IND-CPA-D and KR-D Security With Reduced Noise from the HintLWE Problem

Tabitha Ogilvie
tabitha.l.ogilvie@gmail.com

Royal Holloway, University of London, UK

Abstract. Approximate Homomorphic Encryption (AHE), introduced by Cheon et al. [20] offers a powerful solution for encrypting real-valued data by relaxing the correctness requirement and allowing small decryption errors. Existing constructions from (Ring) Learning with Errors achieve standard IND-CPA security, but this does not fully capture scenarios where an adversary observes decrypted outputs. Li and Micciancio [41] showed that when decryptions are passively leaked, these schemes become vulnerable to practical key recovery attacks even against honest-but-curious attackers. They formalise security when decryptions are shared with new notions of IND-CPA-D and KR-D security.
We propose new techniques to achieve provable IND-CPA-D and KR-D security for AHE, while adding substantially less additional decryption noise than the prior provable results. Our approach hinges on refined "game-hopping" tools in the bit-security framework, which allow bounding security loss with a lower noise overhead. We also give a noise-adding strategy independent of the number of oracle queries, removing a costly dependence inherent in the previous solution.
Beyond generic noise-flooding, we show that leveraging the recently introduced HintLWE problem [35] can yield particularly large security gains for AHE ciphertexts that are the result of "rescaling", a common operation in CKKS. Our analysis uses the fact that rescale-induced noise amounts to a linear "hint" on the secret to enable a tighter reduction to LWE (via HintLWE). In many practical parameter regimes where the rescaling noise dominates, our results imply an additional precision loss of as little as two bits is sufficient to restore a high level of security against passive key-recovery attacks for standard parameters.
Overall, our results enable a provably secure and efficient real-world deployment of Approximate Homomorphic Encryption in scenarios with realistic security requirements.

1 Introduction

Approximate Homomorphic Encryption, introduced in the CKKS scheme [20], differs from Homomorphic Encryption by relaxing the strict correctness requirement for decryptions. It has proven to be a very popular primitive in recent years,

This work was partially supported by EPSRC and the UK Government, Grant EP/S021817/1.

due to its efficiency and ability to handle floating point data [37,33,39,49,50]. In practice, Approximate Homomorphic Encryption is often built from assumptions about the hardness of the Learning with Errors (LWE) or Ring Learning with Errors (RLWE) problems [54,43]. From these assumptions, existing constructions for Approximate Homomorphic Encryption can be shown to meet the standard notion for secure encryption, Indistinguishability under Chosen Plaintext Attacks (IND-CPA).

However, in [41], the authors discuss that IND-CPA security is not necessarily sufficient to capture the security of *Homomorphic* Encryption, especially when the decryptor may share the resulting decryptions. Indeed, they consider a passive adversary who sees not only the result of encryptions and homomorphic evaluations, but also observes the result of decryptions. They formalize this idea by introducing IND-CPA-D (IND-CPA with decryption oracles) security, a natural extension to the IND-CPA notion which captures the extra information exposed via these decryptions.

In general, it is not true that a Homomorphic Encryption scheme being approximate implies that it would fail to achieve IND-CPA-D security. Unfortunately, the nature of the approximation in standard LWE based schemes means that the result of decryptions can be used to completely compromise the encryption security. In [41], the authors exhibit a full key recovery attack in the IND-CPA-D model, using the information exposed by decryptions.

Along with the IND-CPA-D model, the authors of [41] define the KR-D (Key Recovery with decryption oracles) security notion, which challenges a passive adversary who observes decryptions to calculate the secret key. KR-D is a weaker model than IND-CPA-D security, since any adversary that can extract the key can definitely decrypt ciphertexts. Many of the attacks which have gained attention in the IND-CPA-D model are actually KR-D attacks [8,16,17,31].

Heuristic mitigations to prevent these passive attacks were implemented in many popular (Approximate) Homomorphic Encryption libraries following the publication of [41]: an overview is given in [19]. However, a mitigation which achieves provable security for Approximate Homomorphic Encryption wasn't provided until [42], where the authors consider adding noise to the result of decryptions to obscure any potential leakage. Their results require adding a significant amount of noise to achieve IND-CPA-D security, requiring either reducing the output precision or choosing larger parameters: to the best of our knowledge, no library has implemented adding the amount of noise required to prevent a significant loss in security between the IND-CPA-D and IND-CPA games.

1.1 Our Results

In this work, we greatly reduce the amount of noise necessary to achieve provable IND-CPA-D and KR-D bit security compared to the state of the art [42] for standard parameters [10], in some settings showing that only two additional bits of noise are sufficient to claim security against key recovery attacks.

To achieve this, we improve existing results for game hopping in the bit security framework, and find that we can reduce the variance required by a factor

of $16t$, where t is the number of permitted decryptions. This gives a precision gain of $2 + \frac{1}{2}\log t$ bits, and crucially allows a parametrization which is independent of the number of permitted queries. We also show how we can argue for IND-CPA-D security from a related hardness assumption, the HintLWE problem [35], which reduces from LWE. We find that we can use a hybrid approach to argue security for ciphertexts which are the result of rescaling, a very common operation in CKKS applications. In settings where the noise added by rescaling dominates the rescaled noise [34,38,40,23], we find this approach can give 20 additional bits of precision for standard parameters. As part of this analysis, we give a reduction from the hardness of LWE to HintLWE in the bit security framework, which we find enables us to set better parameters than permitted by the original proof.

We additionally consider targeting provable security in the KR-D game directly. We give an analysis of game hopping in the bit security framework based on the Rényi divergence, and find that it is sufficient to add $O(\frac{1}{2}\log\kappa - \log l)$ additional bits of noise to achieve a security of $(\kappa - l)$ bits, where κ denotes the security of the scheme in the IND-CPA model, and l denotes a tolerable security loss. Using the HintLWE approach to achieve KR-D security, we find that in settings where the rescale noise dominates, it is sufficient to add 2 bits of noise to achieve 125 bit security for standard parameters.

1.2 Technical Overview and Detailed Contributions

Game Hopping and Bit Security We will show security inside the bit security framework, introduced in [48]. Our first contribution is novel hopping techniques in this framework based on new notions of distance between probability distributions. In more detail, we will "hop" between games which differ in a single oracle: in a game $\mathcal{G}_0$, the oracle outputs a sample from a distribution $\mathcal{X}$, whereas in $\mathcal{G}_1$, it outputs a sample from a distribution $\mathcal{Y}$. We will want to say the distributions $\mathcal{X}$ and $\mathcal{Y}$ are sufficiently "close" that the advantage of any adversary cannot be significantly higher in $\mathcal{G}_1$ than in $\mathcal{G}_0$. A standard notion for "closeness" of distributions is the statistical, or total variational, distance, but this is only one of many metrics. In [48], the authors show a general class of so-called λ-efficient metrics is compatible with jumping between distinguishing games in the bit security framework. We extend this result by giving a result specific to the Kullback-Leibler divergence which improves on the general result by a constant factor, as well as providing a result based on the R_∞ distance between games.

We additionally provide a generic result concerning game hopping based on the Rényi divergence in search games. This result is a function of the Rényi order a, the original security level κ, the acceptable security loss l, and the number of queries to the oracle, t. This explicit formulation enables us to select the a which minimises the noise flooding parameter in our application.

A further technique which allows us to reduce the amount of required noise is the observation that *oracles queries aren't free*. In this introduction, we give an explanation in terms of search games, where the advantage of an adversary

is just the probability of solving the search problem. Then, for any program $\mathcal{A}$,

$$\begin{aligned}\frac{\Pr[\mathcal{A} \text{ wins } \mathcal{G}_1]}{T(\mathcal{A})} &\leq \frac{\Pr[\mathcal{A} \text{ wins } \mathcal{G}_0]}{T(\mathcal{A})} + \frac{\Delta(\mathcal{G}_0, \mathcal{G}_1)}{T(\mathcal{A})} \\ &\leq \frac{\Pr[\mathcal{A} \text{ wins } \mathcal{G}_0]}{T(\mathcal{A})} + \frac{t\delta(\mathcal{X}\|\mathcal{Y})}{T(\mathcal{A})},\end{aligned}$$

for Δ the statistical distance, δ some divergence with certain properties, t the number of queries made to the oracle, and $T(\mathcal{A})$ a measure of the resources used by $\mathcal{A}$. Observe that $\mathcal{A}$ cannot make more queries than it has resources: in other words, $t \leq T(\mathcal{A})$. This means we can bound $\frac{\Pr[\mathcal{A} \text{ wins } \mathcal{G}_1]}{T(\mathcal{A})}$ using only a bound on $\frac{\Pr[\mathcal{A} \text{ wins } \mathcal{G}_0]}{T(\mathcal{A})}$ and a bound on $\delta(\mathcal{X}\|\mathcal{Y})$, and we can proceed independently of t. We use this observation in the IND-CPA-D setting to reduce the amount of added noise by $\frac{1}{2}\log t$ bits, and crucially has parameters which are independent of the number of shared decryptions, allowing for a more flexible deployment.

Noise Flooding Approach. The first route to IND-CPA-D and KR-D security, proposed in [42], is generic to all Approximate Homomorphic Encryption schemes. This approach attempts to add enough noise to the result of decryptions that we can complete a game hop that completely removes the decryption error. More specifically, we consider games in which the decryption oracle outputs $\mathbf{m}+\mathbf{e}+\epsilon$ and $\mathbf{m}+\epsilon$ respectively, for $\mathbf{e}$ the decryption error, and ϵ is drawn from a discrete or continuous Gaussian. Since these decryptions no longer contain any extra information, one can argue this second scheme is as secure as the original scheme in the IND-CPA setting.

The efficiency and security of this solution is therefore determined by the size of the added error ϵ. If ϵ is large, the precision of the result $\mathbf{m}$ is compromised, while if ϵ is small, the distance between the two games will be too large to make an effective game hop. For IND-CPA-D security, we use our result on game-hopping based on the Kullback-Leibler divergence to give a lower bound on the variance of the added noise ϵ of the following form

$$4E^2 2^{\kappa - l}$$

to achieve $(\kappa - l)$ bit security, whenever the IND-CPA game is κ bit secure, and E is an upper bound on the decryption error.

The KR-D game is a search game, so that we argue for security using the Rényi divergence to analyse the closeness of $\mathbf{m}+\mathbf{e}+\epsilon$ and $\mathbf{m}+\epsilon$. This allows adding significantly less noise to achieve the same security level. We give a result which is fully parametrisable in terms of the target security level – we believe this flexibility will prove useful in settings where there are different target security levels in the KR-D and KR settings. We find that noise with variance

$$\sigma^2 = \frac{2\kappa t E^2 \log e}{l^2}$$

is sufficient to achieve a bit security of $(\kappa - l)$, whenever the KR game is κ bit secure. The inverse square relationship with the security loss l is new to this

work, and allows implementors to sacrifice a few bits of security for a significant reduction in the required noise variance.

HintLWE Approach. When using the noise flooding approach, our results are completely generic to the underlying HE scheme, only requiring the plaintext space is some subset of $\mathbb{R}^n$. However in practice, we want to achieve these security notions in the context of the (Ring) Learning with Errors ((R)LWE) problem. Given a concrete construction, it may be possible to quantify exactly how the noise during encryption degrades the security of the construction.

To this end, we consider appealing to the hardness of the HintLWE problem, introduced in [35]. In this problem, the adversary is given access to an LWE sample, as well as some noisy linear hints in the secret:

$$(\mathbf{A}, \mathbf{As} + \mathbf{e} \mod Q, \mathbf{X}, \mathbf{Xs} + \mathbf{y}).$$

In the decision variant, the adversary is asked to distinguish the component $\mathbf{As} + \mathbf{e}$ from uniform, and we can naturally extend this to give a search variant, where the adversary is challenged to find $\mathbf{s}$.

In general, the linear "hints" may result in a much easier problem than if only had access to the LWE sample, even solvable in polynomial time [9]. However, the authors of [35] give a reduction which shows that when $\mathbf{s}$ and $\mathbf{y}$ are drawn from sufficiently wide Gaussians, the HintLWE problem is as hard as a related LWE problem. In this work, we argue for security via a reduction from LWE, rather than estimating the security of HintLWE directly. To stay within the bit security framework, we reframe the reduction as a series of game hops, which additionally means the security lost during the reduction is no longer simply a function of statistical distance. We instead use the infinite order Rényi divergence R_∞ to give a more efficient reduction.

Our goal is to argue that the decryptions in (R)LWE based Approximate Homomorphic Encryption schemes resemble the "hint" component of a HintLWE sample. From here, IND-CPA-D security would follow from security of the decision variant, while KR-D security would follow from the search variant. Unfortunately, we cannot argue that all decryption noise has the form $\mathbf{Xs}$.[1] Indeed, when we evaluate the product of two messages homomorphically, the resulting noise is a non linear combination of both the noise before the product and the encrypted messages. However, the CKKS scheme has an operation called Rescale which reduces the noise in a ciphertext, but adds a new noise which is a known linear function of the secret. More specifically, rescaling a ciphertext $(\mathbf{A}, \mathbf{b} = \mathbf{As} + \mathbf{e} + \mathbf{m})$ by q returns a ciphertext encrypting $\left\lceil \frac{m}{q} \right\rfloor$ which has noise

$$\underbrace{\frac{1}{q}([\mathbf{m}]_q + \mathbf{e})}_{\mathbf{e}^{(1)}} + \underbrace{\mathbf{A}_q\mathbf{s} - \mathbf{b}_q}_{\mathbf{e}^{(2)}}$$

[1] In fact, it would be sufficient for all noise to have the form $\mathbf{X}_1\mathbf{s} + \mathbf{X}_2\mathbf{e}$, as in the original description of the HintLWE problem in [35].

for $\mathbf{A}_q = [\mathbf{A}]_q/q, \mathbf{b}_q = [\mathbf{b}]_q/q$. As before, it's unclear the impact term $\mathbf{e}^{(1)}$ has on security in either the key recovery or distinguishability game. However, $\mathbf{e}^{(2)}$ is a known linear function of the secret: in other words, if this was the only noise, we could argue security from the HintLWE problem applied to $\mathbf{X} = \mathbf{A}_q$. We therefore give a hybrid approach, which starts by jumping between games in which the decryption oracle outputs:

$$(\mathbf{e}^{(1)} + \mathbf{e}^{(2)} + \epsilon) \mapsto (\mathbf{e}^{(2)} + \epsilon)$$

and from here, arguing from the security of the related LWE instance. In settings where the noise introduced by rescaling dominates, this proves to be a useful approach.

1.3 Related Work

The most closely related work to ours is [42], which also considers provable bit security in the IND-CPA-D setting, showing that is is sufficient to add noise with variance $16tE^2 2^{\kappa-l}$ to achieve $(\kappa - l)$ bit security in the IND-CPA-D setting. This work also introduces a new notion of (c, s) security in the same game, which enables adding less noise when $s < c$. However, the condition $s < c$ gives that this is a strictly weaker notion of security than c-bit security, as shown in Lemma 10 of [42].

The bit security framework we use is introduced in [48], but alternative frameworks can be found in [57,58]. Using the Rényi divergence to argue for security in search games is considered in [6,36,53], and the difficulty of extending this to decision games is documented in [11,51].

The utility of modulus switching and rescaling in eliminating noise which might otherwise compromise the security of the scheme was explored in [27]. A popular variant of the CKKS scheme was proposed in [34], where rescaling noise dominates all other sources of noise: this is discussed in Section 5.1 of [42]. The importance of noise from rescaling during CKKS bootstrapping is detailed in [38,40], and a trend of modulus switching noise dominating other noise was observed experimentally in [23].

The HintLWE problem was introduced in [35], where it was shown to be as hard as a related instance of LWE. In that work, this was used to construct a performant Proof of Knowledge from lattices with hardness based on Module Learning with Errors. The hardness of HintLWE was used in a subsequent work [29] to give a provably masked hash-and-sign algorithm. An alternative line of work [9,24,25,45] does not use a reduction from LWE, but instead estimates the hardness of LWE with Hints directly: most of these works consider a greater variety of "hint" beyond the noisy linear hint considered in this work.

The authors of [8] also consider the hardness of recovering the secret key when given access to noisy decryptions of rescaled ciphertexts. This section of our work differs from theirs in the following respects. First, we consider a provable reduction from the hardness of LWE, rather than directly estimating the difficulty of HintLWE at various noise levels. Second, we do not model the term

$\mathbf{e}$ as being zero, as this could never be the case. Lastly we do not model the "hint" matrix $[\mathbf{A}]_q$ (or $[a]_q$) in the ring setting) as having Gaussian entries with the same variance as a uniform modulo q, as we are unable to prove or verify that this property holds, particularly given adversarially selected circuits.

Lastly, although we focus on CKKS as an Approximate Homomorphic Encryption scheme, we emphasise that it is not the case that CKKS and Approximate Homomorphic Encryption are the same concept. Indeed, a recent line of work [4,5,21,28] has considered using CKKS to conduct exact computation, so that noise can be removed during decryption, while another line of work [16,17,31] argues that schemes previously considered "exact" can become approximate, either due to aggressive parameter setting or adversarially selected circuits.

1.4 Paper Outline

In Section 2 we give necessary technical preliminaries. In Section 3 we revisit the security loss due to game hops within the bit security framework, and then apply this to noise flooding for IND-CPA-D and KR-D security in Section 4. We propose an alternative approach based on the HintLWE problem in Section 5. We conclude by presenting concrete parameters and a comparison between the two approaches in Section 6. Some technical details are deferred to the appendices.

2 Preliminaries

2.1 Notation

We write ln() for the natural (base e) logarithm, and log() for the base 2 logarithm.

We represent $\mathbb{Z}_q$ via the centred interval $[-q/2, ..., q/2)$. We write $\mathcal{R}$ for the standard power of two cyclotomic ring $\mathbb{Z}[X]/(X^n+1)$ with n a power of two, and $\mathcal{R}_q$ for $\mathbb{Z}_q[X]/(X^n+1)$. We use $\lceil \cdot \rfloor$ for the rounding operation (rounding half up), and extend this to polynomials coefficient wise. Similarly, we will use $[\cdot]_q$ to denote modular reduction into $[-q/2, ..., q/2)$, and extend this to polynomials coefficient wise, and matrices entry wise.

We write $(X_i)_i$ for the vector $(X_0, X_1, ..., X_n)$, and use column notation for vectors. We denote the infinity norm of a vector via $\|\mathbf{v}\|_\infty$, and extend this for polynomials a by using $\|a\|_\infty$ to denote the infinity norm over the coefficient vector.

For a matrix $\mathbf{M}$, we will use $\|\mathbf{M}\|_2$ to denote the spectral norm of $\mathbf{M}$, which is given by the largest singular value of $\mathbf{M}$.

2.2 Power of Two Cyclotomic Rings and Negacyclic Matrices

We consider homomorphic schemes defined over a ring $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^n+1)$. Elements of this ring can be represented uniquely by a polynomial of degree

at most $n-1$ and coefficients in $\mathbb{Z}_Q$. Observe further that multiplying by a polynomial $m \in \mathcal{R}_Q$ can be calculated by multiplying the coefficient vector by the following matrix modulo Q:

$$U_m = \begin{bmatrix} m_0 & -m_{n-1} & -m_{n-2} & \dots & -m_1 \\ m_1 & m_0 & -m_{n-1} & \dots & -m_2 \\ m_2 & m_1 & m_0 & \dots & -m_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ m_{n-1} & m_{n-2} & m_{n-3} & \dots & m_0 \end{bmatrix}$$

this matrix is negacyclic, or skew circulant, meaning that each row is a right shift of the one before, negating the new first element. For more details on negacyclic matrices, we refer the reader to [26]. We will use the following facts.

Lemma 1 (3.2.1 of [26]). *Suppose that U, V are real negacyclic matrices. Then U^T, UV and U^{-1} are all negacyclic.*

Lemma 2 (3.2.1 of [26]). *All negacyclic matrices are simultaneously diagonalizable, with the eigenvalues of the negacyclic matrix U_m for $m \in \mathcal{R}$ given by $m\left(\exp\left(\frac{(2j+1)\pi i}{n}\right)\right)$ for $j = 0, ..., n-1$.*

Elsewhere in the Ring Learning with Errors literature, the vector of these eigenvalues are also known as the canonical embedding of the polynomial m: see [44] for further details.

We will also use the following result, which bounds the largest eigenvalue of a polynomial in $\mathcal{R}$ in terms of its largest coefficient.

Lemma 3 (Theorem 1 and Lemma 4 [22]). *Let $m \in \mathcal{R}$. Then*

$$\left| m\left(\exp\left(\frac{(2j+1)\pi i}{n}\right)\right)\right| \leq \sqrt{\left(\frac{2n}{\pi}\right)^2 + 1}\, \|m\|_\infty$$

for $j = 0, ..., n-1$. In canonical embedding notation, this can be written $\|m\|_\infty^{\text{can}} \leq \sqrt{\left(\frac{2n}{\pi}\right)^2 + 1}\, \|m\|_\infty$.

