Privacy Loss Classes: The Central Limit Theorem in Differential Privacy

David Sommer
ETH Zurich
david.sommer@inf.ethz.ch

Sebastian Meiser
UCL
s.meiser@ucl.ac.uk

Esfandiar Mohammadi
ETH Zurich
mohammadi@inf.ethz.ch

September 3, 2018

Abstract

In recent years, privacy enhancing technologies have gained tremendous momentum and they are expected to keep a sustained importance. Quantifying the degree of privacy offered by any mechanism working on potentially sensitive data is a complex and well-researched topic; ε-differential privacy (DP) and its slightly weaker and more versatile variant (ϵ, δ)-approximate differential privacy (ADP) have become the de-facto standard for privacy measures in the literature. Recently, novel variants of (A)DP focused on giving tighter privacy bounds under continual observation. In this paper, we unify many of these previous works in a common core theory, focused on the *privacy loss* of a mechanism. We show that in sequential composition of the mechanism, the privacy loss (represented as a distribution) undergoes a convolution, which in turn enables us to show the central limit theorem for differential privacy: the privacy loss of any mechanism will converge to a Gauss distribution. This observation leads us to several practically relevant insights: 1) we show that several of the novel DP-variants are equally expressive as ADP, 2) we improve existing bounds, such as the moments accountant bound, 3) we derive *exact* ADP guarantees for the Gauss mechanism, i.e., an analytical and simple formula to directly calculate ADP (not an over-approximating bound), 4) we derive *exact* ADP guarantees for the Randomized Response, and, 5) we characterize the privacy guarantees of a mechanism by the Gauss distribution to which it converges, its *privacy class*, and using normal approximation theorems derive novel upper and lower ADP bounds for arbitrary mechanisms.

Contents

1 Introduction

Privacy-preservation of personal data is an increasingly important design goal of data processing systems, in particular with recently enacted strong privacy regulations [23]. Modern systems, however, are increasingly reliant on personal data to provide the expected utility. Hence, privacy and utility are often diametrical, rendering perfect solutions impossible but leaving space for solutions that provide privacy under limited usage.

For such scenarios recent work [8] proposes to prove that the algorithms that process sensitive data (called mechanism) satisfy a quantitative strong privacy notion, called (ε, δ)-approximate differential privacy (ADP), for a single usage. The literature provides many successful examples that prove ADP guarantees [2, 27, 25, 11]. The privacy guarantees naturally deteriorate against attackers that can repeatedly observe the mechanism, i.e., (ε, δ) increase, to a point where using the mechanism is considered insecure. Determining exactly how ADP's (ε, δ) parameters increase under repeated observation turns out to be a challenging task about which there is a rich body of work [15, 2, 6, 10, 20, 19]. Many of these bounds have been shown to be loose [19], which can lead to underestimating the number of observations under which a mechanism provides strong privacy guarantees. While recent work [19] proposed a method for finding tight upper and lower bounds under repeated observation, that work relies on an iterative numerical approach, which can fall prey to numerical errors, memory limitations and discretization problems.

1.1 Contribution

This work examines how fast privacy deteriorates under independent repetitive invocations of these mechanisms. We take a viewpoint that has been proposed by a seminal work by Dinur and Nissim [7], the *privacy loss* of a mechanism, and construct a probability distribution out of it, the *privacy loss distribution*. This privacy loss distribution enables us to precisely argue about sequential composition and to prove the following results.

(a) The privacy loss distribution is sufficient for deriving many differential privacy metrics, including novel bounds such as concentrated differential privacy (CDP), Rényi differential privacy (RDP), as well as the classical notions of pure differential privacy (DP), approximate differential privacy (ADP) and probabilistic differential privacy (PDP) and the Kullback-Leibler divergence. We additionally show that there is a close connection between ADP, the privacy loss distribution, and RDP.

(b) For non-adaptive mechanisms, we prove that this privacy loss distribution evolves under sequential composition, as a convolution of privacy loss distributions. Using the central limit theorem, the privacy loss distribution of any non-adaptive mechanism converges to a Gauss distribution under sufficiently many compositions. This Gauss distribution can be predicted from the privacy loss distribution before composition (i.e., convolution). We characterize mechanisms by the Gauss distribution (i.e., the mean and variance) to which their privacy loss distribution converges, which we call their *privacy loss class*. As an example we derive the privacy loss class of the randomized response mechanism, the Gauss mechanism, and the Laplace mechanism.

(c) We prove that the privacy loss distribution of the Gauss mechanism is again a Gauss distribution, and the privacy loss distribution of the randomized response mechanism is a binomial distribution. As these distributions remain Gauss / binomial distributions under self-convolution, we derive analytical formulas for ADP and PDP for the Gauss mechanism and the randomized response mechanism. For these mechanism, we hence do not only provide ADP- and PDP-parameters under sequential composition that can be efficiently calculated but that are, in particular, tight.

(d) We prove ADP and PDP upper and lower bounds for any mechanism (for which we know the privacy loss distribution before composition) after n-fold sequential composition. Using the Berry-Esseen and Nagaev normal approximation theorems, we can approximate the privacy loss distribution after n convolutions (i.e., for n-fold sequential composition).

(e) We prove that any ADP bounds on a differentially private mechanism can be translated to bound a variation of this mechanism that includes distinguishing events. As an example, we generalize the RDP bounds [2, 20] for the Gauss mechanism to RDP bounds for the truncated Gauss mechanism.

(f) We apply our exact characterization of the Gauss mechanism to show that it clearly outperforms the Laplace mechanism under composition in terms of a variance to privacy trade-off: A Gauss

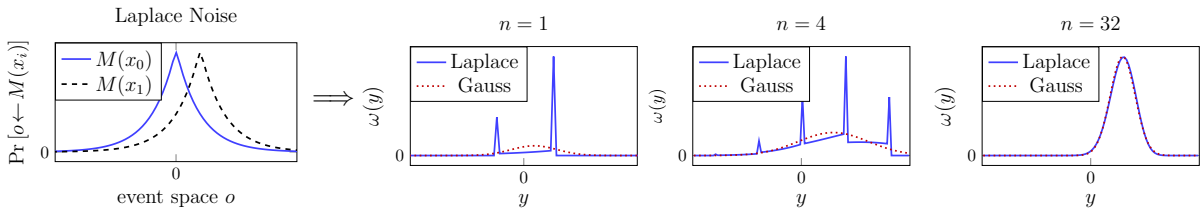

Figure 1: Laplace in Privacy Loss Space for different number of compositions n. Recall that a composition of two independent mechanisms corresponds to a convolution of the privacy loss distribution. As illustration of the privacy loss class and in the spirit of the central limit theorem for differential privacy, a Gauss with identical μ and σ^2 as the shown privacy loss distribution has been plotted.

