Simple and General Counterexamples for Private-Coin Evasive LWE

Nico Döttling
CISPA

Abhishek Jain
NTT Research & JHU

Giulio Malavolta
Bocconi

Surya Mathialagan
MIT

Vinod Vaikuntanathan
MIT

February 27, 2025

Abstract

We present a simple counterexample to all known variants of the private-coin evasive learning with errors (LWE) assumption. Unlike prior works, our counterexample is direct, it does not use heavy cryptographic machinery (such as obfuscation or witness encryption), and it applies to *all variants* of the assumption. Our counterexample can be seen as a "zeroizing" attack against evasive LWE, calling into question the soundness of the underlying design philosophy.

1 Introduction

Modern lattice-based cryptography, initiated by the work of Ajtai [Ajt96], relies on *average-case* problems that come with a reduction from well-studied, and widely believed to be hard, *worst-case* problems on integer lattices. In other words, this line of research gives us a number of instances of lattice-based (average-case) hardness. The most popular such average-case problems are the Short Integer Solution (SIS) problem [MR04] and the Learning With Errors (LWE) problem [Reg05], as well as their ring and module counterparts [PR06, LPR10, LS15].

However, in recent years, cryptographic constructions have emerged, notably in the contexts of broadcast encryption [BV22] and witness encryption [CVW18], that draw from tools and techniques that were developed in the context of lattice-based cryptography, but for which proving security from lattice-based hardness assumptions (and indeed, any non-vacuous assumption at all) is out of the reach of our current understanding. While such constructions may reasonably be conjectured to be secure, this is an unsatisfying state of affairs as a break of such schemes may not lead to interesting new mathematical insights. This is in stark contrast to the win-win paradigm of (lattice-based) cryptography, where any attack directly gives us improved (worst-case) algorithms for well-studied mathematical problems.

Evasive LWE. To remedy this situation and provide a way to systematically assess the security of such cryptographic constructions, Wee [Wee22] and Tsabary [Tsa22] introduced the *evasive LWE* framework. Concretely, these works used it to argue the security of a new broadcast encryption and a new witness encryption scheme, respectively.

Assume in the following that we are working over the ring $\mathbb{Z}_q$. Evasive LWE is commonly stated as follows:[1] fix a sampler Samp which outputs $(\mathbf{P}, \mathbf{S}, \mathsf{aux})$, where $\mathbf{P}$ and $\mathbf{S}$ are matrices and aux is a bit-string, referred to as *auxiliary information*. Further let $\mathbf{B}$ be a *narrower* matrix (than $\mathbf{P}$) chosen uniformly at random, $\mathbf{E}, \mathbf{E}'$ be Gaussian error matrices of suitable size, and $\mathbf{U}, \mathbf{U}'$ be uniformly random matrices of suitable size. We will let the notation $\mathbf{B}^{-1}(\mathbf{P})$ denote a short Gaussian pre-image of $\mathbf{P}$ with respect to $\mathbf{B}$. In particular, $\mathbf{B}^{-1}(\mathbf{P})$ is a matrix with small entries such that $\mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P}) = \mathbf{P} \pmod q$.

The evasive LWE assumption postulates that if the pre-condition

$$(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathsf{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \mathbf{U}, \mathbf{U}', \mathsf{aux}) \tag{1.1}$$

holds, i.e. the two ensembles are computationally indistinguishable, then also the post-condition

$$(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \mathbf{U}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \tag{1.2}$$

holds.

It is not hard to see that if the pre-condition is false, so is the post-condition; i.e. a distinguisher for the pre-condition implies a distinguisher for the post-condition. Indeed, to construct a distinguisher for the post-condition, observe that using $\mathbf{D} := \mathbf{B}^{-1}(\mathbf{P})$, we can *expand* $(\mathbf{SB} + \mathbf{E})$ into

$$(\mathbf{SB} + \mathbf{E}) + \mathbf{E}' \cdot \mathbf{B}^{-1}(\mathbf{P}) = \mathbf{SP} + (\mathbf{E}\mathbf{B}^{-1}(\mathbf{P}) + \mathbf{E}') \approx_s \mathbf{SP} + \mathbf{E}',$$

where $\mathbf{E}'$ is a Gaussian "flooding" noise. Now invoking the distinguisher for the pre-condition does the job for us.[2] This is an instance of the so-called "zeroizing attacks" [CHL+15, MSZ16].

Indeed, zeroizing attacks, namely multiplying $\mathbf{SB} + \mathbf{E}$ and $\mathbf{B}^{-1}(\mathbf{P})$ and trying to distinguish the product from random, are the only known type of attacks against the post-condition[3]. This state of affairs inspired the formulation of the evasive LWE assumption which says that *the only way* to distinguish the post-condition is to come up with a distinguisher for the pre-condition. Specifically, it postulates that for every attack on the post-condition, there is a matching attack on the pre-condition with polynomially related success probability. Stated differently, the hope is that if there is an attack against the post-condition which does not translate into an attack against the pre-condition, it would reveal *interesting* structural insights which let us go beyond zeroizing attacks and could potentially even lead to a win-win situation.[4]

Public-Coin vs. Private-Coin Evasive LWE. Let us now dive into the details of the evasive LWE assumption a little further. The sampler Samp is a critical aspect of the evasive LWE assumption. In Wee's formulation [Wee22], the sampler generates the matrices $\mathbf{B}$, $\mathbf{P}$ and the auxiliary information aux, whereas the secret $\mathbf{S}$ is chosen to be independently and uniformly random. Futhermore, [Wee22] only considers a restricted class of *public-coin* samplers Samp where aux contains all the random coins used by Samp. Public-coin evasive LWE has seen a range of applications, including

[1] We describe the assumption following a variant from the recent work of Brzuska, Ünal and Woo [BUW24].

[2] We are ignoring some technicalities, e.g. to argue that $\mathbf{U} \cdot \mathbf{B}^{-1}(\mathbf{P})$ is indeed statistically close to uniform, invoking the leftover hash lemma, one needs $\mathbf{B}$ to be sufficiently wide. This is indeed the regime that LWE and evasive LWE are used.

[3] That is, short of ignoring $\mathbf{B}^{-1}(\mathbf{P})$ and distinguishing the rest from random which would be an attack on plain LWE.

[4] Compare this situation to the recent attacks on the SIDH problem [CD23]; the algorithm underlying the attack led to the construction of a provably secure signature scheme from isogenies [DLRW24].

optimal broadcast encryption [Wee22], multi-authority ABE [WWW22], and unbounded-depth ABE [HLL23]. However, public-coin evasive LWE only let us get a little beyond LWE-land; in particular, it was completely unclear whether simple, natural, and apparently secure constructions of witness encryption (such as [CVW18]) can be proven secure.

Motivated by this state of affairs, Vaikuntanathan, Wee and Wichs [VWW22] introduced a *private-coin* variant of evasive LWE, where the random coins used by Samp are not included in aux and are hence kept hidden in both the pre- and the post-conditions. This variant has proven to be extremely fruitful in constructing several highly sought-after cryptographic primitives such as witness encryption and null-IO [Tsa22, VWW22], adaptively sound SNARGs for UP [MPV24, JKLM25], non-adaptive SNARG for NP [JKLM25], and many more. A recent work [BDJ+24] realizes a novel notion of *pseudorandom obfuscation* from private-coin evasive LWE and additional standard assumptions. They further show that one can construct IO itself assuming evasive LWE and bilinear maps. Additionally, [BDJ+24] shows how to construct succinct witness encryption, an object that has so far been out of reach from even IO, assuming evasive LWE. Finally, Agrawal, Kumari and Yamada [AKY24] construct the related notion of compact pseudorandom functional encryption from private-coin evasive LWE.

At a technical level, all the above-mentioned works (both in the public- and the private-coin setting) use evasive LWE to compress noise: In the evasive LWE pre-condition, the Gaussian error $\mathbf{E}'$ is used to *noise-flood* some residual, low-norm, term which usually depends on a secret. Then, by appealing to the evasive LWE assumption, one hopes that the *expanded error* $\mathbf{ED}$ in the post-condition (where $\mathbf{D} := \mathbf{B}^{-1}(\mathbf{P})$) has the same effect, i.e. it *pseudo-drowns* the problematic secret-dependent residual term.