2.3 Gaussian Distributions

Central to our work are the continuous and discrete Gaussian distributions. These are parametrised by a vector $\mu \in \mathbb{R}^n$ and a symmetric positive definite matrix $\Sigma \in \mathbb{R}^{n \times n}$. We denote the continuous Gaussian distribution via $\mathcal{N}(\mu, \Sigma)$, and the discrete Gaussian with $\mathcal{D}_{\mu,\Sigma}$, and $\mathcal{D}_\Sigma$ in the case $\mu = 0$.

We use the unnormalized version of these distributions[2], so that both distributions have probability mass (resp. density) directly proportional to

$$\rho_{\mu,\Sigma}(x) \propto \exp\left(-\frac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu)\right)$$

in the case of the discrete Gaussian, the support of x is some discrete subset of $\mathbb{R}^n$: in our work, this will always be $\mathbb{Z}^n$.

2.4 Divergences between Probability Distributions

When arguing for security, it is useful to claim that two distributions are "close", regardless of how many resources the adversary uses. In this work, we will use the following notions of divergence between probability distributions.

Definition 1 (Statistical Distance). *Let $\mathcal{P}, \mathcal{Q}$ be random variables with shared support. Then the Statistical Distance (or Total Variation Distance) is defined as*

$$\Delta(\mathcal{P}, \mathcal{Q}) = \sup_{E \subseteq \mathrm{Supp}(\mathcal{P})} |\mathcal{P}(E) - \mathcal{Q}(E)|$$

We will also use the Kullback-Leibler divergence, defined as follows.

Definition 2 (KL Divergence). *The Kullback-Leibler divergence between two random variables $\mathcal{P}, \mathcal{Q}$ with shared support is defined as*

$$D(\mathcal{P}\|\mathcal{Q}) = \sum_{x \in \mathrm{Supp}(\mathcal{P})} \mathcal{P}(x) \ln \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}$$

whenever $\mathcal{P}, \mathcal{Q}$ are discrete random variables, or

$$D(\mathcal{P}\|\mathcal{Q}) = \int_{x \in \mathrm{Supp}(\mathcal{P})} \mathcal{P}(x) \ln \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}$$

for continuous $\mathcal{P}, \mathcal{Q}$.

Definition 3 (Rényi Divergence [53]). *Let $\mathcal{P}, \mathcal{Q}$ be two distributions such that $\mathrm{Supp}(\mathcal{P}) \subset \mathrm{Supp}(\mathcal{Q})$. For $a > 0, a \neq 1$ we define the Rényi divergence of order a by*

$$R_a(\mathcal{P}\|\mathcal{Q}) = \left(\sum_{x \in \mathrm{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)^a}{\mathcal{Q}(x)^{a-1}}\right)^{\frac{1}{a-1}}$$

whenever $\mathcal{P}, \mathcal{Q}$ are discrete random variables, or

$$R_a(\mathcal{P}\|\mathcal{Q}) = \left(\int_{x \in \mathrm{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)^a}{\mathcal{Q}(x)^{a-1}}\right)^{\frac{1}{a-1}}$$

for continuous $\mathcal{P}, \mathcal{Q}$.

[2] We can convert to the normalized version via $\Sigma \mapsto 2\pi\Sigma$: we refer to [7], footnote 7, for details.

This definition can be extended to $a = \infty$ by taking a limit as follows.

Definition 4 (Rényi Divergence for $a = \infty$). *Let $\mathcal{P}, \mathcal{Q}$ be two distributions such that $\mathrm{Supp}(\mathcal{P}) \subset \mathrm{Supp}(\mathcal{Q})$. Define*

$$R_\infty(\mathcal{P}\|\mathcal{Q}) = \sup_{x \in \mathrm{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}.$$

We are using the definition standard to the cryptographic literature [53,11,6]: elsewhere, the Rényi divergence is defined as the (natural) logarithm of this value. We give properties of the KL and Rényi divergences in Section A.1.

2.5 Bit Security

We will consider security within the bit security framework of [48]; we review the components we will use below, following [42].

Definition 5 (Indistinguishability Game). *Let $\{\mathcal{D}^0_\theta\}_\theta, \{\mathcal{D}^1_\theta\}_\theta$ be two distribution ensembles. In the indistinguishability game, a challenger chooses $b \leftarrow \mathcal{U}(\{0,1\})$, and provides the adversary oracle access to $\mathcal{D}^b_\theta$. The adversary wins the game if it outputs $b' = b$.*

Definition 6 (Indistinguishability Game Bit Security). *Given an adversary $\mathcal{A}$ playing an indistinguishability game $\mathcal{G}$, define:*

- *the probability $\mathcal{A}$ has an output, $\alpha^{\mathcal{A}} = \Pr[\mathcal{A} \neq \perp]$;*
- *the conditional success probability of $\mathcal{A}$, $\beta^{\mathcal{A}} = \Pr[b' = b | \mathcal{A} \neq \perp]$,*

with these probabilities taken over the randomness of both the adversary and the indistinguishability game. From here, we define the advantage of $\mathcal{A}$ as

$$\mathsf{adv}^{\mathcal{A}} = \alpha^{\mathcal{A}}(2\beta^{\mathcal{A}} - 1)^2.$$

The bit security of the indistinguishability game is $\min_{\mathcal{A}} \log_2 \frac{T(\mathcal{A})}{\mathsf{adv}^{\mathcal{A}}}$, where $T(\mathcal{A}) \geq 1$ measures the running time of $\mathcal{A}$.

Crucially, within this framework advantage in distinguishing games is *not* a probability.

Definition 7 (Search Game). *Let $\{\mathcal{D}^x_\theta\}_\theta$ be a distribution ensemble parametrized by an n-bit value x. In the search game, a challenger chooses an n-bit $x \leftarrow \mathcal{D}_X$, and provides the adversary oracle access to $\mathcal{D}^x_\theta$. The adversary wins the game if it outputs $x' = x$.*

Definition 8 (Search Game Bit Security). *Given an adversary $\mathcal{A}$ playing a search game $\mathcal{G}$, define:*

- *the probability $\mathcal{A}$ has an output, $\alpha^{\mathcal{A}} = \Pr[\mathcal{A} \neq \perp]$;*
- *the conditional success probability of $\mathcal{A}$, $\beta^{\mathcal{A}} = \Pr[x' = x | \mathcal{A} \neq \perp]$,*

with these probabilities taken over the randomness of both the adversary and the search game. From here, we define the advantage of $\mathcal{A}$ as

$$\mathsf{adv}^{\mathcal{A}} = \alpha^{\mathcal{A}}\beta^{\mathcal{A}}.$$

The bit security of the search game is $\min_{\mathcal{A}} \log_2 \frac{T(\mathcal{A})}{\mathsf{adv}^{\mathcal{A}}}$*, where* $T(\mathcal{A}) \geq 1$ *measures the running time of* $\mathcal{A}$*.*

We will also use the following lemma, which closely corresponds to either Lemma 5 of [42] or Lemma 3.2 of [47].

Lemma 4. *Consider a game* $\mathcal{G}^{\mathcal{P}}$ *with an oracle that accepts an input* θ *and outputs a sample from a distribution* $\mathcal{P}_\theta$*. Then for an adversary* $\mathcal{A}$ *playing either* $\mathcal{G}^{\mathcal{P}}$ *or* $\mathcal{G}^{\mathcal{Q}}$ *which makes at most* t *calls to the oracle,*

$$D(\mathcal{G}^{\mathcal{P}} \| \mathcal{G}^{\mathcal{Q}}) \leq t \max_\theta D(\mathcal{P}_\theta \| \mathcal{Q}_\theta), \tag{1}$$

$$R_a(\mathcal{G}^{\mathcal{P}} \| \mathcal{G}^{\mathcal{Q}}) \leq \left(\max_\theta D(\mathcal{P}_\theta \| \mathcal{Q}_\theta) \right)^t. \tag{2}$$

Proof. The proof for the KL divergence corresponds to applying sub-additivity. In more detail, let the queries, response pairs from the oracle be given by $((\theta_i, X_i))_i$ and $((\theta'_i, X'_i)_i$ where $X_i \sim \mathcal{P}_{\theta_i}$ and $X'_i \sim \mathcal{Q}_{\theta_i}$ respectively. Then

$$\begin{aligned}
D(\mathcal{G}^{\mathcal{P}} \| \mathcal{G}^{\mathcal{Q}}) &= D(((\theta_i, X_i))_i \| ((\theta'_i, X'_i))_i) \\
&\leq \sum_i \max_x D((\theta_i, X_i) | ((\theta_j, X_j))_{j<i} = x \| (\theta'_i, X'_i) | ((\theta'_j, X'_j))_{j<i} = x) \\
&= \sum_i \max_x D((\theta_i, X_i) | ((\theta_j, X_j))_{j<i} = x \| (\theta_i, X'_i) | ((\theta'_j, X'_j))_{j<i} = x) \\
&\leq \sum_i \max_\theta D(\mathcal{P}_\theta \| \mathcal{Q}_\theta) \\
&= t \max_\theta D(\mathcal{P}_\theta \| \mathcal{Q}_\theta).
\end{aligned}$$

The proof for the Rényi divergence follows in the same way, using sub-multiplicativity rather than sub-additivity.

The following lemma enables lower bounding bit security via a hybrid argument. Our result is 1 bit tighter than given in the literature – this can be obtained using the existing proof technique, as detailed in Section A.2.

Lemma 5 (Lemma 2 of [48]). *Let* $\mathcal{H}_i$ *be a sequence of* k *distributions, and let* $\mathcal{G}_{ij}$ *be the indistinguishability game instantiated with* $\mathcal{H}_i$ *and* $\mathcal{H}_j$*. Let* C *be a fixed constant. Let* $\epsilon_{ij} = \max_{\mathcal{A}} \mathsf{adv}^{\mathcal{A}}$ *over all adversaries with* $\mathcal{A}$ *with* $T(\mathcal{A}) \leq C$*. Then* $\epsilon_{1k} \leq \frac{3}{2} k \sum \epsilon_{i,i+1}$*.*

2.6 Homomorphic Encryption and Security Notions

We briefly review Homomorphic Encryption and related security notions. For a more thorough treatment, we refer the reader to [42]. We outline the symmetric variant, as this more closely corresponds to the setting we will consider later.

Definition 9 (HE Scheme). *A (symmetric key) homomorphic encryption scheme with circuit, secret key, evaluation key, plaintext, and ciphertext spaces $\mathcal{L}$, $\mathcal{S}$, $\mathcal{EVK}$, $\mathcal{M}$, $\mathcal{C}$ respectively, consists of a tuple of four algorithms:*

$$\begin{aligned}\mathsf{KeyGen} &: 1^{\mathbb{N}} \to \mathcal{EVK} \times \mathcal{SK}\\ \mathsf{Enc} &: \mathcal{M} \times \mathcal{SK} \to \mathcal{C}\\ \mathsf{Dec} &: \mathcal{C} \times \mathcal{SK} \to \mathcal{M}\\ \mathsf{Eval} &: \mathcal{EVK} \times \mathcal{L} \times \mathcal{C} \to \mathcal{C}\end{aligned}$$

The idea of a homomorphic encryption scheme is that the evaluation function Eval can evaluate any valid circuit $C \in \mathcal{L}$ on ciphertexts and produce a ciphertext which encrypts the result if the circuit was applied to the underlying messages. This is formalised by the notion of correctness.

Definition 10 (Correctness (Definition 9 on [42])). *An HE scheme $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ is* correct *if for all $m_1, ..., m_k \in \mathcal{M}$, $C \in \mathcal{L}$, $(\mathsf{evk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$ we have that*

$$\mathsf{Dec}\left(\mathsf{Eval}(\mathsf{evk}, \mathsf{Enc}(m_1, \mathsf{sk}), ..., \mathsf{Enc}(m_k, \mathsf{sk})), \mathsf{sk}\right) = C(m_1, ..., m_k)$$

In some settings however, perfect correctness may not be necessary: the CKKS scheme [20] for example permits a relaxation of this requirement, introducing a notion of "Approximate Homomorphic Encryption". For normed message spaces $\mathcal{M}$, we can understand the quality of this approximation as the distance between the 'true' result $C(m_1, ..., m_k)$ and the result obtained after decryption.

Definition 11 (Ciphertext Error (Definition 10 of [42])). *Let Π be an HE scheme with message space $\mathcal{M}$ which admits a norm $\|\cdot\|$. For any ciphertext ct, secret key sk, and message $m \in \mathcal{M}$, the* ciphertext error *of a tuple $(\mathsf{ct}, m, \mathsf{sk})$ is defined via*

$$\mathsf{Error}(\mathsf{ct}, m, \mathsf{sk}) = \|\mathsf{Dec}(\mathsf{ct}, \mathsf{sk}) - m\|$$

We extract from Definition 11 of [42] the notion of a static noise estimate.

Definition 12 (Static Noise Estimate). *Let Π be an HE scheme with message space $\mathcal{M}$ which admits a norm $\|\cdot\|$. We say $E \in \mathbb{R}$ is a* static noise estimate *for a circuit C if for all $m_1, ..., m_k \in \mathcal{M}$, for all $(\mathsf{evk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$, we have that if*

$$m = C(m_1, ..., m_k), \quad \mathsf{ct} = \mathsf{Eval}_{\mathsf{evk}}(C, \mathsf{Enc}_{\mathsf{sk}}(m_1), ..., \mathsf{Enc}_{\mathsf{sk}}(m_k)),$$

then

$$\mathsf{Error}(\mathsf{ct}, m, \mathsf{sk}) \leq E.$$

This notion is static (rather than dynamic) in that it does not depend on the particular values of the messages $m_1, ..., m_k$. Looking ahead, this is an important requirement as it means these noise estimates are publicly computable.

We will consider four separate security notions for homomorphic encryptions schemes: two in which the adversary must determine if the encryption oracle is encrypting on the left or right (IND-CPA and IND-CPA-D), and two where the adversary must extract the secret key sk (KR and KR-D). The difference between IND-CPA and IND-CPA-D (as well as KR and KR-D) is that the adversary is given access to a restricted decryption oracle: the adversary can request decryptions for which the underlying message is the same on both the left and the right. We formalise these games in Section A.3, following [41,42].

It was observed in [41] that approximate HE schemes based on lattices lose security when we give the adversary access to standard decryptions in the IND-CPA-D and KR-D games. In this work, we consider mitigating this vulnerability by post processing decryptions by adding a sample from an appropriate centred Gaussian. We follow [42] in using the notation $\tilde{\Pi}$ to refer to the scheme obtained by applying this post processing to the result of decryption in scheme Π. We give more detail in Algorithm 1.

Algorithm 1 Decryption Function for the Scheme $\tilde{\Pi}$

Input: HE scheme $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ with message space $\mathcal{M} \subset \mathbb{R}^n$.
$\mathsf{Dec}(\mathsf{ct}_i, \mathsf{sk}, \sigma_i^2) :=$
  $\mathsf{m} \leftarrow \mathsf{Dec}(\mathsf{ct}_i, \mathsf{sk})$
  $\epsilon \leftarrow \mathcal{D}^n_{\sigma_i^2}$
  return $\mathsf{m} + \epsilon$

2.7 LWE and HintLWE

We will justify the security of our proposal by appealing to the security of the Learning with Errors (LWE) problem, which is defined as follows.

Definition 13 (The LWE Problem (Decision)). *Let m, n, Q be positive integers, and let $\chi_{\mathbf{s}}$, $\chi_{\mathbf{e}}$ be distributions over $\mathbb{Z}_Q^n$ and $\mathbb{Z}_Q^m$ respectively. The LWE problem, denoted by $\mathsf{LWE}_{m,n,Q,\chi_{\mathbf{s}},\chi_{\mathbf{e}}}$ asks the adversary $\mathcal{A}$ to distinguish between the following cases:*

1. $(\mathbf{A}, \mathbf{As} + \mathbf{e} \mod Q)$ *for* $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_Q^{m \times n}), \mathbf{s} \leftarrow \chi_{\mathbf{s}}, \mathbf{e} \leftarrow \chi_{\mathbf{e}}$
2. $(\mathbf{A}, \mathbf{u})$ *for* $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_Q^{m \times n}), \mathbf{u} \leftarrow \mathcal{U}(\mathbb{Z}_Q^m)$.

We can extend this definition in the natural way to the module setting by replacing $\mathbb{Z}_Q$ with some ring $\mathcal{R}$ to give the Module Learning with Errors problem. In the case $d = 1$, so that the secret $\mathbf{s}$ consists of a single ring element s, we call this the Ring Learning With Errors problem (RLWE).

We can define a search version of these problems by challenging the adversary to find the secret vector $\mathbf{s}$ or secret ring element s.

We can use the LWE problem to build a homomorphic encryption scheme [14]. We refer to a scheme which uses plain LWE for encryption as LWE-HE: in more detail, the encryption oracle for this scheme returns

$$(\mathbf{A}, \mathbf{As} + \mathbf{e} + \mathbf{m}^{(b)}).$$

We additionally define the HintLWE problem, originally introduced in Definition 7 of [35]. There, the exact characterisation was inspired by considering the transcripts in lattice based proof of knowledge protocols. We slightly modify the definition to more naturally match the IND-CPA-D setting: we detail these modifications below.

Definition 14 (The HintLWE Problem (Decision)). *Let m, n, t, Q be positive integers, and let $\chi_{\mathbf{s}}$, $\chi_{\mathbf{e}}$ and ξ_i be distributions over $\mathbb{Z}_Q^n$, $\mathbb{Z}_Q^m$ and $\mathbb{Z}^{k_i}$ respectively for $i = 1, ..., t$. The HintLWE problem, denoted by $\mathsf{HintLWE}^{\xi,X,t}_{m,n,Q,\chi_{\mathbf{s}},\chi_{\mathbf{e}}}$, asks the adversary $\mathcal{A}$ to distinguish between the following two cases:*

1. $(\mathbf{A}, \mathbf{As}+\mathbf{e} \mod Q, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s}+\mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s}+\mathbf{y}_{t-1})$ *for* $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_Q^{m\times n})$, $\mathbf{s} \leftarrow \chi_{\mathbf{s}}, \mathbf{e} \leftarrow \chi_{\mathbf{e}}, \mathbf{y}_i \leftarrow \xi_i$, *and each* $\mathbf{X}_i \in \mathbb{R}^{k_i \times n}$ *is some function of (at most)* $\mathbf{A}$.
2. $(\mathbf{A}, \mathbf{u}, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s} + \mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s} + \mathbf{y}_{t-1})$ *for* $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_Q^{m\times n}), \mathbf{u} \leftarrow \mathcal{U}(\mathbb{Z}_Q^m), \mathbf{s} \leftarrow \chi_{\mathbf{s}}, \mathbf{e} \leftarrow \chi_{\mathbf{e}}, \mathbf{y}_i \leftarrow \xi_i$, *and each* $\mathbf{X}_i \in \mathbb{R}^{k_i \times n}$ *is some function of (at most)* $\mathbf{A}$.

We can understand this definition as the LWE problem with parameters m, n, Q, $\chi_{\mathbf{s}}, \chi_{\mathbf{e}}$ but where the adversary also has access to noisy linear hints on the secret, $\mathbf{Xs} + \mathbf{y}$. The parameter t roughly corresponds to how many hints we have access to, while ξ_i determines how noisy the hints are.

Our definition differs from [35] in the following respects. First, for clarity we introduce the problem with hints on the secret $\mathbf{s}$ only, rather than on the extended secret $(\mathbf{s}, \mathbf{e})$. This is for convenience in our application, where we find that our hints are on the secret only. Secondly, we allow the matrices $\mathbf{X}_i$ to be generic, and the noise $\mathbf{y}$ to vary with i: we will use this detail to add noise tuned to the specific matrix $\mathbf{X}_i$.

When $\chi_{\mathbf{s}}$ and $\chi_{\mathbf{e}}$ are spherical discrete or continuous Gaussians with variances $\sigma_{\mathbf{s}}^2$ and $\sigma_{\mathbf{e}}^2$ respectively, we will refer instead to $\mathsf{HintLWE}^{\xi,X,t}_{m,n,Q,\sigma_{\mathbf{s}}^2,\sigma_{\mathbf{e}}^2}$.

We can extend this definition to define the $\mathsf{HintRLWE}^{\xi,X,t}_{\mathcal{R},m,n,q,\chi_s,\chi_e}$ problem over the ring $\mathcal{R}$ by letting $\mathbf{A} \in \mathcal{R}_q^m$, and $\mathbf{s} = s \in \mathcal{R}_q$, and considering the products $\mathbf{Xs}$ to be with respect to the coordinate vector for s.

We can define a search version of these problems by challenging the adversary to find the secret vector $\mathbf{s}$ or secret ring element s.