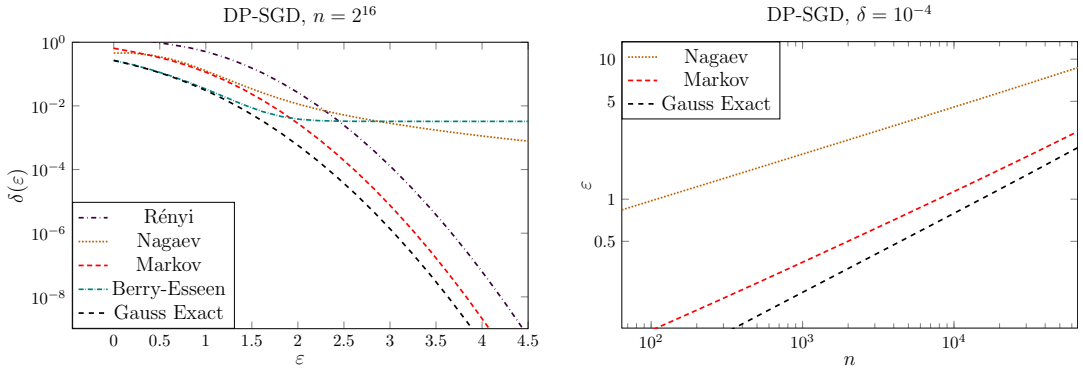

Figure 2: Comparing bounds for differentially private stochastic gradient descent mechanism with noise parameter $q = 0.01$ and $\sigma = 4$. Left: after $n = 2^{16}$ compositions, right: minimal ε values over the number of compositions n for $\delta \leq 10^{-4}$. In the right graph, the Rényi-DP and Berry-Esseen bound did not fall into the plotting range and were omitted.

mechanism with half the variance as the Laplace mechanism provides the same privacy guarantees, for ADP, PDP and even for (almost) pure DP, except for a tiny overhead in delta, which in our example ($\sigma = 40$) can be considered negligible even by cryptographic standards: $< 2^{-80}$ after 128 compositions and $< 2^{-150}$ after 256 compositions.

2 Overview

We illustrate a selection our results to highlight some key ideas. Dwork and Rothblum defined the *privacy loss* of any observable outcome of a mechanism M on inputs x_0 or x_1 as the logarithmic ratio between the probability to observe the outcome y if x_0 is the input compared to if x_1 is the input.

$$\mathcal{L}_{M(x_0)/M(x_1)}(y) = \ln\left(\frac{\Pr[M(x_0) = y]}{\Pr[M(x_1) = y]}\right).$$

This privacy loss spans a real-valued random variable obtained by sampling $y \sim M(x_0)$ and outputting $\mathcal{L}_{M(x_0)/M(x_1)}(y)$, which in turn defines the privacy loss distribution, i.e. the distribution of events o which ratio is equal to $\mathcal{L}(o)$.

In this work, we quantify the differential privacy of probabilistic mechanisms M by analyzing a concrete pair of distributions, given privacy parameters and sensitivity of M.[1] In order to derive differential privacy bounds from the analysis of concrete pairs of distributions A, B, we consider so-called worst-case distributions, for which the differential privacy bounds are worse than the bounds for $M(x_0), M(x_1)$, for any input pair x_0, x_1. As discussed in a recent work [19], for non-adaptive mechanisms there is always such a pair of worst-case distributions [9, 2]. As an illustrative example, consider the simple and commonly used case of a database-query-response system with real-valued queries $q : X \to \mathbb{R}$; to

[1] As an example, the privacy parameter of the Gauss mechanism is the standard deviation σ and the sensitivity is the difference of the means, and of the Laplace mechanism is the scale parameter λ and the sensitivity is difference of the means as well.

preserve the privacy of users' entries, the system outputs adds noise to the answers before releasing them: $M(x) := q(x) + N$, where N is a symmetrically distributed random variable with mean zero, e.g., given by the Laplace distribution or the Gauss distribution. If we define the sensitivity of q for a pair of inputs x_0, x_1 as $|q(x_0) - q(x_1)|$, then for a given sensitivity s, the corresponding distribution for the random variable $M(0)$ and $M(s)$ are worst-case distribution. The privacy loss of any output $y \in \mathbb{R}$ is given by the logarithmic ratio between the probabilities that the output occurs: $\mathcal{L}_{M(0)/M(s)}(y) = \ln\left(\frac{M(0)=y}{M(s)=y}\right)$ and the privacy loss distribution is given by sampling $y \approx M(0)$ and outputting $\mathcal{L}_{M(0)/M(s)}(y)$. Since the noise N in our example is symmetrically distributed, only considering $\mathcal{L}_{M(0)/M(s)}(y)$ suffices; for non-symmetrically distributed noise we would additionally investigate $\mathcal{L}_{M(s)/M(0)}(y)$.

Given a pair of distribution, we can consider the corresponding privacy loss distribution. For non-adaptive mechanisms, this privacy loss distribution naturally evolves under sequential composition as a convolution of privacy loss distributions (Theorem 1), as Figure 1 illustrates for the Laplace mechanism. In other words, understanding the privacy loss distribution for one adversarial observation is sufficient to compute the privacy loss under an arbitrary number of observations. By the central limit theorem, the privacy loss distribution of any non-adaptive mechanism converges to a predictable Gauss distribution under sufficiently many compositions (Theorem 4).

The privacy loss distribution of the Gauss mechanism turns out to be another Gauss distribution. As the convolution of two Gauss distributions is again a Gauss distribution, we can give an analytical and efficiently computable formula for any number of compositions for ADP and PDP. Note that these are not approximate bounds, but indeed precise characterizations (Theorem 5, Figure 4).

For arbitrary mechanisms (for which a worst-case reduction exists), we can offer bounds that are in some cases better than previous work, in particular for a very large number n of compositions ($n > 2^{22}$). Our representation in the privacy loss space directly shows that the moments accountant and the RDP bound, actually are an application of the Markov inequality to compute PDP. With our representation, we can naturally extend that bound to ADP, which results in tighter bounds (Markov-ADP bound: Theorem 3). At the same time, we can apply normal approximation theorems (the Berry-Esseen Theorem and Nagaev-Bound) to achieve tight bounds for a very large number of observations and very small epsilons, as is, e.g., needed for timing leakage analyses as in CoverUp [26], see Figure 5. The minimum of these normal approximation bounds and the ADP-version of the Markov inequality, achieves a very competitive bound, in particular for a very large number of observations. We offer an efficient implementation for computing this minimum.

Figure 2 illustrates our results. The left graph plots for a recent mechanism for training deep neural networks [2] for each ε the minimal $\delta(\varepsilon)$ after 2^{16} compositions. The right graph shows the minimal ε for which $\delta(\varepsilon) < 10^{-4}$ over the number of compositions n. The figure displays the performance of our improved Markov-ADP bound and the performance of our normal approximation bounds, Berry-Esseen and Nagaev. The figure even displays that our exact bound for the Gauss distribution that matches the privacy loss class of the mechanism is very close to the other bounds. Section 6 provides strong evidence that the privacy loss class is actually an accurate characterization of the privacy-preservation of a mechanism and even closer to the tight bounds.

