A Note on Obfuscation-based Attacks on Private-coin Evasive LWE

Tzu-Hsiang Huang[1], Wei-Hsiang Hung[1], Shota Yamada[2]

[1]Academia Sinica, Taipei, Taiwan
jimmy@iis.sinica.edu.tw, arniehung@iis.sinica.edu.tw
[2]National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, Japan
yamada-shota@aist.go.jp

Abstract

The evasive learning with errors (evasive LWE) assumption is a new assumption recently introduced by Wee [Wee22] and Tsabary [Tsa22] independently, as a significant strengthening of the standard LWE assumption. While the assumption is known to imply various strong primitives including witness encryption [Tsa22, VWW22], the assumption in the most general case (i.e., the private coin variant) is considered quite implausible due to the obfuscation based attack mentioned in [Wee22]. This obfuscation based attack is then later formalized by Vaikuntanathan, Wee, and Wichs [VWW22]. In this note, we revisit their attack and show that the attack actually does not work by showing a concrete counterexample. We then show that their attack can be made valid with some modifications. Along the way, we also improve the counterexample by making it provable. Specifically, our counterexample is valid assuming the (plain) LWE assumption and the existence of instance-hiding witness encryption, whereas their original counterexample was dependent on the heuristic assumption of the existence of an ideal obfuscation.

Contents

1 Introduction

The evasive learning with errors (evasive LWE) assumption is a new assumption recently introduced by Wee [Wee22] and Tsabary [Tsa22] independently, as a significant strengthening of the standard LWE assumption. This assumption enables various constructions of cryptographic primitives, which are not known to be possible by assuming the LWE assumption alone for many years. Notably, this includes the constructions of (optimal) broadcast encryption [Wee22] and witness encryption (WE) [Tsa22]. There are many subsequent works that use evasive LWE to construct various primitives including multi-authority ABE [WWW22], multi-input ABE [ARYY23], key-policy ABE for unbounded depth circuits [HLL23], ciphertext-policy ABE [HLL24], ABE for Turing machines [AKY24a], SNARKs for UP [MPV24], functional encryption and obfuscation for pseudorandom functionalities [AKY24b, AKY24c, BDJ+24], and succinct witness encryption [BDJ+24].

The primitives mentioned above, except for succinct witness encryption, can be based on indistinguishability obfuscation (IO) [GGH+13], which can be constructed from a combination of well-founded assumptions [JLS21, JLS22, RVV24]. Current IO constructions, however, are vulnerable to quantum computers due to reliance on bilinear maps [JLS21, JLS22, RVV24] or depend on non-standard assumptions [BDGM20, GP21, WW21, DQV+21], some of which are broken [HJL21, JLLS23]. The evasive LWE assumption relies solely on lattices, and thus appears to be post-quantum secure. Furthermore, it circumvents the zeroizing attacks [Wee21], which are typically applicable to obfuscation candidates. These characteristics make the evasive LWE assumption a promising starting point for constructing advanced primitives for which we do not know how to construct without IO, in the post-quantum regime.

The assumption involves a sampler, denoted as Samp, which outputs $(\mathbf{S}, \mathbf{P}, \mathrm{aux})$. The assumption states that if $(\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathrm{aux}) \approx_c (\$, \$, \mathrm{aux})$ holds, then $(\mathbf{SB} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}), \mathrm{aux}) \approx_c (\$, \mathbf{B}^{-1}(\mathbf{P}), \mathrm{aux})$ should also hold, where $\mathbf{B}$ and $\mathbf{S}$ are random matrices, $\mathbf{E}$ and $\mathbf{E}'$ are short Gaussian distributions, $\mathbf{P}$ is possibly structured matrix, $\mathbf{B}^{-1}(\mathbf{P})$ is a short matrix satisfying $\mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P}) = \mathbf{P}$ sampled from discrete Gaussian distribution, and $\$$ stands for a random matrix with the corresponding size.[1] Wee [Wee22] discusses that the assumption may not be valid for all samplers. In more detail, he introduces the concept of public coin versus private coin versions of the assumption, where the distinction is made based on whether the sampler's coin is given to the distinguisher or not. He then notes that the private coin version may suffer from an "obfuscation-based counterexample." This motivates the research community to rely on the public coin version of the assumption when it is possible. However, Wee's paper does not provide any technical explanation of this obfuscation-based counterexample. A description of the obfuscation-based counterexample is later given by Vaikuntanthan, Wee, and Wichs [VWW22] (henceforth VWW in the following), who showed a sampler for which pre-condition distributions are indistinguishable, but the post-condition distributions can be distinguished efficiently.

In this note, we show that the pre-condition distributions given by them are actually distinguishable, as opposed to their claim. This invalidates their attack against the private-coin evasive LWE, since evasive LWE only refers to the case where pre-condition distributions are indistinguishable and claims nothing else otherwise. Nevertheless, we are able to fix their counterexample to offer a new valid counterexample. Inspired by the recent work by Branco et al. [BDJ+24], we also improve the VWW counterexample by eliminating the heuristic assumption that ideal obfuscation exists that was used by VWW. Specifically, our counterexample is provable, showing indistinguishability of the pre-condition by assuming the existence of instance-hiding WE. Goyal et al. [GKW17] and Wichs et al. [WZ17] obtain instance-hiding WE by

[1] Multiple variants of this assumption exist. For example, in some variants, $\mathbf{B}$ and $\mathbf{P}$ are provided to the adversary in both pre-condition and post-condition in some versions such as [Wee22, ARYY23].

combining (plain) WE with compute-and-compare obfuscation (a.k.a lockable obfuscation) and also show that compute-and-compare obfuscation can be constructed from LWE. Additionally, VWW and [BDJ+24] construct WE from private-coin evasive LWE. Consequently, assuming the hardness of LWE, our work shows that many variants of private-coin evasive LWE cannot be secure by providing a concrete and provable counterexample.