2.8 The Rescale Operation

The rescale operation in the CKKS scheme generalises the modulus switching operation for BGV and B/FV schemes [13,12,30]. Throughout, we will consider

the Residue Number System variant [18]. In this variant, ciphertexts are of the form $(a, b = as + m + e) \in \mathcal{R}_{Q_l} \times \mathcal{R}_{Q_l}$, where $Q_l = \prod_{i=0}^{l} q_i$ with the q_i prime. A rescale operation takes (a, b) with respect to Q_l and returns (a', b') that are with respect to Q_{l-1}, and is calculated as follows[3]:

$$(a', b') := \left(\left\lceil \frac{a}{q_l} \right\rfloor \mod Q_{l-1}, \left\lceil \frac{b}{q_l} \right\rfloor \mod Q_{l-1} \right) = \left(\frac{a - [a]_{q_l}}{q_l}, \frac{b - [b]_{q_l}}{q_l} \right)$$

For the purposes of the IND-CPA-D game, we consider the rescale operation as the gate $m \leftarrow \left\lceil \frac{m}{q_l} \right\rfloor$, i.e. our messages will have integer coefficients[4]. This gives that our noise has integer coefficients, and can be calculated as:

$$b' - a's - \left\lceil \frac{m}{q_l} \right\rfloor = \frac{e + [m]_{q_l}}{q_l} + \frac{[a]_{q_l}}{q_l} s - \frac{[b]_{q_l}}{q_l}$$

so that the rescale operations achieve both a reduction in noise and a reduction in message magnitude. Switching down to the minimal Q_l necessary is also beneficial for lowering bandwidth requirements [32].

We can extend the definition of rescaling to apply to LWE ciphertexts – in this case, rescaling $(\mathbf{A}, \mathbf{b} = \mathbf{As} + \mathbf{e} + \mathbf{m})$ by q_l returns a ciphertext $(\mathbf{A}', \mathbf{b}')$ encrypting $\left\lceil \frac{\mathbf{m}}{q_l} \right\rfloor$ with noise

$$\mathbf{b}' - \mathbf{A}'\mathbf{s} - \left\lceil \frac{\mathbf{m}}{q_l} \right\rfloor = \frac{\mathbf{e} + [\mathbf{m}]_{q_l}}{q_l} + \frac{[\mathbf{A}]_{q_l}}{q_l} \mathbf{s} - \frac{[\mathbf{b}]_{q_l}}{q_l}$$

2.9 Gaussians on Lattices

We recall the definition of the smoothing parameter of a lattice, as well as a tight bound on this quantity for $\mathbb{Z}^n$.

Definition 15 (Smoothing Parameter $\eta_\varepsilon(\Lambda)$ [52]). *Let $\Sigma \geq 0$ be a positive semi-definite matrix. We say that $\sqrt{\Sigma} \geq \eta_\varepsilon(\Lambda)$ if $\rho_{\Sigma^{-1}}(\Lambda^* \backslash \{\mathbf{0}\}) \leq \varepsilon$.*

Lemma 6 (Smoothing Parameter $\eta_\varepsilon(\mathbb{Z}^n)$ [46,52]). *For any $\epsilon > 0$,*

$$\eta_\varepsilon(\mathbb{Z}^n) \leq \frac{1}{\pi} \sqrt{\frac{\ln(2n(1 + 1/\varepsilon))}{2}}$$

Lemma 7 (Lemma 5 of [35]). *For a positive definite matrix Σ, $\sqrt{\Sigma} \geq \eta_\varepsilon(\mathbb{Z}^n)$ if $\left\| \Sigma^{-1} \right\|_2 \leq \eta_\varepsilon(\mathbb{Z}^n)^{-2}$.*

[3] we use that for $x \in \mathbb{Z}$, $\left\lceil \frac{x}{q} \right\rfloor = \frac{x - [x]_q}{q}$, which follows by definition of rounding half up and modular reduction into the centred interval.

[4] if we instead consider a rescale ciphertext as an encryption of $\frac{m}{q}$, our noise no longer has real coefficients. This situation requires flooding with a continuous (rather than discrete) Gaussian, but otherwise is very similar.

We will use the following convolution lemma for discrete Gaussians. For completeness, we provide a proof in Section A.4.

Lemma 8 (Simplified Convolution Lemma [52,35]). *Let $\Sigma_1, \Sigma_2 \in \mathbb{R}^{n\times n}$ be symmetric positive definite matrices such that $\Sigma_3^{-1} = \Sigma_1^{-1} + \Sigma_2^{-1}$ satisfies $\sqrt{\Sigma_3} \geq \eta_\varepsilon(\mathbb{Z}^n)$ for $0 < \varepsilon < 1$. Then for $\mu_1, \mu_2 \in \mathbb{R}^n$ the random variables*

$$\begin{aligned} Z = X_1 + X_2, \quad & X_1 \leftarrow \mathcal{D}_{\mu_1,\Sigma_1}, X_2 \leftarrow \mathcal{D}_{\mu_2,\Sigma_2} \\ Z', \quad & Z' \leftarrow \mathcal{D}_{\mu_1+\mu_2,\Sigma_1+\Sigma_2} \end{aligned}$$

satisfy, for $a \neq 1$,

$$\begin{aligned} D(Z\|Z'), D(Z'\|Z) &\in \left[\ln\frac{1-\varepsilon}{1+\varepsilon}, \ln\frac{1+\varepsilon}{1-\varepsilon}\right], \\ R_a(Z\|Z'), R_a(Z'\|Z) &\in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right]. \end{aligned}$$

3 Game Hop Bit Security

In this section, we give a variety of new results that enable bounding the loss in security when we replace a distribution $\mathcal{P}$ with a similar distribution $\mathcal{Q}$ inside of an oracle.

3.1 Indistinguishability Games

The following theorem resembles Theorem 8 of [48] (restated as Theorem 1 in [42]) for arbitrary λ efficient measures. However, by targeting the KL divergence directly, we are able to give a tighter bound.

Theorem 1. *Consider a distinguishing game $\mathcal{G}^{\mathcal{P}}$ in which the adversary has black box access to a probability distribution $\mathcal{P}$. Suppose further that $D(\mathcal{P}\|\mathcal{Q}) \leq 2^{-\lambda-3}$ for $\lambda \leq \kappa - 4$. Then if $\mathcal{G}^{\mathcal{P}}$ is κ-bit secure, $\mathcal{G}^{\mathcal{Q}}$ is λ-bit secure.*

Proof. Deferred to Section B.1.

In the previous best known result, stated as Theorem 1 of [42], $D(\mathcal{P}\|\mathcal{Q}) \leq 2^{1-\kappa}$ gives that $\mathcal{G}^{\mathcal{Q}}$ is $(\kappa-8)$-bit secure. Using our result with $\lambda = \kappa - 4$, we find that the same divergence implies that $\mathcal{G}^{\mathcal{Q}}$ achieves $(\kappa-4)$-bit security, a 4 bit gain.

We will also use the following theorem in settings where we can bound the distance $R_\infty\left(\mathcal{G}^{\mathcal{Q}}\|\mathcal{G}^{\mathcal{P}}\right)$. This is close to a special case of Theorem 8 of [48], by observing that the max log distance $\Delta_{\mathrm{ML}}(X,Y) = \max(\ln R_\infty(X,Y), \ln R_\infty(Y,X))$. However, we only require the one-way bound $R_\infty\left(\mathcal{G}^{\mathcal{Q}}\|\mathcal{G}^{\mathcal{P}}\right)$, with no need for a bound on $R_\infty\left(\mathcal{G}^{\mathcal{P}}\|\mathcal{G}^{\mathcal{Q}}\right)$. In addition, letting $l = 8$ shows that our result implies $\ln R_\infty \leq 2^{3-\kappa/2}$, whereas the prior result requires $\ln R_\infty \leq 2^{-\kappa/2}$, a three bit gain.

Theorem 2. *Consider two distinguishing games $\mathcal{G}^{\mathcal{P}}, \mathcal{G}^{\mathcal{Q}}$. Then if $\mathcal{G}^{\mathcal{P}}$ is κ-bit secure and $R_\infty\left(\mathcal{G}^{\mathcal{Q}} \| \mathcal{G}^{\mathcal{P}}\right) \leq 2^{(l-\kappa-2)/2}+1$ for $\kappa > l \geq 4$ then $\mathcal{G}^{\mathcal{Q}}$ is $(\kappa - l)$-bit secure.*

Proof. Given in Section B.1.

3.2 Search Games

We now turn our attention to search games, where we follow [53,11] in considering the Rényi divergence. In later sections, we will see that this enables us to give tight bounds.

Lemma 9. *Suppose a search game S_0 is κ-bit secure, and let $R(a) = R_a(S_1 \| S_0)$. Then if there is $a > 1$ such that $\kappa \geq l \geq \left(1 - \frac{1}{a}\right) \log R(a) + \frac{\kappa}{a}$, we have that S_1 is $(\kappa - l)$-bit secure.*

Proof. Given in Section B.2.

Our lemma makes explicit the trade off discussed in Section 3.3 of [53], where the author discusses that small values of a decrease the tightness of a reduction, while large values impact the efficiency. In our result, we can characterise this as follows: a small a translates to smaller values for $R(a)$, while a larger a give smaller values for $\frac{\kappa}{a}$. Writing this trade off down explicitly will enable us to select an a which minimises our security loss.

We will also use the Rényi divergence of order $a = \infty$ to bound the distance between games, for which we have the following lemma.

Lemma 10. *Suppose a search game S_0 is κ-bit secure, and that $R_\infty(S_1 \| S_0) \leq 2^l$ with $\kappa > l$. Then S_1 is $(\kappa - l)$-bit secure.*

Proof. Given in Section B.2.

4 IND-CPA-D and KR-D Security via Noise Flooding

In this section we use the results of the previous section to give smaller noise flooding parameters. We consider sampling noise from either the discrete or continuous Gaussian.

4.1 IND-CPA-D Security

Theorem 3. *Let Π be an FHE scheme with plaintext space $\mathcal{M} \subset \mathbb{R}^n$, and consider the scheme $\tilde{\Pi}$ which uses a discrete or continuous Gaussian noise flooding mechanism with variance*

$$\sigma^2 \geq 4E^2 2^\lambda$$

where E is a static upper bound on the error $\|e\|_2$, and $\lambda \leq \kappa - 4$. Then if Π is κ-bit secure in the IND-CPA game, $\tilde{\Pi}$ is λ-bit secure in the IND-CPA-D game.

Proof. Given in Section C.

Remark 1. In [42], the authors choose to bound noise in the infinity norm, and then relate to the 2-norm via $\|e\|_2^2 \leq n\|e\|_\infty^2$, which accounts for the extra factor of n in their expression. We choose to leave our results in the two norm so that this noise estimate can be targeted directly; if practitioners prefer to estimate noise in the infinity norm, our variance can be modified with an extra factor of n.

Our theorem makes essential use of the observation that if an adversary makes t queries to an oracle, then $t \leq T(\mathcal{A})$, found in [47,48], and uses this observation to eliminate the factor t from the flooding variance in the bit security framework.

4.2 KR-D Security

We can now give the key recovery analogue to Theorem 3. The proof of this result uses the Rényi divergence to argue for security with significantly less noise. To the best of our knowledge, the inverse reciprocal relationship with security loss is new to this work.

Theorem 4. *Let Π be an HE scheme with plaintext space $\mathcal{M} \subset \mathbb{Z}^n$, and consider the scheme $\tilde{\Pi}$ which uses the discrete or continuous Gaussian noise flooding mechanism with variance*

$$\sigma^2 \geq \frac{2\kappa t E^2 \log e}{l^2}$$

where the adversary is allowed at most t queries to its decryption oracle, E is a static upper bound on the error $\|e\|_2$, and $1 \leq l < \kappa$. Then if Π is κ-bit secure in the KR game, $\tilde{\Pi}$ is $(\kappa - l)$-bit secure in the KR-D game.

Proof. Given in Section C.

Remark 2. The corresponding variance in the IND-CPA-D setting(Theorem 3) is of the form $\sigma^2 = 2^{\kappa - l + 2}E^2$. We therefore have variance that is exponential in our desired final security level, whereas in the key recovery game we are able to find a variance which is linear in the starting security, and proportional to the inverse square of the tolerated security loss. However, our search game variance is linear in the number of permitted queries to the decryption oracle, whereas our decision game variance is independent of this value.

5 IND-CPA-D and KR-D Security for Rescaled Ciphertexts from HintLWE

So far, our results hold for generic approximate homomorphic encryption schemes: we only require that it is possible to upper bound the decryption error. However, in practice, we would want to consider securing schemes that are based on the LWE or RLWE problem. If we add extra noise to decryptions, the most closely related lattice problem is the HintLWE problem: in this section, we consider a method for relating IND-CPA-D security to Hint(R)LWE hardness by restricting our decryption oracle to rescaled ciphertexts.

5.1 Rescaled Ciphertexts as HintLWE Samples

Consider the view of an adversary who is given a challenge ciphertext

$$\left(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{m}^{(b)}\right),$$

and who requests the decryption of an LWE ciphertext which have been rescaled by q and then flooded with a sample $\epsilon \leftarrow \mathcal{D}^n_{\sigma_\epsilon^2}$. If the ciphertext before rescale is given by $\mathtt{ct}.\mathbf{m} = (\mathbf{A}', \mathbf{b}' = \mathbf{A}'\mathbf{s} + \mathbf{e}' + \mathbf{m})$, with $\left\lceil \frac{\mathbf{m}}{q} \right\rfloor$ known to the adversary, then from this decryption the adversary learns

$$\frac{\mathbf{e}' + [\mathbf{m}]_q}{q} + \frac{[\mathbf{A}']_q}{q}\mathbf{s} - \frac{[\mathbf{b}']_q}{q} + \epsilon$$

The key observation is that $q, [\mathbf{A}']_q$ and $[\mathbf{b}']_q$ *are known to the adversary*, so we can summarise the view of the adversary as –

$$\left(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{m}^{(b)} \mod Q, \frac{\mathbf{e}' + [\mathbf{m}]_q}{q} + \underbrace{\frac{[\mathbf{A}']_q}{q}\mathbf{s} + \epsilon}_{\text{hint}}\right),$$

and the adversary is required to determine whether $b = 0$ or $b = 1$, which requires distinguishing $\mathbf{A}\mathbf{s} + \mathbf{e} \mod Q$ from uniform.

If $\mathbf{e}' + [\mathbf{m}]_q = 0$, we would already be able to appeal to the hardness of HintLWE: we could take $\mathbf{X} = \frac{[\mathbf{A}']_q}{q}$. However, this situation occurs with negligible probability, and indeed, $[\mathbf{m}]_q$ may leak b, while it's unclear exactly how $\mathbf{e}'$ will impact our security. For example, $\mathbf{e}'$ could be the result of a previous rescale, and so leak information on the secret $\mathbf{s}$. Alternatively, the ciphertext could be the result of a ciphertext-ciphertext multiplication, and so leak information on $\mathbf{m}^{(b)}$. Given this ambiguity, we prefer a hybrid approach where we use the noise flooding approach to first "remove" the term $\frac{\mathbf{e}'+[\mathbf{m}]_q}{q}$, and then appeal to the hardness of HintLWE. In more detail, we consider games $\mathcal{G}_0$ and $\mathcal{G}_1$ where the decryption oracle outputs

$$\mathcal{G}_0 : \frac{\mathbf{e}' + [\mathbf{m}]_q}{q} + \frac{[\mathbf{A}']_q}{q}\mathbf{s} - \frac{[\mathbf{b}']_q}{q} + \epsilon \quad \mathcal{G}_1 : \frac{[\mathbf{A}']_q}{q}\mathbf{s} - \frac{[\mathbf{b}']_q}{q} + \epsilon,$$

and argue that ϵ is from a sufficient wide distribution that these distributions are close together.

This hybrid approach clearly has more utility when the rescaled error $\frac{\mathbf{e}'}{q}$ is small, and in particular will give an advantage over the approach in the previous section when the rescaled error is smaller than the noise introduced by rescale.

For the remainder of this section, we consider adding noise from a *continuous* Gaussian. This is for two reasons: first, the hint matrices $\mathbf{X} = \frac{[\mathbf{A}]_q}{q}$ do not have integer entries, so to make the reduction from LWE go through, it is necessary

to consider a continuous Gaussian; secondly, the noise after rescale has integer entries, but $\frac{\mathbf{e}}{q}$ does not – adding noise from a continuous Gaussian ensures we do not have integral decryptions in one game and rational decryptions in the next.

We now give a formal argument for IND-CPA-D bit security from the security of HintLWE. The proof reduces to combining two results: first, we use Lemma 5 to relate IND-CPA-D security to HintLWE; second, we use Theorem 1 to argue that the security loss from removing $\frac{\mathbf{e}}{q}$ is small.

Theorem 5. *Let $\kappa' - \lambda \geq 4 + \log 12$, and let Π be an LWE-HE scheme with parameters $m, n, Q, \chi_{\mathbf{s}}, \chi_{\mathbf{e}}$. Assume further that the HintLWE (Decision) problem with the same parameters and with flooding variance σ^2 for hint matrix $\frac{1}{q}\left[\mathbf{A}_j\right]_q$ is κ'-bit secure. Then if we use flooding noise sampled from a continuous Gaussian with variance*

$$\sigma_j^2 \geq \max\left(4E^2 2^\lambda, \sigma^2\right)$$

on input $\mathsf{Rescale}\left(\left(\mathbf{A}_j, \mathbf{b}_j = \mathbf{A}_j\mathbf{s} + \mathbf{e}_j + \mathbf{m}_j\right), q\right)$ with $E \geq \frac{1}{q}\left\|\mathbf{e}_j + \left[\mathbf{m}_j\right]_q\right\|_2$ a static upper bound, then $\tilde{\Pi}$ is λ-bit secure in the IND-CPA-D game where we only decrypt rescaled ciphertexts.

Proof. Given in Section D.

In the key recovery setting, we find the following result.

Theorem 6. *Let $\kappa' - \lambda \geq 1$, and let Π be an LWE-HE scheme with parameters $m, n, Q, \chi_{\mathbf{s}}, \chi_{\mathbf{e}}$. Assume that the HintLWE (Search) problem with the same parameters and with flooding variance σ^2 for hint matrix $\frac{1}{q}\left[\mathbf{A_j}\right]_q$ is κ'-bit secure. Then if we use flooding noise sampled from a continuous Gaussian with variance*

$$\sigma_j^2 \geq \max\left(\frac{2t\kappa' E^2 \log e}{(\kappa' - \lambda)^2}, \sigma^2\right)$$

on input $\mathsf{Rescale}\left(\left(\mathbf{A}_j, \mathbf{b}_j = \mathbf{A}_j\mathbf{s} + \mathbf{e}_j + \mathbf{m}_j\right), q\right)$ with $E \geq \frac{1}{q}\left\|\mathbf{e}_j + \left[\mathbf{m}_j\right]_q\right\|$ a static upper bound. Then $\tilde{\Pi}$ is λ-bit secure in the KR-D game where we only decrypt rescaled ciphertexts, and the adversary is allowed at most t queries to the decryption oracle.

Proof. Given in Section D.

5.2 HintLWE Hardness

Theorem 5 assumes we have a bit security level κ' for HintLWE. For our results, we do not estimate this level directly, instead deriving the security using a related LWE instance.

In more detail, we relate the security of $\mathsf{HintLWE}^{l,\xi,X}_{m,n,Q,\chi}$ with Gaussian secret and hint errors to $\mathsf{LWE}_{m,n,Q,\chi'}$ for a narrower secret distribution χ', while leaving the parameters m, n, Q and $\sigma_{\mathbf{e}}$ unchanged. Our proof follows the proof strategy

of Theorem 1 of [35], but tightens the reduction for the bit security framework. We additionally generalise their result, allowing the width of the error on the i^{th} hint to depend on the i^{th} hint matrix $\mathbf{X}_i$. As our hint matrix $\mathbf{X}$ has real values, we also sample the vector $\mathbf{y}$ from a *continuous* Gaussian. Since we follow the same strategy as [35], we defer the proof to Section E.

Theorem 7 (Reduction LWE to HintLWE.). *Let $\chi_{\mathbf{s}}$ be a centred spherical discrete Gaussian distribution over $\mathbb{Z}^n$ with variances σ^2, and χ_i be a centred spherical continuous Gaussian with variance σ_i^2 respectively, where*

$$\frac{\sigma_i^2}{\rho^2} \geq \left\|\mathbf{X}_i^T \mathbf{X}_i\right\|_2 \tag{3}$$

for all indices i, for some fixed constant $\rho > 0$. Let σ'^2 satisfy

$$\frac{1}{\sigma'^2} \geq 2\left(\frac{1}{\sigma^2} + \frac{t}{\rho^2}\right),$$

and suppose that there is $\varepsilon \in (0,1)$ with $\sigma'^2 \geq 2\eta_\varepsilon(\mathbb{Z}^n)^2$. Then, if $\chi_{\mathbf{s}'} = \mathcal{D}^n_{\sigma'^2}$, there is an efficient reduction from $\mathsf{LWE}_{m,n,Q,\chi_{\mathbf{s}'},\chi_{\mathbf{e}}}$ to $\mathsf{HintLWE}^{\xi,\mathbf{X},t}_{m,n,Q,\chi_{\mathbf{s}},\chi_{\mathbf{e}}}$ with security loss as follows:

- *If the decisional LWE game is κ-bit secure, and $\frac{2\varepsilon}{1-\varepsilon} \leq 2^{(-\kappa'-2)/2}$ for $\kappa' \leq \kappa - 4$, then the decisional HintLWE game is κ'-bit secure.*
- *If the search LWE game is κ-bit secure, and $\frac{1+\varepsilon}{1-\varepsilon} \leq 2^{\kappa-\kappa'}$ for $\kappa' < \kappa$, then the search HintLWE game is κ'-bit secure.*

Proof. Given in Section E.

