathbf{d}_1|$ is bounded by $q/4$. Lemma 9 guarantees that we have

$$|f_{i, m+1} - \mathbf{f}'_i \cdot \mathbf{d}_1| < m \chi^2 < \frac{q}{4}$$

with overwhelming probability $\geq 1 - 2^{-\lambda}$.

Hence, assume that decryption succeeds, which happens with overwhelming probability, and that \mathcal{B} can recover $\mathbf{u}' = \mathbf{u}_2$. Assume that $\mathbf{d}_2 - \mathbf{u}' = \mathbf{d}_2 - \mathbf{u}_2$ is not zero. We now distinguish two cases:

If $\mathbf{c} = \mathbf{B}^T \mathbf{s} + \mathbf{e}_0$, then we have

$$\begin{aligned} r &= \mathbf{c}^T \cdot (\mathbf{d}_2 - \mathbf{u}') = (\mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T) \cdot (\mathbf{d}_2 - \mathbf{u}') \bmod q \\ &= \mathbf{s}^T \mathbf{B} \cdot (\mathbf{d}_2 - \mathbf{u}') + \mathbf{e}_0^T \cdot (\mathbf{d}_2 - \mathbf{u}') \bmod q \\ &= \mathbf{s}^T \cdot (\mathbf{p}_2 - \mathbf{p}_2) + \mathbf{e}_0^T \cdot (\mathbf{d}_2 - \mathbf{u}') = \mathbf{e}_0^T \cdot (\mathbf{d}_2 - \mathbf{u}') \bmod q. \end{aligned}$$

According to Lemma 9, the quantity $\mathbf{e}_0^T \cdot (\mathbf{d}_2 - \mathbf{u}_2)$ is bounded $\lambda \chi (\lambda \chi + 1) \cdot m$ with overwhelming probability $\geq 1 - 2^{-\lambda}$. Hence, in this case, \mathcal{B} will output 0 with overwhelming probability.

Else, if $\mathbf{c} \leftarrow \mathbb{Z}_q^n$, the value $r = \mathbf{c}^T \cdot (\mathbf{d}_2 - \mathbf{u}_2) \bmod q$ is uniform whenever $\mathbf{d}_2 - \mathbf{u}_2 \neq \mathbf{0}$. By Lemma 8 this is the case with overwhelming probability $\geq 1 - q^{-n} - 2^{-m+2}$. In this case, the probability that r lies in $\{-\lambda \chi (\lambda \chi + 1) \cdot m, \dots, \lambda \chi (\lambda \chi + 1) \cdot m\}$ is bounded by $2 \frac{\lambda \chi (\lambda \chi + 1) m}{q} \leq \frac{2(\lambda \chi + 1)}{\lambda^2 \cdot \chi \cdot m} \in O(1/(\lambda m))$. Hence, \mathcal{B} will output 1 in this case with probability $1 - O(1/(\lambda m))$. It follows that the advantage of \mathcal{B} lies in $1 - O(1/(\lambda m))$. \square

6 Evasive LWE Instance in [VWW22]

We show that we can apply the PrivateHideEvLWE assumption (Definition 9) to the security proof of [VWW22]. Consequently, assuming PrivateHideEvLWE, the GGM15-encoding in [VWW22], as well as its witness encryption and null-iO constructions, remains secure.

Let $n, w, m, h \in \text{poly}(\lambda)$, let $\hat{n} = wn$, $t = 2^{j-1}\hat{n}$, and fix some $j \in [h]$. To recall, the PPT sampler used in the proof of [VWW22, Lemma 5.2], which we denote by Samp_{VWW} , outputs the following:

$$\begin{aligned} \mathbf{S} &:= \left\{ \hat{\mathbf{S}}_{i,b} \right\}_{i \in [h], b \in \{0,1\}} \in \mathbb{Z}_q^{t \times \hat{n}}, \\ \mathbf{P} &:= \left(\hat{\mathbf{S}}_{j,0} \mathbf{A}_j + \mathbf{E}_{j,0}, \hat{\mathbf{S}}_{j,1} \mathbf{A}_j + \mathbf{E}_{j,1} \right) \bmod q \in \mathbb{Z}_q^{\hat{n} \times 2m}, \\ \text{aux} &:= \left\{ \mathbf{A}_{i-1}^{-1} (\hat{\mathbf{S}}_{i,b} \mathbf{A}_i + \mathbf{E}_{i,b} \bmod q) \right\}_{i \geq j+1, b \in \{0,1\}}, \quad \left\{ \hat{\mathbf{S}}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}, \end{aligned}$$

where

- $\hat{\mathbf{S}}_{i,b} \in \mathbb{Z}_q^{\hat{n} \times \hat{n}}$ for $i \in [h], b \in \{0,1\}$ are arbitrary matrices (in the context of [VWW22] representing a branching program), and $\left\{ \hat{\mathbf{S}}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}$ denotes stacking the 2^h matrices vertically,
- $\mathbf{E}_{j,0}, \mathbf{E}_{j,1} \leftarrow \mathfrak{s}(\mathcal{D}_{\mathbb{Z}, \chi'})^{\hat{n} \times m}$ are Gaussian with parameter $\chi' \geq \lambda^{\omega(1)} \lambda^h O(n)$,
- $\mathbf{A}_i \in \mathbb{Z}_q^{\hat{n} \times m}$ for $i \geq j+1$ are uniformly random matrices,
- each $\mathbf{A}_{i-1}^{-1} (\hat{\mathbf{S}}_{i,b} \mathbf{A}_i + \mathbf{E}_{i,b} \bmod q)$ for $i \geq j+1, b \in \{0,1\}$ denotes a Gaussian preimage w.r.t. \mathbf{A}_{i-1} for the image $\hat{\mathbf{S}}_{i,b} \mathbf{A}_i + \mathbf{E}_{i,b} \bmod q$, with parameter $\chi'' = O(2\sqrt{nw \log q})$.

The proof of [VWW22, Lemma 5.2] showed that, assuming the condition of [VWW22, Equation 6], there exists no PPT \mathcal{A} distinguishing the Pre1^b experiments in Definition 9 with respect to Samp_{VWW} with non-negligible probability. In Proposition 7 below, we show that, for a large ℓ , there exists no PPT \mathcal{A} that can win Pre2 with respect to Samp_{VWW} with non-negligible probability. Invoking the PrivateHideEvLWE_{param} assumption with $\text{param} = (\mathbb{Z}, q, n, m, 2m, t, \mathcal{U}(\mathbb{Z}_q^{n \times m}), (\mathcal{D}_{\mathbb{Z}, \chi})^{t \times m}, (\mathcal{D}_{\mathbb{Z}, \chi'})^{t \times 2m}, \ell, (\chi'')^2 \mathbf{I})$ for $\chi = \lambda^{\omega(1)} \chi'$ completes the proof of [VWW22, Lemma 5.2]. The existence of a secure witness encryption and null-iO, under the (sub-exponential) LWE and private-coin hiding evasive LWE assumptions, then follows from [VWW22, Theorem 5.1, Sections 6,7].

Proposition 7. *Let $\ell = \lambda^h$. With respect to Samp_{VWW} , there exists no PPT \mathcal{A} distinguishing Pre2^b in Definition 9 with non-negligible probability in λ .*

Proof. Note that the Gaussian parameter of $\mathbf{E}_{j,0}, \mathbf{E}_{j,1}$ is $\chi' \geq \lambda^{\omega(1)} \lambda^h O(n)$, implying $\chi' \geq \lambda^{\omega(1)} \ell$. Also, aux contains no information on $\mathbf{E}_{j,0}, \mathbf{E}_{j,1}$. Therefore, conditioned on aux , for $\mathbf{R} \leftarrow \mathfrak{s} \mathcal{U}(\{0,1, \dots, \ell\}^{n \times 2m})$, we have

$$\begin{aligned} \mathbf{P} + \mathbf{R} &= (\hat{\mathbf{S}}_{j,0} \mathbf{A}_j + \mathbf{E}_{j,0}, \hat{\mathbf{S}}_{j,1} \mathbf{A}_j + \mathbf{E}_{j,1}) + \mathbf{R} \bmod q \\ &\approx_s (\hat{\mathbf{S}}_{j,0} \mathbf{A}_j + \mathbf{E}_{j,0}, \hat{\mathbf{S}}_{j,1} \mathbf{A}_j + \mathbf{E}_{j,1}) \bmod q, \end{aligned}$$

where the second line is due to $(\mathbf{E}_{j,0}, \mathbf{E}_{j,1}) + \mathbf{R} \approx_s (\mathbf{E}_{j,0}, \mathbf{E}_{j,1})$ by noise flooding (Lemma 4). Therefore, $(\mathbf{P}, \text{aux}) \approx_s (\mathbf{P} + \mathbf{R} \bmod q, \text{aux})$. The claim follows. \square

7 Obfuscation-Based Counterexample

We provide an obfuscation-based counterexample against private-coin evasive LWEs in general, which applies simultaneously to variants in prior works and both the binding and hiding evasive LWE from Section 4. The idea is similar to that originally sketched in [Wee22, VWW22]; We show that upon minor tweaks, such counterexample can be readily proven assuming only LWE and the existence null-iO, the latter implied by LWE and the existence of witness encryption [GKW17, WZ17] (see also Remark 13).

We will sketch a proof of the counterexample against binding evasive LWE (Definition 8). The arguments can be easily adapted to the setting of hiding evasive LWE (Definition 9).

We start by recalling the definition of null-iO.

Definition 10 (Null-iO). An *indistinguishability obfuscation for null-circuits* (null-iO) is a PPT algorithm iO mapping circuits to circuits with the same input domain and the following properties:

Correctness. If $C : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit, then $1 - \Pr_{\tilde{C} \leftarrow \text{iO}(C, 1^\lambda)} [\forall x \in \{0, 1\}^\ell : \tilde{C}(x) = C(x)]$ is negligible.

Security. For all PPT adversaries $(\mathcal{D}, \mathcal{A})$ we have the following. If \mathcal{D} is such that on input 1^λ , with overwhelming probability, it outputs (C_0, C_1) of equal size, domain $\{0, 1\}^{\ell_\lambda}$ and such that for all $x \in \{0, 1\}^{\ell_\lambda}$, $C_0(x) = C_1(x) = 0$, then

$$|\Pr[1 = \mathcal{A}(1^\lambda, \text{iO}(C_0, 1^\lambda), C_0, C_1)] - \Pr[1 = \mathcal{A}(1^\lambda, \text{iO}(C_1, 1^\lambda), C_0, C_1)]|$$

is negligible, where the probability is over $(C_0, C_1) \leftarrow \mathcal{D}(1^\lambda)$, the random coins of iO and the random coins of \mathcal{A} .

Definition 11 (Factorisation Circuits). For a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times m}$ and control parameters n, t define the circuit $C_{\mathbf{A}, n, t}$ as follows: $C_{\mathbf{A}, n, t} : (\mathbb{Z}_q^{m \times n})^2 \rightarrow \{0, 1\}$ takes as input two matrices $\mathbf{M}_1, \mathbf{M}_2 \in \mathbb{Z}_q^{m \times n}$. If $\mathbf{A} - \mathbf{M}_1 \mathbf{M}_2^\top \bmod q \in \{-t, \dots, t\}^{m \times m}$, it outputs 1. Otherwise, it outputs 0.

Lemma 10. Let $m > n$ and $t < q$. Draw $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times m}$ uniformly at random.

1. We have $\Pr [\exists \mathbf{M}_1, \mathbf{M}_2 \in \mathbb{Z}_q^{m \times n} : \mathbf{A} - \mathbf{M}_1 \mathbf{M}_2^\top \in \{-t, \dots, t\}^{m \times m}] \leq q^{2mn} \cdot \left(\frac{2t+1}{q}\right)^{m^2}$.
2. With probability $\geq 1 - q^{-2mn} \cdot \left(\frac{2t+1}{q}\right)^{m^2}$, every PPT adversary has negligible advantage in distinguishing $\text{iO}(C_{\mathbf{A}, n, t})$ and $\text{iO}(\mathbf{0})$, where $\mathbf{0}$ is an appropriately padded circuit that always outputs 0 and is devoid of any information).

Lemma 11. Let $q > 1$ be any modulus and let $I \subseteq \mathbb{Z}_q$ be a consecutive interval. Let $\mathbf{u} \in \mathbb{Z}_q^n$ be non-zero and draw $\mathbf{r} \leftarrow \mathbb{Z}_q^n$ uniformly at random. We have

$$\Pr [\mathbf{u}^\top \cdot \mathbf{r} \bmod q \in I] \leq \frac{\#I + \|\mathbf{u}\|_\infty}{q}.$$

We provide the proofs of Lemmas 10 and 11 in Appendix B.

Attack. Consider the parameters

$$\begin{aligned} m_P &\geq 2m \log(q) + 1, & m &\geq n \log(q), \\ q &\geq 2(2t + 1), & t &= \chi' + \lambda^2 \chi^2 \cdot m, & \chi' &\geq \lambda^{\omega(1)} \chi, & \chi &\geq 2\sqrt{n}. \end{aligned}$$

On input 1^λ , the sampler proceeds as follows: It draws

$$\mathbf{S} \leftarrow \mathbb{Z}_q^{m_P \times n}, \quad \mathbf{P} \leftarrow \mathbb{Z}_q^{n \times m_P}, \quad \mathbf{E}'' \leftarrow \{-\chi', \dots, \chi'\}^{m_P \times m_P}$$

and outputs \mathbf{S}, \mathbf{P} and as auxiliary information the obfuscated circuit

$$\text{aux} = \text{iO}(C_{\mathbf{SP} + \mathbf{E}'', m, t}).$$

Breaking the Post-Challenge. The post adversary is given $\mathbf{B}, \mathbf{P}, \mathbf{U}, \text{aux}$ where \mathbf{U} is sampled from $D_{\mathbb{Z}^{m \times m_P}, \chi}$ conditioned on

$$\mathbf{BU} = \mathbf{P} \bmod q.$$

It has to distinguish $\mathbf{R} = \mathbf{SB} + \mathbf{E} \bmod q$ from uniform randomness $\mathbf{R} \leftarrow \mathbb{Z}_q^{m_P \times m}$. It inputs $\mathbf{R}, \mathbf{U}^\top \in \mathbb{Z}_q^{m_P \times m}$ into the circuit $\text{aux} = \text{iO}(C_{\mathbf{PS} + \mathbf{E}'', m, t})$. If $\mathbf{R} = \mathbf{SB} + \mathbf{E}$, we have

$$\mathbf{SP} + \mathbf{E}'' - \mathbf{RU} = \mathbf{SP} + \mathbf{E}'' - (\mathbf{SB} + \mathbf{E})\mathbf{U} = \mathbf{E}'' - \mathbf{EU} \bmod q.$$

With overwhelming probability $\geq 1 - 4m_P^2 \exp(-\lambda^2)$, the largest entries of \mathbf{E} and \mathbf{U} are bounded by $\chi\lambda$. Since the entries of \mathbf{E}'' are bounded by χ' , the entries of $\mathbf{E}'' - \mathbf{E}\mathbf{U}$ are bounded by $t = \chi' + \lambda^2\chi^2 \cdot m$ with overwhelming probability. In this case, $\text{iO}(C_{\mathbf{SP}+\mathbf{E}'',m,t})$ outputs 1.