We note that [BDJ+24] constructs pseudorandom obfuscation (PRO) based on (a slight variant of) private-coin evasive LWE and show counter-example against PRO for general samplers. Combination of these results constitutes counter-example against the variant of private-coin evasive LWE. Our counter-example is with respect to a different sampler than theirs and would be more direct and simpler in that we do not have to go through the construction of PRO, which involves the computation of FHE etc. In addition, we directly refutes the original version of evasive LWE instead of the slight variant (See discussion in Sec 2.4 of [BDJ+24] for details).

Discussion and Open question Observe that the original VWW attack intends to work even against a sampler who chooses $\mathbf{S}$ independently from anything else (i.e., $\mathbf{P}$ and aux). If it works, this is a stronger attack than what we show in Section 4, since the sampler in our attack uses the information of $\mathbf{S}$ to generate aux. However, as we showed in Section 3, their attack does not work as intended. It remains open to remove the dependence of aux on $\mathbf{S}$ from our attack.

Independent Work Brzsuka, Ünal, and Woo [BÜW24] present several counterexamples to the private coin evasive LWE, identifying plausible classes of this assumption. Their work overlaps somewhat with ours since they provide a formal counterexample using ideas similar to ours.[2] For example, Lemma 4.1 is similar to their Lemma 10, and we use a similar Samp as a counterexample in Section 4, compared to their Section 7. However, all of our proofs differ in some details. We emphasize that these overlapping results are developed independently and note that our work uniquely includes the specific counterexample against the VWW attack on the private coin evasive LWE.

2 Preliminaries

2.1 Notations

We use dollar signs $\$$, $\$'$ and $\$''$ to denote the samples drawn from uniformly random matrices over $\mathbb{Z}_q$. Denote $|\cdot|$ as the maximum absolute value over all entries of a matrix. For a matrix $\mathbf{M}$, we denote $(\mathbf{M})_{i,j}$ as the i-th row, j-th column entry of $\mathbf{M}$. For two distributions χ and χ', we denote the statistical distance between them by $\Delta(\chi, \chi')$.

2.2 Lattice background

Gaussian Distributions. For an integer $m > 0$, let $\mathcal{D}_{\mathbb{Z}^m,\sigma}$ be the discrete Gaussian distribution over $\mathbb{Z}^m$ with deviation $\sigma > 0$.

Gaussian Preimage. Let $n, m, q \in \mathbb{N}$ and consider a matrix $\mathbf{B} \in \mathbb{Z}_q^{n\times m}$. For all $\mathbf{P} \in \mathbb{Z}_q^{n\times m'}$, we denote $\mathbf{B}^{-1}(\mathbf{P}, \sigma)$ as the Gaussian distribution $\mathcal{D}_{\mathbb{Z}^{m\times m'},\sigma}$ conditioned on $\mathbf{B}\cdot\mathbf{B}^{-1}(\mathbf{P}, \sigma) = \mathbf{P}$. We suppress σ when the context is clear.

[2] In more detail, one of the authors saw the primary version of [BÜW24] in private circulation before this work is done. However, the primary version did not contain the part that overlaps with ours.

Lemma 2.1 ([Reg09], Lemma 2.5). *We have* $\Pr[\|\mathbf{x}\|_\infty > \sigma\sqrt{m} : x \leftarrow \mathcal{D}_{m,\sigma}] \leq 2^{-2m}$

Lemma 2.2 ([MR04], Lemma 4.4). *For* $\mathbf{u} \in \mathbb{Z}_q^n$*, it follows* $\Pr[\|\mathbf{B}^{-1}(\mathbf{u}, \sigma)\|_\infty > \sqrt{m}\sigma] = \mathsf{negl}(n)$

Lemma 2.3 ([KYY18], Lemma 6). *Let* $\sigma > 16\sqrt{\log 2m/\pi}$ *and* $\mathbf{u}$ *be any vector in* $\mathbb{Z}_q^n$*. Then, for all but* q^{-n} *fraction of* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$*, we have*

$$H_\infty(\mathbf{A}^{-1}(\mathbf{u}, \sigma)) \geq m - 1.$$

Lemma 2.4 (Leftover hash lemma, [Reg09]). *For a random variable* $\mathbf{x}$ *over* $\mathbb{Z}_q^m$ *with min-entropy* h*, the statistical distance between* $(\mathbf{a}, \langle \mathbf{a}, \mathbf{x}\rangle)$ *and* $(\mathbf{a}, u)$ *is less than* $\sqrt{q^m/2^h}$ *for* $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_q^m$ *and* $u \xleftarrow{\$} \mathbb{Z}_q$*.*

