Independent Leakage Does Not Save Evasive LWE

Refuting Folklore Fixes in the Private-Coin Setting

Yao-Ching Hsieh
University of Washington
ychsieh@cs.washington.edu

Abstract

The evasive learning with errors (evasive LWE) assumption has enabled a wide range of advanced cryptographic constructions from lattices in the past few years. The simple intuition that LWE remains secure even when additional "non-zeroizing" short preimages are given has proved highly versatile in applications. Unfortunately, recent counterexamples have refuted its general private-coin formulations, showing that having the short preimages along with arbitrary auxiliary leakage about the secret and target matrices can invalidate this intuition.

To preserve important applications, particularly witness encryption, the community has sought to repair the assumption by restricting the allowed leakage. One folklore proposal requires the secret and target matrices, together with their respective leakage, to be generated independently. This restriction appears to preserve the application to witness encryption while ruling out all previously known counterexamples.

In this work, we refute this proposal. Assuming LWE and indistinguishability obfuscation, we show that both the binding and hiding variants of private-coin evasive LWE are false even with the independent sampler restriction. Our counterexamples exploit only auxiliary leakage about the target matrix. We sample the secret honestly from a standard uniform or Gaussian distribution, independently of the target and its leakage, and we use no auxiliary information about the secret.

Disclosure of AI Use. The initial ideas behind the presented counterexamples arose as a byproduct of a research discussion on a different topic between the author and ChatGPT 6 Astra. The author then relied heavily on the assistance of ChatGPT 6.1 Sol, 6 Astra, and Opus 5.5 to simplify the constructions and prepare the manuscript. All results have been independently verified by the author, and the author takes full responsibility for the content of the manuscript.

Contents

1 Introduction **3**
- 1.1 Our Results . . . 5
- 1.2 Related Work . . . 5

2 Technical Overview **6**
- 2.1 Revisiting the Evasive LWE Formulations . . . 6
- 2.2 General Counterexample Strategy . . . 8
- 2.3 Counterexample for Binding Evasive LWE . . . 9
- 2.4 Counterexample for Hiding Evasive LWE . . . 10

3 Preliminaries **13**
- 3.1 Gaussian Preimage Sampling . . . 14
- 3.2 Gaussian Tails and Flooding . . . 15
- 3.3 Rank of Gaussian Samples . . . 15
- 3.4 Rational Right Inverses . . . 16
- 3.5 Learning with Errors . . . 17
- 3.6 Leftover Hash Lemma . . . 18
- 3.7 Indistinguishability Obfuscation . . . 18

4 Evasive LWE **18**
- 4.1 Evasive LWE Sampler . . . 18
- 4.2 Evasive LWE Formulations . . . 19

5 Counterexample to Private-coin Binding Evasive LWE **21**
- 5.1 Sampler Description . . . 21
- 5.2 Postcondition Adversary . . . 23
- 5.3 Pseudorandomness of the Precondition Distribution . . . 24

6 Counterexample to Private-coin Hiding Evasive LWE **27**
- 6.1 Sampler Description . . . 27
- 6.2 Postcondition Adversary . . . 28
- 6.3 Proof of Preconditions . . . 31

A Additional Binding Counterexample **40**
- A.1 Sampler Description . . . 40
- A.2 Postcondition Adversary . . . 41
- A.3 Pseudorandomness of the Precondition Distribution . . . 42

B Additional Hiding Counterexample **44**
- B.1 Sampler Description . . . 44
- B.2 Helper lemmas . . . 45
- B.3 Postcondition Adversary . . . 46
- B.4 Proof of Preconditions . . . 48

1 Introduction

The evasive learning with errors (evasive LWE) assumption, introduced by Wee [Wee22], states that LWE remains hard even when the adversary receives "benign" short preimages of target matrices. Building on this assumption, Wee resolved the long-standing open problem of constructing CP-ABE and broadcast encryption from lattices. At a high level, in addition to the LWE sample, the adversary is given a short Gaussian preimage $\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P})$ satisfying $\mathbf{BK} = \mathbf{P}$. Multiplying an LWE sample of $\mathbf{B}$ by $\mathbf{K}$ gives a structured LWE sample of $\mathbf{P}$, i.e.,

$$(\mathbf{SB} + \mathbf{E}_B)\mathbf{K} = \mathbf{SP} + \mathbf{E}_B\mathbf{K}.$$

For a given distribution of the target matrix $\mathbf{P}$, the *precondition* requires $(\mathbf{SB} + \mathbf{E}_B, \mathbf{SP} + \mathbf{E}_P)$ to be jointly pseudorandom for fresh, independent errors, i.e., the preimage does not simply zero out the $\mathbf{SB}$ term. Under this condition, the assumption treats $\mathbf{K}$ as benign and asserts the *postcondition* that $\mathbf{SB} + \mathbf{E}_B$ remains pseudorandom even given $\mathbf{K}$. The assumption can also be viewed as an "idealized" model of lattice preimages. In this view, the only useful way to exploit $\mathbf{K}$ is to multiply it by the corresponding LWE sample $\mathbf{SB} + \mathbf{E}_B$, and the structure of the resulting error provides no additional distinguishing advantage over fresh noise.

The idea of allowing benign preimages in constructions soon led to substantial follow-up work on various applications, including witness encryption [Tsa22, VWW22], multi-input ABE [ARYY23], ABE with unbounded-depth policies [HLL23, AKY24, CW25], and pseudorandom FE and obfuscation [BDJ+25, AKY25], to name a few. To support these applications, these works proposed substantially stronger variants of evasive LWE, allowing much greater freedom in the underlying distributions, as well as auxiliary leakage about the LWE secret or the target matrix's trapdoor. These choices are specified by an efficient *sampler*, which jointly generates the secret, target, and auxiliary information from an arbitrary distribution.[1] The auxiliary information is given to the adversary in both the precondition and the postcondition. All these variants are justified by appealing to the same intuition about benign preimages.

Unfortunately, enthusiasm for these new applications faded as a series of counterexamples appeared, exploiting the extra freedom given to the sampler by these stronger evasive LWE formulations. Recent works construct samplers for which the precondition holds but the postcondition fails [BÜW24, DJM+25, HJL25, AMYY25, HHY26]. These counterexamples show that the community was too optimistic about the evasive LWE intuition. Whether this intuition can be trusted under suitable restrictions remains unclear.

Existing counterexamples to evasive LWE roughly fall into the following two categories[2]:

- **Decomposition-based counterexamples** [BÜW24, HHY26]: These counterexamples exploit the short preimage $\mathbf{K}$ given in the postcondition, while finding such a preimage is hard in the precondition. With $\mathbf{K}$, the adversary can obtain a nontrivial approximate decomposition

  $$\mathbf{SP} + \mathbf{E} \approx (\mathbf{SB} + \mathbf{E}_B)\mathbf{K}.$$

  These counterexamples include an auxiliary obfuscation that hardcodes the target LWE sample $\mathbf{SP} + \mathbf{E}$ and checks for such a decomposition, so that an adversary can distinguish an LWE sample from a uniform sample once it receives $\mathbf{K}$.

[1] This version of evasive LWE is often referred to as the private-coin evasive LWE assumption, though the community has never reached a consensus on what precisely counts as private-coin.

[2] Here we exclude the counterexamples by Brzuska, Ünal, and Woo [BÜW24] to evasive LWE with hidden target matrices $\mathbf{P}$, which should be viewed more as refinements leading to valid formulations in this setting.

- **Cancellation-based counterexamples** [DJM+25, HJL25, AMYY25]: These counterexamples exploit the difference in error structure between the postcondition and the precondition, generating leakage that reveals structure only when combined with the postcondition's LWE samples. The general strategy is to allow the computation of a hint $\mathbf{C}^* = \mathbf{SP} + \mathbf{R}$, where $\mathbf{R}$ is (pseudo)random in its high-order bits but structured in its low-order bits. In the precondition, the difference between $\mathbf{C}^*$ and the fresh LWE sample $\mathbf{C}_P = \mathbf{SP} + \mathbf{E}_P$ is close to uniform, as $\mathbf{E}_P$ conceals the low-order structure. However, in the postcondition, the difference between $\mathbf{C}^*$ and the structured LWE sample $\mathbf{C}_B\mathbf{K} = \mathbf{SBK} + \mathbf{E}_B\mathbf{K} = \mathbf{SP} + \mathbf{E}_B\mathbf{K}$ reveals the low-order structure, enabling an attack on the postcondition. In the literature, the hint $\mathbf{C}^*$ is either included directly in the auxiliary information [DJM+25] or computed via the dual-use technique [BTVW17] applied to LWE samples with circularly defined public matrices [HJL25, AMYY25].

In both cases, the construction seems to rely on either generating the target matrix using the secret or generating auxiliary information that depends jointly on the secret and target.[3]

In response to these counterexamples, the community has sought alternative assumptions that support the important applications of private-coin evasive LWE. One promising proposal is to restrict the sampler's freedom to exclude known counterexamples. Indeed, "plain" evasive LWE [Wee22], where the sampler's only freedom is to sample a target matrix using public randomness, has no known counterexamples. This gives hope that intermediate restrictions on the sampler's freedom could preserve security while still allowing enough flexibility for the desired applications.

One folklore proposal for such an intermediate restriction, mentioned in a few research talks [Wee25b, Woo26], is to require an *independent sampler*, in which the LWE secret $\mathbf{S}$ and the target $\mathbf{P}$, together with their respective leakage, are generated independently. Formally, the sampler is split into two *independent* components: Samp_S generating the secret $\mathbf{S}$ along with its auxiliary information aux_S, and Samp_P generating the target $\mathbf{P}$ along with its auxiliary information aux_P. Each component may still use an arbitrary distribution and generate arbitrary auxiliary leakage about its own output. The independence requirement rules out both forms of joint sampling described above and therefore seems to exclude all previously proposed counterexamples. Furthermore, this restriction seems sufficiently versatile, with folklore suggesting that the multi-input ABE construction [ARYY23] and the witness encryption construction [VWW22] could be proved secure under evasive LWE with independent samplers.[4]

In this work, however, we refute this folklore proposal and show that the independent-sampler restriction is not sufficient to save evasive LWE from counterexamples. More precisely, we prove the following:

> Assuming indistinguishability obfuscation and LWE, private-coin evasive LWE with independent samplers is *false*.

In fact, our counterexamples require no auxiliary leakage about the secret. The secret is sampled honestly from a uniform or Gaussian distribution, while the target and its auxiliary information are generated independently of the secret.

[3] A heuristic counterexample proposed in [VWW22] uses ideal obfuscation and auxiliary information that does not depend on the secret. However, this counterexample was later invalidated in [HHY26].

[4] The unbounded-depth KP-ABE application [HLL23] and its follow-ups rely crucially on circular LWE samples and do not fit within the independent-sampler restriction. Indeed, a recent repair uses pairings in addition to lattice assumptions [ABM+26].

1.1 Our Results

In the following, we summarize our counterexamples against evasive LWE with independent samplers. We use the *private-coin binding evasive LWE* and *private-coin hiding evasive LWE* formulations from [BÜW24], which are the mainstream formulations used in recent works on (private-coin) evasive LWE. We refer the reader to Section 2.1 for a brief introduction to these evasive LWE variants, and Section 4 for the formal definitions.

For binding evasive LWE, we give a counterexample under LWE and indistinguishability obfuscation (iO).

Theorem 1 (Informal, see Section 5)**.** *Assuming LWE and iO, there exists an independent sampler* Samp *where the binding evasive LWE assumption is false. Furthermore, the secret sampler* $\mathsf{Samp}_{\mathsf{S}}$ *can be described as* $\mathsf{Samp}_{\mathsf{S}} : (\mathbf{S} \leftarrow \mathbb{Z}_q^n, \mathsf{aux}_{\mathsf{S}} = \bot)$.

For hiding evasive LWE, we give a counterexample under LWE and null indistinguishability obfuscation (null-iO).

Theorem 2 (Informal, see Section 6)**.** *Assuming LWE and null-iO, there exists an independent sampler* Samp *where the hiding evasive LWE assumption is false. Furthermore, the secret sampler* $\mathsf{Samp}_{\mathsf{S}}$ *can be described as* $\mathsf{Samp}_{\mathsf{S}} : (\mathbf{S} \leftarrow \mathcal{D}_\sigma^n, \mathsf{aux}_{\mathsf{S}} = \bot)$*, where* $\mathcal{D}_\sigma$ *is the discrete Gaussian distribution.*

Combined with the folklore belief that the null-iO construction of [VWW22] remains secure under (subexponential-sized) hiding evasive LWE with independent samplers, this suggests that the hiding evasive LWE assumption is false assuming only subexponentially hard LWE.

As a complementary result, we also give counterexamples to both binding and hiding evasive LWE for independent samplers that, instead of relying on the leakage $\mathsf{aux}_{\mathsf{P}}$ about the target matrix, use only the auxiliary information $\mathsf{aux}_{\mathsf{S}}$ generated by the secret sampler. Both counterexamples sample $\mathbf{P}$ without a trapdoor and with no auxiliary leakage ($\mathsf{aux}_{\mathsf{P}} = \bot$), and are proven assuming LWE and null-iO. See Sections A and B for details.

1.2 Related Work

In this section, we review the applications of evasive LWE and related assumptions, and briefly discuss their status in light of recent counterexamples.

Applications of unbroken evasive LWE variants. Several evasive LWE variants with strict restrictions on the samplers have not been refuted by the recent counterexamples.[5] Their applications include CP-ABE for bounded-depth circuits [Wee22, HLL24], multi-authority ABE [WWW22], multi-client ABE [CLW24], and broadcast and trace schemes [AKYY23]. Subsequent work has obtained or improved many of these results under simpler, more concrete assumptions, such as succinct LWE [Wee24, Wee25a] and decomposed LWE [AMR25].

Broken assumptions with unbroken constructions. Most of the stronger applications of evasive LWE rely on variants that allow much greater freedom in the samplers. Recent counterexamples have refuted these general formulations. These applications include witness encryption [Tsa22, VWW22], constant-arity MI-ABE [ARYY23], ABE for unbounded-depth circuits [HLL23] and Turing machines [AKY24, CW25], registered ABE [ZZC+25], SNARGs for UP [MPV24] and NP [JKLM25], and auxiliary-input point obfuscation [CM25].

[5] These variants are often classified as the "public-coin" variants of evasive LWE.

No attacks on these constructions are currently known, and they may still be secure despite the counterexamples to their underlying assumptions. However, most of these results have yet to be recovered from simpler lattice assumptions (one exception is registered ABE from succinct and decomposed LWE [AMR25, CHW25, WW25]). Finding principled, falsifiable assumptions that recover the other results is an interesting open problem.

Pseudorandom Obfuscation. Pseudorandom obfuscation and pseudorandom functional encryption are among the strongest applications of evasive LWE [BDJ+25, AKY25]. These primitives have also motivated proposals for iO in the pseudorandom oracle model [SBP25] and registered pseudorandom FE [PST26]. However, the strong pseudorandomness guarantees of these primitives turn out to be impossible in full generality [BDJ+25, AMYY25]. Their proposed downstream applications rely on heuristic security claims tailored to the particular samplers arising in their security proofs.

Lattice-based indistinguishability obfuscation. A related line of work constructs indistinguishability obfuscation (iO) from new lattice-based assumptions [BDGM20, GP21, WW21, DQV+21, BDGM22, CLW25, HJL25, AMYY26]. Several of these underlying assumptions have since been broken [HJL21, JLLS23]. These assumptions share the basic structure of evasive LWE: they assert the security of LWE samples in the presence of additional hints. They demand more in terms of functionality, since the hints must also support non-evasive computation. At the same time, the components of these assumptions are generally more concrete and tailored to the constructions, so counterexamples to general evasive LWE do not automatically apply. To our knowledge, several recent proposals have not been broken [CLW25, HJL25, AMYY26], though little is understood about the hardness of their underlying assumptions.

2 Technical Overview

We start in Section 2.1 by revisiting evasive LWE, the binding and hiding variants from [BÜW24], and the independent-sampler restriction. Readers familiar with these formulations can skip directly to the general ideas behind our counterexamples in Section 2.2. We then outline the binding and hiding counterexamples in Sections 2.3 and 2.4, respectively.

2.1 Revisiting the Evasive LWE Formulations

Plain evasive LWE. We start with the original (public-coin) evasive LWE assumption introduced by Wee [Wee22]. An efficient sampler Samp outputs the parameters and matrices $(\mathbf{A}, \mathbf{P})$. The secret $\mathbf{S}$, errors, and source matrix $\mathbf{B}$ are generated honestly and independently of the sampler, with a trapdoor for $\mathbf{B}$. The assumption asserts that the following precondition implies the postcondition:

Sampler.

$$(\mathsf{param}, \mathbf{A}, \mathbf{P}) \leftarrow \mathsf{Samp}(1^\lambda; r_{\mathrm{Samp}}), \quad \mathsf{aux} := r_{\mathrm{Samp}}.$$

Precondition.

$$(\mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\mathrm{A}} = \mathbf{SA} + \mathbf{E}_{\mathrm{A}}, \mathbf{C}_{\mathrm{B}} = \mathbf{SB} + \mathbf{E}_{\mathrm{B}}, \mathbf{C}_{\mathrm{P}} = \mathbf{SP} + \mathbf{E}_{\mathrm{P}}, \mathsf{aux})$$
$$\approx (\mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{U}_{\mathrm{A}} \leftarrow \$, \mathbf{U}_{\mathrm{B}} \leftarrow \$, \mathbf{U}_{\mathrm{P}} \leftarrow \$, \mathsf{aux}).$$

Postcondition.

$$(\mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{C}_A = \mathbf{SA} + \mathbf{E}_A, \mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B, \mathbf{K} = \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx (\mathbf{A}, \mathbf{B}, \mathbf{P}, \mathbf{U}_A \leftarrow \$, \mathbf{U}_B \leftarrow \$, \mathbf{K}, \mathsf{aux}).$$

Here aux is exactly the sampler's random coins, and $\mathbf{U} \leftarrow \$$ denotes an independent uniform matrix of the appropriate dimensions.

Private-coin variations. Subsequent works substantially relaxed the sampler restrictions, allowing $(\mathbf{S}, \mathbf{P}, \mathsf{aux})$ to follow an arbitrary efficiently samplable joint distribution without revealing all the sampling coins. The auxiliary information can now include additional LWE samples such as $\mathbf{SA} + \mathbf{E}_A$, so these formulations need not keep $\mathbf{A}$ as a separate component. These are commonly called *private-coin* variants. The formulation in [VWW22] further omits the matrices $\mathbf{B}, \mathbf{P}$ from both the precondition and postcondition views. The hiding of $\mathbf{P}$ is meaningful since neither the auxiliary information nor the sample $\mathbf{SP} + \mathbf{E}_P$ necessarily reveals it.

It turns out that simply hiding the matrices is problematic. The preimage $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{P})$ can leak information about $\mathbf{P}$. For instance, $\mathbf{BK} = \mathbf{P}$ implies that the row space of $\mathbf{K}$ contains the row space of $\mathbf{P}$. Such leakage may not be available in the precondition, where the adversary is given only the noisy sample $\mathbf{SP} + \mathbf{E}_P$. Brzuska, Ünal, and Woo [BÜW24] showed that this difference can indeed be exploited to construct counterexamples.

In light of this, they formulated *binding* and *hiding* evasive LWE to address this issue. These formulations cover most earlier private-coin applications and have been widely adopted since.

Binding evasive LWE. The binding variant makes $\mathbf{B}$ and $\mathbf{P}$ public in both conditions, so the preimage is no longer an additional leakage about $\mathbf{P}$ in the postcondition. The formulation is as follows.

Sampler.

$$(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda).$$

Precondition.

$$(\mathbf{B}, \mathbf{P}, \mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B, \mathbf{C}_P = \mathbf{SP} + \mathbf{E}_P, \mathsf{aux}) \approx (\mathbf{B}, \mathbf{P}, \mathbf{U}_B \leftarrow \$, \mathbf{U}_P \leftarrow \$, \mathsf{aux}).$$

Postcondition.

$$(\mathbf{B}, \mathbf{P}, \mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B, \mathbf{K} = \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx (\mathbf{B}, \mathbf{P}, \mathbf{U}_B \leftarrow \$, \mathbf{K}, \mathsf{aux}).$$

Hiding evasive LWE. The hiding variant omits the matrices from both views. To address the information about $\mathbf{P}$ that the preimage might reveal, it adds a *target-hiding precondition*: adding a bounded random perturbation to $\mathbf{P}$ must be undetectable even given aux, forcing $\mathbf{P}$ to be noisy with small fresh noise. This can be motivated by enhanced lattice trapdoors [GKW18]: under LWE and suitable parameters, when $\mathbf{B}$ is hidden, the preimage $\mathbf{B}^{-1}(\mathbf{P} + \mathbf{\Delta})$ with fresh error $\mathbf{\Delta}$ hides $\mathbf{P}$.

The hiding evasive LWE assumption asserts that if both preconditions hold, the postcondition follows.

Sampler.

$$(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda).$$

Precondition 1.

$$(\mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B, \mathbf{C}_P = \mathbf{SP} + \mathbf{E}_P, \mathsf{aux}) \approx (\mathbf{U}_B \leftarrow \$, \mathbf{U}_P \leftarrow \$, \mathsf{aux}).$$

Precondition 2.

$$(\mathbf{P}, \mathsf{aux}) \approx (\mathbf{P} + \mathbf{\Delta}, \mathsf{aux}), \qquad \mathbf{\Delta} \leftarrow \{-B, \ldots, B\}^*.$$

Postcondition.

$$(\mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B, \mathbf{K} = \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx (\mathbf{U}_B \leftarrow \$, \mathbf{K}, \mathsf{aux}).$$

Hiding evasive LWE is suited to the witness encryption application of [VWW22], where $\mathbf{P}$ encodes intermediate computation values that must remain hidden. The required target matrices in the construction naturally come with fresh noise, allowing the target-hiding precondition to hold, as shown in [BÜW24]. We give the formal binding and hiding definitions in Definitions 5 and 6.

Independent sampler. In either the hiding or binding formulation, the joint sampler can correlate the secret, target, and auxiliary information. Both formulations are broken by known counterexamples, which exploit this joint correlation. The independent-sampler restriction is proposed to rule out these correlations. It requires the secret and its auxiliary information to be generated independently of the target and its auxiliary information, apart from shared public parameters. We therefore formalize an independent sampler as three algorithms:

$$\begin{aligned} \mathsf{Samp}_{\mathsf{param}}(1^\lambda; r_{\mathsf{Samp}}) &\longrightarrow \mathsf{param}, \\ \mathsf{Samp}_S(\mathsf{param}; r_S) &\longrightarrow (\mathbf{S}, \mathsf{aux}_S), \\ \mathsf{Samp}_P(\mathsf{param}; r_P) &\longrightarrow (\mathbf{P}, \mathsf{aux}_P). \end{aligned}$$

The first algorithm generates only public lattice parameters. The two remaining samplers use independent randomness, and

$$\mathsf{aux} = (r_{\mathsf{Samp}}, \mathsf{aux}_S, \mathsf{aux}_P).$$

The full sampler interface is given in Definition 4.

2.2 General Counterexample Strategy

We are now ready to present our counterexamples to evasive LWE with independent samplers.

Our counterexamples combine ideas from the decomposition-based counterexamples [BÜW24, HHY26] and the cancellation-based counterexamples [DJM+25, HJL25, AMYY25]. Recall that in the former, the auxiliary input contains an obfuscation that tests whether a candidate sample $\mathbf{C}$ and a short $\mathbf{K}$ satisfy $\mathbf{CK} \approx \mathbf{C}^*_P$, for a hardcoded target sample $\mathbf{C}^*_P = \mathbf{SP} + \mathbf{E}^*$. Hardcoding this sample introduces precisely the joint dependence on $\mathbf{S}$ and $\mathbf{P}$ that the independent-sampler restriction rules out. One might try to remove the dependence on $\mathbf{S}$ by using a trapdoor of the target matrix $\mathbf{P}$ to test for LWE membership. This, however, obstructs the use of LWE for proving the precondition. With both the trapdoor $\mathsf{td}_\mathbf{P}$ and the sample $\mathbf{C}_P = \mathbf{SP} + \mathbf{E}_P$ in view, LWE can no longer be applied with the secret $\mathbf{S}$.

Our counterexamples take this idea one step further. We design an obfuscation that tests decompositions only for *LWE samples with structured error*. Given a candidate $\mathbf{C}_B$ and $\mathbf{K}$, the auxiliary circuit forms $\mathbf{C}_B\mathbf{K} = \mathbf{SP} + \mathbf{E}_B\mathbf{K}$ in the real postcondition branch and tries to leverage the structure of the error using $\mathbf{K}$ and hints about $\mathbf{P}$. To preserve the precondition, we choose a hint which allows

only partial recovery of the *target error*. Writing $\mathbf{SP} + \mathbf{E}_P = [\mathbf{SP}_0 + \mathbf{E}_{P,0} \mid \mathbf{SP}_1 + \mathbf{E}_{P,1}]$, the hint allows (approximate) recovery of $\mathbf{E}_{P,0}$ but not $\mathbf{E}_{P,1}$. When the error is structured as in the postcondition, $\mathbf{E}_{P,0} = \mathbf{E}_B\mathbf{K}_0$ is highly correlated with $\mathbf{E}_{P,1}$, allowing a partial recovery to be transferred to a full recovery. But when the error is fresh, $\mathbf{E}_{P,0}$ reveals no information about $\mathbf{E}_{P,1}$, keeping the leakage only within the 0-th slot.

Of course, this is not the full story. With a naive design, the recovery of $\mathbf{E}_{P,0}$ also enables the recovery of $\mathbf{S}$ from $\mathbf{SP}_0$. Just like the full trapdoor attempt, this would break the precondition. The trick is to design $\mathbf{P}_0$ such that $\mathbf{SP}_0$ does not reveal the secret $\mathbf{S}$ fully. We will see how this is done in the binding and hiding counterexamples below respectively.