What about utility and sensitivity? By considering pairs of distributions, we abstract away from utility and sensitivity. As argued above, for any sensitivity and any utility function non-static mechanisms have worst-case distributions.

3 Related work

Meiser and Mohammadi [19] have recently introduced a novel numerical method for computing ADP bounds, based on a pair of distributions. Their work investigated the privacy loss of mechanisms and approximated this loss to give very good ADP bounds (including lower bounds) under continual observation. Their work is exploratory and interesting, but lacks the mathematical insights provided here. Moreover, they have higher computation requirements, in particular for a very large number n of observations. For the Gauss mechanism our results (Theorem 5) clearly show tighter results for very large n. When repeating their CoverUp analysis, our approach also shows far better results for very high n values, which is highly relevant for a system like CoverUp.

Kairouz et al. [15] derive tight ADP bounds for the approximate randomized response mechanism (ARR) and use these bounds to prove upper ADP bounds for any mechanism. Their work, however, characterizes set of bounds for the ARR mechanism that contains the tight bounds. This results in a non-trivial optimization problem to find the minimal bounds in this set of bounds. We derive a formula

$M(x_0)$, $M(x_1)$	random variable of a probabilistic mechanism applied to input x_0 and x_1, often abbreviated as A and B
$\Pr[o \leftarrow A]$	probability of o in A
$\mathcal{X}$	set of mechanism-inputs
$\mathcal{U}$	universe of the mechanisms' the atomic events
o	atomic event in $\mathcal{U}$
$\mathcal{L}_{A/B}(o)$	privacy loss of observation o of A and B
ω	privacy loss distribution (PLD)
y	privacy loss (i.e., atomic event) in the PLD
$\omega(y)$	privacy loss pdf/pmf for y
$\mathcal{Y}$	set of atomic events in the PLD, the image of $\mathcal{L}_{A/B}(\mathcal{U})$
ω, $\omega(y)$, $\mathcal{C}$	dual PLD of ω (Definition 4.3)

Table 1: Notation table

(Example 1) for the ARR under sequential composition that directly computes such minimal bounds.

Recent work on concentrated differential privacy (CDP) [10, 6] directly focuses on the privacy loss for deriving tighter ADP and PDP bounds. That work provides interesting insights into differential privacy and into improved bounds for the Gauss mechanism, but for other mechanism that work provides at most very loose bounds. Our work, in contrast, identifies the variance, the mean, and the mass of the distinguishing events of the privacy loss distribution before composition (which we call the privacy loss class) as a valuable characterization for the degree of privacy that a mechanism provides. We illustrate that this characterization is accurate and derive upper and lower ADP and PDP bounds from it.

Rényi differential privacy [20] introduces a privacy notion that is based on the log normalized-moments of the privacy loss distribution (the Rényi divergence). It is a generalization of the moments account bound [2]. We evaluate the moments accountant bound in Section 6, and show that there is a close connection between Rényi differential privacy and ADP (Theorem 2).

4 Privacy Loss Space

We review the privacy loss, a representation of the privacy leakage introduced by Dinur and Nissim [7]. We define a probability distribution from it, the *privacy loss distribution* (PLD), and show that it is useful for defining many privacy notions from the literature: approximate differential privacy [8], probabilistic differential privacy [17, 13], and Rényi differential privacy [20]. We further prove that a sequential composition translates to convolution of the respective privacy loss distributions.

4.1 Privacy Loss Variables and Distributions

At the core of this work lies the representation of privacy leakage as the privacy loss. The privacy loss $\mathcal{L}$ of any one output of the mechanism with respect to two potential inputs is the logarithmic ratio between the probabilities to observe the output for each input. This ratio is of course not defined if this probability is 0 for either the nominator or the denominator. For a more uniform treatment of realistic mechanisms, we introduce distinct symbols ∞ and $-\infty$ that behave similar to infinity and minus infinity. If the nominator is 0, we define the privacy loss $\mathcal{L}$ to be $-\infty$, and analogously if only the denominator is 0 we define it to be ∞. This captures distinguishing events, which, if observed, reveal which of the two inputs was used.

Definition 4.1. *Given a probabilistic mechanism $M : \mathcal{X} \to \mathcal{U}$, let $o \in \mathcal{U}$ be any potential output of M and let x_0, $x_1 \in \mathcal{X}$ be two inputs. We define the* privacy loss *random variable of o for x_0, x_1 as*

$$\mathcal{L}_{M(x_0)/M(x_1)}(o) = \begin{cases} \infty & \text{if } \Pr[o \leftarrow M(x_0)] \neq 0 \text{ and } \Pr[o \leftarrow M(x_1)] = 0 \\ \ln\left(\frac{\Pr[o \leftarrow M(x_0)]}{\Pr[o \leftarrow M(x_1)]}\right) & \text{if } \Pr[o \leftarrow M(x_i)] \neq 0 \ \forall i \in \{0,1\} \\ -\infty & \text{else,} \end{cases}$$

where we consider ∞ and $-\infty$ to be distinct symbols.

For readability, we write $A := M(x_0)$ and $B := M(x_1)$ for the output distributions of M on two particular inputs x_0 and x_1 and then write $\mathcal{L}_{A/B}(o) = \ln\left(\frac{\Pr[o \leftarrow A]}{\Pr[o \leftarrow B]}\right)$ for the privacy loss of the observation o.

The privacy loss $\mathcal{L}$ naturally gives rise to a probability distribution over the privacy losses, *the privacy loss distribution* (PLD), for two given probability distributions A and B. The set of privacy losses $\mathcal{Y} := \bigcup_{o \in \mathcal{U}} \{\mathcal{L}_{A/B}(o)\}$ are the atomic events of the distribution. The respective probability density/mass function ω of a privacy loss y is defined as the cumulative weight of all observations o in A with privacy loss y: $\omega(y) := \sum_{\{o \mid \mathcal{L}_{A/B}(o)=y,\, o \in \mathcal{U}\}} \Pr[o \leftarrow A]$ with $y \in \mathcal{Y}$. Formally, the PLD is the compound probability distribution of the random variable $\mathcal{L}$. To be able to sum over all events, we require the universe $\mathcal{U}$ to be countable. For continuous distributions, this restriction can be generalized to Lebesgue measurable sets which we will do in Section 5.2.

Definition 4.2 (Privacy Loss Distribution (PLD))**.** *A* privacy loss distribution ω *of A over B is defined as follows: there exist two probability distributions A, B over the countable universe $\mathcal{U}$ such that*

$$\mathcal{Y} = \bigcup_{o \in \mathcal{U}} \{\mathcal{L}_{A/B}(o)\} \subset \mathbb{R} \tag{1}$$

$$\omega(y) = \sum_{\{o \mid \mathcal{L}_{A/B}(o)=y,\, o \in \mathcal{U}\}} \Pr[o \leftarrow A] \quad \textit{with} \quad y \in \mathcal{Y} \tag{2}$$

The support $\mathcal{Y}$ of ω additionally includes the symbol[2] *$-\infty$: $\operatorname{supp}(\omega) := \{y \mid \omega(y) \neq 0\} \cup \{-\infty\}$. We define $\forall y \in \mathbb{R} : -\infty < y < \infty$, $y + \infty = \infty$, $-\infty + y = -\infty$, $-\infty + \infty = -\infty$.*

Next, we prove basic properties about the PLD.

