Analysis of Shuffling Beyond Pure Local Differential Privacy*

SHUN TAKAGI, LY Corporation, Japan
SENG PEI LIEW, LY Corporation, Japan

Shuffling is a powerful way to amplify privacy of a local randomizer in private distributed data analysis. Most existing analyses of how shuffling amplifies privacy are based on the pure local differential privacy (DP) parameter ε_0. This paper raises the question of whether ε_0 adequately captures the privacy amplification. For example, since the Gaussian mechanism does not satisfy pure local DP for any finite ε_0, does it follow that shuffling yields weak amplification? To solve this problem, we revisit the privacy blanket bound of Balle et al. (the *blanket divergence*) and develop a direct asymptotic analysis that bypasses ε_0. Our key finding is that, asymptotically, the blanket divergence depends on the local mechanism only through a single scalar parameter χ and that this dependence is monotonic. Therefore, this parameter serves as a proxy for shuffling efficiency, which we call the *shuffle index*. By applying this analysis to both upper and lower bounds of the shuffled mechanism's privacy profile, we obtain a band for its privacy guarantee through shuffle indices. Furthermore, we derive a simple structural, necessary and sufficient condition on the local randomizer under which this band collapses asymptotically. k-RR families with $k \geq 3$ satisfy this condition, while for generalized Gaussian mechanisms the condition may not hold but the resulting band remains tight. Finally, we complement the asymptotic theory with an FFT-based algorithm for computing the blanket divergence at finite n, which offers rigorously controlled relative error and near-linear running time in n, providing a practical numerical analysis for shuffle DP.

CCS Concepts: • **Security and privacy**;

Additional Key Words and Phrases: Differential Privacy, Shuffle Model

ACM Reference Format:
Shun Takagi and Seng Pei Liew. 2026. Analysis of Shuffling Beyond Pure Local Differential Privacy. *Proc. ACM Manag. Data* 4, 2 (PODS), Article 96 (May 2026), 26 pages. https://doi.org/10.1145/3801892

1 Introduction

Local privacy [11], in which each user applies a local randomizer to their data before sending it, removes the need for a trusted curator and thus enables distributed private data analysis. However, the price of eliminating trust is often a substantial loss in accuracy. Shuffling, viewed as an anonymization layer that breaks the link between users and their messages, offers a compelling middle ground. Since it only anonymizes messages across users, it preserves the utility of the analytics tasks already supported in the local model, such as frequency estimation, histograms/heavy hitters, and mean estimation. Moreover, a growing line of work shows that this simple step can dramatically amplify privacy, allowing protocols to approach trusted-curator privacy-utility trade-offs [4, 13, 14]. Given its practical utility, shuffling has become a key primitive for private distributed data collection and analytics [19, 20].

*This is the conference version. Please refer to the full version of this paper for appendices that include full proofs [27].

Authors' Contact Information: Shun Takagi, LY Corporation, Tokyo, Japan, shutakag@lycorp.co.jp; Seng Pei Liew, LY Corporation, Tokyo, Japan, sliew@lycorp.co.jp.

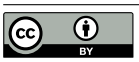

ACM 2836-6573/2026/5-ART96
https://doi.org/10.1145/3801892

From a practical data-management perspective, we need precise analysis of how and to what extent shuffling amplifies privacy to choose the best local randomizer for a given task. However, existing tools are local randomizer-agnostic and largely tailored to the pure local differential privacy (DP) setting [4, 13] (i.e., $(\varepsilon_0, 0)$-local DP) and this ε_0-centric viewpoint has two important limitations.

First, the parameter ε_0 summarizes the local privacy guarantee, but it is a rather crude descriptor from the viewpoint of shuffling: it largely ignores the structural properties of the local mechanism that actually govern its privacy amplification. In particular, existing general bounds typically do not distinguish between, say, a Laplace mechanism with a given noise scale, a k-randomized response (k-RR) mechanism with a given k, or other natural families of local mechanisms. As a consequence, generic upper bounds expressed solely in terms of ε_0 can be quite loose. Indeed, several works have shown that by exploiting the specific structure of a given mechanism one can obtain significantly tighter bounds on the effect of shuffling (e.g., for k-RR one can show that privacy amplification becomes stronger as k increases, even when ε_0 is kept fixed [7, 14]). Yet, we currently lack a general understanding of what intrinsic properties of a local mechanism govern its "shuffle efficiency."

Second, the restriction to pure local DP mechanisms excludes a wide range of natural and practically relevant local randomizers. Pure local DP is a convenient sufficient condition for privacy amplification by shuffling, but it is far from necessary. In practice, one often encounters approximate-DP local mechanisms or mechanisms that do not satisfy any DP condition. Surprisingly, even for the Gaussian mechanism, arguably one of the most prominent mechanisms in central DP, it is not known how to precisely characterize privacy amplification by shuffling. Balle et al. [4] already pointed out that the Gaussian case is technically challenging, and subsequent work by Koskela et al. [17] observed that numerical accounting for Gaussian shuffling is nontrivial. As a result, to date, existing results for shuffling of Gaussian mechanisms are essentially limited to lower bounds [10, 18]. While there are attempts to generalize shuffling analysis to approximate local DP (i.e., $(\varepsilon_0, \delta_0 > 0)$-local DP) in a mechanism-agnostic fashion [9, 14], such general bounds tend to be overly pessimistic; they often show little visible amplification for moderate values of n, because they are driven by worst-case mechanisms that do not reflect the structure of the specific mechanism.

The privacy blanket and blanket divergence. A central tool in the analysis of shuffle DP is the *privacy blanket* introduced by Balle et al. [4]. Given a local randomizer $\mathcal{R}$, the blanket distribution is a data-independent distribution $\mathcal{R}_{\mathrm{BG}}$ defined using $\mathcal{R}$. Balle et al. showed that the privacy profile $\delta(\varepsilon)$ [3] of the shuffled mechanism can be upper-bounded by a divergence computed under the blanket distribution; we refer to this quantity as the *blanket divergence*. This blanket-divergence bound is currently the best-known general upper bound on privacy amplification by shuffling [25].

However, the blanket divergence itself is still poorly understood, largely due to its analytical and numerical intractability. In particular, Balle et al. [4] mainly derive generic upper bounds via Hoeffding- and Bennett-type concentration inequalities based on ε_0. Consequently, these bounds can be quite loose and fail to capture how the blanket divergence depends on the structure of the underlying local randomizer.

This is partly because Balle et al. explicitly aim to obtain finite-n upper bounds. More generally, carrying out a precise finite-n analysis of shuffling through the blanket divergence is intrinsically challenging: DP requires a worst-case guarantee over all neighboring datasets (i.e., establishing a uniform upper bound over all neighboring input pairs). Identifying the pairs that maximize the blanket divergence at finite n is highly nontrivial in general. In contrast, in the ε_0-local DP setting, the single parameter ε_0 already summarizes the worst case over all neighboring input pairs, so concentration-based bounds can be expressed directly in terms of ε_0 without explicitly resolving the supremum over all neighboring input pairs.

Our goal and approach. In light of the above limitations, our goal is to circumvent the apparent finite-n bottlenecks of generic concentration-based analyses by taking a different route, namely an asymptotic but direct analysis of the blanket divergence. While this sacrifices exactness at any fixed n, it is tight and allows us to address the conceptual and practical limitations discussed above.

Concretely, the relevant quantity in blanket divergence can be written as a sum of n i.i.d. random variables, so that its behavior is driven by a central limit theorem (CLT)-type phenomenon. We leverage asymptotic expansions of the CLT [21] to obtain a sharp leading-order characterization of the blanket divergence. Formally, in the regime[1] $\varepsilon_n = \omega(n^{-1/2})$ and $\varepsilon_n = O(\sqrt{\log n/n})$, the blanket divergence admits the asymptotic expansion

$$\mathcal{D}^{\text{blanket}} = \varphi\big(\chi\,\varepsilon_n\sqrt{n}\big)\left(\frac{1}{\chi^3}\,\frac{1}{\varepsilon_n^2 n^{3/2}}\right)\big(1+o(1)\big),$$

where φ denotes the PDF of the standard normal distribution and χ is a mechanism-dependent constant easily computed from the structure of the local randomizer. Thus χ fully governs how well a given local randomizer interacts with shuffling. Crucially, this dependence is monotonic: a larger χ results in a smaller divergence, thereby implying stronger privacy guarantees. This means that χ serves as a single-number index of privacy after shuffling; hence, we call it the *shuffle index*, which substitutes for ε_0 as a more refined descriptor of the local randomizer's shuffle efficiency. In other words, we can choose the best local randomizer for shuffling by choosing the one with the largest shuffle index χ.

Moreover, we obtain an asymptotic characterization of the worst-case blanket divergence and hence the shuffled mechanism's $(\varepsilon \approx \frac{1}{\chi}\sqrt{\frac{\log n}{n}}, \delta \approx \frac{\alpha}{n})$ guarantee of DP for any $\alpha > 0$. Because we can express not only upper bounds but also lower bounds in terms of the blanket divergence, we can quantitatively assess the tightness of the upper bound. In particular, we derive a simple structural necessary and sufficient condition on the local randomizer under which the blanket divergence yields an asymptotically optimal characterization (i.e., matches the upper bound and the lower bound) of shuffling. This condition is satisfied, for example, by the k-RR family with $k \geq 3$. Moreover, we show that the Laplace and Gaussian mechanisms do not satisfy the condition but the upper and lower bounds narrow in higher privacy regimes.

On the practical side, asymptotic expansions alone are not sufficient for numerical accounting at finite n. We therefore complement our asymptotic analysis with a finite-n algorithm for computing the blanket divergence. Our algorithm is based on an FFT approximation of the distribution, combined with explicit control of the truncation, discretization, and aliasing errors. Leveraging the asymptotic expansion, we show that the parameters can be tuned by a parameter η so that the relative error is $O(\eta)$, while the running time scales as $\widetilde{O}(n/\eta)$.

Related Work and Our Contributions

A line of work seeks to go beyond the ε_0-centric view by exploiting additional structure of the local randomizer, typically in the form of a small number of scalar parameters. Feldman et al. [15] introduced a stronger clone paradigm and showed that, for mechanisms such as k-RR, one can parameterize the effect of shuffling using ε_0 together with two additional scalar quantities, yielding substantially tighter bounds than ε_0-only analyses, although even in the k-RR case exact optimality is not established. Building on this, Wang et al. [28, 29] refine the stronger clone-based analysis, using two scalar parameters, and provide sufficient conditions on the local randomizer under which their bounds become optimal. A subtle point is that the proof of Lemma 4.5 in [28] implicitly assumes the existence of a nontrivial post-processing map (see the errata of [15]). Outside specific

[1]This corresponds to the moderate deviation regime required to apply the asymptotic expansion of the CLT.

mechanism classes such as k-RR where no such map is needed (see Theorem 7.1 of [15]), the results should therefore be interpreted with some care. Two concurrent and independent works derive quantities corresponding to the shuffle index [22, 26]. However, Shvets [22] considers only binary datasets with finite output alphabets, while Su et al. [26] focus on mutual information bounds rather than DP guarantees. Our results are formulated in a more general framework based on DP notions, encompassing a broader class of local randomizers and thus subsuming these settings.

Several works have proposed numerical methods for evaluating shuffle DP guarantees. Koskela et al. [17] use the clone paradigm together with FFT to approximate the privacy loss of shuffled mechanisms. Their approach, however, is restricted to pure local DP mechanisms and inherits the looseness of the clone paradigm, as shown by Su et al. [25]. Su et al. [25] propose a numerical analysis of the blanket divergence based on FFT, similar in spirit to ours. However, they still assume pure local DP and do not provide rigorous control of the relative error; in addition, the running time scales as $O(n^2)$ in the worst case when targeting a constant additive error. Biswas et al. [7] give an exact combinatorial method for computing the privacy guarantees of shuffled k-RR, but their analysis focuses on fixed neighboring database pairs rather than general (ε, δ)-DP guarantees.

Our contributions. Compared with the above works, our contributions are as follows:

- We are the first to give a unified shuffle DP analysis that does not assume pure local DP and applies to arbitrary local randomizers under mild regularity assumptions.
- We theoretically analyze the blanket divergence including the exact leading constant factor and summarize its leading behavior by a single shuffle index χ, yielding mechanism-aware bounds and a necessary-and-sufficient optimality condition.
- We develop an FFT-based blanket-divergence accountant with rigorous relative-error $O(\eta)$ guarantees and near-linear $\widetilde{O}(n/\eta)$ running time.
- We empirically validate our asymptotic and FFT-based analysis, and in a distribution-estimation task we show that generalized Gaussian mechanisms achieve favorable privacy-utility trade-offs compared with pure local DP mechanisms.

2 Preliminaries

Notations. Φ and φ denote the CDF and PDF of the standard normal distribution. $\Gamma(\cdot)$ denotes the Gamma function, $\Gamma(t) = \int_0^\infty u^{t-1}\mathrm{e}^{-u}\,du$ for $t > 0$. $W(\cdot)$ denotes the principal branch of the Lambert W-function. Throughout the paper, whenever an expression can be written in terms of Φ, Γ, W, and elementary functions, we will also refer to it as being available in closed form. $\mathbb{E}[X]$ and $\mathrm{Var}(X)$ denote the expectation and variance of a random variable X. $\mathbb{R}$ and $\mathbb{N}$ denote the set of real numbers and natural numbers, respectively. $\mathbf{1}\{\cdot\}$ denotes the indicator function, which equals 1 if the condition inside the brackets is satisfied, and 0 otherwise. We use standard asymptotic notations $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, $\omega(\cdot)$, and $\Theta(\cdot)$.

2.1 Differential Privacy (DP)

We first recall the notion of DP and express it in a form based on the hockey-stick divergence.

Definition 2.1 (Hockey-stick divergence). Let $\mathcal{Y}$ be an output space equipped with a reference measure (counting measure in the discrete case, Lebesgue measure in the continuous case). Let P and Q be probability distributions on $\mathcal{Y}$ that admit densities p and q with respect to this reference measure, and fix a parameter $\alpha \geq 1$. The *hockey-stick divergence* of P from Q of order α is

$$\mathcal{D}_\alpha(P\|Q) := \int_{\mathcal{Y}} \big[p(y) - \alpha\, q(y)\big]_+ \, dy,$$

where $[u]_+ := \max\{u, 0\}$. In the discrete case, the integral should be read as a sum over $y \in \mathcal{Y}$.