2.3 Counterexample for Binding Evasive LWE

Sampler. The counterexample sampler is as follows.

$$\mathsf{Samp}_S: \quad \mathbf{S} = [\mathbf{s}_0^\intercal \mid \mathbf{s}_1^\intercal] \leftarrow \mathbb{Z}_q^{1\times 2\lambda},$$
$$\mathsf{Samp}_P: \quad \mathbf{P} = \begin{bmatrix} \mathbf{A}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_1 \end{bmatrix}, \quad \mathbf{A}_0, \mathbf{A}_1 \leftarrow \mathsf{TrapGen}, \quad \mathsf{aux}_P = \mathsf{iO}(\mathsf{Vfy}_{0,\mathbf{P},\mathsf{td}_{\mathbf{A}_0}}).$$

We describe the obfuscated circuit Vfy_0 later. For now, we only need to know that it hardcodes the trapdoor $\mathsf{td}_{\mathbf{A}_0}$ but not $\mathsf{td}_{\mathbf{A}_1}$.

We can already see why this block structure of $\mathbf{P}$ is useful. The LWE sample $\mathbf{SB} + \mathbf{E}_B$ can be written as $\mathbf{s}_0^\intercal\mathbf{B}_0 + \mathbf{s}_1^\intercal\mathbf{B}_1 + \mathbf{E}_B$, where $\mathbf{B} = \begin{bmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{bmatrix}$, while $\mathbf{SP} + \mathbf{E}_P$ can be written as the concatenation $[\mathbf{s}_0^\intercal\mathbf{A}_0 + \mathbf{E}_{P,0} \mid \mathbf{s}_1^\intercal\mathbf{A}_1 + \mathbf{E}_{P,1}]$. When the errors $\mathbf{E}_{P,0}$ and $\mathbf{E}_{P,1}$ are fresh and independent, having $\mathsf{td}_{\mathbf{A}_0}$ in the view reveals information about $\mathbf{s}_0$ but not $\mathbf{s}_1$. It therefore does not affect the pseudorandomness of the tuple $(\mathbf{s}_1^\intercal\mathbf{B}_1 + \mathbf{E}_B,\ \mathbf{s}_1^\intercal\mathbf{A}_1 + \mathbf{E}_{P,1})$ under the secret $\mathbf{s}_1$. On the other hand, when the errors are structured, $\mathbf{E}_{P,0} = \mathbf{E}_B\mathbf{K}_0$ and $\mathbf{E}_{P,1} = \mathbf{E}_B\mathbf{K}_1$, having $\mathsf{td}_{\mathbf{A}_0}$ allows us to recover $\mathbf{E}_B\mathbf{K}_0$. If $\mathbf{K}_0$ has full row rank over $\mathbb{Q}$, which can be ensured with overwhelming probability by appropriate parameter choices, we can then recover $\mathbf{E}_B$ and subsequently the full secret $\mathbf{S}$.

We can now describe the circuit Vfy_0 with one-bit output, which recovers $\mathbf{E}_B$ and distinguishes LWE samples from uniform samples.

$\mathsf{Vfy}_{0,\mathbf{P},\mathsf{td}_{\mathbf{A}_0}}(\mathbf{B},\ \mathbf{C}_B,\ \mathbf{K} = [\mathbf{K}_0 \mid \mathbf{K}_1])$:

1. First, verify that $\mathbf{BK} = \mathbf{P}$ and $\mathbf{K}$ is short.
2. Second, compute $\mathbf{C}_B\mathbf{K}_0$ and decode its error using $\mathsf{td}_{\mathbf{A}_0}$ to recover a candidate $\mathbf{E}_B$. Check that $\mathbf{E}_B$ is a short integer vector.
3. Finally, check that $\mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B$ for some $\mathbf{S}$. Output 1 if all checks pass, and 0 otherwise.

The circuit Vfy_0 can be used to break the postcondition. Given an honest LWE sample $\mathbf{C}_B = \mathbf{SB}+\mathbf{E}_B$ and the sampled preimage $\mathbf{K}$, it outputs 1 with overwhelming probability. On the other hand, a uniform $\mathbf{U}_B$ is unlikely to admit a representation $\mathbf{SB} + \mathbf{E}_B$ with short $\mathbf{E}_B$, so the circuit rejects with overwhelming probability. It remains to argue that the precondition holds for this sampler.

Proving the precondition. To prove the precondition, we need to show that the LWE samples $\mathbf{C}_B$ and $\mathbf{C}_P$ are pseudorandom in the following view:

$$\mathcal{H}_0 : \quad \begin{pmatrix} \mathbf{B},\ \mathbf{P},\ \mathbf{C}_B = \mathbf{s}_0^\intercal \mathbf{B}_0 + \mathbf{s}_1^\intercal \mathbf{B}_1 + \mathbf{E}_B, \\ \mathbf{C}_P = [\mathbf{s}_0^\intercal \mathbf{A}_0 + \mathbf{E}_{P,0} \mid \mathbf{s}_1^\intercal \mathbf{A}_1 + \mathbf{E}_{P,1}],\ \mathsf{aux} = \mathsf{iO}(\mathsf{Vfy}_{0,\mathbf{P},\mathsf{td}_{\mathbf{A}_0}}) \end{pmatrix}.$$

Since $\mathsf{td}_{\mathbf{A}_0}$ reveals no information about $\mathbf{s}_1$ and all errors are fresh in the precondition view, we can apply LWE with secret $\mathbf{s}_1$ to obtain

$$\mathcal{H}_1 : \quad (\mathbf{B},\ \mathbf{P},\ \mathbf{C}_B = \$,\ \mathbf{C}_P = [\mathbf{s}_0^\intercal \mathbf{A}_0 + \mathbf{E}_{P,0} \mid \$],\ \mathsf{aux} = \mathsf{iO}(\mathsf{Vfy}_{0,\mathbf{P},\mathsf{td}_{\mathbf{A}_0}})).$$

The remaining challenge is to replace $\mathbf{s}_0^\intercal \mathbf{A}_0 + \mathbf{E}_{P,0}$ with a uniform sample, which cannot be done by the same LWE argument while $\mathsf{td}_{\mathbf{A}_0}$ remains in view. The key observation is that the circuit can instead be implemented using $\mathsf{td}_{\mathbf{A}_1}$. The alternative circuit $\mathsf{Vfy}_{1,\mathbf{P},\mathsf{td}_{\mathbf{A}_1}}$ has the same structure, but recovers $\mathbf{E}_B$ from $\mathbf{C}_B\mathbf{K}_1 = \mathbf{s}_1^\intercal \mathbf{A}_1 + \mathbf{E}_B\mathbf{K}_1$. The two circuits are *functionally equivalent*: every input accepted by Vfy_0 satisfies $\mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B$ for some $\mathbf{S}$ and short $\mathbf{E}_B$. Together with the checks that $\mathbf{K}$ is short and $\mathbf{BK} = \mathbf{P}$, this ensures that Vfy_1 can decode $\mathbf{E}_B\mathbf{K}_1$, recover the same $\mathbf{E}_B$, and accept. The reverse direction holds as well. Therefore, by iO security, we can replace $\mathsf{iO}(\mathsf{Vfy}_{0,\mathbf{P},\mathsf{td}_{\mathbf{A}_0}})$ with $\mathsf{iO}(\mathsf{Vfy}_{1,\mathbf{P},\mathsf{td}_{\mathbf{A}_1}})$ in the view. We can then apply LWE with secret $\mathbf{s}_0$ to replace $\mathbf{s}_0^\intercal \mathbf{A}_0 + \mathbf{E}_{P,0}$ with a uniform sample. This completes the argument for the precondition.

We refer the reader to Section 5 for the detailed analysis of the binding counterexample.

2.4 Counterexample for Hiding Evasive LWE

Next we move to the hiding evasive LWE counterexample. In the hiding setting, the target matrix $\mathbf{P}$ has to be noisy in order to satisfy the target-hiding precondition. This complicates the design.

Adaptation for noisy targets. To support the target-hiding precondition $\mathbf{P} \approx \mathbf{P} + \mathbf{\Delta}$, the natural approach is to sample $\mathbf{P}$ with noise, i.e., $\mathbf{P} = \widetilde{\mathbf{P}} + \mathbf{F}$, where $\widetilde{\mathbf{P}}$ is a structured matrix and $\mathbf{F}$ is fresh Gaussian noise. With a preimage $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{P})$ for the noisy target, one can obtain

$$(\mathbf{SB} + \mathbf{E}_B)\mathbf{K} = \mathbf{SP} + \mathbf{E}_B\mathbf{K} = \mathbf{S}\widetilde{\mathbf{P}} + (\mathbf{SF} + \mathbf{E}_B\mathbf{K}).$$

We work in the setting where the secret $\mathbf{S}$ has small norm, so that the resulting expression is a noisy LWE sample for the matrix $\widetilde{\mathbf{P}}$. This is also the typical use case of hiding evasive LWE.

Compared with the binding counterexample, it is not immediate how to recover $\mathbf{E}_B$ from the joint noise $\mathbf{N} = \mathbf{SF} + \mathbf{E}_B\mathbf{K}$. Indeed, the solution is not even unique, as any $\mathbf{E}' = \mathbf{E}_B + \mathbf{SD}$, $\mathbf{F}' = \mathbf{F} - \mathbf{DK}$ for some matrix $\mathbf{D}$ also satisfies $\mathbf{SF}' + \mathbf{E}'\mathbf{K} = \mathbf{SF} + \mathbf{E}_B\mathbf{K}$. The trick here is to use the difference in scale between the terms. In typical applications of hiding evasive LWE, we set $\mathbf{E}_B$ to have a relatively large norm compared to the other terms, as it typically serves as the "smudging" term to flood structured noise terms. Therefore, we can hope to approximately recover $\mathbf{E}_B$ from the joint noise $\mathbf{N}$. In particular, for a wide preimage matrix $\mathbf{K}$ sampled from a discrete Gaussian, with high probability we can find a rational right inverse $\mathbf{J}$ such that $\mathbf{KJ} = \mathbf{I}$ and $\mathbf{J}$ has a relatively small norm (Corollary 5). Therefore, we can approximately recover $\mathbf{E}_B$ as $\lfloor \mathbf{NJ} \rceil$. We rely on this approximate recovery in our obfuscation strategy.

Sampler. We now describe the sampler for our counterexample. While it seems likely that we can adapt the binding counterexample to this setting using the above approximate error recovery technique, we instead redesign the sampler to exploit the full power of the hiding setting. In particular, we can reduce the assumption to just null-iO instead of full-fledged iO. Our design crucially relies on the fact that $\mathbf{P}$ is hidden from the adversary's view. The counterexample sampler is as follows.

$$\begin{aligned}
&\mathsf{Samp}_{\mathrm{S}}: \quad \mathbf{S} \leftarrow \mathcal{D}^{1\times n},\\
&\mathsf{Samp}_{\mathrm{P}}: \quad \mathbf{P} = \Big[\underbrace{\mathbf{T}\mathbf{A}_0 + \mathbf{F}_0}_{\text{Error solving slot}} \,\Big|\, \underbrace{\mathbf{A}_1 + \mathbf{F}_1 \,|\, \mathbf{A}_1\mathbf{R} + a\mathbf{I} + \mathbf{F}_2}_{\text{Fragile trapdoor}} \,\Big|\, \underbrace{\mathbf{A}_3 + \mathbf{F}_3}_{\text{Sparsifier}}\Big], \quad \mathsf{aux}_{\mathrm{P}} = \mathsf{iO}(\mathsf{Vfy}_{\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\dots}).
\end{aligned}$$

where $\mathbf{T}$ is a random invertible square matrix, $\mathbf{A}_0$ is sampled with a trapdoor, $\mathbf{A}_1, \mathbf{A}_3$ are sufficiently wide uniform matrices sampled without trapdoors, $\mathbf{R}$ is a binary matrix serving as the fragile trapdoor of the matrix $[\mathbf{A}_1 \mid \mathbf{A}_1\mathbf{R} + a\mathbf{I}]$, and a is an appropriate scalar such that $\|\mathbf{E}\|_\infty \gg a \gg \|\mathbf{F}\|_\infty$, i.e., it is much larger than the perturbation noise, but much smaller than the evasive LWE sample noise. Again, we ignore the circuit Vfy for now, and remember only that it hardcodes the trapdoor information for $\mathbf{A}_0$ and $\mathbf{R}$.

Now, looking at the LWE sample $\mathbf{C}_{\mathrm{P}} = \mathbf{S}[\mathbf{P}_0 \mid \mathbf{P}_1 \mid \mathbf{P}_2 \mid \mathbf{P}_3] + [\mathbf{E}_{\mathrm{P},0} \mid \mathbf{E}_{\mathrm{P},1} \mid \mathbf{E}_{\mathrm{P},2} \mid \mathbf{E}_{\mathrm{P},3}]$, we see that the trapdoor $\mathsf{td}_{\mathbf{A}_0}$ lets us decode the first slot $\mathbf{S}\mathbf{P}_0 + \mathbf{E}_{\mathrm{P},0} = \mathbf{S}\mathbf{T}\mathbf{A}_0 + \mathbf{N}_0$ and recover the joint noise $\mathbf{N}_0 = \mathbf{S}\mathbf{F}_0 + \mathbf{E}_{\mathrm{P},0}$. Furthermore, since $\mathbf{T}$ is private and hidden from the remaining view, recovering $\mathbf{S}\mathbf{T}$ reveals no information about $\mathbf{S}$ itself. In other words, the power of $\mathsf{td}_{\mathbf{A}_0}$ is limited to recovering the noise.

In the precondition, the independent fresh error $\mathbf{E}_{\mathrm{P},0}$ floods $\mathbf{S}\mathbf{F}_0$, so recovering $\mathbf{N}_0$ reveals no information about the remaining slots of $\mathbf{C}_{\mathrm{P}}$. Furthermore, in the fragile trapdoor slot, the fresh error $\mathbf{E}_{\mathrm{P},2}$ smudges the fragile structure from $a\mathbf{I}$, so $\mathbf{R}$ also leaks no information about the secret $\mathbf{S}$.

On the other hand, in the postcondition, the decoded noise $\mathbf{N}_0 = \mathbf{S}\mathbf{F}_0 + \mathbf{E}_{\mathrm{B}}\mathbf{K}_0$ has a structured form. Using a right inverse $\mathbf{J}$ of $\mathbf{K}_0$, we can recover the estimate $\widehat{\mathbf{E}}_{\mathrm{B}} = \lfloor \mathbf{N}_0\mathbf{J} \rceil$ of the source error. With this, we can "prune" the LWE sample $\mathbf{C}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}$ by subtracting the estimate and obtaining $\widetilde{\mathbf{C}}_{\mathrm{B}} = \mathbf{C}_{\mathrm{B}} - \widehat{\mathbf{E}}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \widetilde{\mathbf{E}}_{\mathrm{B}}$, where the residual error $\widetilde{\mathbf{E}}_{\mathrm{B}} = \mathbf{E}_{\mathrm{B}} - \widehat{\mathbf{E}}_{\mathrm{B}}$ is small. This in turn allows us to compute $\widetilde{\mathbf{C}}_{\mathrm{B}}[\mathbf{K}_1 \mid \mathbf{K}_2] = [\mathbf{S}\mathbf{A}_1 \mid \mathbf{S}\mathbf{A}_1\mathbf{R} + a\mathbf{S}]$ + small noise, with the noise term much smaller than the scalar a. In this setting, the secret $\mathbf{S}$ can be recovered using the fragile trapdoor $\mathbf{R}$.

This gives our first attempt at defining the verification circuit Vfy, which approximately recovers $\mathbf{E}_{\mathrm{B}}$ and then recovers $\mathbf{S}$ for a consistency check.

$\mathsf{Vfy}^{*}_{\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{A}_3}(\mathbf{C}_{\mathrm{B}}, \mathbf{K} = [\mathbf{K}_0 \mid \mathbf{K}_1 \mid \mathbf{K}_2 \mid \mathbf{K}_3])$:

1. First, verify that $\mathbf{K}$ is short.
2. Second, decode $\mathbf{C}_{\mathrm{B}}\mathbf{K}_0$ using $\mathsf{td}_{\mathbf{A}_0}$ to recover the noise $\mathbf{N}_0$, and set $\widehat{\mathbf{E}}_{\mathrm{B}} = \lfloor \mathbf{N}_0\mathbf{J} \rceil$, where $\mathbf{K}_0\mathbf{J} = \mathbf{I}$. Check that $\widehat{\mathbf{E}}_{\mathrm{B}}$ is a short integer vector, and prune the LWE sample $\widetilde{\mathbf{C}}_{\mathrm{B}} = \mathbf{C}_{\mathrm{B}} - \widehat{\mathbf{E}}_{\mathrm{B}}$.
3. Use the fragile trapdoor $\mathbf{R}$ on $\widetilde{\mathbf{C}}_{\mathrm{B}}[\mathbf{K}_1 \mid \mathbf{K}_2]$ to recover a candidate secret $\widehat{\mathbf{S}}$. Reject if $\widehat{\mathbf{S}} = \mathbf{0}$.
4. Check that $\widehat{\mathbf{S}}$ is consistent with the final slot, i.e., $\widehat{\mathbf{S}}\mathbf{A}_3 \approx \widetilde{\mathbf{C}}_{\mathrm{B}}\mathbf{K}_3$. Output 1 if all checks pass, and 0 otherwise.

The above auxiliary input is sufficient for breaking the postcondition: the consistency check accepts a real postcondition sample with high probability and rejects a uniform sample with overwhelming probability. On the other hand, we have informally established that in the precondition, recovering $\mathbf{S}$ from $\mathsf{td}_{\mathbf{A}_0}$ and $\mathbf{R}$ is computationally hard. However, inversion hardness is not enough to apply iO security. We therefore apply a standard trick of sparsification. Instead of checking $\widehat{\mathbf{S}}\mathbf{A}_3 \approx \widetilde{\mathbf{C}}_{\mathrm{B}}\mathbf{K}_3$ directly, we hardcode an LWE sample $\mathbf{M} = \mathbf{A}_3\mathbf{W} + \mathbf{F}_{\mathrm{M}}$ with binary public matrix $\mathbf{W}$ and small error $\mathbf{F}_{\mathrm{M}}$, and check that $\widehat{\mathbf{S}}\mathbf{M} \approx \widetilde{\mathbf{C}}_{\mathrm{B}}\mathbf{K}_3\mathbf{W}$, leading to the final verification circuit

$\mathsf{Vfy}_{\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}=\mathbf{A}_3\mathbf{W}+\mathbf{F}_{\mathrm{M}}}(\mathbf{C}_{\mathrm{B}},\ \mathbf{K} = [\mathbf{K}_0 \mid \mathbf{K}_1 \mid \mathbf{K}_2 \mid \mathbf{K}_3])$:

1–3. Same as Vfy^*.

4. Check that $\widehat{\mathbf{S}}$ is consistent with the final slot via $\widehat{\mathbf{S}}\mathbf{M} \approx \widetilde{\mathbf{C}}_{\mathrm{B}}\mathbf{K}_3\mathbf{W}$.

Intuitively, an LWE argument with secret $\mathbf{S}$ hides $\mathbf{A}_3$ in the precondition samples, allowing us to apply LWE again with $\mathbf{A}_3$ as the secret and replace $\mathbf{M}$ by a uniform matrix independent of $\mathbf{W}$. Once this replacement is made, with overwhelming probability the check $\widehat{\mathbf{S}}\mathbf{M} \approx \widetilde{\mathbf{C}}_{\mathrm{B}}\mathbf{K}_3\mathbf{W}$ will have no non-trivial solution, allowing us to replace the obfuscated circuit with an identically zero circuit via null-iO security.

Proving the precondition. With all the intuition in place, we now walk through how the precondition is established assuming null-iO. We refer the reader to Section 6 for the detailed analysis.

To prove the precondition, we need to show that the LWE samples $\mathbf{C}_{\mathrm{B}}$ and $\mathbf{C}_{\mathrm{P}}$ are pseudorandom in the following view:

$$\mathcal{H}_0 : \begin{pmatrix} \mathbf{C}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}, \\ \mathbf{C}_{\mathrm{P}} = \left[\mathbf{S}\mathbf{T}\mathbf{A}_0 + \mathbf{S}\mathbf{F}_0 \mid \mathbf{S}\mathbf{A}_1 + \mathbf{S}\mathbf{F}_1 \mid \mathbf{S}\mathbf{A}_1\mathbf{R} + a\mathbf{S} + \mathbf{S}\mathbf{F}_2 \mid \mathbf{S}\mathbf{A}_3 + \mathbf{S}\mathbf{F}_3\right] \\ \qquad + [\mathbf{E}_{\mathrm{P},0} \mid \mathbf{E}_{\mathrm{P},1} \mid \mathbf{E}_{\mathrm{P},2} \mid \mathbf{E}_{\mathrm{P},3}], \\ \mathsf{aux} = \mathsf{iO}\left(\mathsf{Vfy}_{\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}=\mathbf{A}_3\mathbf{W}+\mathbf{F}_{\mathrm{M}}}\right) \end{pmatrix}.$$

The fact that $\mathbf{E}$ has a much larger norm allows us to smudge all the lower-bit structures, leading to the statistically close distribution

$$\mathcal{H}_1 : \begin{pmatrix} \mathbf{C}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}, \\ \mathbf{C}_{\mathrm{P}} = \left[\mathbf{S}\mathbf{T}\mathbf{A}_0 + \mathbf{E}_{\mathrm{P},0} \mid \mathbf{S}\mathbf{A}_1 + \mathbf{E}_{\mathrm{P},1} \mid \mathbf{S}\mathbf{A}_1\mathbf{R} + \mathbf{E}_{\mathrm{P},2} \mid \mathbf{S}\mathbf{A}_3 + \mathbf{E}_{\mathrm{P},3}\right] \\ \mathsf{aux} = \mathsf{iO}\left(\mathsf{Vfy}_{\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}=\mathbf{A}_3\mathbf{W}+\mathbf{F}_{\mathrm{M}}}\right) \end{pmatrix}.$$

As elaborated previously, the term $\mathbf{S}\mathbf{T}$ perfectly hides $\mathbf{S}$. Therefore we can simply apply LWE with secret $\mathbf{S}$ to get

$$\mathcal{H}_2 : \begin{pmatrix} \mathbf{C}_{\mathrm{B}} \leftarrow \$, \\ \mathbf{C}_{\mathrm{P}} = \left[\mathbf{S}\mathbf{T}\mathbf{A}_0 + \mathbf{E}_{\mathrm{P},0} \mid \mathbf{U}_1 \leftarrow \$ \mid \mathbf{U}_1\mathbf{R} + \mathbf{E}_{\mathrm{P},2} \mid \mathbf{U}_3 \leftarrow \$\right] \\ \mathsf{aux} = \mathsf{iO}\left(\mathsf{Vfy}_{\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}=\mathbf{A}_3\mathbf{W}+\mathbf{F}_{\mathrm{M}}}\right) \end{pmatrix}.$$

In this view, $\mathbf{A}_3$ is no longer used in the LWE samples $\mathbf{C}_{\mathrm{P}}$. Therefore we can apply LWE with $\mathbf{A}_3$ as the secret and replace $\mathbf{M}$ by a uniform matrix. Once this replacement is made, Vfy becomes a

null circuit (with overwhelming probability), and can be replaced with a dummy circuit with no information about $\mathsf{td}_{\mathbf{A}_0}$ and $\mathbf{R}$ hardcoded. The pseudorandomness of the remaining components is then straightforward.

3 Preliminaries

Notations. The security parameter is denoted as λ. Unless stated otherwise, all algorithms are probabilistic polynomial-time. We use $\approx$ to denote computational indistinguishability of ensembles, and use $\approx_s$ to denote statistical indistinguishability of ensembles. The symbol $\perp$ denotes the empty string.

We use $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ for the integers, rationals, and reals, respectively. Throughout this paper, the modulus q is an odd prime, and $\mathbb{Z}_q$ denotes the integers modulo q. We operate over $\mathbb{Z}_q$ and $\mathbb{Z}$ interchangeably when clear from context: integers are reduced modulo q, while elements of $\mathbb{Z}_q$ are lifted to their centered representatives in $[-\lfloor q/2 \rfloor, \lfloor q/2 \rfloor]$. We use log for the base-two logarithm and ln for the natural logarithm. We write $\lfloor x \rceil := \lfloor x + 1/2 \rfloor$ for rounding to the nearest integer, applied entrywise to vectors and matrices. In rounded quotients, a numerator over $\mathbb{Z}_q$ is lifted to its centered integer representative before division.

All vectors are column vectors unless stated otherwise. We use boldface lowercase letters for vectors and boldface uppercase letters for matrices, and write $\mathbf{A}^\intercal$ for the transpose of $\mathbf{A}$. We write $\mathbf{0}_n$ for the zero vector in dimension n, $\mathbf{0}_{n\times m}$ for the $n \times m$ zero matrix, and $\mathbf{I}_n$ for the $n \times n$ identity matrix. We omit these dimension subscripts when clear from context. For a field $\mathbb{F}$, $\mathsf{rank}_{\mathbb{F}}(\mathbf{A})$ denotes the dimension of the column span of $\mathbf{A}$ over $\mathbb{F}$. The notation $\|\mathbf{k}\|_\infty$ denotes the infinity norm; for a matrix, we use $\|\mathbf{K}\|_\infty$ to denote the maximum absolute entry. For $\mathbb{Z}_q$ vectors and matrices, this norm is computed using the centered representatives in $\mathbb{Z}$.