Lemma 1. *For two distributions A and B, let $\mathcal{Y}$ and $\omega(y)$ be as in Definition 4.2, we have*

1. *The set $\mathcal{Y}$ is countable.*
2. *$\forall y \in \mathcal{Y} : \omega(y) \geq 0$*
3. *$\sum_{y \in \mathcal{Y}} \omega(y) = 1$*
4. *$\omega(\infty) = \sum_{\{x \mid \Pr[x \leftarrow B]=0\}} \Pr[x \leftarrow A]$*
5. *$\omega(-\infty) = 0$*

Proof. The proofs directly follow from Definitions 4.1 and 4.2.

1. $\mathcal{Y}$ is a mapping from the countable set $\mathcal{U}$ and is therefore countable as well.
2. Follows from $\Pr[o \leftarrow A] \geq 0 \; \forall o \in \mathcal{U}$.
3. $\sum_{y \in \mathcal{Y}} \omega(y) = \sum_{o \in \mathcal{U}} \Pr[o \leftarrow A] = 1$.
4. Follows by the definition of the privacy loss $\mathcal{L}$.
5. By definition of $\mathcal{L}$, $o \in \mathcal{U}$:

$$\begin{aligned} \omega(-\infty) &= \textstyle\sum_{\{o \mid \mathcal{L}_{A/B}(o)=-\infty\}} \Pr[o \leftarrow A] \\ &= \textstyle\sum_{\{o \mid \Pr[o \leftarrow A]=0\}} \Pr[o \leftarrow A] \\ &= 0 \end{aligned}$$

□

With these properties at hand, we can prove that the privacy loss distribution of a pair of independent product distributions $A \times C$ vs. $B \times D$ is the same as the convolution of the privacy loss distributions of the pair of single distributions A vs. B and C vs. D. This theorem is vital because sequential composition of non-adaptive mechanisms, translates to the independent product distributions of the respective mechanisms.

[2] We are aware that the support of a probability mass function $\omega(y)$ is usually defined as the set of y with $\omega(y) = 0$. The extension simplifies notation.

Theorem 1 (Composition)**.** *Let $M : \mathcal{X} \to \mathcal{U}$ and $M' : \mathcal{X}' \to \mathcal{U}'$ be independent probabilistic mechanisms, and let $x_0, x_1 \in \mathcal{X}$ and $x'_0, x'_1 \in \mathcal{X}'$. Let ω be the privacy loss distribution created by $M(x_0)$ over $M(x_1)$ with support $\mathcal{Y}$, and ω' with support $\mathcal{Y}'$ by $M'(x'_0)$ over $M'(x'_1)$ respectively. Let ω_c with support $\mathcal{Y}_c$ be the privacy loss distribution created by $M(x_0) \times M'(x'_0)$ over $M(x_1) \times M'(x'_1)$ where $\times$ denotes the independent distribution product. Then, ω_c can be derived from ω and ω' as follows:*

$$\mathcal{Y}_c = \{y_c \mid y_c = y + y' \ \ \forall y \in \mathcal{Y}, \forall y' \in \mathcal{Y}'\} \tag{3}$$

$$\begin{aligned}\omega_c(y) &= (\omega * \omega')(y) \qquad \forall y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\} \\ &= \sum_{\{y,\, y' \mid y_1 + y_2 = y\}} \omega(y) \cdot \omega(y') \end{aligned} \tag{4}$$

$$\omega_c(\infty) = 1 - [1 - \omega(\infty)] \cdot [1 - \omega'(\infty)] \tag{5}$$

$$\omega_c(\text{-}\infty) = 0 \tag{6}$$

where $\omega * \omega'$ is a convolution, and the set $\mathcal{Y}_c$ is countable.

Proof. For ease of readability, we assume that M and M' are defined on the input set $\mathcal{X}$ and on the same output set $\mathcal{U}$. The same proof applies if they are defined on different sets. We put emphasis on the difference between $M(x_0)$ and $M(x_1)$, as well as between $M'(x'_0)$ and $M'(x'_1)$ respectively, which leads to four different probability-terms, namely $\Pr[o \leftarrow M(x_0)]$, $\Pr[o \leftarrow M(x_1)]$, $\Pr[o' \leftarrow M'(x'_0)]$, and $\Pr[o' \leftarrow M'(x'_1)]$ all defined on $o, o' \in \mathcal{U}$. Let us split $\mathcal{U}^2 = \mathcal{U} \times \mathcal{U}$ in three sets

$$\begin{aligned}\mathcal{U}_+^2 &= \{(o,o') \mid \forall (o,o') \in \mathcal{U}^2, \forall i \in \{0,1\} : \Pr[o = M(x_i)] \neq 0 \wedge \Pr[o' = M'(x'_i)] \neq 0\} \\ \mathcal{U}_0^2 &= \{(o,o') \mid \forall (o,o') \in \mathcal{U}^2, \forall i \in \{0,1\} : \Pr[o = M(x_i)] = 0 \wedge \Pr[o' = M'(x'_i)] = 0\} \\ \mathcal{U}_\infty^2 &= \mathcal{U}^2 \setminus (\mathcal{U}_+^2 \cup \mathcal{U}_0^2) \quad \text{(one to three probabilities are 0)}\end{aligned}$$

Obviously, they are pairwise distinct and contain together all elements in $\mathcal{U}^2 = \mathcal{U}_+^2 \cup \mathcal{U}_\infty^2 \cup \mathcal{U}_0^2$. Therefore, this proof examines these sets separately: first, the set $\mathcal{U}_+^2$ (leading to the convolution property), second $\mathcal{U}_\infty^2$ (for $\omega(\infty)$ and partly $\omega(\text{-}\infty)$), and last $\mathcal{U}_0^2$ (leftover $\omega(\text{-}\infty)$).