Definition 2.2 (Differential privacy via hockey-stick divergence [6, 12]). Let $\mathcal{X}$ be an input domain and $\mathcal{Z}$ an output space. A randomized mechanism is a map $\mathcal{M} : \mathcal{X}^n \to \mathcal{Z}$ that takes a dataset $x_{1:n} \in \mathcal{X}^n$ as input and outputs a random element of $\mathcal{Z}$. For each dataset $x_{1:n}$ we write $\mathcal{M}(x_{1:n})$ for the corresponding output distribution on $\mathcal{Z}$, and we denote its density by the same symbol when convenient. Two datasets $x_{1:n}, x'_{1:n} \in \mathcal{X}^n$ are said to be neighboring ($x_{1:n} \simeq x'_{1:n}$) if they differ in exactly one element. For a mechanism $\mathcal{M}$, $\varepsilon \geq 0$, and $\delta \in [0, 1]$, $\mathcal{M}$ is (ε, δ)-DP if and only if

$$\delta_{\mathcal{M}}(\varepsilon) = \sup_{x_{1:n} \simeq x'_{1:n}} \mathcal{D}_{e^\varepsilon}\left(\mathcal{M}(x_{1:n}) \,\middle\|\, \mathcal{M}(x'_{1:n})\right) \leq \delta.$$

The function $\delta_{\mathcal{M}}(\varepsilon)$ is referred to as the privacy profile of $\mathcal{M}$ [3]. Sometimes, we say that $\mathcal{M}$ satisfies (ε, δ)-DP for $x_{1:n} \simeq x'_{1:n}$ if $\mathcal{D}_{e^\varepsilon}\left(\mathcal{M}(x_{1:n}) \,\middle\|\, \mathcal{M}(x'_{1:n})\right) \leq \delta$ [24].

In the shuffle model, the mechanism $\mathcal{M}$ is typically constructed by applying a local randomizer independently to each user's input.

Definition 2.3 (Local randomizer and blanket distribution). Let $\mathcal{X}$ be an input space and $\mathcal{Y}$ an output space. A *local randomizer* is a randomized mechanism $\mathcal{R} : \mathcal{X} \to \mathcal{Y}$. For each input $x \in \mathcal{X}$ we write $\mathcal{R}_x$ for the distribution of the random output $\mathcal{R}(x)$. We assume that there is a reference measure on $\mathcal{Y}$ (counting measure in the discrete case, Lebesgue measure in the continuous case) such that each $\mathcal{R}_x$ admits a density, and we write $\mathcal{R}_x(y)$ for this probability mass function or density, depending on the context. The *blanket distribution* $\mathcal{R}_{\mathrm{BG}}$ associated with $\mathcal{R}$ is defined for $y \in \mathcal{Y}$ by,

$$\underline{\mathcal{R}}(y) := \inf_{x \in \mathcal{X}} \mathcal{R}_x(y), \qquad \gamma := \int_{\mathcal{Y}} \underline{\mathcal{R}}(y)\, dy \in (0, 1], \qquad \mathcal{R}_{\mathrm{BG}}(y) := \frac{1}{\gamma}\underline{\mathcal{R}}(y).$$

where, in the purely discrete case, the integral should be read as a sum over $y \in \mathcal{Y}$. The scalar γ is called the *blanket mass* of $\mathcal{R}$.

Following Balle et al. [4], we work in the single-message *randomize-then-shuffle* model. That is, we assume black-box access to an ideal functionality *shuffler*[2].

Definition 2.4 (Single-message shuffling and shuffled mechanism). Let $\mathcal{R}$ be a local randomizer. Given a dataset $x_{1:n} \in \mathcal{X}^n$, each user $i \in [n]$ applies the local randomizer to obtain a single message $Y_i := \mathcal{R}(x_i) \in \mathcal{Y}$. The *shuffler* is a randomized map $\mathcal{S} : \mathcal{Y}^n \to \mathcal{Y}^n$ which samples a permutation π uniformly at random from the symmetric group on $[n]$ and outputs $\mathcal{S}(y_1, \ldots, y_n) := (y_{\pi(1)}, \ldots, y_{\pi(n)})$. A shuffled mechanism is $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$, which shuffles the outputs of $\mathcal{R}$.

Shuffled mechanisms are the main object of study in this paper and the following upper bound is known.

Lemma 2.5 (Lemma 5.3 of Balle et al. [4]). *Fix $\varepsilon \geq 0$, let $x_{1:n} \simeq x'_{1:n}$ be neighboring inputs where only the first element differs (i.e., $x_1 \neq x'_1$), and $\mathcal{R}$ has blanket mass γ. Then, $\mathcal{S} \circ \mathcal{R}^n$ satisfies $(\varepsilon, \delta(x_{1:n}, x'_{1:n}))$-DP for $x_{1:n}$ and $x'_{1:n}$, where*

$$\delta(x_{1:n}, x'_{1:n}) \leq \mathcal{D}^{\mathrm{blanket}}_{\mathrm{e}^\varepsilon, n, \mathcal{R}_{\mathrm{BG}}, \gamma}(\mathcal{R}_{x_1} \| \mathcal{R}_{x'_1}) := \frac{1}{n\gamma}\mathbb{E}_{\substack{Y_{1:M} \overset{\text{i.i.d.}}{\sim} \mathcal{R}_{\mathrm{BG}} \\ M \sim \mathrm{Bin}(n,\gamma)}}\left[\left(\sum_{i=1}^{M} l_\varepsilon(Y_i)\right)_+\right],$$

where $(\cdot)_+$ denotes $\max\{0, \cdot\}$.

[2]In contrast to the multi-message shuffling model, which typically requires access to a stronger ideal functionality $\mathcal{S}$ capable of splitting or decorrelating multiple messages from the same user (for example by hiding timing and grouping information), the single-message shuffling model only relies on a comparatively simple "anonymization" ideal functionality that permutes user identities. This makes the single-message model significantly easier to implement in practice, and therefore in this work we restrict attention to the single-message shuffling setting.

Here, Bin *is a binomial distribution and* $l_\varepsilon(y) := \frac{\mathcal{R}_{x_1}(y) - e^\varepsilon \mathcal{R}_{x_1'}(y)}{\mathcal{R}_{\mathrm{BG}}(y)}$ *is called the privacy amplification random variable when* $y = Y \sim \mathcal{R}_{\mathrm{BG}}$.

We term this upper bound *blanket divergence* $\mathcal{D}^{\mathrm{blanket}}_{\mathrm{e}^\varepsilon, n, \mathcal{R}_{\mathrm{BG}}, \gamma}$. We omit the parameters and simply write $\mathcal{D}^{\mathrm{blanket}}$ when they are clear from the context. Moreover, for the same mechanism, the following lower bound, also expressed in terms of the blanket divergence, is known.

Theorem 2.6 (Theorem 14 of Su et al. [25])**.** *For any* $x, a_1, a_1' \in \mathcal{X}, \varepsilon, n$, *and local randomizer* $\mathcal{R}$, *it holds that:*

$$\delta_{\mathcal{S}\circ\mathcal{R}^n}(\varepsilon) \geq \mathcal{D}^{\mathit{blanket}}_{\mathrm{e}^\varepsilon, n, \mathcal{R}_x, 1}(\mathcal{R}_{a_1} \| \mathcal{R}_{a_1'}).$$

Here, to provide a unified treatment of the blanket divergences corresponding to the upper and lower bounds, we generalized the definition of the privacy amplification random variable as follows.

Definition 2.7 (Privacy amplification random variable). Fix the pair $x_1 \neq x_1' \in \mathcal{X}$. Let $\mathcal{R}_{\mathrm{ref}} \in \{\mathcal{R}_{\mathrm{BG}}\} \cup \{\mathcal{R}_x : x \in \mathcal{X}\}$. We define the *privacy amplification random variable* with respect to $\mathcal{R}_{\mathrm{ref}}$ by

$$l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}}) := \frac{\mathcal{R}_{x_1}(Y) - e^\varepsilon \mathcal{R}_{x_1'}(Y)}{\mathcal{R}_{\mathrm{ref}}(Y)}.$$

Throughout the paper, whenever we write $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ without further comment, the random variable Y is understood to be drawn from the reference distribution $\mathcal{R}_{\mathrm{ref}}$. When the dependence on $(x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ is clear from context, we simply write $l_\varepsilon(Y)$ or l_ε. Two choices of the reference distribution $\mathcal{R}_{\mathrm{ref}}$ will be particularly important in this work:

- If $\mathcal{R}_{\mathrm{ref}} = \mathcal{R}_{\mathrm{BG}}$ is the blanket distribution, then $l_\varepsilon(Y)$ is the privacy amplification random variable underlying the blanket-divergence *upper bound* of Balle et al. [4].
- If $\mathcal{R}_{\mathrm{ref}} = \mathcal{R}_x$ for some $x \in \mathcal{X}$, then $l_\varepsilon(Y)$ is the privacy amplification random variable appearing in the *lower bound* of Su et al. [25].

2.2 Local Randomizers Beyond Pure Local DP

Most existing analyses of shuffling focus on local randomizers that satisfy *pure* local DP (i.e., $(\varepsilon_0, 0)$-DP at the user level). One of the main goals of this paper is to remove this restriction and develop a more general analysis that applies to arbitrary local randomizers satisfying the following mild regularity assumptions.

Assumption 2.8 (Regularity conditions for the local randomizer)**.** Consider the local randomizer $\mathcal{R} : \mathcal{X} \to \mathcal{Y}$. We assume that there exists $\rho_0 > 0$ such that, for every pair $x_1 \neq x_1' \in \mathcal{X}$ and every reference distribution $\mathcal{R}_{\mathrm{ref}} \in \{\mathcal{R}_{\mathrm{BG}}\} \cup \{\mathcal{R}_x : x \in \mathcal{X}\}$, the following conditions hold uniformly over all $|\varepsilon| \leq \rho_0$.

(1) **Uniform moment bounds.** For every integer $k \geq 1$, $\mathbb{E}_{Y\sim\mathcal{R}_{\mathrm{ref}}}\left[\, |l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})|^k \,\right] < \infty$.

(2) **Non-degenerate variance.** The variance $\sigma^2 := \mathrm{Var}_{Y\sim\mathcal{R}_{\mathrm{ref}}}\big(l_0(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})\big)$ is strictly positive, and the following conditions hold: $\int \frac{\mathcal{R}_{x_1}(y)^2}{\mathcal{R}_{\mathrm{ref}}(y)}\, dy < \infty$ and $\int \frac{\mathcal{R}_{x_1'}(y)^2}{\mathcal{R}_{\mathrm{ref}}(y)}\, dy < \infty$.

(3) **Structural condition.** Either of the following holds.

(Cont) $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ has a nontrivial absolutely continuous component with respect to Lebesgue measure whose support contains a fixed nondegenerate bounded interval $I \subset \mathbb{R}$ independent of ε.

(Bound) There exists a constant $C < \infty$ such that $|l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})| \leq C$ almost surely.

We denote by $\mathfrak{R}$ the class of local randomizers $\mathcal{R}$ that satisfy Assumption 2.8.

Assumption 2.8 is mild: any nontrivial local randomizer that satisfies pure local DP automatically satisfies these conditions, and even when pure local DP does not hold, if the privacy amplification random variable $l_0(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ is non-degenerate (i.e., not almost surely constant), then one can enforce the bounded case (Bound) in (3) by truncating the output space to a large but bounded interval (i.e., by slightly modifying the local randomizer).

Accordingly, we formulate our results in a mechanism-agnostic way and do not commit to a specific local randomizer. Nevertheless, it is useful to keep in mind a concrete running example satisfying Assumption 2.8 without pure local DP and truncation. In particular, we will often refer to the *generalized Gaussian mechanism*, which forms a flexible family of location mechanisms (parameterized by the shape β) and simultaneously encompasses the Laplace mechanism ($\beta = 1$) and the Gaussian mechanism ($\beta = 2$). The latter case is especially noteworthy, as the Gaussian mechanism has been explicitly highlighted in prior work (e.g., Balle et al. [4] and Koskela et al. [17]) as technically challenging to analyze in the shuffle model.

Definition 2.9 (Generalized Gaussian local randomizer (β-Gaussian mechanism)). Let $\mathcal{X} := [0, 1]$ be the input domain and fix parameters $c > 0$ and $\beta \in [1, 2]$. The *generalized Gaussian local randomizer* $\mathcal{R} : \mathcal{X} \to \mathcal{Y}$ is defined by $\mathcal{Y} := \mathbb{R}, \mathcal{R}(x) = Y$, where $Y \sim \phi_x^\beta$ and ϕ_x^β denotes the generalized Gaussian density

$$\phi_x^\beta(y) := \frac{\beta}{2c\,\Gamma(1/\beta)} \exp\left(-\left|\frac{y-x}{c}\right|^\beta\right), \qquad y \in \mathbb{R}.$$

In other words, for each input $x \in \mathcal{X}$, the output $\mathcal{R}(x)$ is obtained by adding generalized Gaussian noise (with parameters c, β) to the input x. The variance of this distribution is given by $\mathrm{Var}(Y) = c^2 \frac{\Gamma(3/\beta)}{\Gamma(1/\beta)}$; by choosing c appropriately, we can calibrate $\mathrm{Var}(Y)$. We call this the β-Gaussian mechanism.

3 Analysis of Shuffling via the Central Limit Theorem (CLT)

Here, we first review the blanket divergence and its interpretation. Then, we analyze the blanket divergence to obtain an asymptotic expansion. Next, we analyze the privacy profile of the shuffled mechanism based on this expansion. Then, we interpret and explore the shuffle indices of various local randomizers. Finally, we derive a necessary and sufficient condition under which the blanket divergence yields an asymptotically optimal characterization.

3.1 Review of Blanket Divergence

Given a shuffled mechanism $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$, the direct computation of the privacy profile requires the two output distributions of $\mathcal{M}(x_{1:n})$ and $\mathcal{M}(x'_{1:n})$ where $x_{1:n}$ and $x'_{1:n}$ are the worst-case neighboring datasets. Given a dataset $x_{1:n}$, the output distribution of the shuffled mechanism $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$ can be written as the permutation-mixture

$$\Pr[\mathcal{M}(x_{1:n}) = y_{1:n}] = \frac{1}{n!} \sum_{\pi \in S_n} \prod_{i=1}^{n} \mathcal{R}_{x_i}\big(y_{\pi(i)}\big), \qquad y_{1:n} \in \mathcal{Y}^n.$$

That is, evaluating the shuffled density naively involves a mixture with $n!$ components. Moreover, we need to apply a nonlinear transformation to this and integrate the result, together with a maximization over all neighboring datasets to identify the worst case. Altogether, this direct computation quickly becomes intractable. We therefore turn to the privacy blanket approach proposed by Balle et al. [4].

Here, we briefly recall it. The local randomizer admits the mixture decomposition $\mathcal{R}_x = \gamma\,\mathcal{R}_{\mathrm{BG}} + (1-\gamma)\,\mathcal{R}_x^{\mathrm{LO}}$, where $\mathcal{R}_x^{\mathrm{LO}}$ is the corresponding leftover distribution. Based on this decomposition,

we consider the extended mechanism $\widetilde{\mathcal{M}}$ depicted in Fig. 1. We can interpret a blanket divergence as an upper bound on the hockey-stick divergence of $\widetilde{\mathcal{M}}$.

