"DP as a Discount": A Hybrid Differential Privacy Approach with Provable Privacy and Security for Approximate Homomorphic Encryption

Angus He [1]

Abstract

Homomorphic encryption enables computations on ciphertexts directly without the need for decryption. Approximate homomorphic encryption is a type of homomorphic encryption scheme that permits data-dependent noise introduced during the computation. Recent work has shown that approximate homomorphic encryption algorithms can offer certain differential privacy guaranties. However, approximate homomorphic encryption schemes are vulnerable to chosen plaintext attacks with decryption oracles (IND-CPA$^{\mathrm{D}}$), and recent work mitigates this by adding Gaussian noise at the decryption stage. In this project, we integrate the result of these two prior works and introduce a hybrid framework called "DP as a Discount". We show that, under certain explicit assumptions, this hybrid mechanism achieves strong cryptographic security (IND-CPA$^{\mathrm{D}}$) and differential privacy ((ε, δ)-DP) at the same time.

1. Introduction

1.1. Homomorphic Encryption

Homomorphic encryption (Gentry, 2009) is a class of cryptographic schemes that allows computations to be performed directly on ciphertexts, eliminating the need of decrypting the ciphertexts, performing computations on the decrypted plaintexts, and re-encrypting them again. This property allows a third party to process encrypted data while preserving confidentiality as the plaintext is never exposed to the third party. Homomorphic encryption ensures that decrypting a computed ciphertext yields the same result as performing the corresponding operations on the original plaintext. Homomorphic encryption is suitable for situations where sensitive data should remain confidential during outsourced computation, such as secure voting systems, processing sensitive healthcare or financial data, privacy-preserving machine learning, and so on.

Homomorphic encryption schemes can be classified as partially homomorphic encryption or fully homomorphic encryption (Acar et al., 2018), depending on whether they support limited or arbitrary homomorphic computations. Alternatively, they can be classified as exact or approximate homomorphic encryption (Cheon et al., 2017), depending on whether the computations on ciphertexts yield precise or approximate results.

Definition 1.1 (Homomorphic Encryption)**.** Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be a homomorphic encryption scheme where $\mathsf{KeyGen} \to (pk, sk)$ is a function that generates a public key pk and a secret key sk, $\mathsf{Enc}(pk, m) \to c$ encrypts a plaintext m into ciphertext c, $\mathsf{Dec}(sk, c) \to m$ decrypts a ciphertext c to recover plaintext m, and $\mathsf{Eval}(pk, f, c_1, \ldots, c_n) \to c_f$ evaluates a function f over ciphertexts. Then, Π is said to be a homomorphic encryption scheme if for any function f in a supported class of functions and plaintexts $m_1, \ldots, m_n$, the following correctness property always holds:

$$\mathsf{Dec}\big(sk, \mathsf{Eval}(pk, f, \mathsf{Enc}(pk, m_1), \ldots, \mathsf{Enc}(pk, m_n))\big) \simeq f(m_1, \ldots, m_n) \quad (1)$$

Homomorphic encryption is increasingly used in systems that perform computations on sensitive data (Cheon et al., 2017; Acar et al., 2018). Recently, homomorphic encryption has gained increasing popularity in cloud computing, allowing users to outsource computation on financial, behavioral, or proprietary datasets while maintaining confidentiality. Despite its increased practicality, homomorphic encryption remains challenging to deploy due to significant performance issue (Alaya et al., 2020). Thus, a subclass of homomorphic encryption called approximate homomorphic encryption is proposed to solve this issue by allowing the final decrypted output to be approximate.

[1]Computer Science Department, University of Wisconsin - Madison, Madison, WI, United States. Correspondence to: Angus He <cahe@wisc.edu>.

Proceedings of the 43^{rd} International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026.

1.2. Approximate Homomorphic Encryption

Approximate homomorphic encryption is a type of homomorphic encryption that allows the decrypted results to contain small, bounded errors. This relaxation allows for more computationally efficient implementations compared to fully homomorphic encryption. Among approximate homomorphic encryption schemes, the CKKS (Cheon–Kim–Kim–Song) scheme (Cheon et al., 2017) is the most commonly used approach, as it supports efficient approximate arithmetic on real and complex numbers.

Approximate homomorphic encryption is useful for applications where precision is less important than computational efficiency, such as privacy-preserving machine learning and statistical analysis. In approximate homomorphic encryption schemes, errors build up with each computation performed, which can affect the final result. Thus, parameters need to be chosen carefully to keep the noise under control, ensuring that the decrypted output is still accurate.

Mathematically, approximate homomorphic encryption is defined similarly to the original definition of homomorphic encryption, but the correctness is relaxed to allow bounded error.

Definition 1.2 (Approximate Homomorphic Encryption)**.** Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be a homomorphic encryption scheme. Then, Π is said to be an *approximate homomorphic encryption* scheme if for any function f in a supported class of functions and plaintexts $m_1, \ldots, m_n$, the following approximate correctness property always holds:

$$\mathsf{Dec}\big(sk, \mathsf{Eval}(pk, f, \mathsf{Enc}(pk, m_1), \ldots, \mathsf{Enc}(pk, m_n))\big) \approx f(m_1, \ldots, m_n) \quad (2)$$

Here, $\approx$ denotes that the decrypted result is within a predefined error bound. To formalize the definition of $\approx$, we begin by defining *ciphertext error*, which quantifies how far a homomorphic computation deviates from the expected correct result.

Definition 1.3 (Ciphertext Error)**.** Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be a homomorphic encryption scheme with message space $\widetilde{M}$, where $\widetilde{M}$ is the ambient space to accommodate the output of an approximate decryption function. For any ciphertext ct, secret key sk, and message m, the ciphertext error $\mathsf{Error}(ct, m, sk)$ is defined as

$$\mathsf{Error}(ct, m, sk) := \|\mathsf{Dec}_{sk}(ct) - m\| \quad (3)$$

With the above definition in place, we can now formally define the approximate correctness:

Definition 1.4 (Approximate Correctness)**.** Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be an approximate homomorphic encryption scheme with message space $M \subseteq \widetilde{M}$. Let $\mathcal{L}$ be a space of circuits, with $\mathcal{L}_k \subseteq \mathcal{L}$ representing the subset of parity k circuits. Let $\mathsf{Estimate} := \bigsqcup_{k\in\mathbb{N}} \mathcal{L}_k \times \mathbb{R}^k_{\geq 0} \to \mathbb{R}_{\geq 0}$ be an efficiently computable function. Then, the approximate correctness of Π is defined as:

$\forall k \in \mathbb{N}$, $\forall \mathcal{C} \in \mathcal{L}_k$, and $\forall (pk, sk) \leftarrow \mathsf{KeyGen}(1^\kappa)$ where pk is the public key and sk is the secret key. If $\forall i \in \{1, 2, \ldots, k\}$, $\mathsf{Error}(ct_i, m_i, sk) \leq t_i$. Then,

$$\mathsf{Error}(\mathsf{Eval}_{pk}(\mathcal{C}, ct_1, \ldots, ct_k), \mathcal{C}(m_1, \ldots, m_k), sk) \leq \mathsf{Estimate}(\mathcal{C}, t_1, \ldots, t_k) \quad (4)$$

As we can see, Estimate is defined as a function that takes an input computation circuit $\mathcal{C}$ and an error bound t_i for each of the k input wires to the circuit $\mathcal{C}$. With this definition, Estimate is in fact static because it only depends on the computation circuit and error bound. This property is crucial for our analysis becuase the error bound is publicly known, so any additional noise added during decryption can be calibrated as a function of the error bound without introducing additional dependencies.