On the other hand, if \mathbf{R} is uniformly random, we claim that $\mathbf{SP} + \mathbf{E}'' - \mathbf{R}\mathbf{U} \bmod q$ does not lie in $\{-t, \dots, t\}^{m \times m}$ with probability $\geq 1 - \frac{2t+1+\lambda\chi}{q}$. Denote by $\mathbf{r} \in \mathbb{Z}_q^m$ the first row of \mathbf{R} and by $\mathbf{u} \in \mathbb{Z}_q^m$ the first column of \mathbf{U} . The first entry of $\mathbf{R}\mathbf{U}$ is given by $\mathbf{r}^T \mathbf{u}$. With overwhelming probability, the entries of \mathbf{u} are bounded by $\lambda\chi$. Hence, Lemma 11 implies that

$$\Pr[\mathbf{r}^T \mathbf{u} \bmod q \in \{z-t, \dots, z+t\}] \leq \frac{2t+1}{q} + \frac{\lambda\chi}{q},$$

where z denotes the first entry of $\mathbf{SP} + \mathbf{E}'' \bmod q$. It follows $\text{iO}(C_{\mathbf{SP}+\mathbf{E}'',m,t})$ outputs 0 with overwhelming probability if \mathbf{R} is uniformly random.

In conclusion, the adversary has a distinguishing advantage of $\geq 1 - \frac{2t+1+\lambda\chi}{q} - \text{negl}(\lambda)$.

On the Security of the Pre-Challenge. The a priori adversary is given $\mathbf{B}, \mathbf{P}, \text{aux} = \text{iO}(C_{\mathbf{SP}+\mathbf{E}'',m,t})$ and is tasked with distinguishing $\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}'$ from uniform randomness. We prove indistinguishability by a sequence of hybrids:

$$\begin{aligned} & (\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{P}, \mathbf{SP} + \mathbf{E}', \text{iO}(C_{\mathbf{SP}+\mathbf{E}'',m,t})) \\ & \stackrel{\text{flood}}{\approx} (\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{P}, \mathbf{SP} + \mathbf{E}', \text{iO}(C_{\mathbf{SP}+\mathbf{E}'+\mathbf{E}'',m,t})) \\ & \stackrel{\text{LWE}}{\approx} (\mathbf{B}, \mathbf{R}_1, \mathbf{P}, \mathbf{R}_2, \text{iO}(C_{\mathbf{R}_2+\mathbf{E}'',m,t})) \\ & \stackrel{\text{iO}}{\approx} (\mathbf{B}, \mathbf{R}_1, \mathbf{P}, \mathbf{R}_2, \text{iO}(\mathbf{0})) \\ & \stackrel{\text{iO}}{\approx} (\mathbf{B}, \mathbf{R}_1, \mathbf{P}, \mathbf{R}_2, \text{iO}(C_{\mathbf{R}_3+\mathbf{E}'',m,t})) \\ & \stackrel{\text{LWE}}{\approx} (\mathbf{B}, \mathbf{R}_1, \mathbf{P}, \mathbf{R}_2, \text{iO}(C_{\mathbf{SP}+\mathbf{E}'+\mathbf{E}'',m,t})) \\ & \stackrel{\text{flood}}{\approx} (\mathbf{B}, \mathbf{R}_1, \mathbf{P}, \mathbf{R}_2, \text{iO}(C_{\mathbf{SP}+\mathbf{E}'',m,t})) \end{aligned}$$

By a noise-flooding argument with respect to uniform distributions over bounded sets, we can replace the hard-wired $\mathbf{SP} + \mathbf{E}'' \bmod q$ in the obfuscated circuit by $\mathbf{SP} + \mathbf{E}'' + \mathbf{E}' \bmod q$ and only incur a negligible statistical difference. By LWE, we can replace $\mathbf{SB} + \mathbf{E} \bmod q$ and $\mathbf{SP} + \mathbf{E}' \bmod q$ by random matrices $\mathbf{R}_1 \leftarrow_{\$} \mathbb{Z}_q^{m_P \times m}$, $\mathbf{R}_2 \leftarrow_{\$} \mathbb{Z}_q^{m_P \times m_P}$. Now, we invoke null-iO security twice: first, with overwhelming probability, $C_{\mathbf{R}_2+\mathbf{E}'',m,t}$ outputs always zero, hence, we can replace aux by $\text{iO}(\mathbf{0})$. With the same argument, we can then replace $\text{iO}(\mathbf{0})$ by $\text{iO}(C_{\mathbf{R}_3+\mathbf{E}'',m,t})$ for an independent uniformly random matrix $\mathbf{R}_3 \leftarrow_{\$} \mathbb{Z}_q^{m \times m}$. We apply LWE again, and replace \mathbf{R}_3 by $\mathbf{SP} + \mathbf{E}' \bmod q$ in $C_{\mathbf{R}_3+\mathbf{E}'',m,t}$. Finally, by flooding, we can remove \mathbf{E}' and get the same auxiliary information $\text{iO}(C_{\mathbf{SP}+\mathbf{E}'',m,t})$ in the joint distribution again.

By using the flooding lemma, we incur a statistical difference of twice

$$\frac{\chi}{\chi'} = \lambda^{-\omega(1)}.$$

To use the security of null-iO, we need that $\mathbf{R}_2 + \mathbf{E}''$ and $\mathbf{R}_3 + \mathbf{E}''$ do not admit an approximated factorisation. Since both matrices are uniformly random, this probability is bounded by

$$\leq \frac{q^{2m_P m} \cdot (2t+1)^{m_P^2}}{q^{m_P^2}} = \left(q^{2m} \cdot \left(\frac{2t+1}{q} \right)^{m_P} \right)^{m_P} \leq (q^{2m} \cdot 2^{-m_P})^{m_P} \leq (2^{-1})^{m_P} = 2^{-m_P}.$$

This proves the hardness of the a priori challenge against PPT adversaries.

Remark 11 (Applicability to Hiding Evasive LWE). It is easy to see that the above counterexample also applies to the private-coin hiding evasive LWE (Definition 9). On the one hand, the above distinguisher of then does not rely on the knowledge of \mathbf{B}, \mathbf{P} at all. On the other hand, the if condition of private-coin hiding evasive LWE requires that \mathbf{P} hides its lower bits in the presence of aux , which can be argued in a similar way as above: one can argue that $(\mathbf{P}, \text{aux}) = (\mathbf{P}, \text{iO}_0(C_{\mathbf{P}\mathbf{S}+\mathbf{E}''}))$ is indistinguishable from $(\mathbf{P}, \text{iO}(\mathbf{0}))$. Since \mathbf{P} is uniformly random, it is now apparent that $(\mathbf{P}, \text{iO}(\mathbf{0}))$ is equally distributed as $(\mathbf{P} + \mathcal{U}(\{0, 1, \dots, \ell\}^{n \times m_P}), \text{iO}(\mathbf{0}))$, which is indistinguishable from $(\mathbf{P} + \mathcal{U}(\{0, 1, \dots, \ell\}^{n \times m_P}), \text{iO}(C_{\mathbf{P}\mathbf{S}+\mathbf{E}''}))$.

Remark 12 (Applicability to Variation of Binding Evasive LWE in [BDJ⁺24]). In [BDJ⁺24], a variant of private-coin binding evasive LWE that is more conservative than Definition 8 has been put forth. In their definition, **Samp** is not allowed to generate **P** on its own. Instead, it takes **P** as input (which is indistinguishable from uniform randomness in the view of the sampler) and outputs only **S** and **aux**. Since the above counterexample samples **P** uniformly at random, **Samp** can easily be adapted to satisfy the restrictions of [BDJ⁺24].

Remark 13 (Assuming Witness Encryption instead of Null-iO). Null-iO can be constructed from witness encryption when additionally assuming LWE. Concretely, Goyal, Koppula, and Waters [GKW17] construct lockable encryption from LWE and then show that lockable encryption and witness encryption together imply null-iO. In a concurrent work, Wichs and Zirdelis [WZ17] build obfuscators for compute-and-compare programs from LWE and similarly show that these, together with witness encryption, imply null-iO.

Thus, the above obfuscation-based counterexample can be proven from assuming only the hardness of LWE and the existence of witness encryption. Interestingly, constructions of witness encryption are known under several different variants of private-coin evasive LWE. As mentioned, the witness encryption (and null-iO) of Vaikuntanathan, Wee and Wichs [VWW22] can be proven assuming private-coin hiding evasive LWE. More recently, Branco, Döttling, Jain, Malavolta, Mathialagan, Peters, and Vaikuntanathan [BDJ⁺24] also construct witness encryption from pseudorandom obfuscation (PRO) assuming LWE and a variant of private-coin binding evasive LWE. However, it should be emphasised that PRO is impossible to exist, as the authors point out themselves.

Remark 14 (On Auxiliary Information). The above counterexample provides the obfuscated program **iO** in the auxiliary information **aux**. The same idea would equally apply (up to changing the assumption's dimension parameters) when **iO** is instead embedded in the target image matrix, for example, **Samp** outputs $\mathbf{P}' = \left(\mathbf{P}, \begin{pmatrix} \mathbf{R} \\ \mathbf{iO} \end{pmatrix} \right)$ where **R** is uniformly random, i.e. right-extend the target matrix and let **iO** be embedded as the bottom chunk of the extension, and **aux** = \perp is empty. That $\mathbf{S} \begin{pmatrix} \mathbf{R} \\ \mathbf{iO} \end{pmatrix} \bmod q$ (with or without error) is randomly distributed is obvious. The distinguisher of **then** simply ignores the addition preimages w.r.t. the extension. Overall, the (im)plausibility of private-coin evasive LWE seems to only loosely depend on the complexity of **aux**, but more critically on the mere fact the **Samp** is private-coin.

8 Counterexample against Non-Uniform Matrix **B**

This section outlines a simple counterexample against evasive LWE with non-uniformly random matrix **B**, which we discovered subsequently to publishing the proceedings version. This counterexample applies to all families Definitions 7 to 9 introduced in Section 4. Since allowing $\mathbf{B} \leftarrow \mathcal{D}$ to follow arbitrary public distribution \mathcal{D} is insecure, it seems advisable to use evasive LWE only with uniform **B**, until the role of the matrix **B** distribution is better understood.

Counterexample. Consider $\mathbf{B} = \begin{pmatrix} \mathbf{B}_0 & \\ & \mathbf{B}_1 \end{pmatrix}$ to be a 2×2 block diagonal matrix, where $\mathbf{B}_0, \mathbf{B}_1 \leftarrow \mathbb{Z}_q^{n \times m}$ are independent and uniformly random. Let $\mathbf{s} \leftarrow \mathbb{Z}_q^{2n}$ be uniformly random. Let **Samp**(1^λ) output

$$\mathbf{P} = \begin{pmatrix} \mathbf{P}_0 & \\ & \mathbf{P}_1 \end{pmatrix}, \quad \mathbf{aux} = \perp,$$

that is, **P** is also 2×2 block diagonal, and where $\mathbf{P}_0, \mathbf{P}_1 \leftarrow \mathbb{Z}_q^{n \times m_P}$ are independent and uniformly random.

We can directly show the hardness of the if challenge. Indeed, writing $\mathbf{s} = (\mathbf{s}_0, \mathbf{s}_1) \in \mathbb{Z}_q^n \times \mathbb{Z}_q^n$, we have

$$\begin{aligned} (\mathbf{s}^\top \mathbf{B} + \mathbf{e}_B^\top, \quad \mathbf{s}^\top \mathbf{P} + \mathbf{e}_P^\top) &= (\mathbf{s}_0^\top \mathbf{B}_0 + \mathbf{e}_{B,0}^\top, \quad \mathbf{s}_1^\top \mathbf{B}_1 + \mathbf{e}_{B,1}^\top, \quad \mathbf{s}_0^\top \mathbf{P}_0 + \mathbf{e}_{P,0}^\top, \quad \mathbf{s}_1^\top \mathbf{P}_1 + \mathbf{e}_{P,1}^\top) \bmod q \\ &\approx_c (\mathbf{c}_0^\top, \quad \mathbf{c}_1^\top, \quad \mathbf{d}_0^\top, \quad \mathbf{d}_1^\top) \bmod q \end{aligned}$$

where in the first line we express the error $\mathbf{e}_B^\top = (\mathbf{e}_{B,0}^\top, \mathbf{e}_{B,1}^\top)$, similarly for \mathbf{e}_P^\top , and in the second line we invoke $\text{LWE}_{\mathbb{Z}, n, m+m_P, q, \chi}$ twice to swap both $\mathbf{s}_0^\top(\mathbf{B}_0, \mathbf{P}_0) + (\mathbf{e}_{B,0}^\top, \mathbf{e}_{P,0}^\top) \bmod q$ and $\mathbf{s}_1^\top(\mathbf{B}_1, \mathbf{P}_1) + (\mathbf{e}_{B,1}^\top, \mathbf{e}_{P,1}^\top) \bmod q$ to uniformly random.

To show that the then condition does not hold, we observe that a preimage matrix \mathbf{U} satisfies

$$\begin{pmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{pmatrix} \underbrace{\begin{pmatrix} \mathbf{U}_{00} & \mathbf{U}_{01} \\ \mathbf{U}_{10} & \mathbf{U}_{11} \end{pmatrix}}_{\mathbf{U}} = \begin{pmatrix} \mathbf{P}_0 \\ \mathbf{P}_1 \end{pmatrix} \bmod q$$

where we pay attention to that $\mathbf{U}_{01}, \mathbf{U}_{10}$ are non-zero with overwhelming probability, by a standard entropy argument of the Gaussian distribution. Therefore we obtain

$$\begin{aligned} \mathbf{B}_0 \mathbf{U}_{01} &= \mathbf{0} \pmod q, \\ \mathbf{B}_1 \mathbf{U}_{10} &= \mathbf{0} \pmod q. \end{aligned}$$

This allows to recover an Ajtai-trapdoor of both $\mathbf{B}_0, \mathbf{B}_1$, allowing to distinguish both $\mathbf{s}_0^\top \mathbf{B}_0 + \mathbf{e}_{B,0}^\top \bmod q$ and $\mathbf{s}_1^\top \mathbf{B}_1 + \mathbf{e}_{B,1}^\top \bmod q$ from uniform randomness.