One might wonder if the low-rankness of $\mathbf{ED}$ leads to an attack right away. This is not the case: $\mathbf{ED}$ is not given to the distingisher by itself; it comes with the "LWE term" $\mathbf{SP}$ which hides the low-rank structure in the expanded error. Indeed, this gets to the heart of why LWE does not suffer from a trivial rank attack. Nevertheless, looking ahead, we will attack precisely this pseudo-drowning intuition, and will crucially use the fact that $\mathbf{ED}$ is low-rank *over the integers* and not just mod q.

Prior Attacks on Evasive LWE, Take 1. It was already established in [VWW22] that private-coin evasive LWE, in its full generality, is false. The core of the issue is that the auxiliary information aux can encrypt information about $\mathbf{S}$ or the matrix $\mathbf{B}$, and this information might be inaccessible in the pre-condition but accessible in the post-condition. Specifically, the Gaussian pre-image $\mathbf{B}^{-1}(\mathbf{P})$ can serve as a witness or a decryption hint to decrypt this information. In [VWW22], a contrived example of such an auxiliary information was provided under heuristic obfuscation assumptions. This counterexample has since been simplified and refined to rely on null-IO in [BUW24] or even witness encryption in [BDJ+24].

Prior Attacks on Evasive LWE, Take 2. While it seemed reasonable to conjecture that we can avoid these types of obfuscation-based attacks by restricting ourselves to "natural" distributions of $\mathbf{P}, \mathbf{S}$ and aux, the recent work of Brzuska, Ünal and Woo [BUW24] demonstrated simple algebraic attacks against variants of private-coin evasive LWE, as summarized in Fig. 1. In particular, they give counterexamples in the following settings:

- The matrix $\mathbf{B}$ is hidden, but $\mathbf{P}$ is only partially available to the distinguisher, in both the pre- and post-conditions. This regime is captured by the blue cross in Fig. 1.

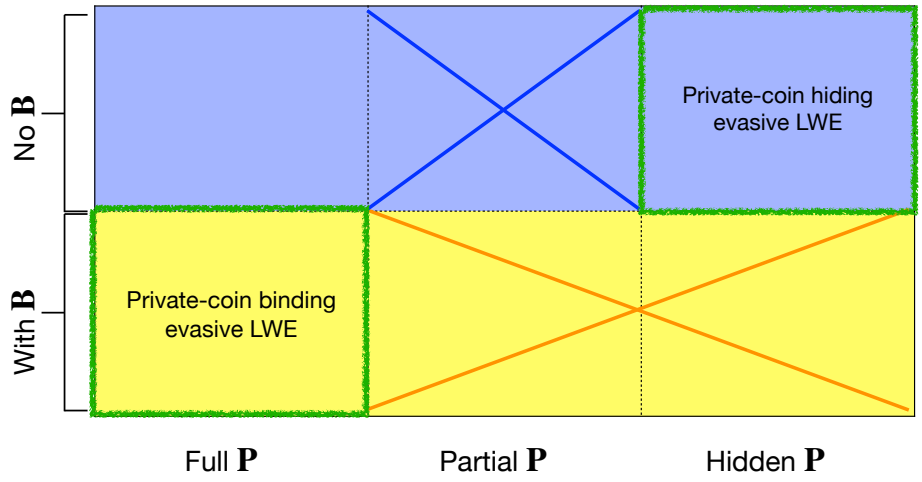

Figure 1: The above figure is adapted from [BUW24, Figure 1(b)] (BUW). BUW partition the private-coin evasive LWE based on whether $\mathbf{B}$ is public or hidden (in both the pre- and post-conditions); and whether $\mathbf{P}$ is public, hidden or "partially" hidden (in both the pre- and post-conditions). BUW then give simple counterexamples which rule out the different private-coin variants. Each counterexample is denoted by a cross of different color. Finally, the proposed classes of private-coin evasive LWE that survive their attack are indicated by the green outline. For simplicity of exposition, we do not mention an additional attack of BUW that allows the Samp algorithm takes $\mathbf{B}$ as input, although we note that our attack subsumes that setting as well.

- The matrix $\mathbf{B}$ is available, but $\mathbf{P}$ is either partially or fully hidden from the distinguisher, in both the pre- and post-conditions. This regime is captured by the yellow cross in Fig. 1.

To summarize, their work leaves two variants of private-coin evasive LWE that we do not yet have counterexamples for. We state them here informally.

Private-coin "binding" evasive LWE. This corresponds to the bottom-left corner of the grid in Fig. 1 (and is identical to the version described in the beginning of the paper, equations (1.1) and (1.2)). In this variant, both $\mathbf{B}$ and $\mathbf{P}$ are available in the clear in both the pre- and post-conditions. The assumption then states that for any sampler Samp which takes a security parameter λ as input, and outputs $(\mathbf{P}, \mathbf{S}, \mathsf{aux})$:

$$\begin{aligned} &\textbf{if} && (\mathbf{B}, \mathbf{P}, \mathbf{SB}+\mathbf{E}, \mathbf{SP}+\mathbf{E}', \mathsf{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{C}', \mathsf{aux}), \\ &\textbf{then} && (\mathbf{B}, \mathbf{P}, \mathbf{SB}+\mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \end{aligned}$$

Private-coin hiding evasive LWE. This corresponds to the top-right corner of the grid in Fig. 1. In this variant, both $\mathbf{B}$ and $\mathbf{P}$ are *hidden* in both the pre- and post-conditions. The assumption then states that for any sampler Samp which takes a security parameter λ as input, and outputs $(\mathbf{P}, \mathbf{S}, \mathsf{aux})$, if $(\mathbf{P}, \mathsf{aux}) \approx_c (\mathbf{P}+\mathbf{R}, \mathsf{aux})$ (where $\mathbf{R}$ is sampled from some bounded norm

distribution, e.g. uniform over an interval), we have that:

$$\begin{aligned} &\textbf{if} && (\mathbf{SB}+\mathbf{E}, \mathbf{SP}+\mathbf{E}', \mathsf{aux}) \approx_c (\mathbf{C}, \mathbf{C}', \mathsf{aux}), \\ &\textbf{then} && (\mathbf{SB}+\mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{C}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \end{aligned}$$

Note that neither version is stronger than the other; hiding both $\mathbf{P}$, for example, hurts the distinguisher in the pre-condition just as well as it hurts the distinguisher in the post-condition. The hiding version suffices to prove the security of several protocols, e.g. witness encryption [VWW22], SNARGs for UP [MPV24], and SNARG for NP [JKLM25]; and the binding version suffices for constructions of pseudorandom obfuscation [BDJ+24].

While the obfuscation-based counterexamples still apply to these variants, it seemed reasonable to assume security with essentially any "natural distribution" of aux that is not specifically constructed to contain a counterexample program.

Implications to Constructions. None of the known attacks, including ours, break the constructions of witness encryption and null-IO [VWW22, Tsa22], adaptively sound SNARGs for UP [MPV24, JKLM25], or non-adaptive SNARG for NP [JKLM25]. We will update this section as more information becomes available.

2 Our Results

In this work, we show a simple and general counterexample to evasive LWE that *simultaneously* rules out *all* private-coin variants captured in Fig. 1. We show a counterexample which satisfies the "strongest" **if** condition (where $\mathbf{B}$ and $\mathbf{P}$ are available to the distinguisher) assuming LWE, and we demonstrate an attack on the "weakest" **then** condition (where $\mathbf{B}$ and $\mathbf{P}$ are hidden from the distinguisher).

In a nutshell, our counterexample demonstrates that the *pseudo-drowning* heuristic, which is at the core of all security proofs that rely on evasive LWE, is not sound in the case of private-coin evasive LWE. This is the case even for very simple and benign-looking distributions of the auxiliary information aux, and for uniformly random $\mathbf{P}$ and $\mathbf{B}$, and a Gaussian $\mathbf{S}$.