In $\widetilde{\mathcal{M}}$, each user (id $>$ 1) draws a message from $\mathcal{R}_{\mathrm{BG}}$ with probability γ, and from $\mathcal{R}_{x_{\mathrm{id}}}^{\mathrm{LO}}$ with probability $1-\gamma$. Messages sampled from the leftover part are revealed together with their user indices, whereas messages sampled from the blanket part are released only after being shuffled. User 1 outputs $\mathcal{R}_{x_1}$ (resp. $\mathcal{R}_{x_1'}$) and Y_1 is always placed in the shuffled blanket multiset. The hockey-stick divergence of $\widetilde{\mathcal{M}}$ is upper-bounded by applying its joint convexity to the mixture over the number of blanket messages, which yields the blanket divergence (see Section A for details). Importantly, the shuffled mechanism $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$ can be obtained by shuffling all outputs of $\widetilde{\mathcal{M}}$. Therefore, by the data processing inequality, the hockey-stick divergence of $\mathcal{M}$ is also upper-bounded by the blanket divergence.

Next, in Fig. 1, we consider a fixed $x_2 \in \mathcal{X}$ and restrict attention to the neighboring-dataset instance $(x_1, x_2, \ldots, x_2) \simeq (x_1', x_2, \ldots, x_2)$. In this setting, if we take $\gamma = 1$ and choose the reference distribution as $\mathcal{R}_{\mathrm{BG}} = \mathcal{R}_{x_2}$, then the resulting construction of $\widetilde{\mathcal{M}}$ coincides in distribution with the shuffled mechanism output for these two datasets (i.e., $\mathcal{M}((x_1, x_2, \ldots, x_2))$ and $\mathcal{M}((x_1', x_2, \ldots, x_2))$). The inequalities introduced earlier, namely the convexity step in γ and the data processing inequality, do not appear here, so that the blanket divergence matches the hockey-stick divergence of $\widetilde{\mathcal{M}}$, and thus that of the shuffled mechanism $\mathcal{M}$. Since this neighboring instance is feasible in the supremum that defines the privacy profile, it follows that this blanket divergence provides a lower bound.

Remark 3.1 (Relationship to the clone paradigm)**.** The blanket divergence viewpoint with $\gamma = 1$ is closely related to the clone paradigm [14, 15]. The clone paradigm reduces shuffle privacy analysis to an auxiliary mechanism; shuffling of n categorical distributions $(\mathcal{R}_b^{\mathrm{Cat}}, \mathcal{R}_{\mathrm{ref}}^{\mathrm{Cat}}, \ldots, \mathcal{R}_{\mathrm{ref}}^{\mathrm{Cat}})$, where $b \in \{0, 1\}$ changes the target user. In this sense, this reduced mechanism can be interpreted as a particular $\gamma = 1$ instance of $\widetilde{\mathcal{M}}$, with an appropriate choice of $\mathcal{R}_b^{\mathrm{Cat}}$ and $\mathcal{R}_{\mathrm{ref}}^{\mathrm{Cat}}$. Note that Su et al. [25] showed that the blanket-based reduction can be at least as tight as the clone reductions. Motivated by this and by the mechanism-aware nature of the blanket divergence, we focus primarily on the blanket-divergence upper bound of Balle et al. [4] in what follows.

3.2 Analysis of Blanket Divergence

As shown in Lemma 2.5, the blanket divergence no longer involves an explicit permutation. This is because, conditioned on the number of blanket messages $M = m$ in $\widetilde{\mathcal{M}}$, the $m-1$ blanket outputs are i.i.d. draws from the blanket distribution $\mathcal{R}_{\mathrm{BG}}$; moreover, thanks to the symmetry attained by shuffling, the analysis of the hockey-stick divergence reduces to that of a sum of m i.i.d. random variables. Exploiting this tractability via the CLT, we can derive the following asymptotic expansion.

Lemma 3.2 (Asymptotic expansion of the blanket divergence). *Given $\mathcal{R} \in \mathfrak{R}$, fix a pair $x_1 \neq x_1' \in \mathcal{X}$ and let $\mathcal{R}_{\mathrm{ref}} \in \{\mathcal{R}_{\mathrm{BG}}\} \cup \{\mathcal{R}_x : x \in \mathcal{X}\}$. For $\gamma \in (0, 1]$, define $\sigma^2 := \mathrm{Var}_{Y \sim \mathcal{R}_{\mathrm{ref}}}\big(l_{\varepsilon=0}(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})\big)$ and $\chi := \sqrt{\gamma}/\sigma$. Assume that as $n \to \infty$, $\varepsilon_n = \omega(\sqrt{1/n})$ and $\varepsilon_n = O\Big(\sqrt{\log n/n}\Big)$. Then, as $n \to \infty$,*

$$\mathcal{D}^{\mathrm{blanket}}_{e^{\varepsilon_n}, n, \mathcal{R}_{\mathrm{ref}}, \gamma}(\mathcal{R}_{x_1} \| \mathcal{R}_{x_1'}) \;=\; \varphi\big(\chi\varepsilon_n\sqrt{n}\big)\left(\frac{1}{\chi^3}\,\frac{1}{\varepsilon_n^2 n^{3/2}}\right)\big(1 + o(1)\big). \tag{1}$$

Here, we present an informal derivation based on the CLT to build intuition of how the result arises. The rigorous proof with the asymptotic expansion of CLT [1, 21, 23] is in the full version [27].

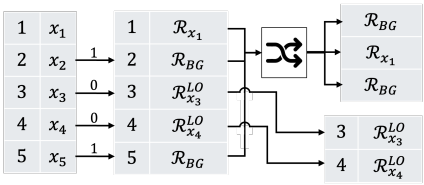

Fig. 1. $\widetilde{\mathcal{M}}$, which yields the blanket divergence.

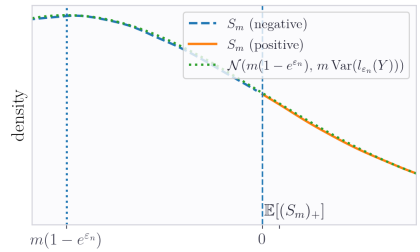

Fig. 2. The intuition behind the CLT of the blanket divergence with $m = 100$.

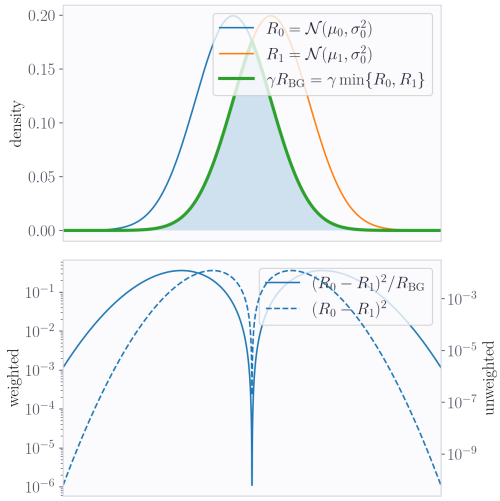

Fig. 3. The intuition of σ with the Gaussian mechanism where $\sigma_0 = 2$. x axis represents the output of the local randomizer and is common for both images.

Recall that the blanket divergence can be written as

$$\mathcal{D}^{\text{blanket}} = \frac{1}{n\gamma}\,\mathbb{E}_M\left[\mathbb{E}\left[\left(\sum_{i=1}^{M} l_{\varepsilon_n}(Y_i)\right)_+ \middle| M\right]\right], \qquad M \sim \mathrm{Bin}(n,\gamma), \quad Y_i \overset{\text{i.i.d.}}{\sim} \mathcal{R}_{\text{ref}}.$$

Let us condition on $M = m$ and define $S_m := \sum_{i=1}^{m} l_{\varepsilon_n}(Y_i)$. The inner expectation becomes $\mathbb{E}[(S_m)_+]$, whose behavior is shown in Fig. 2. Since S_m is a sum of i.i.d. random variables, the CLT suggests

$$S_m \approx N_m := \mathcal{N}(\mu_m, s_m^2), \qquad \mu_m := m\,\mathbb{E}[l_{\varepsilon_n}(Y)] = m(1-e^{\varepsilon_n}), \quad s_m^2 := m\,\mathrm{Var}(l_{\varepsilon_n}(Y)),$$

for large m. For a Gaussian $N_m \sim \mathcal{N}(\mu_m, s_m^2)$ one has the identity $\mathbb{E}[(N_m)_+] = s_m\,\varphi\left(\frac{\mu_m}{s_m}\right) + \mu_m\,\Phi\left(\frac{\mu_m}{s_m}\right)$. Using Mills' ratio (in the regime $|\mu_m/s_m| \to \infty$), we obtain

$$\mathbb{E}[(N_m)_+] \;\approx\; \frac{s_m^3}{\mu_m^2}\,\varphi\left(\frac{\mu_m}{s_m}\right) \;=\; \frac{\mathrm{Var}(l_{\varepsilon_n}(Y))^{3/2}}{(e^{\varepsilon_n}-1)^2}\,\frac{1}{\sqrt{m}}\,\varphi\left(\frac{(e^{\varepsilon_n}-1)\sqrt{m}}{\sqrt{\mathrm{Var}(l_{\varepsilon_n}(Y))}}\right),$$

$M \sim \mathrm{Bin}(n,\gamma)$ concentrates around its mean $m = n\gamma$, so we plug in $m = n\gamma$ in the leading term. Moreover, for small ε_n, $\mathrm{Var}(l_{\varepsilon_n}(Y)) \approx \mathrm{Var}(l_0(Y)) =: \sigma^2$. With $\chi := \sqrt{\gamma}/\sigma$, we obtain

$$\mathbb{E}_M[(N_M)_+] \approx \frac{\sigma^3}{(e^{\varepsilon_n}-1)^2}\,\frac{1}{\sqrt{n\gamma}}\,\varphi\big(\chi(e^{\varepsilon_n}-1)\sqrt{n}\big).$$

Finally, recalling the prefactor $1/(n\gamma)$, and using $\sigma^3/\gamma^{3/2} = 1/\chi^3$, we arrive at

$$\mathcal{D}^{\text{blanket}} \approx \frac{1}{n\gamma}\,\mathbb{E}_M[(N_M)_+] \approx \varphi\big(\chi(e^{\varepsilon_n}-1)\sqrt{n}\big)\left(\frac{1}{\chi^3}\,\frac{1}{(e^{\varepsilon_n}-1)^2\,n^{3/2}}\right),$$

which matches the leading term in (1) using $e^{\varepsilon_n} \approx 1 + \varepsilon_n$. The rigorous proof in the full version [27] shows that this normal approximation is indeed valid under Assumption 2.8 (i.e., $\mathcal{R} \in \mathfrak{R}$) in the regime $\varepsilon_n = \omega(\sqrt{1/n})$ and $\varepsilon_n = O(\sqrt{\log n/n})$.

Our key conceptual finding from Lemma 3.2 is that the leading term depends on the local randomizer only through the single parameter $\chi = \sqrt{\gamma}/\sigma$. That is, at the level of the leading asymptotics, the dependence of the blanket divergence on the local randomizer $\mathcal{R}$ is fully encoded by the one-dimensional index χ. Furthermore, the leading term is monotone in χ: a larger χ yields a smaller leading term in the blanket divergence, and hence corresponds to stronger privacy amplification by shuffling. This motivates interpreting χ as a measure of the shuffle efficiency of the local randomizer. Owing to its central role, we refer to χ as the *shuffle index*.

3.3 Main Result

Up to this point, we have analyzed the blanket divergence for a fixed pair $(x_1, x_1', \mathcal{R}_{\mathrm{ref}})$. In contrast, $(\varepsilon_n, \delta_n)$-DP requires a guarantee that holds uniformly over all neighboring datasets $x_{1:n} \simeq x'_{1:n}$ (i.e., upper bounding the privacy profile $\delta_n \geq \delta_{\mathcal{M}}(\varepsilon_n)$). This section aims to relate the result of the blanket divergence to the privacy profile of the shuffled mechanism.

Recall from Section 3.1 that the privacy profile can be sandwiched between two blanket divergences. By combining this property with the monotonicity of the leading term with respect to the shuffle index χ, evaluating the privacy profile reduces to identifying the worst-case value of χ over all pairs (x_1, x_1'). To derive the upper bound, under the reference distribution $\mathcal{R}_{\mathrm{ref}} = \mathcal{R}_{\mathrm{BG}}$, we must choose the pair (x_1, x_1') that minimizes the shuffle index. Conversely, to establish the lower bound, under $\mathcal{R}_{\mathrm{ref}} = \mathcal{R}_x$ for some $x \in \mathcal{X}$ (with $\gamma = 1$), we are free to select any pair (x_1, x_1'); naturally, to maximize this lower bound, we should again choose the pair that minimizes the shuffle index. Formally, we use the following specific shuffle indices for the upper and lower bounds, respectively.

Definition 3.3. Let $\gamma \in (0, 1]$ be the blanket mass of $\mathcal{R} \in \mathfrak{R}$. The *lower shuffle index* is defined by $\chi_{\mathrm{lo}}(\mathcal{R}) := \sqrt{\gamma}/\sup_{x_1 \simeq x_1' \in \mathcal{X}} \sqrt{\mathrm{Var}_{Y \sim \mathcal{R}_{\mathrm{BG}}}\left[l_0\left(Y; x_1, x_1', \mathcal{R}_{\mathrm{BG}}\right)\right]}$. The *upper shuffle index* is defined by $\chi_{\mathrm{up}}(\mathcal{R}) := 1/\sup_{x_1 \simeq x_1' \in \mathcal{X}} \sup_{x \in \mathcal{X}} \sqrt{\mathrm{Var}_{Y \sim \mathcal{R}_x}\left[l_0\left(Y; x_1, x_1', \mathcal{R}_x\right)\right]}$.

With these specific shuffle indices in hand, we can formally establish the upper and lower bounds for the privacy profile of the shuffled mechanism, as stated in the following corollary.

Corollary 3.4 (Privacy profile bounds). *Let $\mathcal{R} \in \mathfrak{R}$. Assume that as $n \to \infty$, $\varepsilon_n = \omega(\sqrt{1/n})$ and $\varepsilon_n = O\left(\sqrt{\log n/n}\right)$. Then, there exist sequences $e_n^{\mathrm{up}}, e_n^{\mathrm{lo}} \to 0$ as $n \to \infty$ such that*

$$f_{n,\varepsilon_n}(\chi_{\mathrm{up}}(\mathcal{R}))\left(1 + e_n^{\mathrm{up}}\right) \leq \delta_{\mathcal{S} \circ \mathcal{R}^n}(\varepsilon_n) \leq f_{n,\varepsilon_n}(\chi_{\mathrm{lo}}(\mathcal{R}))\left(1 + e_n^{\mathrm{lo}}\right),$$

for all sufficiently large n where $f_{n,\varepsilon_n}(\chi) := \varphi\left(\chi\, \varepsilon_n \sqrt{n}\right)\left(\frac{1}{\chi^3}\, \frac{1}{\varepsilon_n^2 n^{3/2}}\right)$.

This corollary is a direct consequence of applying the asymptotic expansion established in Lemma 3.2 to the lower and upper bounds (Lemma 2.5 and Theorem 2.6).

In DP, it is common to set δ_n to be α/n for some fixed constant $\alpha > 0$. Our next result characterizes the corresponding ε_n.