1.3. Chosen Plaintext Attacks

Traditionally, the security of homomorphic encryption schemes is evaluated using the concrete-security definition of indistinguishability under chosen plaintext attacks (IND–CPA) (Katz & Lindell, 2007). In an indistinguishability game, the advantage of an adversary A, denoted as adv_A, is a metric that determines how often the adversary provides an answer and how accurate the answer is.

In the concrete-security setting, the advantage of an adversary A is expressed in terms of two quantities: its output probability and its conditional distinguishing advantage. In this setting, the adversary is allowed to abort and output $\perp$. The output probability α_A captures how often the adversary produces a valid guess, while the conditional distinguishing advantage δ_A captures how well the adversary distinguishes the challenge bit, conditioned on not aborting the game.

Definition 1.5 (Output Probability)**.** Let A be an adversary in the IND–CPA experiment. The output probability of A is defined as

$$\alpha_A := \Pr[A \neq \perp] \quad (5)$$

Definition 1.6 (Conditional Distinguishing Advantage)**.** Let A be an adversary in the IND–CPA experiment. The *conditional distinguishing advantage* of A is defined as

$$\delta_A := 2 \cdot \Pr[b' = b \mid A \neq \perp] - 1 \quad (6)$$

where b is the challenge bit and b' is the adversary's output.

Since random guessing succeeds with probability $\frac{1}{2}$, we measure the advantage relative to this baseline. Thus, one

way of defining adv_A can be $\alpha_A \cdot \left(\beta_A - \frac{1}{2}\right)$. Yet, for convenience, we scale this quantity by a factor of 2 so that the advantage ranges in $[-1, 1]$.

Definition 1.7 (Advantage Formula)**.** Given the output probability α_A and the conditional distinguishing advantage δ_A, the *advantage formula* adv_A is defined as

$$\mathsf{adv}_A := \alpha_A \cdot (\delta_A)^2 \tag{7}$$

Note that the above equation quadratically scales δ_A. Typically, the advantage is defined linearly as $\mathsf{adv}_A = \alpha_A \cdot \delta_A$. The rationale for using the quadratic term is motivated by Micciancio and Walter (Micciancio & Walter, 2018), who argue that this better captures the notion of advantage in decision problems.

We now use this notion of advantage to formalize the concrete-security definition of indistinguishability under chosen plaintext attacks:

Definition 1.8 (IND–CPA Security)**.** Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be an approximate homomorphic encryption scheme. The IND–CPA game is an indistinguishability game where a challenger chooses a bit $b \leftarrow Unif(\{0, 1\})$ and an adversary A is given access to an encryption oracle $E_b^{pk}(m_0, m_1)$ which returns $ct \leftarrow \mathsf{Enc}_{pk}(m_b)$. The scheme Π is κ-bit IND-CPA-secure if for any adversary A, the ratio of the running time T_A to its advantage adv_A satisfies

$$\mathsf{adv}_A \leq \frac{T_A}{2^\kappa} \tag{8}$$

Essentially, this definition guarantees that dversaries get only tiny advantage unless they run in exponential time.

Rearranging the inequality of 1.8 gives the equivalent bound $\kappa \leq \log_2\left(\frac{T_A}{\mathsf{adv}_A}\right)$. Hence, we then define another notation called the *bit security* of the indistinguishable game to quantify the resources required by an adversary to successfully win the game.

Definition 1.9 (Bit Security of Indistinguishable Game)**.** For an adversary A participating in the indistinguishability game, let T_A denotes the runtime and adv_A denotes the advantage of the adversary. Then, the bit security κ_s is defined as

$$\kappa_s := \min_A \log_2\left(\frac{T_A}{\mathsf{adv}_A}\right) \tag{9}$$

1.4. Norm KL Differential Privacy

In this project, we use the notion of Rényi differential privacy (Mironov, 2017). We specifically focus on the case where the Rényi parameter $\alpha = 1$, as this yields the tightest bounds for our security framework. Following this paradigm, we define the ρ-KL differential privacy (ρ-KLDP) (Wang et al., 2016) to be a mechanism based on its ability to bound the information leakage relative to the distance between the inputs.

Definition 1.10 (Norm KL Differential Privacy)**.** For $t \geq 0$, let M_t be a family of mechanisms. Let $\rho \in \mathbb{R}$ be a privacy bound. Then, we say that the family of mechanism M_t is *ρ-KL differentially private (ρ-KLDP)* if, for all $x, x' \in \mathcal{X}$ with $\|x - x'\| \leq t$:

$$\mathsf{KL}(M_t(x) \| M_t(x')) \leq \rho \tag{10}$$

where KL denotes the KL-divergence.

Lemma 1.11 (Gaussian Mechanism is ρ-KLDP)**.** *Let $\rho > 0$, and $n \in \mathbb{N}$. Define the Gaussian noise transformation $\mathsf{Gauss} : \mathbb{Z}^n \to \mathbb{Z}^n$ to be the mechanism that, on input $x \in \mathbb{Z}^n$, outputs a sample from $\mathcal{N}_{\mathbb{Z}^n}\left(x, \frac{t^2}{2\rho} I_n\right)$. Then, for any $\rho > 0$ and $n \in \mathbb{N}$, Gauss is ρ-KLDP.*

Proof. The proof can be found in (Li et al., 2022), so we omit it. □

2. Related Work

2.1. Adding External Noise to Approximate Homomorphic Encryption

Recent research by Li and Micciancio (Li & Micciancio, 2021) showed that the traditional definition of IND–CPA is not enough to capture the security of approximate homomorphic encryption. Under IND–CPA, an adversary has access to only an encryption oracle. However, the definition of IND–CPA does not accurately describe other important aspects of homomorphic encryption systems. First, in typical homomorphic encryption schemes, ciphertexts are evaluated homomorphically multiple times before the final result is decrypted. A passive adversary can therefore observe the approximate decrypted outputs of these "honest" computations. Secondly, an adversary may know or choose to perform homomorphic computation on some passively collected "honest" ciphertexts.

In approximate homomorphic encryption schemes, decryption returns a value of the form $m + e$, where the error term e may depend on the ciphertext, secret key, or the underlying data (Cheon et al., 2017). Observing these decrypted outputs may introduce information leakage that is not captured by IND–CPA. In fact, this leakage can be exploited in practice. For example, in the CKKS scheme, Li and Micciancio (Li & Micciancio, 2021) demonstrated an key recovery attack by exploiting the aforementioned leakage on the CKKS scheme. Thus, while it is true that approximate homomorphic encryption schemes are IND–CPA secure, IND–CPA alone is not enough to accurately describe approximate homomorphic encryption schemes. The need

for stronger notions that explicitly models this leakage is required.