For a matrix $\mathbf{B} \in \mathbb{Z}_q^{n\times m}$, we define lattice cosets

$$\Lambda_{\mathbf{p}}^{\perp}(\mathbf{B}) = \{\mathbf{k} \in \mathbb{Z}^m : \mathbf{B}\mathbf{k} = \mathbf{p} \pmod q\}, \quad \Lambda^{\perp}(\mathbf{B}) := \Lambda_{\mathbf{0}}^{\perp}(\mathbf{B}).$$

For a lattice coset $L \subseteq \mathbb{R}^m$ and width $\sigma > 0$, $\mathcal{D}_{L,\sigma}$ denotes the discrete Gaussian distribution on L defined by

$$\mathcal{D}_{L,\sigma}(\mathbf{x}) := \frac{\exp(-\pi\|\mathbf{x}\|_2^2/\sigma^2)}{\sum_{\mathbf{y}\in L}\exp(-\pi\|\mathbf{y}\|_2^2/\sigma^2)}, \qquad \mathbf{x} \in L.$$

We abbreviate $\mathcal{D}_{\mathbb{Z},\sigma}$ as $\mathcal{D}_\sigma$. The notation $\mathcal{D}_\sigma^n$ or $\mathcal{D}_\sigma^{n\times m}$ denotes a vector or matrix whose entries are sampled independently from $\mathcal{D}_\sigma$.

For a full-rank lattice $\Lambda \subseteq \mathbb{R}^m$, let $\Lambda^* = \{\mathbf{y} \in \mathbb{R}^m : \langle \mathbf{y}, \mathbf{x}\rangle \in \mathbb{Z}\ \forall \mathbf{x} \in \Lambda\}$ be its dual lattice. For $\varepsilon \in (0,1)$, the smoothing parameter $\eta_\varepsilon(\Lambda)$ is the smallest $s > 0$ such that

$$\sum_{\mathbf{y}\in\Lambda^*\setminus\{\mathbf{0}\}} \exp(-\pi s^2\|\mathbf{y}\|_2^2) \le \varepsilon.$$

For the integer lattice, the standard smoothing bound [MR04, Lemma 3.3] gives

$$\eta_\varepsilon(\mathbb{Z}^m) \le \sqrt{\frac{\ln(2m(1+1/\varepsilon))}{\pi}}.$$

3.1 Gaussian Preimage Sampling

We start by reviewing the standard lattice trapdoors [GPV08, MP12].

Lemma 1 (Adapted from [MP12])**.** *Fix the thresholds*

$$m_0(n, q) = 6n\lceil \log_2 q \rceil, \qquad \rho_0(n, m) = 32m.$$

There are efficient algorithms $\mathsf{TrapGen}$, SampD, *and* Decode *satisfying the following conditions:*

- $\mathsf{TrapGen}(1^n, 1^m, q)$ *takes as input* $n \geq 1$, $q \geq 2$, *and* $m \geq m_0(n, q)$. *It outputs* $(\mathbf{B}, \mathsf{td}_{\mathbf{B}})$ *such that* $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ *and* $\mathbf{B}$ *is* $\mathsf{negl}(n)$*-close to uniform over* $\mathbb{Z}_q^{n \times m}$.
- $\mathsf{SampD}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{p}, \rho)$ *takes as input* $\mathbf{B}, \mathsf{td}_{\mathbf{B}}$ *from* $\mathsf{TrapGen}$, *some* $\mathbf{p} \in \mathbb{Z}_q^n$, *and* $\rho \geq \rho_0(n, m)$. *It outputs* $\mathbf{k} \in \mathbb{Z}^m$ *such that* $\mathbf{B}\mathbf{k} = \mathbf{p} \pmod q$, $\|\mathbf{k}\|_\infty \leq \rho\sqrt{m}$, *and* $\mathbf{k}$ *is* $\mathsf{negl}(n)$*-close to* $\mathcal{D}_{\Lambda_{\mathbf{p}}^\perp(\mathbf{B}), \rho}$.
- $\mathsf{Decode}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{c})$ *is a deterministic LWE decoding algorithm that, for every pair* $(\mathbf{B}, \mathsf{td}_{\mathbf{B}})$ *output by* $\mathsf{TrapGen}$, *every* $\mathbf{s} \in \mathbb{Z}_q^n$, *and every* $\mathbf{e} \in \mathbb{Z}^m$ *with* $\|\mathbf{e}\|_\infty < q/(8m)$, *recovers* $(\mathbf{s}, \mathbf{e})$ *from the column vector* $\mathbf{c} \in \mathbb{Z}_q^m$ *satisfying* $\mathbf{c}^\mathsf{T} = \mathbf{s}^\mathsf{T}\mathbf{B} + \mathbf{e}^\mathsf{T} \pmod q$. *It returns* $\perp$ *if no such pair* $(\mathbf{s}, \mathbf{e})$ *exists.*

We write $\mathbf{k} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{p})$ to denote a sample output by $\mathsf{SampD}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{p}, \rho)$. When clear from the context, or when the specific width is not relevant to the discussion, we may omit the subscript ρ.

It is convenient to extend SampD to process multiple $\mathbf{p}$'s in one shot. For a matrix $\mathbf{P} = [\mathbf{p}_1, \ldots, \mathbf{p}_\ell]$, viewed as a batch of vectors, we define $\mathsf{SampD}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{P}, \rho)$ as

$$\mathbf{K} \leftarrow \big[\mathsf{SampD}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{p}_1, \rho), \ \ldots, \ \mathsf{SampD}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{p}_\ell, \rho)\big],$$

with fresh randomness for each call to SampD on the right-hand side. Similarly, we denote the output of the batch sampling by $\mathbf{B}_\rho^{-1}(\mathbf{P})$.

We also write $\mathsf{Decode}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{C})$ for batch LWE decoding of $\mathbf{C} = \mathbf{S}\mathbf{B} + \mathbf{E} \in \mathbb{Z}_q^{t \times m}$. Write its jth row as $\mathbf{c}_j^\mathsf{T}$, where $\mathbf{c}_j \in \mathbb{Z}_q^m$ is a column vector. The algorithm computes

$$(\mathbf{s}_j, \mathbf{e}_j) = \mathsf{Decode}(\mathbf{B}, \mathsf{td}_{\mathbf{B}}, \mathbf{c}_j), \qquad j \in \{1, \ldots, t\},$$

and returns $(\mathbf{S}, \mathbf{E})$, where

$$\mathbf{S} = [\mathbf{s}_1, \ldots, \mathbf{s}_t]^\mathsf{T} \in \mathbb{Z}_q^{t \times n}, \qquad \mathbf{E} = [\mathbf{e}_1, \ldots, \mathbf{e}_t]^\mathsf{T} \in \mathbb{Z}^{t \times m}.$$

It rejects if any individual call rejects.

We also recall the following lemmas on matrices with trapdoors.

Lemma 2 (Small smoothing parameter [GPV08, Lemma 5.3])**.** *Let* q *be prime and* $m \geq m_0(n, q)$, *and sample* $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$. *There is an* $\varepsilon \leq \mathsf{negl}(m)$ *such that, except with probability* $\mathsf{negl}(n)$,

$$\eta_\varepsilon\big(\Lambda_{\mathbf{0}}^\perp(\mathbf{B})\big) \leq \log m. \tag{1}$$

Lemma 3 (Gaussian preimages of uniform targets [GPV08, Lemma 5.2])**.** *Let* $(\mathbf{B}, \mathsf{td}_{\mathbf{B}})$ *be an output of* $\mathsf{TrapGen}(1^n, 1^m, q)$ *with* $m \geq m_0(n, q)$ *satisfying the smoothing bound (Equation (1)) for some* $\varepsilon \leq \mathsf{negl}(m)$. *For every* $\rho \geq \rho_0(n, m)$, *sampling* $\mathbf{p} \leftarrow \mathbb{Z}_q^n$ *uniformly and then* $\mathbf{k} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{p})$ *gives a marginal distribution of* $\mathbf{k}$ *that is* $\mathsf{negl}(n)$*-close to* $\mathcal{D}_\rho^m$.

3.2 Gaussian Tails and Flooding

We recall some basic properties of the discrete Gaussian distribution.

Lemma 4 (Gaussian tails [MP12, Lemma 2.8]). *Let $\Lambda \subseteq \mathbb{R}^m$ be a full-rank lattice, $L = \Lambda + \mathbf{c}$ for $\mathbf{c} \in \mathbb{R}^m$, $\varepsilon \in (0,1)$, and $\sigma \geq \eta_\varepsilon(\Lambda)$. For every $r > 0$,*

$$\Pr_{\mathbf{x} \leftarrow \mathcal{D}_{L,\sigma}}[\|\mathbf{x}\|_\infty > r\sigma] \leq 2m\frac{1+\varepsilon}{1-\varepsilon}e^{-\pi r^2}.$$

Corollary 1 (Tail bound for $\mathbf{B}^{-1}(\mathbf{P})$). *Let $m \geq m_0(n,q)$ and $\rho \geq \rho_0(n,m)$. Sample $(\mathbf{B}, \mathsf{td}_\mathbf{B}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$, choose any target $\mathbf{P} \in \mathbb{Z}_q^{n\times\ell}$, and sample $\mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P})$. Then*

$$\Pr[\|\mathbf{K}\|_\infty > \rho \log n] \leq \mathsf{negl}(n).$$

Lemma 5 (Gaussian flooding [BDE+18, Lemma 3.2]). *Let $X \leftarrow \mathcal{D}_\sigma$ for $\sigma \geq 1.60$. For every $u \in \mathbb{Z}$,*

$$\mathrm{SD}(\mathcal{D}_\sigma, \mathcal{D}_\sigma + u) \leq \frac{|u|}{\sqrt{2\mathbb{E}[X^2]}} = O\left(\frac{|u|}{\sigma}\right).$$

Lemma 6 (Gaussian point probabilities [PR06, Lemma 2.11]). *Let $\Lambda \subseteq \mathbb{R}^m$ be a full-rank lattice, $L = \Lambda + \mathbf{c}$ for $\mathbf{c} \in \mathbb{R}^m$, $\varepsilon \in (0,1)$, and $\sigma \geq 2\eta_\varepsilon(\Lambda)$. For every $\mathbf{y} \in L$,*

$$\Pr_{\mathbf{x} \leftarrow \mathcal{D}_{L,\sigma}}[\mathbf{x} = \mathbf{y}] \leq 2^{-m}\frac{1+\varepsilon}{1-\varepsilon}.$$

Corollary 2 (Nonzero Gaussian samples). *Let $\sigma > \sqrt{4\ln(8n)/\pi}$ and sample $\mathbf{s} \leftarrow \mathcal{D}_\sigma^n$. Then*

$$\Pr[\mathbf{s} = \mathbf{0}] \leq \mathsf{negl}(n).$$

Proof. The smoothing bound for $\mathbb{Z}^n$ gives $\sigma > 2\eta_{1/3}(\mathbb{Z}^n)$. Applying Lemma 6 with $\varepsilon = 1/3$ bounds the probability by $2^{1-n} \leq \mathsf{negl}(n)$. □

3.3 Rank of Gaussian Samples

In our analysis, we repeatedly need to invert matrices over both $\mathbb{Z}_q$ and $\mathbb{Q}$. We therefore recall a few results on the rank of Gaussian matrices.

First, following [ABB10b], we argue that Gaussian preimages of a uniform target are likely to have full rank modulo q.

Lemma 7 ([ABB10b, Theorem 30]). *Let $\rho > \omega(\sqrt{\log m})$ and $q > \rho\omega(\sqrt{\log m})$. Then the probability that m vectors in $\mathbb{Z}^m$ sampled independently from $\mathcal{D}_{\mathbb{Z}^m,\rho}$ are $\mathbb{Z}_q$-linearly independent is at least $1 - \frac{1+\varepsilon}{\rho-1}$ for some negligible function ε.*

Corollary 3 (Full rank of $\mathbf{B}^{-1}(\mathbf{P})$ over $\mathbb{Z}_q$ of uniform targets). *Let $m \geq m_0(n,q)$, $\ell \geq m$, $\rho \geq \rho_0(n,m)$, and $q > \rho\omega(\sqrt{\log m})$. Sample*

$$(\mathbf{B}, \mathsf{td}_\mathbf{B}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q), \qquad \mathbf{P} \leftarrow \mathbb{Z}_q^{n\times\ell}, \qquad \mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P}),$$

where $\mathbf{P}$ is uniform and independent of $(\mathbf{B}, \mathsf{td}_\mathbf{B})$. Then

$$\Pr[\mathsf{rank}_{\mathbb{Z}_q}(\mathbf{K}) < m] \leq \frac{1}{\rho-1} + \mathsf{negl}(n).$$

Proof. By Lemmas 2 and 3, the marginal distribution of $\mathbf{K}$ is $\mathsf{negl}(n)$-close to $\mathbf{X} \leftarrow \mathcal{D}_{\rho}^{m\times\ell}$. The rest follows from Lemma 7. □

Next, following [PV08], we argue that in the rationals, Gaussian preimages of any target are likely to have full rank.

Lemma 8 ([PV08, Lemma 2.12])**.** *For any m-dimensional lattice Λ, any coset $L = \Lambda + \mathbf{c}$ with $\mathbf{c} \in \mathbb{R}^m$, any $\varepsilon \in (0,1)$ and $\rho \geq \sqrt{2}\,\eta_\varepsilon(\Lambda)$, and any fixed $(m-1)$-dimensional affine hyperplane $H \subset \mathbb{R}^m$,*

$$\Pr_{\mathbf{x}\leftarrow\mathcal{D}_{L,\rho}}[\mathbf{x} \in H] < \frac{1}{\sqrt{2}}\,\frac{1+\varepsilon}{1-\varepsilon}.$$

Corollary 4 (Full rank of $\mathbf{B}^{-1}(\mathbf{P})$ over $\mathbb{Q}$ for arbitrary targets)**.** *Let $m \geq m_0(n,q)$, $\ell \geq 4m$, and $\rho \geq \rho_0(n,m)$. Sample $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$, choose any target $\mathbf{P} \in \mathbb{Z}_q^{n\times\ell}$, and sample $\mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P})$. Then*

$$\Pr[\mathsf{rank}_{\mathbb{Q}}(\mathbf{K}) < m] \leq \mathsf{negl}(n).$$

Proof. Let $\Lambda = \Lambda_{\mathbf{0}}^{\perp}(\mathbf{B})$ and let $\varepsilon \leq \mathsf{negl}(m)$ be as in Lemma 2. The smoothing bound gives

$$\sqrt{2}\,\eta_\varepsilon(\Lambda) \leq \sqrt{2}\log m \leq \rho_0(n,m) \leq \rho$$

with probability at least $1 - \mathsf{negl}(n)$.

Now condition on the smoothing bound holding. For any target $\mathbf{P} = [\mathbf{p}_1|\cdots|\mathbf{p}_\ell]$, consider independent samples $\mathbf{k}_j \leftarrow \mathcal{D}_{\Lambda_{\mathbf{p}_j}^{\perp}(\mathbf{B}),\rho}$ in order. Condition on any values of the preceding columns whose span W over $\mathbb{Q}$ is a proper subspace of $\mathbb{Q}^m$, and choose a real hyperplane H containing W. Applying Lemma 8 to $\Lambda_{\mathbf{p}_j}^{\perp}(\mathbf{B})$ gives

$$\Pr[\mathbf{k}_j \in W \mid \mathbf{k}_1,\ldots,\mathbf{k}_{j-1}] \leq \Pr[\mathbf{k}_j \in H \mid \mathbf{k}_1,\ldots,\mathbf{k}_{j-1}] < \frac{1+\varepsilon}{\sqrt{2}(1-\varepsilon)}.$$

Thus, until the span is all of $\mathbb{Q}^m$, each column increases its dimension with conditional probability at least $p = 1 - \frac{1+\varepsilon}{\sqrt{2}(1-\varepsilon)}$. Since ε is negligible, $p - 1/4$ is bounded below by a positive constant for all sufficiently large n. The failure probability is therefore at most $\Pr[Z < m]$ for $Z \leftarrow \mathrm{Bin}(\ell, p)$. Hoeffding's inequality gives

$$\Pr[\mathsf{rank}_{\mathbb{Q}}([\mathbf{k}_1 \mid \cdots \mid \mathbf{k}_\ell]) < m] \leq \Pr[Z < m] \leq \Pr[Z < \ell/4] \leq e^{-2\ell(p-1/4)^2} = e^{-\Omega(\ell)}.$$

Accounting for the $\mathsf{negl}(n)$ statistical error of $\mathbf{B}_\rho^{-1}(\mathbf{P})$ proves the claim. □

3.4 Rational Right Inverses

We recall a lower bound on the smallest singular value of a discrete Gaussian matrix and use it to bound the norm of its rational right inverse.

Lemma 9 (Rational right inverse)**.** *Let $\mathbf{K} \in \mathbb{Q}^{m\times L}$ have full row rank, and set*

$$\mathbf{J} = \mathbf{K}^{\mathsf{T}}(\mathbf{K}\mathbf{K}^{\mathsf{T}})^{-1} \in \mathbb{Q}^{L\times m}.$$

Then

$$\mathbf{K}\mathbf{J} = \mathbf{I}_m, \qquad \|\mathbf{J}\|_\infty \leq \|\mathbf{J}\|_{\mathrm{op}} = \frac{1}{s_{\min}(\mathbf{K})},$$

where $s_{\min}(\mathbf{K})$ denotes the smallest singular value of $\mathbf{K}$, and $\|\cdot\|_{\mathrm{op}}$ denotes the Euclidean operator norm.

Proof. Full row rank makes $\mathbf{K}\mathbf{K}^\intercal$ invertible, so $\mathbf{K}\mathbf{J} = \mathbf{I}_m$. Moreover, $\mathbf{J}^\intercal\mathbf{J} = (\mathbf{K}\mathbf{K}^\intercal)^{-1}$, whose largest eigenvalue is $s_{\min}(\mathbf{K})^{-2}$. This proves the operator-norm identity. Each matrix entry is bounded in absolute value by the operator norm, giving the infinity-norm bound. □

The following bound applies to discrete Gaussians on any full-rank lattice.

Lemma 10 ([AGHS13, Lemma 8])**.** *There are absolute constants $c, C > 0$ with the following property. Let $\Lambda \subseteq \mathbb{R}^m$ be a full-rank lattice, $L = \mathsf{poly}(m) \geq 2m$, and $\varepsilon \leq \mathsf{negl}(m)$. For $\rho > C\eta_\varepsilon(\Lambda)$, sample the columns of $\mathbf{X} \in \mathbb{R}^{m\times L}$ independently from $\mathcal{D}_{\Lambda,\rho}$. Except with probability $\mathsf{negl}(m)$,*

$$s_{\min}(\mathbf{X}) \geq c\rho\sqrt{L}.$$

Corollary 5 (Small right inverse of $\mathbf{B}^{-1}(\mathbf{P})$ for uniform targets)**.** *Let $m \geq m_0(n,q)$, $L \geq m^3$, and $\rho \geq \rho_0(n,m)$. Sample*

$$(\mathbf{B}, \mathsf{td}_\mathbf{B}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q), \qquad \mathbf{P} \leftarrow \mathbb{Z}_q^{n\times L}, \qquad \mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P}),$$

where $\mathbf{P}$ is uniform and independent of $(\mathbf{B}, \mathsf{td}_\mathbf{B})$. Except with probability $\mathsf{negl}(n)$, the matrix $\mathbf{K}$ has full row rank over $\mathbb{Q}$ and its rational right inverse $\mathbf{J} = \mathbf{K}^\intercal(\mathbf{K}\mathbf{K}^\intercal)^{-1}$ satisfies

$$\mathbf{K}\mathbf{J} = \mathbf{I}_m, \qquad \|\mathbf{J}\|_\infty \leq \frac{1}{\rho}.$$

Proof. By Lemmas 2 and 3, the marginal distribution of $\mathbf{K}$ is $\mathsf{negl}(n)$-close to $\mathcal{D}_\rho^{m\times L}$. For $\varepsilon = e^{-m}$, the smoothing bound for $\mathbb{Z}^m$ gives $\eta_\varepsilon(\mathbb{Z}^m) = O(\sqrt{m})$, so $\rho \geq 32m$ satisfies the hypothesis of Lemma 10. Thus $s_{\min}(\mathbf{K}) \geq c\rho\sqrt{L}$ except with probability $\mathsf{negl}(n)$. In particular, $\mathbf{K}$ has full row rank, and the right-inverse bound (Lemma 9) gives

$$\|\mathbf{J}\|_\infty \leq \frac{1}{c\rho\sqrt{L}} \leq \frac{1}{\rho}.$$

□

3.5 Learning with Errors

Definition 1 (Learning with errors [Reg05])**.** For dimension n, sample count m, prime modulus q, and Gaussian width $\sigma > 0$, the assumption $\mathsf{LWE}_{n,m,q,\sigma}$ states that

$$(\mathbf{A}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal) \approx (\mathbf{A}, \mathbf{u}^\intercal),$$

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

Standard reductions show that LWE is hard assuming the worst-case hardness of SIVP with approximation factor $\mathsf{poly}(n) \cdot q/\sigma$ [Reg05, Pei09, BLP+13]. We take σ to be a sufficiently large polynomial and $q/\sigma \leq 2^{n^{1-\delta}}$ for some constant $\delta \in (0,1)$, so that the resulting SIVP approximation factor lies in a regime for which no polynomial-time algorithm is known [LLL82, GN08, ABB+23].

We also use the following two consequences of LWE. Lemma 11 allows a discrete Gaussian secret, while Lemma 12 allows a binary public matrix.

Lemma 11 (Gaussian-secret LWE [ACPS09, Lemma 2])**.** *Assume $\mathsf{LWE}_{n,m+n+\lambda,q,\sigma}$. Then*

$$(\mathbf{A}, \mathbf{s}^\intercal\mathbf{A} + \mathbf{e}^\intercal) \approx (\mathbf{A}, \mathbf{u}^\intercal),$$

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

Lemma 12 (Binary-public-matrix LWE [BLMR13, Theorem 4.3])**.** *Let $n' \geq n(\lceil \log_2 q \rceil + \lambda)$ and assume $\mathsf{LWE}_{n,m,q,\sigma}$. Then*

$$(\mathbf{A}, \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal) \approx (\mathbf{A}, \mathbf{u}^\intercal),$$

where $\mathbf{A} \leftarrow \{0,1\}^{n' \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^{n'}$, $\mathbf{e} \leftarrow \mathcal{D}_\sigma^m$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$.

The argument for binary public matrices in [BLMR13, Corollary 4.5] extends to arbitrary q by using $\lceil \log_2 q \rceil + \lambda$-bit representations of uniformly sampled congruent lifts.

3.6 Leftover Hash Lemma

In this section, we recall the leftover hash lemma [ILL89], focusing on the variant for matrix multiplication.

Lemma 13 ([ABB10a, Lemma 13])**.** *Let $q > 2$ be prime, $m \geq 2n\lceil \log_2 q \rceil$, and $k = \mathsf{poly}(m)$. Then*

$$(\mathbf{A}, \mathbf{AR}) \approx_s (\mathbf{A}, \mathbf{U}),$$

with statistical distance $\mathsf{negl}(m)$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{R} \leftarrow \{0,1\}^{m \times k}$, and $\mathbf{U} \leftarrow \mathbb{Z}_q^{n \times k}$ are independent.

3.7 Indistinguishability Obfuscation

Definition 2 (iO and null-iO [GGH+13, WZ17])**.** An indistinguishability obfuscation iO is a PPT algorithm that maps a circuit C to a polynomial-size circuit $\mathsf{iO}(C)$ computing the same function. For every pair of polynomial-size circuit sequences $\{C_{0,\lambda}\}_\lambda$ and $\{C_{1,\lambda}\}_\lambda$ such that $C_{0,\lambda}$ and $C_{1,\lambda}$ have equal size and compute the same function for every λ, the following ensembles are computationally indistinguishable to all PPT adversaries:

$$\{\mathsf{iO}(C_{0,\lambda})\}_\lambda \approx \{\mathsf{iO}(C_{1,\lambda})\}_\lambda.$$

Null-iO requires this security condition only when both circuits compute the identically zero function.

We pad circuits to a common public size before applying this definition.

4 Evasive LWE

4.1 Evasive LWE Sampler

We formalize the sampler interface and the independent-sampler restriction described in Section 2.1.

Definition 3 (Evasive LWE Sampler)**.** An evasive LWE sampler $\mathsf{Samp}(1^\lambda; r)$ is an efficient algorithm that, on input the security parameter 1^λ and randomness r, outputs $(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux})$, where

$$\mathsf{param} = (1^\lambda, 1^n, 1^m, q, 1^t, 1^\ell, \sigma_{\mathrm{pre}}, \sigma_{\mathrm{post}}, \rho).$$

Here λ is the security parameter and q is the modulus. The dimensions specify the source matrix $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$, secret matrix $\mathbf{S} \in \mathbb{Z}_q^{t \times n}$, and target matrix $\mathbf{P} \in \mathbb{Z}_q^{n \times \ell}$. The parameters σ_{pre} and σ_{post} are the Gaussian error widths in the precondition and postcondition, respectively, and ρ is the Gaussian preimage sampling width. The auxiliary information is some arbitrary string $\mathsf{aux} \in \{0,1\}^*$.

The parameters must satisfy

$$\sigma_{\mathrm{post}} \geq \sigma_{\mathrm{pre}} > 0, \qquad m \geq m_0(n, q), \qquad \rho \geq \rho_0(n, m),$$

with the trapdoor sampling thresholds from Lemma 1.

We next define the independent-sampler restriction.