2.1 Our Counterexample

We now describe our counterexample. Let q be an odd prime which is superpolynomial in the security parameter λ; we let $\mathbb{Z}_q := \{0, 1, \ldots, q-1\}$. We define the following sampler Samp:

- Sample a matrix $\mathbf{P} \leftarrow \mathbb{Z}_q^{n \times k}$.
- Sample a matrix $\mathbf{S} \leftarrow D_{\mathbb{Z},\sigma}^{\ell \times n}$, where $D_{\mathbb{Z},\sigma}$ is the discrete Gaussian over $\mathbb{Z}$ with standard deviation σ.
- Sample an $\ell \times k$ matrix $\mathbf{T} \in \mathbb{Z}_q^{\ell \times k}$ such that all the entries are uniformly random numbers from $[0, 1, \ldots, \lfloor q/2 \rfloor]$.
- Let $\mathsf{aux} = \mathbf{SP} - 2\mathbf{T} \bmod q$, and output $\mathbf{P}, \mathbf{S}, \mathsf{aux}$.

Pre-condition: We argue that the pre-condition holds in the *strongest* setting where $\mathbf{B}$ and $\mathbf{P}$ are available to the distinguisher. We sketch the argument here via the following hybrid distributions:

$$\begin{aligned}
&(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}'', \mathsf{aux} = \mathbf{SP} - 2\mathbf{T}) \\
&\approx_s (\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + 2\widetilde{\mathbf{E}} + \mathbf{E}'', \mathsf{aux} = \mathbf{SP} + 2\widetilde{\mathbf{E}} - 2\mathbf{T}) \qquad (2.1) \\
&\approx_c (\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{V} + \mathbf{E}'', \mathsf{aux} = \mathbf{V} - 2\mathbf{T}). \qquad (2.2)
\end{aligned}$$

where $\mathbf{U}$ and $\mathbf{V}$ are random matrices over $\mathbb{Z}_q$. Here, (2.1) follows via a noise flooding argument by picking $|\mathbf{E}''|_\infty \gg |\widetilde{\mathbf{E}}|_\infty$, and (2.2) follows from the fact that LWE with even error is as hard as LWE which in turn follows from the fact that 2 is invertible mod q (since q is odd).[5]

At this point, it suffices to show that

$$(\mathbf{V} + \mathbf{E}'', \mathbf{V} - 2\mathbf{T}) \approx_c (\mathbf{C}', \mathbf{V} - 2\mathbf{T})$$

for a uniformly random matrix $\mathbf{C}'$. Essentially, we use the fact that the difference between the two quantities, given by $2\mathbf{T} + \mathbf{E}''$, is statistically close to uniform over $\mathbb{Z}_q$. Intuitively, the error term $\mathbf{E}''$ "smudges" $2\mathbf{T}$ to hide the fact that the term is even. Then, by a tail bound, we can show that $|2\mathbf{T} + \mathbf{E}''| < q$ with high probability by choosing the norm of $\mathbf{E}''$ to be superpolynomially smaller than q. Therefore, we have that $2\mathbf{T} + \mathbf{E}''$ is statistically close to uniformly random over $\mathbb{Z}_q$. We refer the reader to Lemma 5.2 for a detailed analysis.

To establish private-coin hiding evasive LWE, we additionally need to show that $(\mathbf{P}, \mathsf{aux}) \approx_c (\mathbf{P} + \mathbf{R}, \mathsf{aux})$ for bounded-norm $\mathbf{R}$. We show this via a similar argument as the main pre-condition. We refer the reader to Lemma 5.3 for a detailed analysis.

Post-condition: We now show that there exists a distinguisher which distinguishes $(\mathbf{SB}+\mathbf{E}, \mathbf{D}, \mathsf{aux})$ from $(\mathbf{C}, \mathbf{D}, \mathsf{aux})$, where $\mathbf{C}$ is a uniformly random matrix and $\mathbf{D} = \mathbf{B}^{-1}(\mathbf{P})$. The distinguisher does the following: compute

$$\mathbf{W} = (\mathbf{CD} - \mathsf{aux} \mod q) \mod 2$$

and check if the first row of $\mathbf{W}$ is in the span of the rows of $\mathbf{D} \mod 2$. If yes, output 1, else output 0.

When the distinguisher gets $\mathbf{SB} + \mathbf{E}$, it computes

$$(\mathbf{SB} + \mathbf{E}) \cdot \mathbf{D} - \mathsf{aux} = (\mathbf{SB} + \mathbf{E}) \cdot \mathbf{D} - (\mathbf{SP} - 2\mathbf{T}) = \mathbf{ED} + 2\mathbf{T} \pmod q$$

Since $\mathbf{D}, \mathbf{E}$ are Gaussian with small variance, with high probability,

$$||\mathbf{ED} + 2\mathbf{T}||_\infty < q\,,$$

that is $\mathbf{ED} + 2\mathbf{T} \pmod q$ is the same as $\mathbf{ED} + 2\mathbf{T}$ as an integer. Hence, reducing modulo 2 we obtain $\mathbf{W} = \mathbf{ED} \pmod 2$. It is easy to see that the first row of $\mathbf{W}$ is in fact in the row span of $\mathbf{D} \pmod 2$. Hence, the distinguisher will output 1 with high probability.

[5] We note that the use of LWE to establish the pre-condition in our counterexample is expected since the pre-condition asserts a computational statement.

On the other hand, when the distinguisher gets a uniformly random matrix $\mathbf{C}$, we can argue via leftover-hash lemma that the expression

$$\mathbf{CD} - (\mathbf{SP} - 2\mathbf{T}) \bmod q$$

is in fact statistically close to *uniform* over $\mathbb{Z}_q$. Therefore, taking the result mod 2 will give a uniformly random matrix mod 2. The first row of this matrix is not in the span of $\mathbf{D} \pmod 2$ with high probability since $\mathbf{D}$ is a wide matrix. We refer the reader to Lemma 5.1 for a detailed analysis.

In a nutshell, even though the distribution in the pre-condition is pseudorandom, the distribution in the post-condition allows us to mount a "zeroizing" attack, recovering a quantity that does not wrap around mod q.

2.2 Comparison to Prior Attacks

We first recap the idea of the prior attacks on private-coin evasive LWE.

Obfuscation-based counterexample. VWW [VWW22] noted that the private-coin evasive LWE assumption was already prone to heuristic obfuscation-based counterexamples. The work of Brzuska, Ünal and Woo [BUW24] then strengthened this counterexample to only rely on null-IO and LWE. We briefly describe the counterexample, following [BUW24, Remark 3]. Consider the sampler Samp that works as follows:

- Sample a wide matrix $\mathbf{P}$ in $\mathbb{Z}_q^{n\times k}$.
- Sample a random tall matrix $\mathbf{S} \leftarrow \mathbb{Z}_q^{\ell\times n}$ and a random error term $\mathbf{E}''$ (with variance super-polynomially larger than $\mathbf{E}'$).
- Set $\mathbf{W} = \mathbf{SP} + \mathbf{E}''$.
- Construct a circuit $C_{\mathbf{W}}$ which on input $\mathbf{M}_1 \in \mathbb{Z}_q^{\ell\times m}$ and $\mathbf{M}_2 \in \mathbb{Z}_q^{m\times k}$ outputs 1 if $\mathbf{W} \approx_c \mathbf{M}_1 \cdot \mathbf{M}_2$, and 0 otherwise.
- Output $(\mathbf{S}, \mathsf{aux} = \mathcal{O}(C_{\mathbf{W}}))$, where $\mathcal{O}$ is a null-IO scheme.

In the post-condition, the distinguisher can simply run aux on inputs $(\mathbf{C}, \mathbf{D})$. If $\mathbf{C} = \mathbf{SB} + \mathbf{E}$, then the program outputs 1. If $\mathbf{C}$ is random, it outputs 0 with high probability. In the pre-condition, one can argue via LWE that $\mathbf{W}$ is indistinguishable from random. For random $\mathbf{W}$, note that $C_{\mathbf{W}}$ is equal to the zero function if we choose $\ell, k \gg m$. Therefore, one can invoke null-IO to replace the obfuscation with an all-zeroes function. It is then clear that aux does not help in distinguishing the pre-condition.