Because of this, Li and Micciancio (Li & Micciancio, 2021) introduced a stronger security formulation known as IND–CPA$^{\mathsf{D}}$ to address this problem. The main difference is that IND–CPA$^{\mathsf{D}}$ gives the adversary access to a "restricted" decryption oracle where the adversary can observe decrypted results of some "honestly" generated ciphertexts.

Definition 2.1 (IND–CPA$^{\mathsf{D}}$)**.** Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ be a homomorphic encryption scheme. The IND–CPA$^{\mathsf{D}}$ game is parameterized by three stateful oracles

1. The encryption oracle $E_b^{pk}(m_0, m_1)$: returns $ct \leftarrow \mathsf{Enc}_{pk}(m_b)$ and records the triplet (m_0, m_1, ct) in the global state σ.
2. The Evaluation Oracle $H_b^{pk}(f, J)$: for indices $J = (j_1, \ldots, j_k)$, it computes $ct \leftarrow \mathsf{Eval}_{pk}(f, \sigma[j_1].ct, \ldots, \sigma[j_k].ct)$ and records the resulting messages and ciphertext in σ.
3. The decryption oracle $D_b^{sk}(i)$: if the recorded messages for index i satisfy $S[i].m_0 = S[i].m_1$, it returns the approximate decryption $\mathsf{Dec}_{sk}(\sigma[i].ct)$; otherwise, it returns $\perp$.

Then, the scheme Π is κ-bit IND–CPA$^{\mathsf{D}}$ secure if for any adversary A, the ratio of the running time T_A to its advantage adv_A satisfies

$$\mathsf{adv}_A \leq \frac{T_A}{2^\kappa} \tag{11}$$

Meanwhile, we define q-IND–CPA$^{\mathsf{D}}$ as IND–CPA$^{\mathsf{D}}$ scheme but restricted to adversaries that can only make at most q queries to the decryption oracle. We will see later that this definition can serve as a "smooth transition" from IND–CPA to IND–CPA$^{\mathsf{D}}$ in our proof.

Essentially, 2.1 models a realistic scenario where an adversary can perform homomorphic computation on ciphertexts and request approximate decryption of the evaluated ciphertexts, provided that the resulting plaintexts are independent of the challenge bit b. For homomorphic encryption schemes with exact decryption, Li and Micciancio (Li & Micciancio, 2021) proved that IND–CPA$^{\mathsf{D}}$ is equivalent to the traditional IND–CPA notation. However, in the context of approximate homomorphic encryption, the definition of IND–CPA$^{\mathsf{D}}$ is necessary to make sure that the ciphertext errors do not leak additional information. Consequently, additional post-processing mechanisms are required in order to make approximate homomorphic encryption secure against IND–CPA$^{\mathsf{D}}$.

In fact, an interesting result proposed by Li and Micciancio (Li et al., 2022) stated that adding simple Gaussian noises allows a transform from IND–CPA to IND–CPA$^{\mathsf{D}}$ with only 8-bits of security loss.

Definition 2.2 (Gaussian Noise Transformation)**.** Let $\widetilde{\Pi} = (\Pi, \mathsf{Estimate})$ be an approximate homomorphic encryption scheme where $\mathsf{Estimate}$ is a publicly computable function providing an upper bound on ciphertext error. The transformed scheme $\mathsf{Gauss}[\widetilde{\Pi}]$ is defined with a modified decryption function as such:

$$\mathsf{Dec}'_k(ct) := \mathsf{Dec}_k(ct) + Z_t \tag{12}$$

where $Z_t \sim \mathcal{N}(0, \sigma_t^2)$, and the variance σ_t^2 is chosen as a function of the error bound t so that the Gaussian noise transformation satisfies the ρ-KLDP condition.

Before we get into the proof, we mention the following lemma stated in (Micciancio & Walter, 2018), which bounds the loss of bit security when replacing one distribution with another.

Lemma 2.3 (Approximate Sampler)**.** *Given two indinstinguishability games G_P and G_Q with access to a probability ensemble P_θ and Q_θ respectively. If G_P is κ-bit secure, and*

$$\sup_\theta \mathsf{KL}(P_\theta || Q_\theta) \leq 2^{-\kappa+1}$$

Then G_Q is $(\kappa - 8)$-bit secure. Here, KL denotes the KL-divergence.

Theorem 2.4 (Gaussian Noise Transformation on Approximate Homomorphic Encryption)**.** *Let Π be a homomorphic encryption scheme with message space $\mathcal{M} \subseteq \widetilde{\mathcal{M}}$. Let Gauss be the Gaussian noise transformation with $\rho = \frac{2^{-\kappa-7}}{q}$ 1.10. If Π is $(\kappa + 8)$-bit secure in the IND-CPA game, then the transformed scheme $\mathsf{Gauss}[\Pi]$ is κ-bit secure in the $q - \mathsf{IND} - \mathsf{CPA}^{\mathsf{D}}$ game.*

Proof. Let $\widetilde{\Pi} = (\Pi, \mathsf{Estimate})$ be an approximate homomorphic encryption scheme where $\mathsf{Estimate}$ is a publicly computable function providing an upper bound on ciphertext error. Let $\mathsf{Gauss}[\widetilde{\Pi}]$ be its variant by adding Gaussian noise to the decrypted output, where Gauss is a ρ-KLDP mechanism. We evaluate the security of $\mathsf{Gauss}[\widetilde{\Pi}]$ against an adversary A making q decryption queries through a sequence of games

1. **Game** G_0: The real-world q-IND-CPAD experiment. The decryption oracle $D_{sk}(i)$ returns $y_{0i} = \mathsf{Gauss}(\mathsf{Dec}_{sk}(ct_i))$ where ct_i is an honestly generated ciphertext.
2. **Game** G_1: An idealized q-IND-CPAD experiment. The decryption oracle is replaced by an idealized version $D^*(i)$ that returns $y_{1i} = \mathsf{Gauss}(m_i)$, where m_i is the exact message. This oracle ignores the specific ciphertext error.

3. **Game** G_2: The standard IND-CPA game for the underlying scheme Π.

Let $X_\theta = \{y_{01}, \ldots, y_{0q}\}$ and $Y_\theta = \{y_{11}, \ldots, y_{1q}\}$ be the ensembles of decryption responses of G_0 and G_1. Applying the result from (Li et al., 2022) regarding adaptive adversaries and the sub-additivity of KL divergence, we have:

$$\mathsf{adv}_A \leq \frac{q}{2} \max_{\theta \in \Theta} \mathsf{KL}(X_\theta \| Y_\theta) \leq \frac{\rho q}{2}$$

Then by 2.3, provided that $\frac{\rho q}{2} \leq 2^{-\kappa-8}$, we conclude that if G_1 is $(\kappa+8)$-bit q-$\mathsf{IND\text{–}CPA}^{\mathsf{D}}$ secure, then G_0 is κ-bit q-$\mathsf{IND\text{–}CPA}^{\mathsf{D}}$ secure.