Definition 4 (Independent Sampler)**.** An evasive LWE sampler $\mathsf{Samp}(1^\lambda; r)$ is called an *independent sampler* if it can be described with three algorithms $\mathsf{Samp}_{\mathrm{param}}, \mathsf{Samp}_{\mathrm{S}}, \mathsf{Samp}_{\mathrm{P}}$ such that:

- $\mathsf{Samp}_{\mathrm{param}}(1^\lambda; r_{\mathrm{Samp}}) \to \mathsf{param}$.
- $\mathsf{Samp}_{\mathrm{S}}(\mathsf{param}; r_{\mathrm{S}}) \to (\mathbf{S}, \mathsf{aux}_{\mathrm{S}}) \in \mathbb{Z}_q^{t \times n} \times \{0,1\}^*$.
- $\mathsf{Samp}_{\mathrm{P}}(\mathsf{param}; r_{\mathrm{P}}) \to (\mathbf{P}, \mathsf{aux}_{\mathrm{P}}) \in \mathbb{Z}_q^{n \times \ell} \times \{0,1\}^*$.
- The auxiliary information $\mathsf{aux} = (r_{\mathrm{Samp}}, \mathsf{aux}_{\mathrm{S}}, \mathsf{aux}_{\mathrm{P}})$.

The random tapes $r_{\mathrm{Samp}}, r_{\mathrm{S}}, r_{\mathrm{P}}$ are assumed to be independent and uniformly distributed.

4.2 Evasive LWE Formulations

Following [BÜW24], we define binding and hiding evasive LWE. We refer the reader to Section 2.1 for the intuition behind these definitions.

The binding formulation gives both the source and target matrices to the adversary.

Definition 5 (Binding evasive LWE)**.** For any evasive LWE sampler Samp and any adversary $\mathcal{A}$, the binding precondition advantage $\mathsf{Adv}^{\text{bind-pre}}_{\mathsf{Samp},\mathcal{A}}(\lambda)$ is defined by the following sampling procedure:

- Sample $(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda)$.
- Sample $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$.
- Sample LWE errors $\mathbf{E}_{\mathrm{B}} \leftarrow \mathcal{D}^{t \times m}_{\sigma_{\mathrm{pre}}}$, $\mathbf{E}_{\mathrm{P}} \leftarrow \mathcal{D}^{t \times \ell}_{\sigma_{\mathrm{pre}}}$.
- Compute LWE samples $\mathbf{C}_{\mathrm{B}} = \mathbf{SB} + \mathbf{E}_{\mathrm{B}}$, $\mathbf{C}_{\mathrm{P}} = \mathbf{SP} + \mathbf{E}_{\mathrm{P}}$.
- Generate random samples $\mathbf{U}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{t \times m}$, $\mathbf{U}_{\mathrm{P}} \leftarrow \mathbb{Z}_q^{t \times \ell}$.

The advantage is defined as

$$\mathsf{Adv}^{\text{bind-pre}}_{\mathsf{Samp},\mathcal{A}}(\lambda) = \left|\Pr\left[\mathcal{A}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\mathrm{B}}, \mathbf{C}_{\mathrm{P}}, \mathsf{aux}) = 1\right] - \Pr\left[\mathcal{A}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{U}_{\mathrm{B}}, \mathbf{U}_{\mathrm{P}}, \mathsf{aux}) = 1\right]\right|.$$

The binding postcondition advantage $\mathsf{Adv}^{\text{bind-post}}_{\mathsf{Samp},\mathcal{A}}(\lambda)$ is defined by the following sampling procedure:

- Sample $(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda)$.
- Sample $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$.
- Sample LWE errors $\mathbf{E}_{\mathrm{B}} \leftarrow \mathcal{D}^{t \times m}_{\sigma_{\mathrm{post}}}$.
- Sample preimage $\mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P})$.
- Compute LWE samples $\mathbf{C}_{\mathrm{B}} = \mathbf{SB} + \mathbf{E}_{\mathrm{B}}$.
- Generate random samples $\mathbf{U}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{t \times m}$.

The advantage is defined as

$$\mathsf{Adv}^{\text{bind-post}}_{\mathsf{Samp},\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{U}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux}) = 1]|.$$

We say that the binding evasive LWE assumption holds with respect to a sampler Samp if there exists a polynomial function $p(\lambda)$ and a negligible function negl such that for all PPT adversaries $\mathcal{A}_{\text{post}}$, there exists a PPT adversary $\mathcal{A}_{\text{pre}}$ such that

$$\mathsf{Adv}^{\text{bind-pre}}_{\mathsf{Samp},\mathcal{A}_{\text{pre}}}(\lambda) \geq \mathsf{Adv}^{\text{bind-post}}_{\mathsf{Samp},\mathcal{A}_{\text{post}}}(\lambda)/p(\lambda) - \mathsf{negl}(\lambda).$$

The hiding formulation omits the source and target matrices from the LWE views and adds a target-hiding precondition requiring bounded random perturbations of the target to be undetectable given the auxiliary information.

Definition 6 (Hiding evasive LWE)**.** For any evasive LWE sampler Samp and any adversary $\mathcal{A}$, the hiding precondition advantage $\mathsf{Adv}^{\text{hide-pre}}_{\mathsf{Samp},\mathcal{A}}(\lambda)$ is defined by the following sampling procedure:

- Sample $(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda)$.
- Sample $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$.
- Sample LWE errors $\mathbf{E}_{\mathrm{B}} \leftarrow \mathcal{D}^{t\times m}_{\sigma_{\text{pre}}}$, $\mathbf{E}_{\mathrm{P}} \leftarrow \mathcal{D}^{t\times \ell}_{\sigma_{\text{pre}}}$.
- Compute LWE samples $\mathbf{C}_{\mathrm{B}} = \mathbf{SB} + \mathbf{E}_{\mathrm{B}}$, $\mathbf{C}_{\mathrm{P}} = \mathbf{SP} + \mathbf{E}_{\mathrm{P}}$.
- Generate random samples $\mathbf{U}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{t\times m}$, $\mathbf{U}_{\mathrm{P}} \leftarrow \mathbb{Z}_q^{t\times \ell}$.

The advantage is defined as

$$\mathsf{Adv}^{\text{hide-pre}}_{\mathsf{Samp},\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}(\mathsf{param}, \mathbf{C}_{\mathrm{B}}, \mathbf{C}_{\mathrm{P}}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathsf{param}, \mathbf{U}_{\mathrm{B}}, \mathbf{U}_{\mathrm{P}}, \mathsf{aux}) = 1]|.$$

The target-hiding precondition advantage, parameterized by an integer perturbation level $B \geq 1$, is defined by the following sampling procedure:

- Sample $(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda)$.
- Sample perturbation $\mathbf{\Delta} \leftarrow \{-B, \ldots, B\}^{n\times \ell}$ with independent uniform entries.

The advantage is defined as

$$\mathsf{Adv}^{\text{hide-pre2}}_{\mathsf{Samp},\mathcal{A}}(\lambda, B) = |\Pr[\mathcal{A}(\mathsf{param}, \mathbf{P}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathsf{param}, \mathbf{P} + \mathbf{\Delta}, \mathsf{aux}) = 1]|.$$

The hiding postcondition advantage $\mathsf{Adv}^{\text{hide-post}}_{\mathsf{Samp},\mathcal{A}}(\lambda)$ is defined by the following sampling procedure:

- Sample $(\mathsf{param}, \mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda)$.
- Sample $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$.
- Sample LWE errors $\mathbf{E}_{\mathrm{B}} \leftarrow \mathcal{D}^{t\times m}_{\sigma_{\text{post}}}$.
- Sample preimage $\mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P})$.

- Compute LWE samples $\mathbf{C}_{\mathrm{B}} = \mathbf{SB} + \mathbf{E}_{\mathrm{B}}$.
- Generate random samples $\mathbf{U}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{t \times m}$.

The advantage is defined as

$$\mathsf{Adv}^{\text{hide-post}}_{\mathsf{Samp},\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}(\mathsf{param}, \mathbf{C}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux}) = 1] - \Pr[\mathcal{A}(\mathsf{param}, \mathbf{U}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux}) = 1]|.$$

We say that the hiding evasive LWE assumption holds with respect to a sampler Samp and a perturbation level $B = B(\lambda)$ if there exists a polynomial function $p(\lambda)$ and a negligible function negl such that for all PPT adversaries $\mathcal{A}_{\mathrm{post}}$, there exists a pair of PPT adversaries $\mathcal{A}_{\mathrm{pre}}$ and $\mathcal{A}_{\mathrm{pre2}}$ such that

$$\mathsf{Adv}^{\text{hide-pre}}_{\mathsf{Samp},\mathcal{A}_{\mathrm{pre}}}(\lambda) + \mathsf{Adv}^{\text{hide-pre2}}_{\mathsf{Samp},\mathcal{A}_{\mathrm{pre2}}}(\lambda, B) \geq \mathsf{Adv}^{\text{hide-post}}_{\mathsf{Samp},\mathcal{A}_{\mathrm{post}}}(\lambda)/p(\lambda) - \mathsf{negl}(\lambda).$$

We say that binding evasive LWE holds for a class of samplers if it holds for every sampler in the class. We use the analogous convention for hiding evasive LWE at a fixed perturbation level $B = B(\lambda)$.

5 Counterexample to Private-coin Binding Evasive LWE

In this section, we present a counterexample to the private-coin binding evasive LWE assumption over independent samplers. We present it in three parts: a specification of an independent sampler, an adversary $\mathcal{A}_{\mathrm{post}}$ with large postcondition advantage, and a proof that no PPT adversary $\mathcal{A}_{\mathrm{pre}}$ can achieve a non-negligible precondition advantage.

5.1 Sampler Description

The independent sampler Samp consists of the following three algorithms.

$\mathsf{Samp}_{\mathrm{param}}(1^\lambda; r_{\mathrm{Samp}})$**:** Sample a prime q from $[\lambda^{24}, 2\lambda^{24}]$, and set

$$n = 2\lambda, \qquad t = 1, \qquad m = 8m_0(n, q), \qquad \ell = 8m, \qquad \sigma_{\mathrm{pre}} = \sigma_{\mathrm{post}} = \lambda^3, \qquad \rho = m^3. \tag{2}$$

Output $\mathsf{param} = (1^\lambda, 1^n, 1^m, q, 1^t, 1^\ell, \sigma_{\mathrm{pre}}, \sigma_{\mathrm{post}}, \rho)$.

$\mathsf{Samp}_{\mathrm{S}}(\mathsf{param}; r_{\mathrm{S}})$**:** Sample independent uniform column vectors $\mathbf{s}_0, \mathbf{s}_1 \leftarrow \mathbb{Z}_q^\lambda$ and output

$$\mathbf{S} = [\mathbf{s}_0^\intercal \mid \mathbf{s}_1^\intercal] \in \mathbb{Z}_q^{1 \times n}, \qquad \mathsf{aux}_{\mathrm{S}} = \bot.$$

$\mathsf{Samp}_{\mathrm{P}}(\mathsf{param}; r_{\mathrm{P}})$**:** Sample

$$(\mathbf{A}_i, \mathsf{td}_{\mathbf{A}_i}) \leftarrow \mathsf{TrapGen}(1^\lambda, 1^{4m}, q), \qquad i \in \{0, 1\},$$

where $\mathbf{A}_i \in \mathbb{Z}_q^{\lambda \times 4m}$. Set

$$\mathbf{P} = \begin{bmatrix} \mathbf{A}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_1 \end{bmatrix} \in \mathbb{Z}_q^{n \times \ell}, \qquad \mathsf{aux}_{\mathrm{P}} = \mathsf{iO}(\mathsf{Vfy}_{0,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_0}}).$$

The circuit Vfy is specified in Figure 1. The sampler then outputs $(\mathbf{P}, \mathsf{aux}_{\mathrm{P}})$.

$\mathsf{Vfy}_{i,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_i}}(\mathbf{B}, \mathbf{C}, \mathbf{K})$: The verification circuit hardcodes the selection bit i, the evasive LWE parameters $\mathsf{param} = (1^\lambda, 1^n, 1^m, q, 1^t, 1^\ell, \sigma_{\mathrm{pre}}, \sigma_{\mathrm{post}}, \rho)$, the target matrix $\mathbf{P}$, and the trapdoor $\mathsf{td}_{\mathbf{A}_i}$.
The inputs have dimensions

$$\mathbf{B} \in \mathbb{Z}_q^{n\times m}, \qquad \mathbf{C} \in \mathbb{Z}_q^{1\times m}, \qquad \mathbf{K} = [\mathbf{K}_0 \mid \mathbf{K}_1], \quad \mathbf{K}_0, \mathbf{K}_1 \in \mathbb{Z}^{m\times 4m}.$$

Set

$$\beta_{\mathrm{E}} = \lambda\sigma_{\mathrm{post}}, \qquad \beta_{\mathrm{K}} = \lambda\rho.$$

The circuit performs the following checks. It outputs 0 for rejection whenever a check fails, and 1 for acceptance if all checks pass.

1. Check the input format and dimensions. Check if

$$\mathbf{BK} = \mathbf{P} \pmod q, \|\mathbf{K}\|_\infty \le \beta_{\mathrm{K}}, \qquad \mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_0) = \mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_1) = m.$$

2. Run batch LWE decoding $\mathsf{Decode}(\mathbf{A}_i, \mathsf{td}_{\mathbf{A}_i}, \mathbf{CK}_i \bmod q)$ from Lemma 1 to decode $\mathbf{CK}_i = \mathbf{S}_i\mathbf{A}_i + \mathbf{N}_i$. Reject if decoding fails.

3. Over $\mathbb{Q}$, compute the right inverse of $\mathbf{K}_i$ by

$$\mathbf{J}_i = \mathbf{K}_i^\intercal(\mathbf{K}_i\mathbf{K}_i^\intercal)^{-1} \in \mathbb{Q}^{4m\times m}$$

   and solve $\widehat{\mathbf{E}} = \mathbf{N}_i\mathbf{J}_i$. Check if

$$\widehat{\mathbf{E}} \in \mathbb{Z}^{1\times m}, \qquad \left\|\widehat{\mathbf{E}}\right\|_\infty \le \beta_{\mathrm{E}}.$$

4. Accept if and only if there exists $\widehat{\mathbf{S}} \in \mathbb{Z}_q^{1\times n}$ such that

$$\mathbf{C} = \widehat{\mathbf{S}}\mathbf{B} + \widehat{\mathbf{E}} \pmod q.$$

   Note that since $\widehat{\mathbf{E}}$ is fixed, the linear system is efficiently solvable.

Figure 1: Verification circuit for the private-coin binding counterexample.

5.2 Postcondition Adversary

Next, we propose an adversary $\mathcal{A}_{\text{post}}$ with binding postcondition advantage $1 - \mathsf{negl}(\lambda)$.

The adversary $\mathcal{A}_{\text{post}}$ receives $(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\text{B}}, \mathbf{K}, \mathsf{aux})$ in the binding postcondition experiment of Definition 5. It extracts the obfuscated circuit $\widetilde{\mathsf{Vfy}} = \mathsf{aux}_{\text{P}}$ from aux and outputs

$$\mathcal{A}_{\text{post}}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\text{B}}, \mathbf{K}, \mathsf{aux}) = \widetilde{\mathsf{Vfy}}(\mathbf{B}, \mathbf{C}_{\text{B}}, \mathbf{K}).$$

We show that $\mathcal{A}_{\text{post}}$ accepts a real LWE sample $\mathbf{C}_{\text{B}} = \mathbf{SB} + \mathbf{E}_{\text{B}}$ with probability at least $1 - \mathsf{negl}(\lambda)$, and accepts an independent uniform sample $\mathbf{U}_{\text{B}} \leftarrow \mathbb{Z}_q^{1\times m}$ with negligible probability.

Theorem 3 (Binding postcondition advantage)**.** *Consider the sampler defined above. Use $\beta_{\text{E}}, \beta_{\text{K}}$ from the obfuscated circuit in Figure 1. Suppose that, for all sufficiently large λ,*

$$m \geq m_0(n, q), \qquad \rho \geq \rho_0(n, m), \qquad q \geq 4\beta_{\text{E}} + 2, \qquad m\beta_{\text{E}}\beta_{\text{K}} < \frac{q}{32m}.$$

Then

$$\mathsf{Adv}^{\text{bind-post}}_{\mathsf{Samp},\mathcal{A}_{\text{post}}}(\lambda) \geq 1 - \mathsf{negl}(\lambda).$$

We prove Theorem 3 by analyzing the two branches of the postcondition experiment separately.

Lemma 14 (LWE branch)**.** *Under the parameter setting of Theorem 3,*

$$\Pr[\mathcal{A}_{\text{post}}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\text{B}} = \mathbf{SB} + \mathbf{E}_{\text{B}}, \mathbf{K}, \mathsf{aux}) = 1] \geq 1 - \mathsf{negl}(\lambda).$$

Proof. It suffices to analyze the probability that $\mathsf{Vfy}(\mathbf{B}, \mathbf{C}_{\text{B}}, \mathbf{K}) = 1$ when $\mathbf{C}_{\text{B}}$ is an LWE sample. We argue that each check in the verification circuit passes with overwhelming probability.

For Step 1, correctness of preimage sampling (Lemma 1) ensures that $\mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P})$ satisfies $\mathbf{BK} = \mathbf{P}$, and the preimage tail bound (Corollary 1), with $\ell = 8m$, gives

$$\|\mathbf{K}\|_\infty \leq \rho \log n \leq \lambda\rho = \beta_{\text{K}}$$

except with probability $\mathsf{negl}(n) \leq \mathsf{negl}(\lambda)$. Furthermore, by the full-rank bound over $\mathbb{Q}$ (Corollary 4), applied to the target blocks

$$\mathbf{P}_0 = \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{0} \end{bmatrix}, \qquad \mathbf{P}_1 = \begin{bmatrix} \mathbf{0} \\ \mathbf{A}_1 \end{bmatrix} \in \mathbb{Z}_q^{n\times 4m},$$

the preimages $\mathbf{K}_i \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P}_i)$ satisfy

$$\Pr[\mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_0) < m \ \vee \ \mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_1) < m] \leq \mathsf{negl}(\lambda).$$

Thus all checks in Step 1 pass except with negligible probability.

Next, the Gaussian tail bound (Lemma 4) gives

$$\|\mathbf{E}_{\text{B}}\|_\infty \leq \sigma_{\text{post}} \log n \leq \lambda\sigma_{\text{post}} = \beta_{\text{E}}$$

except with negligible probability. Therefore, whenever the previous checks pass, for each $i \in \{0, 1\}$,

$$\mathbf{C}_{\text{B}}\mathbf{K}_i = \mathbf{s}_i^{\mathsf{T}}\mathbf{A}_i + \mathbf{E}_{\text{B}}\mathbf{K}_i \pmod q, \qquad \|\mathbf{E}_{\text{B}}\mathbf{K}_i\|_\infty \leq m\beta_{\text{E}}\beta_{\text{K}} < \frac{q}{32m}$$

holds with overwhelming probability. Batch LWE decoding (Lemma 1) thus returns $(\mathbf{s}_i^\top, \mathbf{E}_\mathrm{B}\mathbf{K}_i)$. The recovered integer error is then $\widehat{\mathbf{E}} = \mathbf{E}_\mathrm{B}\mathbf{K}_i\mathbf{J}_i = \mathbf{E}_\mathrm{B}$, which, as elaborated above, has norm bounded by β_E with overwhelming probability.

Finally, the final check succeeds whenever $\mathbf{E}_\mathrm{B}$ is correctly recovered.

Combining all the above, we conclude that $\mathsf{Vfy}(\mathbf{B}, \mathbf{C}_\mathrm{B}, \mathbf{K}) = 1$ with overwhelming probability in the LWE branch. □

Lemma 15 (Uniform branch)**.** *Under the parameter setting of Theorem 3,*

$$\Pr[\mathcal{A}_{\mathrm{post}}(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{U} \leftarrow \mathbb{Z}_q^{1\times m}, \mathbf{K}, \mathsf{aux}) = 1] \le \mathsf{negl}(\lambda).$$

Proof. Again, it suffices to analyze the probability that $\mathsf{Vfy}(\mathbf{B}, \mathbf{U}, \mathbf{K}) = 1$ when $\mathbf{U}$ is a random $\mathbb{Z}_q$ matrix.

First, observe that, for Vfy to output 1, the input must pass the final check, which requires the existence of some $\widehat{\mathbf{S}}$ and $\widehat{\mathbf{E}}$ such that

$$\mathbf{U} = \widehat{\mathbf{S}}\mathbf{B} + \widehat{\mathbf{E}}, \quad \left\|\widehat{\mathbf{E}}\right\|_\infty \le \beta_\mathrm{E}$$

There are at most $q^n(2\lfloor\beta_\mathrm{E}\rfloor+1)^m$ such pairs. Therefore, fixing any $\mathbf{B}$, there exist at most $q^n(2\lfloor\beta_\mathrm{E}\rfloor+1)^m$ possible accepting values of $\mathbf{U}$ among all q^m possibilities.

Since $\mathbf{U}$ is sampled independently, the probability that Vfy outputs one is bounded by

$$\frac{q^n \cdot (2\lfloor\beta_\mathrm{E}\rfloor + 1)^m}{q^m} \le q^n 2^{-m} \le 2^{-5m/6} \le \mathsf{negl}(\lambda),$$

where we use $m \ge m_0(n, q) = 6n\lceil\log_2 q\rceil$. This completes the proof. □

Proof of Theorem 3. The theorem follows immediately by combining Lemmas 14 and 15.

The parameters in Equation (2) satisfy the theorem's conditions. Indeed, $m = 96\lambda\lceil\log_2 q\rceil = O(\lambda\log\lambda)$, $\rho = m^3 \ge 32m$, $\beta_\mathrm{E} = \lambda^4$, and $\beta_\mathrm{K} = \lambda m^3$. Since $q \ge \lambda^{24}$, we have $q \ge 4\beta_\mathrm{E} + 2$ and $32\lambda^5 m^5 < q$, where the latter is equivalent to $m\beta_\mathrm{E}\beta_\mathrm{K} < q/(32m)$. □

5.3 Pseudorandomness of the Precondition Distribution

We first show that the two verification circuits compute the same function on every input, allowing an iO switch between circuits hardcoding one of the two trapdoors.

Lemma 16 (Functional equivalence of the verification circuits)**.** *Suppose $m\beta_\mathrm{E}\beta_\mathrm{K} < q/(32m)$. For every pair $(\mathbf{A}_i, \mathsf{td}_{\mathbf{A}_i})$, $i \in \{0, 1\}$, generated by the sampler given in Section 5.1, the circuits $\mathsf{Vfy}_{0,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_0}}$ and $\mathsf{Vfy}_{1,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_1}}$ in Figure 1 are functionally equivalent.*

Proof. It suffices to show that for any input, both circuits output the same value, or equivalently, that for every $b \in \{0, 1\}$, every tuple $(\mathbf{B}, \mathbf{C}, \mathbf{K})$ accepted by Vfy_b is also accepted by Vfy_{1-b}.

Suppose Vfy_b accepts $(\mathbf{B}, \mathbf{C}, \mathbf{K})$, and let $\widehat{\mathbf{E}}$ be the error computed in Step 3. By Step 4, there is a witness $\widehat{\mathbf{S}} = [\widehat{\mathbf{s}}_0^\top \mid \widehat{\mathbf{s}}_1^\top]$ such that

$$\mathbf{C} = \widehat{\mathbf{S}}\mathbf{B} + \widehat{\mathbf{E}} \pmod q, \qquad \left\|\widehat{\mathbf{E}}\right\|_\infty \le \beta_\mathrm{E}.$$

The check in Step 1 gives

$$\mathbf{BK} = \mathbf{P} \pmod q, \qquad \|\mathbf{K}\|_\infty \le \beta_{\mathrm{K}}, \qquad \mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_{1-b}) = m.$$

Consequently,

$$\mathbf{CK}_{1-b} = \widehat{\mathbf{s}}_{1-b}^{\mathsf{T}}\mathbf{A}_{1-b} + \widehat{\mathbf{E}}\mathbf{K}_{1-b} \pmod q,$$

with noise bounded by

$$\left\|\widehat{\mathbf{E}}\mathbf{K}_{1-b}\right\|_\infty \le m\beta_{\mathrm{E}}\beta_{\mathrm{K}} < \frac{q}{32m}.$$

By LWE decoding (Lemma 1), Vfy_{1-b} recovers $\mathbf{N}_{1-b} = \widehat{\mathbf{E}}\mathbf{K}_{1-b}$. The rank check ensures $\mathbf{K}_{1-b}\mathbf{J}_{1-b} = \mathbf{I}_m$, so its recovered error is

$$\mathbf{N}_{1-b}\mathbf{J}_{1-b} = \widehat{\mathbf{E}}\mathbf{K}_{1-b}\mathbf{J}_{1-b} = \widehat{\mathbf{E}}.$$

This is the same error derived from Vfy_b and passes the same norm check. Thus Vfy_{1-b} solves $\widehat{\mathbf{S}}$ just like Vfy_b and accepts as well. □

With the lemma, we can now prove the pseudorandomness of the precondition distribution. The idea is straightforward: in hybrids where Vfy_b is in view, $\mathsf{td}_{\mathbf{A}_{1-b}}$ is not in the view, and thus one can apply LWE over the public matrix $\mathbf{A}_{1-b}$.

Theorem 4 (Binding precondition)**.** *Consider the sampler with parameters from Equation (2). Assume the security of iO (Definition 2) and* $\mathsf{LWE}_{\lambda,5m,q,\sigma_{\mathrm{pre}}}$ *(Definition 1). Then, for every PPT adversary* $\mathcal{A}_{\mathrm{pre}}$*,*

$$\mathsf{Adv}^{\text{bind-pre}}_{\mathsf{Samp},\mathcal{A}_{\mathrm{pre}}}(\lambda) \le \mathsf{negl}(\lambda).$$

Proof. The parameters satisfy $m\beta_{\mathrm{E}}\beta_{\mathrm{K}} < q/(32m)$, so the functional-equivalence lemma (Lemma 16) applies. We argue that the precondition advantage is negligible via the following sequence of hybrids.