Theorem 3.5 (Shuffling delta band). *Let $\mathcal{R} \in \mathfrak{R}$ with blanket mass γ. Fix $\alpha > 0$. For $n \geq 2$ and $\chi > 0$, set*[3]

$$\varepsilon_n(\alpha, \chi) := \log\left(1 + \sqrt{\frac{2}{\chi^2 n} W\left(\frac{\sqrt{n}}{2\alpha\chi\sqrt{2\pi}}\right)}\right) \tag{2}$$

Then, there exists a sequence $(\varepsilon_n^* \in [\varepsilon_n(\alpha, \chi_{\mathrm{up}}(\mathcal{R})), \varepsilon_n(\alpha, \chi_{\mathrm{lo}}(\mathcal{R}))])_{n\geq 1}$ such that $\delta_{\mathcal{S}\circ\mathcal{R}^n}(\varepsilon_n^*) = \frac{\alpha}{n}(1 + o(1))$ for all sufficiently large n.

This is obtained by inverting the relationship between ε_n and δ_n from the bounds established in Corollary 3.4. Specifically, by equating the leading term of the privacy profile $f_{n,\varepsilon_n}(\chi)$ to the target $\delta_n = \alpha/n$, we can solve for ε_n to derive the explicit formula $\varepsilon_n(\alpha, \chi)$ for both the upper and lower bounds. Informally, Theorem 3.5 shows that for any fixed target constant $\alpha > 0$ we construct a sequence $(\varepsilon_n^*)_{n\geq 2}$ whose values lie between two explicit curves $\varepsilon_n(\alpha, \chi_{\mathrm{up}}(\mathcal{R}))$ and $\varepsilon_n(\alpha, \chi_{\mathrm{lo}}(\mathcal{R}))$, and for which $\delta_n = \frac{\alpha}{n}(1 + o(1))$.[4] In other words, the band precisely captures ε_n^* that realizes $\delta_n \approx \alpha/n$.

Another important implication of Theorem 3.5 is that the asymptotic expansion of (2) $= \frac{1}{\chi}\sqrt{\frac{\log n}{n}}\,(1 + o(1))$ as $n \to \infty$ reveals a simple relationship. The amplified privacy parameter ε_n is inversely proportional to the shuffle index χ. Consequently, the width of the band is asymptotically captured by the ratio

$$\frac{\varepsilon_n(\alpha, \chi_{\mathrm{lo}}(\mathcal{R}))}{\varepsilon_n(\alpha, \chi_{\mathrm{up}}(\mathcal{R}))} = \frac{\chi_{\mathrm{lo}}(\mathcal{R})}{\chi_{\mathrm{up}}(\mathcal{R})}\,(1 + o(1))\,.$$

Hence, if $\chi_{\mathrm{lo}}(\mathcal{R})/\chi_{\mathrm{up}}(\mathcal{R})$ is close to 1, the band is narrow, indicating that the shuffle index provides a tight characterization of the amplified privacy parameter.

3.4 Shuffle Index

In this subsection, we explore the shuffle index χ: its interpretation, examples, and the analysis of the tightness of the blanket-divergence bounds through the shuffle index.

Interpretation of Shuffle Index. Technically, by the CLT as shown in Section 3.2, the leading asymptotics of the blanket divergence with $(x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ can be summarized by the two parameters γ and σ. Here we defined $\sigma^2 := \mathrm{Var}_{Y\sim\mathcal{R}_{\mathrm{ref}}}\left(\frac{\mathcal{R}_{x_1}(Y) - \mathcal{R}_{x_1'}(Y)}{\mathcal{R}_{\mathrm{ref}}(Y)}\right) = \int_{\mathcal{Y}} \frac{\left(\mathcal{R}_{x_1}(y) - \mathcal{R}_{x_1'}(y)\right)^2}{\mathcal{R}_{\mathrm{ref}}(y)}\,dy$ so that σ is a χ^2-type distance between $\mathcal{R}_{x_1}$ and $\mathcal{R}_{x_1'}$ with respect to the reference distribution $\mathcal{R}_{\mathrm{ref}}$. This differs from the usual χ^2 divergence, whose reference distribution is $\mathcal{R}_{x_1}$ or $\mathcal{R}_{x_1'}$; here the reference distribution is the third distribution $\mathcal{R}_{\mathrm{ref}}$. In Fig. 3, we take $\mathcal{R}$ to be Gaussian and plot the integrand of σ^2, namely $(\mathcal{R}_{x_1} - \mathcal{R}_{x_1'})^2/\mathcal{R}_{\mathrm{BG}}$. Compared to the ordinary squared distance $(\mathcal{R}_{x_1} - \mathcal{R}_{x_1'})^2$, we see that regions where the blanket is thin are emphasized. This distinction represents the shuffled setting: we can interpret σ as distinguishability of $\mathcal{R}_{x_1}$ and $\mathcal{R}_{x_1'}$ from the blanket $\mathcal{R}_{\mathrm{BG}}$ view.

To reiterate, the blanket divergence can be interpreted as the natural upper bound of the hockey-stick divergence between $\mathcal{S}(\mathcal{R}_{x_1}, \mathcal{R}_{\mathrm{ref}}, \ldots, \mathcal{R}_{\mathrm{ref}})$ and $\mathcal{S}(\mathcal{R}_{x_1'}, \mathcal{R}_{\mathrm{ref}}, \ldots, \mathcal{R}_{\mathrm{ref}})$. γ represents the fraction of messages that fall into the blanket, and hence the effective message size is $M \approx n\gamma$. On the other hand, since $1/\sigma$ is the reciprocal of the distinguishability measure in the above setting, it quantifies

[3]The explicit formula for $\varepsilon_n(\alpha, \chi)$ in Theorem 3.5 is obtained by solving the leading-order approximation (1) $= \alpha/n$ in closed form via the Lambert W-function. For our numerical experiments and practical use, however, we instead compute ε_n by numerically solving the more accurate equation based on the full asymptotic expression, using a simple bisection method.

[4]More generally, one can obtain any prescribed polynomial rate $\delta_n = \Theta(n^{-\beta})$ with $\beta > 1$ in the leading term. In this work we focus on the regime $\delta_n = \alpha/n(1 + o(1))$ because in the DP literature, setting δ_n to be α/n is a common practice.

Table 1. Shuffle index asymptotics of β-Gauss with $\mathrm{Var}(\mathcal{R}(x)) = \sigma_0^2$ where $c_\beta = \frac{\Gamma(1/\beta)}{\beta\sqrt{\Gamma(3/\beta)\Gamma(2-1/\beta)}}$ and k-RR.

β-Gauss	$\sigma_0 \to \infty$	$\sigma_0 \to 0$
χ_{lo}	$c_\beta\sigma_0\,(1+o\,(1))$	$\exp\!\left(-\Omega(\sigma_0^{-\beta})\right)$
$\chi_{\mathrm{lo}}/\chi_{\mathrm{up}}$	$1-O\left(\frac{1}{\sigma_0}\right)$	$\frac{1}{\sqrt{2}}+o(1)$

kRR | 2RR	$\varepsilon_0 \to 0$	$\varepsilon_0 \to \infty$
χ_{lo}	$\sqrt{\frac{k}{2}}\frac{1}{\varepsilon_0}\,(1+O(\varepsilon_0))$	$\frac{e^{-\frac{\varepsilon_0}{2}}}{\sqrt{2}}\,(1+O(e^{-\varepsilon_0}))$
$\chi_{\mathrm{lo}}/\chi_{\mathrm{up}}$	$1 \mid 1-O(\varepsilon_0)$	$1 \mid \frac{1}{\sqrt{2}}+O(e^{-\varepsilon_0})$

the per-message indistinguishability. Because the contribution of γ to indistinguishability appears through the $\sqrt{n}$ scaling of fluctuations, the index aggregates these effects as $\sqrt{\gamma}/\sigma$.

Examples of Shuffle Indices. We illustrate the shuffle index of two local randomizers: the β-Gaussian mechanism and k-RR. Concretely, we examine (i) the asymptotics of the lower shuffle index χ_{lo}, and (ii) the ratio $\chi_{\mathrm{lo}}/\chi_{\mathrm{up}}$ as a proxy for the tightness of the blanket-divergence sandwich. Closed-form derivations and some plots of shuffle indices are deferred to Appendix B. The resulting asymptotics are summarized in Table 1: for β-Gaussian we parametrize by the local noise scale σ_0, while for k-RR we parametrize by the pure local-DP level ε_0, and in both cases we report limits as the parameter tends to 0 and ∞. These asymptotics precisely capture the fundamental behavior established in theoretical bounds [13] and explain recent empirical observations [8] that shuffling yields highly efficient privacy amplification in the high-privacy regime, while providing negligible amplification in the low-privacy regime.

β-Gaussian. For large noise ($\sigma_0 \to \infty$), the lower shuffle index grows linearly, $\chi_{\mathrm{lo}} \approx c_\beta\sigma_0$, where c_β depends only on the shape parameter β. Concretely, c_β increases from $1/\sqrt{2}$ at $\beta = 1$ (Laplace) up to 1 at $\beta = 2$ (Gaussian), and then decreases for $\beta > 2$; hence, in the high-noise regime, the Gaussian case $\beta = 2$ yields the largest shuffle index. Moreover, the tightness ratio converges to 1, so the blanket-divergence bounds essentially coincide asymptotically when σ_0 is sufficiently large. That is, in the mean estimation task, since σ_0 monotonically determines the accuracy of the estimate, $\beta = 2$ induces the best privacy-utility trade-off in the high-noise regime (see Section 5.3 for the exploration of this topic). In contrast, in the low-noise regime ($\sigma_0 \to 0$), χ_{lo} decays exponentially to 0, which reflects that shuffling provides little amplification when the local noise is too small.

Remark 3.6 (Dimensional dependence)**.** For the Gaussian case ($\beta = 2$), it is natural to extend the input domain to d dimensions. Concretely, consider the Gaussian local randomizer $\mathcal{R}(x) = x + \mathcal{N}(0, \sigma_0^2 I_d)$ with fixed σ_0 and inputs in an ℓ_2-ball of fixed radius. Then the blanket mass satisfies $\gamma = \exp(-\Theta(\sqrt{d}))$, i.e., it decays exponentially fast in $\sqrt{d}$. Intuitively, this decay is driven by the thin-shell phenomenon of high-dimensional Gaussians: a typical output exponentially concentrates on a thin shell around its mean x (i.e., $\|\mathcal{R}(x) - x\| \approx \sigma_0\sqrt{d}$). On this typical shell, the blanket distribution takes an infimum over all inputs in the ℓ_2-ball, so the blanket becomes exponentially thin in $\sqrt{d}$ (see the full version [27] for details). Moreover, while χ_{up} is not directly affected by γ, χ_{lo} is suppressed by the factor $\sqrt{\gamma}$ and thus shrinks exponentially with $\sqrt{d}$; consequently, the gap between the corresponding upper and lower bounds becomes large. Closing this gap appears to require techniques that go beyond the blanket divergence, which remains an open problem.

k-RR. In large noise ($\varepsilon_0 \to 0$), we have $\chi_{\mathrm{lo}} \approx \sqrt{\frac{k}{2}}\frac{1}{\varepsilon_0}$. Importantly, for $k \geq 3$ the tightness ratio satisfies $\chi_{\mathrm{lo}}/\chi_{\mathrm{up}} = 1$, i.e., the two indices coincide and the blanket-divergence sandwich is tight at leading order; for 2-RR, the ratio is also close to 1 as $\varepsilon_0 \to 0$. In the low-noise regime ($\varepsilon_0 \to \infty$), χ_{lo} again vanishes exponentially, and thus shuffling becomes ineffective.

Remark 3.7 (Comparison to the clone paradigm)**.** As noted in Remark 3.1, the clone paradigm can be understood in terms of a blanket divergence with $\gamma = 1$ and an appropriate choice of the categorical distributions. Hence, we can induce a shuffle index of the clone paradigm. For k-RR, the explicit reduction of Feldman et al. [15, Theorem 7.4] leads to a shuffle index that coincides with our χ_{lo} (see Section B.2.3 for details). While their work only demonstrates empirical closeness to known lower bounds, our framework provides a theoretical and asymptotic explanation for the closeness.

Tightness of Blanket-Divergence Analysis. We have seen that $\chi_{\mathrm{lo}}/\chi_{\mathrm{up}} = 1$ for k-RR with $k \geq 3$. Here, we explore the structural conditions under which this equality holds, which implies that the blanket-divergence bounds are asymptotically tight in the sense of Theorem 3.5. Formally, we have the following characterization.

Theorem 3.8. *Let $\mathcal{R} \in \mathfrak{R}$ with blanket mass γ and assume that the suprema in the definitions of $\chi_{\mathrm{lo}}(\mathcal{R})$ and $\chi_{\mathrm{up}}(\mathcal{R})$ are attained by some pairs. Then, $\chi_{\mathrm{up}}(\mathcal{R}) = \chi_{\mathrm{lo}}(\mathcal{R})$ if and only if there exists a pair which attains the suprema for the lower shuffle index $(x_1^*, x_1'^*)$ and an input $x^* \in \mathcal{X}$ such that*

$$\mathcal{R}_{x^*}(y) = \gamma \mathcal{R}_{\mathrm{BG}}(y) \;\; \text{for almost every } y \in \left\{ y \in \mathcal{Y} : \mathcal{R}_{x_1^*}(y) \neq \mathcal{R}_{x_1'^*}(y) \right\}.$$

In particular, when these equivalent conditions hold, the lower and upper curves in Theorem 3.5 coincide asymptotically, and there exists an asymptotic privacy curve $\varepsilon_n^*(\alpha)$ realizing $\delta_{S_\circ \mathcal{R}^n}(\varepsilon_n^*(\alpha)) = \frac{\alpha}{n}(1 + o(1))$ such that $\varepsilon_n^*(\alpha) = \varepsilon_n\big(\alpha, \chi_{\mathrm{lo}}(\mathcal{R})\big)\,(1 + o(1))$ as $n \to \infty$.

In particular, the theorem applies to the case of k-RR ($k \geq 3$), and hence to mechanisms that are built on top of k-RR or share the same blanket structure, such as local hashing [30] and (for sufficiently small caps) PrivUnit [2]. Thus, in many practically relevant shuffle protocols, the blanket-divergence analysis is asymptotically optimal in the sense of Theorem 3.8.

4 Analysis of Blanket Divergence via the Fast Fourier Transform (FFT)

The results in the previous section provide an asymptotic characterization of the blanket divergence. However, in applications one is ultimately interested in the finite-n behavior of the shuffled mechanism. To this end, we develop a complementary, non-asymptotic approach for computing the blanket divergence at finite n.