Finally, note that the decryption oracle of G_1 is perfectly simulatable. Thus, any adversary against $\tilde{\Pi}$ in the IND-CPA game yields an adversary of the same advantage and running time against $\mathsf{Gauss}[\tilde{\Pi}]$ in the q-$\mathsf{IND\text{–}CPA}^{\mathsf{D}}$ game. As a result, if $\tilde{\Pi}$ has $(\kappa+8)$ bits of $\mathsf{IND\text{–}CPA}$ security in G_2, then $\mathsf{Gauss}[\tilde{\Pi}]$ has $(\kappa+8)$ bits of security in game G_1, and therefore $\mathsf{Gauss}[\tilde{\Pi}]$ has κ bits of q-$\mathsf{IND\text{–}CPA}^{\mathsf{D}}$ security in G_0. □

2.2. Intrinsic Noise

In most modern homomorphic encryption schemes, error terms are intentionally introduced during encryption to ensure security. Traditionally, fully homomorphic encryption algorithms are designed to remove this noise after decrypting the plaintext. However, most approximate homomorphic encryption algorithms operate in a way which the noise is not eliminated but persisted. As an example, in the CKKS algorithm, the error is persisted in the least significant bits of the decrypted result. This aises the question of whether the internal noise generated by the error can be leveraged to guarantee privacy. Recently, Ogilvie (Ogilvie, 2024) explored this problem and concluded that this error itself can indeed provide DP "for free" under certain circumstances.

We now evaluate the differential privacy guarantee of an approximate homomorphic encryption scheme from the internal noise. Because all intermediate values remain encrypted during the homomorphic computation process, it suffices to analyze the privacy of the final decrypted output. Following the analysis of (Costache et al., 2023), the output of a homomorphic computation can be modeled as a Gaussian distribution

$$\beta + \mathcal{N}(0, \sigma^2)$$

A natural approach is therefore to treat the internal noise as an instance of the Gaussian noise transformatio, and to require that σ is sufficiently large to mask the distance $\|D - D'\|$ between two adjacent databases $D \sim D'$. Recall the following well-known guarantee for the Gaussian noise transformatio

Lemma 2.5. *for $\varepsilon \in (0,1)$ and $c > \sqrt{2\ln(\frac{1.25}{\delta})}$, the mechanism is (ε, δ)-differentially private provided that*

$$\sigma \geq \frac{c\,\Delta_f}{\varepsilon} \tag{13}$$

where Δ_f denotes the ℓ_2-sensitivity of the function.

Applying 2.5 to the homomorphic setting, suppose that at iteration k, we have a bound $\|\beta - \beta'\| \leq \delta_k$, and that the output noise has variance at least σ^2. One might then expect that the condition

$$\sigma > \frac{\sqrt{2\ln\left(\frac{1.25}{\delta}\right)}\Delta_f}{\varepsilon}$$

is sufficient to guarantee (ε, δ)-differential privacy at iteration k (Dwork et al., 2006). However, this argument overlooks an important feature of noise growth in approximate homomorphic encryption: the variance of this internal noise may be dependent on the underlying plaintext. As an example, the homomorphic multiplication of two ciphertexts $c_1 = m_1 + e_1$ and $c_2 = m_2 + e_2$ yields the following result:

$$\begin{aligned} c_{\text{eval}} &= c_1 \cdot c_2 \\ &= (m_1 + e_1)(m_2 + e_2) \\ &= \underbrace{m_1 m_2}_{m_{\text{eval}}} + \underbrace{m_1 e_2 + m_2 e_1 + e_1 e_2}_{e_{\text{eval}}} \end{aligned}$$

As we can see, the resulting error term becomes dependent on the underlying plaintext after homomorphic multiplication. In addition, one can easily verify that for the ciphertext c_1 and c_2, the result ciphertext has variance

$$N\sigma_1^2\sigma_2^2 + \sigma_1^2\|m_2\|^2 + \sigma_2^2\|m_1\|^2 \tag{14}$$

Where N is the dimension of the space on which the noise distributions are defined. In other words, the variance also depends on the input data m_1 and m_2. This means that must account for the effect of this message-dependent variance when analyzing differential privacy. As a result, Ogilvie (Ogilvie, 2024) proposed that the output of the homomorphic computation should be more accurately modeled as

$$\beta_D + \mathcal{N}(0, \sigma_D^2) \tag{15}$$

where both the mean β_D and the variance σ_D^2 depend on the database (message) D. This is different from the standard Gaussian noise transformation in which the noise distribution is independent of the data.

Under this more accurate model, we first look at the privacy budget requirement to guarantee differential privacy from the internal error noise:

Theorem 2.6 (Approximate Homomorphic Encryption Gaussian DP Theorem)**.** *Suppose we use a Gaussian noise transformation with mean β_D and variance σ_D^2 dependent on the underlying database D. Then the resulting mechanism is (ϵ, δ)-DP if*

$$\epsilon > T^2 K\sqrt{L} + \frac{1}{2}T^2K^2 + \frac{1}{2}(T^2-1)L + \ln T \quad (16)$$

where $L = 2\ln\left(\frac{1}{\delta}\right)$, $\sup_{D\sim D'} \frac{\sigma_D}{\sigma_{D'}} \leq T$, *and* $\sup_{D\sim D'} \frac{|\beta_D - \beta_{D'}|}{\sigma_D} \leq K$

Intuitively, the parameter T quantifies the stability of the noise scale across adjacent datasets. Specifically, it upper bounds the multiplicative change in the standard deviation. Thus, T measures how much the variance of the Gaussian noise can vary between adjacent datasets. As T approaches 1, the noise variance between two adjacent datasets becomes increasingly similar, making the outputs harder to distinguish. Intuitively, this implies stronger differential privacy guarantees.

Meanwhile, the parameter K measures how the sensitivity changes compared to the magnitude of the noise. As K approaches 0, the bias term β_D approaches constant across all adjacent datasets. This means that differences between the output distributions come purely from the random Gaussian noise rather than message-dependent bias, implying stronger privacy.