First, we examine the set $\mathcal{U}_+^2$. This will lead to the convolution property for $y \neq \text{-}\infty, \infty$. As the tree sets are separated in a way that no event (o, o') in $\mathcal{U}_+^2$ has a probability of zero, we do not need to consider $\omega_c(\infty)$ or $\omega_c(\text{-}\infty)$ in this part. For all events $o, o' \in \mathcal{U}_+$, the privacy loss is additive under composition:

$$\begin{aligned}&\forall (o,o) \in \mathcal{U}_+^2 : \\ &\mathcal{L}_{(M(x_0), M'(x'_0))/(M(x_1), M'(x'_1))}(o,o') \\ &\quad = \log\left(\frac{\Pr[(o,o') \leftarrow (M(x_0), M'(x'_0))]}{\Pr[(o,o') \leftarrow (M(x_1), M'(x'_1))]}\right) \\ &\quad = \log\left(\frac{\Pr[o \leftarrow M(x_0)]}{\Pr[o \leftarrow M(x_1)]} \frac{\Pr[o' \leftarrow M'(x'_0)]}{\Pr[o' \leftarrow M'(x'_1)]}\right) \\ &\quad = \log\left(\frac{\Pr[o \leftarrow M(x_0)]}{\Pr[o \leftarrow M(x_1)]}\right) + \log\left(\frac{\Pr[o' \leftarrow M'(x_0)]}{\Pr[o' \leftarrow M'(x'_1)]}\right) \\ &\quad = \mathcal{L}_{M(x_0)/M(x_1)}(o) + \mathcal{L}_{M'(x_1)/M'(x'_1)}(o')\end{aligned}$$

where we have used that M and M' are independent. Let us define

$$\mathcal{Y}_+ = \{y_c \mid y_c = y + y' \ \ \forall y \in \mathcal{Y}, \forall y' \in \mathcal{Y}', y, y' \neq \text{-}\infty, \infty\}$$

As $\mathcal{Y}$ and $\mathcal{Y}'$ are countable, their composition $\mathcal{Y}_+$ is countable as well.

For readability, let us define

$$\begin{aligned}\mathcal{L}_c(o,o') &:= \mathcal{L}_{(M(x_0), M'(x'_0))/(M(x_1), M'(x'_1))}(o,o') \\ \mathcal{L}(o) &:= \mathcal{L}_{M(x_0)/M(x_1)}(o) \\ \mathcal{L}'(o') &:= \mathcal{L}_{M'(x'_0)/M'(x'_1)}(o')\end{aligned}$$

With $y_c \in \mathcal{Y}_+$

$$\begin{aligned}
\omega_c(y_c) &= \sum_{\{(o,o') \mid \mathcal{L}_c(o,o')=y_c\}} \Pr\left[(o,o') \leftarrow (M(x_0), M'(x'_0))\right] \\
&= \sum_{\{(o,o') \mid \mathcal{L}(o)+\mathcal{L}'(o')=y_c\}} \Pr\left[o \leftarrow M(x_0)\right] \cdot \Pr\left[o' \leftarrow M'(x'_0)\right] \\
&= \sum_{\{(y,y') \mid y+y'=y_c\}} \sum_{\{o \mid \mathcal{L}(o)=y\}} \Pr\left[o \leftarrow M(x_0)\right] \cdot \\
&\qquad \left(\sum_{\{o' \mid \mathcal{L}'(o')=y'\}} \Pr\left[o' \leftarrow M'(x'_0)\right] \right) \\
&= \sum_{\{(y,y') \mid y+y'=y_c\}} \omega(y) \cdot \omega'(y')
\end{aligned}$$

Which is a convolution. We have used that the sums considered converge absolutely; thus, the sum-product is a Cauchy product and thereby the last equality is valid. For the second equality, we have used the independence of M and M'. As there are no events (o, o') in $\mathcal{U}_+^2$ for which one of the four probabilities $\Pr[o \leftarrow M(x_i)]$, $\Pr[o' \leftarrow M'(x_i)]$ with $i \in \{0, 1\}$ equals to zero, we do not need to consider $\omega_c(\infty)$ or $\omega_c(\text{-}\infty)$ here.

In the second part, we prove the composition of $\omega_c(\infty)$ and show that all events in $\mathcal{U}_\infty^2$ which add to $\omega_c(\text{-}\infty)$ are zero. Let us split the set $\mathcal{U}_\infty^2$ in four subsets

$$\begin{aligned}
\mathcal{U}_\infty &= \{o \mid \Pr[o \leftarrow M(x_1)] = 0,\ (o, o') \in \mathcal{U}_\infty^2\} \\
\mathcal{U}'_\infty &= \{o' \mid \Pr[o' \leftarrow M'(x'_1)] = 0,\ (o, o') \in \mathcal{U}_\infty^2\} \\
\mathcal{U}_+ &= \{o_1 \mid \Pr[o \leftarrow M(x_1)] \neq 0,\ (o, o') \in \mathcal{U}_\infty^2\} \\
\mathcal{U}'_+ &= \{o_2 \mid \Pr[o' \leftarrow M'(x'_1)] \neq 0,\ (o, o') \in \mathcal{U}_\infty^2\} \\
\mathcal{U}_\perp^2 &= \mathcal{U}_\infty^2 \setminus (\mathcal{U}_+ \times \mathcal{U}'_\infty) \cup (\mathcal{U}_\infty \times \mathcal{U}'_+) \cup (\mathcal{U}_\infty \times \mathcal{U}'_\infty)
\end{aligned}$$

First, let us argue about $\omega(\text{-}\infty)$: It is always zero as for any corresponding events of $M(x_0)$ have occurrence probability 0 as in Lemma 1. By construction, the sets $\mathcal{U}_+$ and $\mathcal{U}'_+$ contain all events o, o' for which the corresponding $\Pr[o \leftarrow M(x_i)] \neq 0$ and $\Pr[o' \leftarrow M'(x'_i)] \neq 0$ for $i \in \{0, 1\}$. Therefore $\sum_{o \in \mathcal{U}_+} \Pr[o \leftarrow M(x_0)] = 1 - \omega(\infty)$ (analogously for M'). Moreover, all the leftover events in $\mathcal{U}_\perp^2$ have either $\Pr[o \leftarrow M(x_0)] = 0$ or $\Pr[o' \leftarrow M'(x'_0)] = 0$ or both and are captured in the third and fourth statement. By construction, if and only if $(o, o') \in (\mathcal{U}_+ \times \mathcal{U}'_\infty) \cup (\mathcal{U}_\infty \times \mathcal{U}'_+) \cup (\mathcal{U}_\infty \times \mathcal{U}'_\infty)$, then $\Pr[(o, o') \leftarrow (M(x_1), M'(x'_1))] = 0$ and thus the event is within $\omega_c(\infty)$.

$$\begin{aligned}
\omega_c(\infty) &= \sum_{\{(o,o') \mid \Pr\left[(o,o') \leftarrow (M(x_1), M'(x'_1))\right]=0\}} \Pr\left[(o,o') \leftarrow (M(x_0), M'(x'_0))\right] \\
&= \sum_{\{(o,o') \mid \Pr[o \leftarrow M(x_1)] \cdot \Pr\left[o' \leftarrow M'(x'_1)\right]=0\}} \Pr\left[o \leftarrow M(x_0)\right] \cdot \Pr\left[o' \leftarrow M'(x'_0)\right] \\
&= \sum_{(o,o') \in (\mathcal{U}_+ \times \mathcal{U}'_\infty)} \Pr\left[o \leftarrow M(x_0)\right] \cdot \Pr\left[o' \leftarrow M'(x'_0)\right] \\
&\quad + \sum_{(o,o') \in (\mathcal{U}_\infty \times \mathcal{U}'_+)} \Pr\left[o \leftarrow M(x_0)\right] \cdot \Pr\left[o' \leftarrow M'(x'_0)\right] \\
&\quad + \sum_{(o,o') \in (\mathcal{U}_\infty \times \mathcal{U}'_\infty)} \Pr\left[o \leftarrow M(x'_0)\right] \cdot \Pr\left[o' \leftarrow M'(x'_0)\right] \\
&= [1 - \omega(\infty)]\omega'(\infty) + \omega(\infty)[1 - \omega'(\infty)] \\
&\quad + \omega(\infty)\omega'(\infty) \\
&= 1 - [1 - \omega(\infty)][1 - \omega'(\infty)]
\end{aligned}$$

where we have separated the infinite sums as before (independence and Cauchy products) and we have used $\sum_{o \in \mathcal{U}_+} \Pr[o \leftarrow M(x_0)] = 1 - \omega(\infty)$ (and analogously for M').