Lemma 2.5 (Lattice Trapdoor, [MP11]). *There is an efficient algorithm* $\mathsf{GenTrap}(1^n, 1^m, q)$ *outputs* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ *and a trapdoor* $\mathbf{T}$ *such that the distribution of* $\mathbf{A}$ *is statistically* $\mathsf{negl}(n)$*-close to uniform. Moreover, there is an efficient algorithm* Invert *that with overwhelming probability over all random choices, do the following:*

- *Let* $\mathbf{b}^\top = \mathbf{s}^\top\mathbf{A} + \mathbf{e}^\top$ *,where* $\mathbf{s} \in \mathbb{Z}_q^n$ *is arbitrary, and* $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\alpha q}$ *for* $1/\alpha \geq \sqrt{n \log q} \cdot r$ *with some* $r = \omega(\sqrt{\log n})$*. The deterministic algorithm* $\mathsf{Invert}(\mathbf{T}, \mathbf{A}, \mathbf{b})$ *outputs* $\mathbf{s}$ *and* $\mathbf{e}$*.*

Definition 2.6 (Learning with Errors, [Reg09]). *Let* $n, m, q, \sigma \in \mathbb{N}$*. The* $\mathsf{LWE}_{n,m,q,\sigma}$ *assumption states that*

$$(\mathbf{A}, \mathbf{s}^\top\mathbf{A} + \mathbf{e}^\top) \approx_c (\mathbf{A}, \mathbf{c})$$

where

$$\mathbf{A} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma}, \mathbf{c} \leftarrow \mathbb{Z}_q^m$$

We rely on the LWE assumption with subexponential modulus-to-noise ratio, i.e., $q/\sigma \leq 2^{n^\delta}$*, for some* $\delta > 0$*.*

Definition 2.7 (Private-coin Evasive LWE [VWW22]). *Let* Samp *be a PPT algorithm that on input* 1^λ*, outputs*

$$\mathbf{S} \in \mathbb{Z}_q^{n'\times n}, \mathbf{P} \in \mathbb{Z}_q^{n\times t}, aux \in \{0,1\}^*.$$

We define the following advantage functions:

$$\mathsf{Adv}_{\mathcal{A}_0}^{Pre} := \Pr[\mathcal{A}_0(\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', aux) = 1] - \Pr[\mathcal{A}_0(\mathbf{C}, \mathbf{C}', aux) = 1]$$
$$\mathsf{Adv}_{\mathcal{A}_1}^{Post} := \Pr[\mathcal{A}_1(\mathbf{SB} + \mathbf{E}, \mathbf{D}, aux) = 1] - \Pr[\mathcal{A}_1(\mathbf{C}, \mathbf{D}, aux) = 1],$$

where

$$(\mathbf{S}, \mathbf{P}, aux) \leftarrow \mathsf{Samp}(1^\lambda)$$
$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}^{n'\times m},\sigma}, \mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}^{n'\times t},\sigma'}$$
$$\mathbf{C} \leftarrow \mathbb{Z}_q^{n'\times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{n'\times t}, \mathbf{D} \leftarrow \mathbf{B}^{-1}(\mathbf{P}, \sigma).$$

We say that the evasive LWE assumption holds if for every PPT Samp *there exists some polynomial* $poly(\cdot)$ *such that for every PPT* $\mathcal{A}_1$*, there exists another PPT* $\mathcal{A}_0$ *such that*

$$\mathsf{Adv}_{\mathcal{A}_0}^{Pre} \geq \mathsf{Adv}_{\mathcal{A}_1}^{Post}/\mathsf{poly}(\lambda) - \mathsf{negl}(\lambda).$$

Lemma 2.8 (Noise Flooding). *Let* $\chi \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}$ *and* $t \in \mathbb{Z}$*. We have that*

$$\Delta(\chi, \chi + t) \leq \sqrt{\frac{\pi}{2}} \cdot \frac{\|t\|}{\sigma}.$$

2.3 (Instance-Hiding) Witness Encryption

Definition 2.9 (Witness Encryption[GGSW13]). *A witness encryption scheme for an NP language L (with corresponding witness relation R) consists of the following two PPT algorithms:*

- $\mathsf{Enc}(1^\lambda, x, m)$*: takes as input the security parameter* 1^λ*, an instance* $x \in L$ *and a message* $m \in \{0,1\}^*$*, and outputs a ciphertext* ct*.*
- $\mathsf{Dec}(ct, w)$*: takes as input a ciphertext* ct *and a witness* w*, and outputs a message* m *or the symbol* $\perp$*.*

These algorithms satisfy the following properties:

- ***Correctness.*** *For any* $\lambda, m \in \{0,1\}^{\mathsf{poly}(\lambda)}$ *and* $x \in L$ *s.t.* $R(x,w)$ *holds,*

$$\Pr\left[\mathsf{Dec}(\mathsf{Enc}(1^\lambda, x, m), w)\right] = 1 - \mathsf{negl}(\lambda).$$

- ***Soundness.*** *For any PPT* $\mathcal{D}$ *there exists a negligible function* $\mathsf{negl}(\cdot)$ *s.t. for any* $x \notin L, m_0$ *and* $m_1 \in \{0,1\}^{\mathsf{poly}(\lambda)}$*,*

$$\left|\Pr[\mathcal{D}(\mathsf{Enc}(1^\lambda, x, m_0)) = 1] - \Pr[\mathcal{D}(\mathsf{Enc}(1^\lambda, x, m_1)) = 1]\right| = \mathsf{negl}(\lambda).$$