The result from [35] is outside the bit security framework, and considers the decision game exclusively, finding the requirement $\frac{2\varepsilon}{1-\varepsilon} \leq O\left(2^{l-\kappa}\right)$. We have moved the argument into the bit security framework, and, by consideration of the R_∞ divergence, have been able to square root this constraint.

This theorem gives the following corollaries, which should be easier to interpret in the context of noise flooding.

Corollary 1. *Let $\kappa' \leq \kappa - 4$ and $\sigma^2 > 0$, and suppose*

$$\sigma'^2 := \frac{\ln 4n + \ln\left(2^{(\kappa'+2)/2} + 1\right)}{\pi^2} < \frac{1}{2}\sigma^2. \tag{4}$$

Then if $\mathsf{LWE}_{m,n,Q,\chi_{\mathbf{e}},\chi'_{\mathbf{s}}}$ for $\chi_{\mathbf{s}'} = \mathcal{D}^n_{\sigma'^2}$ is κ-bit secure in the decision game, then $\mathsf{HintLWE}_{m,n,Q,\chi_{\mathbf{e}},\chi_{\mathbf{s}}}$ with $\chi_{\mathbf{s}} = \mathcal{D}^n_{\sigma^2}$ is κ'-bit secure in the decision game whenever $\mathbf{y}_i \leftarrow \mathcal{N}(0, \sigma_i^2 \mathbf{I}_{k_i})$ with

$$\sigma_i^2 \geq t \left\|\mathbf{X}_i^T \mathbf{X}_i\right\|_2 \left(\frac{1}{2\sigma'^2} - \frac{1}{\sigma^2}\right)^{-1}. \tag{5}$$

Proof. Let $\rho^2 = t\left(\frac{1}{2\sigma'^2} - \frac{1}{\sigma^2}\right)^{-1}$, and observe that

$$2\left(\frac{1}{\sigma^2} + \frac{t}{\rho^2}\right) = \frac{1}{\sigma'^2},$$

so that it suffices to show that there exists $\varepsilon \in (0,1)$ such that both $\sigma'^2 \geq 2\eta_\varepsilon(\mathbb{Z}^n)^2$ and $\frac{2\varepsilon}{1-\varepsilon} \leq 2^{(-\kappa'-2)/2}$. This is achieve by setting $\varepsilon = \frac{1}{2^{(\kappa'+4)/2}+1}$, and using Lemma 6.

For the search game, we have the following.

Corollary 2. *Let $\kappa' \leq \kappa - 1$ and $\sigma^2 > 0$, and suppose*

$$\sigma'^2 := \frac{\ln 4n - \ln\left(1 - 2^{\kappa'-\kappa}\right)}{\pi^2} < \frac{1}{2}\sigma^2. \tag{6}$$

Then if the $\mathsf{LWE}_{m,n,Q,\chi_{\mathbf{e}},\chi'_{\mathbf{s}}}$ for $\chi_{\mathbf{s}'} = \mathcal{D}^n_{\sigma'^2}$ is κ-bit secure in the key recovery game, then the $\mathsf{HintLWE}_{m,n,Q,\chi_{\mathbf{e}},\chi_{\mathbf{s}}}$ with $\chi_{\mathbf{s}} = \mathcal{D}^n_{\sigma^2}$ is κ'-bit secure in the key recovery game whenever $\mathbf{y}_i \leftarrow \mathcal{N}(0, \sigma_i^2\mathbf{I}_{k_i})$ with

$$\sigma_i^2 \geq t\left\|\mathbf{X}_i^T\mathbf{X}_i\right\|_2\left(\frac{1}{2\sigma'^2} - \frac{1}{\sigma^2}\right)^{-1}. \tag{7}$$

Proof. Let $l = \kappa - \kappa'$. We proceed as in the proof of Corollary 1, this time letting and $\varepsilon = \frac{2^l-1}{2^l+1}$, and showing that $\frac{1+\varepsilon}{1-\varepsilon} \leq 2^l$.

We now discuss the constraints on the HintLWE secret if we want to use these reductions to argue for security from LWE security. As we will see, these translate directly to constraints on the secret key distribution of our HE scheme.

First, our secret must come from a discrete Gaussian. Such secrets do appear in the HE security standard [10], but are less frequently used in practice compared to bounded distributions. No corresponding reduction is known for the ternary and sparse ternary distributions that are popular in applications of CKKS, and obtaining such a reduction seems challenging. The first step to obtaining a reduction that follows the same pattern would be express exactly how the distribution of such a secret changes conditional on knowledge of $\mathbf{Xs} + \mathbf{y}$, corresponding to the Gaussian case in Lemma 20.

Second, we need the secret variance σ^2 to be sufficiently wide to accommodate the reduction. For the search version, this requirement roughly translates to needing $\sigma^2 > 2\eta_1(\mathbb{Z}^n)^2$. However, for the decision variant, our requirement is of the form $\sigma^2 \geq \frac{(\kappa'+2)\ln 2}{\pi^2}$. For small secrets, this represents a significant limitation to the amount of security we can claim once we incorporate hints. However, giving a reduction in the bit security framework and using an R_∞ based game hop rather than one based on the statistical distance, we have relaxed this restriction from $\sigma^2 \geq \frac{2(\kappa'+2)\ln 2}{\pi^2}$ to $\sigma^2 \geq \frac{(\kappa'+2)\ln 2}{\pi^2}$, enabling us to claim bit security levels that are greater by a factor of two.

The exact same results hold in the ring setting, where we take n to be the ring dimension.

5.3 From LWE to IND-CPA-D and KR-D Security

If we derive κ' and σ using our reduction (Corollary 1), we can remove the appeal to HintLWE hardness from Theorem 5, instead claiming our scheme is at least as secure as an instance of the better studied LWE problem. This is achieved by the following theorem.

Corollary 3. *Let Π be an LWE-HE scheme with parameters m, n, q, error distribution $\chi_{\mathbf{e}}$ and secret distribution $\mathcal{D}^n_{\sigma^2}$. Let $l \geq 4, \kappa - l - \lambda \geq 4 + \log 12$ and suppose*

$$\sigma'^2 := \frac{\ln 4n + \ln\left(1 + 2^{(\kappa - l + 2)/2}\right)}{\pi^2} < \frac{1}{2}\sigma^2.$$

Then if $\mathsf{LWE}_{m,n,q,\chi_{\mathbf{e}},\chi_{\mathbf{s}'}}$ (Decision) is κ-bit secure with $\chi_{\mathbf{s}'} = \mathcal{D}^n_{\sigma'^2}$ then $\tilde{\Pi}$ is λ-bit secure in the IND-CPA-D model provided we only decrypt rescaled ciphertexts, and use a continuous Gaussian noise flooding mechanism such that on input $\mathsf{Rescale}\left(\left(\mathbf{A}_j, \mathbf{b}_j = \mathbf{A}_j\mathbf{s} + \mathbf{e}_j + \mathbf{m}_j\right), q_i\right)$,

$$\sigma_j^2 \geq \max\left(4E^2 2^\lambda, \frac{t}{q_i^2}\left\|\left[\mathbf{A}_j^T\right]_{q_i}\left[\mathbf{A}_j\right]_{q_i}\right\|_2\left(\frac{1}{2\sigma'^2} - \frac{1}{\sigma^2}\right)^{-1}\right),$$

where $E \geq \frac{1}{q_i}\left\|\mathbf{e}_j + [\mathbf{m}_j]_{q_i}\right\|_2$ is a static upper bound, and the adversary is allowed at most t queries to the decryption oracle.

Proof. This proof amounts to letting $\kappa' = \kappa - l$ in Corollary 1, and then concluding with Theorem 5.

We could simplify the presentation of this corollary if we set the parameter l to its smallest possible value (and so tightest possible reduction), $l = 4$. However the relationship between $(\kappa - l)$ means that a small secret cannot necessarily accommodate this reduction, as we found when we applied our result to concrete parameters. Therefore we leave the result in terms of the loss l.

In the search setting, we obtain the following.

Corollary 4. *Let Π be an LWE-HE scheme with parameters m, n, q, error distribution $\chi_{\mathbf{e}}$ and secret distribution $\mathcal{D}^n_{\sigma^2}$. Let $l \geq 1, \kappa - l - \lambda \geq 1$ and suppose*

$$\sigma'^2 := \frac{\ln 4n - \ln\left(1 - 2^{-l}\right)}{\pi^2} < \frac{1}{2}\sigma^2.$$

Then if $\mathsf{LWE}_{m,n,q,\chi_{\mathbf{e}},\chi_{\mathbf{s}'}}$ (Search) is κ-bit secure with $\chi_{\mathbf{s}'} = \mathcal{D}^n_{\sigma'^2}$ then $\tilde{\Pi}$ is λ-bit secure in the KR-D model provided we only decrypt rescaled ciphertexts, and use a continuous Gaussian noise flooding mechanism such that on input $\mathsf{Rescale}\left(\left(\mathbf{A}_j, \mathbf{b}_j = \mathbf{A}_j\mathbf{s} + \mathbf{e}_j + \mathbf{m}\right), q_i\right)$,

$$\sigma_j^2 \geq \frac{t}{\max}\left(\frac{2(\kappa - l)E^2 \log e}{(\kappa - l - \lambda)^2}, \frac{1}{q_i^2}\left\|\left[\mathbf{A}_j^T\right]_{q_i}\left[\mathbf{A}_j\right]_{q_i}\right\|_2\left(\frac{1}{2\sigma'^2} - \frac{1}{\sigma^2}\right)^{-1}\right),$$

where $E \geq \frac{1}{q_I}\left\|\mathbf{e}_j + [\mathbf{m}_j]_{q_i}\right\|_2$ is a static upper bound, and the adversary is allowed at most t queries to the decryption oracle.

We note that these final results make no mention of the HintLWE problem, instead appealing to the hardness of LWE. We make this choice so that we can use security estimates for the well studied LWE problem, provided by the community sourced LWE Estimator [1], rather than estimating the hardness of HintLWE directly. However, given an estimate for the bit security for HintLWE with a given flooding noise, we can instead make use of Theorems 5 and 6 directly. We give a high level overview of our results, as well as their hardness assumptions and restrictions, in Table 1.

Hardness Assumption	IND-CPA-D	KR-D	Restrictions
IND-CPA/KR Security	Theorem 3	Theorem 4	None
HintLWE	Theorem 5	Theorem 6	Only decrypt rescaled ciphertexts
LWE	Corollary 3	Corollary 4	Only decrypt rescaled ciphertexts, Gaussian secret

Table 1: Our results for achieving IND-CPA-D and KR-D security from different hardness assumptions.

To give the ring analogues of these theorems, we would instead argue from the security of RLWE, and make the substitution:

$$\left\| \left[\mathbf{A}_j^T\right]_{q_i} [\mathbf{A}_j]_{q_i} \right\|_2 \mapsto \left\| U_{a_j}^T U_{a_j} \right\|_2$$

In power of two cyclotomics, this is equivalent to $\left(\|a_j\|_2^{\text{can}}\right)^2$.

We have left our results in terms of $\left\|\mathbf{A}_j^T \mathbf{A}_j\right\|_2$ and $\|a_j\|_2^{\text{can}}$, to enable practitioners to set noise flooding parameters in an "application aware" way, as described in [2]. For example, for some circuits, $\mathbf{A}_j$ is always uniformly distributed, so that $[\mathbf{A}_j]_q$ is also uniform. In this case, we could use [9], Theorem 4.2, to give a high probability bound.

However, once our circuit includes rotations or plaintext multiplications, it becomes difficult to claim that these entries (or coefficients) are independently uniformly distributed. We therefore provide worst case bounds; these are given by the following lemma.

Lemma 11. *Let $q \in \mathbb{Z}$ be arbitrary, and m and n be positive integers. Then*

1. *Let $\mathbf{A} \in \mathbb{Z}^{m \times n}$. Then $\left\|[\mathbf{A}^T]_q [\mathbf{A}]_q\right\|_2 \leq \frac{q^2 nm}{4}$.*
2. *Let $a \in \mathcal{R}$. Then $\left(\|a\|_2^{\text{can}}\right)^2 \leq \left(\left(\frac{2n}{\pi}\right)^2 + 1\right) \frac{q^2}{4}$.*

Proof. We observe that the entries of $[\mathbf{A}]_q$ (resp. coefficients of $[a]_q$) are all of magnitude at most $\frac{q}{2}$, so that the first result follows from $\left\|\mathbf{M}^T\mathbf{M}\right\|_2 \leq \text{Trace}\left(\mathbf{M}^T\mathbf{M}\right)$, while the second follows from Lemma 3.

6 Concrete Parameters and Comparison

So far, we have presented two methods for achieving IND-CPA-D and KR-D security: a pure bit security approach, and a hybrid between this approach and HintLWE. Here we give an insight into how to use these two approaches, and comparisons between them. We additionally provide an analysis of the parameters used in the OpenFHE library [3] in Section F. The source code for these experiments can be found here[5].

6.1 HintLWE Security

We consider building on the parameters given in the most recent HE Community Standard [10], in particular the parameters given in Table 5.2 which have a secret drawn from a discrete Gaussian with standard deviation 3.19, and which target 128-bit security in the classical model. These are given in Table 2.

$\log n$	$\log q$
10	28
11	55
12	108
13	216
14	432
15	870
16	1749
17	3525

Table 2: The RLWE parameters we consider in this section, taken from the HE Community Standard [10], Table 5.2. For all parameters, the secret and error are taken from a Discrete Gaussian distribution with standard deviation 3.19. For all parameters, we consider an adversary with access to arbitrarily many RLWE samples.

As discussed in Section 5.2, the reduction from LWE to HintLWE requires the secret key to be from a sufficiently wide distribution: in the Decision case (and so IND-CPA-D setting), we require the secret to have variance

$$\sigma^2 > 2\eta_1(\mathbb{Z}^n)^2 + \frac{\ln\left(1+2^{\lambda/2+1}\right)}{\pi^2} \approx 2\eta_1(\mathbb{Z}^n)^2 + \frac{\lambda+2}{\pi^2}\ln(2),$$

[5] https://github.com/TabOg/hintLWE-IND-CPA-D

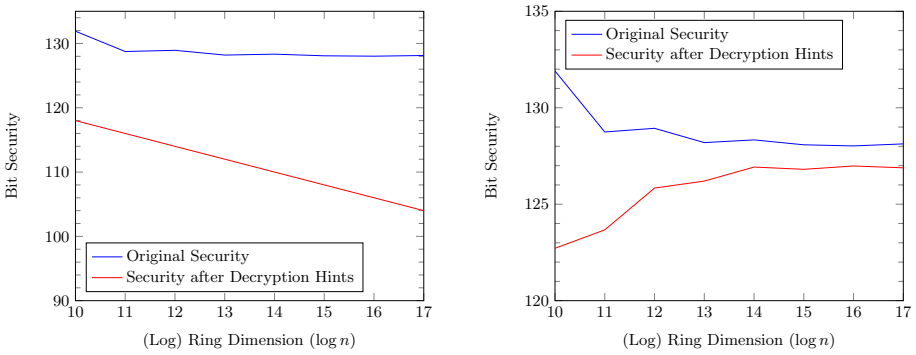
to argue that the HintLWE problem is λ-bit secure. This linearity in the final security is quite restrictive for small secrets, and we find that in the decision setting, this restriction dictates the bit security of HintLWE, rather than the security of the related LWE instance.

For the Search problem (and so KR-D setting), the requirement is instead

$$\sigma^2 > 2\eta_1(\mathbb{Z}^n)^2 - \frac{2\ln\left(1-2^{-l}\right)}{\pi^2} \approx 2\eta_1(\mathbb{Z}^n)^2.$$

To achieve a security of $(\kappa - l)$, where κ is the bit security of the related LWE instance. This restriction is much more permissive. In fact we find that for standard parameters, it is possible to achieve only a 1 bit loss in security due to hints.

In Fig. 1, we give the highest possible bit security of the HintLWE Decision and Search problems according when we use a reduction from LWE, following Corollaries 1 and 2. We use the same cost model for LWE as was used to derive the community standard parameters, and rely on the popular community tool the Lattice Estimator [1]. We consider these security levels reasonable, but for applications which require building on the decision problem and desire higher security, this can easily be achieved by increasing the secret key deviation from 3.19: in fact, we find that increasing this deviation only as far as 3.5 allows bit security levels of 123 bit in the Decision case. We additionally remark that these security levels would not be possible without our new approach using the R_∞ divergence: the prior method is only able to claim security levels for the Decision problem of half these when the original secret has deviation 3.19.

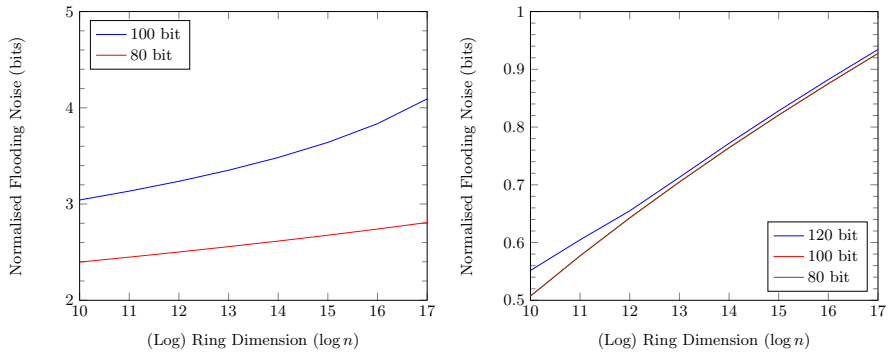


(a) Bit Security HintLWE (Decision)

(b) Bit Security of HintLWE (Search).

Fig. 1: Security of HintLWE Decision and Search as in Corollaries 1 and 2 for Community Standard Parameters [10] using the Lattice Estimator [1].

We additionally give an indication of how much flooding noise is necessary with these parameters to achieve a given HintLWE security level in Fig. 2.



(a) Decision (b) Search

Fig. 2: The amount of flooding noise σ necessary to achieve a given security level in the HintLWE game, normalised by $\sqrt{t \left\| \mathbf{X}_i^T \mathbf{X}_i \right\|_2}$. Parameters from Table 2.

6.2 IND-CPA-D Parameters

We now compare the amount of flooding noise required for each of our approaches, as well as the prior best known result. We recap the noise required to achieve λ bits of security in the IND-CPA-D setting in Table 3.

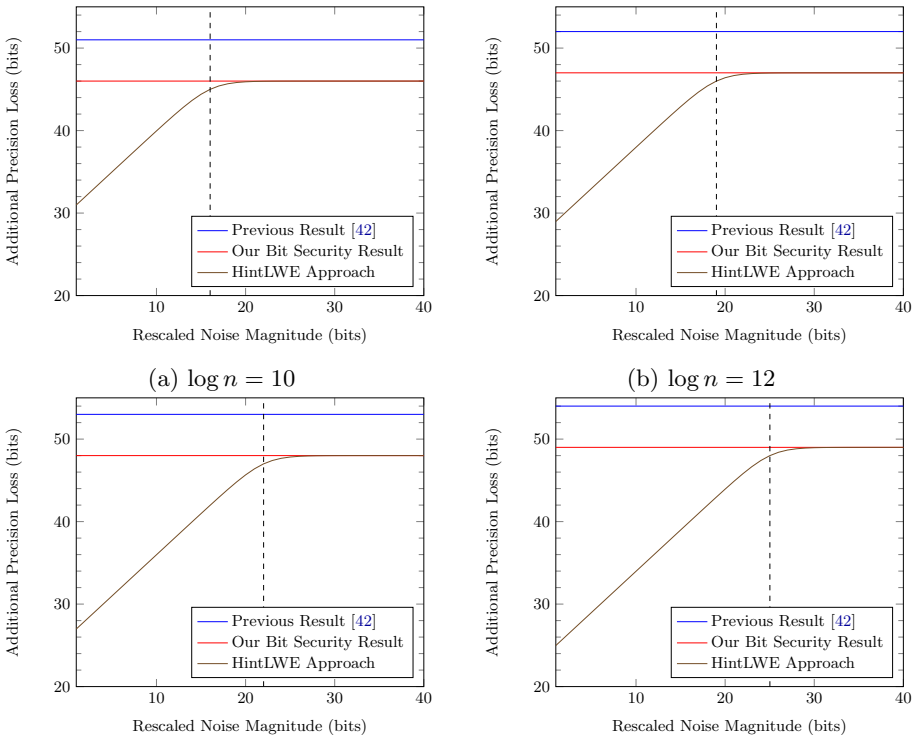
Method	Requirements	Flooding Variance
Prior Result [42]	$\kappa - \lambda \geqslant 8$	$64tE^2 2^\lambda$
Bit Security (Theorem 3)	$\kappa - \lambda \geqslant 4$	$4E^2 2^\lambda$
HintLWE Approach (Theorem 5)	Only decrypt rescaled ciphertexts, HintLWE (Decision) with flooding noise σ^2 is κ'-bit secure, and $\kappa' - \lambda \geqslant 4 + \log 12$	$\max\left\{4\tilde{E}^2 2^\lambda, \sigma^2\right\}$

Table 3: The amount of noise required to achieve λ bits of security in the IND-CPA-D setting using each approach. We use κ for the IND-CPA security, t for the number of decryption queries allowed, E for an upper bound on the total noise (2 norm) and $\tilde{E}$ for an upper bound on the rescaled noise (2 norm).