$\boxed{\mathcal{H}_0}$ This is the LHS distribution of the precondition in Definition 5, consisting of

$$(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\mathrm{B}}, \mathbf{C}_{\mathrm{P}}, \mathsf{aux}),$$

where

$$\begin{aligned}
\mathbf{C}_{\mathrm{B}} &= \mathbf{SB} + \mathbf{E}_{\mathrm{B}} = \mathbf{s}_0^{\mathsf{T}}\mathbf{B}_0 + \mathbf{s}_1^{\mathsf{T}}\mathbf{B}_1 + \mathbf{e}_{\mathrm{B}}^{\mathsf{T}},\\
\mathbf{C}_{\mathrm{P}} &= \mathbf{SP} + \mathbf{E}_{\mathrm{P}} = (\mathbf{C}_{\mathrm{P},0} \mid \mathbf{C}_{\mathrm{P},1}),\\
\mathbf{C}_{\mathrm{P},i} &= \mathbf{s}_i^{\mathsf{T}}\mathbf{A}_i + \mathbf{e}_{\mathrm{P},i}^{\mathsf{T}}, \qquad i \in \{0,1\},\\
\mathsf{aux}_{\mathrm{P}} &= \mathsf{iO}(\mathsf{Vfy}_{0,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_0}}), \qquad \mathsf{aux} = (r_{\mathsf{Samp}}, \perp, \mathsf{aux}_{\mathrm{P}}).
\end{aligned}$$

Here $\mathbf{B} = \begin{bmatrix}\mathbf{B}_0\\ \mathbf{B}_1\end{bmatrix}$ with $\mathbf{B}_i \in \mathbb{Z}_q^{\lambda\times m}$ is sampled via $\mathsf{TrapGen}$; $\mathbf{P} = \begin{bmatrix}\mathbf{A}_0 & \mathbf{0}\\ \mathbf{0} & \mathbf{A}_1\end{bmatrix}$, and $\mathbf{A}_0, \mathbf{A}_1$ are generated via $\mathsf{TrapGen}$.

$\boxed{\mathcal{H}_1}$ Same as Hybrid $\mathcal{H}_0$, except that $\mathbf{B}$ and $\mathbf{A}_1$ are sampled uniformly at random.

$\boxed{\mathcal{H}_2}$ Same as Hybrid $\mathcal{H}_1$, except that this hybrid samples $\mathbf{C}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{1\times m}$ and $\mathbf{C}_{\mathrm{P},1} \leftarrow \mathbb{Z}_q^{1\times 4m}$ instead of computing them as LWE samples. The sample $\mathbf{C}_{\mathrm{P},0}$ remains as an LWE sample $\mathbf{C}_{\mathrm{P},0} = \mathbf{s}_0^{\mathsf{T}}\mathbf{A}_0 + \mathbf{e}_{\mathrm{P},0}^{\mathsf{T}}$.

$\boxed{\mathcal{H}_3}$ Same as Hybrid $\mathcal{H}_2$, but $\mathbf{A}_1$ is again sampled via $\mathsf{TrapGen}$.

$\boxed{\mathcal{H}_4}$ Same as Hybrid $\mathcal{H}_3$, except that the auxiliary component $\mathsf{aux}_{\mathrm{P}}$ is replaced with

$$\mathsf{aux}_{\mathrm{P}} = \mathsf{iO}(\mathsf{Vfy}_{1,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_1}}).$$

The two verification circuits are functionally equivalent by Lemma 16.

$\boxed{\mathcal{H}_5}$ Same as Hybrid $\mathcal{H}_4$, but now $\mathbf{A}_0$ is sampled uniformly at random.

$\boxed{\mathcal{H}_6}$ Same as Hybrid $\mathcal{H}_5$, but now $\mathbf{C}_{\mathrm{P},0}$ is also sampled uniformly at random.

$\boxed{\mathcal{H}_7}$ Same as Hybrid $\mathcal{H}_6$, but $\mathbf{B}$ and $\mathbf{A}_0$ are again sampled via $\mathsf{TrapGen}$.

$\boxed{\mathcal{H}_8}$ Same as Hybrid $\mathcal{H}_7$, but the auxiliary component $\mathsf{aux}_{\mathrm{P}}$ is replaced back with $\mathsf{iO}(\mathsf{Vfy}_{0,\mathsf{param},\mathbf{P},\mathsf{td}_{\mathbf{A}_0}})$. This is the RHS distribution of the precondition in Definition 5.

We now state the indistinguishability of each adjacent pair of hybrids.

Claim 1. *The distributions in Hybrids $\mathcal{H}_0$ and $\mathcal{H}_1$ are statistically indistinguishable.*

Proof. This follows from the public-matrix guarantee of $\mathsf{TrapGen}$ (Lemma 1): matrices generated with a trapdoor are $\mathsf{negl}(n)$-close to uniform. Neither $\mathsf{td}_{\mathbf{B}}$ nor $\mathsf{td}_{\mathbf{A}_1}$ is included in the view in Hybrid $\mathcal{H}_0$. □

Claim 2. *Assuming $\mathsf{LWE}_{\lambda,5m,q,\sigma_{\mathrm{pre}}}$, the distributions in Hybrids $\mathcal{H}_1$ and $\mathcal{H}_2$ are computationally indistinguishable.*

Proof. The indistinguishability follows from applying LWE to the public matrix $[\mathbf{B}_1 \mid \mathbf{A}_1]$ with secret $\mathbf{s}_1$. Note that $[\mathbf{B}_1 \mid \mathbf{A}_1]$ is generated uniformly at random in these hybrids. □

Claim 3. *The distributions in Hybrids $\mathcal{H}_2$ and $\mathcal{H}_3$ are statistically indistinguishable.*

Proof. Same as Claim 1. □

Claim 4. *Assuming the security of iO, the distributions in Hybrids $\mathcal{H}_3$ and $\mathcal{H}_4$ are computationally indistinguishable.*

Proof. Since the two verification circuits in Hybrid $\mathcal{H}_3$ and Hybrid $\mathcal{H}_4$ are functionally equivalent (Lemma 16), the security of iO implies the indistinguishability of the two hybrids. □

Claim 5. *The distributions in Hybrids $\mathcal{H}_4$ and $\mathcal{H}_5$ are statistically indistinguishable.*

Proof. Same as Claim 1. □

Claim 6. *Assuming $\mathsf{LWE}_{\lambda,5m,q,\sigma_{\mathrm{pre}}}$, the distributions in Hybrids $\mathcal{H}_5$ and $\mathcal{H}_6$ are computationally indistinguishable.*

Proof. The indistinguishability follows from applying LWE to the public matrix $\mathbf{A}_0$ with secret $\mathbf{s}_0$. Note that $\mathbf{A}_0$ is generated uniformly at random in this hybrid. □

Claim 7. *The distributions in Hybrids $\mathcal{H}_6$ and $\mathcal{H}_7$ are statistically indistinguishable.*

Proof. Same as Claim 1. □

Claim 8. *Assuming the security of iO, the distributions in Hybrids $\mathcal{H}_7$ and $\mathcal{H}_8$ are computationally indistinguishable.*

Proof. Same as Claim 4. □

Combining Claims 1 to 8, the distributions in Hybrids $\mathcal{H}_0$ and $\mathcal{H}_8$ are computationally indistinguishable, proving the theorem. □

6 Counterexample to Private-coin Hiding Evasive LWE

In this section, we present a counterexample to the private-coin hiding evasive LWE assumption over independent samplers at perturbation level B. The parameters of our counterexample depend on B, but the construction supports every $B = 2^{\mathsf{poly}(\lambda)}$. Throughout this section, we fix $B = B(\lambda) \geq 1$ and write $\kappa = \log_2 B$, where κ is polynomially bounded in λ.

In the following, we again present the counterexample in three parts: a specification of a counterexample sampler, an adversary $\mathcal{A}_{\mathrm{post}}$ with large distinguishing advantage against the postcondition, and a proof that both preconditions with respect to the counterexample sampler hold against all PPT adversaries.

6.1 Sampler Description

The independent sampler Samp consists of the following three algorithms.

$\mathsf{Samp}_{\mathrm{param}}(1^\lambda; r_{\mathrm{Samp}})$**:** Sample a prime q from $[2^{3\lambda+\kappa}, 2^{3\lambda+\kappa+1})$ and set

$$\begin{gathered} t = 1, \qquad n = \lceil \lambda + \kappa \rceil^2, \qquad m = 8m_0(n, q), \\ L = m^3, \qquad \ell = L + 2m + n, \qquad d = 2(m+n)\lceil \log_2 q \rceil, \\ \eta = 2^\lambda B = 2^{\lambda+\kappa}, \qquad a = m^8 \sigma_{\mathrm{s}} \eta, \qquad \rho = \rho_0(n, m) = 32m, \\ \sigma_{\mathrm{s}} = \mathsf{poly}(n) > n, \qquad \sigma_{\mathrm{M}} = \mathsf{poly}(n) > n, \qquad \sigma_{\mathrm{pre}} = \sigma_{\mathrm{post}} = 2^\lambda a. \end{gathered} \tag{3}$$

Here σ_{s} and σ_{M} are sufficiently large integer-valued polynomials in n, chosen for the LWE hardness regime (Section 3.5). Output $\mathsf{param} = (1^\lambda, 1^n, 1^m, q, 1^t, 1^\ell, \sigma_{\mathrm{pre}}, \sigma_{\mathrm{post}}, \rho)$.

$\mathsf{Samp}_{\mathrm{S}}(\mathsf{param}; r_{\mathrm{S}})$**:** Sample a short column vector $\mathbf{s} \leftarrow \mathcal{D}_{\sigma_{\mathrm{s}}}^n$ and output

$$\mathbf{S} = \mathbf{s}^\intercal \in \mathbb{Z}_q^{1\times n}, \qquad \mathsf{aux}_{\mathrm{S}} = \bot.$$

$\mathsf{Samp}_{\mathrm{P}}(\mathsf{param}; r_{\mathrm{P}})$**:** Sample

$$(\mathbf{A}_0, \mathsf{td}_{\mathbf{A}_0}) \leftarrow \mathsf{TrapGen}(1^n, 1^L, q), \qquad \mathbf{A}_0 \in \mathbb{Z}_q^{n\times L},$$

and sample $\mathbf{T}$ uniformly from the invertible matrices in $\mathbb{Z}_q^{n\times n}$, together with

$$\mathbf{A}_1, \mathbf{A}_3 \leftarrow \mathbb{Z}_q^{n\times m}, \qquad \mathbf{R} \leftarrow \{0,1\}^{m\times n}, \qquad \mathbf{W} \leftarrow \{0,1\}^{m\times d}, \qquad \mathbf{F}_{\mathrm{M}} \leftarrow \mathcal{D}_{\sigma_{\mathrm{M}}}^{n\times d}.$$

Sample the target errors

$$\mathbf{F}_0 \leftarrow \mathcal{D}_\eta^{n\times L}, \qquad \mathbf{F}_1, \mathbf{F}_3 \leftarrow \mathcal{D}_\eta^{n\times m}, \qquad \mathbf{F}_2 \leftarrow \mathcal{D}_\eta^{n\times n}.$$

Set

$$\mathbf{P} = [\mathbf{T}\mathbf{A}_0 + \mathbf{F}_0 \mid \mathbf{A}_1 + \mathbf{F}_1 \mid \mathbf{A}_1\mathbf{R} + a\mathbf{I}_n + \mathbf{F}_2 \mid \mathbf{A}_3 + \mathbf{F}_3] \in \mathbb{Z}_q^{n\times\ell},$$
$$\mathbf{M} = \mathbf{A}_3\mathbf{W} + \mathbf{F}_\mathrm{M} \in \mathbb{Z}_q^{n\times d}, \qquad \mathsf{aux}_\mathrm{P} = \mathsf{iO}(\mathsf{Vfy}_{\mathsf{param},\mathbf{A}_0,\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}}).$$

The circuit Vfy is specified in Figure 2. The sampler then outputs $(\mathbf{P}, \mathsf{aux}_\mathrm{P})$.

6.2 Postcondition Adversary

Next, we propose an adversary $\mathcal{A}_\mathrm{post}$ with non-negligible hiding postcondition advantage.

The adversary $\mathcal{A}_\mathrm{post}$ is straightforward: on receiving $(\mathsf{param}, \mathbf{C}_\mathrm{B}, \mathbf{K}, \mathsf{aux})$ in the hiding postcondition experiment of Definition 6, it extracts the obfuscated circuit $\widetilde{\mathsf{Vfy}} = \mathsf{aux}_\mathrm{P}$ from aux and outputs

$$\mathcal{A}_\mathrm{post}(\mathsf{param}, \mathbf{C}_\mathrm{B}, \mathbf{K}, \mathsf{aux}) = \widetilde{\mathsf{Vfy}}(\mathbf{C}_\mathrm{B}, \mathbf{K}).$$

Theorem 5 (Hiding postcondition advantage)**.** *Consider the sampler with parameters from Equation (3). Then*

$$\mathsf{Adv}^{\text{hide-post}}_{\mathsf{Samp},\mathcal{A}_\mathrm{post}}(\lambda) \geq 1 - \frac{1}{\rho - 1} - \mathsf{negl}(\lambda).$$

We prove Theorem 5 by analyzing the two branches of the postcondition experiment separately.

Lemma 17 (LWE branch)**.** *Under the parameter setting of Theorem 5, sample the LWE branch of the hiding postcondition experiment (Definition 6), with* $\mathbf{C}_\mathrm{B} = \mathbf{S}\mathbf{B} + \mathbf{E}_\mathrm{B}$*. Then*

$$\Pr[\mathcal{A}_\mathrm{post}(\mathsf{param}, \mathbf{C}_\mathrm{B}, \mathbf{K}, \mathsf{aux}) = 1] \geq 1 - \frac{1}{\rho - 1} - \mathsf{negl}(\lambda).$$

Proof. It suffices to analyze the probability that $\mathsf{Vfy}(\mathbf{C}_\mathrm{B}, \mathbf{K}) = 1$ when $\mathbf{C}_\mathrm{B}$ is an LWE sample. We check each step of the verification circuit.

The preimage tail bound (Corollary 1) and the Gaussian tail bounds (Lemma 4) guarantee that, except with negligible probability, it holds that

$$\|\mathbf{K}\|_\infty \leq \rho \log n \leq \lambda\rho, \qquad \|\mathbf{S}\|_\infty \leq \lambda\sigma_\mathrm{s}, \qquad \|\mathbf{E}_\mathrm{B}\|_\infty \leq \lambda\sigma_\mathrm{post},$$
$$\left\|\mathbf{F}_j\right\|_\infty \leq \lambda\eta \quad (0 \leq j \leq 3), \qquad \|\mathbf{F}_\mathrm{M}\|_\infty \leq \lambda\sigma_\mathrm{M}.$$

The nonzero Gaussian-sample bound (Corollary 2) and the tail bounds above give $\Pr[\mathbf{S} = \mathbf{0}] \leq \mathsf{negl}(\lambda)$ over $\mathbb{Z}_q$, since $\sigma_\mathrm{s} > n$ and $\lambda\sigma_\mathrm{s} < q$.

For Step 1, the public-matrix guarantee of TrapGen (Lemma 1) implies that $\mathbf{P}_0 = \mathbf{T}\mathbf{A}_0 + \mathbf{F}_0$ is $\mathsf{negl}(n)$-close to uniform and independent of $\mathbf{B}$. By the preimage rank bound (Corollary 3),

$$\Pr[\mathsf{rank}_{\mathbb{Z}_q}(\mathbf{K}_0) < m] \leq \frac{1}{\rho - 1} + \mathsf{negl}(\lambda).$$

Since $\mathbf{P}_0$ is nearly uniform, the right-inverse norm bound for uniform targets (Corollary 5), applied with $L = m^3$, also gives $\|\mathbf{J}\|_\infty \leq 1/\rho$ except with negligible probability.

Condition on all these bounds. In Step 2 the product to be decoded is

$$\mathbf{C}_\mathrm{B}\mathbf{K}_0 = \mathbf{S}\mathbf{T}\mathbf{A}_0 + \mathbf{N}_0, \qquad \mathbf{N}_0 = \mathbf{E}_\mathrm{B}\mathbf{K}_0 + \mathbf{S}\mathbf{F}_0,$$

and its integer error satisfies

$$\|\mathbf{N}_0\|_\infty \leq m\,\|\mathbf{E}_\mathrm{B}\|_\infty\,\|\mathbf{K}_0\|_\infty + n\,\|\mathbf{S}\|_\infty\,\|\mathbf{F}_0\|_\infty \leq \lambda^2 m\rho\sigma_\mathrm{post} + \lambda^2 n\eta\sigma_\mathrm{s} \leq \beta_\mathrm{N} < \frac{q}{8L}.$$

$\mathsf{Vfy}_{\mathsf{param},\mathbf{A}_0,\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}}(\mathbf{C},\mathbf{K})$: The verification circuit hardcodes the parameters param, the matrices $\mathbf{A}_0, \mathbf{R}, \mathbf{W}, \mathbf{M}$, and the trapdoor $\mathsf{td}_{\mathbf{A}_0}$.
The inputs have dimensions

$$\mathbf{C} \in \mathbb{Z}_q^{t\times m}, \qquad \mathbf{K} = [\mathbf{K}_0 \mid \mathbf{K}_1 \mid \mathbf{K}_2 \mid \mathbf{K}_3] \in \mathbb{Z}^{m\times \ell},$$

where the four blocks have widths L, m, n, m, respectively. Compute the remaining parameters using Equation (3), and set

$$\beta_{\mathrm{N}} = 2\lambda^2 m\rho\sigma_{\mathrm{post}}, \qquad \beta_{\mathrm{M}} = 2m^7\eta\sigma_{\mathrm{s}}.$$

The circuit outputs 0 for rejection whenever a check fails, and 1 for acceptance if all checks pass.

1. Check the input format and dimensions. Check that $\mathsf{rank}_{\mathbb{Z}_q}(\mathbf{K}_0) = m$. Over $\mathbb{Q}$, compute the right inverse of $\mathbf{K}_0$ from Lemma 9:

$$\mathbf{J} = \mathbf{K}_0^{\mathsf{T}}(\mathbf{K}_0\mathbf{K}_0^{\mathsf{T}})^{-1} \in \mathbb{Q}^{L\times m}.$$

2. Run the LWE decoding algorithm from Lemma 1:

$$(\widehat{\mathbf{S}}_0, \mathbf{N}_0) = \mathsf{Decode}(\mathbf{A}_0, \mathsf{td}_{\mathbf{A}_0}, \mathbf{C}\mathbf{K}_0) \in \mathbb{Z}_q^{t\times n} \times \mathbb{Z}^{t\times L}.$$

   Reject if decoding fails or $\|\mathbf{N}_0\|_\infty > \beta_{\mathrm{N}}$.

   Set the approximate error $\widehat{\mathbf{E}} = \lfloor \mathbf{N}_0\mathbf{J} \rceil \in \mathbb{Z}^{t\times m}$ and the pruned sample $\widetilde{\mathbf{C}} = \mathbf{C} - \widehat{\mathbf{E}} \in \mathbb{Z}_q^{t\times m}$.

3. Compute $\mathbf{Y}_1 = \widetilde{\mathbf{C}}\mathbf{K}_1$ and $\mathbf{Y}_2 = \widetilde{\mathbf{C}}\mathbf{K}_2$ over $\mathbb{Z}_q$, and solve for the secret

$$\widehat{\mathbf{S}} = \left\lfloor \frac{\mathbf{Y}_2 - \mathbf{Y}_1\mathbf{R}}{a} \right\rceil, \qquad \widehat{\mathbf{S}} \in \mathbb{Z}^{t\times n}.$$

   Reject if $\widehat{\mathbf{S}} = \mathbf{0}$.

4. Set $\mathbf{Z} = \widetilde{\mathbf{C}}\mathbf{K}_3 \in \mathbb{Z}_q^{t\times m}$, and accept if and only if

$$\left\| \mathbf{Z}\mathbf{W} - \widehat{\mathbf{S}}\mathbf{M} \right\|_\infty \le \beta_{\mathrm{M}}.$$

Figure 2: Verification circuit for the private-coin hiding counterexample.

LWE decoding (Lemma 1) thus returns $(\mathbf{ST}, \mathbf{N}_0)$.

Set $\mathbf{D} = \mathbf{F}_0\mathbf{J} \in \mathbb{Q}^{n\times m}$ and $\mathbf{\Delta} = \lfloor \mathbf{SD} \rceil - \mathbf{SD}$. Since $\mathbf{K}_0\mathbf{J} = \mathbf{I}_m$, the recovered error is

$$\widehat{\mathbf{E}} = \mathbf{E}_\mathrm{B} + \lfloor \mathbf{SD} \rceil = \mathbf{E}_\mathrm{B} + \mathbf{SD} + \mathbf{\Delta}, \qquad \|\mathbf{\Delta}\|_\infty \le \tfrac{1}{2}.$$

For $j \in \{1, 2, 3\}$, the pruned sample $\widetilde{\mathbf{C}}\mathbf{K}_j$ satisfies

$$\widetilde{\mathbf{C}}\mathbf{K}_j = \mathbf{S}(\mathbf{P}_j - \mathbf{F}_j) + \underbrace{\mathbf{S}(\mathbf{F}_j - \mathbf{D}\mathbf{K}_j) - \mathbf{\Delta}\mathbf{K}_j}_{\widetilde{\mathbf{N}}_j}.$$

Here $\mathbf{P}_j - \mathbf{F}_j$ is the target block without its perturbation error. Using $\|\mathbf{D}\|_\infty \le L\lambda\eta/\rho$, we obtain

$$\left\|\widetilde{\mathbf{N}}_j\right\|_\infty \le \lambda^2 n\eta\sigma_\mathrm{s} + \lambda^3 nmL\eta\sigma_\mathrm{s} + \lambda m\rho/2 \le m^6\eta\sigma_\mathrm{s},$$

where the last inequality uses $L = m^3$, $\rho = 32m$, $\lambda^2 \le n$, and $m = 48n\lceil \log_2 q \rceil \ge 144\lambda n$.

With the bounds above, the products for Step 3 can be expressed as

$$\mathbf{Y}_1 = \widetilde{\mathbf{C}}\mathbf{K}_1 = \mathbf{S}\mathbf{A}_1 + \widetilde{\mathbf{N}}_1, \qquad \mathbf{Y}_2 = \widetilde{\mathbf{C}}\mathbf{K}_2 = \mathbf{S}\mathbf{A}_1\mathbf{R} + a\mathbf{S} + \widetilde{\mathbf{N}}_2,$$

so $\mathbf{Y}_2 - \mathbf{Y}_1\mathbf{R} = a\mathbf{S} + \widetilde{\mathbf{N}}_2 - \widetilde{\mathbf{N}}_1\mathbf{R}$ modulo q. Since $\mathbf{R}$ is binary and $a = m^8\eta\sigma_\mathrm{s}$,

$$\left\|\widetilde{\mathbf{N}}_2 - \widetilde{\mathbf{N}}_1\mathbf{R}\right\|_\infty \le (m+1)m^6\eta\sigma_\mathrm{s} \le 2m^7\eta\sigma_\mathrm{s} < \frac{a}{2}.$$

Moreover, the unreduced integer expression satisfies

$$\left\|a\mathbf{S} + \widetilde{\mathbf{N}}_2 - \widetilde{\mathbf{N}}_1\mathbf{R}\right\|_\infty \le a\lambda\sigma_\mathrm{s} + 2m^7\eta\sigma_\mathrm{s} < q/2,$$

since $a = \mathsf{poly}(\lambda, \kappa)\eta$ and $q/\eta \ge 2^{2\lambda}$. Thus taking centered representatives and rounding recovers $\widehat{\mathbf{S}} = \mathbf{S}$, which passes the nonzero check.

Finally, for Step 4, we have $\mathbf{Z} = \widetilde{\mathbf{C}}\mathbf{K}_3 = \mathbf{S}\mathbf{A}_3 + \widetilde{\mathbf{N}}_3$, and hence

$$\mathbf{ZW} - \widehat{\mathbf{S}}\mathbf{M} = \widetilde{\mathbf{N}}_3\mathbf{W} - \mathbf{S}\mathbf{F}_\mathrm{M}.$$

Since $\mathbf{W}$ is binary, the bound above on $\left\|\widetilde{\mathbf{N}}_3\right\|_\infty$ gives

$$\left\|\widetilde{\mathbf{N}}_3\mathbf{W} - \mathbf{S}\mathbf{F}_\mathrm{M}\right\|_\infty \le m^7\eta\sigma_\mathrm{s} + \lambda^2 n\sigma_\mathrm{s}\sigma_\mathrm{M} \le \beta_\mathrm{M}.$$

Combining the bounds above, all checks pass except with probability at most $1/(\rho - 1) + \mathsf{negl}(\lambda)$. Correctness of iO completes the proof. □

Lemma 18 (Uniform branch)**.** *Under the parameter setting of Theorem 5, sample* $(\mathsf{param}, \mathbf{K}, \mathsf{aux})$ *as in Lemma 17, and independently sample* $\mathbf{U}_\mathrm{B} \leftarrow \mathbb{Z}_q^{1\times m}$*. Then*

$$\Pr[\mathcal{A}_{\mathrm{post}}(\mathsf{param}, \mathbf{U}_\mathrm{B}, \mathbf{K}, \mathsf{aux}) = 1] \le \mathsf{negl}(\lambda).$$

Proof. It suffices to bound the probability that $\mathsf{Vfy}(\mathbf{U}_{\mathrm{B}}, \mathbf{K}) = 1$ for an independent uniform $\mathbf{U}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{1\times m}$. Fix $\mathbf{K}$ and the verification circuit. If $\mathbf{K}_0$ has deficient rank modulo q, the verification circuit rejects. Otherwise choose a set I of m columns such that $\mathbf{K}_{0,I}$ is invertible modulo q. Acceptance in Step 2 requires

$$\mathbf{U}_{\mathrm{B}}\mathbf{K}_{0,I} = \widehat{\mathbf{S}}_0\mathbf{A}_{0,I} + \mathbf{N}_{0,I}, \qquad \|\mathbf{N}_{0,I}\|_\infty \le \beta_{\mathrm{N}}.$$

Each pair $(\widehat{\mathbf{S}}_0, \mathbf{N}_{0,I})$ determines at most one $\mathbf{U}_{\mathrm{B}}$. The acceptance probability is consequently at most

$$\frac{q^n \cdot (2\lfloor\beta_{\mathrm{N}}\rfloor + 1)^m}{q^m} \le 2^{-\Omega(m\lambda)} \le \mathsf{negl}(\lambda),$$

where we use $\log_2 q = 3\lambda + \kappa + O(1)$, $\log_2(2\lfloor\beta_{\mathrm{N}}\rfloor + 1) = 2\lambda + \kappa + O(\log(\lambda + \kappa))$, and $m = 48n\lceil\log_2 q\rceil$. Averaging over the fixed values and applying correctness of iO proves the claim. □

Proof of Theorem 5. The theorem follows immediately by combining Lemmas 17 and 18. □

6.3 Proof of Preconditions

We show that the sampler satisfies both preconditions in Definition 6: target hiding under bounded perturbations and pseudorandomness of the joint LWE samples.