It is easy to see that this counterexample can be generalised to, for example, \mathbf{B}, \mathbf{P} being $k \times k$ block diagonal matrices.

Discussion. The above counterexample has two direct consequences. First, our original claim that the proposed definitions capture all prior/concurrent evasive LWE variants is false. In particular, for the public-coin family (Definition 7), unless the instantiation yields $\mathbf{B}, \mathbf{P}, \mathbf{U}$ all being block diagonal matrices, which happens with negligible probability, Definition 7 is unable to capture the public-coin evasive LWE variant of [WWW22, CLW24], where the assumption involves multiple LWE samples $\mathbf{s}_i^\top \mathbf{B}_i + \mathbf{e}_i^\top \bmod q$ in the joint distributions. We do not know an alternative way to capture this type of variants, except by directly stating them individually in the assumption, as they were also stated in [WWW22, CLW24].

Second, the above counterexample shows that our proposal of allowing $\mathbf{B} \leftarrow \mathcal{D}$ to follow arbitrary public distribution \mathcal{D} is insecure in its full generality. Fortunately, all prior works have adopted uniformly distributed \mathbf{B} where such counterexample clearly does not apply. At the moment, we do not have a comprehensive understanding of how different choices of the distribution of the matrix \mathbf{B} affect the hardness of evasive LWE. Concretely, it is easy to come up with alternative non-uniform distributions of \mathbf{B} where no similarly trivial counterexample seems to apply, but at the same time, similarly trivial counterexamples still apply to slight variations of the block-diagonal case. At the moment, we are unable to meaningfully characterise distributions which do not suffer from similar counterexamples. Thus, at the moment, we advise only using evasive LWE with uniformly random \mathbf{B} , as it has been in prior works.

Acknowledgements

The authors thank the anonymous reviewers for insightful comments which very much improved this work, in particular, sharing with us the counterexamples against a prior version of Hiding Evasive LWE, and against public-coin Evasive LWE when the sampler inputs \mathbf{B} . Chris Brzuska and Ivy K. Y. Woo are supported by Research Council of Finland grant 358950. We thank Russell W. F. Lai and Hoeteck Wee for helpful discussions.

References

- ACFP14. Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, and Ludovic Perret. Algebraic algorithms for LWE. Cryptology ePrint Archive, Report 2014/1018, 2014. C.2
- AG11. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP 2011, Part I*, volume 6755 of *LNCS*, pages 403–415. Springer, Berlin, Heidelberg, July 2011. 8, C.2
- AKY24. Shweta Agrawal, Simran Kumari, and Shota Yamada. Attribute based encryption for turing machines from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 352–386. Springer, Cham, August 2024. 1, 2.1, 2.3, 2.4, 2.4, 4.2

- ARYY23. Shweta Agrawal, Mélissa Rossi, Anshu Yadav, and Shota Yamada. Constant input attribute based (and predicate) encryption from evasive and tensor LWE. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 532–564. Springer, Cham, August 2023. 1, 1.1, 1.1, 2.1, 2.3, 2.4, 2.4, 4.2
- BD20. Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 551–575. Springer, Cham, May 2020. 4, 4.1
- BDE⁺18. Jonathan Bootle, Claire Delaplace, Thomas Espitau, Pierre-Alain Fouque, and Mehdi Tibouchi. LWE without modular reduction and improved side-channel attacks against BLISS. *Cryptology ePrint Archive*, Report 2018/822, 2018. 3.1
- BDJ⁺24. Pedro Branco, Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Pseudorandom obfuscation and applications. *Cryptology ePrint Archive*, Paper 2024/1742, 2024. 12, 13
- BDK⁺11. Boaz Barak, Yevgeniy Dodis, Hugo Krawczyk, Olivier Pereira, Krzysztof Pietrzak, François-Xavier Standaert, and Yu Yu. Leftover hash lemma, revisited. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 1–20. Springer, Berlin, Heidelberg, August 2011. 3
- CKPS00. Nicolas Courtois, Alexander Klimov, Jacques Patarin, and Adi Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 392–407. Springer, Berlin, Heidelberg, May 2000. C.2
- CLW24. Valerio Cini, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based multi-authority/client attribute-based encryption for circuits. *CiC*, 4, 2024. To appear. 2.3, 4.1, 4.1, 8
- CM24. Yilei Chen and Xinyu Mao. Universal computational extractors from lattice assumptions. *Cryptology ePrint Archive*, 2024. <https://ia.cr/2024/225>. 2.1, 2.4
- DBM⁺08. Jintai Ding, Johannes Buchmann, Mohamed Saied Emam Mohamed, Wael Said Abd Elmageed Mohamed, and Ralf-Philipp Weinmann. Mutantxl. *SCC*, (TUD-CS-2009-0142):16–22, 01 2008. C.2
- DL78. Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978. A.2
- GGH15. Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Berlin, Heidelberg, March 2015. 2.1
- GKW17. Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In Chris Umans, editor, *58th FOCS*, pages 612–621. IEEE Computer Society Press, October 2017. 7, 13
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 1, 3.1, A.1, 14, 15
- HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. 3
- HLL23. Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 415–434. IEEE, 2023. 6, 4.2, D, D
- HLL24. Yao-Ching Hsieh, Huijia Lin, and Ji Luo. A general framework for lattice-based abe using evasive inner-product functional encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 433–464. Springer, 2024. 1, 2.1, 2.3, 4.1, 4.1
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. 3.2, 1
- MP13. Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 21–39. Springer, Berlin, Heidelberg, August 2013. C.2
- MPV24. Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge snarks for up. *Cryptology ePrint Archive*, 2024. <https://ia.cr/2024/227>. 2.1, 2.4
- MR04. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th FOCS*, pages 372–381. IEEE Computer Society Press, October 2004. 5, 2, A.1
- NMSÚ24. Miguel Cueto Noval, Simon-Philipp Merz, Patrick Stählin, and Akin Ünäl. On the soundness of algebraic attacks against code-based assumptions. 2024. 8, C.2, C.2
- Pei07. Chris Peikert. Limits on the hardness of lattice problems in lp norms. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC07)*, pages 333–346, 2007. 1, 3.1, A.1, 13
- PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Berlin, Heidelberg, March 2006. 2
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. 10, A.1, 12

- Sal23. Flavio Salizzoni. An upper bound for the solving degree in terms of the degree of regularity, 2023. C.2
- Sch80. J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 10 1980. A.2
- STA20. Chao Sun, Mehdi Tibouchi, and Masayuki Abe. Revisiting the hardness of binary error LWE. In Joseph K. Liu and Hui Cui, editors, *ACISP 20*, volume 12248 of *LNCS*, pages 425–444. Springer, Cham, November / December 2020. C.2
- Ste24. Matthias Johann Steiner. The complexity of algebraic algorithms for LWE. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part III*, volume 14653 of *LNCS*, pages 375–403. Springer, Cham, May 2024. C.2, C.2
- Tsa22. Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 535–559. Springer, Cham, August 2022. 1, 1.1, 1.1, 2.1, 2.4
- VWW22. Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 195–221. Springer, Cham, December 2022. 1, 1, 1.1, 1.1, 2.1, 2.1, 2.1, 2, 3, 2.4, 4.2, 4.3, 7, 6, 7, 13
- Wee22. Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Cham, May / June 2022. 1, 3, 3, 2.3, 2.3, 4.1, 4.1, 7
- Wee24. Hoeteck Wee. Circuit ABE with $\text{poly}(\text{depth}, \lambda)$ -sized ciphertexts and keys from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 178–209. Springer, Cham, August 2024. 2.3, 4.1
- WWW22. Brent Waters, Hoeteck Wee, and David J. Wu. Multi-authority ABE from lattices without random oracles. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 651–679. Springer, Cham, November 2022. 1, 2.1, 2.3, 4.1, 4.1, 8
- WZ17. Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In Chris Umans, editor, *58th FOCS*, pages 600–611. IEEE Computer Society Press, October 2017. 7, 13
- Zip79. Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation*, pages 216–226, Berlin, Heidelberg, 1979. Springer Berlin Heidelberg. A.2

A On Regular Preimages

We will prove here the following theorem:

Theorem 2. *Let $\chi \in \omega(\sqrt{n})$, $q > 2 \cdot \chi$ be prime and $m \geq 2n \cdot \log q$.*

Draw $\mathbf{B}, \mathbf{P} \leftarrow \mathbb{Z}_q^{n \times m}$ uniformly at random. Draw $\mathbf{U} \leftarrow D_{\mathbb{Z}^m \times m, \chi}$ conditioned on $\mathbf{BU} = \mathbf{P} \bmod q$. Then, we have

$$\Pr[\det(\mathbf{U}) = 0 \bmod q] \leq e^\pi \cdot \frac{m}{2\lfloor \chi \rfloor + 1} + n \cdot 4 \exp\left(-\pi \frac{(q-1)^2}{4\chi^2}\right) + 2q^{-n} + 2m \cdot 2^{-n}.$$

To prove Theorem 2, we will proceed in several steps: Appendix A.1 revisits useful facts about lattices and proves that sampling a *normal* Gaussian vector $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$ and sampling a Gaussian preimage \mathbf{e}' under \mathbf{B} for a random point $\mathbf{u} \in \mathbb{Z}_q^n$ is statistically close. This implies that \mathbf{U} is close to being sampled from $D_{\mathbb{Z}^m \times m, \chi}$ without conditioning its distribution on $\mathbf{BU} = \mathbf{P}$. In Appendix A.2, we revisit the Schwartz-Zippel lemma and prove a variant of it for points $\mathbf{x} \leftarrow D_{\mathbb{Z}^m, \chi}$ of spherical discrete Gaussian distribution. In Appendix A.3, we conclude the proof by observing that we can bound the probability of the determinant vanishing on \mathbf{U} by our variant of the Schwartz-Zippel lemma.

A.1 More Lattice Preliminaries

We now revisit additional useful notions and facts from lattice theory and prove several corollaries. Recall that the smoothing parameter [MR04] for a full-rank lattice $\Lambda \subset \mathbb{R}^m$ and $\epsilon > 0$ is given by

$$\eta_\epsilon(\Lambda) := \inf\{\chi > 0 \mid \rho_{1/\chi}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon\}.$$

Further, denote by $\lambda_1^\infty(\Lambda)$ the infinity-norm of the shortest non-zero vector of $\Lambda \subset \mathbb{R}^m$, i.e.,

$$\lambda_1^\infty(\Lambda) := \min_{\mathbf{x} \in \Lambda \setminus \{\mathbf{0}\}} \|\mathbf{x}\|_\infty = \min_{(x_1, \dots, x_m) \in \Lambda \setminus \{\mathbf{0}\}} \left(\max_{i=1, \dots, m} |x_i| \right).$$

We import the following lemmas from [Reg05, Pei07, GPV08]:

Lemma 12 ([Reg05]). *Let $m \geq 2n \cdot \log q$. If we draw $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, then we have with probability at least $1 - q^{-n}$*

$$\{\mathbf{A} \cdot \mathbf{e} \bmod q \mid \mathbf{e} \in \{0, 1\}^m\} = \mathbb{Z}_q^n.$$

Lemma 13 ([Pei07]). *Let $\Lambda \subset \mathbb{R}^m$ be a full-rank lattice and $\epsilon > 0$. We have*

$$\eta_\epsilon(\Lambda) \leq \frac{\sqrt{\log(2m) + \log(1 + 1/\epsilon)}}{\lambda_1^\infty(\Lambda^*) \cdot \sqrt{\pi}}.$$

Lemma 14 ([GPV08]). *Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be s.t. $\mathbf{A} \cdot \mathbb{Z}_q^m = \mathbb{Z}_q^n$. Let $\epsilon \in (0, 0.5)$, $\chi \geq \eta_\epsilon(\Lambda^\perp(\mathbf{A}))$.*

1. *For $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$ and $\mathbf{u} \leftarrow \mathbb{Z}_q^n$, the statistical distance between $(\mathbf{A}, \mathbf{Ae} \bmod q)$ and (\mathbf{A}, \mathbf{u}) is bounded by 2ϵ .*
2. *Fix $\mathbf{t} \in \mathbb{Z}^m$ and set $\mathbf{u} = \mathbf{At} \bmod q$. Draw $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$ conditioned on $\mathbf{Ae} = \mathbf{u}$. Then, $\mathbf{e} - \mathbf{t}$ is distributed according to $D_{\Lambda^\perp(\mathbf{A}), \chi, -\mathbf{t}}$.*

Lemma 15 ([GPV08]). *Let q be prime, $m \geq 2n \cdot \log q$. Draw $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ uniformly at random. Then, we have with probability at least $1 - q^{-n}$*

$$\lambda_1^\infty(\Lambda(\mathbf{A})) \geq \frac{q}{4}.$$

We now prove the following:

Lemma 16. *Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be s.t. $\mathbf{A} \cdot \mathbb{Z}_q^m = \mathbb{Z}_q^n$. Let $\epsilon \in (0, 0.5)$, $\chi \geq \eta_\epsilon(\Lambda^\perp(\mathbf{A}))$.*

Draw $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$, set $\mathbf{u} = \mathbf{Ae} \bmod q$ and draw $\mathbf{e}' \leftarrow D_{\mathbb{Z}^m, \chi}$ conditioned on $\mathbf{Ae}' = \mathbf{u} \bmod q$. Then, \mathbf{e}' (without seeing \mathbf{e}) is distributed according to $D_{\mathbb{Z}^m, \chi}$.