To understand how the bit security and HintLWE approaches compare, we need to split the total noise (E) into rescaled noise ($\tilde{E}$) and noise from rescaling. We achieve this via

$$E = \frac{1}{q} \|e_i + [m_i]_q + [a_i]_q s - [b_i]_q\|_2$$
$$\leq \tilde{E} + \frac{1}{q} \|[a_i]_q\|_\infty^{\text{can}} \sqrt{n}\sigma_s + \frac{1}{2}.$$

where we have used the definition of the operator norm, and the approximation $\|s\|_2 \approx \sqrt{n}\sigma_s$. We consider all of these to be pessimistic choices, and may change drastically given a target application where we can make more concrete assessments.



(a) $\log n = 10$ (b) $\log n = 12$

(c) $\log n = 14$ (d) $\log n = 16$

Fig. 3: Additional noise required to achieve IND-CPA-D security. The number of bits of noise we add under the bit security and HintLWE approaches to achieve IND-CPA-D security, as a function of the magnitude of the rescaled noise ($\tilde{E} \geq \frac{1}{q} \|e + [m]_q\|_2$) for $t = 2^6$ and other parameters from Table 2. We use $\frac{\sqrt{n}\sigma}{\|E\|}$ as a measure of additional precision loss. The dotted line indicates when the rescaled noise has equal magnitude to the noise from rescaling ($E = 2\tilde{E}$). All parameters achieve 80 bits of security in the IND-CPA-D game.

In Fig. 3, we plot the amount of additional noise required by each of these approaches, for a range of parameters. These are calculated via $\frac{\sqrt{n}\sigma}{E}$. We have used the reduction from LWE to determine the hardness of HintLWE (Decision). Throughout, we assume a worst case bound on $\|[a]_q\|^{\text{can}}$.

We can see that our noise flooding analysis gives a constant factor improvement over the state of the art. With the HintLWE approach, we find that we can achieve IND-CPA-D security by adding enough noise to flood out the rescaled noise ($\tilde{E}$) only, so that our approach gives significant savings when this noise is a small fraction of the total noise. In the region where the rescaling noise dominates (to the left of the dotted line), the HintLWE approach requires adding much less noise.

6.3 KR-D Parameters

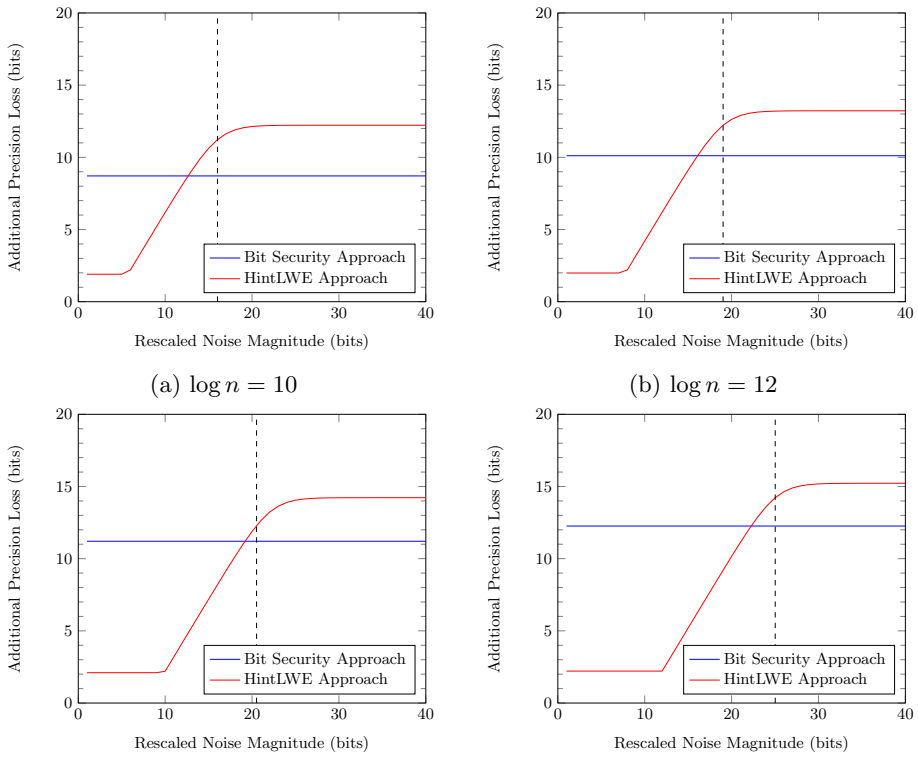
We conclude by comparing the noise we need to add to achieve KR-D security using the approaches we have outlined in this paper. We recap the amount of noise needed to achieve λ bits of security in this setting in Table 4. We only give parameters for the most efficient reduction from HintLWE. We again limit our attention to the parameters considered in the community guidelines for Homomorphic Encryption [10] as given in Table 2.

Method	Requirements	Flooding Variance
Bit Security (Theorem 4)	$\kappa - \lambda \geqslant 1$	$\frac{2\kappa t E^2 \log e}{(\kappa-\lambda)^2}$
HintLWE Approach (Theorem 6)	Only decrypted rescaled ciphertexts, HintLWE(Search) with flooding noise σ^2 is κ'-bit secure, and $\kappa' - \lambda \geqslant 1$	$\max\left\{\frac{2\kappa t \tilde{E}^2}{(\kappa'-\lambda)^2}, \sigma^2\right\}$

Table 4: The amount of noise required to achieve λ bits of security in the KR-D setting using each approach. We use κ for the KR security, t for the number of permitted decryption queries, E for an upper bound on the total noise (2 norm) and $\tilde{E}$ for an upper bound on the rescaled noise (2 norm).

As in the IND-CPA-D setting, we seek to understand how much *additional* noise is required to achieve a given security level as we vary $\tilde{E} \geq \frac{1}{q}\|e_i + [m_i]_q\|$. We use the same method as the previous section to compare the quantities E and $\tilde{E}$. The absolute and relative noise magnitudes are given in Fig. 4.

We see that both methods have basically constant flooding noise while the noise from rescaling ($\frac{1}{q}([a]_q s - [b]_q)$) dominates the rescaled noise ($\frac{1}{q}(e + [m]_q)$), with the HintLWE approach adding 10 fewer bits of noise compared to the pure bit security approach. However, once these two sources of noise are of similar



(a) $\log n = 10$ (b) $\log n = 12$

(c) $\log n = 14$ (d) $\log n = 16$

Fig. 4: Additional noise required to achieve KR-D security. The number of bits of noise we add under the bit security and HintLWE approaches to achieve KR-D security, as a function of the magnitude of the rescaled noise ($\tilde{E} \geq \frac{1}{q} \left\| e + [m]_q \right\|_2$) for $t = 2^6$ and other parameters from Table 2. We use $\frac{\sqrt{n}\sigma}{\|E\|}$ as a measure of additional precision loss. The dotted line indicates when the rescaled noise has equal magnitude to the noise from rescaling ($E = 2\tilde{E}$). This amount of noise flooding achieves 120 bit security in the KR-D game.

magnitude, (as indicated by the dashed line) the two approaches require similar levels of precision loss. We remark that the amount of precision loss required to achieve only KR-D security is significantly lower than IND-CPA-D security: indeed, from this graph we can see that in this setting, only two additional bits of precision loss are sufficient to give 120 bits of security against key recovery attacks.

7 Conclusion

In this work, we have shown that we can achieve λ-bit security IND-CPA-D for any approximate homomorphic encryption scheme by adding an additional $\lambda/2 + O(1)$ bits of noise, independent of the number of permitted decryptions. In the KR-D setting, when we allow t decryption queries we have shown that it is enough to add $\frac{1}{2}\log\lambda + O(\log t)$ additional bits of noise, meaning that this security can be achieved with minimal overhead.

By considering the specific case of rescaled (R)LWE ciphertexts, we have proposed a hybrid approach which appeals to the hardness of HintLWE. This approach splits the noise into two parts, the rescaled noise and the noise added by the rescale operation. We find that when the second of these dominates, this approach permits adding only $O(\log t)$ additional bits of noise in both the IND-CPA-D and KR-D settings.

It is beyond the scope of this paper to suggest what is the "right" security model that HE schemes should meet in practice. However, we do note that no implementation of Homomorphic Encryption attempts to achieve a meaningful, provable level of bit security in the IND-CPA-D setting, and in fact some libraries contain disclaimers that the library does not remain secure in this model[6], and it is declared beyond the scope of the recent community standard [10]. We hope that the competitive parameters we are able to achieve which give provable KR-D security will lead practitioners to consider adopting at least this security model.

8 Further Work

Hint LWE In this work, we have established the hardness of our HintLWE instances by considering the hardness of a related LWE instance, following [35]. However, this reduction is limited in two respects. First, it requires the secret to be drawn from a sufficiently wide distribution. In fact, the limitation means we obtain a significant 20 bit security gap between the HintLWE(Search) and HintLWE(Decision) problems for standard parameters. Future work could close this gap, either by revisiting the reduction, or by giving a reduction between the decision and search problems in the bit security framework. Alternatively,

[6] Consider ConcreteML, SEAL. One exception is OpenFHE [3], which targets the much weaker notion of (c, s) security introduced in [42].

future standards could investigate standardising wider secret distributions to ensure this gap does not materialise.

The second limitation is that we can only consider discrete Gaussian secrets: in practice, applications of HE or CKKS often use a uniform ternary secret distribution. Showing HintLWE is as hard as LWE in this setting would greatly enhance the utility of our results.

An alternative approach would be to omit the reduction from LWE entirely, and to estimate the hardness of HintLWE directly. An estimator which compiles attacks from the literature would then serve a similar purpose to the LWE Estimator, which is currently used to develop standardised parameters.

Robust Noise Estimates In our results, we have assumed that it is possible to upper bound the decryption noise statically. In practice, and in particular in lattice based HE, this can be difficult to achieve with performant parameters, while staying secure against adversaries who can select the circuit to be evaluated. This subtlety is exploited in [16,17,31], and a framework for resolution is proposed in [2]. Future work may wish to consider how we can upper bound noise in a way that is robust against adversarially manipulated circuits.

Sparse Key LWE When reducing LWE to HintLWE, we appeal to the hardness of an LWE instance with a smaller key. For standard parameters, this requires us to appeal to the hardness of LWE with a secret key that is quite sparse. The security estimates provided by the LWE Estimator suggest that this sparsification does not degrade security by much compared to standard parameters: more research into the sparse secret setting is necessary to ensure that this is indeed the case.

Acknowledgements

I am grateful to Duhyeong Kim for introducing me to the HintLWE problem, and to Fernando Virdia for discussions on its concrete hardness. In addition, thanks to Dan Jones, Elizabeth Quaglia, and Jelle Vos for their invaluable feedback on an initial draft of this work.

References

1. Martin R Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
2. Andreea Alexandru, Ahmad Al Badawi, Daniele Micciancio, and Yuriy Polyakov. Application-aware approximate homomorphic encryption: Configuring FHE for practical use. Cryptology ePrint Archive, Paper 2024/203, 2024.
3. Ahmad Al Badawi, Andreea Alexandru, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Carlo Pascoe, Yuriy Polyakov, Ian Quah, Saraswathy R.V., Kurt Rohloff, Jonathan Saylor, Dmitriy Suponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. OpenFHE: Open-source fully homomorphic encryption library. Cryptology ePrint Archive, Paper 2022/915, 2022. https://eprint.iacr.org/2022/915.

4. Youngjin Bae, Jung Hee Cheon, Jaehyung Kim, and Damien Stehlé. Bootstrapping bits with CKKS. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 94–123. Springer, Cham, May 2024.
5. Youngjin Bae, Jaehyung Kim, Damien Stehlé, and Elias Suvanto. Bootstrapping small integers with CKKS. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part I*, volume 15484 of *LNCS*, pages 330–360. Springer, Singapore, December 2024.
6. Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 3–24. Springer, Berlin, Heidelberg, November / December 2015.
7. Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: using the rényi divergence rather than the statistical distance. *Journal of Cryptology*, 31:610–640, 2018.
8. Flávio Bergamaschi, Anamaria Costache, Dana Dachman-Soled, Hunter Kippen, Lucas LaBuff, and Rui Tang. Revisiting the security of approximate FHE with noise-flooding countermeasures. In Tibor Jager and Jiaxin Pan, editors, *PKC 2025, Part V*, volume 15678 of *LNCS*, pages 102–134. Springer, Cham, May 2025.
9. 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.
10. Jean-Philippe Bossuat, Rosario Cammarota, Ilaria Chillotti, Benjamin R. Curtis, Wei Dai, Huijing Gong, Erin Hales, Duhyeong Kim, Bryan Kumara, Changmin Lee, Xianhui Lu, Carsten Maple, Alberto Pedrouzo-Ulloa, Rachel Player, Yuriy Polyakov, Luis Antonio Ruiz Lopez, Yongsoo Song, and Donggeon Yhee. Security guidelines for implementing homomorphic encryption. *CiC*, 1(4):26, 2024.
11. Katharina Boudgoust and Peter Scholl. Simple threshold (fully homomorphic) encryption from LWE with polynomial modulus. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part I*, volume 14438 of *LNCS*, pages 371–404. Springer, Singapore, December 2023.
12. Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886. Springer, Berlin, Heidelberg, August 2012.
13. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012.
14. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on computing*, 43(2):831–871, 2014.
15. Clément L Canonne, Gautam Kamath, and Thomas Steinke. The discrete gaussian for differential privacy. *Advances in Neural Information Processing Systems*, 33:15676–15688, 2020.
16. Marina Checri, Renaud Sirdey, Aymen Boudguiga, and Jean-Paul Bultel. On the practical CPA^D security of "exact" and threshold FHE schemes and libraries. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part III*, volume 14922 of *LNCS*, pages 3–33. Springer, Cham, August 2024.

17. Jung Hee Cheon, Hyeongmin Choe, Alain Passelègue, Damien Stehlé, and Elias Suvanto. Attacks against the IND-CPA$^{\mathrm{D}}$ security of exact FHE schemes. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024*, pages 2505–2519. ACM Press, October 2024.
18. Jung Hee Cheon, Kyoohyung Han, Andrey Kim, Miran Kim, and Yongsoo Song. A full RNS variant of approximate homomorphic encryption. In Carlos Cid and Michael J.: Jacobson, Jr., editors, *SAC 2018*, volume 11349 of *LNCS*, pages 347–368. Springer, Cham, August 2019.
19. Jung Hee Cheon, Seungwan Hong, and Duhyeong Kim. Remark on the security of CKKS scheme in practice. Cryptology ePrint Archive, Report 2020/1581, 2020.
20. Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic encryption for arithmetic of approximate numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 409–437. Springer, Cham, December 2017.
21. Heewon Chung, Hyojun Kim, Young-Sik Kim, and Yongwoo Lee. Amortized large look-up table evaluation with multivariate polynomials for homomorphic encryption. Cryptology ePrint Archive, Report 2024/274, 2024.
22. Anamaria Costache, Benjamin R. Curtis, Erin Hales, Sean Murphy, Tabitha Ogilvie, and Rachel Player. On the precision loss in approximate homomorphic encryption. In Claude Carlet, Kalikinkar Mandal, and Vincent Rijmen, editors, *SAC 2023*, volume 14201 of *LNCS*, pages 325–345. Springer, Cham, August 2024.
23. Anamaria Costache, Lea Nürnberger, and Rachel Player. Optimisations and tradeoffs for HElib. In Mike Rosulek, editor, *CT-RSA 2023*, volume 13871 of *LNCS*, pages 29–53. Springer, Cham, April 2023.
24. Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi. LWE with side information: Attacks and concrete security estimation. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 329–358. Springer, Cham, August 2020.
25. Dana Dachman-Soled, Huijing Gong, Tom Hanson, and Hunter Kippen. Revisiting security estimation for LWE with hints from a geometric perspective. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 748–781. Springer, Cham, August 2023.
26. Philip J Davis. *Circulant matrices*, volume 120. Wiley New York, 1979.
27. Leo de Castro, Chiraag Juvekar, and Vinod Vaikuntanathan. Fast vector oblivious linear evaluation from ring learning with errors. In *Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, pages 29–41, 2021.
28. Nir Drucker, Guy Moshkowich, Tomer Pelleg, and Hayim Shaul. BLEACH: Cleaning errors in discrete computations over CKKS. *Journal of Cryptology*, 37(1):3, January 2024.
29. Muhammed F. Esgin, Thomas Espitau, Guilhem Niot, Thomas Prest, Amin Sakzad, and Ron Steinfeld. Plover: Masking-friendly hash-and-sign lattice signatures. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VII*, volume 14657 of *LNCS*, pages 316–345. Springer, Cham, May 2024.
30. Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012.
31. Qian Guo, Denis Nabokov, Elias Suvanto, and Thomas Johansson. Key recovery attacks on approximate homomorphic encryption with non-worst-case noise flooding countermeasures. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024.

32. Zhicong Huang, Wen-jie Lu, Cheng Hong, and Jiansheng Ding. Cheetah: Lean and fast secure {Two-Party} deep neural network inference. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 809–826, 2022.
33. Xiaoqian Jiang, Miran Kim, Kristin E. Lauter, and Yongsoo Song. Secure outsourced matrix computation and application to neural networks. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1209–1222. ACM Press, October 2018.
34. Andrey Kim, Antonis Papadimitriou, and Yuriy Polyakov. Approximate homomorphic encryption with reduced approximation error. In Steven D. Galbraith, editor, *CT-RSA 2022*, volume 13161 of *LNCS*, pages 120–144. Springer, Cham, March 2022.
35. Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-MLWE. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 549–580. Springer, Cham, August 2023.
36. Adeline Langlois, Damien Stehlé, and Ron Steinfeld. Gghlite: More efficient multilinear maps from ideal lattices. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 239–256. Springer, 2014.
37. Joon-Woo Lee, HyungChul Kang, Yongwoo Lee, Woosuk Choi, Jieun Eom, Maxim Deryabin, Eunsang Lee, Junghyun Lee, Donghoon Yoo, Young-Sik Kim, et al. Privacy-preserving machine learning with fully homomorphic encryption for deep neural network. *iEEE Access*, 10:30039–30054, 2022.
38. Joon-Woo Lee, Eunsang Lee, Yongwoo Lee, Young-Sik Kim, and Jong-Seon No. High-precision bootstrapping of RNS-CKKS homomorphic encryption using optimal minimax polynomial approximation and inverse sine function. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 618–647. Springer, Cham, October 2021.
39. Seewoo Lee, Garam Lee, Jung Woo Kim, Junbum Shin, and Mun-Kyu Lee. Hetal: efficient privacy-preserving transfer learning with homomorphic encryption. In *International Conference on Machine Learning*, pages 19010–19035. PMLR, 2023.
40. Yongwoo Lee, Joon-Woo Lee, Young-Sik Kim, Yongjune Kim, Jong-Seon No, and HyungChul Kang. High-precision bootstrapping for approximate homomorphic encryption by error variance minimization. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 551–580. Springer, Cham, May / June 2022.
41. Baiyu Li and Daniele Micciancio. On the security of homomorphic encryption on approximate numbers. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 648–677. Springer, Cham, October 2021.
42. Baiyu Li, Daniele Micciancio, Mark Schultz, and Jessica Sorrell. Securing approximate homomorphic encryption using differential privacy. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 560–589. Springer, Cham, August 2022.
43. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010.
44. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013.

45. Alexander May and Julian Nowakowski. Too many hints - when LLL breaks LWE. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part IV*, volume 14441 of *LNCS*, pages 106–137. Springer, Singapore, December 2023.
46. 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.
47. Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 455–485. Springer, Cham, August 2017.
48. Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 3–28. Springer, Cham, April / May 2018.
49. Tabitha Ogilvie, Rachel Player, and Joe Rowell. Improved privacy-preserving training using fixed-hessian minimisation. In Michael Brenner, Tancrède Lepoint (Eds.), proceedings of the 8th Workshop on Encrypted Computing and Applied Homomorphic Cryptography (WAHC '20), 2020. https://doi.org/10.25835/0072999.
50. Saerom Park, Jaewook Lee, Jung Hee Cheon, Juhee Lee, Jaeyun Kim, and Junyoung Byun. Security-preserving support vector machine with fully homomorphic encryption. *SafeAI@ AAAI*, 2301, 2019.
51. Alain Passelègue and Damien Stehlé. Low communication threshold fully homomorphic encryption. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 297–329. Springer, 2025.
52. Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 80–97. Springer, Berlin, Heidelberg, August 2010.
53. Thomas Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 347–374. Springer, Cham, December 2017.
54. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009.
55. Mélissa Rossi. *Extended security of lattice-based cryptography*. Theses, Université Paris sciences et lettres, September 2020.
56. Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
57. Shun Watanabe and Kenji Yasunaga. Bit security as computational cost for winning games with high probability. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 161–188. Springer, Cham, December 2021.
58. Shun Watanabe and Kenji Yasunaga. Unified view for notions of bit security. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VI*, volume 14443 of *LNCS*, pages 361–389. Springer, Singapore, December 2023.