For the third set $\mathcal{U}_0^2$, the observation that for any $(o, o') \in \mathcal{U}_0^2$ the loss function evaluates to -∞, but any occurrence-probabilities are zero leads to the conclusion that its contribution to any bucket is 0.

To prove the structure of $\mathcal{Y}_c$

$$\mathcal{Y}_c = \{y_c \mid y_c = y + y' \ \ \forall y \in \mathcal{Y}, \forall y' \in \mathcal{Y}'\}$$

Note that for all events in $\mathcal{U}_\infty^2 \setminus \mathcal{U}_\perp^2$ we can set $y = \infty$ and for all events in $\mathcal{U}^2 \setminus (\mathcal{U}_+^2 \cup \mathcal{U}_\infty^2)$ we can set y = -∞. Together with the addition rules in Definition 4.2, it is valid to define $\mathcal{Y}_c = \mathcal{Y}_+ \cup \{\text{-}\infty, \infty\}$. Again, we neglect the set $\mathcal{U}_\perp^2$ and $\mathcal{U}_0^2$ as they do not contribute to the privacy loss distribution. $\mathcal{Y}_c$ is countable as $\mathcal{Y}$ and $\mathcal{Y}'$ and $\{\text{-}\infty, \infty\}$ are countable. this concludes the proof. □

4.2 Dual Privacy Loss Distribution

The ADP definition is symmetric, but the notion of a privacy loss distribution (PLD) of A over B is inherently asymmetric, since $\omega(y)$ is defined by probabilities in A. We show that it is possible to derive the PLD of B over A, the *dual PLD*, directly from the PLD of A over B.

Definition 4.3 (Dual PLD)**.** *Given a mechanism $M : \mathcal{X} \to \mathcal{U}$, for a privacy loss distribution ω with support $\mathcal{Y}$ created by $M(x_0)$ over $M(x_1)$, the* dual privacy loss distribution *(dual PLD) $\curlywedge$ with support $\mathcal{K}$ is defined as*

$$\mathcal{K} = \{\text{-}y \mid y \in \mathcal{Y}\} \tag{7}$$
$$\curlywedge(y) = \omega(\text{-}y)\ e^y \tag{8}$$
$$\curlywedge(\infty) = 1 - \textstyle\sum_{y \in \mathcal{K} \setminus \{\text{-}\infty,\, \infty\}} \curlywedge(y) \tag{9}$$
$$\curlywedge(\text{-}\infty) = 0 \tag{10}$$

Lemma 2. *Given a mechanism $M : \mathcal{X} \to \mathcal{U}$, for a privacy loss distribution ω created by $M(x_0)$ over $M(x_1)$, then the privacy loss distribution created by $M(x_1)$ over $M(x_0)$ is the dual PLD $\curlywedge$ as defined in Definition 4.3.*

Proof. Let us split $\mathcal{U}$ in three sets

$$\mathcal{U}_+ = \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) \in \mathcal{K} \setminus \{\text{-}\infty, \infty\}, \forall o \in \mathcal{U}\}$$
$$\mathcal{U}_\infty = \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) = \infty, \forall o \in \mathcal{U}\}$$
$$\mathcal{U}_0 = \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) = \text{-}\infty, \forall o \in \mathcal{U}\}$$

Note that the sets $\mathcal{U}_+, \mathcal{U}_\infty, \mathcal{U}_0$ are pairwise distinct and $\mathcal{U} = \mathcal{U}_+ \cup \mathcal{U}_\infty \cup \mathcal{U}_0$. We look at each set individually. First, the set $\mathcal{U}_+$: As for for all events $o \in \mathcal{U}_+$ neither $\Pr[o \leftarrow M(x_0)]$ nor $\Pr[o \leftarrow M(x_1)]$ evaluates to zero, we can use the logarithmic nature of the privacy loss $\mathcal{L}_{M(x_0)/M(x_1)}(o) = -\mathcal{L}_{M(x_1)/M(x_0)}(o)$ which gives us

$$\mathcal{K}^+ = \{\text{-}y \mid \forall y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}\}$$

So,

$$\begin{aligned}
\curlywedge(y) &= \sum_{\{o \,|\, \mathcal{L}_{M(x_1)/M(x_0)}(o) = y\}} \Pr[o \leftarrow M(x_1)], \qquad \forall y \in \mathcal{K}^+ \\
&= \sum_{\{x \,|\, \mathcal{L}_{M(x_0)/M(x_1)}(o) = \text{-}y\}} \Pr[o \leftarrow M(x_0)] \cdot \frac{\Pr[x \leftarrow M(x_1)]}{\Pr[o \leftarrow M(x_0)]} \\
&= \sum_{\{x \,|\, \mathcal{L}_{M(x_0)/M(x_1)}(o) = \text{-}y\}} \Pr[o \leftarrow M(x_0)] \cdot e^{\mathcal{L}_{M(x_1)/M(x_0)}(o)} \\
&= \omega(\text{-}y)\ e^y
\end{aligned}$$

There are no events in $\mathcal{U}_+$ which could go into $\curlywedge(\text{-}\infty)$ or $\curlywedge(\infty)$. Next, we look at $\mathcal{U}_0$. We use the fact that $\Pr[o \leftarrow M(x_0)] = 0$ and for all $o \in \mathcal{U}_0$, $\mathcal{L}_{M(x_1)/M(x_0)}(o) = \text{-}\infty$. In this case, according to Lemma 1:

$\omega(-\infty) = 0$. Next, for the set $\mathcal{U}_\infty$, we use

$$\begin{aligned}\omega(\infty) &= \sum_{o \in \mathcal{U}_\infty} \Pr\left[o \leftarrow M(x_1)\right] \\ &= \sum_{o \in \mathcal{U} \setminus \mathcal{U}_0, \mathcal{U}_+} \Pr\left[o \leftarrow M(x_1)\right] \\ &= 1 - \underbrace{\omega(-\infty)}_{=0} - \sum_{y \in \mathcal{L}^+} \omega(y)\end{aligned}$$

Finally, note that the support of ω namely $\mathcal{L}$ coincides with $\mathcal{L}^+ \cup \{-\infty, \infty\}$. This concludes the proof. □

4.3 Approximate Differential Privacy

We first present the definition from the literature and then prove that our PLD-based definition is equivalent.