We call it instance-hiding witness encryption if it additionally satisfies the following:

- ***Instance-hiding.*** *For any PPT* $\mathcal{D}$ *there exists a negligible function* $\mathsf{negl}(\cdot)$ *s.t. for any* $x_0, x_1 \notin L$ *and* $m \in \{0,1\}^{\mathsf{poly}(\lambda)}$*,*

$$\left|\Pr[\mathcal{D}(\mathsf{Enc}(1^\lambda, x_0, m)) = 1] - \Pr[\mathcal{D}(\mathsf{Enc}(1^\lambda, x_1, m)) = 1]\right| = \mathsf{negl}(\lambda).$$

As is shown in [GKW17, WZ17], witness encryption without instance-hiding property can be converted into instance-hiding witness encryption using the compute-and-compare obfuscation (a.k.a lockable obfuscation), which can be constructed from the LWE assumption.

3 Counterexample for VWW Obfuscation-based attack

We show that the VWW obfuscation-based attack on private-coin evasive LWE (Definition 2.7) fails.[3] Specifically, we invalidate the precondition by constructing a distinguisher that separates the two distributions induced by the attack sampler in the precondition.

We present a hard-to-decompose lemma that will be used later in this section.

Lemma 3.1 (Hard-to-Decompose Lemma). *Let* $k, m, n \in \mathbb{N}$ *and* $m > n$*, for any* $\mathbf{P} \in \mathbb{Z}_q^{n \times m}$

$$\Pr\left[\exists\ \mathbf{S_0} \in \mathbb{Z}_q^{k \times n} \text{ s.t. } |\mathbf{S_0 P} - \mathbf{U}| \leq \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{k \times m}\right] \leq \frac{q^{k \times n} \times (2\epsilon + 1)^{k \times m}}{q^{k \times m}}.$$

[3] Although we use Definition 2.7, obfuscation-based attacks apply to general private-coin evasive LWE variants, since the attacks only use $\mathbf{SB} + \mathbf{E}$ and $\mathbf{B}^{-1}(\mathbf{P})$ to invalidate the post-condition.

Proof.

$$\begin{aligned}
&\Pr\left[\exists \mathbf{S_0} \in \mathbb{Z}_q^{k\times n} \text{ s.t. } |\mathbf{S_0P} - \mathbf{U}| \le \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{k\times m}\right] \\
=& \Pr\left[\exists \mathbf{M} \in Img(\mathbf{P}) s.t. |\mathbf{M} - \mathbf{U}| \le \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{k\times m}\right] \\
\le& \frac{|Img(\mathbf{P})| \times \left|\{A \in \mathbb{Z}_q^{k\times m} : |A| \le \epsilon\}\right|}{|\mathbb{Z}_q^{k\times m}|} \\
\le& \frac{q^{k\times n} \times (2\epsilon+1)^{k\times m}}{q^{k\times m}},
\end{aligned}$$

where we denote the image of $\mathbf{P}$ as $Img(\mathbf{P}) = \{\mathbf{S_0P} : \forall \mathbf{S_0} \in \mathbb{Z}_q^{k\times n}\}$. The last inequality holds since the size of the image is less than or equal to the size of the domain.

□

Recall that in VWW obfuscation-based attack, Samp on input 1^λ outputs $\mathbf{S} \in \mathbb{Z}_q^{2m\times n}, \mathbf{P} \in \mathbb{Z}_q^{n\times 2m}$ and $aux \in \{0,1\}^*$, where

$$\mathbf{S} \leftarrow \mathbb{Z}_q^{2m\times n}, (\mathbf{P}, \mathbf{T}) \leftarrow \mathsf{GenTrap}(1^n, 1^{2m}, q), aux \leftarrow \mathsf{Obf}(\Pi_{\mathbf{P},\mathbf{T}})$$

such that Obf is an ideal obfuscation, $\mathbf{T}$ is a trapdoor of $\mathbf{P}$ and $\Pi_{\mathbf{P},\mathbf{T}}$ is the program shown as follows:

Algorithm 1 $\Pi_{\mathbf{P},\mathbf{T}}$

Input: $\mathbf{C} \in \mathbb{Z}_q^{2m\times m}, \mathbf{D} \in \mathbb{Z}_q^{m\times 2m}$
  Use $\mathbf{T}$ to solve $\mathbf{S_0}$ s.t. $|\mathbf{C}\cdot\mathbf{D} - \mathbf{S_0}\cdot\mathbf{P}| < \epsilon$
  if $|\mathbf{D}| < \epsilon'$ and $\mathbf{S_0}$ exists **then**
    return 1
  else
    return 0
  end if

We show that the precondition of the evasive LWE instantiated by Samp does not hold.

Theorem 3.2. *There is a PPT distinguisher $\mathcal{D}$ s.t.* $\mathsf{Adv}_{\mathcal{D}}^{Pre} \ge \mathsf{non\text{-}negl}(\lambda)$.