6.3.1 Target hiding

The target hiding precondition is direct by design. The perturbation errors $\mathbf{F}_j$ added to $\mathbf{P}_j$ flood any perturbation of norm at most B, while the auxiliary information is independent of these errors.

Theorem 6 (Target-hiding precondition)**.** *Consider the sampler with parameters from Equation (3). Then, for every adversary $\mathcal{A}_{\mathrm{pre2}}$,*

$$\mathsf{Adv}^{\text{hide-pre2}}_{\mathsf{Samp},\mathcal{A}_{\mathrm{pre2}}}(\lambda, B) \le \mathsf{negl}(\lambda).$$

Proof. Condition on the public parameter tape r_{Samp}. The current sampler defines the target matrix $\mathbf{P}$ by

$$\mathbf{P} = \mathbf{P}^* + \mathbf{F}, \qquad \mathbf{F} \leftarrow \mathcal{D}_\eta^{n\times\ell},$$

where the pair $(\mathbf{P}^*, \mathsf{aux}_{\mathrm{P}})$ is independent of $\mathbf{F}$.

Let $\mathbf{\Delta}$ be the independent perturbation from the target-hiding experiment (Definition 6). It satisfies $\|\mathbf{\Delta}\|_\infty \le B$. By the Gaussian flooding bound (Lemma 5), the statistical distance between $\mathbf{F}$ and $\mathbf{F} + \mathbf{\Delta}$ is at most

$$O\left(\frac{n\ell B}{\eta}\right) = O(n\ell\, 2^{-\lambda}) \le \mathsf{negl}(\lambda).$$

□

6.3.2 Pseudorandomness of the precondition distribution

We first show the core lemma for our null-iO hybrid. Once the LWE sample $\mathbf{M}$ is replaced by a uniform matrix, the circuit Vfy becomes a null circuit with overwhelming probability.

Lemma 19 (Sparsified rowspace consistency)**.** *Let $\beta \geq 1$, $d \geq 2(m+n)\lceil \log_2 q \rceil$, and $q \geq 4\beta + 2$. Fix $\mathbf{W} \in \mathbb{Z}_q^{m \times d}$ and sample $\mathbf{M} \leftarrow \mathbb{Z}_q^{n \times d}$ uniformly. The probability that there exist $\mathbf{Z} \in \mathbb{Z}_q^{t \times m}$ and $\widehat{\mathbf{S}} \in \mathbb{Z}_q^{t \times n}$ satisfying*

$$\widehat{\mathbf{S}} \neq \mathbf{0}, \qquad \left\| \mathbf{ZW} - \widehat{\mathbf{S}}\mathbf{M} \right\|_\infty \leq \beta$$

is at most $\mathsf{negl}(n)$.

Proof. Any such pair contains a nonzero row $\widehat{\mathbf{s}}^\intercal$ of $\widehat{\mathbf{S}}$ and the corresponding row $\mathbf{z}^\intercal$ of $\mathbf{Z}$ passing the norm check. For each fixed row pair $(\widehat{\mathbf{s}}, \mathbf{z})$ with $\widehat{\mathbf{s}} \neq \mathbf{0}$, the vector $\widehat{\mathbf{s}}^\intercal \mathbf{M}$ is uniform over $\mathbb{Z}_q^{1 \times d}$. Thus

$$\Pr_{\mathbf{M}}\left[\|\mathbf{z}^\intercal \mathbf{W} - \widehat{\mathbf{s}}^\intercal \mathbf{M}\|_\infty \leq \beta \right] \leq \frac{(2\beta+1)^d}{q^d}.$$

A union bound over at most q^{m+n} such row pairs bounds the probability that at least one pair passes the norm check by

$$\frac{q^{m+n} \cdot (2\beta+1)^d}{q^d} \leq q^{m+n} 2^{-d} \leq q^{-(m+n)} \leq \mathsf{negl}(n).$$

□

With this lemma, we can now prove the pseudorandomness of the precondition distribution.

Theorem 7 (Hiding precondition)**.** *Consider the sampler with parameters from Equation (3). Assume null-iO (Definition 2) and* $\mathsf{LWE}_{n,d,q,\min\{\sigma_s,\sigma_M\}}$ *(Definition 1). Then, for every PPT adversary* $\mathcal{A}_{\mathrm{pre}}$*,*

$$\mathsf{Adv}^{\text{hide-pre}}_{\mathsf{Samp},\mathcal{A}_{\mathrm{pre}}}(\lambda) \leq \mathsf{negl}(\lambda).$$

Proof. We argue that the precondition advantage is negligible via the following sequence of hybrids. Throughout, $\mathsf{aux} = (r_{\mathsf{Samp}}, \perp, \mathsf{aux}_{\mathrm{P}})$, and the public parameter is generated as in $\mathsf{Samp}_{\mathrm{param}}$.

$\boxed{\mathcal{H}_0}$ This is the LHS distribution of the precondition in Definition 6, consisting of

$$(\mathsf{param}, \mathbf{C}_{\mathrm{B}}, \mathbf{C}_{\mathrm{P}}, \mathsf{aux}),$$

where $\mathbf{C}_{\mathrm{P}} = [\mathbf{C}_{\mathrm{P},0} \mid \mathbf{C}_{\mathrm{P},1} \mid \mathbf{C}_{\mathrm{P},2} \mid \mathbf{C}_{\mathrm{P},3}]$ and

$$\begin{aligned}
\mathbf{C}_{\mathrm{B}} &= \mathbf{SB} + \mathbf{E}_{\mathrm{B}}, \\
\mathbf{C}_{\mathrm{P},0} &= \mathbf{STA}_0 + \mathbf{SF}_0 + \mathbf{E}_0, & \mathbf{C}_{\mathrm{P},1} &= \mathbf{SA}_1 + \mathbf{SF}_1 + \mathbf{E}_1, \\
\mathbf{C}_{\mathrm{P},2} &= \mathbf{SA}_1\mathbf{R} + a\mathbf{S} + \mathbf{SF}_2 + \mathbf{E}_2, & \mathbf{C}_{\mathrm{P},3} &= \mathbf{SA}_3 + \mathbf{SF}_3 + \mathbf{E}_3, \\
\mathsf{aux}_{\mathrm{P}} &= \mathsf{iO}(\mathsf{Vfy}_{\mathsf{param},\mathbf{A}_0,\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{M}}).
\end{aligned}$$

Here $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$, $\mathbf{E}_{\mathrm{B}} \leftarrow \mathcal{D}^{1\times m}_{\sigma_{\mathrm{pre}}}$, and $[\mathbf{E}_0 \mid \mathbf{E}_1 \mid \mathbf{E}_2 \mid \mathbf{E}_3] \leftarrow \mathcal{D}^{1\times \ell}_{\sigma_{\mathrm{pre}}}$, with block widths L, m, n, m. The remaining values are generated by the sampler.

$\boxed{\mathcal{H}_1}$ Same as Hybrid $\mathcal{H}_0$, except that $\mathbf{C}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{1\times m}$ is uniform and $\mathbf{ST}$ in $\mathbf{C}_{\mathrm{P},0}$ is replaced by an independent uniform nonzero row $\mathbf{Q} \leftarrow \mathbb{Z}_q^{1\times n} \setminus \{\mathbf{0}\}$.

$\boxed{\mathcal{H}_2}$ Same as Hybrid $\mathcal{H}_1$, except that we sample $\mathbf{E}_{\mathrm{A}_1}, \mathbf{E}_{\mathrm{A}_3} \leftarrow \mathcal{D}_{\sigma_s}^{1\times m}$, set $\mathbf{U}_1 = \mathbf{S}\mathbf{A}_1 + \mathbf{E}_{\mathrm{A}_1}$ and $\mathbf{U}_3 = \mathbf{S}\mathbf{A}_3 + \mathbf{E}_{\mathrm{A}_3}$, and replace the target responses by

$$\mathbf{C}_{\mathrm{P},0} = \mathbf{Q}\mathbf{A}_0 + \mathbf{E}_0, \quad \mathbf{C}_{\mathrm{P},1} = \mathbf{U}_1 + \mathbf{E}_1,$$
$$\mathbf{C}_{\mathrm{P},2} = \mathbf{U}_1\mathbf{R} + \mathbf{E}_2, \quad \mathbf{C}_{\mathrm{P},3} = \mathbf{U}_3 + \mathbf{E}_3.$$

$\boxed{\mathcal{H}_3}$ Same as Hybrid $\mathcal{H}_2$, except that $\mathbf{U}_1$ and $\mathbf{U}_3$ are sampled independently and uniformly from $\mathbb{Z}_q^{1\times m}$.

$\boxed{\mathcal{H}_4}$ Same as Hybrid $\mathcal{H}_3$, except that $\mathbf{M} = \mathbf{A}_3\mathbf{W} + \mathbf{F}_{\mathrm{M}}$ inside the obfuscation is replaced by $\mathbf{U}_{\mathrm{M}} \leftarrow \mathbb{Z}_q^{n\times d}$.

$\boxed{\mathcal{H}_5}$ Same as Hybrid $\mathcal{H}_4$, except that $\mathsf{aux}_{\mathrm{P}} = \mathsf{iO}(C_\perp)$, where $C_\perp$ is the identically zero circuit with the same input domain and padded size as Vfy.

$\boxed{\mathcal{H}_6}$ Same as Hybrid $\mathcal{H}_5$, except that $\mathbf{C}_{\mathrm{P},1}, \mathbf{C}_{\mathrm{P},2}, \mathbf{C}_{\mathrm{P},3}$ are sampled independently and uniformly.

$\boxed{\mathcal{H}_7}$ Same as Hybrid $\mathcal{H}_6$, except that $\mathbf{C}_{\mathrm{P},0} \leftarrow \mathbb{Z}_q^{1\times L}$ is also uniform. Thus $\mathbf{C}_{\mathrm{B}}$ and $\mathbf{C}_{\mathrm{P}}$ are independent uniform matrices.

$\boxed{\mathcal{H}_8}$ Same as Hybrid $\mathcal{H}_7$, except that the auxiliary input is replaced with

$$\mathsf{aux}_{\mathrm{P}} = \mathsf{iO}(\mathsf{Vfy}_{\mathsf{param},\mathbf{A}_0,\mathsf{td}_{\mathbf{A}_0},\mathbf{R},\mathbf{W},\mathbf{U}_{\mathrm{M}}}),$$

where $(\mathbf{A}_0, \mathsf{td}_{\mathbf{A}_0}, \mathbf{R}, \mathbf{W}, \mathbf{U}_{\mathrm{M}})$ are sampled as in Hybrid $\mathcal{H}_4$, independently of the uniform responses.

$\boxed{\mathcal{H}_9}$ Same as Hybrid $\mathcal{H}_8$, except that $\mathbf{U}_{\mathrm{M}}$ is replaced back with $\mathbf{M} = \mathbf{A}_3\mathbf{W} + \mathbf{F}_{\mathrm{M}}$, where $\mathbf{A}_3, \mathbf{F}_{\mathrm{M}}$ are sampled as in $\mathsf{Samp}_{\mathrm{P}}$. This is the RHS distribution of the precondition in Definition 6.

We now state the indistinguishability of each adjacent pair of hybrids.

Claim 9. *The distributions in Hybrids $\mathcal{H}_0$ and $\mathcal{H}_1$ are statistically indistinguishable.*

Proof. The public-matrix guarantee of $\mathsf{TrapGen}$ (Lemma 1) allows replacing the hidden $\mathbf{B}$ by uniform. It appears in the view only through $\mathbf{C}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}$. For every fixed nonzero $\mathbf{S}$, this response is then uniform, independently of the remaining view. By Corollary 2 and lemma 4, the Gaussian secret $\mathbf{S}$ is nonzero over $\mathbb{Z}_q$ except with negligible probability.

Similarly, for every fixed nonzero $\mathbf{S}$, the product $\mathbf{S}\mathbf{T}$ is uniform over the nonzero rows of $\mathbb{Z}_q^{1\times n}$. Since $\mathbf{T}$ has no other use in the view and this distribution does not depend on $\mathbf{S}$, we can replace $\mathbf{S}\mathbf{T}$ by an independent uniform nonzero row $\mathbf{Q}$. □

Claim 10. *The distributions in Hybrids $\mathcal{H}_1$ and $\mathcal{H}_2$ are statistically indistinguishable.*

Proof. This follows by standard flooding arguments.

Note that between Hybrids $\mathcal{H}_1$ and $\mathcal{H}_2$, the four LWE samples differ by

$$\boldsymbol{\Delta}_0 = \mathbf{S}\mathbf{F}_0, \qquad \boldsymbol{\Delta}_1 = \mathbf{S}\mathbf{F}_1 - \mathbf{E}_{\mathrm{A}_1},$$
$$\boldsymbol{\Delta}_2 = a\mathbf{S} + \mathbf{S}\mathbf{F}_2 - \mathbf{E}_{\mathrm{A}_1}\mathbf{R}, \quad \boldsymbol{\Delta}_3 = \mathbf{S}\mathbf{F}_3 - \mathbf{E}_{\mathrm{A}_3}.$$

By the Gaussian tail bounds (Lemma 4), these shifts have infinity norm at most

$$u = a\lambda\sigma_s + \lambda^2 n\eta\sigma_s + (m+1)\lambda\sigma_s \le 2a\lambda\sigma_s$$

except with negligible probability. Condition on all sampling randomness other than the fresh Gaussian errors $\mathbf{E}_j$ of width σ_{pre}. The Gaussian flooding bound (Lemma 5), followed by averaging over the conditioning, bounds the statistical distance between the two views by

$$O\left(\frac{\ell u}{\sigma_{\mathrm{pre}}}\right) + \mathsf{negl}(\lambda) = O(\ell\lambda\sigma_s 2^{-\lambda}) + \mathsf{negl}(\lambda) \le \mathsf{negl}(\lambda).$$

□

Claim 11. *Assuming* $\mathsf{LWE}_{n,2m+n+\lambda,q,\sigma_s}$*, the distributions in Hybrids* $\mathcal{H}_2$ *and* $\mathcal{H}_3$ *are computationally indistinguishable.*

Proof. Note that in Hybrid $\mathcal{H}_2$, all elements except for $\mathbf{U}_1$ and $\mathbf{U}_3$ are either independent of $\mathbf{s}$ or can be generated via $\mathbf{U}_1$ and $\mathbf{U}_3$. Therefore, we can apply (Gaussian-secret) LWE (Lemma 11) to the public matrix $[\mathbf{A}_1 \mid \mathbf{A}_3]$ with secret $\mathbf{s}$ to switch $\mathbf{U}_1$ and $\mathbf{U}_3$ to uniform, as required by Hybrid $\mathcal{H}_3$. □

Claim 12. *Assuming* $\mathsf{LWE}_{n,d,q,\sigma_M}$*, the distributions in Hybrids* $\mathcal{H}_3$ *and* $\mathcal{H}_4$ *are computationally indistinguishable.*

Proof. Note that in Hybrid $\mathcal{H}_3$, the only element that depends on the secret $\mathbf{A}_3$ is $\mathbf{M} = \mathbf{A}_3\mathbf{W} + \mathbf{F}_M$. Therefore, we can apply (binary-public-matrix) LWE (Lemma 12) to $\mathbf{M}$ with public matrix $\mathbf{W}$ and secret $\mathbf{A}_3$ to switch $\mathbf{M}$ to uniform $\mathbf{U}_M$, as required by Hybrid $\mathcal{H}_4$. □

Claim 13. *Assuming null-iO security, the distributions in Hybrids* $\mathcal{H}_4$ *and* $\mathcal{H}_5$ *are computationally indistinguishable.*

Proof. In Hybrid $\mathcal{H}_4$, $\mathbf{U}_M$ is uniform and independent of every other variable in the view. The parameter setting gives $d = 2(m+n)\lceil\log_2 q\rceil$ and $q \ge 4\beta_M + 2$. By the sparsified rowspace consistency lemma (Lemma 19), except with negligible probability over $\mathbf{U}_M$, there are no $\mathbf{Z}$ and $\widehat{\mathbf{S}} \ne \mathbf{0}$ satisfying

$$\left\|\mathbf{Z}\mathbf{W} - \widehat{\mathbf{S}}\mathbf{U}_M\right\|_\infty \le \beta_M.$$

For every input to Vfy, centered rounding gives

$$\left\|\widehat{\mathbf{S}}\right\|_\infty \le \frac{q-1}{2a} + \frac{1}{2} < q,$$

so the nonzero check in Step 3 ensures that $\widehat{\mathbf{S}} \ne \mathbf{0}$ over $\mathbb{Z}_q$. Every accepting input must also pass the consistency check in Step 4. Thus Vfy is identically zero with overwhelming probability, and null-iO security gives the claim after accounting for its negligible complement. □

Claim 14. *The distributions in Hybrids* $\mathcal{H}_5$ *and* $\mathcal{H}_6$ *are statistically indistinguishable.*

Proof. As aux no longer hardcodes private information, the only remaining use of $\mathbf{R}$ is in $(\mathbf{U}_1 + \mathbf{E}_1, \mathbf{U}_1\mathbf{R} + \mathbf{E}_2)$. By the leftover hash lemma (Lemma 13), $(\mathbf{U}_1, \mathbf{U}_1\mathbf{R})$ is statistically close to independent uniform matrices $(\mathbf{U}_1, \mathbf{U}_2)$, since $m \ge 2\lceil\log_2 q\rceil$. □

Claim 15. *The distributions in Hybrids* $\mathcal{H}_6$ *and* $\mathcal{H}_7$ *are statistically indistinguishable.*

Proof. As aux no longer hardcodes private information, the only remaining use of $\mathbf{A}_0$ is in $\mathbf{QA}_0 + \mathbf{E}_0$. The public-matrix guarantee of TrapGen (Lemma 1) ensures that $\mathbf{A}_0$ is statistically close to uniform. Since $\mathbf{Q} \neq \mathbf{0}$, $\mathbf{C}_{P,0}$ is statistically close to uniform, completing the proof. □

Claim 16. *Assuming null-iO security, the distributions in Hybrids $\mathcal{H}_7$ and $\mathcal{H}_8$ are computationally indistinguishable.*

Proof. Same as Claim 13, but in the reverse direction. □

Claim 17. *Assuming $\mathsf{LWE}_{n,d,q,\sigma_M}$, the distributions in Hybrids $\mathcal{H}_8$ and $\mathcal{H}_9$ are computationally indistinguishable.*

Proof. Same as Claim 12, but in the reverse direction. □

Combining Claims 9 to 17, the distributions in Hybrids $\mathcal{H}_0$ and $\mathcal{H}_9$ are computationally indistinguishable, proving the theorem. □

References

[ABB10a] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Berlin, Heidelberg, May / June 2010.

[ABB10b] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 98–115. Springer, Berlin, Heidelberg, August 2010.

[ABB+23] Divesh Aggarwal, Huck Bennett, Zvika Brakerski, Alexander Golovnev, Rajendra Kumar, Zeyong Li, Spencer Peters, Noah Stephens-Davidowitz, and Vinod Vaikuntanathan. Lattice problems beyond polynomial time. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1516–1526. ACM Press, June 2023.

[ABM+26] Shweta Agrawal, Supriyo Banerjee, Anuja Modi, Rajeev Raghunath, Anshu Yadav, and Shota Yamada. ABE for unbounded circuits from circular LWE and pairings. Cryptology ePrint Archive, Report 2026/1439, 2026.

[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, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Berlin, Heidelberg, August 2009.

[AGHS13] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian leftover hash lemma over infinite domains. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 97–116. Springer, Berlin, Heidelberg, December 2013.

[AKY24] Shweta Agrawal, Simran Kumari, and Shota Yamada. Attribute based encryption for turing machines from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 352–386. Springer, Cham, August 2024.

[AKY25] Shweta Agrawal, Simran Kumari, and Shota Yamada. Pseudorandom FE, iO and applications. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part II*, volume 16269 of *LNCS*, pages 222–258. Springer, Cham, December 2025.

[AKYY23] Shweta Agrawal, Simran Kumari, Anshu Yadav, and Shota Yamada. Broadcast, trace and revoke with optimal parameters from polynomial hardness. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 605–636. Springer, Cham, April 2023.

[AMR25] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Key-homomorphic computations for RAM: Fully succinct randomised encodings and more. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 236–268. Springer, Cham, August 2025.

[AMYY25] Shweta Agrawal, Anuja Modi, Anshu Yadav, and Shota Yamada. Zeroizing attacks against evasive and circular evasive LWE. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part II*, volume 16269 of *LNCS*, pages 259–290. Springer, Cham, December 2025.

[AMYY26] Shweta Agrawal, Anuja Modi, Anshu Yadav, and Shota Yamada. $\rangle O$ from LWE with structured preimage. Cryptology ePrint Archive, Report 2026/2547, 2026.

[ARYY23] Shweta Agrawal, Mélissa Rossi, Anshu Yadav, and Shota Yamada. Constant input attribute based (and predicate) encryption from evasive and tensor LWE. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 532–564. Springer, Cham, August 2023.

[BDE+18] Jonathan Bootle, Claire Delaplace, Thomas Espitau, Pierre-Alain Fouque, and Mehdi Tibouchi. LWE without modular reduction and improved side-channel attacks against BLISS. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 494–524. Springer, Cham, December 2018.

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate iO from homomorphic encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 79–109. Springer, Cham, May 2020.

[BDGM22] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for IO: Circular-secure LWE suffices. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *ICALP 2022*, volume 229 of *LIPIcs*, pages 28:1–28:20. Schloss Dagstuhl, July 2022.

[BDJ+25] Pedro Branco, Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Pseudorandom obfuscation and applications. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part V*, volume 16004 of *LNCS*, pages 663–698. Springer, Cham, August 2025.

[BF11] Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 1–16. Springer, Berlin, Heidelberg, March 2011.

[BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Berlin, Heidelberg, August 2013.

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013.

[BTVW17] Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained PRFs (and more) from LWE. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 264–302. Springer, Cham, November 2017.

[BÜW24] Chris Brzuska, Akin Ünal, and Ivy K. Y. Woo. Evasive LWE assumptions: Definitions, classes, and counterexamples. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IV*, volume 15487 of *LNCS*, pages 418–449. Springer, Singapore, December 2024.

[CHW25] Jeffrey Champion, Yao-Ching Hsieh, and David J. Wu. Registered ABE and adaptively-secure broadcast encryption from succinct LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 3–34. Springer, Cham, August 2025.

[CLW24] Valerio Cini, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based multi-authority/client attribute-based encryption for circuits. *CiC*, 1(4):1, 2024.

[CLW25] Valerio Cini, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based obfuscation from NTRU and equivocal LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 39–72. Springer, Cham, August 2025.

[CM25] Yilei Chen and Xinyu Mao. Universal computational extractors and multi-bit AIPO from lattice assumptions. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 63–93. Springer, Cham, May 2025.

[CW25] Valerio Cini and Hoeteck Wee. Faster ABE for turing machines from circular evasive LWE. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 94–125. Springer, Cham, May 2025.

[DJM+25] Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, and Vinod Vaikuntanathan. Simple and general counterexamples for private-coin evasive LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 73–92. Springer, Cham, August 2025.

[DQV+21] Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Succinct LWE sampling, random polynomials, and obfuscation. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 256–287. Springer, Cham, November 2021.

[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.

[GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th ACM STOC*, pages 660–670. ACM Press, June 2018.

[GN08] Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within Mordell's inequality. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 207–216. ACM Press, May 2008.

[GP21] Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 736–749. ACM Press, June 2021.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.

[GV13] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, fourth edition, 2013.

[HHY26] Tzu-Hsiang Huang, Wei-Hsiang Hung, and Shota Yamada. A note on obfuscation-based attacks on private-coin evasive LWE. *CiC*, 2(4):20, 2026.

[HJL21] Samuel B. Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying iO. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 673–700, Virtual Event, August 2021. Springer, Cham.

[HJL25] Yao-Ching Hsieh, Aayush Jain, and Huijia Lin. Lattice-based post-quantum iO from circular security with random opening assumption. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 3–38. Springer, Cham, August 2025.

[HLL23] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In *64th FOCS*, pages 415–434. IEEE Computer Society Press, November 2023.

[HLL24] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. A general framework for lattice-based ABE using evasive inner-product functional encryption. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 433–464. Springer, Cham, May 2024.

[ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24. ACM Press, May 1989.

[JKLM25] Zhengzhong Jin, Yael Tauman Kalai, Alex Lombardi, and Surya Mathialagan. Universal SNARGs for NP from proofs of correctness. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 933–943. ACM Press, June 2025.

[JLLS23] Aayush Jain, Huijia Lin, Paul Lou, and Amit Sahai. Polynomial-time cryptanalysis of the subspace flooding assumption for post-quantum $i\mathcal{O}$. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 205–235. Springer, Cham, April 2023.

[LLL82] AK Lenstra, HW Lenstra, and L Lovász. Factoring polynomials with rational coefficients. *Math. ann*, 261(4):515–534, 1982.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012.