A Additional Preliminaries

A.1 Divergence Properties

The KL Divergence verifies the following properties.

Lemma 12 (KL Divergence Properties (Lemma 1 of [42])).

1. **Subadditivity.** *For vectors of random variables* $\mathcal{P} = (P_i)_i$ *and* $\mathcal{Q} = (Q_i)_i$ *with shared support,*

$$D(\mathcal{P}\|\mathcal{Q}) \leq \sum_i \max_x D([P_i|P_{<i} = x]\|[Q_i|Q_{<i} = x])$$

2. **Data Processing Inequality.** *For any function* f *and random variables* $\mathcal{P}, \mathcal{Q}$ *with shared support,*

$$D(\mathcal{P}\|\mathcal{Q}) \leq D(f(\mathcal{P})\|f(\mathcal{Q}))$$

3. **Pinsker's Inequality.** *For any two random variables* $\mathcal{P}, \mathcal{Q}$*,*

$$\Delta(\mathcal{P}, \mathcal{Q}) \leq \sqrt{D(\mathcal{P}\|\mathcal{Q})/2}.$$

We will also use the following properties of the Rényi divergence, which are analogous to the corresponding properties of the KL divergence. For the first property, we refer the reader to [11,55], while for the final properties we refer the reader to [56,36]: observe that the continuous case of the probability preservation property can be shown as in [36] using Hölder's inequality.

Lemma 13 (Rényi Divergence Properties).

– **Submultiplicativity.** *For vectors of random variables* $\mathcal{P} = (P_i)_i$ *and* $\mathcal{Q} = (Q_i)_i$ *with* $\mathrm{Supp}(\mathcal{P}) \subseteq \mathrm{Supp}(\mathcal{Q})$*,*

$$R_a(\mathcal{P}\|\mathcal{Q}) \leq \prod_i \max_x R_a([P_i|P_{<i} = x]\|[Q_i|Q_{<i} = x])$$

– **Data Processing Inequality.** *For any function* f *and random variables* $\mathcal{P}, \mathcal{Q}$*,*

$$R_a(f(\mathcal{P})\|f(\mathcal{Q})) \leq R_a(\mathcal{P}\|\mathcal{Q})$$

– **Probability Preservation.** *For any event* $E \subseteq \mathrm{Supp}(\mathcal{Q})$ *and* $a > 1$*,*

$$\mathcal{Q}(E) \geq \mathcal{P}(E)/R_\infty(\mathcal{P}\|\mathcal{Q});$$
$$\mathcal{Q}(E) \geq \mathcal{P}(E)^{a/(a-1)}/R_a(\mathcal{P}\|\mathcal{Q}).$$

We will also use the following lemmas, bounding the KL and Rényi divergence between both discrete and continuous Gaussians with the same standard deviation but different means.

Lemma 14 (Proposition 5, [15], full version). *Let* $\mathcal{D}$ *be the discrete Gaussian distribution. For* $\sigma^2 > 0$ *and* $\infty > a > 1$*,*

$$D(\mathcal{D}_{\mu,\sigma^2 I_n}\|\mathcal{D}_{\mu',\sigma^2 I_n}) \leq \frac{\|\mu - \mu'\|_2^2}{2\sigma^2},$$
$$R_a(\mathcal{D}_{\mu,\sigma^2 I_n}\|\mathcal{D}_{\mu',\sigma^2 I_n}) \leq \exp\left(\frac{a\|\mu - \mu'\|_2^2}{2\sigma^2}\right)$$

Lemma 15 ([56]). *For $\sigma^2 > 0$ and $\infty > a > 1$,*

$$D(\mathcal{N}(\mu, \sigma^2 I_n) \| \mathcal{N}(\mu', \sigma^2 I_n)) = \frac{\|\mu - \mu'\|_2^2}{2\sigma^2}$$
$$R_a(\mathcal{N}(\mu, \sigma^2 I_n) \| \mathcal{N}(\mu', \sigma^2 I_n)) = \exp\left(\frac{a\,\|\mu - \mu'\|_2^2}{2\sigma^2}\right)$$

A.2 Hybrid Advantage Details

In this Appendix we explain why we believe the advantage can be bounded with a constant $\frac{3}{2}$, rather than 3.

Lemma (Lemma 2 of [48]). Let $\mathcal{H}_i$ be a sequence of k distributions, and let $\mathcal{G}_{ij}$ be the indistinguishability game instantiated with $\mathcal{H}_i$ and $\mathcal{H}_j$. Let C be a fixed constant. Let $\epsilon_{ij} = \max_{\mathcal{A}} \mathsf{adv}^{\mathcal{A}}$ over all adversaries with $\mathcal{A}$ with $T(\mathcal{A}) \leq C$. Then $\epsilon_{1k} \leq \frac{3}{2} k \sum \epsilon_{i,i+1}$.

In Appendix B, the authors verify that

$$\sum_i^k \mathrm{adv}_{i,i+1} \geq \frac{\alpha_1 + \alpha_k}{\alpha_1 + 2\sum_{i=2}^{k-1} \alpha_i + \alpha_k} \mathrm{adv}_{1,k}\,.$$

A close inspection shows that the authors' proof also shows that

$$\sum_i^k \mathrm{adv}_{i,i+1} \geq \frac{2(\alpha_1 + \alpha_k)}{\alpha_1 + 2\sum_{i=2}^{k-1} \alpha_i + \alpha_k} \mathrm{adv}_{1,k}\,.$$

In more detail, using their notation from Appendix B, we have that

$$\begin{aligned} f(\gamma_{k-1}) &\geq \frac{(2(\gamma_1 - \gamma_k) - \alpha_1 + \alpha_k)^2}{\alpha_1 + 2\sum_{i=2}^{k-1} \alpha_i + \alpha_k} \\ &= \frac{2(\alpha_1 + \alpha_k)}{\alpha_1 + 2\sum_{i=2}^{k-1} \alpha_i + \alpha_k} \frac{(2(\gamma_1 - \gamma_k) - \alpha_1 + \alpha_k)^2}{2(\alpha_1 + \alpha_k)} \\ &= \frac{2(\alpha_1 + \alpha_k)}{\alpha_1 + 2\sum_{i=2}^{k-1} \alpha_i + \alpha_k} \mathrm{adv}_{1,k} \end{aligned}$$

as claimed. This factor of two can be replicated throughout the remainder of the proof to give the claimed result. The best possible constant factor is 1, as argued in [48].

A.3 IND-CPA-D and KR-D Games

In this appendix we give the formal definition of IND-CPA-D and KR-D security, following [41,42].

Algorithm 2 Oracles for the IND-CPA-D game

initialization:
  $(\mathsf{evk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$
  $S \leftarrow \emptyset$
  $i \leftarrow 0$

$E^b(m_0, m_1) :=$
  $\mathsf{ct} \leftarrow \mathsf{Enc}(m_b, \mathsf{sk})$
  $S[i] \leftarrow (m_0, m_1, \mathsf{ct})$
  $i \leftarrow i + 1$
  return ct

$H(C, J = (j_1, \ldots, j_k)) :=$
  $\mathsf{ct} \leftarrow \mathsf{Eval}\,(C, S[j_1].\mathsf{ct}, \ldots, S[j_k].\mathsf{ct})$
  $Cm_0 \leftarrow C\,(S[j_1].m_0, \ldots, S[j_k].m_0)$
  $Cm_1 \leftarrow C\,(S[j_1].m_1, \ldots, S[j_k].m_1)$
  $S[i] \leftarrow (Cm_0, Cm_1, \mathsf{ct})$
  $i \leftarrow i + 1$
  return ct

$D(i) :=$
  if $S[i].m_0 = S[i].m_1$:
    return $\mathsf{Dec}(S[i].\mathsf{ct}, \mathsf{sk})$
  else:
    return $\perp$

Algorithm 3 Oracles for the KR-D game

initialization:
  $(\mathsf{evk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$
  $S \leftarrow \emptyset$
  $i \leftarrow 0$

$E(m) :=$
  $\mathsf{ct} \leftarrow \mathsf{Enc}(m, \mathsf{sk})$
  $S[i] \leftarrow \mathsf{ct}$
  $i \leftarrow i + 1$
  return ct

$H(C, J = (j_1, \ldots, j_k)) :=$
  $\mathsf{ct} \leftarrow \mathsf{Eval}\,(C, S[j_1].\mathsf{ct}, \ldots, S[j_k].\mathsf{ct})$
  $S[i] \leftarrow \mathsf{ct}$
  $i \leftarrow i + 1$
  return ct

$D(i) :=$
  return $\mathsf{Dec}(S[i].\mathsf{ct}, \mathsf{sk})$

Definition 16. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be an HE scheme, and E^b, D, H be as in Algorithm 2. Define the IND-CPA game to be the indistinguishability game instantiated with distributions $\{(E^0, H)\}$ and $\{(E^1, H)\}$, and the IND-CPA-D game to be the indistinguishability game instantiated with distributions $\{(E^0, H, D)\}$ and $\{(E^1, H, D)\}$.*

Definition 17. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be an HE scheme, and E, D, H be as in Algorithm 3. We define the KR game as the search game instantiated with $x = \mathsf{sk}$ and oracle access to (E, H), while we define the KR-D game as the search game with $x = \mathsf{sk}$ and oracle access to (E, H, D).*

A.4 Convolution Lemma Divergence Proof

In this Appendix we prove Lemma 8, following [52].

Lemma. Let $\Sigma_1, \Sigma_2 \in \mathbb{R}^{n \times n}$ be symmetric positive definite matrices such that $\Sigma_3^{-1} = \Sigma_1^{-1} + \Sigma_2^{-1}$ satisfies $\sqrt{\Sigma_3} \geq \eta_\varepsilon(\mathbb{Z}^n)$ for $0 < \varepsilon < 1$. Then for $\mu_1, \mu_2 \in \mathbb{R}^n$ the random variables

$$Z = X + Y, \quad X \leftarrow \mathcal{D}_{\mu_1, \Sigma_1}, Y \leftarrow \mathcal{D}_{\mu_2, \Sigma_2}$$
$$Z', \quad Z' \leftarrow \mathcal{D}_{\mu_1 + \mu_2, \Sigma_1 + \Sigma_2}$$

satisfy, for $a \neq 1$,

$$D(Z\|Z'), D(Z'\|Z) \in \left[\ln\frac{1-\varepsilon}{1+\varepsilon}, \ln\frac{1+\varepsilon}{1-\varepsilon}\right],$$
$$R_a(Z\|Z'), R_a(Z\|Z') \in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right],$$
$$R_\infty(Z\|Z'), R_\infty(Z'\|Z) \in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right].$$

In this section we use the notation $\rho_\Sigma(S)$ for a lattice coset S to denote $\sum_{x\in S}\rho_\Sigma(x)$. We will use the following fact about the smoothing parameter, which is implicit in the proof of Lemma 4.4 of [46] and given as Lemma 2.4 of [52].

Lemma 16. *Let Λ be an n-dimensional lattice. For $0 < \varepsilon < 1$, if $\Sigma > 0$ and $\sqrt{\Sigma} \geq \eta_\varepsilon(\Lambda)$, then for any $c \in \mathbb{R}^n$,*

$$\rho_\Sigma(\Lambda + c) \in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1\right]\rho_\Sigma(\Lambda)$$

Finally, we observe the following (a special case of Fact 2.1 of [52]).

Lemma 17. *Let $x, z \in \mathbb{Z}^n$, $\Sigma_1, \Sigma_2 > 0$. Then*

$$\rho_{\Sigma_1}(x-\mu_1)\rho_{\Sigma_2}(z-x-\mu_2) = \rho_{\Sigma_0}(z-(\mu_1+\mu_2))\rho_{\Sigma_3}(x-c(z)),$$

where $\Sigma_0 = \Sigma_1+\Sigma_2$, $\Sigma_3^{-1} = \Sigma_1^{-1}+\Sigma_2^{-1}$, and $c(z) = \Sigma_3\left(\Sigma_1^{-1}\mu_1 + \Sigma_2^{-1}(z-\mu_2)\right)$.

We now proceed to prove the main convolution lemma. We start by calculating:

$$\begin{aligned}
\Pr[X_1 = \mathbf{x}_1 \wedge Z = \mathbf{z}] &= \Pr[X_1 = \mathbf{x}_1 \wedge X_2 = (\mathbf{z}-\mathbf{x}_1)] \\
&= \Pr[X_1 = \mathbf{x}_1]\Pr[X_2 = (\mathbf{z}-\mathbf{x}_1)] \\
&= \frac{\rho_{\Sigma_1}(\mathbf{x}_1-\mu_1)\rho_{\Sigma_2}(\mathbf{z}-\mathbf{x}_1-\mu_2)}{\rho_{\Sigma_1}(\mu_1+\mathbb{Z}^n)\rho_{\Sigma_2}(\mu_2+\mathbb{Z}^n)} \\
&= \frac{\rho_{\Sigma_0}(\mathbf{z}-(\mu_1+\mu_2))\rho_{\Sigma_3}(\mathbf{x}_1-c(\mathbf{z}))}{\rho_{\Sigma_1}(\mu_1+\mathbb{Z}^n)\rho_{\Sigma_2}(\mu_2+\mathbb{Z}^n)} \quad \text{(by Lemma 17)}
\end{aligned}$$

so that, letting $\mu_0 = \mu_1+\mu_2$, for any $\mathbf{z} \in \mathbb{Z}^n$ we have that

$$\Pr[Z = \mathbf{z}] = \sum_{\mathbf{x}_1\in\mathbb{Z}^n}\Pr[Z = \mathbf{z} \wedge X = \mathbf{x}_1] = \frac{\rho_{\Sigma_0}(\mathbf{z}-\mu_0)\rho_{\Sigma_3}(c(\mathbf{z})+\mathbb{Z}^n)}{\rho_{\Sigma_1}(\mu_1+\mathbb{Z}^n)\rho_{\Sigma_2}(\mu_2+\mathbb{Z}^n)}.$$

As this is a probability distribution, it must sum to 1, which implies

$$\begin{aligned}
\rho_{\Sigma_1}(\mu_1+\mathbb{Z}^n)\rho_{\Sigma_2}(\mu_2+\mathbb{Z}^n) &= \sum_{\mathbf{z}\in\mathbb{Z}^n}\rho_{\Sigma_3}(c(\mathbf{z})+\mathbb{Z}^n)\rho_{\Sigma_0}(\mathbf{z}-\mu_0) \\
&\in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1\right]\rho_{\Sigma_3}(\mathbb{Z}^n)\sum_{\mathbf{z}\in\mathbb{Z}^n}\rho_{\Sigma_0}(\mathbf{z}-\mu_0) \\
&= \left[\frac{1-\varepsilon}{1+\varepsilon}, 1\right]\rho_{\Sigma_3}(\mathbb{Z}^n)\rho_{\Sigma_0}(\mu_0+\mathbb{Z}^n),
\end{aligned}$$

where we have used that since Σ_3 is larger than the smoothing parameter for ε, we can use Lemma 16.

To prove D, R_a and R_∞ are bounded as claimed, we bound the ratio $\frac{\Pr[Z=z]}{\Pr[Z'=z]}$. Observe –

$$\begin{aligned}
\frac{\Pr[Z=\mathbf{z}]}{\Pr[Z'=\mathbf{z}]} &= \frac{\rho_{\Sigma_0}(\mathbf{z}-\mu_0)\rho_{\Sigma_3}(c(\mathbf{z})+\mathbb{Z}^n)}{\rho_{\Sigma_1}(\mu_1+\mathbb{Z}^n)\rho_{\Sigma_2}(\mu_2+\mathbb{Z}^n)} \cdot \frac{\rho_{\Sigma_0}(\mu_0+\mathbb{Z}^n)}{\rho_{\Sigma_0}(\mathbf{z}-\mu_0)} \\
&= \frac{\rho_{\Sigma_0}(\mu_0+\mathbb{Z}^n)\rho_{\Sigma_3}(c(\mathbf{z})+\mathbb{Z}^n)}{\rho_{\Sigma_1}(\mu_1+\mathbb{Z}^n)\rho_{\Sigma_2}(\mu_2+\mathbb{Z}^n)} \\
&\in \left[1, \frac{1+\varepsilon}{1-\varepsilon}\right] \frac{\rho_{\Sigma_0}(\mu_0+\mathbb{Z}^n)\rho_{\Sigma_3}(c(\mathbf{z})+\mathbb{Z}^n)}{\rho_{\Sigma_0}(\mu_0+\mathbb{Z}^n)\rho_{\Sigma_3}(\mathbb{Z}^n)} \\
&= \left[1, \frac{1+\varepsilon}{1-\varepsilon}\right] \frac{\rho_{\Sigma_3}(c(\mathbf{z})+\mathbb{Z}^n)}{\rho_{\Sigma_3}(\mathbb{Z}^n)} \\
&\in \left[\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right].
\end{aligned}$$

The bounds on $R_\infty(Z\|Z'), R_\infty(Z'\|Z)$ follow immediately. For the KL divergence, by definition we have:

$$\begin{aligned}
D(Z\|Z') &= \sum_{\mathbf{z}} \Pr[Z=\mathbf{z}] \ln \frac{\Pr[Z=\mathbf{z}]}{\Pr[Z'=\mathbf{z}]} \\
&\in \left[\ln \frac{1-\varepsilon}{1+\varepsilon}, \ln \frac{1+\varepsilon}{1-\varepsilon}\right] \sum_{\mathbf{z}} \Pr[Z=\mathbf{z}] \\
&= \left[\ln \frac{1-\varepsilon}{1+\varepsilon}, \ln \frac{1+\varepsilon}{1-\varepsilon}\right],
\end{aligned}$$

While for the Rényi divergence,

$$\begin{aligned}
(a-1)\log R_a(Z\|Z') &= \log \sum_{\mathbf{z}} \Pr[Z=\mathbf{z}] \left(\frac{\Pr[Z=\mathbf{z}]}{\Pr[Z'=\mathbf{z}]}\right)^{a-1} \\
&\in (a-1)\left[\log \frac{1-\varepsilon}{1+\varepsilon}, \log \frac{1+\varepsilon}{1-\varepsilon}\right] + \log \sum_{\mathbf{z}} \Pr[Z=\mathbf{z}] \\
&= (a-1)\left[\log \frac{1-\varepsilon}{1+\varepsilon}, \log \frac{1+\varepsilon}{1-\varepsilon}\right],
\end{aligned}$$

and similarly for $Z'\|Z$. From here the claim follows.

B Proofs for Security for Game Hops

B.1 Indistinguishability Games

We start with a lemma which enables lower bounding $1-(1-\alpha_{\mathcal{Q}})^n$ when we set $n = \left\lfloor \frac{1}{\alpha_{\mathcal{Q}}} \right\rfloor$. In [47], the authors set $n = \frac{1}{\alpha_{\mathcal{Q}}}$, and so are able to use a lower

bound of $1 - e^{-1}$. In cases that $\frac{1}{\alpha_{\mathcal{Q}}} \notin \mathbb{Z}$, we prefer our more conservative lower bound, and find it does not affect the strength of the overall result.

Lemma 18. *Let $x \in (0, 1]$. Then $\frac{1}{2} \leq 1 - (1-x)^{\left\lfloor \frac{1}{x} \right\rfloor}$, and this bound is best possible.*

Proof. Let $n = \left\lfloor \frac{1}{x} \right\rfloor \geq 1$, so that by definition $\frac{1}{n+1} < x \leq \frac{1}{n}$. Now

$$1 - \left(1 - \frac{1}{n+1}\right)^n < 1 - (1-x)^n,$$

and we claim that $1-(1-\frac{1}{n+1})^n \geq \frac{1}{2}$. Indeed, first observe that for $n = 1$, we have equality, while for $n > 1$, $n \ln\left(1 - \frac{1}{n+1}\right)$ is decreasing, so that $\left(1 - \frac{1}{n+1}\right)^n \leq \frac{1}{2}$, and the result follows.

To show the bound is best possible, consider $x \to \frac{1}{2}^+$. Then $\left\lfloor \frac{1}{x} \right\rfloor = 1$, and $1 - (1-x)^1 \to \frac{1}{2}$.

We are now able to prove the theorem, which we restate below.

Theorem. Consider a distinguishing game $\mathcal{G}^{\mathcal{P}}$ in which the adversary has black box access to a probability distribution $\mathcal{P}$. Suppose further that $D(\mathcal{P}\|\mathcal{Q}) \leq 2^{-\lambda-3}$ for $\lambda \leq \kappa - 4$. Then if $\mathcal{G}^{\mathcal{P}}$ is κ-bit secure, $\mathcal{G}^{\mathcal{Q}}$ is λ-bit secure.