Our goal is to design an algorithm that, for a given local randomizer $\mathcal{R}$ and parameters (n, ε), produces rigorous upper and lower bounds on $\mathcal{D}^{\mathrm{blanket}}_{\mathrm{e}^\varepsilon, n, \mathcal{R}_{\mathrm{ref}}, \gamma}(\mathcal{R}_{x_1} \| \mathcal{R}_{x_1'})$ with explicitly controlled error and provable running time guarantees. The main idea is to exploit the same structure that underlies our CLT-based analysis: the blanket divergence can be expressed in terms of sums of i.i.d. privacy amplification random variables. This observation allows us to approximate the distribution of the sum by means of the Fast Fourier Transform (FFT). To obtain rigorous bounds rather than heuristic approximations, we carefully track the errors introduced by truncation of the summands, discretization on a grid, and FFT wrap-around. By combining these error bounds with the asymptotic expansion from Lemma 3.2, we can both tune the numerical parameters (window size, grid resolution, truncation level) and analyze the overall complexity of the method.

The starting point is another representation of the blanket divergence.

Lemma 4.1. *Let $M \sim 1 + \mathrm{Bin}(n-1, \gamma)$. Then,*

$$\mathcal{D}^{\mathrm{blanket}}_{\mathrm{e}^\varepsilon, n, \mathcal{R}_{\mathrm{ref}}, \gamma}(\mathcal{R}_{x_1} \| \mathcal{R}_{x_1'}) = \Pr_{\substack{Y_1 \sim \mathcal{R}_{x_1} \\ Y_{2:n} \overset{\text{i.i.d.}}{\sim} \mathcal{R}_{\mathrm{ref}}}} \left[\textstyle\sum_{i=1}^M l_\varepsilon(Y_i) > 0 \right] - \mathrm{e}^\varepsilon \Pr_{\substack{Y_1 \sim \mathcal{R}_{x_1'} \\ Y_{2:n} \overset{\text{i.i.d.}}{\sim} \mathcal{R}_{\mathrm{ref}}}} \left[\textstyle\sum_{i=1}^M l_\varepsilon(Y_i) > 0 \right].$$

Here, we work with a representation of the blanket divergence that does not involve explicit expectations or integrals, which is particularly convenient for deriving rigorous bounds on the

truncation error. In particular, in our FFT-based computation we must truncate the support of the random variable in order to approximate an unbounded distribution on a finite grid. This issue is relevant for non-pure local DP (see, e.g., the discussion of the Gaussian mechanism in Balle et al. [4]); hence this formulation is crucial for our purposes.

Su et al. [25] also propose an FFT-based accountant for the blanket divergence, but their running time scales as $O(n^2)$ in the number of users when targeting a constant additive error. In contrast, by adopting the concentration-inequality-based error control of Gopi et al. [16] with our results in Section 3, we can tune the truncation and discretization parameters so that our algorithm runs in near-linear time $\widetilde{O}(n)$ while achieving a prescribed relative error.

4.1 FFT Algorithm

Algorithms 1 and 2 together form our FFT-based procedure. As we show in Theorem 4.2 below, the output of these algorithms yields rigorous upper and lower bounds on the blanket divergence.

Algorithm 1 CalculatePMF

Input:

n, γ: Parameters for the number of summands.
$F_{\mathrm{ref}}(t)$: The CDF of the summands $l(Y_i)$ for $i \geq 2$.
$[s, s + w^{\mathrm{in}}]$: Truncation interval for $l(Y_i)$, $i \geq 2$.
h: Bin width for discretization.
w^{out}: The window size of FFT.

▷ **Step 1: Create PMF for a single truncated and discretized summand**

$p_{\mathrm{out}} \leftarrow F_{\mathrm{ref}}(s) + (1 - F_{\mathrm{ref}}(s + w^{\mathrm{in}}))$.
$p_{\mathrm{in}} \leftarrow 1 - p_{\mathrm{out}}$.
Let Z^{tr} be a random variable for a single summand conditioned on being in $[s, s + w^{\mathrm{in}}]$.
The CDF of Z^{tr} is $F_{Z^{\mathrm{tr}}}(t) = \left(F_{\mathrm{ref}}(t) - F_{\mathrm{ref}}(s)\right)/p_{\mathrm{in}}$.
$\mu_{Z^{\mathrm{tr}}} \leftarrow \left((s + w^{\mathrm{in}}) F_{\mathrm{ref}}(s + w^{\mathrm{in}}) - s F_{\mathrm{ref}}(s) - \int_s^{s+w^{\mathrm{in}}} F_{\mathrm{ref}}(z)\, dz \right) / \left(F_{\mathrm{ref}}(s + w^{\mathrm{in}}) - F_{\mathrm{ref}}(s) \right)$.
Define a discrete grid $\mathcal{G} = \{x_j = j \cdot h \mid j \in \mathbb{Z}\}$.
For each grid point x_j in the truncated range, compute the PMF of the discretized variable Z^{di}: $\mathbf{p}_{Z^{\mathrm{di}}}[j] \leftarrow F_{Z^{\mathrm{tr}}}(x_j + h/2) - F_{Z^{\mathrm{tr}}}(x_j - h/2)$.
$\mu_{Z^{\mathrm{di}}} \leftarrow \sum_j x_j \mathbf{p}_{Z^{\mathrm{di}}}[j]$

▷ **Step 2: Calculate the characteristic function of the sum** S

$\mu_{S^{\mathrm{tr}}} \leftarrow (n-1)\,\gamma\,\mu_Z^{\mathrm{tr}}$, $\mu_{S^{\mathrm{di}}} \leftarrow (n-1)\,\gamma\,\mu_Z^{\mathrm{di}}$
Set $N \leftarrow \mathrm{Round}(w^{\mathrm{out}}/h)$, construct zero-padded $\mathbf{p}_{\mathrm{pad}}$ of length N from $\mathbf{p}_{Z^{\mathrm{di}}}$, and let $\omega_k = \frac{2\pi k}{Nh}$.
$\psi_{Z^{\mathrm{di}}} \leftarrow \mathtt{FFT}\left(\mathtt{FFTSHIFT}(\mathbf{p}_{\mathrm{pad}})\right)$
$\psi_S \leftarrow \left((1-\gamma)\,\mathbf{1} + \gamma\,\psi_{Z^{\mathrm{di}}}\right)^{\odot(n-1)}$, $\psi \leftarrow \left[e^{-i\,\omega_k\,\mu_{S^{\mathrm{di}}}}\right]_{k=0}^{N-1}$ where $\odot$ denotes element-wise operations.
$\mathbf{p}_{S-\mu_{S^{\mathrm{di}}}} \leftarrow \mathrm{Re}\left(\mathtt{IFFTSHIFT}\left(\mathtt{IFFT}(\psi \odot \psi_S)\right)\right)$ ▷ PMF of the mean-centered sum $S - \mu_{S^{\mathrm{di}}}$. $\mathrm{Re}(\cdot)$ denotes the real part.
return $\mathbf{p}_{S-\mu_{S^{\mathrm{di}}}}, \mu_{S^{\mathrm{tr}}}$

Algorithm 2 CalculateMainTerm

Input:

$\mathbf{p}_{S-\mu_{S^{\mathrm{di}}}}, \mu_{S^{\mathrm{tr}}}$: The PMF and the mean of the truncated sum.
$F(t)$: The CDF of $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ under $\mathcal{R}_{x_1}$ or $\mathcal{R}_{x_1'}$.
c: The constant for the discretization error bound.
$p_{\mathrm{main}} \leftarrow 0$.
for each grid point z_j with corresponding PMF value $\mathbf{p}_{S-\mu_{S^{\mathrm{di}}}}[j]$ **do**
    $p_{\mathrm{main}} \leftarrow p_{\mathrm{main}} + \mathbf{p}_{S-\mu_{S^{\mathrm{di}}}}[j] \cdot \big(1 - F(-c - z_j - \mu_{S^{\mathrm{tr}}})\big)$.
return p_{main}

Theorem 4.2. [5] *Fix a distinct pair $x_1, x_1' \in \mathcal{X}$ and let $\mathcal{R}$ be a local randomizer. Let the reference distribution be $\mathcal{R}_{\mathrm{ref}} \in \{\mathcal{R}_{\mathrm{BG}}\} \cup \{\mathcal{R}_x : x \in \mathcal{X}\}$. For $Y \sim \mathcal{R}_{\mathrm{ref}}$, denote by F_{ref} the CDF of the privacy amplification random variable $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$, and for $j \in \{x_1, x_1'\}$ denote by F_j the CDF of $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ under $Y \sim \mathcal{R}_j$. Let $(B_i)_{i=2}^n$ be i.i.d.* Bernoulli(γ) *and $(Y_i)_{i=2}^n$ be i.i.d. with $\mathcal{R}_{\mathrm{ref}}$, independent of $(B_i)_{i=2}^n$.*

Given parameters $(n, \gamma, F_{\mathrm{ref}}, s, w^{\mathrm{in}}, h, w^{\mathrm{out}})$, we define the truncation event $\perp := \Big\{\exists i \in \{2, \dots, n\} : B_i = 1,\ l_\varepsilon\big(Y_i; x_1, x_1', \mathcal{R}_{\mathrm{ref}}\big) \notin [s, s+w^{\mathrm{in}}]\Big\}$ and let $\big(\mathbf{p}_{S-\mu_S^{\mathrm{di}}}, \mu_S^{\mathrm{tr}}\big) := \mathit{CalculatePMF}\big(n, \gamma, F_{\mathrm{ref}}, s, w^{\mathrm{in}}, h, w^{\mathrm{out}}\big)$. For $x \in \{x_1, x_1'\}$ and $c \geq 0$, define $P(c; x) := \big(1 - \Pr[\perp]\big)\ \mathit{CalculateMainTerm}\big(\mathbf{p}_{S-\mu_S^{\mathrm{di}}}, \mu_S^{\mathrm{tr}}, F_x, c\big)$.

Then, for any $c \geq 0$,

$$P(-c; x_1) - \mathrm{e}^\varepsilon P(c; x_1') - \delta_{\mathrm{err}}^{\mathrm{low}}(c, h, s, w^{\mathrm{in}}, w^{\mathrm{out}}) \leq D_{\mathrm{e}^\varepsilon, n, \mathcal{R}_{\mathrm{ref}}, \gamma}^{\mathrm{blanket}}(\mathcal{R}_{x_1} \| \mathcal{R}_{x_1'}) \leq P(c; x_1) - \mathrm{e}^\varepsilon P(-c; x_1') + \delta_{\mathrm{err}}^{\mathrm{up}}(c, h, s, w^{\mathrm{in}}, w^{\mathrm{out}}),$$

where $\delta_{\mathrm{err}}^{\mathrm{low}}$ and $\delta_{\mathrm{err}}^{\mathrm{up}}$ are the total error terms , each of which admits a rigorous explicit bound.

At a high level, Algorithm 1 (CalculatePMF) constructs, via FFT, a discrete approximation to the distribution of the sum $\sum_{i=2}^M l_\varepsilon(Y_i)$, where $M \sim 1 + \mathrm{Bin}(n-1, \gamma)$ as in Lemma 4.1 and Y_i are i.i.d. draws from $\mathcal{R}_{\mathrm{ref}}$. More precisely, we first truncate the privacy amplification random variable to a bounded interval, discretize it on a uniform grid, and then use the FFT to compute the probability mass function of the resulting discretized sum. This step replaces the sum of i.i.d. (continuous) variables by a finitely supported discrete distribution that approximates $\sum_{i=2}^M l_\varepsilon(Y_i)$ up to controlled truncation, discretization and wrap-around errors.

Algorithm 2 (CalculateMainTerm) then uses this discrete distribution approximating $\sum_{i=2}^M l_\varepsilon(Y_i)$ together with the CDFs of $l_\varepsilon(Y)$ under $\mathcal{R}_{x_1}$ and $\mathcal{R}_{x_1'}$ to approximate the probabilities $\Pr\Big[\sum_{i=2}^M l_\varepsilon(Y_i) > -l_\varepsilon(Y_1)\Big]$. Because the distribution of $\sum_{i=2}^M l_\varepsilon(Y_i)$ has been discretized, these probabilities reduce to finite sums over grid points, and are therefore straightforward to evaluate numerically.

4.2 Analysis of FFT Algorithm

Algorithms 1 and 2 depend on several numerical parameters $(c, s, w^{\mathrm{in}}, h, w^{\mathrm{out}})$ that control truncation, discretization, and FFT wrap-around. Using the asymptotic characterization of the blanket divergence from Lemma 3.2, we can tune these parameters to control both the relative error and the computational cost of the FFT-based approximation. The following theorem makes this precise.

[5] Our error analysis assumes exact real arithmetic. While practical implementations relying on floating-point arithmetic introduce round-off errors, analyzing their propagation is orthogonal to the algorithmic errors characterized in this work.

Theorem 4.3. *Fix a distinct pair $x_1, x_1' \in \mathcal{X}$ and let $\mathcal{R} \in \mathfrak{R}$. Let the reference distribution be $\mathcal{R}_{\mathrm{ref}} \in \{\mathcal{R}_{\mathrm{BG}}\} \cup \{\mathcal{R}_x : x \in \mathcal{X}\}$. Fix target knobs $\eta_{\mathrm{main}}, \eta_{\mathrm{trunc}}, \eta_{\mathrm{disc}}, \eta_{\mathrm{alias}} \in \mathbb{R}$. Let $\varepsilon_n \to 0$ be a sequence of privacy levels such that $\varepsilon_n = \Theta\big(\sqrt{\log n/n}\big)$. We denote by L_n and U_n the lower and upper bounds from Theorem 4.2, that is, $L_n := P(-c; x_1) - \mathrm{e}^{\varepsilon} P(c; x_1') - \delta_{\mathrm{err}}^{\mathrm{low}}$ and $U_n := P(c; x_1) - \mathrm{e}^{\varepsilon} P(-c; x_1') + \delta_{\mathrm{err}}^{\mathrm{up}}$ We define our numerical approximation as the midpoint $\widehat{D}_n := \frac{U_n + L_n}{2}$.*

Then we can tune parameters $(c_n, s_n, w_n^{\mathrm{in}}, h_n, w_n^{\mathrm{out}})$ depending on $(\eta_{\mathrm{main}}, \eta_{\mathrm{trunc}}, \eta_{\mathrm{disc}}, \eta_{\mathrm{alias}})$ such that

$$\frac{|\widehat{D}_n - \mathcal{D}^{\mathrm{blanket}}|}{\mathcal{D}^{\mathrm{blanket}}} = O\big(\eta_{\mathrm{main}} + \eta_{\mathrm{trunc}} + \eta_{\mathrm{disc}} + \eta_{\mathrm{alias}}\big) \qquad (n \to \infty),$$

when $|\mu_{S^{\mathrm{di}}} - \mu_{S^{\mathrm{tr}}}| = O(c_n)$[6] and $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ is a non-lattice distribution[7]. The overall running time of Algorithms 1–2 with these parameters is $O\left(\frac{n}{\eta_{\mathrm{main}}}(\log \frac{n}{\eta_{\mathrm{main}}})^3\right)$, where log log *factors are omitted.*

In particular, in the regime $\varepsilon_n = \Theta(\sqrt{\log n/n})$, the theorem shows that we can prescribe a target relative error level $\eta > 0$ and choose the numerical knobs so that the relative error becomes $O(\eta)$, which leads to an overall running time $\widetilde{O}\left(\frac{n}{\eta}\right)$, up to polylogarithmic factors in n and $1/\eta$. We provide details on parameter selection and implementation in the full version of the paper [27].