[MPV24] Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge SNARKs for UP. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 38–71. Springer, Cham, August 2024.

[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.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Berlin, Heidelberg, March 2006.

[PST26] Tapas Pal, Robert Schädlich, and Erkan Tairi. Registered functional encryption for pseudorandom functionalities from lattices: Registered ABE for unbounded depth circuits and turing machines, and more. *IACR Communications in Cryptology*, 3(3), 2026.

[PV08] Chris Peikert and Vinod Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 536–553. Springer, Berlin, Heidelberg, August 2008.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[SBP25] Sora Suegami, Enrico Bottazzi, and Gayeong Park. Diamond iO: A straightforward construction of indistinguishability obfuscation from lattices. Cryptology ePrint Archive, Report 2025/236, 2025.

[Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 535–559. Springer, Cham, August 2022.

[VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 195–221. Springer, Cham, December 2022.

[Wee22] Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Cham, May / June 2022.

[Wee24] Hoeteck Wee. Circuit ABE with poly(depth, λ)-sized ciphertexts and keys from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 178–209. Springer, Cham, August 2024.

[Wee25a] Hoeteck Wee. Almost optimal KP and CP-ABE for circuits from succinct LWE. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 34–62. Springer, Cham, May 2025.

[Wee25b] Hoeteck Wee. Obfuscation via GGH15 encodings, 10 years later. Talk at the Obfuscation workshop, part of the program Cryptography 10 Years Later: Obfuscation, Proof Systems, and Secure Computation, Simons Institute for the Theory of Computing, 2025. June 24, 2025.

[Woo26] Ivy K. Y. Woo. Characterisation of evasive LWE and their hardness (or: The lack of it). Talk at the Novel Lattice Assumptions workshop, Edinburgh, 2026. July 16, 2026.

[WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 127–156. Springer, Cham, October 2021.

[WW25] Hoeteck Wee and David J. Wu. Unbounded distributed broadcast encryption and registered ABE from succinct LWE. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part III*, volume 16002 of *LNCS*, pages 204–235. Springer, Cham, August 2025.

[WWW22] Brent Waters, Hoeteck Wee, and David J. Wu. Multi-authority ABE from lattices without random oracles. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 651–679. Springer, Cham, November 2022.

[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.

[ZZC+25] Ziqi Zhu, Kai Zhang, Zhili Chen, Junqing Gong, and Haifeng Qian. Black-box registered ABE from lattices. Cryptology ePrint Archive, Report 2025/051, 2025.

A Additional Binding Counterexample

We give an additional counterexample to private-coin binding evasive LWE. As opposed to the construction in Section 5, here the target sampler does not use $\mathsf{TrapGen}$, and $\mathsf{aux}_{\mathrm{P}} = \perp$. The obfuscation-style auxiliary information is included in $\mathsf{aux}_{\mathrm{S}}$.

The idea is almost identical to the construction in Section 5, except that the trapdoor structure is now over the secret $\mathbf{S}$ rather than the public matrix $\mathbf{P}$. Since the secret $\mathbf{S}$ is hidden, we can also apply the sparsifying technique from the hiding counterexample in Section 6, weakening the obfuscation assumption from iO to null-iO.

A.1 Sampler Description

$\mathsf{Samp}_{\mathrm{param}}(1^\lambda; r_{\mathrm{Samp}})$**:** Sample a prime q from $[\lambda^{24}, 2\lambda^{24}]$, and set

$$n = 3\lambda, \qquad t = 1, \qquad m = 8m_0(n, q), \qquad \ell = 4m, \qquad \sigma_{\mathrm{pre}} = \sigma_{\mathrm{post}} = \lambda^3, \qquad \rho = m^3. \tag{4}$$

$\mathsf{Vfy}_{\mathbf{X},\mathbf{s}_1}(\mathbf{B},\mathbf{P},\mathbf{C},\mathbf{K})$: The circuit hardcodes $\mathbf{X} \in \mathbb{Z}_q^{\lambda\times 2\lambda}$ and $\mathbf{s}_1 \in \mathbb{Z}_q^{2\lambda}$. Its inputs have dimensions

$$\mathbf{B} = \begin{bmatrix}\mathbf{B}_0\\ \mathbf{B}_1\end{bmatrix} \in \mathbb{Z}_q^{n\times m}, \quad \mathbf{P} \in \mathbb{Z}_q^{n\times \ell}, \quad \mathbf{C} \in \mathbb{Z}_q^{1\times m}, \quad \mathbf{K} \in \mathbb{Z}^{m\times \ell},$$

where $\mathbf{B}_0$ has λ rows. Set $\beta_{\mathrm{E}} = \lambda\sigma_{\mathrm{post}}$ and $\beta_{\mathrm{K}} = \lambda\rho$. Reject whenever a check fails, and accept if all checks pass.

1. Check the input format and dimensions, and that

$$\mathbf{P} = \begin{bmatrix}\mathbf{0}\\ \mathbf{P}_1\end{bmatrix}, \qquad \mathbf{BK} = \mathbf{P}, \qquad \|\mathbf{K}\|_\infty \le \beta_{\mathrm{K}}, \qquad \mathsf{rank}_{\mathbb{Q}}(\mathbf{K}) = m, \qquad \mathsf{rank}_{\mathbb{Z}_q}(\mathbf{B}_0) = \lambda.$$

2. Compute the noise $\mathbf{N} = \mathbf{CK} - \mathbf{s}_1^\intercal\mathbf{P}_1 \in \mathbb{Z}^{1\times\ell}$. Over $\mathbb{Q}$, compute

$$\mathbf{J} = \mathbf{K}^\intercal(\mathbf{K}\mathbf{K}^\intercal)^{-1}, \qquad \widehat{\mathbf{E}} = \mathbf{NJ}.$$

   Check that $\widehat{\mathbf{E}} \in \mathbb{Z}^{1\times m}$ and $\left\|\widehat{\mathbf{E}}\right\|_\infty \le \beta_{\mathrm{E}}$.

3. Solve for $\widehat{\mathbf{s}}_0 \in \mathbb{Z}_q^\lambda$ in

$$\mathbf{C} - \widehat{\mathbf{E}} - \mathbf{s}_1^\intercal\mathbf{B}_1 = \widehat{\mathbf{s}}_0^\intercal\mathbf{B}_0.$$

   Reject if no solution exists. Finally, check that $\left\|\mathbf{s}_1^\intercal - \widehat{\mathbf{s}}_0^\intercal\mathbf{X}\right\|_\infty \le \beta_{\mathrm{E}}$.

Figure 3: Verification circuit for the additional binding counterexample.

Output $\mathsf{param} = (1^\lambda, 1^n, 1^m, q, 1^t, 1^\ell, \sigma_{\mathrm{pre}}, \sigma_{\mathrm{post}}, \rho)$.

$\mathsf{Samp}_{\mathrm{S}}(\mathsf{param}; r_{\mathrm{S}})$**:** Sample $\mathbf{s}_0 \leftarrow \mathbb{Z}_q^\lambda$ and $\mathbf{X} \leftarrow \mathbb{Z}_q^{\lambda\times 2\lambda}$ uniformly, and sample $\mathbf{e}_{\mathrm{M}} \leftarrow \mathcal{D}_{\sigma_{\mathrm{pre}}}^{2\lambda}$. Set

$$\mathbf{s}_1^\intercal = \mathbf{s}_0^\intercal\mathbf{X} + \mathbf{e}_{\mathrm{M}}^\intercal, \qquad \mathbf{S} = [\mathbf{s}_0^\intercal \mid \mathbf{s}_1^\intercal].$$

Output

$$\left(\mathbf{S}, \mathsf{aux}_{\mathrm{S}} = (\mathbf{X}, \mathsf{iO}(\mathsf{Vfy}_{\mathbf{X},\mathbf{s}_1}))\right),$$

where the circuit Vfy is specified in Figure 3.

$\mathsf{Samp}_{\mathrm{P}}(\mathsf{param}; r_{\mathrm{P}})$**:** Sample $\mathbf{P}_1 \leftarrow \mathbb{Z}_q^{2\lambda\times\ell}$ uniformly and output

$$\mathbf{P} = \begin{bmatrix}\mathbf{0}_{\lambda\times\ell}\\ \mathbf{P}_1\end{bmatrix}, \qquad \mathsf{aux}_{\mathrm{P}} = \bot.$$

A.2 Postcondition Adversary

Given $(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux})$, the adversary extracts the obfuscated circuit $\widetilde{\mathsf{Vfy}}$ from $\mathsf{aux}_{\mathrm{S}}$ and outputs $\widetilde{\mathsf{Vfy}}(\mathbf{B}, \mathbf{P}, \mathbf{C}_{\mathrm{B}}, \mathbf{K})$.

Theorem 8 (Binding postcondition advantage)**.** *Consider the sampler above and the bounds in Figure 3. Suppose*

$$m \geq m_0(n,q), \qquad \rho \geq \rho_0(n,m), \qquad m\beta_E\beta_K < q/2, \qquad q \geq (2\beta_E+1)^2, \qquad m \geq 4\lambda.$$

Then $\mathsf{Adv}^{\text{bind-post}}_{\mathsf{Samp},\mathcal{A}_{\text{post}}}(\lambda) \geq 1 - \mathsf{negl}(\lambda)$.

The parameters in Equation (4) satisfy these conditions: $m = O(\lambda \log \lambda)$, $m\beta_E\beta_K = \lambda^5 m^4 = o(q)$, and $2\beta_E + 1 = O(\lambda^4)$.

We prove Theorem 8 by analyzing the two branches of the postcondition experiment separately.

Lemma 20 (LWE branch)**.** *Under the parameter setting of Theorem 8, sample the LWE branch of the binding postcondition experiment (Definition 5). The adversary accepts* $\mathbf{C}_B = \mathbf{SB} + \mathbf{E}_B$ *with probability* $1 - \mathsf{negl}(\lambda)$.

Proof. The preimage tail and rank bounds (Corollaries 1 and 4) give $\|\mathbf{K}\|_\infty \leq \rho \log n \leq \beta_K$ and $\mathsf{rank}_{\mathbb{Q}}(\mathbf{K}) = m$, except with negligible probability. Since $\mathbf{B}$ is statistically close to uniform, its first λ rows have full row rank except with probability $q^{\lambda-m} + \mathsf{negl}(\lambda)$. The Gaussian tail bounds (Lemma 4) give

$$\|\mathbf{E}_B\|_\infty \leq \sigma_{\text{post}} \log n \leq \beta_E, \qquad \|\mathbf{e}_M\|_\infty \leq \sigma_{\text{pre}} \log n \leq \beta_E,$$

which immediately implies that

$$\mathbf{C}_B\mathbf{K} - \mathbf{s}_1^\intercal\mathbf{P}_1 = (\mathbf{s}_0^\intercal\mathbf{B}_0 + \mathbf{s}_1^\intercal\mathbf{B}_1 + \mathbf{E}_B)\mathbf{K} - \mathbf{s}_1^\intercal\mathbf{P}_1 = \mathbf{E}_B\mathbf{K} \pmod q, \qquad \|\mathbf{E}_B\mathbf{K}\|_\infty \leq m\beta_E\beta_K < q/2.$$

Consequently, $\mathbf{N} = \mathbf{E}_B\mathbf{K}$ over the integers and $\widehat{\mathbf{E}} = \mathbf{E}_B\mathbf{K}\mathbf{J} = \mathbf{E}_B$. The unique solution in Step 3 is therefore $\widehat{\mathbf{s}}_0 = \mathbf{s}_0$, and the final difference is $\mathbf{e}_M^\intercal$, which has norm bounded by β_E. Therefore all checks pass and the adversary outputs 1. □

Lemma 21 (Uniform branch)**.** *Under the parameter setting and sampling of Lemma 20, the adversary accepts an independent* $\mathbf{U}_B \leftarrow \mathbb{Z}_q^{1\times m}$ *with probability* $\mathsf{negl}(\lambda)$.

Proof. Fix $\mathbf{B}, \mathbf{P}, \mathbf{K}, \mathbf{X}, \mathbf{s}_1$ and the obfuscation coins. Acceptance implies that $\mathbf{U}_B = \widehat{\mathbf{s}}_0^\intercal\mathbf{B}_0 + \mathbf{s}_1^\intercal\mathbf{B}_1 + \widehat{\mathbf{E}}$ for some $\widehat{\mathbf{s}}_0 \in \mathbb{Z}_q^\lambda$ and integer $\widehat{\mathbf{E}}$ with $\left\|\widehat{\mathbf{E}}\right\|_\infty \leq \beta_E$. Thus the probability is at most

$$\frac{q^\lambda \cdot (2\beta_E+1)^m}{q^m} \leq q^{\lambda - m/2} \leq q^{-\lambda} \leq \mathsf{negl}(\lambda),$$

which completes the proof. □

Proof of Theorem 8. The theorem follows by combining Lemmas 20 and 21. □

A.3 Pseudorandomness of the Precondition Distribution

Theorem 9 (Binding precondition)**.** *Consider the sampler with parameters from Equation (4). Assume null-iO (Definition 2) and the LWE assumption* $\mathsf{LWE}_{\lambda,\ell,q,\sigma_{\text{pre}}}$ *with the sampler's public modulus distribution. Then, for every PPT adversary* $\mathcal{A}_{\text{pre}}$,

$$\mathsf{Adv}^{\text{bind-pre}}_{\mathsf{Samp},\mathcal{A}_{\text{pre}}}(\lambda) \leq \mathsf{negl}(\lambda).$$

Proof. Consider the following sequence of hybrids.

$\boxed{\mathcal{H}_0}$ This is the LHS of the precondition distribution, consisting of

$$(\mathsf{param}, \mathbf{B}, \mathbf{P}, \mathbf{C}_B = \mathbf{s}_0^\intercal \mathbf{B}_0 + \mathbf{s}_1^\intercal \mathbf{B}_1 + \mathbf{E}_B, \mathbf{C}_P = \mathbf{s}_1^\intercal \mathbf{P}_1 + \mathbf{E}_P, \mathsf{aux})$$

Here $\mathbf{B} = \begin{bmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{bmatrix}$ is sampled with a trapdoor, and both $\mathbf{E}_B$ and $\mathbf{E}_P$ have width σ_{pre}.

$\boxed{\mathcal{H}_1}$ Same as Hybrid $\mathcal{H}_0$, but sample $\mathbf{B}$ uniformly.

$\boxed{\mathcal{H}_2}$ Same as Hybrid $\mathcal{H}_1$, but replace $(\mathbf{s}_1^\intercal = \mathbf{s}_0^\intercal \mathbf{X} + \mathbf{e}_M^\intercal, \mathbf{s}_0^\intercal \mathbf{B}_0 + \mathbf{E}_B)$ with independent uniform rows $(\mathbf{u}_1^\intercal, \mathbf{U}_0)$. Thus $\mathbf{C}_B = \mathbf{U}_0 + \mathbf{u}_1^\intercal \mathbf{B}_1$, $\mathbf{C}_P = \mathbf{u}_1^\intercal \mathbf{P}_1 + \mathbf{E}_P$, and $\mathsf{aux}_S = (\mathbf{X}, \mathsf{iO}(\mathsf{Vfy}_{\mathbf{X},\mathbf{u}_1}))$.

$\boxed{\mathcal{H}_3}$ Same as Hybrid $\mathcal{H}_2$, but replace the obfuscation by $\mathsf{iO}(C_\perp)$, where $C_\perp$ is the identically zero circuit with the same input domain and padded size as Vfy.

$\boxed{\mathcal{H}_4}$ Same as Hybrid $\mathcal{H}_3$, but replace $\mathbf{C}_B$ with an independent uniform $\mathbf{U}_B$.

$\boxed{\mathcal{H}_5}$ Same as Hybrid $\mathcal{H}_4$, but replace $\mathbf{C}_P$ with an independent uniform $\mathbf{U}_P$.

$\boxed{\mathcal{H}_6}$ Same as Hybrid $\mathcal{H}_5$, but restore $\mathsf{iO}(\mathsf{Vfy}_{\mathbf{X},\mathbf{u}_1})$.

$\boxed{\mathcal{H}_7}$ Same as Hybrid $\mathcal{H}_6$, but restore $\mathbf{s}_1^\intercal = \mathbf{s}_0^\intercal \mathbf{X} + \mathbf{e}_M^\intercal$ in the auxiliary information. Both challenge samples remain uniform.

$\boxed{\mathcal{H}_8}$ Same as Hybrid $\mathcal{H}_7$, but sample $\mathbf{B}$ using $\mathsf{TrapGen}$. This is the RHS of the precondition distribution.

Claim 18. *Hybrids $\mathcal{H}_0$ and $\mathcal{H}_1$ are statistically indistinguishable.*

Proof. The public-matrix guarantee of $\mathsf{TrapGen}$ (Lemma 1) states that $\mathbf{B}$ is statistically close to uniform. This applies because $\mathsf{td}_{\mathbf{B}}$ is not used in the precondition view. □

Claim 19. *Assuming $\mathsf{LWE}_{\lambda,\ell,q,\sigma_{\mathrm{pre}}}$, Hybrids $\mathcal{H}_1$ and $\mathcal{H}_2$ are computationally indistinguishable.*

Proof. This follows by applying LWE with secret $\mathbf{s}_0$ and public matrix $[\mathbf{X} \mid \mathbf{B}_0]$. □

Claim 20. *Assuming null-iO, Hybrids $\mathcal{H}_2$ and $\mathcal{H}_3$ are computationally indistinguishable.*

Proof. By counting the choices of $\widehat{\mathbf{s}}_0$ and the bounded residual, for any fixed $\mathbf{X}$, a uniform $\mathbf{u}_1$ admits some $\widehat{\mathbf{s}}_0 \in \mathbb{Z}_q^\lambda$ with $\left\|\mathbf{u}_1^\intercal - \widehat{\mathbf{s}}_0^\intercal \mathbf{X}\right\|_\infty \le \beta_E$ with probability at most

$$\frac{q^\lambda \cdot (2\beta_E + 1)^{2\lambda}}{q^{2\lambda}} = \left(\frac{(2\beta_E + 1)^2}{q}\right)^\lambda \le \mathsf{negl}(\lambda).$$

Therefore, the circuit is identically zero with overwhelming probability. The indistinguishability then follows from null-iO security. □

Claim 21. *Hybrids $\mathcal{H}_3$ and $\mathcal{H}_4$ are identically distributed.*

Proof. Obvious since $\mathbf{U}_0$ is uniform and independent. □

Claim 22. *Assuming $\mathsf{LWE}_{2\lambda,\ell,q,\sigma_{\mathrm{pre}}}$, Hybrids $\mathcal{H}_4$ and $\mathcal{H}_5$ are computationally indistinguishable.*

Proof. Note that after replacing the obfuscation with null, no other component of the view depends on $\mathbf{u}_1$. Therefore the claim follows by applying LWE with secret $\mathbf{u}_1$ and public matrix $\mathbf{P}_1$. □

Claim 23. *Assuming null-iO, Hybrids $\mathcal{H}_5$ and $\mathcal{H}_6$ are computationally indistinguishable.*

Proof. Same as Claim 20. □

Claim 24. *Assuming $\mathsf{LWE}_{\lambda,\ell,q,\sigma_{\mathrm{pre}}}$, Hybrids $\mathcal{H}_6$ and $\mathcal{H}_7$ are computationally indistinguishable.*

Proof. Same as Claim 19. □

Claim 25. *Hybrids $\mathcal{H}_7$ and $\mathcal{H}_8$ are statistically indistinguishable.*

Proof. Same as Claim 18. □

Combining all the above proves indistinguishability of Hybrids $\mathcal{H}_0$ and $\mathcal{H}_8$, as required. □

B Additional Hiding Counterexample

We give an additional counterexample to private-coin hiding evasive LWE. As opposed to the construction in Section 6, here the target sampler does not use $\mathsf{TrapGen}$, and $\mathsf{aux}_{\mathrm{P}} = \perp$. The obfuscation-style auxiliary information is included in $\mathsf{aux}_{\mathrm{S}}$.

The construction is pretty much a direct composition of the binding counterexample in Section A and the hiding counterexample in Section 6, where we start from the binding counterexample and add extra columns to $\mathbf{P}$ to introduce the fragile trapdoor structure from the hiding counterexample. One subtlety is that the verification circuit is not directly given $\mathbf{P}$ in the hiding setting. We therefore add structured rows to $\mathbf{S}$ to obtain noisy samples of the necessary part of the target matrix $\mathbf{P}$.

As in Section 6, fix a perturbation level $B \geq 1$ and write $\kappa = \log_2 B = \mathsf{poly}(\lambda)$.

B.1 Sampler Description

$\mathsf{Samp}_{\mathrm{param}}(1^\lambda; r_{\mathrm{Samp}})$**:** Sample a prime q from $[2^{3\lambda+\kappa}, 2^{3\lambda+\kappa+1})$, and set

$$\begin{gathered} k = \lceil \lambda + \kappa \rceil^2, \qquad n = 2k, \qquad t = k+1, \qquad m = 8m_0(n,q), \\ L = m^3, \qquad \ell = L + k, \qquad d = 2k, \qquad \sigma_{\mathrm{s}} = \mathsf{poly}(k) > k, \\ \eta = 2^\lambda B, \qquad a = m^7 \sigma_{\mathrm{s}} \eta, \qquad \rho = \rho_0(n,m) = 32m, \qquad \sigma_{\mathrm{pre}} = \sigma_{\mathrm{post}} = 2^\lambda a. \end{gathered} \tag{5}$$

Here σ_{s} is an integer-valued polynomial chosen for the LWE hardness regime in Section 3.5. Output $\mathsf{param} = (1^\lambda, 1^n, 1^m, q, 1^t, 1^\ell, \sigma_{\mathrm{pre}}, \sigma_{\mathrm{post}}, \rho)$.

$\mathsf{Samp}_{\mathrm{S}}(\mathsf{param}; r_{\mathrm{S}})$**:** Sample independent column vectors $\mathbf{s}_0, \mathbf{s}_1 \leftarrow \mathcal{D}_{\sigma_{\mathrm{s}}}^k$, and sample

$$\mathbf{X} \leftarrow \mathbb{Z}_q^{k \times d}, \qquad \mathbf{e}_{\mathrm{M}} \leftarrow \mathcal{D}_{\sigma_{\mathrm{s}}}^d, \qquad \mathbf{M} = \mathbf{s}_0^{\mathsf{T}} \mathbf{X} + \mathbf{e}_{\mathrm{M}}^{\mathsf{T}} \in \mathbb{Z}_q^{1 \times d}.$$

Output

$$\mathbf{S} = \begin{bmatrix} \mathbf{s}_0^{\mathsf{T}} & \mathbf{s}_1^{\mathsf{T}} \\ \mathbf{0}_{k \times k} & \mathbf{I}_k \end{bmatrix}, \qquad \mathsf{aux}_{\mathrm{S}} = (\mathbf{X}, \mathsf{iO}(\mathsf{Vfy}_{\mathbf{s}_1, \mathbf{X}, \mathbf{M}})),$$

where the circuit is specified in Figure 4.

$\mathsf{Vfy}_{\mathbf{s}_1,\mathbf{X},\mathbf{M}}(\mathbf{C},\mathbf{K})$: The circuit hardcodes $\mathbf{s}_1 \in \mathbb{Z}^k$, $\mathbf{X} \in \mathbb{Z}_q^{k\times d}$, and $\mathbf{M} \in \mathbb{Z}_q^{1\times d}$. The inputs have dimensions

$$\mathbf{C} = \begin{bmatrix}\mathbf{C}_0\\ \mathbf{C}_1\end{bmatrix} \in \mathbb{Z}_q^{t\times m}, \qquad \mathbf{K} = [\mathbf{K}_0 \mid \mathbf{K}_1] \in \mathbb{Z}^{m\times \ell},$$

where $\mathbf{C}_0$ has one row and $\mathbf{K}_0$ has L columns. Set

$$\beta_{\mathrm{N}} = m^4\sigma_{\mathrm{s}}\rho\sigma_{\mathrm{post}}, \qquad \beta_{\mathrm{M}} = \lambda\sigma_{\mathrm{s}}.$$

Reject whenever a check fails, and accept if all checks pass.

1. Check the input format and dimensions. Check $\mathsf{rank}_{\mathbb{Z}_q}(\mathbf{K}_0) \geq k$ and $\mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_0) = m$. Over $\mathbb{Q}$, compute $\mathbf{J} = \mathbf{K}_0^{\intercal}(\mathbf{K}_0\mathbf{K}_0^{\intercal})^{-1}$.
2. Set $\mathbf{R} = \mathbf{C}_0 - \mathbf{s}_1^{\intercal}\mathbf{C}_1 \in \mathbb{Z}_q^{1\times m}$. Compute the noise $\mathbf{N}_0 = \mathbf{R}\mathbf{K}_0 \in \mathbb{Z}^{1\times L}$, and check $\|\mathbf{N}_0\|_\infty \leq \beta_{\mathrm{N}}$. Set
$$\widehat{\mathbf{E}} = \lfloor \mathbf{N}_0\mathbf{J} \rceil \in \mathbb{Z}^{1\times m}, \qquad \widetilde{\mathbf{C}} = \mathbf{R} - \widehat{\mathbf{E}} \in \mathbb{Z}_q^{1\times m}.$$
3. Using the centered representative of $\widetilde{\mathbf{C}}\mathbf{K}_1$, set
$$\widehat{\mathbf{s}}_0^{\intercal} = \left\lfloor \frac{\widetilde{\mathbf{C}}\mathbf{K}_1}{a} \right\rceil.$$
Check that $\left\|\mathbf{M} - \widehat{\mathbf{s}}_0^{\intercal}\mathbf{X}\right\|_\infty \leq \beta_{\mathrm{M}}$.

Figure 4: Verification circuit for the additional hiding counterexample.