Algebraic attacks of Brzuska-Ünal-Woo. In the algebraic attacks of BUW [BUW24], the counterexamples rely heavily on "planting a short vector". To illustrate this, we recap one of their counterexamples in the setting where $\mathbf{P}$ is hidden (this corresponds to the yellow cross in Fig. 1). The sampler Samp on input 1^λ does the following:

- Sample $\mathbf{P}_1$ along with a short vector $\mathbf{u}$ such that $\mathbf{P}_1\mathbf{u} = \mathbf{0}$.

- Sample a random matrix $\mathbf{R}$.
- Output

$$\mathbf{P} = \left(\mathbf{P}_1, \quad \mathbf{P}_2 = \begin{pmatrix} \mathbf{u}^T \\ \mathbf{R} \end{pmatrix}\right).$$

The pre-condition can be established via routine application of LWE. For the post-condition, given

$$(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P})),$$

note that one can compute $\mathbf{P}_2 \leftarrow \mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P})$ to obtain the vector $\mathbf{u}$. and

$$(\mathbf{SB} + \mathbf{E}) \cdot \mathbf{B}^{-1}(\mathbf{P}) = (\mathbf{V}, \mathbf{W}) \approx (\mathbf{SP}_1, \mathbf{SP}_2)$$

Now, one can test if

$$\mathbf{V} \cdot \mathbf{u} \approx \mathbf{SP}_2 \cdot \mathbf{u} = 0 \bmod q\,.$$

Here, the counterexample relies heavily on planting a short vector $\mathbf{u}$ that enables a zeroizing attack. The vector $\mathbf{u}$ lives in the matrix $\mathbf{P}$ which is hidden in the precondition, but is revealed in the post-condition. Our attack does not plant any short vectors, works when $\mathbf{P}$ is completely out in the open, and still manages to find a zeroizing attack.

3 Preliminaries

We write $\lambda \in \mathbb{N}$ to denote the security parameter. Given a finite set W, $\mathcal{U}(W)$ denotes the uniform distribution over W. We write $x \leftarrow W$ for $x \leftarrow \mathcal{U}(W)$. Given distributions $\mathcal{D}_1, \mathcal{D}_2$, we write $\mathcal{D}_1 \approx_c \mathcal{D}_2$ to denote computational indistinguishability and $\mathcal{D}_1 \approx_s \mathcal{D}_2$ to denote statistical indistinguishability. $D_{\mathbb{Z},\sigma}$ denotes the discrete Gaussian over $\mathbb{Z}$ with parameter σ; that is, the distribution which assigns mass proportional to $\exp(-\pi x^2/\sigma^2)$ to each $x \in \mathbb{Z}$.

We recall some inequalities that we will rely on in our analysis.

Lemma 3.1. *For any $\sigma > 0$ and $n \in \mathbb{N}$,*

$$\Pr_{x \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}}[|x| \geq \sigma\sqrt{n}] \leq 2^{-n}\,.$$

Lemma 3.2 ([CVW18, Lemma 3.2])**.** *For all $y \in \mathbb{Z}$ and $\sigma \in \mathbb{R}$, the statistical distance between $\mathcal{D}_{\mathbb{Z},\sigma}$ and $\mathcal{D}_{\mathbb{Z},\sigma} + y$ is at most y/σ.*

3.1 LWE Assumption and Trapdoor Sampling

Definition 3.3 (Learning With Errors (LWE))**.** Given $n, m, q \in \mathbb{N}$ and $\sigma > 0$ with $n, m \in \mathsf{poly}(\lambda)$, the LWE assumption $\mathsf{LWE}_{n,m,q,\sigma}$ asserts that

$$(\mathbf{A}, \mathbf{s}^T\mathbf{A} + \mathbf{e}) \approx_c (\mathbf{A}, \mathbf{b}),$$

where $\mathbf{s} \leftarrow \mathcal{U}(\mathbb{Z}_q^n)$, $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times m})$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^m$, and $\mathbf{b} \leftarrow \mathcal{U}(\mathbb{Z}_q^m)$.

In this work, we also need the following variant of LWE which was proven in [ACPS09].

Lemma 3.4 (LWE with small secrets, [ACPS09]). *Given $n, m, q \in \mathbb{N}$ and $\sigma > 0$ with $n, m \in \mathsf{poly}(\lambda)$, we have*

$$(\mathbf{A}, \mathbf{s}^T\mathbf{A} + \mathbf{e}) \approx_c (\mathbf{A}, \mathbf{b}),$$

where $\mathbf{s} \leftarrow \mathcal{D}^n_{\mathbb{Z},\sigma}$, $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{n\times m})$, $\mathbf{e} \leftarrow \mathcal{D}^m_{\mathbb{Z},\sigma}$, and $\mathbf{b} \leftarrow \mathcal{U}(\mathbb{Z}_q^m)$ assuming $\mathsf{LWE}_{n,m,q,\sigma}$.

Given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ for $m \geq 2n \log q$, a vector $\mathbf{y} \in \mathbb{Z}_q^n$, and $\sigma > 0$, we use $\mathbf{A}^{-1}(\mathbf{y}, \sigma)$ to denote the distribution of a vector $\mathbf{d}$ sampled from $\mathcal{D}^m_{\mathbb{Z},\sigma}$ conditioned on $\mathbf{Ad} = \mathbf{y} \pmod q$. (Vectors satisfying the condition exist except with probability $\mathsf{negl}(\mu)$.) We extend this notation to matrices $\mathbf{Y} \in \mathbb{Z}_q^{n\times k}$ in the natural way (i.e., columnwise). We sometimes suppress σ when it is clear from context.

3.2 Min-Entropy and Leftover Hash Lemma

The *min-entropy* of a discrete variable X is defined as $H_\infty(X) = -\log(\max_x \Pr[X = x])$. We will also use the notion of *conditional min-entropy* of a random variable X conditioned on a variable Y is defined as follows:

$$\widetilde{H}_\infty(X|Z) = -\log \mathbb{E}_z[\max_x \Pr[X = x|Z = z]] = -\log \mathbb{E}_z\left[2^{-H_\infty(X|Z=z)}\right].$$

Lemma 3.5 (Chain rule, [DORS03]). *Let X, Y, Z be random variables. Then,*

$$\tilde{H}_\infty(X|Y, Z) \leq \tilde{H}_\infty(X, Y|Z) - |Y|$$

where $|Y|$ is the bit length of the value of Y.

Definition 3.6 (Extractor). A function $\mathsf{Ext} : \{0,1\}^d \times \mathcal{X} \to \{0,1\}^\ell$ is a *strong seeded average-case (k, ε)-extractor*, if it holds for all random variables X with support $\mathcal{X}$ and Z defined on some finite support that if $\widetilde{H}_\infty(X|Z) \geq k$, then it holds that the statistical distance between the following distributions is at most ε:

$$\{\mathsf{seed}, \mathsf{Ext}(\mathsf{seed}, X), Z\} \approx_s \{\mathsf{seed}, U, Z\}$$

where $\mathsf{seed} \leftarrow \{0,1\}^d$ and $U \leftarrow \{0,1\}^\ell$.

Recall that a hash function family $\mathcal{H}$ of functions $h : \mathcal{X} \to \mathcal{Y}$ is a *universal hash family* if for all $x \neq x' \in \mathcal{X}$, it holds that:

$$\Pr_{h\leftarrow\mathcal{H}}[h(x) = h(x')] \leq \frac{1}{|\mathcal{Y}|}.$$

Lemma 3.7 (Leftover hash-lemma, [DORS03]). *Let X be a random-variable supported on a finite set $\mathcal{X}$ and let Z be a (possibly correlated) random variable supported on a finite set $\mathcal{Z}$ such that $\widetilde{H}_\infty(X|Z) \geq k$. Let $\mathcal{H}$ be a universal hash family with functions $h : \mathcal{X} \to \{0,1\}^\ell$, where $\ell \geq k - 2\log\left(\frac{1}{\varepsilon}\right)$. Then $\mathcal{H}$ is a seeded strong average-case (k, ε)-extractor.*

4 Private-Coin Evasive LWE Variants

In this section, we recap two variants of private-coin evasive LWE from [BUW24].