Proof. Fix adjacent databases $D \sim D'$. Let the corresponding output distributions be

$$A \sim \mathcal{N}(\beta, \sigma^2)$$
$$A' \sim \mathcal{N}(\beta', \sigma'^2)$$

We analyze the ratio of their probability density functions at a point $\alpha \in \mathbb{R}$:

$$\frac{f_A(\alpha)}{f_{A'}(\alpha)} = \frac{\sigma'}{\sigma}\exp\left(\frac{(\alpha-\beta')^2}{2\sigma'^2} - \frac{(\alpha-\beta)^2}{2\sigma^2}\right)$$

Define

$$t := \frac{\sigma}{\sigma'}, \qquad k := \frac{\beta'-\beta}{\sigma}$$

and rewrite the ratio as

$$\frac{f_A(\alpha)}{f_{A'}(\alpha)} = \frac{1}{t}\exp\left(\frac{t^2}{2}\left(\frac{\alpha-\beta}{\sigma} - k\right)^2 - \frac{1}{2}\left(\frac{\alpha-\beta}{\sigma}\right)^2\right)$$

Let

$$Z := \frac{\alpha-\beta}{\sigma}, \quad Z \sim \mathcal{N}(0,1)$$

Then the log-likelihood ratio becomes

$$\ln\frac{f_A(\alpha)}{f_{A'}(\alpha)} = -\ln t + \frac{t^2}{2}(Z-k)^2 - \frac{1}{2}Z^2$$

We seek to bound this quantity within $[-\varepsilon, \varepsilon]$ with probability at least $1-\delta$ over Z. Equivalently, we require

$$-\varepsilon \leq -\ln t + \frac{t^2}{2}(Z-k)^2 - \frac{1}{2}Z^2 \leq \varepsilon$$

Which is equivalent to

$$\varepsilon \geq -\frac{t^2}{2}(Z-k)^2 + \frac{1}{2}Z^2 + \ln t$$
$$\varepsilon \geq \frac{t^2}{2}(Z-k)^2 - \frac{1}{2}Z^2 - \ln t$$

In any case,

$$\varepsilon \geq \frac{t^2}{2}(Z-k)^2 + \frac{1}{2}Z^2 + \ln t)$$

would suffice.

Now, we need to find the suitable Z. Using standard Gaussian tail bounds, for any $u > 0$,

$$\Pr[|Z| > u] \leq 2\exp\left(-\frac{u^2}{2}\right)$$

Setting $u = \sqrt{L}$ we obtain

$$\Pr[|Z| > \sqrt{L}] \leq \delta$$

Since $\frac{t^2}{2}(Z-k)^2 + \frac{1}{2}Z^2 + \ln t$ is a quadratic function that opens upward, and that $|Z| \leq \sqrt{L}$, the maximum happens at either $Z = \sqrt{L}$ or $Z = -\sqrt{L}$. In any case, we have

$$\varepsilon \geq \pm t^2 k\sqrt{L} + \frac{1}{2}t^2k^2 + \frac{1}{2}(t^2-1)L + \ln t$$

which means that

$$\varepsilon \geq t^2 k\sqrt{L} + \frac{1}{2}t^2k^2 + \frac{1}{2}(t^2-1)L + \ln t$$

would suffice.

Now, our proof shows that within a "good" region R where

$$R := \left\{\alpha \in \mathbb{R} : e^{-\varepsilon} \leq \frac{f_A(\alpha)}{f_{A'}(\alpha)} \leq e^{\varepsilon}\right\}$$

we have

$$\Pr[A \in R] \geq 1-\delta$$

We now show that this implies (ε, δ)-differential privacy. Let $S \subseteq \mathbb{R}$ be any measurable set. Then

$$\Pr[A \in S] = \Pr[A \in S \cap R] + \Pr[A \in S \cap R^c]$$

Since $\Pr[A \in R^c] \leq \delta$, we have

$$\Pr[A \in S] \leq \Pr[A \in S \cap R] + \delta$$

Moreover, for every $\alpha \in R$,

$$f_A(\alpha) \le e^{\varepsilon} f_{A'}(\alpha)$$

Therefore,

$$\begin{aligned}\Pr[A \in S \cap R] &= \int_{S\cap R} f_A(\alpha)\, d\alpha \\ &\le e^{\varepsilon} \int_{S\cap R} f_{A'}(\alpha)\, d\alpha \\ &\le e^{\varepsilon} \Pr[A' \in S]\end{aligned}$$

Combining the two bounds gives

$$\Pr[A \in S] \le e^{\varepsilon} \Pr[A' \in S] + \delta$$

Finally, applying the bounds $t \le T$ and $k \le K$, we obtain the stated condition for (ε, δ)-differential privacy. □

3. Result

In this section, we combine the external and internal noise mentioned in the two prior works (Li et al., 2022) (Ogilvie, 2024) and study how to simultaneously obtain cryptographic security and differential privacy in approximate homomorphic encryption. As we have seen, existing approaches treat these goals separately. On one hand, q-IND-CPAD security can be achieved by adding data-independent noise at decryption. On the other hand, recent analyses show that the internal noise arising from the homomorphic computation can guarantee DP, although this noise is message-dependent. This section proves that: under certain assumptions on approximate decryption and the distribution of the released output, it is possible to simultaneously achieves q-IND-CPAD security and (ε, δ)-differential privacy. Furthermore, we prove that due to the presence of internal noise, the amount of external noise required to achieve the same level of security and privacy is reduced.

3.1. Model and Assumptions

Assumption 3.1 (Approximate Homomorphic Encryption Model Assumption)**.** Let $\widetilde{\Pi} = (\Pi, \mathsf{Estimate})$ be an approximate homomorphic encryption scheme where $\mathsf{Estimate}$ is a publicly computable function providing an upper bound on ciphertext error. Let ct be an evaluated ciphertext with exact plaintext value $m \in \mathbb{R}$ and associated public error bound $t = \mathsf{Estimate}(ct)$. Then, there exists a decomposition of approximate decryption:

$$\mathsf{Dec}_{sk}(ct) = m + \beta + Z_{\text{inter}}$$

where: $|\beta| \le t$, $Z_{\text{inter}} \sim \mathcal{N}(0, \sigma^2_{\text{inter}})$. Note that by this definition, Z_{inter} depends only on the upper bound of ciphertext error t.

Assumption 3.2 (Gaussian Output Model)**.** Let $f : \mathcal{D} \to \mathbb{R}$ be a function computed homomorphically. Denote the released value before adding external Gaussian noise as Y_{raw}. Then, for every database D, the released value before adding external Gaussian noise satisfies:

$$Y_{\text{raw}} = f(D) + Z_{\text{inter}}, \quad Z_{\text{inter}} \sim \mathcal{N}(0, \sigma^2_{\text{inter}}) \tag{17}$$

We define the "DP as a Discount" method (DPAD) as:

$$\mathsf{DPAD}(D) = Y_{\text{raw}} + Z_{\text{exter}}, \quad Z_{\text{exter}} \sim \mathcal{N}(0, \sigma^2_{\text{exter}}) \tag{18}$$

Equivalently,

$$\mathsf{DPAD}(D) \sim \mathcal{N}\left(f(D), \sigma^2_{\text{inter}} + \sigma^2_{\text{exter}}\right) \tag{19}$$

3.2. q-IND-CPAD

This section proves that under the abovementioned model, DPAD satisfies IND–CPA$^{\mathsf{D}}$.