Proof. For a lattice Λ and a vector \mathbf{x} , set

$$\mathbf{1}_{\mathbf{x} \in \Lambda} = \begin{cases} 1, & \mathbf{x} \in \Lambda, \\ 0, & \mathbf{x} \notin \Lambda. \end{cases}$$

Because of the second claim of Lemma 14, \mathbf{e}' is distributed according to

$$\mathbf{e} + D_{\Lambda^\perp(\mathbf{A}), \chi, -\mathbf{e}}.$$

Now, fix $\mathbf{x} \in \mathbb{Z}^m$. We will compute the probability that \mathbf{e}' equals \mathbf{x} :

$$\begin{aligned} \Pr[\mathbf{e}' = \mathbf{x}] &= \sum_{\mathbf{t} \in \mathbb{Z}^m} \Pr[\mathbf{e} = \mathbf{t}] \cdot \Pr[\mathbf{x} \leftarrow \mathbf{t} + D_{\Lambda^\perp(\mathbf{A}), \chi, -\mathbf{t}}] \\ &= \sum_{\mathbf{t} \in \mathbb{Z}^m} \frac{\rho_\chi(\mathbf{t})}{\rho_\chi(\mathbb{Z}^m)} \cdot \frac{\rho_{\chi, -\mathbf{t}}(\mathbf{x} - \mathbf{t}) \cdot \mathbf{1}_{\mathbf{x} - \mathbf{t} \in \Lambda^\perp(\mathbf{A})}}{\rho_{\chi, -\mathbf{t}}(\Lambda^\perp(\mathbf{A}))} \\ &= \sum_{\mathbf{t} \in \mathbb{Z}^m} \frac{\rho_\chi(\mathbf{t})}{\rho_\chi(\mathbb{Z}^m)} \cdot \frac{\rho_\chi(\mathbf{x}) \cdot \mathbf{1}_{\mathbf{x} - \mathbf{t} \in \Lambda^\perp(\mathbf{A})}}{\rho_\chi(\mathbf{t} + \Lambda^\perp(\mathbf{A}))} \\ &= \frac{\rho_\chi(\mathbf{x})}{\rho_\chi(\mathbb{Z}^m)} \cdot \sum_{\mathbf{t} \in \mathbb{Z}^m} \frac{\rho_\chi(\mathbf{t}) \cdot \mathbf{1}_{\mathbf{x} - \mathbf{t} \in \Lambda^\perp(\mathbf{A})}}{\rho_\chi(\mathbf{t} + \Lambda^\perp(\mathbf{A}))}. \end{aligned}$$

We claim that $\sum_{\mathbf{t} \in \mathbb{Z}^m} \frac{\rho_\chi(\mathbf{t}) \cdot \mathbf{1}_{\mathbf{x} - \mathbf{t} \in \Lambda^\perp(\mathbf{A})}}{\rho_\chi(\mathbf{t} + \Lambda^\perp(\mathbf{A}))}$ is one. Indeed, we have

$$\begin{aligned} &\sum_{\mathbf{t} \in \mathbb{Z}^m} \frac{\rho_\chi(\mathbf{t}) \cdot \mathbf{1}_{\mathbf{x} - \mathbf{t} \in \Lambda^\perp(\mathbf{A})}}{\rho_\chi(\mathbf{t} + \Lambda^\perp(\mathbf{A}))} = \sum_{\mathbf{t} \in \Lambda^\perp(\mathbf{A}) + \mathbf{x}} \frac{\rho_\chi(\mathbf{t})}{\rho_\chi(\mathbf{t} + \Lambda^\perp(\mathbf{A}))} \\ &= \sum_{\mathbf{t} \in \Lambda^\perp(\mathbf{A}) + \mathbf{x}} \frac{\rho_\chi(\mathbf{t})}{\rho_\chi(\mathbf{x} + \Lambda^\perp(\mathbf{A}))} = \frac{\rho_\chi(\Lambda^\perp(\mathbf{A}) + \mathbf{x})}{\rho_\chi(\mathbf{x} + \Lambda^\perp(\mathbf{A}))} = 1. \end{aligned}$$

Hence, we have $\Pr[\mathbf{e}' = \mathbf{x}] = \frac{\rho_\chi(\mathbf{x})}{\rho_\chi(\mathbb{Z}^m)} = D_{\mathbb{Z}^m, \chi}(\mathbf{x})$. \square

We can now show the following:

Corollary 1. *Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be s.t. $\mathbf{A} \cdot \mathbb{Z}_q^m = \mathbb{Z}_q^n$. Let $\varepsilon \in (0, 0.5)$, $\chi \geq \eta_\varepsilon(\Lambda^\perp(\mathbf{A}))$.*

Let D_1 be the distribution that samples and outputs $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$. Let D_2 be the distribution that first samples $\mathbf{u} \leftarrow \mathbb{Z}_q^n$ and outputs $\mathbf{e}' \leftarrow D_{\mathbb{Z}^m, \chi}$ conditioned on $\mathbf{A}\mathbf{e} = \mathbf{u} \bmod q$.

Then, the statistical distance between D_1 and D_2 is bounded by 2ε . Concretely, we have

$$\Delta((\mathbf{A}, \mathbf{e}), (\mathbf{A}, \mathbf{e}')) \leq 2\varepsilon.$$

Proof. The first part of Lemma 14 implies that the statistical distance between $\mathbf{u} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{A}\mathbf{e}$ for $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$ is bounded by 2ε .

Hence, we can replace D_2 by the distribution D_3 that samples $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$, sets $\mathbf{u} = \mathbf{A}\mathbf{e} \bmod q$ and outputs $\mathbf{e}' \leftarrow D_{\mathbb{Z}^m, \chi}$ conditioned on $\mathbf{A}\mathbf{e}' = \mathbf{u} \bmod q$. Lemma 16 implies now that D_1 and D_3 are identical. \square

Corollary 2. *Let $\chi \in \omega(\sqrt{n})$, let q be prime. Draw $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ uniformly at random. With probability $\geq 1 - 2q^{-n}$ over the randomness of \mathbf{A} , we have*

$$\Delta((\mathbf{A}, \mathbf{e}), (\mathbf{A}, \mathbf{e}')) \leq 2^{-n+1}$$

where $\mathbf{e} \leftarrow D_{\mathbb{Z}^m, \chi}$, $\mathbf{u} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{e}' \leftarrow D_{\mathbb{Z}^m, \chi}$ conditioned on $\mathbf{A}\mathbf{e}' = \mathbf{u} \bmod q$.

Proof. Set $\varepsilon = 2^{-n}$. Lemma 15 implies for $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$

$$\Pr \left[\lambda_1^\infty(\Lambda(\mathbf{A})) \geq \frac{q}{4} \right] \geq 1 - q^{-n}.$$

Note that we have

$$\lambda_1^\infty(\Lambda(\mathbf{A})) = \lambda_1^\infty(q \cdot \Lambda^\perp(\mathbf{A})^*) = q \cdot \lambda_1^\infty(\Lambda^\perp(\mathbf{A})^*).$$

Hence, $\lambda_1^\infty(\Lambda^\perp(\mathbf{A})^*)$ is bounded from below by $1/4$ with overwhelming probability.

According to Lemma 13, the smoothing parameter of $\Lambda^\perp(\mathbf{A})$ is bounded by

$$\eta_\epsilon(\Lambda^\perp(\mathbf{A})) \leq \frac{\sqrt{\log(2m) + \log(1 + 2^n)}}{\lambda_1^\infty(\Lambda^\perp(\mathbf{A})^*) \cdot \sqrt{\pi}} \in O\left(\frac{\sqrt{n}}{\lambda_1^\infty(\Lambda^\perp(\mathbf{A})^*)}\right).$$

Hence, for almost all n , with probability at least $1 - q^{-n}$, $\chi \in \omega(\sqrt{n})$ will be larger than $\eta_\epsilon(\Lambda^\perp(\mathbf{A})) \in O(\sqrt{n})$.

Because of Lemma 12, \mathbf{A} generates \mathbb{Z}_q^n with probability at least $1 - q^{-n}$. If both conditions are satisfied

$$\chi \geq \eta_\epsilon(\Lambda^\perp(\mathbf{A})) \quad \text{and} \quad \mathbf{A} \cdot \mathbb{Z}_q^m = \mathbb{Z}_q^n,$$

then Corollary 1 implies

$$\Delta((\mathbf{A}, \mathbf{e}), (\mathbf{A}, \mathbf{e}')) \leq 2\epsilon = 2^{-n+1}.$$

Both conditions will hold simultaneously with probability at least $1 - 2q^{-n}$. Hence, the claim follows. \square

A.2 Schwartz-Zippel over Spherical Discrete Gaussians

We now introduce a variant of the seminal Schwartz-Zippel lemma [DL78, Zip79, Sch80] over the spherical discrete Gaussian distribution $D_{\mathbb{Z}^m, \chi}$. We use this variation to prove that a square matrix with entries distributed according to a (spherical) discrete Gaussian distribution has full rank with high probability, if the variance of the distribution is large enough.

We first show the following variant of the Schwartz-Zippel lemma for any discrete distribution:

Lemma 17 (Schwartz-Zippel for Any Distributions). *Let \mathcal{X} be any discrete distribution over any field \mathcal{K} . Fix a non-zero polynomial $f \in \mathcal{K}[X_1, \dots, X_n]$. We have*

$$\Pr_{\mathbf{x} \leftarrow \mathcal{X}^n} [f(\mathbf{x}) = 0] \leq \deg(f) \cdot p$$

where $p = \max_{x \in \mathcal{K}} \mathcal{X}(x)$.

Proof. As usual, we prove the claim by induction on the number n of variables:

For $n = 1$, the non-zero polynomial f can have at most $\deg(f)$ roots in \mathcal{K} . The probability that $x \leftarrow \mathcal{X}$ lies in $f^{-1}(0)$ hence is bounded by $\deg(f) \cdot p$.

For the induction step, write f as

$$f(X) = \sum_{i=0}^{\deg f} f_i(X_1, \dots, X_{n-1}) \cdot X_n^i$$

with $\deg f_i \leq \deg(f) - i$. Since $f \neq 0$, one of the $f_0, \dots, f_{\deg(f)}$ must be non-zero, too. Let $d \in \{0, \dots, \deg(f)\}$ be maximal s.t. $f_d \neq 0$. Our induction hypothesis implies

$$\Pr_{\mathbf{x}' \leftarrow \mathcal{X}^{n-1}} [f_d(\mathbf{x}') = 0] \leq d \cdot p.$$

If $f_d(\mathbf{x}')$ is non-zero, then $f(\mathbf{x}', X_n)$ is a univariate polynomial of degree $\deg(f) - d$. Hence, we have

$$\Pr_{x'' \leftarrow \mathcal{X}} [f(\mathbf{x}', x'') = 0] \leq (\deg(f) - d) \cdot p.$$

By a union bound, the claim follows. \square

We want to apply the above lemma on the discrete Gaussian distribution $D_{\mathbb{Z}^n, \chi}$. For this end, recall that the coordinates of a sample $\mathbf{x} \leftarrow D_{\mathbb{Z}^n, \chi}$ are independently sampled according to $D_{\mathbb{Z}, \chi}$. Indeed, we have

$$\rho_\chi(\mathbf{x}) = \exp\left(-\pi \cdot \frac{\|\mathbf{x}\|_2^2}{\chi^2}\right) = \exp\left(-\pi \cdot \frac{x_1^2}{\chi^2}\right) \cdots \exp\left(-\pi \cdot \frac{x_n^2}{\chi^2}\right) = \rho_\chi(x_1) \cdots \rho_\chi(x_n)$$

and

$$\rho_\chi(\mathbb{Z}^n) = (\rho_\chi(\mathbb{Z}))^n.$$

It now follows:

Corollary 3 (Schwartz-Zippel for Spherical Discrete Gaussian Distributions). *Let $\chi > 1$, and let $q > 2\chi$ be prime. Fix a non-zero polynomial $f \in \mathbb{Z}_q[X_1, \dots, X_n]$. We have*

$$\begin{aligned} \Pr_{\mathbf{x} \leftarrow D_{\mathbb{Z}^n, \chi}} [f(\mathbf{x}) = 0] &\leq \frac{\deg(f)}{\rho_\chi(\{-(q-1)/2, \dots, (q-1)/2\})} + 4n \cdot \rho_\chi\left(\frac{q-1}{2}\right) \\ &\leq e^\pi \cdot \frac{\deg(f)}{2\lfloor \chi \rfloor + 1} + 4n \cdot \rho_\chi\left(\frac{q-1}{2}\right). \end{aligned}$$

Proof. Let \mathcal{X} be the restriction of $D_{\mathbb{Z}, \chi}$ to $\{-(q-1)/2, \dots, (q-1)/2\}$. Because of the subgaussianity of $D_{\mathbb{Z}, \chi}$, the statistical distance of \mathcal{X} and $D_{\mathbb{Z}, \chi}$ is bounded by

$$\Delta(\mathcal{X}, D_{\mathbb{Z}, \chi}) \leq 2 \cdot \Pr_{x \leftarrow D_{\mathbb{Z}, \chi}} \left[|x| > \frac{q-1}{2}\right] \leq 4 \exp\left(-\pi \frac{(q-1)^2}{4\chi^2}\right).$$

Hence, we can replace $D_{\mathbb{Z}^n, \chi}$ in the claim by \mathcal{X}^n and incur a statistical error bounded by $\leq 4n \cdot \rho_\chi((q-1)/2)$.

\mathcal{X} assumes its maximum value at zero. For $\mathcal{X}(0)$, we have

$$\mathcal{X}(0) = \frac{\rho_s(0)}{\rho_\chi(\{-(q-1)/2, \dots, (q-1)/2\})} \leq \frac{1}{\rho_\chi(\{-\lfloor \chi \rfloor, \dots, \lfloor \chi \rfloor\})}.$$

For $|x| \leq \chi$, we have $\rho_\chi(x) = \exp(-\pi \cdot x^2/\chi^2) \geq \exp(-\pi)$. Hence,

$$\mathcal{X}(0) \leq \frac{1}{\rho_\chi(\{-\lfloor \chi \rfloor, \dots, \lfloor \chi \rfloor\})} \leq \frac{e^\pi}{2\lfloor \chi \rfloor + 1}.$$

The claim follows now by applying Lemma 17. □

Theorem 3. *Let $m \in \mathbb{N}$ and $\chi > 1$. Let $q > 2\chi$ be a prime. Draw $\mathbf{M} \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}$. We have*

$$\Pr [\det(\mathbf{M}) = 0 \pmod q] \leq e^\pi \cdot \frac{m}{2\lfloor \chi \rfloor + 1} + 4n \cdot \rho_\chi\left(\frac{q-1}{2}\right).$$

Proof. Note that the determinant \det is a polynomial of degree m over m^2 variables with coefficients in \mathbb{Z} . Projecting the coefficients modulo q , we get the polynomial f that computes the map

$$f(\mathbf{M}) = \det(\mathbf{M}) \pmod q.$$

f is non-zero, since regular matrices modulo q do exist. Corollary 3 implies now

$$\Pr_{\mathbf{M} \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}} [\det(\mathbf{M}) = 0 \pmod q] = \Pr_{\mathbf{M} \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}} [f(\mathbf{M}) = 0] \leq e^\pi \cdot \frac{m}{2\lfloor \chi \rfloor + 1} + 4n \cdot \rho_\chi\left(\frac{q-1}{2}\right).$$

Hence, our theorem follows. □

A.3 Wrapping Up

We will now finish the proof of Theorem 2:

Proof (Theorem 2). Draw $\mathbf{B}, \mathbf{P} \leftarrow \mathbb{Z}_q^{n \times m}$ uniformly at random and sample $\mathbf{U} \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}$ conditioned on $\mathbf{B}\mathbf{U} = \mathbf{P}$.