We briefly discuss regimes in which the algorithm can become less effective in practice. First, when the shuffle index χ is small, the relevant distributions for the privacy-amplification sum are effectively more spread out, so representing them on a discrete grid requires a larger window and more bins; heuristically, the required grid size (and hence runtime) can grow like $O(1/\chi^2)$, making computations expensive for small χ. In practice, as discussed in Section 3.4, privacy amplification by shuffling is meaningful only when the shuffle index is at least moderately large; the parameter regime of primary interest typically corresponds to larger χ, where the FFT computation remains well behaved. Second, our parameter-tuning rules are asymptotic, so for small n they may be miscalibrated and yield insufficient accuracy. That said, the method produces certified upper and lower bounds, so one can directly inspect the resulting gap to assess tightness and, if needed, retune the numerical parameters to achieve the desired accuracy. Finally, while floating-point round-off is not captured by the real-arithmetic analysis, prior work shows that the FFT-based computations can remain accurate down to about 10^{-10} when using long double precision [16], which is sufficient for the typical range of δ values of interest in DP.

5 Experiments

In this section, we complement our theoretical results with five experiments. The first two visualize the privacy curves from Theorems 3.5 and 3.8. The next two illustrate the accuracy-runtime trade-offs of our FFT-based accountant from Theorems 4.2 and 4.3. The final experiment compares generalized Gaussian mechanisms with k-RR in a distribution-estimation task.

Unless stated otherwise, all computations are carried out in double precision and we choose $(c, s, w^{\mathrm{in}}, h, w^{\mathrm{out}})$ according to Theorem 4.2 and set $\eta_{\mathrm{trunc}}, \eta_{\mathrm{disc}}, \eta_{\mathrm{alias}} = 10^{-3}$. For each mechanism we first identify the asymptotically worst-case neighboring pair (x_1, x_1') and the corresponding shuffle indices χ_{lo} and χ_{up} as in Theorem 3.5, and then run the FFT-based accountant at this pair.

[6] While this condition is not theoretically guaranteed for discrete mechanisms, we empirically confirmed $|\mu_{S^{\mathrm{di}}} - \mu_{S^{\mathrm{tr}}}| \le c_n$ in the regime of our experiments.

[7] A finitely supported random variable is lattice iff its support is contained in a set of the form $a + h\mathbb{Z}$ for some $a \in \mathbb{R}$ and $h > 0$, that is, all pairwise differences between support points are integer multiples of a common nonzero number. For mechanisms with at least three output symbols this requires special algebraic relations, so for generic choices of parameters $l_\varepsilon(Y; x_1, x_1', \mathcal{R}_{\mathrm{ref}})$ is non-lattice. The 2-RR is a nongeneric exception: in that case l_ε takes only two values and is lattice.

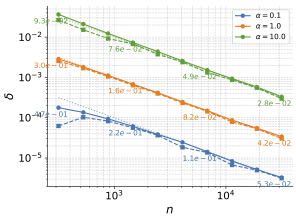

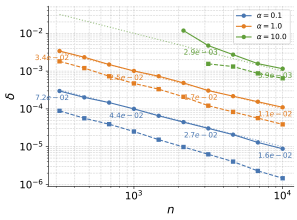

(a) Upper (○) and lower (□) shuffled mechanism's DP δ for $\varepsilon_n(\alpha, \underline{\chi}_{\text{lo}})$ of 3-RR with $\varepsilon_0 = 2$ (left) and the Gaussian mechanism with $\sigma_0 = 2$ (right), both with annotated ε_n; the dotted line indicates α/n.

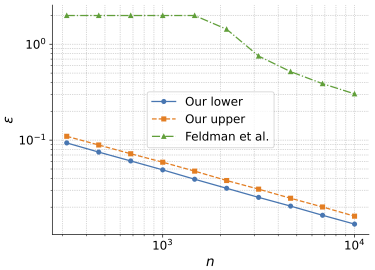

(b) Comparison of the shuffled ε as a function of n between our analysis and that of Feldman et al. [14]. We use the Gaussian mechanism with $\sigma_0 = 2$ and fix $\delta = 10^{-5}$.

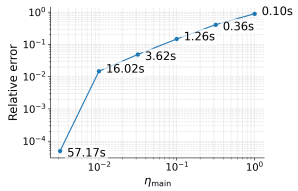

(c) Relative error as a function of the accuracy knob η_{main} with annotated runtime.

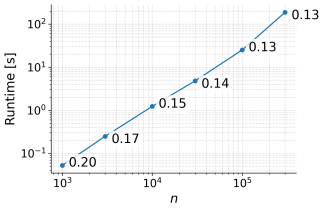

(d) Runtime as a function of the number of users n with annotated relative error.

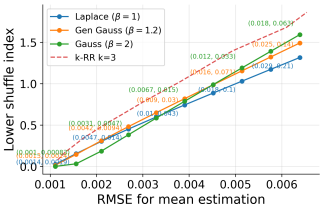

(e) χ_{lo} as a function of the RMSE for mean estimation with annotated RMSE for variance and 3rd moment estimation.

Fig. 4. Experimental validation of the FFT-based blanket-divergence accountant and the shuffle-index analysis.

5.1 Privacy Curves and Shuffle Index

We empirically evaluate the privacy guarantees of Theorem 3.5. For a fixed target constant $\alpha > 0$ and a given local randomizer we set $\varepsilon_n = \varepsilon_n(\alpha, \chi_{\text{lo}})$ and explore the behavior of the shuffled mechanism's hockey-stick divergence δ_n as a function of n for two mechanisms. That is, we use the blanket-based *upper* and *lower* bounds on the *true* shuffled hockey-stick divergence: the Balle et al. upper bound (Lemma 2.5) and the Su et al. lower bound (Theorem 2.6), both evaluated via our FFT accountant at the asymptotically worst-case pair.

The left panel of Fig. 4a shows the results for the case of 3-RR with $\varepsilon_0 = 2$. For each n we plot the upper and lower bounds on the true hockey-stick divergence δ_n of the shuffled mechanism corresponding to $\varepsilon_n(\alpha, \chi_{\text{lo}})$. Across the entire range of n the two curves are practically indistinguishable. This behavior is exactly what Theorem 3.8 predicts for k-RR with $k \geq 3$. On the theoretical side, the lower and upper shuffle indices coincide, $\chi_{\text{lo}} = \chi_{\text{up}}$, so the asymptotic band in Theorem 3.5 collapses. Empirically, we see that this collapse already occurs for moderate values of n, confirming that the blanket-divergence analysis is effectively optimal for this mechanism.

The right panel of Fig. 4a reports the same experiment for the Gaussian local randomizer with noise parameter $\sigma = 2$. Here the upper and lower curves form a visible but relatively narrow band. This matches our theoretical findings: the Gaussian mechanism does not satisfy the structural condition of Theorem 3.8, so the shuffle indices χ_{lo} and χ_{up} need not coincide. Nevertheless, their ratio stays close to one (above 0.7 in our parameter range), and the band remains tight across all n.

Next, we compare our direct analysis of the blanket divergence with the mechanism-agnostic shuffle analysis of Feldman et al. [14] of the calibrated Gaussian local randomizer [5]. Fig. 4b shows the results. Because of the mechanism-agnostic nature and the reliance on a coarse reduction to a pure LDP surrogate, the resulting bounds incur a large constant-factor loss. Empirically, we observe that even at $n = 10^3$, the method predicts essentially no privacy amplification, and for larger n the predicted shuffled ε remains 10 times larger than our direct analysis.

5.2 Accuracy and Runtime of the FFT Accountant

We empirically validate the error and complexity guarantees of Theorems 4.2 and 4.3. Throughout this subsection we fix the local randomizer to be 3-RR with $\varepsilon_0 = 2$ and vary either the accuracy knob η_{main} or the number of users n.

Fig. 4c shows the effect of the parameter η_{main} on both accuracy and runtime. For each value of η_{main} we run the FFT-based algorithm for a fixed (n, ε). Theorem 4.2 provides rigorous lower and upper bounds on the blanket divergence, $D_{\text{low}} \leq \mathcal{D}^{\text{blanket}} \leq D_{\text{up}}$, so the true value is guaranteed to lie in the interval $[D_{\text{low}}, D_{\text{up}}]$. As our notion of relative error, we plot the certified relative uncertainty bandwidth $\frac{D_{\text{up}} - D_{\text{low}}}{D_{\text{up}}}$. As predicted by Theorem 4.3, the relative error shrinks approximately linearly as η_{main} decreases, while the runtime grows roughly like $1/\eta_{\text{main}}$.

In Fig. 4d we fix η_{main} and vary the number of users n. For each n we again compute the interval $[D_{\text{low}}, D_{\text{up}}]$ and plot both the runtime and the relative error. The relative error remains essentially flat as n increases, confirming that our parameter selection keeps the relative error under control uniformly in n. At the same time, the runtime grows almost linearly in n, with only mild logarithmic deviations. This matches the theoretical $\widetilde{O}(n/\eta)$ complexity of Theorem 4.3.

5.3 Distribution Estimation

Finally, we compare the generalized Gaussian mechanisms with k-RR in a distribution estimation task. We fix the data-generating distribution to be $X \sim \text{Unif}[0, 1]$ and set the number of users to $n = 10^5$. For each local mechanism, we evaluate the empirical root-mean-squared error (RMSE) of the mean estimation. Fig. 4e plots the RMSE on the horizontal axis and the lower shuffle index χ_{lo} on the vertical axis. For generalized Gaussian mechanisms we additionally annotate each point with the RMSEs for variance and third-moment estimation, whereas for k-RR we only plot the mean-accuracy vs. shuffle-index trade-off.

In this experiment we focus on k-RR with $k = 3$, which we found to yield the best shuffle index among the values of k under the same mean-estimation accuracy. Under this choice, the figure shows that, for the same level of mean RMSE, 3-RR achieves the largest lower shuffle index in our parameter range, and hence provides the strongest privacy. However, a crucial caveat is that 3-RR (and, more generally, k-RR) does not natively support estimation of higher-order moments: estimating the variance or third central moment typically requires designing a separate mechanism (and hence additional privacy budget). In contrast, generalized Gaussian mechanisms simultaneously support estimation of multiple moments from the same shuffled output. Fig. 4e reveals that the generalized Gaussian family exhibits a nontrivial dependence on the shape parameter β: the value of β that optimizes the shuffle index varies with the target mean-estimation accuracy. In particular, in more noisy regimes (larger local noise), the Gaussian mechanism ($\beta = 2$) attains the largest lower shuffle index among the generalized Gaussian family. Moreover, over a wide range of mean-accuracy levels, the Gaussian mechanism also achieves the smallest RMSE for variance and third-moment estimation among the generalized Gaussian mechanisms, as indicated by the annotations. Taken together, these results suggest that generalized Gaussian mechanisms ($\beta > 1$) provide a better

privacy-utility trade-off than pure local DP mechanisms. Importantly, this kind of optimization is made possible by our unified framework beyond pure local DP.

6 Limitation and Conclusion

Our moderate-deviation guarantees are formulated in terms of an *asymptotically* worst-case neighboring pair, that is, a pair $(x_1^*, x_1'^*)$ that attains the worst-case shuffle index in the limit $n \to \infty$. While this is natural from the point of view of our CLT-based analysis, it does not rule out the possibility that, for finite n, special pairs might induce slightly worse (ε, δ) trade-offs than those induced by the asymptotic worst-case pair. Establishing conditions under which the asymptotically worst-case pair is also worst-case at finite n is a problem that we leave for future work.

Our work is, to the best of our knowledge, the first to go beyond the pure local DP viewpoint and organize privacy amplification by shuffling via the shuffle index. Our framework provides a unified way to reason about shuffle amplification and to obtain tight (ε, δ) guarantees. We believe that these conceptual and technical advantages are valuable both for the theoretical understanding of shuffle DP and for its practical deployment in real systems.

References

[1] Jürgen Angst and Guillaume Poly. 2017. A Weak Cramér Condition and Application to Edgeworth Expansions. *Electronic Journal of Probability* 22, article 59 (2017), 1–24. doi:10.1214/17-EJP77

[2] Hilal Asi, Vitaly Feldman, and Kunal Talwar. 2022. Optimal Algorithms for Mean Estimation under Local Differential Privacy. In *Proceedings of the 39th International Conference on Machine Learning*. PMLR, 1046–1056. https://proceedings.mlr.press/v162/asi22b.html

[3] Borja Balle, Gilles Barthe, and Marco Gaboardi. 2018. Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences. In *Advances in Neural Information Processing Systems*, Vol. 31. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2018/hash/3b5020bb891119b9f5130f1fea9bd773-Abstract.html

[4] Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. 2019. The Privacy Blanket of the Shuffle Model. In *Advances in Cryptology – CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part II*. Springer-Verlag, Berlin, Heidelberg, 638–667. doi:10.1007/978-3-030-26951-7_22

[5] Borja Balle and Yu-Xiang Wang. 2018. Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising. In *Proceedings of the 35th International Conference on Machine Learning*. PMLR, 394–403. https://proceedings.mlr.press/v80/balle18a.html

[6] Gilles Barthe and Federico Olmedo. 2013. Beyond Differential Privacy: Composition Theorems and Relational Logic for f-Divergences between Probabilistic Programs. In *Automata, Languages, and Programming*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Fedor V. Fomin, Rūsiņš Freivalds, Marta Kwiatkowska, and David Peleg (Eds.). Vol. 7966. Springer Berlin Heidelberg, Berlin, Heidelberg, 49–60. doi:10.1007/978-3-642-39212-2_8

[7] Sayan Biswas, Kangsoo Jung, and Catuscia Palamidessi. 2024. Tight Differential Privacy Guarantees for the Shuffle Model with K-Randomized Response. In *Foundations and Practice of Security*, Mohamed Mosbah, Florence Sèdes, Nadia Tawbi, Toufik Ahmed, Nora Boulahia-Cuppens, and Joaquin Garcia-Alfaro (Eds.). Springer Nature Switzerland, Cham, 440–458. doi:10.1007/978-3-031-57537-2_27

[8] Wei-Ning Chen, Dan Song, Ayfer Ozgur, and Peter Kairouz. 2023. Privacy Amplification via Compression: Achieving the Optimal Privacy-Accuracy-Communication Trade-off in Distributed Mean Estimation. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. arXiv. arXiv:2304.01541 [stat] doi:10.48550/arXiv.2304.01541

[9] Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. 2019. Distributed Differential Privacy via Shuffling. In *Advances in Cryptology – EUROCRYPT 2019*, Yuval Ishai and Vincent Rijmen (Eds.). Springer International Publishing, Cham, 375–403. doi:10.1007/978-3-030-17653-2_13

[10] Lynn Chua, Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Amer Sinha, and Chiyuan Zhang. 2024. How Private Are DP-SGD Implementations?. In *Proceedings of the 41st International Conference on Machine Learning*. PMLR, 8904–8918. https://proceedings.mlr.press/v235/chua24a.html