Proof. Our proof is similar to that of Theorem 8 of [48]: we first fix an adversary $\mathcal{A}$ which has runtime $T(\mathcal{A})$, and let $\alpha_{\mathcal{P}}, \alpha_{\mathcal{Q}}, \beta_{\mathcal{P}}, \beta_{\mathcal{Q}}$ be the output and conditional success probabilities in games $\mathcal{G}^{\mathcal{P}}$ and $\mathcal{G}^{\mathcal{Q}}$. By assumption,

$$\frac{\alpha_{\mathcal{P}}(2\beta_{\mathcal{P}} - 1)^2}{T(\mathcal{A})} \leq 2^{-\kappa} \implies \frac{2\beta_{\mathcal{P}} - 1}{\sqrt{T(\mathcal{A})}} \leq \frac{2^{-\kappa/2}}{\sqrt{\alpha_{\mathcal{P}}}}.$$

We also observe that if $\mathcal{A}$ makes t calls to $\mathcal{P}$ or $\mathcal{Q}$, then we must have[7] $t \leq T(\mathcal{A})$. Now,

$$\begin{aligned}
\beta_{\mathcal{Q}} &\leq \beta_{\mathcal{P}} + \Delta(\mathcal{G}^{\mathcal{P}}, \mathcal{G}^{\mathcal{Q}}) && \text{[defn. statistical distance]} \\
&\leq \beta_{\mathcal{P}} + \sqrt{D(\mathcal{G}^{\mathcal{P}}, \mathcal{G}^{\mathcal{Q}})/2} && \text{[Pinsker's inequality]} \\
&\leq \beta_{\mathcal{P}} + \sqrt{T(\mathcal{A}) D(P\|Q)/2} && \text{[Lemma 4 and } t \leq T(\mathcal{A})\text{]}
\end{aligned}$$

We can now use the above inequality to bound $\frac{\mathrm{adv}_{\mathcal{A}}}{T(\mathcal{A})} = \frac{\alpha_{\mathcal{Q}}(2\beta_{\mathcal{Q}}-1)^2}{T(\mathcal{A})}$.

$$\begin{aligned}
\frac{\sqrt{\alpha_{\mathcal{Q}}}(2\beta_{\mathcal{Q}} - 1)}{\sqrt{T(\mathcal{A})}} &\leq \frac{\sqrt{\alpha_{\mathcal{P}}}(2\beta_{\mathcal{P}} - 1)}{\sqrt{T(\mathcal{A})}} + \sqrt{2\alpha_{\mathcal{Q}} D(\mathcal{P}\|\mathcal{Q})} \\
&\leq \sqrt{\frac{\alpha_{\mathcal{Q}}}{\alpha_{\mathcal{P}}}} 2^{-\kappa/2} + \sqrt{2D(\mathcal{P}\|\mathcal{Q})} \\
&\leq \sqrt{\frac{\alpha_{\mathcal{Q}}}{\alpha_{\mathcal{P}}}} 2^{-\kappa/2} + 2^{\frac{-\lambda}{2}-1} \\
&\leq \max\left(\sqrt{\frac{\alpha_{\mathcal{Q}}}{\alpha_{\mathcal{P}}}} 2^{1-\kappa/2}, 2^{\frac{-\lambda}{2}}\right)
\end{aligned}$$

[7] This claim is standard in the literature: consider e.g. [47,7].

If $2^{\frac{-\lambda}{2}} \geq \sqrt{\frac{\alpha_{\mathcal{Q}}}{\alpha_{\mathcal{P}}}} 2^{1-\kappa/2}$, we are done, since now

$$\frac{\alpha_{\mathcal{Q}}(2\beta_{\mathcal{Q}}-1)^2}{T(\mathcal{A})} \leq 2^{-\lambda}$$

as required.

Otherwise assume, seeking contradiction, that $\frac{\alpha_{\mathcal{Q}}(2\beta_{\mathcal{Q}}-1)^2}{T(\mathcal{A})} > 2^{-\lambda}$. Observe we may assume $\alpha_{\mathcal{Q}} > 0$. We will show that $\frac{\alpha_{\mathcal{P}}}{\alpha_{\mathcal{Q}}} > 2^{2-(\kappa-\lambda)}$, which implies $\frac{\alpha_{\mathcal{Q}}}{\alpha_{\mathcal{P}}} 2^{2-\kappa} < 2^{-\lambda}$, giving the required contradiction.

We proceed as in the proof of Lemma 3.3 of [47], considering the hypothetical games $[\mathcal{G}^{\mathcal{P}}]^n$ and $[\mathcal{G}^{\mathcal{Q}}]^n$ which consist of n independent copies of $\mathcal{G}^{\mathcal{P}}$ and $\mathcal{G}^{\mathcal{Q}}$. Let $\alpha_{\mathcal{P}}^{(n)}$ and $\alpha_{\mathcal{Q}}^{(n)}$ be the probability that $\mathcal{A}$ has an output in at least one copy of the game. Then, by Pinsker's inequality and the union bound,

$$\begin{aligned} 1-(1-\alpha_{\mathcal{Q}})^n &= \alpha_{\mathcal{Q}}^{(n)} \\ &\leq \alpha_{\mathcal{P}}^{(n)} + \Delta([\mathcal{G}^{\mathcal{P}}]^n, [\mathcal{G}^{\mathcal{Q}}]^n) \\ &\leq n\alpha_{\mathcal{P}} + \sqrt{nT(\mathcal{A})D(\mathcal{P}\|\mathcal{Q})/2} \end{aligned}$$

now let $n = \left\lfloor \frac{1}{\alpha_{\mathcal{Q}}} \right\rfloor$, so that by Lemma 18, $\frac{1}{2} \leq 1-(1-\alpha_{\mathcal{Q}})^n$, as well as $n\alpha_{\mathcal{P}} \leq \frac{\alpha_{\mathcal{P}}}{\alpha_{\mathcal{Q}}}$. Our contradiction hypothesis $\frac{\alpha_{\mathcal{Q}}(2\beta_{\mathcal{Q}}-1)^2}{T(\mathcal{A})} > 2^{-\lambda}$ also gives that $nT(\mathcal{A}) \leq 2^{\lambda}(2\beta_{\mathcal{Q}}-1)^2 \leq 2^{\lambda}$. Putting all of this together with the bound on $D(\mathcal{P}\|\mathcal{Q})$ gives:

$$\begin{aligned} \frac{1}{2} &< \frac{\alpha_{\mathcal{P}}}{\alpha_{\mathcal{Q}}} + \sqrt{2^{\lambda}2^{-\lambda-4}} \\ &= \frac{\alpha_{\mathcal{P}}}{\alpha_{\mathcal{Q}}} + \frac{1}{4} \end{aligned}$$

so that $\frac{\alpha_{\mathcal{P}}}{\alpha_{\mathcal{Q}}} > \frac{1}{4} \geq 2^{2-l}$ for all $l \geq 4$, as required.

We also prove the following result for the R_∞ distance (Theorem 2).

Theorem. Consider two distinguishing games $\mathcal{G}^{\mathcal{P}}, \mathcal{G}^{\mathcal{Q}}$. Then if $\mathcal{G}^{\mathcal{P}}$ is κ-bit secure and $R_\infty\left(\mathcal{G}^{\mathcal{Q}}\|\mathcal{G}^{\mathcal{P}}\right) \leq 2^{(l-\kappa-2)/2}+1$ for $\kappa > l \geq 4$ then $\mathcal{G}^{\mathcal{Q}}$ is $(\kappa-l)$-bit secure.

Proof. We fix an adversary $\mathcal{A}$ which has runtime $T(\mathcal{A})$, and let $\alpha_{\mathcal{P}}, \alpha_{\mathcal{Q}}, \beta_{\mathcal{P}}, \beta_{\mathcal{Q}}$ be the output and conditional success probabilities in games $\mathcal{G}^{\mathcal{P}}$ and $\mathcal{G}^{\mathcal{Q}}$ respectively, and for clarity write $R = R_\infty\left(\mathcal{G}^{\mathcal{Q}}\|\mathcal{G}^{\mathcal{P}}\right)$. Then by the Probability

Preservation property of R_∞, we have

$$\begin{aligned}
& & \beta_{\mathcal{Q}} &\leq R\beta_{\mathcal{P}} \\
\implies & & 2\beta_{\mathcal{Q}} - 1 &\leq R(2\beta_{\mathcal{P}} - 1) + R - 1 \\
\implies & & \frac{\sqrt{\alpha_{\mathcal{Q}}}(2\beta_{\mathcal{Q}} - 1)}{\sqrt{T(\mathcal{A})}} &\leq \frac{\sqrt{\alpha_{\mathcal{Q}}(2\beta_{\mathcal{P}} - 1)}}{\sqrt{T(\mathcal{A})}} + R - 1 \\
& & &\leq R^{3/2}\frac{\sqrt{\alpha_{\mathcal{P}}(2\beta_{\mathcal{P}} - 1)}}{\sqrt{T(\mathcal{A})}} + R - 1 \\
& & &\leq 2\max\left(R^{3/2}2^{-\kappa/2}, R - 1\right)
\end{aligned}$$

where we have used that by assumption $\frac{\alpha_{\mathcal{P}}(2\beta_{\mathcal{P}}-1)^2}{T(\mathcal{A})} \leq 2^{-\kappa}$. To conclude, we need to argue that both $2R^{3/2} \leq 2^{l/2}$ and $2(R-1) \leq 2^{(l-\kappa)/2}$. The second follows directly from the condition on R. The first follows from the same condition, since this implies that $R \leq 3/2$, so that $R^3 \leq 2^{l-2}$ for any $l \geq 4$.

B.2 Search Games

We start with a proof of Lemma 9.

Lemma. Suppose a search game S_0 is κ-bit secure, and let $R(a) = R_a(S_1\|S_0)$. Then if there is $a > 1$ such that $\kappa \geq l \geq \left(1 - \frac{1}{a}\right)\log R(a) + \frac{\kappa}{a}$, we have that S_1 is $(\kappa - l)$-bit secure.

Proof. Fix an adversary $\mathcal{A}$ with resources $T(\mathcal{A})$, and let $\alpha_0, \beta_0, \alpha_1, \beta_1$ be $\mathcal{A}$'s conditional output and success probabilities in games S_0 and S_1 respectively. We have that $\alpha_i\beta_i = \Pr[\mathcal{A} \text{ succeeds in game } S_i]$, so that, by the probability preservation property,

$$\alpha_0\beta_0 \geq (\alpha_1\beta_1)^{\frac{a}{a-1}}/R(a) \implies (\alpha_0\beta_0 R(a))^{1-\frac{1}{a}} \geq \alpha_1\beta_1.$$

Therefore:

$$\begin{aligned}
\frac{\alpha_1\beta_1}{T(\mathcal{A})} &\leq \frac{(\alpha_0\beta_0)^{1-\frac{1}{a}}}{T(\mathcal{A})}R(a)^{1-\frac{1}{a}} \\
&\leq \left(\frac{\alpha_0\beta_0}{T(\mathcal{A})}\right)^{1-\frac{1}{a}} R(a)^{1-\frac{1}{a}} \\
&\leq 2^{\frac{\kappa}{a}-\kappa}R(a)^{1-\frac{1}{a}} \\
&\leq 2^{l-\kappa},
\end{aligned}$$

where the first and second inequality we have used that $T(\mathcal{A}) \geq 1, a > 0$ implies that $\frac{1}{T(\mathcal{A})} \leq \left(\frac{1}{T(\mathcal{A})}\right)^{1-1/a}$. Since the choice of adversary $\mathcal{A}$ was arbitrary, we can conclude that $\mathcal{G}_1$ is $(l - \kappa)$ bit secure as claimed.

We also prove Lemma 10.

Lemma. Suppose a search game S_0 is κ-bit secure, and that $R_\infty(S_1\|S_0) \leq 2^l$ with $\kappa > l$. Then S_1 is $(\kappa - l)$-bit secure.

Proof. This follows by a routine application of the probability preservation property: we again fix an adversary $\mathcal{A}$ with resources $T(\mathcal{A})$, and let $\alpha_0, \beta_0, \alpha_1, \beta_1$ be $\mathcal{A}$'s conditional output and success probabilities in games S_0 and S_1. Then, since $\alpha_i\beta_i = \Pr[\mathcal{A} \text{ succeeds in game } S_i]$,

$$\alpha_0\beta_0 \geq \alpha_1\beta_1 / R_\infty(S_1\|S_0),$$

so that

$$\frac{\alpha_1\beta_1}{T(\mathcal{A})} \leq \frac{\alpha_0\beta_0}{T(\mathcal{A})} R_\infty(S_1\|S_0) \leq 2^{l-\kappa}$$

as required.

C Noise Flooding Proofs

C.1 IND-CPA-D Setting (Proof of Theorem 3)

Theorem. Let Π be an FHE scheme with plaintext space $\mathcal{M} \subset \mathbb{R}^n$, and consider the scheme $\tilde{\Pi}$ which uses a discrete or continuous Gaussian noise flooding mechanism with variance

$$\sigma^2 \geq 4E^2 2^\lambda$$

where E is a static upper bound on the error $\|e\|_2$, and $\lambda \leq \kappa - 4$. Then if Π is κ-bit secure in the IND-CPA game, $\tilde{\Pi}$ is λ-bit secure in the IND-CPA-D game.

Proof. Our proof follows the same strategy as the proofs of Lemma 10 in [11] or Theorem 2 in [42], namely we consider three key recovery games:

- $\mathcal{G}_0$: the left-or-right game in the IND-CPA model, where the adversary does not have access to a decryption oracle;
- $\mathcal{G}_1$: the left-or-right game in the IND-CPA-D model, where the decryption oracle outputs $m + \epsilon$, for m the true underlying message, and ϵ drawn from the appropriate Gaussian;
- $\mathcal{G}_2$: the left-or-right game in the IND-CPA-D model, where the decryption oracle outputs $\mathrm{Dec}(\mathtt{ct}) + \epsilon$.

Observe that, since the noise estimate is static, $\mathcal{G}_1$ is fully simulatable by any adversary to game $\mathcal{G}_0$, implying that $\mathcal{G}_1$ is κ-bit secure.

From here, we instantiate $\mathcal{P}$ and $\mathcal{Q}$ in Theorem 1 as follows. On input a valid index i, $\mathcal{P}$ outputs $m_i + \epsilon_i$ with $\epsilon_i \leftarrow \mathcal{D}_{\sigma_i^2}$, while $\mathcal{Q}$ outputs $\mathrm{Dec}(\mathtt{ct}_i) + \epsilon_i = m_i + e_i + \epsilon_i$. Then

$$D(\mathcal{P}\|\mathcal{Q}) = D(\mathcal{D}_{m_i,\sigma_i^2}\|\mathcal{D}_{m_i+e_i,\sigma_i^2}) \leq \frac{\|e_i\|_2^2}{2\sigma_i^2} \leq 2^{-\lambda-3},$$

so that Theorem 1 gives that $\mathcal{G}_2$ is λ-bit secure as claimed.

C.2 KR-D Setting (Proof of Theorem 4)

We start by giving a lemma which gives a tight bound on the amount of flooding noise required in search games. We observe that the exact same result applies for flooding noise drawn from either discrete or continuous Gaussians, using Lemma 15 in place of Lemma 14.

Lemma 19. *Let $\kappa > l > 0$. Consider two search games, S_0 and S_1, where the adversary makes at most t queries to an oracle O which takes as input an index i. In S_0, the oracle outputs:*

$$m_i + \epsilon_i, \quad \epsilon_i \leftarrow \mathcal{D}^n_{\sigma_i^2}$$

while in S_1, the oracle outputs

$$m_i + e_i + \epsilon_i, \quad \|e_i\|_2 \leq E_i, \quad \epsilon \leftarrow \mathcal{D}^n_{\sigma_i^2}$$

where again m_i is known to the adversary, and where the bounds E_i holds over all runs of the game S_1. Then if S_0 is κ-bit secure and

$$\sigma_i^2 \geq \frac{2\kappa t E_i^2 \log e}{l^2}$$

then S_1 is $(\kappa - l)$-bit secure.

Proof. Our proof uses Lemmas 9 and 14: we show that there exists $a > 1$ such that

$$\left(1 - \frac{1}{a}\right) \log R_a(S_1 \| S_0) + \frac{\kappa}{a} \leq l;$$

from here we can conclude via Lemma 9. We set

$$a = \frac{\kappa}{l} + \frac{\sqrt{\kappa(\kappa - l)}}{l}.$$

Then $a > 1$ since $\kappa > l$. Now, using Lemmas 4 and 14, we have

$$\begin{aligned}
\left(1 - \frac{1}{a}\right) \log R_a(S_1 \| S_0) + \frac{\kappa}{a} &\leq \left(1 - \frac{1}{a}\right) \sum_{i=1}^{t} \log R_a(e_i + \epsilon_i \| \epsilon_i) + \frac{\kappa}{a} \\
&\leq (a-1) \sum_{i=1}^{t} \frac{\|e_i\|_2^2}{2\sigma_i^2} \log e + \frac{\kappa}{a} \\
&\leq \frac{l^2(a-1)}{4\kappa} + \frac{\kappa}{a} \\
&= l \left(1 - \frac{\sqrt{\kappa - l}\left(4\kappa - (\sqrt{\kappa} + \sqrt{\kappa - l})^2\right)}{4\kappa(\sqrt{\kappa} + \sqrt{\kappa - l})}\right) \\
&\leq l,
\end{aligned}$$

with the final deduction following from $(\sqrt{\kappa} + \sqrt{\kappa - l})^2 \leq 4\kappa$.

We can now prove Theorem 4.

Theorem. Let Π be an HE scheme with plaintext space $\mathcal{M} \subset \mathbb{Z}^n$, and consider the scheme $\tilde{\Pi}$ which uses the discrete or continuous Gaussian noise flooding mechanism with variance

$$\sigma^2 \geq \frac{2\kappa t E^2 \log e}{l^2}$$

where the adversary is allowed at most t queries to its decryption oracle, E is a static upper bound on the error $\|e\|_2$, and $1 \leq l < \kappa$. Then if Π is κ-bit secure in the KR game, $\tilde{\Pi}$ is $(\kappa - l)$-bit secure in the KR-D game.

Proof. Our proof again considers three games:

- $\mathcal{G}_0$: the key recovery game in the IND-CPA model, where the adversary does not have access to a decryption oracle;
- $\mathcal{G}_1$: the key recovery game in the IND-CPA-D model, where the decryption oracle outputs $m + \epsilon$, for m the true underlying message, and ϵ drawn from the appropriate Gaussian;
- $\mathcal{G}_2$: the key recovery game in the IND-CPA-D model, where the decryption oracle outputs $\mathrm{Dec}(ct) + \epsilon$.

First observe that $\mathcal{G}_1$ is fully simulatable by any adversary to game $\mathcal{G}_0$. We therefore have that $\mathcal{G}_1$ is κ-bit secure. Now observe that for an index i, the decryption oracle outputs $m_i + \epsilon_i$, while in $\mathcal{G}_2$ the oracle outputs $m_i + e_i + \epsilon_i$. we therefore apply Lemma 19 and conclude that $\mathcal{G}_1$ is $(k - l)$-bit secure.

D From HintLWE to IND-CPA-D and KR-D Security (Proofs of Theorems 5 and 6)

Theorem. Let $\kappa' - \lambda \geq 4 + \log 12$, and let Π be an LWE-HE scheme with parameters $m, n, Q, \chi_{\mathbf{s}}, \chi_{\mathbf{e}}$. Assume further that the HintLWE (Decision) problem with the same parameters and with flooding variance σ^2 for hint matrix $\frac{1}{q}[\mathbf{A}_j]_q$ is κ'-bit secure. Then if we use flooding noise sampled from a continuous Gaussian with variance

$$\sigma_j^2 \geq \max\left(4E^2 2^\lambda, \sigma^2\right)$$

on input $\mathsf{Rescale}\left((\mathbf{A}_j, \mathbf{b}_j = \mathbf{A}_j\mathbf{s} + \mathbf{e}_j + \mathbf{m}_j), q\right)$ with $E \geq \frac{1}{q}\left\|\mathbf{e}_j + [\mathbf{m}_j]_q\right\|_2$ a static upper bound, then $\tilde{\Pi}$ is λ-bit secure in the IND-CPA-D game where we only decrypt rescaled ciphertexts.