Definition 4.1 (Private-Coin Binding Evasive LWE, [BUW24, Definition 8])**.** Let $m, n, k, \ell > 0$ be integers and let q be a modulus. Let $\tau, \sigma, \sigma' > 0$. Let Samp be an algorithm which takes 1^λ, and outputs a matrix $\mathbf{P} \in \mathbb{Z}_q^{n \times k}$, a matrix $\mathbf{S} \in \mathbb{Z}_q^{\ell \times n}$ and auxiliary information aux. Let

$$\begin{aligned}
&\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m} \\
&\mathbf{D} \leftarrow \mathbf{B}^{-1}(\mathbf{P}, \tau) \\
&(\mathbf{S}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda, \mathbf{P}) \\
&\mathbf{E} \leftarrow D_{\mathbb{Z},\sigma}^{\ell \times m}, \mathbf{E}' \leftarrow D_{\mathbb{Z},\sigma'}^{\ell \times k} \\
&\mathbf{C} \leftarrow \mathbb{Z}_q^{\ell \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{\ell \times k}
\end{aligned}$$

For PPT distinguishers $\mathcal{A}'$ and $\mathcal{A}$ define the following functions:

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{Pre}}_{\mathcal{A}'}(\lambda) &= |\Pr[\mathcal{A}'(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathsf{aux}) = 1] - \Pr[\mathcal{A}'(\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{C}', \mathsf{aux}) = 1]| \\
\mathsf{Adv}^{\mathsf{Post}}_{\mathcal{A}}(\lambda) &= |\Pr[\mathcal{A}'(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{D}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}'(\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{D}, \mathsf{aux}) = 1]|
\end{aligned}$$

We say that the binding evasive LWE assumption $\mathsf{evLWE}(q, m, n, k, \ell, \mathsf{Samp}, \tau, \sigma, \sigma')$ holds, if there exists polynomial Q such that $\deg_\lambda(Q) \leq c \cdot \deg_\lambda(|\mathsf{Samp}|)$ (for some universal constant $c > 0$), such that for every PPT distinguisher $\mathcal{A}$ there exists a PPT distinguisher $\mathcal{A}'$ such that

$$\mathsf{Adv}^{\mathsf{Pre}}_{\mathcal{A}'}(\lambda) \geq \mathsf{Adv}^{\mathsf{Post}}_{\mathcal{A}}(\lambda)/Q(\lambda) - \mathsf{negl}(\lambda)$$

and $\mathsf{time}(\mathcal{A}') \leq \mathsf{time}(\mathcal{A}) \cdot Q(\lambda)$.

Definition 4.2 (Private-Coin Hiding Evasive LWE, [BUW24, Definition 9])**.** Let $m, n, k, \ell > 0$ be integers and let q be a modulus. Let $\tau, \sigma, \sigma' > 0$. Let Samp be an algorithm which takes 1^λ, and outputs a matrix $\mathbf{P} \in \mathbb{Z}_q^{n \times k}$, a matrix $\mathbf{S} \in \mathbb{Z}_q^{\ell \times n}$ and auxiliary information aux. Let

$$\begin{aligned}
&\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m} \\
&\mathbf{D} \leftarrow \mathbf{B}^{-1}(\mathbf{P}, \tau) \\
&(\mathbf{S}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda, \mathbf{P}) \\
&\mathbf{E} \leftarrow D_{\mathbb{Z},\sigma}^{\ell \times m}, \mathbf{E}' \leftarrow D_{\mathbb{Z},\sigma'}^{\ell \times k} \\
&\mathbf{C} \leftarrow \mathbb{Z}_q^{\ell \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{\ell \times k} \\
&\mathbf{R} \leftarrow \mathcal{U}([\kappa])^{n \times k}
\end{aligned}$$

For PPT distinguishers $\mathcal{A}'$ and $\mathcal{A}$ define the following functions:

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{Pre1}}_{\mathcal{A}'}(\lambda) &= |\Pr[\mathcal{A}'(\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathsf{aux}) = 1] - \Pr[\mathcal{A}'(\mathbf{C}, \mathbf{C}', \mathsf{aux}) = 1]| \\
\mathsf{Adv}^{\mathsf{Pre2}}_{\mathcal{A}'}(\lambda) &= |\Pr[\mathcal{A}'(\mathbf{P}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}'(\mathbf{P} + \mathbf{R}, \mathsf{aux}) = 1]| \\
\mathsf{Adv}^{\mathsf{Post}}_{\mathcal{A}}(\lambda) &= |\Pr[\mathcal{A}'(\mathbf{SB} + \mathbf{E}, \mathbf{D}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}'(\mathbf{C}, \mathbf{D}, \mathsf{aux}) = 1]|
\end{aligned}$$

We say that the hiding evasive LWE assumption $\mathsf{evLWE}(q, m, n, k, \ell, \mathsf{Samp}, \kappa, \tau, \sigma, \sigma')$ holds, if there exists polynomial Q such that $\deg_\lambda(Q) \leq c \cdot \deg_\lambda(|\mathsf{Samp}|)$ (for some universal constant $c > 0$), such that for every PPT distinguisher $\mathcal{A}$ there exists a PPT distinguisher $\mathcal{A}'$ such that

$$\mathsf{Adv}^{\mathsf{Pre1}}_{\mathcal{A}'}(\lambda) + \mathsf{Adv}^{\mathsf{Pre2}}_{\mathcal{A}'}(\lambda) \geq \mathsf{Adv}^{\mathsf{Post}}_{\mathcal{A}}(\lambda)/Q(\lambda) - \mathsf{negl}(\lambda)$$

and $\mathsf{time}(\mathcal{A}') \leq \mathsf{time}(\mathcal{A}) \cdot Q(\lambda)$.

5 Our Counterexample

Parameters. Let m, n, k, ℓ be polynomial in λ, and we will assume that m and n are much smaller than k and ℓ. Additionally, $m \geq 2n \log q$. Let σ'/σ be superpolynomial. We also choose q to be an odd prime of size superpolynomial in $\sigma, \sigma', m, n, \ell, k, \lambda, \kappa$ satisfying that $q/(\kappa\sigma)$ is superpolynomial.

Construction 5.1. We define the following sampler $\mathsf{Samp}(1^\lambda)$ for our counterexample as follows:

- Sample a matrix $\mathbf{P} \leftarrow \mathbb{Z}_q^{n \times k}$.
- Sample a matrix $\mathbf{S} \leftarrow D_{\mathbb{Z},\sigma}^{\ell \times n}$.
- Sample a $\ell \times k$ matrix $\mathbf{T} \in \mathbb{Z}_q^{\ell \times k}$ such that all the entries are uniformly random numbers from $[0, 1, \ldots, \lfloor q/2 \rfloor]$.
- Let $\mathsf{aux} = \mathbf{SP} - 2\mathbf{T}$, and output $\mathbf{S}, \mathsf{aux}$.

Analyzing the post-condition. We show that the post-condition for both assumptions do not hold with the following attack.

Lemma 5.1. *There exists an efficient distinguisher $\mathcal{A}$ such that*

$$|\Pr[\mathcal{A}(\mathbf{SB} + \mathbf{E}, \mathbf{D}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathbf{C}, \mathbf{D}, \mathsf{aux}) = 1]| = 1 - \mathsf{negl}(\lambda),$$

where $\mathbf{D} \leftarrow D_{\mathbb{Z},\tau}^{m \times k}\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{P} = \mathbf{B} \cdot \mathbf{D}, \mathbf{E} \leftarrow D_{\mathbb{Z},\sigma}^{\ell \times m}, \mathbf{C} \leftarrow \mathbb{Z}_q^{\ell \times m}$, and $(\mathbf{S}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda, \mathbf{P})$ (as defined in Construction 5.1).

We give the full proof this theorem in Section 5.1

Analyzing the pre-condition. We then show that with respect to Samp in Construction 5.1, the pre-conditions for both variants of private-coin evasive LWE (Definitions 4.1 and 4.2) hold. We show this by establishing Lemmas 5.2 and 5.3 assuming LWE.