$\mathsf{Samp}_{\mathrm{P}}(\mathsf{param}; r_{\mathrm{P}})$**:** Sample

$$\mathbf{A}_0 \leftarrow \mathbb{Z}_q^{k\times L}, \qquad \mathbf{A}_1 \leftarrow \mathbb{Z}_q^{k\times k}, \qquad \mathbf{F} \leftarrow \mathcal{D}_\eta^{n\times \ell}.$$

Output

$$\mathbf{P} = \begin{bmatrix}\mathbf{0}_{k\times L} & a\mathbf{I}_k\\ \mathbf{A}_0 & \mathbf{A}_1\end{bmatrix} + \mathbf{F}, \qquad \mathsf{aux}_{\mathrm{P}} = \perp.$$

Write $\mathbf{F} = \begin{bmatrix}\mathbf{F}_{00} & \mathbf{F}_{01}\\ \mathbf{F}_{10} & \mathbf{F}_{11}\end{bmatrix}$ and $\mathbf{P} = [\mathbf{P}_0 \mid \mathbf{P}_1]$ with dimensions corresponding to the above partitions.

B.2 Helper lemmas

We first extend the right-inverse bound in Corollary 5 from uniform random targets to arbitrary targets. The following discrete convolution lemma lets us apply the Gaussian singular-value bound (Lemma 10) after adding a small independent error.

Lemma 22 (Discrete Gaussian convolution [BF11, Lemma 4.12])**.** *Let $\Lambda \subseteq \mathbb{Z}^m$ be a full-rank lattice, $\mathbf{c} \in \mathbb{Z}^m$, and $\varepsilon \leq \mathsf{negl}(m)$. Let $\rho, s > 0$ satisfy*

$$(\rho^{-2} + s^{-2})^{-1/2} > \eta_\varepsilon(\Lambda).$$

For independent samples

$$\mathbf{x} \leftarrow \mathcal{D}_{\Lambda+\mathbf{c},\rho}, \qquad \mathbf{e} \leftarrow \mathcal{D}_s^m,$$

the distribution of $\mathbf{x}+\mathbf{e}$ *has statistical distance* $\mathsf{negl}(m)$ *from* $\mathcal{D}^m_{\sqrt{\rho^2+s^2}}$.

Lemma 23 (Small right inverse of $\mathbf{B}^{-1}(\mathbf{P})$ for arbitrary targets)**.** *Let* $m \geq m_0(n,q)$, $L \geq m^3$, *and* $\rho \geq \rho_0(n,m)$. *Sample* $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$, *choose any target* $\mathbf{P} \in \mathbb{Z}_q^{n\times L}$, *and sample* $\mathbf{K} \leftarrow \mathbf{B}_\rho^{-1}(\mathbf{P})$. *Except with probability* $\mathsf{negl}(n)$, *the matrix* $\mathbf{K}$ *has full row rank over* $\mathbb{Q}$ *and its rational right inverse* $\mathbf{J} = \mathbf{K}^\intercal(\mathbf{K}\mathbf{K}^\intercal)^{-1}$ *satisfies*

$$\mathbf{K}\mathbf{J} = \mathbf{I}_m, \qquad \|\mathbf{J}\|_\infty \leq \frac{1}{\rho}.$$

Proof. Condition on the smoothing bounds (Lemma 2), and fix $\mathbf{B}, \mathbf{P}$. Write $\Lambda = \Lambda^\perp(\mathbf{B})$ and choose $\varepsilon \leq \mathsf{negl}(m)$ satisfying $\eta_\varepsilon(\Lambda) \leq \log m$. Up to statistical distance $\mathsf{negl}(n)$, the columns of $\mathbf{K}$ are independent Gaussians of width ρ on the prescribed cosets of Λ.

Set $s = 2\log m$, independently sample $\mathbf{E} \leftarrow \mathcal{D}_s^{m\times L}$, and write $\mathbf{X} = \mathbf{K} + \mathbf{E}$. Since $\rho \geq 32m \geq s$,

$$(\rho^{-2} + s^{-2})^{-1/2} \geq s/\sqrt{2} = \sqrt{2}\log m > \eta_\varepsilon(\Lambda).$$

By the discrete Gaussian convolution bound (Lemma 22), $\mathbf{X}$ is $\mathsf{negl}(n)$-close to $\mathcal{D}^{m\times L}_{\sqrt{\rho^2+s^2}}$. The smoothing bound for $\mathbb{Z}^m$ gives $\eta_{e^{-m}}(\mathbb{Z}^m) = O(\sqrt{m})$. Since $\rho \geq 32m$, applying the Gaussian singular-value bound (Lemma 10) therefore gives

$$s_{\min}(\mathbf{X}) \geq c\sqrt{\rho^2+s^2}\sqrt{L} \geq c\rho\sqrt{L}$$

except with negligible probability.

The Gaussian tail bounds (Lemma 4), applied entrywise, give $\|\mathbf{E}\|_\infty \leq s\log m$ except with negligible probability. Consequently,

$$\|\mathbf{E}\|_{\mathrm{op}} \leq \sqrt{mL}\,\|\mathbf{E}\|_\infty \leq s\sqrt{mL}\log m = o(\rho\sqrt{L}),$$

using $s = 2\log m$ and $\rho \geq 32m$. By the singular-value perturbation bound [GV13, Corollary 8.6.2],

$$s_{\min}(\mathbf{K}) = s_{\min}(\mathbf{X}-\mathbf{E}) \geq s_{\min}(\mathbf{X}) - \|\mathbf{E}\|_{\mathrm{op}} \geq \frac{c}{2}\rho\sqrt{L}.$$

Thus $\mathbf{K}$ has full row rank, and the right-inverse bound (Lemma 9) gives

$$\|\mathbf{J}\|_\infty \leq \frac{2}{c\rho\sqrt{L}} \leq \frac{1}{\rho}.$$

Accounting for the exceptional smoothing event and the preimage-sampling error proves the claim. □

B.3 Postcondition Adversary

Now we are ready to describe the postcondition adversary. Given $(\mathsf{param}, \mathbf{C}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux})$, the adversary extracts $\widetilde{\mathsf{Vfy}}$ from $\mathsf{aux}_{\mathrm{S}}$ and outputs $\widetilde{\mathsf{Vfy}}(\mathbf{C}_{\mathrm{B}}, \mathbf{K})$.

Theorem 10 (Hiding postcondition advantage)**.** *Consider the sampler with parameters from Equation* (5)*. Then*

$$\mathsf{Adv}^{\text{hide-post}}_{\mathsf{Samp},\mathcal{A}_{\text{post}}}(\lambda) \geq 1 - \mathsf{negl}(\lambda).$$

We now prove Theorem 10 by analyzing the two branches of the postcondition experiment separately.

Lemma 24 (LWE branch)**.** *Under the parameter setting of Theorem 10, sample the LWE branch of the hiding postcondition experiment (Definition 6), with* $\mathbf{C}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}$*. Then*

$$\Pr[\mathcal{A}_{\text{post}}(\mathsf{param}, \mathbf{C}_{\mathrm{B}}, \mathbf{K}, \mathsf{aux}) = 1] \geq 1 - \mathsf{negl}(\lambda).$$

Proof. Write

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{bmatrix}, \qquad \mathbf{E}_{\mathrm{B}} = \begin{bmatrix} \mathbf{E}_0 \\ \mathbf{E}_1 \end{bmatrix},$$

where $\mathbf{B}_0, \mathbf{B}_1$ each have k rows and $\mathbf{E}_0$ has one row. Cancellation in Step 2 gives

$$\mathbf{R} = \mathbf{s}_0^{\mathsf{T}}\mathbf{B}_0 + \mathbf{E}, \qquad \mathbf{E} = \mathbf{E}_0 - \mathbf{s}_1^{\mathsf{T}}\mathbf{E}_1 \in \mathbb{Z}^{1\times m}.$$

By the Gaussian and preimage tail bounds (Lemma 4 and corollary 1), except with negligible probability,

$$\|\mathbf{s}_i\|_\infty \leq \lambda\sigma_{\mathrm{s}} \quad (i = 0, 1), \qquad \|\mathbf{F}\|_\infty \leq \lambda\eta, \qquad \|\mathbf{E}_{\mathrm{B}}\|_\infty \leq \lambda\sigma_{\text{post}},$$
$$\|\mathbf{K}\|_\infty \leq \rho \log n \leq \lambda\rho, \qquad \|\mathbf{e}_{\mathrm{M}}\|_\infty \leq \lambda\sigma_{\mathrm{s}} = \beta_{\mathrm{M}}.$$

The right-inverse bound for arbitrary targets (Lemma 23), applied to $\mathbf{P}_0$, also gives $\mathsf{rank}_{\mathbb{Q}}(\mathbf{K}_0) = m$ and $\|\mathbf{J}\|_\infty \leq 1/\rho$ except with negligible probability. Finally, the lower block $\mathbf{A}_0 + \mathbf{F}_{10}$ of $\mathbf{P}_0$ is uniform, so it has row rank k except with probability at most q^{k-L}. Since $\mathbf{B}\mathbf{K}_0 = \mathbf{P}_0$, this implies $\mathsf{rank}_{\mathbb{Z}_q}(\mathbf{K}_0) \geq k$. We condition on all these bounds.

The upper block of $\mathbf{B}\mathbf{K}_0 = \mathbf{P}_0$ gives $\mathbf{B}_0\mathbf{K}_0 = \mathbf{F}_{00}$. Consequently,

$$\mathbf{R}\mathbf{K}_0 = \mathbf{s}_0^{\mathsf{T}}\mathbf{F}_{00} + \mathbf{E}\mathbf{K}_0 \pmod q.$$

The integer expression on the right satisfies

$$\left\|\mathbf{s}_0^{\mathsf{T}}\mathbf{F}_{00} + \mathbf{E}\mathbf{K}_0\right\|_\infty \leq k\lambda^2\sigma_{\mathrm{s}}\eta + m\lambda^2\rho\sigma_{\text{post}}(1 + k\lambda\sigma_{\mathrm{s}}) \leq \beta_{\mathrm{N}} < q/2.$$

Thus it is exactly $\mathbf{N}_0$, and the norm check passes. Define $\mathbf{D} = \mathbf{F}_{00}\mathbf{J} \in \mathbb{Q}^{k\times m}$. Since $\mathbf{K}_0\mathbf{J} = \mathbf{I}_m$,

$$\widehat{\mathbf{E}} = \mathbf{E} + \mathbf{s}_0^{\mathsf{T}}\mathbf{D} + \mathbf{\Delta}, \qquad \|\mathbf{\Delta}\|_\infty \leq 1/2.$$

The pruned sample therefore satisfies

$$\widetilde{\mathbf{C}}\mathbf{K}_1 = a\mathbf{s}_0^{\mathsf{T}} + \widetilde{\mathbf{N}}_1 \pmod q, \qquad \widetilde{\mathbf{N}}_1 = \mathbf{s}_0^{\mathsf{T}}(\mathbf{F}_{01} - \mathbf{D}\mathbf{K}_1) - \mathbf{\Delta}\mathbf{K}_1.$$

Here we used $\mathbf{B}_0\mathbf{K}_1 = a\mathbf{I}_k + \mathbf{F}_{01}$. The bounds above give $\|\mathbf{D}\|_\infty \leq L\lambda\eta/\rho$, and hence

$$\left\|\widetilde{\mathbf{N}}_1\right\|_\infty \leq k\lambda^2\sigma_{\mathrm{s}}\eta + kmL\lambda^3\sigma_{\mathrm{s}}\eta + m\lambda\rho/2 \leq m^6\sigma_{\mathrm{s}}\eta < a/2.$$

Here we use $L = m^3$, $\rho = 32m$, $\lambda^2 \leq k$, and $m = 48n\lceil\log_2 q\rceil \geq 288\lambda k$. Also $a\lambda\sigma_{\mathrm{s}} + m^6\sigma_{\mathrm{s}}\eta < q/2$, so the centered representative of $\widetilde{\mathbf{C}}\mathbf{K}_1$ is the displayed integer expression. Rounding recovers $\widehat{\mathbf{s}}_0 = \mathbf{s}_0$, and the final residual is $\mathbf{e}_{\mathrm{M}}^{\mathsf{T}}$, of norm at most β_{M}. □

Lemma 25 (Uniform branch)**.** *Under the parameter setting of Theorem 10, sample* $(\mathsf{param}, \mathbf{K}, \mathsf{aux})$ *as in Lemma 24, and independently sample* $\mathbf{U}_B \leftarrow \mathbb{Z}_q^{t\times m}$*. Then*

$$\Pr[\mathcal{A}_{\text{post}}(\mathsf{param}, \mathbf{U}_B, \mathbf{K}, \mathsf{aux}) = 1] \le \mathsf{negl}(\lambda).$$

Proof. Fix $\mathbf{K}$, the sampler outputs, the obfuscation coins, and the last k rows of $\mathbf{U}_B$. If the rank checks fail, the circuit rejects. Otherwise $\mathbf{R}$ is uniform in $\mathbb{Z}_q^{1\times m}$, and $\mathbf{K}_0$ has at least k linearly independent columns modulo q. On these columns, $\mathbf{R}\mathbf{K}_0$ is uniform in $\mathbb{Z}_q^{1\times k}$. By counting the bounded residuals on these columns, the norm check in Step 2 passes with probability at most

$$\frac{(2\beta_N + 1)^k}{q^k} \le \mathsf{negl}(\lambda).$$

□

Proof of Theorem 10. The theorem follows by combining Lemmas 24 and 25. □

B.4 Proof of Preconditions

We show that the sampler satisfies both preconditions in Definition 6: target hiding under bounded perturbations and pseudorandomness of the joint LWE samples.

B.4.1 Target hiding

The Gaussian error $\mathbf{F}$ added to $\mathbf{P}$ floods any perturbation of norm at most B, while $\mathsf{aux}_P = \perp$.

Theorem 11 (Target-hiding precondition)**.** *Consider the sampler with parameters from Equation (5). Then, for every adversary* $\mathcal{A}_{\text{pre2}}$*,*

$$\mathsf{Adv}^{\text{hide-pre2}}_{\mathsf{Samp},\mathcal{A}_{\text{pre2}}}(\lambda, B) \le \mathsf{negl}(\lambda).$$

Proof. The target sampler computes

$$\mathbf{P} = \mathbf{P}^* + \mathbf{F}, \qquad \mathbf{F} \leftarrow \mathcal{D}_\eta^{n\times\ell},$$

where $\mathbf{P}^*$ is independent of $\mathbf{F}$ and $\mathsf{aux}_P = \perp$. Let $\mathbf{\Delta}$ be the independent perturbation from the target-hiding experiment (Definition 6), with $\|\mathbf{\Delta}\|_\infty \le B$. By the Gaussian flooding bound (Lemma 5), the statistical distance between the two target distributions is at most

$$O\left(\frac{n\ell B}{\eta}\right) = O(n\ell\, 2^{-\lambda}) \le \mathsf{negl}(\lambda).$$

□

B.4.2 Pseudorandomness of the precondition distribution

Theorem 12 (Hiding precondition)**.** *Consider the sampler with parameters from Equation (5). Assume null-iO (Definition 2) and the LWE assumption (Definition 1)*

$$\mathsf{LWE}_{k,\ell+k+\lambda,q,\sigma_s},$$

with the sampler's public modulus distribution. Then, for every PPT adversary $\mathcal{A}_{\text{pre}}$*,*

$$\mathsf{Adv}^{\text{hide-pre}}_{\mathsf{Samp},\mathcal{A}_{\text{pre}}}(\lambda) \le \mathsf{negl}(\lambda).$$

Proof. We argue that the precondition advantage is negligible via the following sequence of hybrids. Throughout, $\mathsf{aux} = (r_{\mathsf{Samp}}, \mathsf{aux}_{\mathsf{S}}, \perp)$, and the public parameter tape is generated as in $\mathsf{Samp}_{\mathsf{param}}$. Let $\overline{\mathbf{A}} = [\mathbf{A}_0 \mid \mathbf{A}_1] + [\mathbf{F}_{10} \mid \mathbf{F}_{11}]$ and $\mathbf{F}_0 = [\mathbf{F}_{00} \mid \mathbf{F}_{01}]$. The matrix $\overline{\mathbf{A}}$ is uniform and independent of $\mathbf{F}_0$. Each hybrid outputs $(\mathsf{param}, \mathbf{C}_{\mathrm{B}}, \mathbf{C}_{\mathrm{P}}, \mathsf{aux})$. Let $\mathbf{E}_{\mathrm{P},0}$ and $\mathbf{E}_{\mathrm{P},1}$ denote the first row and the last k rows of the fresh target error, respectively, all of width σ_{pre}.

$\boxed{\mathcal{H}_0}$ This is the LHS distribution of the precondition in Definition 6, with

$$\begin{aligned} \mathbf{C}_{\mathrm{B}} &= \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}, \\ \mathbf{C}_{\mathrm{P}} &= \begin{bmatrix} \mathbf{s}_1^{\mathsf{T}}\overline{\mathbf{A}} + \mathbf{s}_0^{\mathsf{T}}\mathbf{F}_0 + [\mathbf{0} \mid a\mathbf{s}_0^{\mathsf{T}}] + \mathbf{E}_{\mathrm{P},0} \\ \overline{\mathbf{A}} + \mathbf{E}_{\mathrm{P},1} \end{bmatrix}, \\ \mathsf{aux}_{\mathsf{S}} &= (\mathbf{X}, \mathsf{iO}(\mathsf{Vfy}_{\mathbf{s}_1,\mathbf{X},\mathbf{M}})). \end{aligned}$$

Here $(\mathbf{B}, \mathsf{td}_{\mathbf{B}}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$ and $\mathbf{E}_{\mathrm{B}} \leftarrow \mathcal{D}_{\sigma_{\mathrm{pre}}}^{t\times m}$. The remaining values are generated by the sampler.

$\boxed{\mathcal{H}_1}$ Same as Hybrid $\mathcal{H}_0$, but replace $\mathbf{C}_{\mathrm{B}}$ with an independent uniform $\mathbf{U}_{\mathrm{B}} \leftarrow \mathbb{Z}_q^{t\times m}$.

$\boxed{\mathcal{H}_2}$ Same as Hybrid $\mathcal{H}_1$, but omit $\mathbf{s}_0^{\mathsf{T}}\mathbf{F}_0 + [\mathbf{0} \mid a\mathbf{s}_0^{\mathsf{T}}]$ from the first row of $\mathbf{C}_{\mathrm{P}}$.

$\boxed{\mathcal{H}_3}$ Same as Hybrid $\mathcal{H}_2$, but replace the hardcoded $\mathbf{M} = \mathbf{s}_0^{\mathsf{T}}\mathbf{X} + \mathbf{e}_{\mathrm{M}}^{\mathsf{T}}$ with an independent uniform $\mathbf{U}_{\mathrm{M}} \leftarrow \mathbb{Z}_q^{1\times d}$.

$\boxed{\mathcal{H}_4}$ Same as Hybrid $\mathcal{H}_3$, but replace the obfuscation by $\mathsf{iO}(C_\perp)$, where $C_\perp$ is the identically zero circuit with the same input domain and padded size as Vfy.

$\boxed{\mathcal{H}_5}$ Same as Hybrid $\mathcal{H}_4$, but add an independent $\mathbf{E}_{\mathrm{A}} \leftarrow \mathcal{D}_{\sigma_{\mathrm{s}}}^{1\times \ell}$ to the first row of $\mathbf{C}_{\mathrm{P}}$.

$\boxed{\mathcal{H}_6}$ Same as Hybrid $\mathcal{H}_5$, but replace $\mathbf{s}_1^{\mathsf{T}}\overline{\mathbf{A}} + \mathbf{E}_{\mathrm{A}}$ with an independent uniform row. Thus $\mathbf{C}_{\mathrm{P}}$ is an independent uniform $\mathbf{U}_{\mathrm{P}} \leftarrow \mathbb{Z}_q^{t\times \ell}$.

$\boxed{\mathcal{H}_7}$ Same as Hybrid $\mathcal{H}_6$, but restore $\mathsf{iO}(\mathsf{Vfy}_{\mathbf{s}_1,\mathbf{X},\mathbf{U}_{\mathrm{M}}})$.

$\boxed{\mathcal{H}_8}$ Same as Hybrid $\mathcal{H}_7$, but restore $\mathbf{M} = \mathbf{s}_0^{\mathsf{T}}\mathbf{X} + \mathbf{e}_{\mathrm{M}}^{\mathsf{T}}$ in the circuit. This is the RHS distribution of the precondition in Definition 6.

We now state the indistinguishability of each adjacent pair of hybrids.

Claim 26. *The distributions in Hybrids $\mathcal{H}_0$ and $\mathcal{H}_1$ are statistically indistinguishable.*

Proof. The public-matrix guarantee of $\mathsf{TrapGen}$ (Lemma 1) allows replacing the hidden $\mathbf{B}$ by uniform. It appears in the view only through $\mathbf{C}_{\mathrm{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathrm{B}}$. The nonzero Gaussian bound and the Gaussian tail bounds (Corollary 2 and lemma 4) give $\mathbf{s}_0 \neq \mathbf{0} \pmod q$ except with negligible probability. Hence $\mathsf{rank}_{\mathbb{Z}_q}(\mathbf{S}) = k + 1 = t$, so $\mathbf{C}_{\mathrm{B}}$ is uniform even conditioned on $\mathbf{S}$ and all sampler outputs. □

Claim 27. *The distributions in Hybrids $\mathcal{H}_1$ and $\mathcal{H}_2$ are statistically indistinguishable.*

Proof. By the Gaussian tail bounds (Lemma 4), the omitted integer row has norm at most

$$u = k\lambda^2\sigma_s\eta + a\lambda\sigma_s \le 2a\lambda\sigma_s$$

except with negligible probability. Condition on all sampling randomness other than the fresh error $\mathbf{E}_{P,0}$. The Gaussian flooding bound (Lemma 5), followed by averaging over the conditioning, bounds the statistical distance by

$$O\left(\frac{\ell u}{\sigma_{pre}}\right) + \mathsf{negl}(\lambda) = O(\ell\lambda\sigma_s 2^{-\lambda}) + \mathsf{negl}(\lambda) \le \mathsf{negl}(\lambda).$$

□

Claim 28. *Assuming* $\mathsf{LWE}_{k,\ell+k+\lambda,q,\sigma_s}$*, the distributions in Hybrids* $\mathcal{H}_2$ *and* $\mathcal{H}_3$ *are computationally indistinguishable.*

Proof. In Hybrid $\mathcal{H}_2$, the only remaining use of $\mathbf{s}_0$ is in the LWE sample $\mathbf{M} = \mathbf{s}_0^\intercal\mathbf{X} + \mathbf{e}_M^\intercal$ hardcoded in the circuit. Thus Gaussian-secret LWE (Lemma 11), with public matrix $\mathbf{X}$ and $d \le \ell$ samples, allows replacing $\mathbf{M}$ by $\mathbf{U}_M$. □

Claim 29. *Assuming null-iO security, the distributions in Hybrids* $\mathcal{H}_3$ *and* $\mathcal{H}_4$ *are computationally indistinguishable.*

Proof. In Hybrid $\mathcal{H}_3$, $\mathbf{U}_M$ is uniform and independent of $\mathbf{X}$. For any fixed $\mathbf{X}$, counting the choices of $\widehat{\mathbf{s}}_0 \in \mathbb{Z}_q^k$ and the bounded residual gives

$$\Pr\left[\exists\widehat{\mathbf{s}}_0 \in \mathbb{Z}_q^k : \left\|\mathbf{U}_M - \widehat{\mathbf{s}}_0^\intercal\mathbf{X}\right\|_\infty \le \beta_M\right] \le \frac{q^k \cdot (2\beta_M+1)^d}{q^d} = \left(\frac{(2\beta_M+1)^2}{q}\right)^k \le \mathsf{negl}(\lambda).$$

Every accepting input yields such a $\widehat{\mathbf{s}}_0$ in Step 3. Thus Vfy is identically zero with overwhelming probability, regardless of $\mathbf{s}_1$. Null-iO security then gives the claim. □

Claim 30. *The distributions in Hybrids* $\mathcal{H}_4$ *and* $\mathcal{H}_5$ *are statistically indistinguishable.*

Proof. By the Gaussian tail bounds (Lemma 4), $\|\mathbf{E}_A\|_\infty \le \lambda\sigma_s$ except with negligible probability. Conditioning on $\mathbf{E}_A$ and applying the Gaussian flooding bound (Lemma 5) to the independent fresh error $\mathbf{E}_{P,0}$ therefore bounds the statistical distance by

$$O\left(\frac{\ell\lambda\sigma_s}{\sigma_{pre}}\right) + \mathsf{negl}(\lambda) \le \mathsf{negl}(\lambda).$$

□

Claim 31. *Assuming* $\mathsf{LWE}_{k,\ell+k+\lambda,q,\sigma_s}$*, the distributions in Hybrids* $\mathcal{H}_5$ *and* $\mathcal{H}_6$ *are computationally indistinguishable.*

Proof. In Hybrid $\mathcal{H}_5$, the only remaining use of $\mathbf{s}_1$ is in $\mathbf{s}_1^\intercal\overline{\mathbf{A}} + \mathbf{E}_A$. Thus Gaussian-secret LWE (Lemma 11), with public matrix $\overline{\mathbf{A}}$ and ℓ samples, allows replacing this row by uniform. □

Claim 32. *Assuming null-iO security, the distributions in Hybrids* $\mathcal{H}_6$ *and* $\mathcal{H}_7$ *are computationally indistinguishable.*

Proof. Same as Claim 29. □

Claim 33. *Assuming* $\mathsf{LWE}_{k,\ell+k+\lambda,q,\sigma_s}$*, the distributions in Hybrids* $\mathcal{H}_7$ *and* $\mathcal{H}_8$ *are computationally indistinguishable.*

Proof. Same as Claim 28. □

Combining Claims 26 to 33, the distributions in Hybrids $\mathcal{H}_0$ and $\mathcal{H}_8$ are computationally indistinguishable, proving the theorem. □