Proof. We consider three games, in which the adversary view is as follows:

$$\mathcal{G}_0 : \left(\mathbf{A}, \mathbf{u} + \mathbf{m}^{(b)} \mod Q, \left\{\left\lceil \frac{\mathbf{m}_j}{q} \right\rfloor + \frac{[\mathbf{A}_j]_q}{q}\mathbf{s} - \frac{[\mathbf{b}_j]_q}{q} + \epsilon_j\right\}_{j=0,\ldots,t-1}\right), \text{ for } \mathbf{u} \leftarrow \mathcal{U}(\mathbb{Z}_q^m)$$

$$\mathcal{G}_1 : \left(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{m}^{(b)} \mod Q, \left\{\left\lceil \frac{\mathbf{m}_j}{q} \right\rfloor + \frac{[\mathbf{A}_j]_q}{q}\mathbf{s} - \frac{[\mathbf{b}_j]_q}{q} + \epsilon_j\right\}_{j=0,\ldots,t-1}\right)$$

$$\mathcal{G}_2 : \left(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{m}^{(b)} \mod Q, \left\{ \left\lceil \frac{\mathbf{m}_j}{q} \right\rfloor + \frac{[\mathbf{A}_j]_q}{q}\mathbf{s} - \frac{[\mathbf{b}_j]_q}{q} + \frac{\mathbf{e}_j + [\mathbf{m}_j]_q}{q} + \epsilon_j \right\}_{j=0,\ldots,t-1}\right)$$

where the first components give the output of the encryption oracle, and the final component gives the outputs of the decryption oracle. In each game the adversary is challenged to decide whether $b = 1$ or $b = 0$. By construction of the decryption oracle, the decrypted messages $\left\lceil \frac{\mathbf{m}_j}{q} \right\rfloor$ are known to the adversary.

We recognise $\mathcal{G}_2$ as $\tilde{\Pi}$ in the IND-CPA-D game with a decryption oracle restricted to rescaled ciphertexts. In $\mathcal{G}_1$, the decryption oracle does not perform an LWE decryption, but instead uses the publicly computable $\mathbf{A}_j$, $\mathbf{b}_j$ and the secret key $\mathbf{s}$ to calculate the output. Lastly, we recognise $\mathcal{G}_0$ as a one time pad encryption, so that even an unbounded adversary has zero advantage.

We start by lower bounding the bit security of $\mathcal{G}_1$. Define distributions $\mathcal{H}_i$ for $i = 1, ..., 4$ as follows.

$\mathcal{H}_1$: the view of $\mathcal{G}_0$ in the case $b = 0$
$\mathcal{H}_2$: the view of $\mathcal{G}_1$ in the case $b = 0$
$\mathcal{H}_3$: the view of $\mathcal{G}_1$ in the case $b = 1$
$\mathcal{H}_4$: the view of $\mathcal{G}_0$ in the case $b = 1$

and let $\mathrm{adv}_{i,i+1}$ be the advantage of a bounded adversary $\mathcal{A}$ in distinguishing $\mathcal{H}_i$ and $\mathcal{H}_{i+1}$. Distinguishing $\mathcal{H}_2$ and $\mathcal{H}_3$ is the game $\mathcal{G}_1$, while as discussed, $\mathrm{adv}_{1,4} = 0$. Therefore, Lemma 5 gives that

$$\frac{\mathrm{adv}_{\mathcal{G}_1}}{T(\mathcal{A})} \leq 6\left(\frac{\mathrm{adv}_{1,2}}{T(\mathcal{A})} + \frac{\mathrm{adv}_{3,4}}{T(\mathcal{A})}\right)$$

since $\mathbf{m}^{(0)}$ and $\mathbf{m}^{(1)}$ are known to the adversary and $\mathbf{b}_i$ is publicly computable, any adversary that can distinguish $\mathcal{H}_1$ and $\mathcal{H}_2$ or $\mathcal{H}_3$ and $\mathcal{H}_4$ is an adversary in the HintLWE decision game, with $\mathbf{X}_j = \frac{[\mathbf{A}_j]_q}{q}$, so that by definition of σ_j^2, $\frac{\mathrm{adv}_{\mathcal{G}_1}}{T(\mathcal{A})} \leq 12 \cdot 2^{-\kappa'}$, so that $\mathcal{G}_1$ is $(\kappa' - \log 12)$-bit secure.

Now taking $\mathcal{P}$ and $\mathcal{Q}$ as the decryption oracles in games $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively, then by definition of σ_j^2 and Lemma 15,

$$D(\mathcal{P}\|\mathcal{Q}) \leq \frac{E^2}{2\sigma_i^2} \leq 2^{-\lambda-3}.$$

Since $\mathcal{G}_1$ is $(\kappa' - \log 12)$-bit secure, and $\kappa' - \log 12 - \lambda \geq 4$, Theorem 1 implies that $\mathcal{G}_0$ is λ-bit secure as claimed.

In the search setting, we follow the same strategy, but find it is not necessary to use a hybrid approach.

Theorem. Let $\kappa' - \lambda \geq 1$, and let Π be an LWE-HE scheme with parameters $m, n, Q, \chi_{\mathbf{s}}, \chi_{\mathbf{e}}$. Assume that the HintLWE (Search) problem with the same parameters and with flooding variance σ^2 for hint matrix $\frac{1}{q}[\mathbf{A_j}]_q$ is κ'-bit secure. Then if we use flooding noise sampled from a continuous Gaussian with variance

$$\sigma_j^2 \geq \max\left(\frac{2t\kappa' E^2 \log e}{(\kappa' - \lambda)^2}, \sigma^2\right)$$

on input $\mathsf{Rescale}\left(\left(\mathbf{A}_j, \mathbf{b}_j = \mathbf{A}_j\mathbf{s} + \mathbf{e}_j + \mathbf{m}_j\right), q\right)$ with $E \geq \frac{1}{q}\left\|\mathbf{e}_j + [\mathbf{m}_j]_q\right\|$ a static upper bound. Then $\tilde{\Pi}$ is λ-bit secure in the KR-D game where we only decrypt rescaled ciphertexts, and the adversary is allowed at most t queries to the decryption oracle.

Proof. Our proof is very similar to the proof for the previous result: we consider the following two games –

$$\mathcal{G}_0 : \left(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{m} \mod Q, \left\{\frac{\mathbf{m}_j}{q} + \frac{[\mathbf{A}_j]_q}{q}\mathbf{s} - \frac{[\mathbf{b}_j]_q}{q} + \epsilon_j\right\}_{i=0,\ldots,t-1}\right)$$
$$\mathcal{G}_1 : \left(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{m} \mod Q, \left\{\frac{\mathbf{m}_j}{q} + \frac{[\mathbf{A}_i]_q}{q}\mathbf{s} - \frac{[\mathbf{b}_j]_q}{q} + \frac{[\mathbf{e}_j]_q}{q} + \epsilon_j\right\}_{i=0,\ldots,t-1}\right)$$

Where in each case, the adversary is tasked with finding $\mathbf{s}$.

Since $\frac{[\mathbf{b}_j]_q}{q}$ is publicly computable, any adversary to $\mathcal{G}_0$ gives an adversary in the HintLWE search game, so that $\mathcal{G}_0$ is κ'-bit secure by assumption. From here, we can conclude $\mathcal{G}_1$ is λ-bit secure using Lemma 19.

E HintLWE Hardness

In this appendix, we prove Theorem 7, which we restate below.

Theorem (Reduction LWE to HintLWE.). Let $\chi_{\mathbf{s}}$ be a centred spherical discrete Gaussian distribution over $\mathbb{Z}^n$ with variances σ^2, and χ_i be a centred spherical continuous Gaussian with variance σ_i^2 respectively, where

$$\frac{\sigma_i^2}{\rho^2} \geq \left\|\mathbf{X}_i^T\mathbf{X}_i\right\|_2 \tag{8}$$

for all indices i, for some fixed constant $\rho > 0$. Let σ'^2 satisfy

$$\frac{1}{\sigma'^2} \geq 2\left(\frac{1}{\sigma^2} + \frac{t}{\rho^2}\right),$$

and suppose that there is $\varepsilon \in (0,1)$ with $\sigma'^2 \geq 2\eta_\varepsilon(\mathbb{Z}^n)^2$. Then, if $\chi_{\mathbf{s}'} = \mathcal{D}^n_{\sigma'^2}$, there is an efficient reduction from $\mathsf{LWE}_{m,n,q,\chi_{\mathbf{s}'},\chi_{\mathbf{e}}}$ to $\mathsf{HintLWE}^{\xi,\mathbf{X},t}_{m,n,q,\chi_{\mathbf{s}},\chi_{\mathbf{e}}}$ with security loss as follows:

- If the decisional LWE game is κ-bit secure, and $\frac{2\varepsilon}{1-\varepsilon} \leq 2^{(-\kappa'-2)/2}$ for $\kappa' \leq \kappa - 4$, then the decisional HintLWE game is κ'-bit secure.
- If the search LWE game is κ-bit secure, and $\frac{1+\varepsilon}{1-\varepsilon} \leq 2^{\kappa-\kappa'}$ for $\kappa' < \kappa$, then the search HintLWE game is κ'-bit secure.

In this section, we will use the notation Π to denote the precision matrix of a distribution, defined as the inverse of the covariance matrix Σ. To simplify presentation, we will identify these pairs via sub- and super- scripts: Π' denotes the precision matrix for Σ', Π_1 is the precision matrix for Σ_1, etc.

We start by giving a lemma which captures how the distribution of a Gaussian secret vector changes after seeing a noisy linear transformation of the secret, $\mathbf{Xs} + \mathbf{y}$. This theorem generalises Lemma 7 of [35], and is a particular case of Lemma 6 of [24]. This lemma is standard in a Bayesian analysis of linear regression: we recap the proof for completeness.

Lemma 20 (Conditional Distribution $\mathbf{s}|\mathbf{Xs}+\mathbf{y}$). *Let $\mathbf{s} \in \mathbb{R}^n$ be sampled from a discrete (resp. continuous) Gaussian with mean μ and precision matrix Π, and let $\mathbf{y} \in \mathbb{R}^t$ be sampled from a continuous Gaussian with mean $\mu_{\mathbf{y}}$ and precision matrix $\Pi_{\mathbf{y}}$, and let $\mathbf{X} \in \mathbb{R}^{t \times n}$. Then the distribution of $\mathbf{s}$ conditional on $\mathbf{Xs}+\mathbf{y} = \mathbf{z}$ is also a discrete Gaussian, with covariance matrix and mean*

$$\Sigma' = \left(\mathbf{X}^T \Pi_{\mathbf{y}} \mathbf{X} + \Pi\right)^{-1} \tag{9}$$

$$\mu' = (\Pi')^{-1}(\mathbf{X}^T \Pi_{\mathbf{y}}(\mathbf{z} - \mu_{\mathbf{y}}) + \Pi\mu) \tag{10}$$

Proof. We will show the probability mass (resp. density) function has the required form: since the supports are the same, the result follows.

$$\begin{aligned}\Pr[\mathbf{s}|\mathbf{Xs}+\mathbf{y} = \mathbf{z}] &\propto \rho_{\Sigma}(\mathbf{s}-\mu)\rho_{\Sigma_{\mathbf{y}}}(\mathbf{z}-\mathbf{Xs}) \\ &\propto \rho_{\Sigma'}(\mathbf{s}-\mu'),\end{aligned}$$

where the second line is obtained by completing the square.

This enables us to give the general form of Lemma 7 of [35], which gives a tuple with the same distribution as our secret and hints, but crucially uses different secrets in different components.

Corollary 5. *Let the set up be the same as Lemma 20. Then the following tuples are identically distributed:*

$$(\mathbf{s}, \mathbf{Xs}+\mathbf{y})$$
$$(\mathbf{s}', \mathbf{Xs}+\mathbf{y})$$

where $\mathbf{s}'$ is drawn from a discrete (resp. continuous) Gaussian with covariance matrix Σ' and mean μ' as defined in Lemma 20.

Said another way, we can effectively simulate the knowledge gained about $\mathbf{s}$ from the hints without knowing $\mathbf{s}$.

For our reduction, we use the following corollary to this result:

Corollary 6. *The following tuples are identically distributed:*

$(\mathbf{s}, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s}+\mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s}+\mathbf{y}_{t-1})$ *with* $\mathbf{s} \leftarrow \mathcal{D}^n_{\sigma^2}, \mathbf{y}_i \leftarrow \mathcal{N}(0, \sigma_i^2 \mathbf{I}_{k_i})$
$(\mathbf{s}', \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s}+\mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s}+\mathbf{y}_{t-1})$ *with* $\mathbf{s}' \leftarrow \mathcal{D}_{\mu',\Sigma'}, \mathbf{s} \leftarrow \mathcal{D}^n_{\sigma^2}, \mathbf{y}_i \leftarrow \mathcal{N}(0, \sigma_i^2 \mathbf{I}_{k_i})$

where $\Sigma' = \left(\sum \frac{\mathbf{X}_i^T \mathbf{X}}{\sigma_i^2} + \frac{1}{\sigma_1^2}\right)^{-1}$ *and* $\mu' = \Sigma' \left(\sum \frac{\mathbf{X}_i^T(\mathbf{X}_i\mathbf{s}+\mathbf{y}_i)}{\sigma_i^2}\right)$.

Proof. This follows by taking $\mathbf{X} = \begin{bmatrix} \mathbf{X}_0 \\ \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_{t-1} \end{bmatrix}$, and the covariance matrix for $\mathbf{y}$ as

$$\Sigma_{\mathbf{y}} = \begin{bmatrix} \sigma_0^2 \mathbf{I}_{k_0} & & & \\ & \sigma_1^2 \mathbf{I}_{k_1} & & \\ & & \ddots & \\ & & & \sigma_{t-1}^2 \mathbf{I}_{k_{t-1}} \end{bmatrix}$$

We will also use the following lemma.

Lemma 21. *Let $\sigma, \sigma_i, \sigma', \mathbf{X}_i$ be as in the statement of Theorem 7, and define*

$$\Sigma = \left(\sum_{i=0}^{t-1} \frac{\mathbf{X}_i^T \mathbf{X}_i}{\sigma_i^2} + \frac{1}{\sigma^2} \right), \mu = \Sigma \left(\sum_{i=0}^{t-1} \frac{\mathbf{X}_i^T}{\sigma_i^2} (\mathbf{X}_i \mathbf{s} + \mathbf{y}_i) \right).$$

Then

1. *The matrix $\Sigma - \sigma'^2 \mathbf{I}_n$ is positive definite,*
2. *The matrix Σ_3 defined via $\Sigma_3^{-1} = \frac{1}{\sigma'^2} \mathbf{I}_n + \left(\Sigma - \sigma'^2 \mathbf{I}_n \right)^{-1}$ satisfies $\sqrt{\Sigma_3} \geq \eta_\varepsilon(\mathbb{Z}^n)$.*

Proof. Since the matrix Σ is symmetric positive definite, we can show:

$$\begin{aligned} \lambda_1(\Sigma - \sigma'^2 \mathbf{I}_n) &= \lambda_1(\Sigma) - \sigma'^2 \\ &= \lambda_n \left(\Sigma^{-1} \right)^{-1} - \sigma'^2 \\ &\geq \left(\frac{t}{\rho^2} + \frac{1}{\sigma^2} \right)^{-1} - \sigma'^2 \\ &\geq \sigma'^2 \end{aligned}$$

so that the smallest eigenvalue of $\Sigma - \sigma'^2 \mathbf{I}_n$ is greater than zero, and the first claim follows. For the second:

$$\begin{aligned} \left\| \Sigma_3^{-1} \right\|_2 &= \frac{1}{\sigma'^2} + \left\| \left(\Sigma - \sigma'^2 \mathbf{I}_n \right)^{-1} \right\|_2 \\ &= \frac{1}{\sigma'^2} + \lambda_1 (\Sigma - \sigma'^2 \mathbf{I}_n)^{-1} \\ &\leq \frac{2}{\sigma'^2} \\ &\leq \eta_\varepsilon(\mathbb{Z}^n)^{-2} \end{aligned}$$

so that we can conclude by Lemma 7.

We now give the proof of the main theorem.

Proof. Define

$$\Sigma = \left(\sum_{i=0}^{t-1} \frac{\mathbf{X}_i^T \mathbf{X}_i}{\sigma_i^2} + \frac{1}{\sigma^2}\right), \mu = \Sigma \left(\sum_{i=0}^{t-1} \frac{\mathbf{X}_i^T}{\sigma_i^2} (\mathbf{X}_i \mathbf{s} + \mathbf{y}_i)\right)$$

as functions of $\mathbf{X}_i, \sigma_i$ and $\mathbf{X}_i\mathbf{s} + \mathbf{y}_i$. Let $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_Q^{m \times n}), \mathbf{e} \leftarrow \chi_{\mathbf{e}}$. We consider a series of games, in which the adversary is asked to distinguish the first coordinate from uniform:

- $\mathcal{G}_0 : (\mathbf{A}, \mathbf{A}\mathbf{s}' + \mathbf{e})$, for $\mathbf{s}' \leftarrow \mathcal{D}^n_{\sigma'^2}$,
- $\mathcal{G}_1 : (\mathbf{A}, \mathbf{A}(\mathbf{s}' + \mathbf{s}'') + \mathbf{e}, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s} + \mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s} + \mathbf{y}_{t-1})$, for $\mathbf{s}' \leftarrow \mathcal{D}^n_{\sigma'^2}, \mathbf{s}'' \leftarrow \mathcal{D}_{\mu, \Sigma - \sigma'^2 \mathbf{I}_n}, \mathbf{s} \leftarrow \mathcal{D}^n_{\sigma^2}$
- $\mathcal{G}_2 : (\mathbf{A}, \mathbf{A}\mathbf{s}''' + \mathbf{e}, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s} + \mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s} + \mathbf{y}_{t-1})$, for $\mathbf{s}''' \leftarrow \mathcal{D}_{\mu, \Sigma}, \mathbf{s} \leftarrow \mathcal{D}^n_{\sigma^2}$
- $\mathcal{G}_3 : (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s} + \mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s} + \mathbf{y}_{t-1})$, for $\mathbf{s} \leftarrow \mathcal{D}^n_{\sigma^2}$.

We recognise $\mathcal{G}_0$ as the LWE decision game with secret $\mathbf{s} \leftarrow \mathcal{D}_{\sigma'^2 \mathbf{I}_n}$, which is κ-bit secure by hypothesis, and $\mathcal{G}_3$ as the target HintLWE game.

The game $\mathcal{G}_1$ is well defined, since $\Sigma - \sigma'^2 \mathbf{I}_n > 0$ by Lemma 21. It is also perfectly simulatable by any adversary to $\mathcal{G}_0$, and any adversary that succeeds against $\mathcal{G}_1$ gives an adversary that succeeds against $\mathcal{G}_0$ We can therefore say that $\mathcal{G}_1$ is at least κ-bit secure.

By the data processing inequality, the R_∞ distance between $\mathcal{G}_2$ and $\mathcal{G}_1$ is at most $D(\mathbf{s}'''\|\mathbf{s}' + \mathbf{s}'')$. This is exactly the divergence in the convolution lemma, so that the condition on Σ_3 in Lemma 21 and the data processing inequality implies that $R_\infty(\mathcal{G}_2\|\mathcal{G}_1) \leq \frac{1+\varepsilon}{1-\varepsilon}$. Therefore, using Theorem 1, we can conclude that if $\frac{1+\varepsilon}{1-\varepsilon} - 1 = \frac{2\varepsilon}{1-\varepsilon} \leq 2^{(-\kappa'-2)/2}$ for $\kappa' \leq \kappa - 4$, then $\mathcal{G}_2$ is κ'-bit secure.

From here, we observe that the tuples $(\mathbf{s}''', \mathbf{X}_0, ..., \mathbf{X}_{l-1}, \mathbf{X}_0\mathbf{s} + \mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s} + \mathbf{y}_{t-1})$ and $(\mathbf{s}, \mathbf{X}_0, ..., \mathbf{X}_{t-1}, \mathbf{X}_0\mathbf{s} + \mathbf{y}_0, ..., \mathbf{X}_{t-1}\mathbf{s} + \mathbf{y}_{t-1})$ are identically distributed by Corollary 6, so that the games $\mathcal{G}_2$ and $\mathcal{G}_3$ are identical, so that the claim on the security of the HintLWE decision game follows.

For the search variant, we follow the same series of games, observing that recovering $\mathbf{s}$ enables recovering $\mathbf{s}'$ by calculating $\mathbf{e}$. Since this is now a search game, we use that $R_\infty(\mathcal{G}_2\|\mathcal{G}_1) \leq R_\infty(\mathbf{s}'''\|\mathbf{s}' + \mathbf{s}'') \leq \frac{1+\varepsilon}{1-\varepsilon}$, and the result follows by Lemma 10.

F OpenFHE Parameters

In the OpenFHE library [3], the user has the option to target the (c, s) security introduced in [42]. To this end, they post process decryptions by adding a sample from a discrete Gaussian with variance[8]

$$\sigma^2 = 12t2^s E^2,$$

[8] See here.

where t is an upper bound on the number of decryptions the adversary can request, s is the target statistical security, and E is an upper bound on the infinity canonical embedding of the noise present in the ciphertext $\|e\|_\infty^{\text{can}}$. This noise is added in the canonical embedding space. The notion of (c, s) security is strictly weaker than bit security when $s < c$: indeed, from Theorem 4 of [42] we can conclude that this amount of noise flooding gives a bit security of at most $2s$ in the IND-CPA-D game.

Using Theorem 3, and observing that in the OpenFHE CKKS message space $\mathbb{R}^{n/2}$, $\|e\|_2 \leq \sqrt{\frac{n}{2}} \|e\|_\infty$, we can verify this amount of noise flooding gives a bit security of $s + \log 6t - \log n$. Using Theorem 4, letting $\kappa = 128$, $s \geq \log n + 4$ implies that this noise flooding gives a security loss of at most 1 bit in the key recovery game.