Theorem 3.3 (Hybrid Method Security Guarantee)**.** *Under the assumptions of 3.1 and 3.2, if the base approximate homomorphic encryption scheme is $(\kappa + 8)$-bit IND-CPA secure and the noise parameters satisfy*

$$\sup_t \frac{t^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})} \le \frac{2^{-\kappa-7}}{q} \tag{20}$$

then the sanitized mechanism $\mathsf{DPAD}(D)$ as defined in 3.2 with modified decryption

$$\mathsf{Dec}'_{sk}(ct) = \mathsf{Dec}_{sk}(ct) + Z_{\text{exter}}, \quad Z_{\text{exter}} \sim \mathcal{N}(0, \sigma^2_{\text{exter}}) \tag{21}$$

is κ-bit q-IND–CPA^D secure.
Furthermore, the KL divergence parameter is reduced. That is, for all $\sigma_{\text{exter}} > 0$,

$$\rho' \le \rho$$

Proof. The core idea of the proof is similar to 2.1. We first establish a sequence of games, then we reduce q-IND–CPA$^{\mathsf{D}}$ to IND–CPA on DPAD. The two games established are as follows:

- Game G_0: The real q-IND–CPA$^{\mathsf{D}}$ game. The adversary A interacts with the encryption and evaluation oracles. Upon a decryption query i for a ciphertext ct with public error bound t, the oracle returns $Y_i = m + \beta + Z_{\text{inter}} + Z_{\text{exter}}$.
- Game G_1: A modified game where the decryption oracle is replaced with an "idealized" version. Instead of decrypting the ciphertext, it uses the exact plaintext m and returns $Y'_i = m + Z'_{\text{inter}} + Z'_{\text{exter}}$.

Next, we bound the divergence between G_0 and G_1. By 3.1, the decryption of ct in G_0 satisfies

$$Y_i = m + \beta + Z_{\text{inter}} + Z_{\text{exter}}$$

Since Z_{inter} and Z_{exter} are independent Gaussians,

$$Z_{\text{inter}} + Z_{\text{exter}} \sim \mathcal{N}(0, \sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})$$

Hence, both the outputs Y_i and Y_i' are Gaussian distributions with the same variance

$$\sigma'^2 = \sigma^2_{\text{inter}} + \sigma^2_{\text{exter}}$$

but differ in means:

$$\mu_0 = m + \beta$$
$$\mu_1 = m$$

By KL Divergence of Equal-Variance Gaussians, we get:

$$\mathsf{KL}(Y_i||Y_i') = \frac{(\mu_0 - \mu_1)^2}{2\sigma'^2} = \frac{\beta^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})}$$

From 3.1, given that $|\beta| \leq t$, we have

$$\mathsf{KL}(Y_i||Y_i') \leq \frac{t^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})}$$

By the sub-additivity property of KL-divergence,

$$\mathsf{KL}(G_0||G_1) \leq \sum_{i=1}^{q} \mathsf{KL}(Y_i||Y_i') \leq \sum_{i=1}^{q} \sup_t \frac{t^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})}$$

By definition, we have

$$\sum_{i=1}^{q} \sup_t \frac{t^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})} \leq q \cdot \frac{2^{-\kappa-7}}{q} = 2^{-\kappa-7}$$

In Game G_1, the decryption results Y_i' are only dependent on the exact message m and the upperbound of ciphertext error t. Thus, an adversary A' for the IND–CPA game can perfectly simulate the G_1 decryption oracle because m is known in the context of the challenge. Hence, if G_0 is $(\kappa + 8)$-bit IND-CPA secure, then G_1 is also $(\kappa + 8)$-bit secure.

From 2.3, it stated that if Game G_1 is $(\kappa+8)$-bit secure and the KL divergence

$$\mathsf{KL}(G_0||G_1) \leq 2^{-(\kappa+8)+1} = 2^{-\kappa-7}$$

Then Game G_0 is at least $((\kappa+8)-8) = \kappa$-bit secure. This concludes that the hybrid mechanism is κ-bit q-IND–CPA$^{\mathsf{D}}$ secure.

Finally, we prove that $\rho' \leq \rho$. To obtain the value of ρ, we consider the model where the internal noise is disregarded. In other words, $\sigma_{\text{inter}} = 0$. Assume the same amount of noise with variance σ^2_{exter} is added to this model. Then, given the two same adjacent database $x \sim x'$, we have

$$\rho = \mathsf{KL}(\mathsf{Gauss}(x)||\mathsf{Gauss}(\mathsf{x}')) = \frac{(x-x')^2}{2\sigma^2_{\text{exter}}} \leq \frac{t^2}{2\sigma^2_{\text{exter}}}$$

Since $\sigma_{\text{exter}} > 0$,

$$\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}} > \sigma^2_{\text{exter}}$$

we have

$$\frac{1}{\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}}} < \frac{1}{\sigma^2_{\text{exter}}}$$

Multiplying both sides by $\frac{t^2}{2}$ yields:

$$\frac{t^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})} < \frac{t^2}{2\sigma^2_t}$$

where $Z \sim \mathcal{N}(0, 1)$.

Finally, taking the supremum over t gives:

$$\rho' = \sup_t \frac{t^2}{2(\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}})} \leq \sup_t \frac{t^2}{2\sigma^2_{\text{exter}}} = \rho$$

□

3.3. Differential Privacy

In this subsection, we prove that adding independent noise as mentioned in (Li et al., 2022) reduces the privacy budget ε. To do so, we must demonstrate that the hybrid variance reduces the message-dependence ratio ($\widetilde{T}$) and the sensitivity-to-noise ratio ($\widetilde{K}$). We will also prove that the ε defined in 2.6 is monotonically increasing with respect to $\widetilde{T}$ and $\widetilde{K}$. Finally, we will demonstrate that $\widetilde{T}$ and $\widetilde{K}$ asymptotically approaches 1 and 0 respectively as the added variance increases.

Theorem 3.4 (Reduction of the Message-Dependence Ratio)**.** *Given two adjacent databases D and D' such that $D \sim D'$. In an approximate homomorphic encryption scheme, define the standard deviation of the internal noise of D and D' as σ_D and $\sigma_{D'}$ respectively. Define the inherent message-dependence ratio of noise standard deviations of D and D' as*

$$T = \max \frac{\sigma_D}{\sigma_{D'}}$$

By post-processing the decryption output with independent Gaussian noise of variance σ^2_{exter}, the resulting hybrid message-dependence ratio is

$$\widetilde{T} = \max \sqrt{\frac{\sigma_D^2 + \sigma^2_{\text{exter}}}{\sigma^2_{D'} + \sigma^2_{\text{exter}}}}$$

Then, for all $\sigma_{\text{exter}} > 0$,

$$\widetilde{T} < T \tag{22}$$

In addition,

$$\lim_{\sigma_{\text{exter}} \to \infty} \widetilde{T} = 1 \tag{23}$$

Proof. First, we define a function $f(x)$ that represents squared ratio of the total variances of the two adjacent databases as a function of the sanitization noise $x := \sigma^2_{\text{exter}}$:

$$f(x) = \frac{\sigma_D^2 + x}{\sigma_{D'}^2 + x} \tag{24}$$

In order to determine how the ratio changes as independent noise is added, we calculate the first derivative of $f(x)$ with respect to x using the quotient rule:

$$f'(x) = \frac{(\sigma_{D'}^2 + x) \cdot \frac{d}{dx}(\sigma_D^2 + x) - (\sigma_D^2 + x) \cdot \frac{d}{dx}(\sigma_{D'}^2 + x)}{(\sigma_{D'}^2 + x)^2}$$