Definition 4.4 (ADP)**.** *Given a neighboring relation, let $M : \mathcal{X} \to \mathcal{U}$ be a probabilistic mechanism. We say M is (ε, δ)-differentially private (or (ε, δ)-ADP), if for all neighboring $x_0, x_1 \in \mathcal{X}$ and for all sets $S \subseteq \mathcal{U}$ we have*

$$\Pr\left[M(x_0) \in S\right] \leq e^\varepsilon \Pr\left[M(x_1) \in S\right] + \delta.$$

We say that δ is tight *for ε if there is no $\delta' < \delta$ such that the mechanism is (ε, δ')-ADP. We write $\delta(\varepsilon)$ for this tight δ of an ε. The* ADP-graph *is defined as $(\varepsilon, \delta(\varepsilon))_{\varepsilon \in \mathbb{R}}$.*

The same definition applies if, instead of talking about mechanisms that were based on data universes $\mathcal{X}$, we consider the timing leakage of an algorithm that is based on a secret key, or if we quantify the difficulty of distinguishing two distributions after a single event. For a illustration of ADP on two probability distributions, see Figure 3, following a depiction in [19].

The privacy loss space directly enables us to compute a tight value δ for every value of ε such that (ε, δ)-differential privacy is satisfied. This representation is vital for this work. We connect our definition from above to the definition of tight ADP [19].

Definition 4.5. *For a mechanism $M : \mathcal{X} \to \mathcal{U}$ with neighboring inputs $x_0, x_1 \in \mathcal{X}$ creating a privacy loss distribution ω with support $\mathcal{Y}$ and for $\varepsilon \geq 0$ we define*

$$\begin{aligned}\delta^*_{M(x_0)}(\varepsilon) &= \omega(\infty) + \sum_{y > \varepsilon,\, y \in \mathcal{Y} \setminus \{-\infty, \infty\}} (1 - e^{\varepsilon - y})\, \omega(y) \\ \delta^*_{M(x_1)}(\varepsilon) &= \omega(\infty) + \sum_{y > \varepsilon,\, y \in \mathcal{L} \setminus \{-\infty, \infty\}} (1 - e^{\varepsilon - y})\, \omega(y)\end{aligned}$$

where ω denotes the dual PLD with support $\mathcal{L}$.

We now show that Definitions 4.4 and 4.5 are equivalent.

Lemma 3. *For every probabilistic mechanism $M : \mathcal{X} \to \mathcal{U}$, and for any values $\varepsilon, \delta \geq 0$, M is $(\varepsilon, \delta(\varepsilon))$-tightly differentially private as in Definition 4.4 if and only if, for any two neighboring $x_0, x_1 \in \mathcal{X}^n$, we have $\delta(\varepsilon) = \max\left(\delta^*_{M(x_0)}(\varepsilon), \delta^*_{M(x_1)}(\varepsilon)\right)$ (c.f., Definition 4.5).*

Proof. For simplicity, let us denote $A(o) := \Pr\left[o \leftarrow M(x_0)\right]$ and $B(o) := \Pr\left[o \leftarrow M(x_1)\right]$, and let $\mathcal{L}^{-1}_{A/B}(y) = \{o \,|\, y = \mathcal{L}_{A/B}(o), o \in \mathcal{U}\}$ be the pre-image of y. This proof has two parts. First, we show that

$$\begin{aligned}\sum_{o \in \mathcal{U}} \max(0, A(o) - e^\varepsilon B(o)) &= \omega(\infty) + \sum_{y > \varepsilon,\, y \in \mathcal{Y} \setminus \{-\infty, \infty\}} (1 - e^{\varepsilon - y})\, \omega(y) \\ \sum_{o \in \mathcal{U}} \max(0, B(o) - e^\varepsilon A(o)) &= \omega(\infty) + \sum_{y > \varepsilon,\, y \in \mathcal{L} \setminus \{-\infty, \infty\}} (1 - e^{\varepsilon - y})\ \omega(y)\end{aligned}$$

where ω with support $\mathcal{L}$ denotes the dual distribution of ω. Afterwards, we apply a lemma from prior work to prove the equivalence of the left hand side to tight-ADP.

Let us first consider only the term $\max(0, A(o) - e^{\varepsilon} B(o))$: for any $y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}$ and $\forall o \in \mathcal{L}_{A/B}^{-1}(y)$

$$y = \log \frac{A(o)}{B(o)} \Leftrightarrow B(o) = e^{-y} A(0)$$

This allows us to re-write

$$\begin{aligned}\max(0, A(o) - e^{\varepsilon} B(o)) &= \max(0, (1 - e^{\varepsilon - y}) \cdot A(o)) \\ &= \begin{cases} [1 - e^{\varepsilon - y}] A(o) & \text{if } y > \varepsilon \\ 0 & \text{else} \end{cases}\end{aligned}$$

where we have used the fact that $\forall o \in \mathcal{U}$, $A(0) \geq 0$. After this preparation, we can come to the next step. Keep in mind that the support $\mathcal{Y}$ of ω contains all possible outcomes the loss $\mathcal{L}_{A/B}(o)$ can achieve for all $o \in \mathcal{U}$. Then

$$\begin{aligned}&\sum_{o \in \mathcal{U}} \max(0, A(o) - e^{\varepsilon} B(o)) \\ &\quad = \sum_{o \in \mathcal{L}^{-1}(\infty)} \max(0, A(o) - e^{\varepsilon} \underbrace{B(o)}_{=0}) \\ &\quad + \sum_{o \in \mathcal{L}^{-1}(\text{-}\infty)} \max(0, \overbrace{A(o)}^{=0} - e^{\varepsilon} B(o)) \\ &\quad + \sum_{y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}} \sum_{o \in \mathcal{L}^{-1}(y)} \max(0, [1 - e^{\varepsilon - y}] \cdot A(o)) \\ &\quad = \sum_{o \in \mathcal{L}^{-1}(\infty)} A(o) + \sum_{y > \varepsilon, y \neq \infty} \sum_{o \in \mathcal{L}^{-1}(y)} [1 - e^{\varepsilon - y}] \cdot A(o) \\ &\quad = \omega(\infty) + \sum_{y > \varepsilon,\, y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}} [1 - e^{\varepsilon - y}] \omega(y)\end{aligned}$$

where we have used the definition of $\omega(y) = \sum_{o \in \mathcal{L}^{-1}(y)} A(o)$, the fact that $e^{\varepsilon} > 0$, and $\forall o \in \mathcal{U}$, $A(o), B(o) \geq 0$. By this, we have proven the first equality from the beginning of the proof. The equation for ω is proven identically as we only have to switch $A(x)$ and $B(x)$ to create the dual distribution ω to ω. What is left is the connection to tight-ADP. For this, we use Lemma 1 in [19]:

Lemma (Connection to tight-ADP,[19])**.** *For every ε, two distributions A and B over a finite universe $\mathcal{U}$ are tightly (ε, δ)-ADP with*

$$\begin{aligned}\delta = \max\big(&\textstyle\sum_{o \in U} \max\left(\Pr\left[o \leftarrow A\right] - e^{\varepsilon} \Pr\left[o \leftarrow B\right], 0\right), \\ &\textstyle\sum_{o \in U} \max(\Pr\left[o \leftarrow B\right] - e^{\varepsilon} \Pr\left[o \leftarrow A\right], 0)\big).\end{aligned}$$

which in application directly concludes the proof. □

One immediate corollary is the exact tight-ADP formula for the approximate randomized response mechanism $M_{\varepsilon,\delta}$ (with parameters $\varepsilon \geq 0, \delta \in [0, 1]$), shown to be a worst case mechanism [15] for (ε, δ)-ADP.