Denote the columns of \mathbf{P} by $\mathbf{p}_1, \dots, \mathbf{p}_m \in \mathbb{Z}_q^n$, and the columns of \mathbf{U} by $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{Z}^m$. Note that the columns $\mathbf{u}_1, \dots, \mathbf{u}_m$ are independent of each other. Indeed, each column \mathbf{u}_i is sampled from $D_{\mathbb{Z}^m, \chi}$ conditioned on $\mathbf{B} \cdot \mathbf{u}_i = \mathbf{p}_i$.

Sample new column vectors $\mathbf{u}'_1, \dots, \mathbf{u}'_m \leftarrow D_{\mathbb{Z}^m, \chi}$ and set

$$\mathbf{U}' = (\mathbf{u}'_1 | \dots | \mathbf{u}'_m) \in \mathbb{Z}^{m \times m}.$$

According to Corollary 2, with probability of $\geq 1 - 2q^{-n}$ over the randomness of \mathbf{B} , we have for each $i \in [m]$

$$\Delta(\mathbf{u}'_i, \mathbf{u}_i) \leq 2^{-n+1}.$$

Since the columns $\mathbf{u}_1, \dots, \mathbf{u}_m$ are independent, it follows

$$\Delta(\mathbf{U}', \mathbf{U}) \leq m \cdot 2^{-n+1}.$$

For $\mathbf{U}' \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}$, Theorem 3 states

$$\Pr[\det(\mathbf{U}') = 0 \bmod q] \leq e^\pi \cdot \frac{m}{2[\chi] + 1} + n \cdot 4 \exp\left(-\pi \frac{(q-1)^2}{4\chi^2}\right).$$

By taking the statistical errors $2q^{-n}$ and $m \cdot 2^{-n+1}$ into account, we can replace \mathbf{U} by \mathbf{U}' and get

$$\Pr[\det(\mathbf{U}) = 0 \bmod q] \leq e^\pi \cdot \frac{m}{2[\chi] + 1} + n \cdot 4 \exp\left(-\pi \frac{(q-1)^2}{4\chi^2}\right) + 2q^{-n} + 2m \cdot 2^{-n}.$$

This closes the proof of Theorem 2. \square

B Missing Proofs

We provide the missing proofs for Sections 5 and 7. For ease of reading we also restate the related lemmas.

Lemma 6. Let $(\mathbf{P}_1 \in \mathbb{Z}_q^{n \times m}, \mathbf{u} \in \{0, 1\}^m) \leftarrow \text{Samp}_1(1^\lambda)$. For $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{D}_1 \leftarrow D_{\Lambda_{\mathbf{P}_1}^\perp(\mathbf{B}), \chi}$, we have

$$\Pr[\mathbf{D}_1 \cdot \mathbf{u} = 0 \bmod q] \leq O(m/\chi).$$

Proof. Because of the Leftover-Hash Lemma 3, \mathbf{P}_1 is statistically close to a uniformly random matrix. In the latter case, Lemma 5 implies that \mathbf{D}_1 is regular with probability at least $1 - O(m/\chi)$. Since \mathbf{u} is never zero, $\mathbf{D}_1 \cdot \mathbf{u} \bmod q$ can only be zero if \mathbf{D}_1 is not regular. As we explained, the probability for this event lies in $O(m/\chi)$. \square

Lemma 7. Let $(\mathbf{P}_1 \in \mathbb{Z}_q^{n \times m}, \mathbf{u} \in \{0, 1\}^m) \leftarrow \text{Samp}_1(1^\lambda)$. Sample $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \chi}$ and $\mathbf{D}_1 \leftarrow D_{\Lambda_{\mathbf{P}_1}^\perp(\mathbf{B}), \chi}$. We have

$$\Pr[|\mathbf{e}_0^\top \cdot \mathbf{D}_1 \cdot \mathbf{u}| \geq \lambda^2 m^2 \chi^2] \leq 2^{-\lambda}.$$

Proof. Again, we invoke the Leftover-Hash Lemma 3, to bound the statistical distance of \mathbf{P}_1 to a uniformly random matrix by 2^{-n} . In that case, we prove in Lemma 16 of Appendix A.1 that the columns of \mathbf{D}_1 are distributed according to $D_{\mathbb{Z}^m, \chi}$ if one does not know \mathbf{P}_1 . In particular, the columns of \mathbf{D}_1 are subgaussian with parameter χ . Since \mathbf{e}_0 is subgaussian with parameter χ , too, the absolute values of each entry of \mathbf{D}_1 and \mathbf{e}_0 is bounded by $\chi\lambda$ with probability $1 - \exp(-\pi\lambda^2)$. Since \mathbf{u} only has binary entries, the quantity $|\mathbf{e}_0^\top \cdot \mathbf{D}_1 \cdot \mathbf{u}|$ is bounded by $\lambda^2 m^2 \chi^2$ with probability at least $\geq 1 - (m^2 + m) \exp(-\pi\lambda^2)$. \square

Lemma 8. For $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{u}_2 \leftarrow \{0, 1\}^m$, $\mathbf{p}_2 = \mathbf{B} \cdot \mathbf{u}_2$ and $\mathbf{d}_2 \leftarrow D_{\Lambda_{\mathbf{p}_2}^\perp(\mathbf{B}), \chi}$, we have $\Pr[\mathbf{d}_2 = \mathbf{u}_2] \leq 2^{-m+1} + q^{-n}$.

Proof. By Lemma 1, with probability at least $1 - q^{-n}$, the smoothing parameter of $\Lambda^\perp(\mathbf{B})$ is bounded by $\chi \in \Omega(\sqrt{n})$. In such cases, fix $\mathbf{p}_2 \in \mathbb{Z}_q^n$ and consider

$$D_{\Lambda_{\mathbf{p}_2}^\perp(\mathbf{B}), \chi}(\mathbf{d}_2) = D_{\Lambda^\perp(\mathbf{B}), \chi, \mathbf{u}_2}(\mathbf{d}_2 + \mathbf{u}_2)$$

Lemma 2 bounds the latter term by $2^{-m} \cdot \frac{1+\epsilon}{1-\epsilon}$. Hence, the claim follows. \square

Lemma 9. Draw $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{u}_2 \leftarrow \{0, 1\}^m$ and set $\mathbf{p}_2 = \mathbf{B}\mathbf{u}_2 \bmod q$. Draw $\mathbf{d}_2 \leftarrow D_{\Lambda_{\mathbf{p}_2}^\perp(\mathbf{B}), \chi}$ and $\mathbf{f}'_1, \dots, \mathbf{f}'_m \leftarrow D_{\mathbb{Z}^m, \chi}$.

We have

$$\Pr [\exists i \in [m] : |\mathbf{f}'_i \cdot \mathbf{d}_2| \geq \lambda^2 m \chi^2] \leq 2^{-\lambda}.$$

Additionally, we have for $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}_q^m, \chi}$,

$$\Pr [|\mathbf{e}_0 \cdot (\mathbf{d}_2 - \mathbf{u}_2)| \geq \lambda \chi \cdot (\lambda \chi + 1) \cdot m] \leq 2^{-\lambda}.$$

Proof. Because of the Leftover-Hash Lemma 3, the statistical distance of \mathbf{p}_2 to a uniformly random vector is bounded by 2^{-n} . By using Lemma 16 again, it follows that \mathbf{d}_2 is distributed according to $D_{\mathbb{Z}^m, \chi}$, as long as does not see \mathbf{p}_2 . Hence, the coordinates of \mathbf{d}_2 are subgaussian with parameter χ . The vectors $\mathbf{f}'_1, \dots, \mathbf{f}'_m \leftarrow D_{\mathbb{Z}^m, \chi}$ are subgaussian with parameter χ , too. Ergo, all entries of $\mathbf{f}'_1, \dots, \mathbf{f}'_m, \mathbf{d}_2$ are bounded by $\chi \lambda$ with probability $\geq 1 - (m^2 + m) \exp(-\pi \lambda^2)$. In this case, the quantities $|\mathbf{f}'_i \cdot \mathbf{d}_2|$, $i = 1, \dots, m$, are all bounded by $\lambda^2 m \chi^2$.

Similarly, we can bound $|\mathbf{e}_0 \cdot (\mathbf{d}_2 - \mathbf{u}_2)|$ by $|\mathbf{e}_0 \cdot \mathbf{d}_2| + |\mathbf{e}_0 \cdot \mathbf{u}_2| \leq \lambda^2 \chi^2 m + \lambda \chi m$ with overwhelming probability $\geq 1 - 2^{-\lambda}$. \square

Lemma 10. Let $m > n$ and $t < q$. Draw $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times m}$ uniformly at random.

1. We have $\Pr [\exists \mathbf{M}_1, \mathbf{M}_2 \in \mathbb{Z}_q^{m \times n} : \mathbf{A} - \mathbf{M}_1 \mathbf{M}_2^T \in \{-t, \dots, t\}^{m \times m}] \leq q^{2mn} \cdot \left(\frac{2t+1}{q}\right)^{m^2}$.
2. With probability $\geq 1 - q^{2mn} \cdot \left(\frac{2t+1}{q}\right)^{m^2}$, every PPT adversary has negligible advantage in distinguishing $\text{iO}(C_{\mathbf{A}, n, t})$ and $\text{iO}(\mathbf{0})$, where $\mathbf{0}$ is an appropriately padded circuit that always outputs 0 and is devoid of any information).

Proof. 1. Set

$$F := \{\mathbf{M}_1 \mathbf{M}_2^T \mid \mathbf{M}_1, \mathbf{M}_2 \in \mathbb{Z}_q^{m \times n}\}.$$

$\mathbb{Z}_q^{m \times n}$ has q^{mn} many elements, hence

$$\#F \leq q^{2mn}.$$

Now, for some fixed $\mathbf{A} \in \mathbb{Z}_q^{m \times m}$ consider the set

$$G := \{\mathbf{R} \in \mathbb{Z}_q^{m \times m} : \mathbf{R} - \mathbf{A} \in \{-t, \dots, t\}^{m \times m}\}.$$

Note that we have

$$G = \mathbf{A} + \{-t, \dots, t\}^{m \times m}.$$

Hence,

$$\#G \leq (2t + 1)^{m^2}.$$

Since

$$\begin{aligned} H &:= \{\mathbf{A} \in \mathbb{Z}_q^{m \times m} : \mathbf{A} - \mathbf{M}_1 \mathbf{M}_2^T \in \{-t, \dots, t\}^{m \times m} \text{ for some } \mathbf{M}_1, \mathbf{M}_2 \in \mathbb{Z}_q^{m \times n}\} \\ &= \bigcup_{\mathbf{B} \in F} \{\mathbf{A} \in \mathbb{Z}_q^{m \times m} : \mathbf{A} - \mathbf{B} \in \{-t, \dots, t\}^{m \times m}\} \end{aligned}$$

$$= \bigcup_{\mathbf{B} \in F} \mathbf{B} + \{-t, \dots, t\}^{m \times m},$$

it follows

$$\#H \leq \#F \cdot (2t + 1)^{m^2} = q^{2mn} \cdot (2t + 1)^{m^2}.$$

We now have

$$\begin{aligned} & \Pr [\exists \mathbf{M}_1, \mathbf{M}_2 \in \mathbb{Z}_q^{m \times n} : \mathbf{A} - \mathbf{M}_1 \mathbf{M}_2^T \in \{-t, \dots, t\}^{m \times m}] \\ &= \Pr [\mathbf{A} \in H] \leq \frac{\#H}{\#\mathbb{Z}_q^{m \times m}} \leq \frac{q^{2mn} \cdot (2t + 1)^{m^2}}{q^{m^2}}. \end{aligned}$$

2. With high probability, $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times m}$ will admit no approximated low-rank factorisation. In this case, $C_{\mathbf{A}, n, t}$ will output 0 for each possible input. The claim now follows from the null-iO security. \square

Lemma 11. *Let $q > 1$ be any modulus and let $I \subseteq \mathbb{Z}_q$ be a consecutive interval. Let $\mathbf{u} \in \mathbb{Z}_q^n$ be non-zero and draw $\mathbf{r} \leftarrow \mathbb{Z}_q^n$ uniformly at random. We have*

$$\Pr [\mathbf{u}^T \cdot \mathbf{r} \bmod q \in I] \leq \frac{\#I + \|\mathbf{u}\|_\infty}{q}.$$

Proof. Since the entries of \mathbf{r} are uniformly random, it suffices to consider the case $n = 1$. Hence, let $u \in \mathbb{Z}_q$ be non-zero and draw $r \leftarrow \mathbb{Z}_q$. Let $p = (u, q)$ be the greatest common divisor of u and q . Note that $u \cdot r$ is uniformly distributed in $p \cdot \mathbb{Z}_q = \{0, p, \dots, (\frac{q}{p} - 1)p\}$. Hence, we have

$$\Pr [ur \in I] = \frac{\#(I \cap p\mathbb{Z}_q)}{q/p}.$$

The claim now follows, since $\#(I \cap p\mathbb{Z}_q) \leq \frac{\#I}{p} + 1$ and $p \leq |u|$. \square

C Additional Counterexamples

C.1 Counterexample 4 (Sketch)

We sketch a counterexample for the case $\beta = \gamma = 0$, i.e., Samp_4 will not receive \mathbf{B} as input and neither \mathbf{B} nor \mathbf{P} will be part of the joint distribution of the pre and post challenge. Additionally, Samp_4 will output *no* auxiliary information. However, a key difference will be that the LWE secret \mathbf{S} will now be a matrix of shape $m \times n$ chosen by Samp_4 .

Let $\mathbf{v}_1 \in \mathbb{Z}_q^m$ be some fixed publicly known vector, for example $\mathbf{v}_1 = (1, 0, \dots, 0)$. Samp_4 first samples a uniformly random vector $\mathbf{w}_1 \in \mathbb{Z}_q^m$. Denote the vector of the lowest bits of the entries of \mathbf{w}_1 by $\mathbf{u} = \text{LSB}(\mathbf{w}_1) \in \{0, 1\}^m$. Now, Samp_4 additionally samples a matrix $\mathbf{S} \leftarrow \mathbb{Z}_q^{m \times n}$ which is uniformly random conditioned on $\mathbf{u}^T \mathbf{S} = \mathbf{0}$. Additionally, Samp_4 samples vectors $\mathbf{v}_2, \dots, \mathbf{v}_n, \mathbf{w}_2, \dots, \mathbf{w}_n \leftarrow \mathbb{Z}_q^m$ and sets

$$\mathbf{P} = (\mathbf{P}_1 | \mathbf{P}_2) = \begin{pmatrix} \mathbf{v}_1^T & \mathbf{w}_1^T \\ \vdots & \vdots \\ \mathbf{v}_n^T & \mathbf{w}_n^T \end{pmatrix} \in \mathbb{Z}_q^{n \times (2m)}$$

where $\mathbf{P}_1, \mathbf{P}_2 \in \mathbb{Z}_q^{n \times m}$. Finally, Samp_4 outputs (\mathbf{S}, \mathbf{P}) .