Proof. We set the parameters for evasive LWE as follows:

$$\begin{aligned}
&n, m = \mathsf{poly}(\lambda) \\
&n' = t = 2m && \text{(VWW attack parameters)} \\
&m = O(n\log q), \sigma' = \alpha q, 1/\alpha \ge \sqrt{n\log q}\cdot\omega(\sqrt{\log n}) && \text{(trapdoor inversion)} \\
&\epsilon = \lambda\sigma^2 m, \epsilon' = \sqrt{\lambda}\sigma' && \text{(Gaussian tail bounds)} \\
&q = 2(2\epsilon+1)^2, && \text{(hard-to-decompose lemma)}
\end{aligned}$$

where σ denotes any parameter that satisfies the LWE assumption (Definition 2.6). Additionally, we assume without loss of generality that $\log q \in \mathbb{N}$ to simplify the gadget matrix decomposition. In Algorithm 2, we present a distinguisher $\mathcal{D}$ against the precondition distributions. We then argue that $\mathcal{D}$ has high advantage for distinguishing the precondition distributions.

Algorithm 2 Precondition Distinguisher $\mathcal{D}$

Input: $\mathbf{M_1} \in \mathbb{Z}_q^{2m\times m}, \mathbf{M_2} \in \mathbb{Z}_q^{2m\times 2m}, aux \in \{0,1\}^*$

Let $\mathbf{R} \in \mathbb{Z}_q^{2m\times 2m}$ be the matrix that matches $\mathbf{M_2}$ in the first $(m/\log q)$ rows and consists of zeros in the remaining entries.

Decompose $\mathbf{R}$ into matrices $\mathbf{C} \in \mathbb{Z}_q^{2m\times m}$ and $\mathbf{D} \in \mathbb{Z}_q^{m\times 2m}$ s.t.

$$\mathbf{C} = \begin{bmatrix} \mathbf{I}_{(m/\log q)} \otimes \mathbf{g}^\intercal \\ 0 \end{bmatrix}, \mathbf{g}^\intercal = \left[1, 2^1, ..., 2^{\log q - 1}\right],$$

and $\mathbf{D}$ is a zero-one matrix where $\mathbf{CD} = \mathbf{R}$. Note that such $\mathbf{D}$ can be constructed by replacing each $\mathbb{Z}_q$ element in $\mathbf{R}$ with a binary vector representing the $\mathbb{Z}_q$ element. Each binary vector is ordered such that its least significant bit is at the top, leading to its most significant bit at the bottom.

return $aux(\mathbf{C}, \mathbf{D})$.

If $\mathbf{M_2} = \mathbf{SP} + \mathbf{E}'$, then the non-zero entries of $\mathbf{CD}$ consist of the first $(m/\log q)$ rows of $\mathbf{SP} + \mathbf{E}'$. By the property of lattice trapdoors (Lemma 2.5), the corresponding $\mathbf{S_0}$, which contains the first $(m/\log q)$ rows of $\mathbf{S}$, can be inverted using the trapdoor $\mathbf{T}$ and $\mathbf{P}$ with overwhelming probability. Moreover, $\mathbf{D}$ is a matrix of zeros and ones, we have

$$\Pr\left[|\mathbf{D}| < \epsilon'\right] = \Pr\left[|\mathbf{D}| < \sqrt{\lambda}\sigma'\right] = 1.$$

Furthermore, by Lemma 2.1 and the union bound,

$$\Pr\left[|\mathbf{C}\cdot\mathbf{D} - \mathbf{S_0}\cdot\mathbf{P} < \epsilon\right] \geq \Pr\left[|\mathbf{E}'| < \lambda\sigma^2 m\right] \geq 1 - \mathsf{negl}(\lambda).$$

Thus, $\mathcal{D}$ outputs 1 with overwhelming probability.

However, if $\mathbf{M_2}$ is uniformly random, then the first $(m/\log q)$ rows in $\mathbf{R} = \mathbf{CD}$ are uniformly random. We have

$$\begin{aligned}
&\Pr\left[\exists \mathbf{S_0} \text{ s.t. } |\mathbf{C}\cdot\mathbf{D} - \mathbf{S_0}\cdot\mathbf{P}| < \epsilon\right] \\
&\leq \Pr\left[\exists \mathbf{S}' \in \mathbb{Z}_q^{(m/\log q)\times m} \text{ s.t. } |\mathbf{U} - \mathbf{S}'\mathbf{P}| \leq \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{(m/\log q)\times 2m}\right] \\
&\leq \frac{q^{(m/\log q)\times m} \times (2\epsilon+1)^{(m/\log q)\times 2m}}{q^{(m/\log q)\times 2m}} \\
&= 2^{-(m/\log q)\times m} \\
&= \mathsf{negl}(\lambda).
\end{aligned}$$

The first inequality holds because the norm of a matrix is greater or equal to the norm of its sub-matrix. The second inequality holds by Lemma 3.1.

Therefore, $\mathcal{D}$ outputs 1 with negligible probability. Thus, $\mathcal{D}$ distinguishes the two cases with overwhelming probability.

□

4 Witness Encryption Attack

In this section, we provide an attack against a private-coin variant of evasive LWE based on instance-hiding witness encryption (Definition 2.9) and the LWE assumption with subexponential modulus-to-noise ratio

(Definition 2.6). Our result demonstrates that the assumption does not hold for all samplers in general. Furthermore, our attack is provable in the sense that its success follows from the LWE assumption and the existence of instance-hiding witness encryption. This improves the VWW attack, which relies on the heuristic assumption that the ideal obfuscation exists.