Simplifying the numerator yields

$$f'(x) = \frac{(\sigma_{D'}^2 + x) \cdot 1 - (\sigma_D^2 + x) \cdot 1}{(\sigma_{D'}^2 + x)^2} = \frac{\sigma_{D'}^2 - \sigma_D^2}{(\sigma_{D'}^2 + x)^2}$$

Without loss of generality, assume that $\sigma_D^2 > \sigma_{D'}^2$, which is required for the original ratio T to be greater than 1. Under this assumption, the numerator $(\sigma_{D'}^2 - \sigma_D^2)$ is strictly negative. Since the denominator $(\sigma_{D'}^2 + x)^2$ is always positive for all $x \geq 0$, it follows that

$$f'(x) < 0$$

And because the derivative is strictly negative, $f(x)$ is a strictly decreasing function of the added noise. This implies that for any $\sigma^2_{\text{exter}} > 0$,

$$f(\sigma^2_{\text{exter}}) < f(0) \implies \frac{\sigma_D^2 + \sigma^2_{\text{exter}}}{\sigma_{D'}^2 + \sigma^2_{\text{exter}}} < \frac{\sigma_D^2}{\sigma_{D'}^2}$$

Taking the square root of both sides preserves the inequality, so

$$\widetilde{T} < T$$

Finally, to evaluate the limit as $\sigma_{\text{exter}} \to \infty$, we first divide both the numerator and the denominator within the radical by σ^2_{exter}:

$$\begin{aligned}\lim_{\sigma^2_{\text{exter}} \to \infty} \widetilde{T} &= \lim_{\sigma^2_{\text{exter}} \to \infty} \sqrt{\frac{\sigma_D^2 + \sigma^2_{\text{exter}}}{\sigma_{D'}^2 + \sigma^2_{\text{exter}}}} \\ &= \lim_{\sigma^2_{\text{exter}} \to \infty} \sqrt{\frac{\frac{\sigma_D^2}{\sigma^2_{\text{exter}}} + 1}{\frac{\sigma_{D'}^2}{\sigma^2_{\text{exter}}} + 1}}\end{aligned}$$

Since the inherent variances σ_D^2 and $\sigma_{D'}^2$ are fixed constants for any D and D', their ratios to the increasing noise floor vanish:

$$\lim_{\sigma^2_{\text{exter}} \to \infty} \frac{\sigma_D^2}{\sigma^2_{\text{exter}}} = 0 \ \wedge \ \lim_{\sigma^2_{\text{exter}} \to \infty} \frac{\sigma_{D'}^2}{\sigma^2_{\text{exter}}} = 0$$

Substituting these values into the expression yields:

$$\lim_{\sigma^2_{\text{exter}} \to \infty} \widetilde{T} = \sqrt{\frac{0+1}{0+1}} = \sqrt{1} = 1$$

□

Theorem 3.5 (Reduction of the Sensitivity Ratio)**.** *Given two adjacent databases D and D' such that $D \sim D'$. In an approximate homomorphic encryption scheme, define*

$$K = \frac{\Delta_f}{\sigma_D}$$

as the sensitivity-to-noise ratio, where

$$\Delta_f = \max \|\beta_D - \beta_{D'}\|_2$$

is the l_2 sensitivity of the algorithm and σ_D is the standard deviation of the inherent approximate homomorphic encryption noise. By post-processing the decryption output with independent Gaussian noise of variance σ^2_{exter}, define the hybrid sensitivity ratio as

$$\widetilde{K} = \frac{\Delta_f}{\sqrt{\sigma_D^2 + \sigma^2_{\text{exter}}}}$$

Then, for all $\sigma_{\text{exter}} > 0$,

$$\widetilde{K} < K$$

In addition,

$$\lim_{\sigma_{\text{exter}} \to \infty} \widetilde{K} = 0$$

Proof. Consider the sensitivity ratio as a function $f(x)$ of the total noise standard deviation. Then, Δ_f represents the maximum possible change in the "true" output β across any two adjacent databases D and D'. This value is determined solely by the algorithm and the database constraints, and is independent of the encryption noise. The total standard deviation of the hybrid mechanism, $\widetilde{\sigma_D}$, is defined as the square root of the sum of the inherent variance and the added variance:

$$\widetilde{\sigma_D} = \sqrt{\sigma_D^2 + \sigma^2_{\text{exter}}}$$

Thus, for any $\sigma^2_{\text{exter}} > 0$, it follows that:

$$\sigma_D^2 + \sigma^2_{\text{exter}} > \sigma_D^2 \implies \sqrt{\sigma_D^2 + \sigma^2_{\text{exter}}} > \sigma_D$$

Since the numerator Δ_f remains constant and the denominator has strictly increased ($\widetilde{\sigma_D} > \sigma_D$), the resulting fraction must be strictly smaller. Hence,

$$\frac{\Delta_f}{\widetilde{\sigma_D}} < \frac{\Delta_f}{\sigma_D} \implies \widetilde{K} < K$$

Finally, it is trivial that

$$\lim_{\sigma_{\text{exter}} \to \infty} \frac{\Delta_f}{\sqrt{\sigma_D^2 + \sigma^2_{\text{exter}}}} = 0$$

□

Now that we have proven that $\widetilde{T}$ and $\widetilde{K}$ decrease as σ_{exter} increases, we can bound the resulting reduction in ε.

Theorem 3.6 (Reduction of the Privacy Budget ε)**.** *Let $\varepsilon := \epsilon(T, K)$ be the required privacy budget for an approximate homomorphic encryption scheme. Let $\widetilde{\varepsilon} := \epsilon(\widetilde{T}, \widetilde{K})$ be the privacy budget of "DP as a Discount" with additional noise added to reduce the message-dependence ratio from T to $\widetilde{T}$ and the sensitivity-to-noise ratio from K to $\widetilde{K}$. Then, the resulting mechanism is ($\widetilde{\varepsilon}$, δ)-DP if*

$$\widetilde{\epsilon} > \widetilde{T}^2\widetilde{K}\sqrt{L} + \frac{1}{2}\widetilde{T}^2\widetilde{K}^2 + \frac{1}{2}(\widetilde{T}^2 - 1)L + \ln\widetilde{T}$$

In addition, the total required privacy budget is strictly reduced. That is,

$$\widetilde{\varepsilon} < \varepsilon$$

Proof. According to 2.6, the required privacy budget ε for a one-dimensional output is

$$\epsilon(T, K) = T^2K\sqrt{D} + \frac{1}{2}T^2K^2 + \frac{1}{2}(T^2 - 1)D + \ln T$$

where

$$D = 2\ln\left(\frac{1}{\delta}\right)$$

is a constant determined by the privacy failure probability δ. To determine the behavior of ϵ, we examine its partial derivatives with respect to the two independent variables, K and T.