In the if-challenge, a challenger has to distinguish

$$\mathbf{S}^T \mathbf{B} + \mathbf{E}_0 \bmod q, \quad \mathbf{S}^T \mathbf{P}_1 + \mathbf{E}_1 \bmod q, \quad \mathbf{S}^T \mathbf{P}_2 + \mathbf{E}_2 \bmod q$$

for $(\mathbf{S}, \mathbf{P}_1, \mathbf{P}_2) \leftarrow \text{Samp}_4(1^\lambda)$, $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{E}_0, \mathbf{E}_1, \mathbf{E}_2 \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}$ from three uniformly random matrices of shape $m \times m$. Without loss of generality, assume that the first entry u_1 of \mathbf{u} is one. In this case, we have for $\mathbf{S} = (\mathbf{s}_1 | \dots | \mathbf{s}_n)$

$$\mathbf{s}_1 = - \sum_{i=2}^n u_i \cdot \mathbf{s}_i \pmod q$$

and for $\mathbf{B}^T = (\mathbf{b}_1 | \dots | \mathbf{b}_n)$

$$\begin{aligned} \mathbf{S}\mathbf{B} &= \mathbf{s}_1 \mathbf{b}_1^T + \dots + \mathbf{s}_n \mathbf{b}_n^T = \mathbf{s}_2 (\mathbf{b}_2^T - u_2 \mathbf{b}_1^T) + \dots + \mathbf{s}_n (\mathbf{b}_n^T - u_n \mathbf{b}_1^T) && \pmod q \\ &= (\mathbf{s}_2 | \dots | \mathbf{s}_n) \cdot \left(\begin{pmatrix} \mathbf{b}_2^T \\ \vdots \\ \mathbf{b}_n^T \end{pmatrix} - \begin{pmatrix} u_2 \\ \vdots \\ u_n \end{pmatrix} \cdot \mathbf{b}_1^T \right) && \pmod q. \end{aligned}$$

The important observation is that $(\mathbf{b}_2 | \dots | \mathbf{b}_n) \in \mathbb{Z}_q^{(n-1) \times m}$ is uniformly random, and stays uniformly random if we perturb it by $(u_2, \dots, u_n) \cdot \mathbf{b}_1^T$. Since the last $n-1$ columns of \mathbf{S} are uniformly random and independent of each other, by the LWE assumption with secret-key length $n-1$, it follows that no PPT adversary can distinguish between $\mathbf{S}\mathbf{B} + \mathbf{E}_0 \pmod q$ and a uniformly random matrix. The same argument works analogously for $\mathbf{S}\mathbf{P}_1 + \mathbf{E}_1 \pmod q$ and $\mathbf{S}\mathbf{P}_2 + \mathbf{E}_2 \pmod q$, since the last $n-1$ rows of \mathbf{P}_1 and \mathbf{P}_2 are uniformly random and independent of each other. Hence, the intractability of the if-challenge follows by LWE.

Now, in the then-challenge, an adversary receives

$$\mathbf{C}, \mathbf{D}_1, \mathbf{D}_2$$

for $(\mathbf{S}, \mathbf{P}_1, \mathbf{P}_2) \leftarrow \text{Samp}_4(1^\lambda)$, $\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{E}_0 \leftarrow D_{\mathbb{Z}^{m \times m}, \chi}$, $\mathbf{D}_i \leftarrow D_{\Lambda_{\mathbf{P}_i}^\perp(\mathbf{B}), \chi}$, $i = 1, 2$, and has to decide if $\mathbf{C} = \mathbf{S}^T \mathbf{B} + \mathbf{E}_0 \pmod q$ or if \mathbf{C} has been sampled uniformly at random.

The trick is to extract the vector \mathbf{w}_1 , which lowest bits help at distinguishing, from the preimages $\mathbf{D}_1, \mathbf{D}_2$ by only knowing \mathbf{v}_1 . Because of Lemma 5, $\mathbf{D}_1 \pmod q \in \mathbb{Z}_q^{m \times m}$ will be of full rank with high probability. In that case, the adversary can compute a vector $\mathbf{z} \in \mathbb{Z}_q^m$ s.t.

$$\mathbf{z}^T \mathbf{D}_1 = \mathbf{v}_1^T \pmod q,$$

because \mathbf{v}_1 is a publicly known vector. Since $\mathbf{D}_1 \pmod q$ is regular, the equation system $\mathbf{z}^T \mathbf{D}_1 = \mathbf{v}_1^T \pmod q$ can only have one solution. Hence, \mathbf{z} is unique and must be of shape

$$\mathbf{z}^T = (1, 0, \dots, 0)^T \cdot \mathbf{B} = \mathbf{b}_1^T,$$

since $(1, 0, \dots, 0)^T \cdot \mathbf{P}_1 = \mathbf{v}_1^T$. The adversary can now compute

$$\mathbf{z}^T \cdot \mathbf{D}_2 = (1, 0, \dots, 0)^T \cdot \mathbf{B} \cdot \mathbf{D}_2 = (1, 0, \dots, 0)^T \cdot \mathbf{P}_2 = \mathbf{w}_1 \pmod q.$$

From \mathbf{w}_1 it can extract $\mathbf{u} \in \{0, 1\}^m$, which helps at deciding if \mathbf{C} is random, since

$$\mathbf{u}^T \cdot (\mathbf{S}^T \mathbf{B} + \mathbf{E}_0) = \mathbf{u}^T \cdot \mathbf{S}^T \mathbf{B} + \mathbf{u}^T \cdot \mathbf{E}_0 = \mathbf{u}^T \mathbf{E}_0 \pmod q,$$

where the entries of the vector $\mathbf{u}^T \mathbf{E}_0$ are bounded by $m\lambda\chi$ with overwhelming probability, which can be shown by a usual subgaussianity argument. However, if \mathbf{C} would be uniformly random, then the vector $\mathbf{u}^T \mathbf{C}$ would be statistically close to a uniformly random vector. Hence, we again yield a PPT adversary for the post challenge with high advantage.

We note again that \mathbf{P} is uniformly random except for the first m entries of its first row. However, one can see that such \mathbf{P} fails to achieve indistinguishability as required by Pre2 of private-coin hiding evasive LWE (Definition 9), since a distinguisher can check the consistency of the first row of \mathbf{P}_1 with \mathbf{v}_1 and, for parameter $\ell = 2$, succeed with probability $1 - (1/2)^m$.

$\text{Pre1}'^b_{\mathcal{A}}(1^\lambda)$	$\text{Post}'^b_{\mathcal{B}}(1^\lambda)$
$\mathbf{B} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$	$\mathbf{B} \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$
$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$
if $b = 0$ then	if $b = 0$ then
$\mathbf{E}_B \leftarrow_{\$} \chi_B, \mathbf{E}_P \leftarrow_{\$} \chi_P$	$\mathbf{E}_B \leftarrow_{\$} \chi_B$
$\mathbf{C}_B := \mathbf{S}\mathbf{B} + \mathbf{E}_B \bmod q$	$\mathbf{C}_B := \mathbf{S}\mathbf{B} + \mathbf{E}_B \bmod q$
$\mathbf{C}_P := \mathbf{S}\mathbf{P} + \mathbf{E}_P \bmod q$	$\mathbf{U} \leftarrow_{\$} D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$
if $b = 1$ then	if $b = 1$ then
$\mathbf{C}_B \leftarrow_{\$} \mathbb{Z}_q^{t \times m}$	$\mathbf{C}_B \leftarrow_{\$} \mathbb{Z}_q^{t \times m}$
$\mathbf{C}_P \leftarrow_{\$} \mathbb{Z}_q^{t \times m_P}$	$\mathbf{U} \leftarrow_{\$} D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$
return $\mathcal{A}(\mathbf{C}_B, \mathbf{C}_P, \text{aux})$	return $\mathcal{B}(\mathbf{C}_B, \mathbf{U}, \text{aux})$
<hr/>	
$\text{Pre2}'^b_{\mathcal{A}}(1^\lambda)$	
$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$	
if $b = 0$ then	
return $\mathcal{A}(\mathbf{P} \bmod \ell, \text{aux})$	
if $b = 1$ then	
$\mathbf{R} \leftarrow_{\$} \mathbb{Z}_q^{n \times m_P};$	
return $\mathcal{A}(\mathbf{R} \bmod \ell, \text{aux})$	

Fig. 6: Experiments $\text{Pre1}'$, $\text{Pre2}'$ and Post' for a variation of the private-coin hiding evasive LWE assumption.

C.2 Counterexample 5 (Sketch)

We now sketch here a counterexample¹³ against a **variation** of the private-coin hiding evasive LWE assumption. This variation, summarised in Fig. 6, states that if there is a PPT adversary \mathcal{B} which predicts b in Post' with non-negligible advantage, then there is either a PPT adversary \mathcal{A}_1 which predicts b in $\text{Pre1}'$ with non-negligible advantage, or a PPT adversary \mathcal{A}_2 which predicts b in $\text{Pre2}'$ with non-negligible advantage. Different from Definition 9, the experiment $\text{Pre2}'$ now requires that $\mathbf{P} \bmod \ell$ is computationally indistinguishable from $\mathbf{R} \bmod \ell$, $\mathbf{R} \leftarrow \mathbb{Z}_q^{n \times m_P}$, given aux . This change enables an attack with a time complexity of $\text{poly}(m^\ell)$ so that this variation of private-coin hiding evasive LWE cannot be secure if one allows ℓ to be constant and \mathbf{P} to have up to $O(m^\ell)$ columns. Details of the attack follow.

Let

$$m \geq 3n \log(q), \quad \chi \geq \lambda n \ell \cdot \binom{m + \ell - 1}{\ell}, \quad q \geq \max(2^\lambda \cdot \ell, \lambda^3 m^2 \chi^2).$$

Similarly to counterexample 2 in Section 5.2, the sampler embeds information into aux which helps the post adversary to recover \mathbf{B} when given the short preimage \mathbf{D} . However, this time we have to ensure that aux does not help at distinguishing between $(\mathbf{P} \bmod \ell)$ and $(\mathbf{R} \bmod \ell)$, where \mathbf{R} is a uniformly random matrix of the same shape.

Our attack relies on algorithms for solving the *Learning with Bounded Errors* problem. In this problem, one is given an LWE sample $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)$ and promised that each entry of the noise vector lies in a set of *bounded size*¹⁴. In its simplest form, where each noise-term lies in a set of size 2, this problem is known as *Learning with Binary Errors*. Arora and Ge [AG11] show that this problem can be solved by relinearization when given $m = O(qn^2)$ samples. Algebraic attacks even solve the problem when only given $m = O(n^2)$ samples, cf. [MP13,ACFP14,STA20,Ste24,NMSÜ24]. To make use of these algebraic attacks, we set the number of columns of \mathbf{P} to be

$$m_P = 2m + \binom{m + \ell - 1}{\ell} \in \Theta(m^\ell).$$

¹³ We thank a reviewer of AsiaCrypt 2024 for pointing out this counterexample to us!

¹⁴ In our case, each noise term lies in $\{0, \dots, \ell - 1\}$.

$\text{Samp}_5(1^\lambda)$ outputs the following:

$$(\mathbf{P} = (\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3), \quad \text{aux} = \mathbf{P}_3 - (\mathbf{P}_3 \bmod \ell))$$

where $\mathbf{P}_3 \leftarrow_{\$} \mathbb{Z}_q^{n \times \binom{m+\ell-1}{\ell}}$. Denote by **MSB** the most significant bits of a vector. $\mathbf{P}_1, \mathbf{P}_2$ are sampled similar as follows

$$\begin{aligned} \mathbf{u}' &\leftarrow_{\$} \{0, 1\}^{m-1}, & \mathbf{P}'_1 &\leftarrow_{\$} \mathbb{Z}_q^{n \times (m-1)}, & \mathbf{P}'_2 &\leftarrow_{\$} \mathbb{Z}_q^{(n-1) \times m}, & \mathbf{t} &\leftarrow_{\$} \mathbb{Z}_q^m \\ \mathbf{u}^\top &= ((\mathbf{u}')^\top, 1) \in \{0, 1\}^{1 \times m}, & \mathbf{P}_1 &= (\mathbf{P}'_1 | -\mathbf{P}'_1 \mathbf{u}' \bmod q), & \mathbf{P}_2 &= \left(\begin{array}{c} \mathbf{P}'_2 \\ (\mathbf{t} + 2^{\lceil \log_2(q) \rceil - 1} \cdot (\mathbf{u} - \text{MSB}(\mathbf{t})))^\top \end{array} \right)^\top. \end{aligned}$$

As usual, the last row of \mathbf{P}_2 encodes a binary kernel vector of \mathbf{P}_1 . This time, the vector \mathbf{u} is encoded in the most significant bits of \mathbf{P}_2 . Since the last row of \mathbf{P}_2 is given by $\mathbf{t} + 2^{\lceil \log_2(q) \rceil - 1} \cdot (\mathbf{u} - \text{MSB}(\mathbf{t}))$, all other bits in this row are identically distributed as the corresponding bits of \mathbf{t} .

We have to argue that every PPT adversary has a negligible advantage in $\text{Pre}2'$. In this experiment, an adversary has to distinguish between $\mathbf{P} \bmod \ell$ and $\mathbf{R} \bmod \ell$ for $\mathbf{R} \leftarrow_{\$} \mathbb{Z}_q^{n \times m_P}$ when given $\text{aux} = \mathbf{P}_3 - (\mathbf{P}_3 \bmod \ell)$. Decompose \mathbf{R} into $\mathbf{R}_1, \mathbf{R}_2 \leftarrow_{\$} \mathbb{Z}_q^{n \times m}$ and $\mathbf{R}_3 \leftarrow_{\$} \mathbb{Z}_q^{n \times \binom{m+\ell-1}{\ell}}$. First we note that $\mathbf{P}_2 \bmod \ell$ and $\mathbf{R}_2 \bmod \ell$ are identically distributed. In particular, $\mathbf{P}_2 \bmod \ell$ is devoid of any information of \mathbf{u} . Hence, \mathbf{P}_1 is statistically close to \mathbf{R}_1 . Finally, \mathbf{P}_3 is uniformly random, however $\mathbf{P}_3 - (\mathbf{P}_3 \bmod \ell)$ is known by the distinguisher. Since $q \geq \ell \cdot 2^\lambda$, the statistical distance between $\mathbf{P}_3 \bmod \ell$ and $\mathbf{R}_3 \bmod \ell$ is negligible¹⁵. It follows that the advantage of an adversary in determining b in $\text{Pre}2'$ is negligible, and assuming **LWE**, the advantage of a PPT adversary in distinguishing b in $\text{Pre}1'$ is also negligible. We omit the formal argument here, since it follows along the lines of Section 5.2.