Lemma 5.2. *Let q be odd. Let σ, σ' be such that σ/σ' is superpolynomial and q/σ is superpolynomial. Assuming that $\mathsf{LWE}_{n,k,q,\sigma}$ holds, the pre-condition of evasive LWE holds, i.e. for all distinguishers $\mathcal{A}$,*

$$|\Pr[\mathcal{A}(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{C}', \mathsf{aux}) = 1]| \leq \mathsf{negl}(\lambda)$$

where

$$\begin{aligned}
&\mathbf{D} \leftarrow D_{\mathbb{Z},\tau}^{m \times k}, \mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{P} = \mathbf{B} \cdot \mathbf{D} \\
&(\mathbf{S}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda, \mathbf{P}) \text{ as in Construction 5.1}, \\
&\mathbf{E} \leftarrow D_{\mathbb{Z},\sigma}^{\ell \times m}, \mathbf{E}' \leftarrow D_{\mathbb{Z},\sigma'}^{\ell \times k} \\
&\mathbf{C} \leftarrow \mathbb{Z}_q^{\ell \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{\ell \times k}
\end{aligned}$$

The above shows that the pre-condition of private-coin binding evasive LWE (Definition 4.1) holds. To additionally show that the pre-condition of private-coin hiding evasive LWE (Definition 4.2) holds, we show the following claim.

Lemma 5.3. *Let q be odd. Let σ, σ' be such that σ/σ' is superpolynomial and q/σ is superpolynomial. Let $q/(\sigma\kappa)$ be superpolynomial. Assuming that $\mathsf{LWE}_{n,k,q,\sigma}$, for all distinguishers $\mathcal{A}$,*

$$|\Pr[\mathcal{A}(\mathbf{P}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathbf{P} + \mathbf{R}, \mathsf{aux}) = 1]| \leq \mathsf{negl}(\lambda),$$

where

$$\begin{aligned}
&\mathbf{D} \leftarrow D_{\mathbb{Z},\tau}^{m\times k}, \mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{P} = \mathbf{B}\cdot\mathbf{D},\\
&(\mathbf{S}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda, \mathbf{P}) \text{ as in Construction 5.1},\\
&\mathbf{R} \leftarrow \mathcal{U}([\kappa])^{n\times k}.
\end{aligned}$$

We show the proofs of both of the above claims in Section 5.2. We prove the claims using noise-flooding and LWE. We also crucially rely on the fact that q is prime, and hence 2 is an invertible scalar over $\mathbb{Z}_q$.

5.1 Attack on Post-Condition

In this section, we give the proof of Lemma 5.1.

Proof. The distinguisher on input $\mathbf{C}, \mathbf{D}$, and $\mathsf{aux} = \mathbf{Q}$, does the following:

- Compute $\mathbf{W} = (\mathbf{CD} - \mathbf{Q} \mod q) \mod 2$.
- Check if the first row of $\mathbf{W}$ is in the span of the columns of $\mathbf{D} \mod 2$. If yes, output 1. Else, output 0.

Real case: If $\mathbf{C} = \mathbf{SB} + \mathbf{E}$, then

$$\begin{aligned}
&\mathbf{CD} - \mathbf{Q} \mod q\\
&= (\mathbf{SB} + \mathbf{E})\mathbf{D} - (\mathbf{SP} - 2\mathbf{T}) \mod q\\
&= \mathbf{SP} + \mathbf{ED} - \mathbf{SP} + 2\mathbf{T} \mod q\\
&= \mathbf{ED} + 2\mathbf{T} \mod q.
\end{aligned}$$

By a Gaussian tail bound (Lemma 3.1) and union bound, it is easy to see that $||\mathbf{D}||_\infty \leq \tau\sqrt{t}$ with probability at least $1 - km \cdot 2^{-t}$, and $||\mathbf{E}||_\infty \leq \sigma\sqrt{t}$ with probability at least $1 - m\ell \cdot 2^{-t}$. Therefore, the probability that $||\mathbf{ED}||_\infty \leq m\sigma\tau t$ is at least $1 - O(km \cdot 2^{-t} + k\ell \cdot 2^{-t}) = 1 - O(k\ell \cdot 2^{-t})$ (recall $m \leq \ell$). Choose t to be superpolynomial in m, n, ℓ, k such that q is still superpolynomial in t. Moreover, $2||\mathbf{T}||_\infty \leq q - \alpha$, with probability $1 - O(k\ell\alpha/q)$. By choose q to be superpolynomial in k, ℓ and α, we get that this probability is $1 - \mathsf{negl}(\lambda)$.

Therefore, by setting the values of n, m, ℓ, k, t, q appropriately and choosing $\alpha = m\sigma\tau t + 1$, we have that with probability $1 - \mathsf{negl}(\lambda)$, $\mathbf{ED} + 2\mathbf{T} \pmod q$ is equal to $\mathbf{ED} + 2\mathbf{T}$ (without the modular reduction). In particular,

$$\mathbf{W} = (\mathbf{ED} + 2\mathbf{T} \mod q) \mod 2 = \mathbf{ED} + 2\mathbf{T} \mod 2 = \mathbf{ED} \mod 2.$$

Therefore, it is clear that the first row of $\mathbf{W}$ is in the span of $\mathbf{D} \pmod 2$ and the algorithm will output 1 with probability at least $1 - \mathsf{negl}(\lambda)$.

Random case: If $\mathbf{C}$ is uniformly random, we show that with high probability, the algorithm outputs 0.

For a matrix $\mathbf{M}$, let $\mathbf{M}_i$ denote the ith column, and let $\mathbf{M}^{(-i)}$ denote the matrix obtained by deleting the ith column. We will also denote by $b(X)$ the number of bits in X.

Since a sample from $\mathcal{D}_{\mathbb{Z},\tau}$ has min-entropy $\Theta(\log \tau)$ bits, we have that $H_\infty(\mathbf{D}_i) = \Theta(m \log \tau)$. Now, note that

$$\begin{aligned}
&\widetilde{H}_\infty(\mathbf{D}_i | \mathbf{B}, \mathbf{P}, \mathbf{D} \bmod 2) \\
&\leq \widetilde{H}_\infty(\mathbf{D}_i, \mathbf{P}_i, \mathbf{D}_i \bmod 2 | \mathbf{B}, \mathbf{P}^{(-i)}, \mathbf{D}^{(-i)} \bmod 2) - b(\mathbf{P}_i) - b(\mathbf{D}_i \bmod 2) && (5.1)\\
&= \widetilde{H}_\infty(\mathbf{D}_i | \mathbf{B}, \mathbf{P}^{(-i)}, \mathbf{D}^{(-i)} \bmod 2) - n \log q - m && (5.2)\\
&= H_\infty(\mathbf{D}_i) - n \log q - m && (5.3)\\
&= \Theta(m \log \tau) - n \log q - m = \Theta(m \log \tau).
\end{aligned}$$

The first inequality (5.1) follows from the chain rule (Lemma 3.5). The second equality (5.2) follows from the fact that $\mathbf{P}_i$ and $\mathbf{D}_i \bmod 2$ are determined given $\mathbf{D}$ and $\mathbf{B}$. The third equality (5.3) follows from the fact that $\mathbf{D}_i$ is sampled independently from $\mathbf{B}$ and $\mathbf{P}^{(-i)}$.

Therefore, applying the left-over hash lemma (Lemma 3.7) repeatedly to each column of $\mathbf{D}$, we have that

$$(\mathbf{u}^T \mathbf{D} \mod q, \mathbf{B}, \mathbf{P}, \mathbf{D} \bmod 2) \approx_\varepsilon (\mathbf{v}^T, \mathbf{B}, \mathbf{P}, \mathbf{D} \bmod 2)$$

where $\mathbf{u} \leftarrow \mathbb{Z}_q^{m \times 1}$ and $\mathbf{v} \leftarrow \mathbb{Z}_q^{k \times 1}$, for some $\varepsilon \leq \mathsf{negl}(\lambda)$. Note that the auxiliary information $\mathbf{Q} = \mathbf{SP} - 2\mathbf{T}$ can be simulated using only $\mathbf{B}$ and $\mathbf{P}$, so this indistinguishability holds even in the presence of aux. Therefore, in the presence of $\mathbf{B}, \mathbf{P}, \mathbf{Q}$, the first row of $\mathbf{CD}$, and hence the first row of $\mathbf{CD} - \mathbf{Q}$ is uniformly random in $\mathbb{Z}_q^{1 \times k}$. Since q is superpolynomial sized, the first row of

$$\mathbf{W} = \mathbf{CD} - \mathbf{Q} \bmod 2$$

is statistically close to uniformly random. Recall that for a uniformly random vector in $\mathbb{Z}_2^k$, the probability that it is in the row span of $\mathbf{D} \pmod 2$ is at most $1/2^{k-m} \leq \mathsf{negl}(\lambda)$ since $\mathbf{D} \pmod 2$ has rank at most m.