First of all,

$$\frac{\partial\epsilon}{\partial K} = T^2\sqrt{D} + T^2K$$

Since $T \geq 1$, $K > 0$, and $D > 1$, it follows that

$$\frac{\partial\epsilon}{\partial K} > 0$$

Thus, ϵ is a strictly increasing function of the sensitivity ratio K.

Secondly,

$$\frac{\partial\epsilon}{\partial T} = 2TK\sqrt{D} + TK^2 + TD + \frac{1}{T}$$

Given that T, K, D are all positive, every term in the expression is positive. This means that

$$\frac{\partial\epsilon}{\partial T} > 0$$

Thus, ϵ is also a strictly increasing function of the message-dependence ratio T. As $\epsilon(T, K)$ is strictly increasing in both T and K, any mechanism that simultaneously achieves $\widetilde{T} < T$ and $\widetilde{K} < K$ will result in a smaller required privacy budget. Hence,

$$\widetilde{\varepsilon} < \varepsilon$$

□

This proves that the "DP as a discount" approach can be optimized and hardened by reducing message dependence.

3.4. Result Discussion

The "DP as a Discount" hybrid approach introduces an additional Gaussian noise term with variance σ^2_{exter}, which increases the total noise variance in the system to $\sigma^2_{\text{inter}} + \sigma^2_{\text{exter}}$. This added noise plays an important role in both privacy and security.

From the privacy side, increasing σ_{exter} makes the output distributions for adjacent datasets more similar. As σ_{exter} approaches infinity, the message-dependence ratio $\widetilde{T}$ decreases toward 1, and the sensitivity ratio $\widetilde{K}$ decreases toward 0. This means that the output becomes less dependent on the specific dataset, which improves the differential privacy guarantee.

From the security perspective, the added noise hides the decryption error. Since the adversary only observes noisy outputs, it becomes harder to leak information from the error. In particular, increasing σ_{exter} reduces the KL divergence, resulting in limiting the adversary's advantage.

However, there is a clear drawback as σ_{exter} becomes too large. In the extreme case, as $\sigma_{\text{exter}} \to \infty$ the output becomes almost pure noise and carries little useful information about the true result. This makes the mechanism impractical despite its strong privacy and security guarantees.

In short, σ_{exter} can act as a hyperparameter. Increasing it improves privacy and security, but reduces utility. In practice, it must be chosen carefully to balance these privacy and security with utility.

3.5. Implication

Homomorphic encryption remains difficult to deploy in real-world systems because of its performance limitations. Approximate homomorphic encryption addresses this limitation by allowing small errors in the decrypted result, trading off correctness with efficiency. Yet, the core issue with standard approximate homomorphic encryption schemes is that they are not secure against the stronger attack model $\mathsf{IND\text{–}CPA^D}$, where adversaries can exploit information leakage from approximate decryption outputs. Prior work proposed adding carefully calibrated noise during decryption to mitigate this issue. However, doing so degrades the accuracy of the final output substantially, hindering the practicality of approximate homomorphic encryption.

Nevertheless, our work shows that the intrinsic noise present in approximate homomorphic encryption can be leveraged to reduce the amount of externally added noise. Our approach achieves the same level of q-$\mathsf{IND\text{–}CPA^D}$ security and differential privacy with significantly less additional noise

for a carefully chosen σ_{exter}. We hope that our contribution can make systems that use approximate homomorphic encryption more practical and feasible.

3.6. Future Works

First, our analysis shows that internal noise can reduce the required added external noise to guarantee security. However, the bounds used in this work might not be tight. Future work could focus on deriving tighter bounds on the KL-divergence and the privacy parameters, leading to more precise noise utility trade-off. In addition, our framework assumes that the error introduced in approximate homomorphic encryption is Gaussian. This may not accurately capture real-world behavior when noise distributions might deviate due to implementation. Thus, an important future work direction is to extend the analysis to a more general noise distributions. Finally, our results are theoretical, and validating our proposed framework in practice is extremely crucial. Thus, future work could focus on implementing the DPAD mechanism in real-world homomorphic encryption libraries and evaluating its performance.

4. Conclusion

In summary, this work establishes a connection between cryptographic security and differential privacy in approximate homomorphic encryption. We proposed a hybrid framework that combines internal noise from approximate homomorphic encryption with external Gaussian noise introduced at decryption. This approach provides a novel way to achieve both (ε, δ)-differential privacy and $\mathsf{IND\text{–}CPA^D}$ security simultaneously.

Our analysis shows that internal noise reduces the amount of external noise that needs to be added during decryption. Adding a carefully calculated external noise allows us to reduce both the message-dependence ratio and the sensitivity ratio. This enables the mechanism to achieve strong privacy and security guaranties. However, as more noise is added, the utility of the output degrades significantly. Therefore, noise must be chosen carefully to balance between privacy and security with practicality.

References

Acar, A., Aksu, H., Uluagac, A. S., and Conti, M. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (Csur)*, 51(4):1–35, 2018.

Alaya, B., Laouamer, L., and Msilini, N. Homomorphic encryption systems statement: Trends and challenges. *Computer Science Review*, 36:100235, 2020.

Cheon, J. H., Kim, A., Kim, M., and Song, Y. Homomorphic encryption for arithmetic of approximate numbers. In *International conference on the theory and application of cryptology and information security*, pp. 409–437. Springer, 2017.

Costache, A., Curtis, B. R., Hales, E., Murphy, S., Ogilvie, T., and Player, R. On the precision loss in approximate homomorphic encryption. In *International Conference on Selected Areas in Cryptography*, pp. 325–345. Springer, 2023.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference (TCC)*, pp. 265–284, 2006.

Gentry, C. *A fully homomorphic encryption scheme*. Stanford university, 2009.

Katz, J. and Lindell, Y. *Introduction to Modern Cryptography*. CRC Press, 2007.

Li, B. and Micciancio, D. On the security of homomorphic encryption on approximate numbers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 648–677. Springer, 2021.

Li, B., Micciancio, D., Schultz-Wu, M., and Sorrell, J. Securing approximate homomorphic encryption using differential privacy. In *Annual International Cryptology Conference*, pp. 560–589. Springer, 2022.

Micciancio, D. and Walter, M. On the bit security of cryptographic primitives. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 3–28. Springer, 2018.

Mironov, I. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pp. 263–275. IEEE, 2017.

Ogilvie, T. Differential privacy for free? harnessing the noise in approximate homomorphic encryption. In *Cryptographers' Track at the RSA Conference*, pp. 292–315. Springer, 2024.

Wang, Y.-X., Lei, J., and Fienberg, S. E. On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms. In *International Conference on Privacy in Statistical Databases*, pp. 121–134. Springer, 2016.