Finally, let us turn to the then-challenge. The post adversary receives

$$\mathbf{D}_1 \in \mathbb{Z}_q^{m \times m}, \quad \mathbf{D}_2 \in \mathbb{Z}_q^{m \times m}, \quad \mathbf{D}_3 \in \mathbb{Z}_q^{m \times \binom{m+\ell-1}{\ell}}, \quad \mathbf{W} := \mathbf{P}_3 - (\mathbf{P}_3 \bmod \ell)$$

where \mathbf{D}_i is short with $\mathbf{B} \cdot \mathbf{D}_i = \mathbf{P}_i$. It has to distinguish $\mathbf{s}^\top \mathbf{B} + \mathbf{e}^\top$ from a uniformly random vector. We claim that the adversary can reconstruct \mathbf{B} from \mathbf{D}_3 and \mathbf{W} . Let $\mathbf{b} \in \mathbb{Z}_q^m$ and $\mathbf{w} \in \mathbb{Z}_q^{\binom{m+\ell-1}{\ell}}$ be the i -th row of \mathbf{B} and \mathbf{W} . Note that we have

$$\mathbf{w} - \mathbf{D}_3^\top \cdot \mathbf{b} \in \{0, \dots, \ell - 1\}^{\binom{m+\ell-1}{\ell}}.$$

According to Lemma 18 (shown below), the adversary can solve this problem and extract \mathbf{b} with success probability $\geq 1 - O(1/(n\lambda))$ and time complexity lies in $\text{poly}(m^\ell)$. By a union bound, the adversary can extract all rows of \mathbf{B} from \mathbf{D}_3 and \mathbf{W} with success probability $\geq 1 - O(1/\lambda)$. From here on, the adversary follows the strategy of Section 5.2. I.e., it computes $\mathbf{P}_2 = \mathbf{B}\mathbf{D}_2$, extracts the trapdoor \mathbf{u} from \mathbf{P}_2 and uses $\mathbf{D}_1 \mathbf{u}$ to distinguish $\mathbf{s}^\top \mathbf{B} + \mathbf{e}^\top$ from uniform randomness. We omit the rest of the proof as it is analogous to Section 5.2.

Lemma 18. *Set $M = \binom{m+\ell-1}{\ell}$ and let $q \geq \lambda \cdot \chi$. There is an algorithm \mathcal{B} that runs in time $\text{poly}(m^\ell)$ s.t.*

$$\Pr_{\mathbf{D} \leftarrow D_{\mathbb{Z}_q^M, \chi}} [\forall \mathbf{x} \in \mathbb{Z}_q^m, \mathbf{e} \in \{0, \dots, \ell - 1\}^M : \mathcal{B}(\mathbf{D}, \mathbf{D} \cdot \mathbf{x} - \mathbf{e}) = \mathbf{x}] \geq 1 - O(2^m \ell^M / \chi^{M-m} + \ell M / \chi).$$

Proof. Set $L = \{0, \dots, \ell - 1\}$ and $N = \{-\ell, \dots, \ell\}$. We will first turn to the question of unique solvability. Draw $\mathbf{D} \leftarrow D_{\mathbb{Z}_q^{M \times m}, \chi}$. We claim that it is very unlikely that there exist $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{Z}_q^m, \mathbf{e}_1, \mathbf{e}_2 \in L^M$ s.t. $\mathbf{D}\mathbf{x}_1 - \mathbf{e}_1 = \mathbf{D}\mathbf{x}_2 - \mathbf{e}_2$ and $\mathbf{x}_1 \neq \mathbf{x}_2$. Indeed, this would imply

$$\mathbf{D} \cdot (\mathbf{x}_1 - \mathbf{x}_2) = \mathbf{e}_1 - \mathbf{e}_2 \in N^M,$$

¹⁵ For an explanation, sample a number $x \leftarrow \{0, \dots, q - 1\}$ uniformly at random and decompose it as $x = a + \ell \cdot b$ with $a \in \{0, \dots, \ell - 1\}$ and $b \in \{0, \dots, \lfloor q/\ell \rfloor\}$. Since q is not a power of two, a is sampled uniformly at random conditioned on b iff $b < \lfloor q/\ell \rfloor$. The probability of $b = \lfloor q/\ell \rfloor$ is equal to the probability of $x \geq \lfloor q/\ell \rfloor \cdot \ell$, which is $(q - \lfloor q/\ell \rfloor \cdot \ell) / q \leq 2^{-\lambda}$.

i.e., the *high* matrix \mathbf{D} would have a short image vector. Decompose

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}_1 \in \mathbb{Z}_q^{m \times m} \\ \mathbf{D}_2 \in \mathbb{Z}_q^{(M-m) \times m} \end{pmatrix}, \quad \mathbf{e}_1 - \mathbf{e}_2 = \begin{pmatrix} \mathbf{f}_1 \in \mathbb{Z}_q^m \\ \mathbf{f}_2 \in \mathbb{Z}_q^{M-m} \end{pmatrix}$$

into top and bottom parts. According to Corollary 3, the probability of \mathbf{D}_1 being singular is bounded by

$$\leq e^\pi \cdot \frac{m}{2[\chi] + 1} + 4m^2 \cdot \rho_\chi \left(\frac{q-1}{2} \right) \in O(m/\chi).$$

If \mathbf{D}_1 is regular, the value of $\mathbf{x}_1 - \mathbf{x}_2$ is uniquely determined by \mathbf{f}_1 , since $\mathbf{D}_1(\mathbf{x}_1 - \mathbf{x}_2) = \mathbf{f}_1$. In particular, we have $\mathbf{f}_2 = \mathbf{D}_2 \cdot \mathbf{D}_1^{-1} \mathbf{f}_1$. Now, set $A := \{\mathbf{D}_1 \cdot \mathbf{f} \mid \mathbf{f} \in N^m\}$. Since the entries of \mathbf{D} are independent of each other, we have

$$\begin{aligned} & \Pr_{\mathbf{D}_2 \leftarrow D_{\mathbb{Z}_q^{(M-m) \times m}, \chi}} [\exists \mathbf{y} \in A : \mathbf{D}_2 \mathbf{y} \in N^{M-m}] \\ & \leq \sum_{\mathbf{y} \in A} \Pr_{\mathbf{D}_2 \leftarrow D_{\mathbb{Z}_q^{(M-m) \times m}, \chi}} [\mathbf{D}_2 \mathbf{y} \in N^{M-m}] \\ & = \sum_{\mathbf{y} \in A} \Pr_{\mathbf{d} \leftarrow D_{\mathbb{Z}_q^m, \chi}} [\mathbf{d}^\top \mathbf{y} \in N]^{M-m} \\ & \leq (2\ell - 1)^m \cdot \max_{\mathbf{y} \in A} \left(\Pr_{\mathbf{d} \leftarrow D_{\mathbb{Z}_q^m, \chi}} [\mathbf{d}^\top \mathbf{y} \in N]^{M-m} \right). \end{aligned}$$

Because of Corollary 3, we have

$$\Pr_{\mathbf{d} \leftarrow D_{\mathbb{Z}_q^m, \chi}} [\mathbf{d}^\top \mathbf{y} \in N] \leq e^\pi \cdot \frac{2\ell - 1}{2[\chi] + 1} + 4(2\ell - 1)^2 \cdot \rho_\chi \left(\frac{q-1}{2} \right) \in O(\ell/\chi).$$

Hence,

$$\begin{aligned} & \Pr_{\mathbf{D}_2 \leftarrow D_{\mathbb{Z}_q^{(M-m) \times m}, \chi}} [\exists \mathbf{y} \in A : \mathbf{D}_2 \mathbf{y} \in N^{M-m}] \\ & \leq (2\ell - 1)^m \cdot \max_{\mathbf{y} \in A} \left(\Pr_{\mathbf{d} \leftarrow D_{\mathbb{Z}_q^m, \chi}} [\mathbf{d}^\top \mathbf{y} \in N]^{M-m} \right) \\ & \leq (2\ell - 1)^m \cdot O((\ell/\chi)^{M-m}) = O(2^m \ell^M / \chi^{M-m}). \end{aligned}$$

Taking everything together, it now follows that the probability, that there exist $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{Z}_q^m$, $\mathbf{e}_1, \mathbf{e}_2 \in L^M$ s.t. $\mathbf{D}\mathbf{x}_1 - \mathbf{e}_1 = \mathbf{D}\mathbf{x}_2 - \mathbf{e}_2$ and $\mathbf{x}_1 \neq \mathbf{x}_2$, lies in

$$O(m/\chi + 2^m \ell^M / \chi^{M-m}).$$

Now, set $\mathbf{w} = \mathbf{D}\mathbf{x} - \mathbf{e}$ and let us assume that the problem $\mathbf{D}\mathbf{x} = \mathbf{w} + \mathbf{e}$ with unknowns \mathbf{x} and \mathbf{e} has only one solution. In this case, \mathcal{B} has to solve a linear equation system with $\binom{m+\ell-1}{\ell}$ equations over m variables where each equation is perturbed by a noise value in $\{0, \dots, \ell-1\}$. Introduce formal variables X_1, \dots, X_m that represent the unknown entries of $\mathbf{x} \in \mathbb{Z}_q^m$ and let $X = (X_1, \dots, X_m)$ be the corresponding vector of variables. Let $f \in \mathbb{Z}_q[Z]$ be the following univariate polynomial

$$f(Z) := \prod_{i=0}^{\ell-1} (Z - i).$$

We can transform the noisy linear equation system

$$\mathbf{w} - \mathbf{D} \cdot X = \begin{pmatrix} w_1 \\ \vdots \\ w_M \end{pmatrix} - \begin{pmatrix} \mathbf{d}_1 \\ \vdots \\ \mathbf{d}_M \end{pmatrix} \cdot \begin{pmatrix} X_1 \\ \vdots \\ X_m \end{pmatrix} \in L^M$$

to the following exact polynomial equation system

$$\begin{aligned} f(w_1 - \langle \mathbf{d}_1 \mid X \rangle) &= 0, \\ &\vdots \\ f(w_M - \langle \mathbf{d}_M \mid X \rangle) &= 0. \end{aligned}$$

Since we assumed that this problem has only one solution, a Macaulay matrix-based solving algorithm like Mutant-XL [CKPS00,DBM⁺08] will eventually find the solution \mathbf{x} . To bound the time complexity of this algorithm, we need to bound the *degree of regularity* of the polynomial equation system. This is exactly the degree D at which the ideal generated by the top terms

$$\langle \mathbf{d}_1 \mid X \rangle^\ell, \dots, \langle \mathbf{d}_M \mid B \rangle^\ell$$

contains all monomials of $\mathbb{Z}_q[X]$ of degree D . [NMSÜ24] shows that this is with high probability ℓ . Concretely, let $\tilde{\mathbf{D}}$ be the $M \times M$ -matrix whose i -th row contains the coefficients of the polynomial $\langle \mathbf{d}_i \mid X \rangle^\ell$. The polynomial equation system has degree of regularity ℓ iff $\tilde{\mathbf{D}}$ has full rank. We can consider the entries of $\tilde{\mathbf{D}}$ as polynomials of degree ℓ in the entries of $\mathbf{d}_1, \dots, \mathbf{d}_M$, i.e. in the entries of \mathbf{D} . Then, the polynomial $\det \tilde{\mathbf{D}}$ is of degree $\ell \cdot M$ in the entries of \mathbf{D}_3 . [NMSÜ24] proves that this polynomial is non-zero. Corollary 3 implies now that the probability of $\det \tilde{\mathbf{D}}$ being zero is bounded by

$$\leq e^\pi \cdot \frac{\ell M}{2\lfloor \chi \rfloor + 1} + 4(\ell M)^2 \rho_\chi \left(\frac{q-1}{2} \right) \in O(\ell M / \chi).$$

It is known that the solving degree (for a Mutant-XL approach) of a polynomial equation system is bounded by twice the degree of regularity [Sal23,Ste24]. Hence, with probability $\geq 1 - O(\ell M / \chi)$, \mathcal{B} can extract \mathbf{x} by calculating the Macaulay matrix of the corresponding polynomial equation system up to 2ℓ . The dimensions of this matrix lie in $\text{poly}(m^\ell)$ and performing Gaussian elimination on it has a time complexity of $\text{poly}(m^\ell)$.

With a union bound it follows, the success probability of \mathcal{B} is at least

$$\geq 1 - O(m/\chi + 2^m \ell^M / \chi^{M-m} + \ell M / \chi) = 1 - O(2^m \ell^M / \chi^{M-m} + \ell M / \chi).$$

C.3 Counterexample 6 (Sketch)

We sketch here a counterexample¹⁶ against a variation of public-coin evasive LWE Definition 9 where the sampling algorithm inputs¹⁷ the matrix \mathbf{B} . The key observation here is that the sampling algorithm can enforce algebraic relationships on \mathbf{D} when multiplied from the right to \mathbf{B} . Concretely, we will let \mathbf{B} and \mathbf{P} have $2n$ rows and let \mathbf{P} be

$$\mathbf{P} = \begin{pmatrix} 0 & \mathbf{I}_n \\ \mathbf{I}_n & 0 \end{pmatrix} \cdot \mathbf{B}$$

where \mathbf{I}_n is $n \times n$ -identity matrix. For \mathbf{D} with $\mathbf{B}\mathbf{D} = \mathbf{P}$ it is easy to verify that we must have

$$\mathbf{B}\mathbf{D}^2 = \mathbf{B}.$$

This can be used by the post adversary, assuming that $\mathbf{D}^2 \neq \mathbf{I}_m$ of course. Unfortunately, we could not prove that \mathbf{D} will be not a square root of the identity matrix, since \mathbf{D} does not have a simple spherical distribution. Therefore, to bound the success probability of the following attack, we need to assume the following:

Conjecture 1. Draw $\mathbf{B} \leftarrow \mathbb{Z}_q^{2n \times m}$. For $\mathbf{D} \leftarrow D_{\mathbb{Z}_q^{m \times m}, \chi}$ conditioned on

$$\mathbf{B}\mathbf{D} = \begin{pmatrix} 0 & \mathbf{I}_n \\ \mathbf{I}_n & 0 \end{pmatrix} \mathbf{B},$$

we have

$$\Pr [\mathbf{D}^2 \neq \mathbf{I}_m] \notin \text{negl}(\lambda).$$

¹⁶ We thank a reviewer of AsiaCrypt 2024 for pointing out this counterexample to us!

¹⁷ This demonstrates that public-coin evasive LWE is (presumably) insecure if one allows Samp to take \mathbf{B} as input.

We argue that Conjecture 1 is plausible. For example, if $\mathbf{r} \in \mathbb{Z}_q^m$ is a short kernel vector of \mathbf{B} , then we can add \mathbf{r} to any column of \mathbf{D} without changing its essential property.