Therefore, the algorithm will output 0 with probability $1 - \mathsf{negl}(\lambda)$. □

5.2 Analysis of the Pre-condition

In this section, we prove Lemmas 5.2 and 5.3. We will make use of the following three simple facts:

1. **Gaussian-flooding:** If $\mathbf{E}$ follows $D_{\mathbb{Z},\sigma}^{n \times m}$ then for any $\tilde{\mathbf{E}}$ with $\|\tilde{\mathbf{E}}\|_\infty \leq t$ the statistical distance of $\mathbf{E}$ and $\mathbf{E} + \tilde{\mathbf{E}}$ is at most mnt/σ, which is negligible if σ/t is superpolynomial (Lemma 3.2).

2. **Rectangular Flooding:** If x is uniformly random in $[0, q)$, then the statistical distance of x and $x + t$ is at most $2t/q$.

3. **Bit-Decompositions:** If q is superpolynomial, then a uniformly random integer x sampled from $[0, q)$ is statistically close to $2x' + x''$, where x' is uniformly random in $[0, q/2)$ and x'' is a uniformly random bit.

Proof of Lemma 5.2. We proceed in a few hybrids.

H_1: This is the real distribution $(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathsf{aux} = \mathbf{SP} - 2\mathbf{T})$.

H_2: In this hybrid, we make an LWE-error $2\widetilde{\mathbf{E}}$, where $\widetilde{\mathbf{E}}$ follows $D_{\mathbb{Z},\sigma}^{\ell \times k}$, appear in the third component, i.e. the distribution is

$$(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + 2\widetilde{\mathbf{E}} + \mathbf{E}', \mathbf{SP} - 2\mathbf{T}).$$

Hybrids H_1 and H_2 are statistically close via Gaussian-flooding as σ'/σ is superpolynomial.

H_3: In this hybrid, we make the same LWE-error $\widetilde{\mathbf{E}}$ appear in the auxiliary information, i.e. the distribution is

$$(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + 2\widetilde{\mathbf{E}} + \mathbf{E}', \mathbf{SP} + 2\widetilde{\mathbf{E}} - 2\mathbf{T}).$$

Hybrids H_2 and H_3 are statistically close as $2\mathbf{T}$ and $2\mathbf{T} - 2\widetilde{\mathbf{E}} = 2(\mathbf{T} - \widetilde{\mathbf{E}})$ are statistically close via rectangular flooding as $\|\widetilde{\mathbf{E}}\|$ is $\mathsf{poly}(\lambda) \cdot \sigma'$ bounded and $(q/2)/\sigma'$ is superpolynomial.

H_4: In this hybrid the distribution is

$$(\mathbf{B}, \mathbf{P}, \mathbf{U}, \mathbf{V} + \mathbf{E}', \mathbf{V} - 2\mathbf{T}).$$

Computational indistinguishability of H_3 and H_4 follows routinely by a reduction to small-secret LWE (Lemma 3.4). For this reduction, note that since q is odd, 2 is a unit mod q. Hence 'scaling' LWE samples by a factor of 2 is an invertible operation and does not distort uniform distributions (i.e. multiplying a uniformly random matrix with 2 preserves uniformity modulo q).

H_5: In this hybrid we choose the fourth component uniformly random (and independent of $\mathbf{V}$, the distribution hence is

$$(\mathbf{B}, \mathbf{P}, \mathbf{U}, \mathbf{U}', \mathbf{V} - 2\mathbf{T}).$$

Hybrids H_4 and H_5 are statistically close, as since q is superpolynomial the bit-decomposition property yields that $\mathbf{V}$ is statistically close to $2\mathbf{W} + \mathbf{W}'$, where $\mathbf{W}$ is component-wise uniform in $[0, q/2)$ and $\mathbf{W}'$ component-wise uniform in $\{0, 1\}$. Hence

$$\begin{aligned}
(\mathbf{V} + \mathbf{E}', \mathbf{V} - 2\mathbf{T}) &\approx_s (2\mathbf{W} + \mathbf{W}' + \mathbf{E}', 2\mathbf{W} + \mathbf{W}' - 2\mathbf{T}) && (5.4)\\
&\approx_s (2\mathbf{W} + \mathbf{W}' + \mathbf{E}', \mathbf{W}' - 2\mathbf{T}) && (5.5)\\
&\approx_s (2\mathbf{W} + \mathbf{W}'' + \mathbf{E}', \mathbf{W}' - 2\mathbf{T}) && (5.6)\\
&\approx_s (\mathbf{U}', \mathbf{W}' - 2\mathbf{T}) && (5.7)\\
&\approx_s (\mathbf{U}', \mathbf{W}' + 2\mathbf{R} - 2\mathbf{T}) && (5.8)\\
&\approx_s (\mathbf{U}', \mathbf{V} - 2\mathbf{T}) && (5.9)
\end{aligned}$$

here $\mathbf{W}''$ is component-wise independently uniform in $\{0, 1\}$ and $\mathbf{R}$ is component-wise uniformly random in $[0, q/2)$. Note that since σ is superpolynomial, $\mathbf{E}'$ drowns both $\mathbf{W}'$ and $\mathbf{W}''$, hence (5.4) and (5.5) are statistically close.

H_6: In this hybrid, we replace $\mathbf{V}$ with $\mathbf{SP} + 2\widetilde{\mathbf{E}}$, i.e. the distribution is

$$(\mathbf{B}, \mathbf{P}, \mathbf{U}, \mathbf{U}', \mathbf{SP} + 2\widetilde{\mathbf{E}} - 2\mathbf{T}).$$

Computational indistinguishability follows once more by the LWE assumption.

H_7: In this last hybrid, we drop the LWE error $2\widetilde{\mathbf{E}}$. This follows from the fact that $2\mathbf{T}$ and $2\mathbf{T} - 2\widetilde{\mathbf{E}}$ are statistically close by the rectangular drowning property as $(q/2)/\sigma'$ is superpolynomial.

Therefore, the pre-condition is satisfied. □

Now, we prove Lemma 5.3 via a similar sequence of hybrids.

Proof of Lemma 5.3. We proceed in a sequence of hybrids.

H_1: This is the real distribution $(\mathbf{P}, \mathsf{aux} = \mathbf{SP} - 2\mathbf{T})$.

H_2: In this hybrid, sample $\mathbf{P}$ as $2\widetilde{\mathbf{P}}$ instead for a uniformly random matrix $\widetilde{\mathbf{P}}$, i.e., the distribution is

$$(2\widetilde{\mathbf{P}}, 2\mathbf{S}\widetilde{\mathbf{P}} - 2\mathbf{T}).$$

Since q is odd, 2 is a unit mod q. Therefore, the distribution is identical to the previous hybrid, and we have only made a syntactic change.

H_3: Sample $\mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{\ell \times k}$. Add noise of the form $2\mathbf{E}$ to aux, writing the distribution as

$$(2\widetilde{\mathbf{P}}, 2\mathbf{S}\widetilde{\mathbf{P}} - 2\mathbf{E} - 2\mathbf{T}).$$

This is statistically close because $\mathbf{T} + \mathbf{E}$ and $\mathbf{T}$ are statistically close via rectangular flooding by choosing σ such that $2\sigma/(q/2) = \mathsf{negl}(\lambda)$.