We start by presenting a hard-to-decompose lemma that will be used later in the section.

Lemma 4.1 (Hard-to-decompose Lemma). *Let $m, n \in \mathbb{N}$ and $m > n$.*

$$\Pr\left[\exists(\mathbf{C},\mathbf{D}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^{n\times m} \; s.t. \; |\mathbf{CD}-\mathbf{U}| < \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{m\times m}\right] \leq \frac{(q^{m\times n})^2 \times (2\epsilon+1)^{m\times m}}{q^{m\times m}}.$$

Proof.

$$\begin{aligned}
&\Pr\left[\exists\,(\mathbf{C},\mathbf{D}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^{n\times m} \; s.t. \; |\mathbf{CD}-\mathbf{U}| < \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{m\times m}\right] \\
&= \Pr\left[\exists \mathbf{M} \in \mathcal{S} \; s.t. \; |\mathbf{M}-\mathbf{U}| \leq \epsilon : \mathbf{U} \leftarrow \mathbb{Z}_q^{k\times m}\right], \text{ where } \mathcal{S} = \{\mathbf{CD} \; : (\mathbf{C},\mathbf{D}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^{n\times m}\} \\
&\leq \frac{|\mathcal{S}| \times \left|\{A \in \mathbb{Z}_q^{m\times m} : |A| \leq \epsilon\}\right|}{|\mathbb{Z}_q^{m\times m}|} \\
&\leq \frac{(q^{m\times n})^2 \times (2\epsilon+1)^{m\times m}}{q^{m\times m}}.
\end{aligned}$$

□

Theorem 4.2. *Assuming LWE with subexponential modulus-to-noise ratio (Definition 2.6) and instance-hiding witness encryption (Definition 2.9), there is a sampler with respect to which the private-coin evasive LWE does not hold.*

Proof. We set the parameters as follows:

$$\begin{aligned}
n' = t = 5m, q = (2\epsilon+1)^2 &\qquad \text{(hard-to-decompose lemma)} \\
m = O(n \log q) &\qquad \text{(lattice trapdoor)} \\
\sigma = \sigma'' = \lambda^{\omega(1)}\sigma' &\qquad \text{(noise flooding)} \\
\epsilon = \lambda\sigma^2 m + \sqrt{\lambda}\sigma'', &\qquad \text{(Guassian tail bounds)}
\end{aligned}$$

where σ denotes any parameter that satisfies the LWE assumption (Definition 2.6). We specify a sampler Samp that invalidates the private-coin evasive LWE.

Algorithm 3 Samp

Input: 1^λ
- $\mathbf{S} \leftarrow \mathbb{Z}_q^{5m\times n}$
- $\mathbf{P} \leftarrow \mathbb{Z}_q^{n\times 5m}$
- $\mathbf{E}'' \leftarrow \mathcal{D}_{\mathbb{Z}^{5m\times 5m},\sigma''}$
- $r \leftarrow \{0,1\}^\lambda$
- $aux \leftarrow \mathsf{WE.Enc}(1^\lambda, \mathbf{SP}+\mathbf{E}'', r)||r$
- **return** $\mathbf{S}, \mathbf{P}, aux$

The corresponding NP relation $\mathcal{R}$ for the witness encryption scheme is defined as follows: Let $\mathbf{X} \in \mathbb{Z}_q^{5m\times 5m}$ and $w \in \mathbb{Z}_q^{5m\times m} \times \mathbb{Z}_q^{m\times 5m}$,

$$\mathcal{R}\left(\mathbf{X}, w = (\mathbf{C},\mathbf{D})\right) = 1 \iff |\mathbf{C}\cdot\mathbf{D}-\mathbf{X}| < \epsilon.$$

Precondition holds. We show that the precondition holds under the LWE assumption and the instance-hiding property of the witness encryption scheme.

$$\begin{aligned}\mathbf{SB}+\mathbf{E}, \mathbf{SP}+\mathbf{E}', \mathsf{WE.Enc}(1^\lambda, \mathbf{SP}+\mathbf{E}'', r)||r &\approx_s \mathbf{SB}+\mathbf{E}, \mathbf{SP}+\mathbf{E}', \mathsf{WE.Enc}(1^\lambda, \mathbf{SP}+\mathbf{E}'+\mathbf{E}'', r)||r \\ &\approx_c \$, \$', \mathsf{WE.Enc}(1^\lambda, \$'+\mathbf{E}'', r)||r \\ &\approx_c \$, \$', \mathsf{WE.Enc}(1^\lambda, \$'', r)||r \\ &\approx_c \$, \$', \mathsf{WE.Enc}(1^\lambda, \mathbf{SP}+\mathbf{E}'', r)||r.\end{aligned}$$

The first equation holds because by Lemma 2.1, $|\mathbf{E}'| < \sqrt{\lambda}\sigma'$ with overwhelming probability. Conditioned on this event happens, by Lemma 2.8, the statistical distance between arbitrary entry (i,j) in $\mathbf{E}''$ and $\mathbf{E}'+\mathbf{E}''$ is bounded by

$$\Delta\left((\mathbf{E}'')_{i,j}, (\mathbf{E}'+\mathbf{E}'')_{i,j}\right) \leq \sqrt{\frac{\pi}{2}} \cdot \frac{\sigma'\sqrt{\lambda}}{\sigma''} = \sqrt{\frac{\pi}{2}} \cdot \frac{\sigma'\sqrt{\lambda}}{\lambda^{\omega(1)}\sigma'} = \mathsf{negl}(\lambda).$$

Therefore, by the union bound, the overall statistical distance between $\mathbf{E}''$ and $\mathbf{E}'+\mathbf{E}''$ is also negligible.