Now, Samp_6 works as follows: on input 1^λ and $\mathbf{B} \leftarrow \mathbb{Z}_q^{2n \times m}$ it outputs

$$\mathbf{P} := \begin{pmatrix} 0 & \mathbf{I}_n \\ \mathbf{I}_n & 0 \end{pmatrix} \cdot \mathbf{B}$$

Note that Samp_6 is public-coin as it does not use any randomness. The hinting functions are for this attack defined to be empty (so, there will be no leakage on \mathbf{s}). Also, Samp_6 will not output an \mathbf{A} -matrix (i.e., $m_A = 0$).

In the if-challenge, an adversary has to distinguish

$$\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top, \quad \mathbf{s}^\top \mathbf{P} + \mathbf{e}_1^\top$$

from uniform randomness given \mathbf{B} and \mathbf{P} . \mathbf{B} and \mathbf{P} are strongly dependent, however, we can decompose them as follows

$$\mathbf{s} = \begin{pmatrix} \mathbf{s}_1 \in \mathbb{Z}_q^n \\ \mathbf{s}_2 \in \mathbb{Z}_q^n \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} \mathbf{B}_1 \in \mathbb{Z}_q^{n \times m} \\ \mathbf{B}_2 \in \mathbb{Z}_q^{n \times m} \end{pmatrix}, \quad \mathbf{P} = \begin{pmatrix} \mathbf{B}_2 \in \mathbb{Z}_q^{n \times m} \\ \mathbf{B}_1 \in \mathbb{Z}_q^{n \times m} \end{pmatrix}.$$

Then, we have

$$\begin{aligned} \mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top &= \mathbf{s}_1^\top \mathbf{B}_1 + \mathbf{s}_2^\top \mathbf{B}_2 + \mathbf{e}_0^\top, \\ \mathbf{s}^\top \mathbf{P} + \mathbf{e}_1^\top &= \mathbf{s}_1^\top \mathbf{B}_2 + \mathbf{s}_2^\top \mathbf{B}_1 + \mathbf{e}_1^\top. \end{aligned}$$

The pseudorandomness of $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top, \mathbf{s}^\top \mathbf{P} + \mathbf{e}_1^\top$ (given \mathbf{B} and \mathbf{P}) is implied by the pseudorandomness of $\mathbf{s}_1^\top \mathbf{B}_1 + \mathbf{e}_0^\top, \mathbf{s}_1^\top \mathbf{B}_2 + \mathbf{e}_1^\top$ given $\mathbf{B}_1, \mathbf{B}_2$ and \mathbf{s}_2 . This pseudorandomness follows from LWE for secret vectors of length n and $2m$ samples.

Now, for the post challenge, we are given \mathbf{B}, \mathbf{P} and \mathbf{D} short with

$$\mathbf{B} \cdot \mathbf{D} = \mathbf{P} = \begin{pmatrix} 0 & \mathbf{I}_n \\ \mathbf{I}_n & 0 \end{pmatrix} \cdot \mathbf{B}.$$

In particular, we have $\mathbf{B}\mathbf{D}^2 = \mathbf{B}$. According to Conjecture 1, with some probability, \mathbf{D}^2 will not be the identity. In the then-challenge, we have to decide if a vector $\mathbf{y} \in \mathbb{Z}_q^m$ is uniformly random or of shape $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top$ for short $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}^m, \chi}$. To accomplish that, we compute

$$\mathbf{z} := \mathbf{y}^\top \cdot (\mathbf{D}^2 - \mathbf{I}_m).$$

If \mathbf{y} is uniformly random, then \mathbf{z} will have large entries (with a probability that depends on \mathbf{D}^2 not being the identity). On the other hand, if $\mathbf{y} = \mathbf{s}^\top \mathbf{B} + \mathbf{e}_0^\top$, then $\mathbf{z} = \mathbf{e}_0^\top \cdot (\mathbf{D}^2 - \mathbf{I}_m)$ will be short with high probability. Hence, assuming Conjecture 1, the post adversary has a non-negligible advantage at breaking the public-coin evasive LWE assumption.

D Evasive Circular LWE

We attempt to give an abstraction of the evasive circular LWE assumption introduced by [HLL23] to build key-policy ABE for unbounded-depth circuits. We emphasise that our discussion only serves to understand the structural similarity and differences with its non-circular counterpart, but studying the plausibility of evasive circular LWE is beyond the scope of this work.

Definition 12 (Evasive Circular LWE). *Let the parameters*

$$\text{param} = (\mathcal{R}, q, n, n_A, m, m_P, m_A, m_{\text{circ}}, t, t_A, \mathcal{D}, \mathcal{S}, \chi_B, \chi_P, \chi_A, \chi_{\text{circ}}, f, f_A, \Sigma)$$

be parametrised by λ , where \mathcal{R} is a ring admitting an embedding as a lattice in \mathbb{R}^φ for some $\varphi \in \mathbb{N}$, $\mathcal{D} \sim \mathcal{R}_q^{n \times m}$, $\mathcal{S} \sim \mathcal{R}_q^{t \times n} \times \mathcal{R}_q^{t_A \times n_A}$, $\chi_B \sim \mathcal{R}^{t \times m}$, $\chi_P \sim \mathcal{R}^{t \times m_P}$, $\chi_A \sim \mathcal{R}^{t_A \times m_A}$, and $\chi_{\text{circ}} \sim \mathcal{R}^{t \times m_{\text{circ}}}$

$\text{Pre}_A^b(1^\lambda)$	$\text{Post}_B^b(1^\lambda)$
$\mathbf{B} \leftarrow \mathcal{D}$	$\mathbf{B} \leftarrow \mathcal{D}$
$(\mathbf{A}, \mathbf{P}, f_{\text{circ}}, \text{aux}) \leftarrow \text{Samp}(1^\lambda; \text{rand})$	$(\mathbf{A}, \mathbf{P}, f_{\text{circ}}, \text{aux}) \leftarrow \text{Samp}(1^\lambda; \text{rand})$
assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$	assert $\mathbf{P} \in \mathbf{BR}^{m \times m_P}$
$(\mathbf{S}, \mathbf{S}_A) \leftarrow \mathcal{S}$	$(\mathbf{S}, \mathbf{S}_A) \leftarrow \mathcal{S}$
$\text{hint} := (f(\mathbf{S}), f_A(\mathbf{S}_A))$	$\text{hint} := (f(\mathbf{S}), f_A(\mathbf{S}_A))$
$\mathbf{M} \leftarrow f_{\text{circ}}(\mathbf{S})$	$\mathbf{M} \leftarrow f_{\text{circ}}(\mathbf{S})$
if $b = 0$ then	if $b = 0$ then
$\mathbf{E}_A \leftarrow \chi_A, \mathbf{E}_B \leftarrow \chi_B, \mathbf{E}_P \leftarrow \chi_P$	$\mathbf{E}_A \leftarrow \chi_A, \mathbf{E}_B \leftarrow \chi_B$
$\mathbf{E}_{\text{circ}} \leftarrow \chi_{\text{circ}}$	$\mathbf{E}_{\text{circ}} \leftarrow \chi_{\text{circ}}$
$\mathbf{C}_A := \mathbf{S}_A \mathbf{A} + \mathbf{E}_A \text{ mod } q$	$\mathbf{C}_A := \mathbf{S}_A \mathbf{A} + \mathbf{E}_A \text{ mod } q$
$\mathbf{C}_B := \mathbf{S} \mathbf{B} + \mathbf{E}_B \text{ mod } q$	$\mathbf{C}_B := \mathbf{S} \mathbf{B} + \mathbf{E}_B \text{ mod } q$
$\mathbf{C}_P := \mathbf{S} \mathbf{P} + \mathbf{E}_P \text{ mod } q$	$\mathbf{U} \leftarrow D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$
$\mathbf{C}_{\text{circ}} := \mathbf{S} \mathbf{M} + \mathbf{E}_{\text{circ}} \text{ mod } q$	$\mathbf{C}_{\text{circ}} := \mathbf{S} \mathbf{M} + \mathbf{E}_{\text{circ}} \text{ mod } q$
if $b = 1$ then	if $b = 1$ then
$\mathbf{C}_A \leftarrow \mathcal{R}_q^{t_A \times m_A}$	$\mathbf{C}_A \leftarrow \mathcal{R}_q^{t_A \times m_A}$
$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t_B \times m_B}$	$\mathbf{C}_B \leftarrow \mathcal{R}_q^{t_B \times m_B}$
$\mathbf{C}_P \leftarrow \mathcal{R}_q^{t_P \times m_P}$	$\mathbf{U} \leftarrow D_{A_{\mathbf{P}}^\perp(\mathbf{B}), \Sigma}$
$\mathbf{C}_{\text{circ}} \leftarrow \mathcal{R}_q^{t_{\text{circ}} \times m_{\text{circ}}} \text{ mod } q$	$\mathbf{C}_{\text{circ}} \leftarrow \mathcal{R}_q^{t_{\text{circ}} \times m_{\text{circ}}} \text{ mod } q$
return $\mathcal{A} \left(\begin{array}{c} \mathbf{A}, \mathbf{B}, \mathbf{P}, \\ \mathbf{C}_A, \mathbf{C}_B, \mathbf{C}_P, \mathbf{C}_{\text{circ}}, \end{array} \text{hint, aux, rand} \right)$	return $\mathcal{B} \left(\begin{array}{c} \mathbf{A}, \mathbf{B}, \mathbf{P}, \\ \mathbf{C}_A, \mathbf{C}_B, \mathbf{U}, \mathbf{C}_{\text{circ}}, \end{array} \text{hint, aux, rand} \right)$

Fig. 7: Experiments Pre and Post for evasive circular LWE.

are distributions, $\Sigma \in \mathbb{R}^{\varphi m \times \varphi m}$ is positive definite, and f, f_A are PPT algorithms. Let Samp be a PPT algorithm which, on input 1^λ , outputs

$$(\mathbf{A} \in \mathcal{R}_q^{n_A \times m_A}, \mathbf{P} \in \mathcal{R}_q^{n_P \times m_P}, f_{\text{circ}}, \text{aux} \in \{0, 1\}^*)$$

where f_{circ} is a PPT algorithm. Denote

$$\begin{aligned} \text{Adv}_A^{\text{Pre}}(\lambda) &:= |\Pr[\text{Pre}_A^0(1^\lambda) = 1] - \Pr[\text{Pre}_A^1(1^\lambda) = 1]|, \\ \text{Adv}_B^{\text{Post}}(\lambda) &:= |\Pr[\text{Post}_B^0(1^\lambda) = 1] - \Pr[\text{Post}_B^1(1^\lambda) = 1]|, \end{aligned}$$

where the experiments Pre_A^b and Post_B^b are defined in Fig. 7. The $\text{EvCircLWE}_{\text{param}}$ assumption states that for any PPT Samp and \mathcal{B} there exists a PPT \mathcal{A} such that $\text{Adv}_A^{\text{Pre}}(\lambda) \geq \text{Adv}_B^{\text{Post}}(\lambda)/\text{poly}(\lambda) - \text{negl}(\lambda)$.

Definition 12 is similar to Definition 7, with the following differences (highlighted in Fig. 7 in gray).

1. there are additional circular LWE samples w.r.t. a matrix $\mathbf{M} = f_{\text{circ}}(\mathbf{S})$ which may be correlated with the secret \mathbf{S} , where f_{circ} is chosen by Samp , and
2. the hint functions f_A, f are not required to be public-coin (i.e. randomness used in the evaluation is not necessarily available to the distinguishers).

Using the notation of the original work, where an LWE secret is \mathbf{r} , the evasive circular LWE assumption in [HLL23] can be viewed as a special case of Definition 12 as follows: Let $\mathcal{R} = \mathbb{Z}$, the LWE secret distribution \mathcal{S} is such that $\mathbf{S} = \mathbf{S}_A = \mathbf{r}^T \sim (\mathcal{D}_{\mathbb{Z}, \sigma})^{1 \times n}$ is Gaussian, let the (probabilistic) hint functions

$$f_A = \emptyset, \quad f : \mathbf{r} \mapsto (\overline{\mathbf{A}}_{f_{he}}, \mathbf{r}^T \overline{\mathbf{A}}_{f_{he}} + \mathbf{e}_{f_{he}}^T, g(\mathbf{r})) \text{ mod } q$$

where $\overline{\mathbf{A}}_{fhe}$ is uniformly random, \mathbf{e}_{fhe} is Gaussian noise, and

$$g(\mathbf{r}) := \left(\mathbf{r}^T \overline{\mathbf{A}}_{fhe} + \mathbf{e}_{fhe}^T \right) \mathbf{R} - \text{bits}((\mathbf{r}^T, -1)^T) \otimes \mathbf{G} \bmod q.$$

for uniformly random \mathbf{R} . Let Samp output

$$\left(\overline{\mathbf{A}}', \mathbf{P}, f_{\text{circ}}, \text{aux} = () \right)$$

where

$$f_{\text{circ}} : \mathbf{r} \mapsto \overline{\mathbf{A}}_{\text{circ}} - (1, \text{bits}(g(\mathbf{r}))) \otimes \overline{\mathbf{G}} \bmod q$$

for some $\overline{\mathbf{A}}_{\text{circ}}$ determined by Samp , i.e., the circular matrix $\mathbf{M} = f_{\text{circ}}(\mathbf{r})$ depends on the hint component $g(\mathbf{r})$.

Unlike in our definition of public-coin evasive LWE (Definition 7), here, part of the randomness involved in evaluating the hint function f , specifically the noise \mathbf{e}_{fhe} and the matrix \mathbf{R} , are not available to the distinguisher. Indeed, hiding randomness from the distinguisher is necessary since otherwise the LWE secret \mathbf{r} would be efficiently recoverable from $\mathbf{r}^T \overline{\mathbf{A}}_{fhe} + \mathbf{e}_{fhe}^T \bmod q$, rendering both the Pre^b and Post^b experiments distinguishable (so that the assumption is vacuously true and presumably not useful).

On the other hand, while Definition 12 does not fall into our family of public-coin evasive LWE assumptions (Definition 7), it is easy to see that Definition 12 is indeed a special case of the private-coin binding evasive LWE family (Definition 8): The private-coin Samp computes $f_A(\mathbf{S}), f(\mathbf{S}), \mathbf{C}_A, \mathbf{C}_{\text{circ}}$ itself, and lets aux contain all of these.

Finally, we remark that it is possible to formulate a “public-coin evasive circular LWE assumption” in the style of Definition 7 which involves circularity in analogous sense of Definition 12 via \mathbf{C}_{circ} , but the functions f_A, f and f_{circ} remain public-coin. Since the evasive LWE in [HLL23] requires hint functions with secret random coins, it would also not be captured by such a public-coin evasive circular LWE.