H_4: Replace $2\mathbf{S}\widetilde{\mathbf{P}} - 2\mathbf{E}$ with a random matrix $\mathbf{V} \leftarrow \mathbb{Z}_q^{\ell \times k}$, therefore changing the distribution to

$$(2\widetilde{\mathbf{P}}, \mathbf{V} - 2\mathbf{T}).$$

This indistinguishability follows from the fact that $(\widetilde{\mathbf{P}}, \mathbf{S}\widetilde{\mathbf{P}} - \mathbf{E}) \approx (\widetilde{\mathbf{P}}, 2^{-1} \cdot \mathbf{V})$ assuming small-secret LWE (Lemma 3.4), and the fact that 2 is a unit mod q since q is prime.

H_5: Sample $\mathbf{R} \leftarrow \mathcal{U}([\kappa])^{n \times k}$, and add this quantity to $2\widetilde{\mathbf{P}}$:

$$(2\widetilde{\mathbf{P}} + \mathbf{R}, \mathbf{V} - 2\mathbf{T}).$$

This is statistically indistinguishable from the previous hybrid since $2\widetilde{\mathbf{P}}$ is uniformly random.

H_6: Sample $\mathbf{R} = 2\mathbf{R}' + \mathbf{R}''$, where $\mathbf{R}'' \leftarrow \{0,1\}^{n \times k}$, where $\mathbf{R}' \leftarrow [\rho]^{n \times k}$, where $\rho = \lfloor \frac{\kappa - 1}{2} \rfloor$. It is clear that this is statistically close to the distribution in the previous hybrid. We can write the resulting distribution as follows:

$$(2\widetilde{\mathbf{P}} + 2\mathbf{R}' + \mathbf{R}'', \mathbf{V} - 2\mathbf{T}).$$

H_7: Sample $\mathbf{V} = 2\mathbf{S}(\widetilde{\mathbf{P}} + \mathbf{R}') + 2\mathbf{E}$, therefore changing the distribution to

$$(2\widetilde{\mathbf{P}} + 2\mathbf{R}' + \mathbf{R}'', 2\mathbf{S}\widetilde{\mathbf{P}} + 2\mathbf{S}\mathbf{R}' + 2\mathbf{E} - 2\mathbf{T}).$$

The indistinguishability of these hybrids follows from a few steps:

$$
\begin{aligned}
&(2\widetilde{\mathbf{P}} + 2\mathbf{R}' + \mathbf{R}'', \mathbf{V} - 2\mathbf{T}) \\
&\approx_s (2\mathbf{W} + \mathbf{R}'', \mathbf{V} - 2\mathbf{T}) && (5.10)\\
&\approx_c (2\mathbf{W} + \mathbf{R}'', 2(\mathbf{SW} + \mathbf{E}) - 2\mathbf{T}) && (5.11)\\
&\approx_s (2(\widetilde{\mathbf{P}} + \mathbf{R}') + \mathbf{R}'', 2\mathbf{S}(\widetilde{\mathbf{P}} + \mathbf{R}') + 2\mathbf{E} - 2\mathbf{T}) && (5.12)
\end{aligned}
$$

where (5.10) and (5.12) follow from the fact that $\widetilde{\mathbf{P}}$ is uniformly random, and (5.11) follows from small-secret LWE (Lemma 3.4).

H_8: Sample the distribution instead as

$$(2\widetilde{\mathbf{P}} + 2\mathbf{R}' + \mathbf{R}'', \mathbf{SP} - 2\mathbf{T}).$$

Recall that $||\mathbf{S}||_\infty \leq \sigma\sqrt{\beta}$ with probability at least $1 - n\ell 2^{-\beta}$ (Lemma 3.1). Therefore, $||\mathbf{SR}'||_\infty \leq n\kappa\sigma\sqrt{\beta}$. Choosing β to be superpolynomial such that $q/(n\kappa\sigma\sqrt{\beta})$ is superpolynomial, it follows that $\mathbf{T}$ is statistically close to $\mathbf{T} - \mathbf{SR}' - \mathbf{E}$ by rectangular flooding.

H_9: Rewriting $2\widetilde{\mathbf{P}} = \mathbf{P}$, and $\mathbf{R} = \mathbf{R}' + \mathbf{R}''$, we get the distribution

$$(\mathbf{P} + \mathbf{R}, \mathbf{SP} - 2\mathbf{T}).$$

This completes the proof. □

6 Conclusions

We show a new, simple, attack against all known version of the private-coin evasive LWE assumption. This achieves the same effect as obfuscation-based counterexamples, albeit in a much more elementary way, using benign-looking distributions of the auxiliary input. Conceptually, our attack challenges the *pseudo-drowning* heuristic, which is at the core of all security proofs that rely on evasive LWE.

That said, none of the known attacks, including ours, break the constructions of witness encryption and null-IO [VWW22, Tsa22], adaptively sound SNARGs for UP [MPV24, JKLM25], or non-adaptive SNARG for NP [JKLM25]. Proving the security of these appealingly simple constructions from plausible assumptions remains an open problem.

Where does one go from here vis-a-vis the evasive LWE assumption? The conservative path would be to stick with the *public coin* evasive LWE assumption which has no known attack, and try to expand its (so far limited) reach to more constructions. A different path is to rethink and refine the rationale for (private-coin) evasive LWE, and come up with a version that enables advanced applications yet avoids all known attacks including ours.

Acknowledgements. We thank Rachel Lin for discussions that initiated this research.

References

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Berlin, Heidelberg, Germany. 8, 9

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th Annual ACM Symposium on Theory of Computing*, pages 99–108, Philadephia, PA, USA, May 22–24, 1996. ACM Press. 1

[AKY24] Shweta Agrawal, Simran Kumari, and Shota Yamada. Compact pseudorandom functional encryption from evasive LWE. *IACR Cryptol. ePrint Arch.*, page 1719, 2024. 3

[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. 3, 5

[BUW24] Chris Brzuska, Akin Unal, and Ivy K. Y. Woo. Evasive lwe assumptions: Definitions, classes, and counterexamples. Springer-Verlag, 2024. 2, 3, 4, 7, 9, 10

[BV22] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-inspired broadcast encryption and succinct ciphertext-policy ABE. In Mark Braverman, editor, *ITCS 2022: 13th Innovations in Theoretical Computer Science Conference*, volume 215, pages 28:1–28:20, Berkeley, CA, USA, January 31 – February 3, 2022. LIPIcs. 1

[CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland. 2

[CHL+15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 3–12. Springer, 2015. 2

[CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 577–607, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland. 1, 3, 8

[DLRW24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: New dimensions in cryptography. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part I*, volume 14651 of *Lecture Notes in Computer Science*, pages 3–32, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland. 2

[DORS03] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. Cryptology ePrint Archive, Report 2003/235, 2003. 9

[HLL23] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In *64th Annual Symposium on Foundations of Computer Science*, pages 415–434, Santa Cruz, CA, USA, November 6–9, 2023. IEEE Computer Society Press. 3

[JKLM25] Zhengzhong Jin, Yael Tauman Kalai, Alex Lombardi, and Surya Mathialagan. Universal SNARGs for NP from proofs of correctness. To appear at STOC 2025, 2025. 3, 5, 16

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23, French Riviera, May 30 – June 3, 2010. Springer, Berlin, Heidelberg, Germany. 1

[LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015. 1

[MPV24] Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge SNARKs for UP. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 38–71, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland. 3, 5, 16

[MR04] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th Annual Symposium on Foundations of Computer Science*, pages 372–381, Rome, Italy, October 17–19, 2004. IEEE Computer Society Press. 1

[MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 629–658, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Berlin, Heidelberg, Germany. 2

[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: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166, New York, NY, USA, March 4–7, 2006. Springer, Berlin, Heidelberg, Germany. 1

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press. 1

[Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*,

volume 13507 of *Lecture Notes in Computer Science*, pages 535–559, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland. 1, 3, 5, 16

[VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part I*, volume 13791 of *Lecture Notes in Computer Science*, pages 195–221, Taipei, Taiwan, December 5–9, 2022. Springer, Cham, Switzerland. 3, 5, 7, 16

[Wee22] Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part II*, volume 13276 of *Lecture Notes in Computer Science*, pages 217–241, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland. 1, 2, 3

[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: 20th Theory of Cryptography Conference, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 651–679, Chicago, IL, USA, November 7–10, 2022. Springer, Cham, Switzerland. 3