[11] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. 2013. Local Privacy and Statistical Minimax Rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. 429–438. doi:10.1109/FOCS.2013.53

[12] Cynthia Dwork and Aaron Roth. 2013. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3-4 (2013), 211–407. doi:10.1561/0400000042
[13] Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. 2019. Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '19)*. Society for Industrial and Applied Mathematics, USA, 2468–2479.
[14] Vitaly Feldman, Audra McMillan, and Kunal Talwar. 2022. Hiding Among the Clones: A Simple and Nearly Optimal Analysis of Privacy Amplification by Shuffling. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. arXiv, 954–964. doi:10.1109/FOCS52979.2021.00096
[15] Vitaly Feldman, Audra McMillan, and Kunal Talwar. 2023. Stronger Privacy Amplification by Shuffling for Renyi and Approximate Differential Privacy. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 4966–4981. doi:10.1137/1.9781611977554.ch181
[16] Sivakanth Gopi, Yin Tat Lee, and Lukas Wutschitz. 2021. Numerical Composition of Differential Privacy. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 11631–11642. https://proceedings.neurips.cc/paper/2021/hash/6097d8f3714205740f30debe1166744e-Abstract.html
[17] Antti Koskela, Mikko A. Heikkilä, and Antti Honkela. 2022. Numerical Accounting in the Shuffle Model of Differential Privacy. *Transactions on Machine Learning Research* (Oct. 2022). https://openreview.net/forum?id=11osftjEbF
[18] Seng Pei Liew and Tsubasa Takahashi. 2022. Shuffle Gaussian Mechanism for Differential Privacy. arXiv:2206.09569 [cs] doi:10.48550/arXiv.2206.09569
[19] Seng Pei Liew, Tsubasa Takahashi, Shun Takagi, Fumiyuki Kato, Yang Cao, and Masatoshi Yoshikawa. 2022. Network Shuffling: Privacy Amplification via Random Walks. In *Proceedings of the 2022 International Conference on Management of Data (SIGMOD '22)*. Association for Computing Machinery, New York, NY, USA, 773–787. doi:10.1145/3514221.3526162
[20] Qiyao Luo, Jianzhe Yu, Wei Dong, Quanqing Xu, Chuanhui Yang, and Ke Yi. 2025. RM2: Answer Counting Queries Efficiently under Shuffle Differential Privacy. *Proc. ACM Manag. Data* 3, 3 (June 2025), 210:1–210:24. doi:10.1145/3725415
[21] Valentin V. Petrov. 1975. *Sums of Independent Random Variables*. Springer Berlin Heidelberg, Berlin, Heidelberg. doi:10.1007/978-3-642-65809-9
[22] Alex Shvets. 2026. Universal Shuffle Asymptotics: Sharp Privacy Analysis in the Gaussian Regime. *arXiv preprint arXiv:2602.09029* (2026).
[23] A. D. Slastnikov. 1979. Limit Theorems for Moderate Deviation Probabilities. *Theory of Probability & Its Applications* 23, 2 (March 1979), 322–340. doi:10.1137/1123035
[24] David M. Sommer, Sebastian Meiser, and Esfandiar Mohammadi. 2019. Privacy Loss Classes: The Central Limit Theorem in Differential Privacy. *Proceedings on Privacy Enhancing Technologies* 2019, 2 (April 2019), 245–269. doi:10.2478/popets-2019-0029
[25] Pengcheng Su, Haibo Cheng, and Ping Wang. 2025. Decomposition-Based Optimal Bounds for Privacy Amplification via Shuffling. arXiv:2504.07414 [cs] doi:10.48550/arXiv.2504.07414
[26] Pengcheng Su, Haibo Cheng, and Ping Wang. 2025. Mutual Information Bounds in the Shuffle Model. arXiv:2511.15051 [cs] doi:10.48550/arXiv.2511.15051
[27] Shun Takagi and Seng Pei Liew. 2026. Analysis of Shuffling Beyond Pure Local Differential Privacy. *arXiv preprint arXiv:2601.19154* (2026).
[28] Shaowei Wang, Yun Peng, Jin Li, Zikai Wen, Zhipeng Li, Shiyu Yu, Di Wang, and Wei Yang. 2024. Privacy Amplification via Shuffling: Unified, Simplified, and Tightened. *Proc. VLDB Endow.* 17, 8 (April 2024), 1870–1883. doi:10.14778/3659437.3659444
[29] Shaowei Wang, Sufen Zeng, Jin Li, Shaozheng Huang, and Yuyang Chen. 2025. Shuffle Model of Differential Privacy: Numerical Composition for Federated Learning. *Applied Sciences* 15, 3 (Jan. 2025), 1595. doi:10.3390/app15031595
[30] Tianhao Wang, Jeremiah Blocki, Ninghui Li, and Somesh Jha. 2017. Locally Differentially Private Protocols for Frequency Estimation. In *26th USENIX Security Symposium (USENIX Security 17)*. 729–745. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/wang-tianhao

A The Interpretation of Blanket Divergence

Let $\mathcal{R}$ be a local randomizer with blanket distribution $\mathcal{R}_{\mathrm{BG}}$ and blanket mass γ, so that for every $x \in \mathcal{X}$,

$$\mathcal{R}_x \;=\; \gamma\,\mathcal{R}_{\mathrm{BG}} \;+\; (1-\gamma)\,\mathcal{R}_x^{\mathrm{LO}}, \qquad \mathcal{R}_x^{\mathrm{LO}}(y) := \frac{\mathcal{R}_x(y) - \gamma\mathcal{R}_{\mathrm{BG}}(y)}{1-\gamma}.$$

We define an *extended mechanism* $\widetilde{\mathcal{M}}$ as follows (see Fig. 1). User 1 always outputs $Y_1 \sim \mathcal{R}_{x_1}$ (or $\mathcal{R}_{x_1'}$ under $x'_{1:n}$). For each $i \in \{2, \ldots, n\}$, sample an independent coin $B_i \sim \mathrm{Bernoulli}(\gamma)$. If $B_i = 1$,

output $Y_i \sim \mathcal{R}_{\mathrm{BG}}$; otherwise output $Y_i \sim \mathcal{R}^{\mathrm{LO}}_{x_i}$ and reveal it together with its index i. Let $T := \sum_{i=2}^{n} B_i$ and let $M := 1 + T$ be the number of messages routed to the shuffler. The shuffler is applied only to the M messages $\{Y_1\} \cup \{Y_i : B_i = 1\}$, while the leftover messages $\{(i, Y_i) : B_i = 0\}$ are released without shuffling. Denote the shuffled part of the output by $S_{1:M} \in \mathcal{Y}^M$.

Since the leftover part is identical under $x_{1:n}$ and $x'_{1:n}$, the hockey-stick divergence of $\widetilde{\mathcal{M}}$ reduces to that of the shuffled part. Conditioning on $M = m$, the shuffled part consists of one sample from $\mathcal{R}_{x_1}$ and $m-1$ i.i.d. samples from $\mathcal{R}_{\mathrm{BG}}$ whose order is uniformly randomized. A direct calculation yields the corresponding density

$$p^{(m)}_{x_1}(s_{1:m}) = \frac{1}{m}\sum_{j=1}^{m} \mathcal{R}_{x_1}(s_j) \prod_{k\neq j} \mathcal{R}_{\mathrm{BG}}(s_k) = \Big(\prod_{k=1}^{m} \mathcal{R}_{\mathrm{BG}}(s_k)\Big) \cdot \frac{1}{m}\sum_{j=1}^{m} \frac{\mathcal{R}_{x_1}(s_j)}{\mathcal{R}_{\mathrm{BG}}(s_j)}.$$

Hence, writing

$$l_\varepsilon(y) := \frac{\mathcal{R}_{x_1}(y) - e^\varepsilon \mathcal{R}_{x'_1}(y)}{\mathcal{R}_{\mathrm{BG}}(y)},$$

we have

$$p^{(m)}_{x_1}(s_{1:m}) - e^\varepsilon p^{(m)}_{x'_1}(s_{1:m}) = \Big(\prod_{k=1}^{m} \mathcal{R}_{\mathrm{BG}}(s_k)\Big) \cdot \frac{1}{m}\sum_{j=1}^{m} l_\varepsilon(s_j).$$

Integrating the positive part, this implies the exact identity

$$\mathcal{D}_{e^\varepsilon}\Big(p^{(m)}_{x_1} \Big\| p^{(m)}_{x'_1}\Big) = \mathbb{E}_{Y_{1:m} \overset{\text{i.i.d.}}{\sim} \mathcal{R}_{\mathrm{BG}}}\left[\frac{1}{m}\left(\sum_{j=1}^{m} l_\varepsilon(Y_j)\right)_+\right].$$

Finally, we uncondition over M. In our construction, $T \sim \mathrm{Bin}(n-1, \gamma)$ and $M = 1 + T$, so for $m \in \{1, \dots, n\}$,

$$\Pr[M = m] = \binom{n-1}{m-1}\gamma^{m-1}(1-\gamma)^{n-m}.$$

Therefore,

$$\begin{aligned}\mathcal{D}_{e^\varepsilon}\Big(\widetilde{\mathcal{M}}(x_{1:n}) \Big\| \widetilde{\mathcal{M}}(x'_{1:n})\Big) &= \sum_{m=1}^{n} \Pr[M=m]\, \mathbb{E}_{Y_{1:m}\sim\mathcal{R}_{\mathrm{BG}}}\left[\frac{1}{m}\left(\sum_{j=1}^{m} l_\varepsilon(Y_j)\right)_+\right] \\ &= \sum_{m=1}^{n} \left(\binom{n-1}{m-1}\gamma^{m-1}(1-\gamma)^{n-m} \cdot \frac{1}{m}\right) \mathbb{E}_{Y_{1:m}\sim\mathcal{R}_{\mathrm{BG}}}\left[\left(\sum_{j=1}^{m} l_\varepsilon(Y_j)\right)_+\right].\end{aligned}$$

Using the combinatorial identity $\binom{n-1}{m-1} = \frac{m}{n}\binom{n}{m}$, we obtain

$$\binom{n-1}{m-1}\gamma^{m-1}(1-\gamma)^{n-m} \cdot \frac{1}{m} = \frac{1}{n\gamma}\binom{n}{m}\gamma^{m}(1-\gamma)^{n-m}.$$

Since the $m = 0$ term contributes 0 (because $(\sum_{j=1}^{0} \cdot)_+ = 0$), we may extend the sum to $m = 0$ and recognize the binomial distribution $\mathrm{Bin}(n, \gamma)$:

$$\mathcal{D}_{e^\varepsilon}\Big(\widetilde{\mathcal{M}}(x_{1:n}) \Big\| \widetilde{\mathcal{M}}(x'_{1:n})\Big) = \frac{1}{n\gamma}\, \mathbb{E}_{\substack{M\sim\mathrm{Bin}(n,\gamma) \\ Y_{1:M} \overset{\text{i.i.d.}}{\sim} \mathcal{R}_{\mathrm{BG}}}}\left[\left(\sum_{j=1}^{M} l_\varepsilon(Y_j)\right)_+\right] = \mathcal{D}^{\text{blanket}}_{e^\varepsilon, n, \mathcal{R}_{\mathrm{BG}}, \gamma}\big(\mathcal{R}_{x_1} \| \mathcal{R}_{x'_1}\big).$$

Moreover, the original shuffled mechanism $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$ can be obtained from $\widetilde{\mathcal{M}}$ by post-processing (discarding the revealed indices and leftover messages and retaining only the fully

shuffled multiset of all n messages). By the data processing inequality for the hockey-stick divergence, this yields

$$\mathcal{D}_{e^\varepsilon}\big(\mathcal{M}(x_{1:n})\big\|\mathcal{M}(x'_{1:n})\big) \le \mathcal{D}_{e^\varepsilon}\Big(\widetilde{\mathcal{M}}(x_{1:n})\Big\|\widetilde{\mathcal{M}}(x'_{1:n})\Big) = \mathcal{D}^{\text{blanket}}_{e^\varepsilon,n,\mathcal{R}_{\text{BG}},\gamma}\big(\mathcal{R}_{x_1}\|\mathcal{R}_{x'_1}\big),$$

which recovers the blanket-divergence upper bound.

B The Computation of Shuffle Indices

B.1 Generalized Gaussian Mechanisms

Here, we consider the β-Gaussian local randomizer on $\mathcal{X} = [0,1]$:

$$\mathcal{R}(x) = x + Z, \qquad Z \sim g_{\beta,c}, \qquad g_{\beta,c}(z) := \frac{\beta}{2c\,\Gamma(1/\beta)}\exp\Big(-\Big|\tfrac{z}{c}\Big|^\beta\Big),$$

so that $\phi^\beta_x(y) = g_{\beta,c}(y-x)$.

The blanket density is $\underline{\mathcal{R}}(y) = \inf_{x\in[0,1]}\phi^\beta_x(y)$. Since $x \mapsto |y-x|$ is maximized over $[0,1]$ at the far endpoint,

$$\underline{\mathcal{R}}(y) = \begin{cases}\phi^\beta_1(y), & y \le \frac{1}{2},\\ \phi^\beta_0(y), & y \ge \frac{1}{2},\end{cases} \qquad \gamma := \int_{\mathbb{R}}\underline{\mathcal{R}}(y)\,dy = 2\int_{1/2}^{\infty}\phi^\beta_0(y)\,dy. \tag{3}$$

B.1.1 Lower Shuffle Index. From the definition, for fixed $(x_1, x'_1, \mathcal{R}_{\text{BG}})$, the variance term and the blanket mass in the shuffle index are

$$\sigma^2_{\text{BG},\beta} := \gamma\left\{\int_{-\infty}^{1/2}\frac{\big(\phi^\beta_{x_1}(y)-\phi^\beta_{x'_1}(y)\big)^2}{\phi^\beta_1(y)}\,dy \;+\; \int_{1/2}^{\infty}\frac{\big(\phi^\beta_{x_1}(y)-\phi^\beta_{x'_1}(y)\big)^2}{\phi^\beta_0(y)}\,dy\right\},$$

where $\Gamma(s,x)$ denotes the upper incomplete gamma function. For fixed $\beta > 0$, the quantity $\sigma^2_{\text{BG},\beta}$ (and hence χ_{lo}) can be evaluated efficiently by standard one-dimensional numerical quadrature. As a consequence, searching over (x_1, x'_1) for the lower shuffle index is computationally inexpensive[8].