The second equation holds by invoking the LWE assumption (Definition 2.6). The third equation holds because by Lemma 4.1, the probability that there is a witness $w = (\mathbf{C}, \mathbf{D})$ that decomposes a uniformly random instance $\$'+\mathbf{E}''$ is given by

$$\Pr\left[\exists\, \mathbf{C}, \mathbf{D} \text{ s.t. } \left|\mathbf{CD} - (\$'+\mathbf{E}'')\right| < \epsilon\right] \leq \frac{(q^{5m\times m})^2 \times (2\epsilon+1)^{5m\times 5m}}{q^{5m\times 5m}} = (2\epsilon+1)^{-(5m\times m)} = \mathsf{negl}(\lambda).$$

Therefore, by the instance-hiding property of the witness encryption scheme (Definition 2.9), we can switch the instance $\$'+\mathbf{E}''$ to another hard-to-decompose random matrix $\$''$ except with negligible probability. The last equation holds again by the LWE assumption .

Postcondition is false. We proceed to show a distinguisher in Algorithm 4 that breaks the postcondition with non-negligible advantage.

Algorithm 4 Postcondition Distinguisher $\mathcal{D}$

Input: $\mathbf{M_1} \in \mathbb{Z}_q^{5m\times m}$, $\mathbf{M_2} \in \mathbb{Z}_q^{m\times 5m}$, $aux \in \{0,1\}^*$
  parse aux as $ct||r$
  if $\mathsf{WE.Dec}(ct, (\mathbf{M_1}, \mathbf{M_2})) = r$ **then**
    return 1
  else
    return 0
  end if

If $\mathbf{M_1} = \mathbf{SB} + \mathbf{E}$, then the probability that the distiguisher $\mathcal{D}$ outputs 1 is given by

$$\begin{aligned}
&\Pr\left[\mathcal{D}((\mathbf{SB}+\mathbf{E}), \mathbf{B}^{-1}(\mathbf{P}), aux) = 1\right] \\
&\geq \Pr\left[\left|(\mathbf{SB}+\mathbf{E})\cdot\mathbf{B}^{-1}(\mathbf{P}) - (\mathbf{SP}+\mathbf{E}'')\right| < \epsilon\right] - \mathsf{negl}(\lambda) \\
&= \Pr\left[\left|\mathbf{E}\cdot\mathbf{B}^{-1}(\mathbf{P}) - \mathbf{E}''\right| < \epsilon\right] - \mathsf{negl}(\lambda) \\
&\geq \Pr\left[\left|\mathbf{E}\cdot\mathbf{B}^{-1}(\mathbf{P})\right| + |\mathbf{E}''| < \epsilon\right] - \mathsf{negl}(\lambda) \\
&\geq \Pr\left[\left|\mathbf{E}\cdot\mathbf{B}^{-1}(\mathbf{P})\right| < \lambda\sigma^2 m \wedge |\mathbf{E}''| < \sqrt{\lambda}\sigma''\right] - \mathsf{negl}(\lambda) \\
&\geq 1 - \mathsf{negl}(\lambda),
\end{aligned}$$

where the first inequality holds by the correctness of witness encryption scheme (Definition 2.9), the third inequality follows by the triangle inequality, and the last inequality holds by the union bound and Lemma 2.2.

Whereas if $\mathbf{M_1}$ is uniformly random, the probability that the distiguisher D outputs 1 is given by

$$\begin{aligned}
&\Pr\left[\mathcal{D}(\$, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{WE.Enc}(1^\lambda, \mathbf{SP}+\mathbf{E}'', r)||r) = 1\right] \\
&\leq \Pr\left[\mathcal{D}(\$, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{WE.Enc}(1^\lambda, \$', r)||r) = 1\right] + \mathsf{negl}(\lambda) \\
&\leq \Pr\left[\mathcal{D}(\$, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{WE.Enc}(1^\lambda, \$', \mathbf{0})||r) = 1\right] + \mathsf{negl}(\lambda) \\
&= \Pr\left[\mathsf{WE.Dec}(\mathsf{WE.Enc}(1^\lambda, \$', \mathbf{0})) = r\right] + \mathsf{negl}(\lambda) \\
&= 2^{-\lambda} + \mathsf{negl}(\lambda) \\
&= \mathsf{negl}(\lambda),
\end{aligned}$$

where the first inequality follows by the LWE assumption. The second inequality holds because by the hard-to-decompose lemma (Lemma 4.1), $\$'$ is a no instance except with negligible probability, and by the soundness of WE (Definition 2.9), switching the underlying plaintext from r to a fixed string $\mathbf{0}$ is indistinguishable. The fourth equation holds since r is independent of $\mathsf{WE.Dec}(\mathsf{WE.Enc}(1^\lambda, \$', \mathbf{0}))$.