Here, we adopt the working assumption that the worst-case pair is attained at $(x_1, x'_1) = (0,1)$, and under this assumption, χ_{lo} admits closed-form expressions for the Laplace ($\beta = 1$) and Gaussian ($\beta = 2$) mechanisms when $\text{Var}(Z) = \sigma_0^2$. For Laplace, we have

$$\chi^{(\beta=1)}_{\text{lo}} = \frac{\sqrt{3}}{2}\cdot\frac{t}{\sqrt{1-3t^2+2t^3}}, \qquad t := \exp\left(-\frac{1}{\sqrt{2}\,\sigma_0}\right).$$

For Gaussian, we have

$$\chi^{(\beta=2)}_{\text{lo}} = \frac{1}{\sqrt{2\left(e^{1/\sigma_0^2}\,\Phi\left(\frac{3}{2\sigma_0}\right) - 3\,\Phi\left(\frac{1}{2\sigma_0}\right)+1\right)}},$$

These formulas make the limiting behavior explicit: $\chi^{(\beta=1)}_{\text{lo}}(0,1) \sim \sigma_0/\sqrt{2}$ and $\chi^{(\beta=2)}_{\text{lo}}(0,1) \sim \sigma_0$ as $\sigma_0 \to \infty$, while both decay exponentially fast as $\sigma_0 \to 0$.

[8]In all our numerical experiments for generalized Gaussian mechanisms (including the cases $\beta = 1$ and $\beta = 2$), the pair $(x_1, x'_1) = (0,1)$ consistently attains (or appears to attain) the infimum of the lower shuffle index $\chi_{\text{lo}}(x_1, x'_1)$. This is also intuitively plausible, since $(0,1)$ maximizes the separation between two inputs in the unit interval. However, a general theoretical proof that $(0,1)$ is the unique global minimizer of χ_{lo} for all admissible parameters would require a delicate analysis beyond the scope of this work, and we therefore do not claim such a statement as a theorem.

Fig. 5a illustrates the lower shuffle index χ_{lo} for several values of the shape parameter β. Even when the underlying local randomizers are calibrated to have the same variance (and hence the same accuracy for mean estimation), they exhibit different shuffle indices.

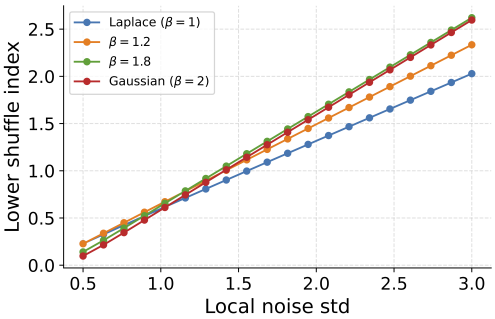

(a) Lower shuffle index for several generalized Gaussian mechanisms.

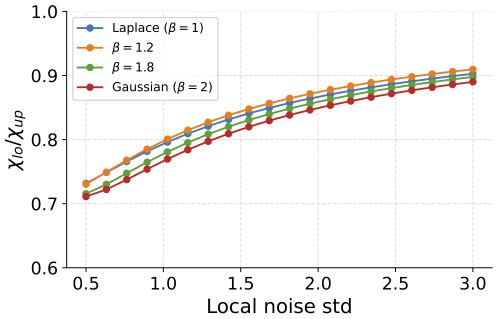

(b) Ratio of lower and upper shuffle indices.

Fig. 5. Shuffle indices for generalized Gaussian mechanisms with different shape parameters β.

B.1.2 Upper Shuffle Index. For the upper shuffle index, we instead take the reference distribution to be one of the local randomizers $\mathcal{R}_x$, and only the variance is needed in our analysis. From the definition, the variance is

$$\sigma^2_{x,\beta} := \int_{-\infty}^{\infty} \frac{\left(\phi^\beta_{x_1}(y) - \phi^\beta_{x'_1}(y)\right)^2}{\phi^\beta_x(y)}\, dy \;=\; c\frac{\beta}{2c\,\Gamma(1/\beta)}\, F_\beta(a_1, a'_1), \tag{4}$$

where

$$F_\beta(a_1, a'_1) := \int_{-\infty}^{\infty} \frac{\left(e^{-|z-a_1|^\beta} - e^{-|z-a'_1|^\beta}\right)^2}{e^{-|z|^\beta}}\, dz, \qquad a_1 := \frac{x_1 - x}{c}, \qquad a'_1 := \frac{x'_1 - x}{c}$$

Although the definition of χ_{up} involves a three-point search over (x_1, x'_1, x), the fact that $\{\phi^\beta_\mu\}_{\mu\in\mathbb{R}}$ forms a location family implies that $\sigma^2_{x,\beta}$ in (4) depends on (x_1, x'_1, x) only through the two offsets $(a_1, a'_1) = ((x_1 - x)/\alpha, (x'_1 - x)/\alpha)$. In particular, the optimization can be reduced to a two-dimensional search over (a_1, a'_1). As in the case of the lower shuffle index, the variance $\sigma^2_{x,\beta}$ can be efficiently evaluated by standard one-dimensional numerical quadrature, so this search remains computationally inexpensive.

Under the same worst-case pair assumption $(x_1, x'_1) = (0, 1)$, the upper shuffle index also admits closed-form expressions for $\beta = 1, 2$ when the noise is parameterized by $\text{Var}(Z) = \sigma_0^2$. For Laplace, taking $\mathcal{R}_{\text{ref}} = \mathcal{R}_0$ yields

$$\chi_{\text{up}}^{(\beta=1)} = \left(\frac{2}{3}e^{2/\sigma_0^2} - 1 + \frac{1}{3}e^{-4/\sigma_0^2}\right)^{-1/2}.$$

For Gaussian, taking $\mathcal{R}_{\text{ref}} = \mathcal{R}_0$ yields

$$\chi_{\text{up}}^{(\beta=2)} = \left(e^{1/\sigma_0^2} - 1\right)^{-1/2}.$$

In both cases, the tightness ratio $\chi_{\mathrm{lo}}/\chi_{\mathrm{up}}$ exhibits the same qualitative limiting behavior. As $\sigma_0 \to \infty$,

$$\frac{\chi_{\mathrm{lo}}^{(\beta=1)}}{\chi_{\mathrm{up}}^{(\beta=1)}} = 1 - \Theta\left(\frac{1}{\sigma_0}\right), \qquad \frac{\chi_{\mathrm{lo}}^{(\beta=2)}}{\chi_{\mathrm{up}}^{(\beta=2)}} = 1 - \Theta\left(\frac{1}{\sigma_0}\right).$$

As $\sigma_0 \to 0$,

$$\frac{\chi_{\mathrm{lo}}^{(\beta=1)}}{\chi_{\mathrm{up}}^{(\beta=1)}} = \frac{1}{\sqrt{2}} + o(1), \qquad \frac{\chi_{\mathrm{lo}}^{(\beta=2)}}{\chi_{\mathrm{up}}^{(\beta=2)}} = \frac{1}{\sqrt{2}} + o(1).$$

The derivation of the asymptotics of χ_{lo} is found in the full version [27].

B.2 k-RR

We next specialize the shuffle indices to the k-RR. For an input $x \in [k] := \{1, \dots, k\}$, the output distribution is

$$\phi_x(y) = \Pr[\mathcal{R}(x) = y] = \begin{cases} p, & y = x, \\ q, & y \neq x, \end{cases} \qquad p = \frac{\mathrm{e}^{\varepsilon_0}}{\mathrm{e}^{\varepsilon_0} + k - 1}, \quad q = \frac{1}{\mathrm{e}^{\varepsilon_0} + k - 1}.$$

The blanket distribution is uniform on $[k]$: $\phi_{\mathrm{BG}}(y) = \frac{1}{k}, \gamma = \sum_y \min_x \phi_x(y) = kq$.

B.2.1 Lower shuffle index. For any distinct pair $x_1 \neq x_1' \in [k]$, the privacy amplification random variable with respect to the blanket distribution is

$$l_0(y; x_1, x_1', \mathcal{R}_{\mathrm{BG}}) = \frac{\phi_{x_1}(y) - \phi_{x_1'}(y)}{\phi_{\mathrm{BG}}(y)} = k\big(\phi_{x_1}(y) - \phi_{x_1'}(y)\big).$$

A simple case distinction shows that, for $Y \sim \mathcal{R}_{\mathrm{BG}}$,

$$l_0(Y; x_1, x_1', \mathcal{R}_{\mathrm{BG}}) = \begin{cases} k(p-q), & Y = x_1, \\ -k(p-q), & Y = x_1', \\ 0, & \text{otherwise}, \end{cases}$$

so that $\mathbb{E}[l_0(Y)] = 0$ and $\sigma_{\mathrm{BG}}^2 := \mathrm{Var}\big(l_0(Y; x_1, x_1', \mathcal{R}_{\mathrm{BG}})\big) = 2k(p-q)^2$. By definition,

$$\chi = \sqrt{\frac{\gamma}{\sigma_{\mathrm{BG}}^2}} = \sqrt{\frac{q}{2(p-q)^2}} = \sqrt{\frac{\mathrm{e}^{\varepsilon_0} + k - 1}{2\big(\mathrm{e}^{\varepsilon_0} - 1\big)^2}}.$$

In particular, χ does not depend on the particular choice of the distinct pair (x_1, x_1'), and hence

$$\chi_{\mathrm{lo}} = \inf_{x_1 \neq x_1'} \chi = \sqrt{\frac{q}{2(p-q)^2}}.$$

B.2.2 Upper shuffle index. For the upper shuffle index we take the reference distribution to be one of the local randomizers $\mathcal{R}_x$. Fix $x_1 \neq x_1'$, and consider

$$l_0(y; x_1, x_1', \mathcal{R}_x) = \frac{\phi_{x_1}(y) - \phi_{x_1'}(y)}{\phi_x(y)}, \qquad Y \sim \mathcal{R}_x.$$

By symmetry, the variance only depends on how x is positioned relative to $\{x_1, x_1'\}$.

- If $x = x_1$ (or symmetrically $x = x_1'$), then

$$l_0(Y; x_1, x_1', \mathcal{R}_{x_1}) = \begin{cases} \dfrac{p-q}{p}, & Y = x_1, \\ -\dfrac{p-q}{q}, & Y = x_1', \\ 0, & \text{otherwise}, \end{cases}$$

  so that $\mathrm{Var}\big(l_0(Y; x_1, x_1', \mathcal{R}_{x_1})\big) = (p-q)^2\Big(\frac{1}{p} + \frac{1}{q}\Big)$.
- If $k \geq 3$ and $x \notin \{x_1, x_1'\}$, then $\mathrm{Var}\big(l_0(Y; x_1, x_1', \mathcal{R}_x)\big) = 2\,\frac{(p-q)^2}{q}$.

Since for $p > q$ we have $2\,\frac{(p-q)^2}{q} > (p-q)^2\Big(\frac{1}{p} + \frac{1}{q}\Big)$, the variance is maximized (and hence χ is minimized) by choosing $x \notin \{x_1, x_1'\}$ whenever $k \geq 3$. Because $\gamma = 1$ for $\mathcal{R}_x$, we obtain, for $k \geq 3$,

$$\chi_{\mathrm{up}} = \inf_{x \in [k]} \sqrt{\frac{1}{\mathrm{Var}(l_0(Y; x_1, x_1', \mathcal{R}_x))}} = \sqrt{\frac{q}{2(p-q)^2}}$$

Thus, χ_{up} also does not depend on the particular choice of the distinct pair (x_1, x_1'), and hence, for k-randomized response with $k \geq 3$, the lower and upper shuffle indices coincide:

$$\chi_{\mathrm{up}} = \chi_{\mathrm{lo}} = \sqrt{\frac{q}{2(p-q)^2}} = \sqrt{\frac{\mathrm{e}^{\varepsilon_0} + k - 1}{2\big(\mathrm{e}^{\varepsilon_0} - 1\big)^2}}.$$

B.2.3 Comparison to the Clone Paradigm. Here, we consider Theorem 7.4 of Feldman [15]: the specific clone reduction of kRR. Let the categorical output alphabet be $\mathcal{A} := \{0, 1, 2, 3\}$. Fix parameters $p = \frac{1}{e^{\varepsilon_0} + k - 1}$ and $q = \frac{k-2}{e^{\varepsilon_0} + k - 1}$. Define the *reference* categorical distribution Q_{ref} on $\mathcal{A}$ by

$$Q_{\mathrm{ref}}(0) = p, \qquad Q_{\mathrm{ref}}(1) = p, \qquad Q_{\mathrm{ref}}(2) = q, \qquad Q_{\mathrm{ref}}(3) = 1 - 2p - q.$$

Define the two *target* categorical distributions Q_0, Q_1 on $\mathcal{A}$:

$$Q_0(0) = e^{\varepsilon_0} p, \qquad Q_0(1) = p, \qquad Q_0(2) = 1 - (1 + e^{\varepsilon_0})p, \qquad Q_0(3) = 0,$$

$$Q_1(0) = p, \qquad Q_1(1) = e^{\varepsilon_0} p, \qquad Q_1(2) = 1 - (1 + e^{\varepsilon_0})p, \qquad Q_1(3) = 0.$$

The reduced shuffled mechanism is then

$$\widehat{\mathcal{M}}_b \;:=\; \mathcal{S}\big(Z_1^{(b)}, Z_2, \ldots, Z_n\big), \qquad Z_1^{(b)} \sim Q_b, \;\; Z_2, \ldots, Z_n \overset{\text{i.i.d.}}{\sim} Q_{\mathrm{ref}}, \qquad b \in \{0, 1\}.$$

That is, one *target* user draws a single categorical message from a distribution Q_b, while each of the remaining $n - 1$ users draws i.i.d. messages from a *reference* categorical distribution Q_{ref}; all n messages are then passed through an ideal shuffler. From the discussion in Section 3.1, the hockey-stick divergence between the two shuffled outputs can be expressed as the blanket divergence with $\gamma = 1$. That is, we can interpret that Theorem 7.4 states that the blanket divergence upper boundeds the hockey-stick divergence of shuffled kRR. Under this interpretation, we may analyze the reduced categorical instance with $\gamma = 1$ and identify the corresponding shuffle index.

Writing $a := e^{\varepsilon_0}$, we obtain the pointwise values of the privacy amplification random variable $l_0(y)$:

$$l_0(0) = \frac{ap - p}{p} = a - 1, \qquad l_0(1) = \frac{p - ap}{p} = 1 - a = -(a-1), \qquad l_0(2) = 0, \qquad l_0(3) = 0,$$

since $Q_0(2) = Q_1(2)$ and $Q_0(3) = Q_1(3) = 0$. Therefore, under $Y \sim Q_{\mathrm{ref}}$,

$$l_0(Y) = \begin{cases} +(a-1) & \text{w.p. } p, \\ -(a-1) & \text{w.p. } p, \\ 0 & \text{w.p. } 1-2p, \end{cases}$$

which implies

$$\sigma^2 = \mathrm{Var}(l_0(Y)) = \mathbb{E}[l_0(Y)^2] = 2p\,(a-1)^2.$$

Consequently, the shuffle index for this $\gamma = 1$ categorical instance is

$$\chi \;=\; \frac{1}{\sqrt{2p}\,(a-1)}.$$

Finally, one has $p = \frac{1}{a+k-1}$ (equivalently, $p = \frac{1}{e^{\varepsilon_0}+k-1}$), hence

$$\chi \;=\; \frac{\sqrt{e^{\varepsilon_0}+k-1}}{\sqrt{2}\left(e^{\varepsilon_0}-1\right)}.$$

This corresponds to χ_{lo} of kRR computed in the previous section.

Received December 2025; accepted February 2026