Thus, the distinguisher $\mathcal{D}$ separates the postcondition with non-negligible advantage

$$\mathsf{Adv}_{\mathcal{D}}^{Post} \geq 1 - \mathsf{negl}(\lambda),$$

invalidating the evasive LWE assumptions.

□

Acknowledgements Tzu-Hsiang Huang and Wei-Hsiang Hung are supported by NSTC QC project, under Grant no. NSTC 113-2119-M-001-009 and the 2025 Academia Sinica Investigator Award (AS-IA-110-M02). Part of this work was completed during their visit to Shota Yamada at the National Institute of Advanced Industrial Science and Technology (AIST). Shota Yamada is supported by JST CREST Grant Number JPMJCR22M1 and JST AIP Acceleration Research JPMJCR22U5. We thank Chris Brzuska, Ivy K. Y. Woo and Akin Ünal for their helpful comments on the presentation of our paper. We also thank Vinod Vaikuntanathan for his helpful comments.

References

[AKY24a] Shweta Agrawal, Simran Kumari, and Shota Yamada. Attribute based encryption for turing machines from lattices. LNCS, pages 352–386, 2024. (Cited on page 3.)

[AKY24b] Shweta Agrawal, Simran Kumari, and Shota Yamada. Compact pseudorandom functional encryption from evasive LWE. Cryptology ePrint Archive, Paper 2024/1719, 2024. (Cited on page 3.)

[AKY24c] Shweta Agrawal, Simran Kumari, and Shota Yamada. Pseudorandom multi-input functional encryption and applications. Cryptology ePrint Archive, Paper 2024/1720, 2024. (Cited on page 3.)

[ARYY23] Shweta Agrawal, Mélissa Rossi, Anshu Yadav, and Shota Yamada. Constant input attribute based (and predicate) encryption from evasive and tensor LWE. LNCS, pages 532–564, 2023. (Cited on page 3.)

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate iO from homomorphic encryption schemes. In Vincent Rijmen and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, LNCS, pages 79–109, May 2020. (Cited on page 3.)

[BDJ+24] Pedro Branco, Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Pseudorandom obfuscation and applications. *IACR Cryptol. ePrint Arch.*, page 1742, 2024. (Cited on page 3, 4.)

[BÜW24] Chris Brzuska, Akin Ünal, and Ivy K. Y. Woo. Evasive LWE assumptions: Definitions, classes, and counterexamples. LNCS, pages 418–449, 2024. (Cited on page 4.)

[DQV+21] Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Succinct LWE sampling, random polynomials, and obfuscation. LNCS, pages 256–287, 2021. (Cited on page 3.)

[GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013. (Cited on page 3.)

[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 467–476, 2013. (Cited on page 6.)

[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. (Cited on page 3, 6.)

[GP21] Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. pages 736–749. ACM Press, 2021. (Cited on page 3.)

[HJL21] Samuel B. Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying iO. LNCS, pages 673–700, 2021. (Cited on page 3.)

[HLL23] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. pages 415–434. IEEE Computer Society Press, 2023. (Cited on page 3.)

[HLL24] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. A general framework for lattice-based ABE using evasive inner-product functional encryption. LNCS, pages 433–464, 2024. (Cited on page 3.)

[JLLS23] Aayush Jain, Huijia Lin, Paul Lou, and Amit Sahai. Polynomial-time cryptanalysis of the subspace flooding assumption for post-quantum $i\mathcal{O}$. LNCS, pages 205–235, 2023. (Cited on page 3.)

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. pages 60–73. ACM Press, 2021. (Cited on page 3.)

[JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb{F}_p$, DLIN, and PRGs in NC^0. LNCS, pages 670–699, 2022. (Cited on page 3.)

[KYY18] Shuichi Katsumata, Shota Yamada, and Takashi Yamakawa. Tighter security proofs for GPV-IBE in the quantum random oracle model. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 253–282, December 2018. (Cited on page 5.)

[MP11] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. Cryptology ePrint Archive, Paper 2011/501, 2011. (Cited on page 5.)

[MPV24] Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge SNARKs for UP. LNCS, pages 38–71, 2024. (Cited on page 3.)

[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. (Cited on page 5.)

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), September 2009. (Cited on page 5.)

[RVV24] Seyoon Ragavan, Neekon Vafa, and Vinod Vaikuntanathan. Indistinguishability obfuscation from bilinear maps and LPN variants. LNCS, pages 3–36, 2024. (Cited on page 3.)

[Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. LNCS, pages 535–559, 2022. (Cited on page 1, 3.)

[VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. LNCS, pages 195–221, 2022. (Cited on page 1, 3, 5.)

[Wee21] Hoeteck Wee. Broadcast encryption with size $N^{1/3}$ and more from k-lin. LNCS, pages 155–178, 2021. (Cited on page 3.)

[Wee22] Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. LNCS, pages 217–241, 2022. (Cited on page 1, 3.)

[WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. LNCS, pages 127–156, 2021. (Cited on page 3.)

[WWW22] Brent Waters, Hoeteck Wee, and David J. Wu. Multi-authority ABE from lattices without random oracles. LNCS, pages 651–679, 2022. (Cited on page 3.)

[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. (Cited on page 3, 6.)