Example 1 (ARR)**.** *Approximate Randomized Response for $\varepsilon \geq 0$, $1 \geq \delta \geq 0$, is defined as follows:* $\Pr[o \leftarrow M(x_0)] = p_0(o), \Pr[o \leftarrow M(x_1)] = p_1(o)$ *with*

$$p_0(o) = \begin{cases} \delta & o = 1 \\ \frac{(1-\delta)e^{\varepsilon}}{e^{\varepsilon}+1} & o = 2 \\ \frac{(1-\delta)}{e^{\varepsilon}+1} & o = 3 \\ 0 & o = 4 \end{cases} \quad p_1(o) = \begin{cases} 0 & o = 1 \\ \frac{(1-\delta)}{e^{\varepsilon}+1} & o = 2 \\ \frac{(1-\delta)e^{\varepsilon}}{e^{\varepsilon}+1} & o = 3 \\ \delta & o = 4 \end{cases}$$

Its privacy loss distribution ω can be seen as a shifted binomial distribution, which has a very simple form under convolution. Using Theorem 1 and Lemma 3, for n compositions, we get the exact result

$$\begin{aligned}\delta(\xi) = &\frac{(1-\delta)^n}{(1+e^{\varepsilon})^n} \cdot \sum_{k=\lceil k_{n,\xi} \rceil}^{n} \binom{n}{k} \left[1 - e^{\xi - \varepsilon(2k-n)}\right] e^{\varepsilon(n-k)} \\ &+ [1 - (1-\delta)^n]\end{aligned}$$

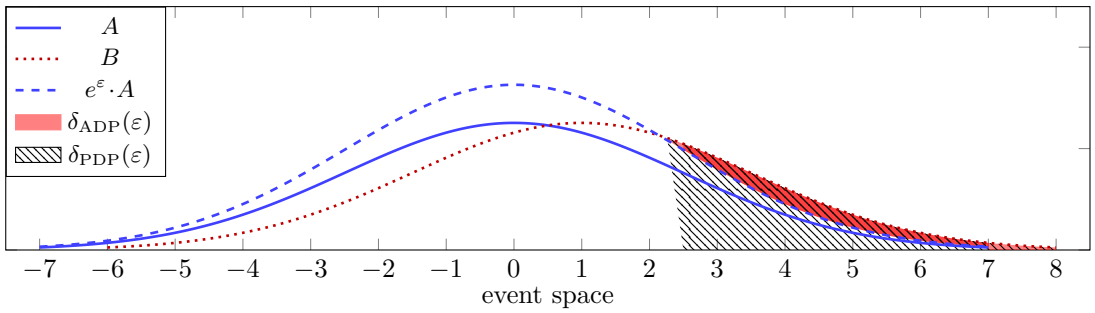

Figure 3: A graphical depiction of the (truncated) Gauss mechanism on two inputs, $A = \mathcal{N}\left(0, \sigma^2\right), B = \mathcal{N}\left(1, \sigma^2\right)$, and of how to compute ADP $\delta_{ADP}(\varepsilon)$ and PDP $\delta_{\text{PDP}}(\varepsilon)$ for a given value ε. Note that $e^{\varepsilon} \cdot A$ is not a probability distribution.

with $\lceil k_{n,\xi} \rceil = \max[0, \min[n, \text{ceil}(\frac{\xi + n\varepsilon}{2\varepsilon})]]$. *For a detailed derivation, see Appendix A.1.*

4.3.1 Equivalence of PLD and ADP-Graphs

We now show that an ADP-graph is as expressive as the privacy loss distribution. For distributions with finite support[3], it is possible to reconstruct the privacy loss distribution from the $(\varepsilon, \delta(\varepsilon))_\varepsilon$ sequence. From Lemma 3 the opposite direction then follows. This is a significant result, as the privacy loss distribution is sufficiently strong for other important privacy notions.

Theorem 2 (Bijection Between ADP and PLD)**.** *Given a set $\mathcal{Y}$ with finite cardinality $|\mathcal{Y}| = k$ (for $k \in \mathbb{N}$), for every PLD ω with support $\mathcal{Y}$, there exists a bijection between the ADP-graph (as in Definition 4.4) and ω.*

Proof. First, define

$$g(y) = \left(1 - e^{-y}\right), y \in \mathbb{R} \quad \text{and} \quad g(\infty) = 1. \tag{11}$$

This function is strong monotonically increasing and therefore there exists maximally one value where $g(y) = 0$, namely $y = 0$, and so $\forall y > 0 \Rightarrow g(y) > 0$.

Second, denote the cardinality $k = |\mathcal{Y}|$, and sort the set $\mathcal{Y}$ in ascending order which is possible as the number of elements in $\mathcal{Y}$ is finite. Denote ε_1 as the smallest element in $\mathcal{Y}$, ε_2 the second, and so forth up to ε_∞ as the largest. Then, for any $i \in \{2, \ldots, n-1, \infty\}$: $\varepsilon_{i-1} < \varepsilon_i$ with $\varepsilon_\infty - \varepsilon_i = \infty$ and $\varepsilon_\infty - \varepsilon_\infty = 0$.

Third, consider the two equations:

$$C = \sum_{y \in \mathcal{Y}} \omega(y) \tag{12}$$

$$\delta_A(\varepsilon_i) = \sum_{\substack{y \geq \varepsilon_i \\ y \in \mathcal{Y}}} \left(1 - e^{\varepsilon_i - y}\right) \omega(y) \quad \text{with} \quad \varepsilon_i \in \mathcal{Y} \tag{13}$$

and noticing that $g(\varepsilon_i - \varepsilon_i) = 0$; they can be written in matrix form:

$$\begin{bmatrix} 1 & 1 & 1 & 1 & \ldots & 1 \\ 0 & g(\varepsilon_2 - \varepsilon_1) & g(\varepsilon_3 - \varepsilon_1) & g(\varepsilon_4 - \varepsilon_1) & \ldots & g(\varepsilon_\infty - \varepsilon_1) \\ 0 & 0 & g(\varepsilon_3 - \varepsilon_2) & g(\varepsilon_4 - \varepsilon_2) & \ldots & g(\varepsilon_\infty - \varepsilon_2) \\ 0 & 0 & 0 & g(\varepsilon_4 - \varepsilon_3) & \ldots & g(\varepsilon_\infty - \varepsilon_3) \\ 0 & 0 & 0 & 0 & \ldots & g(\varepsilon_\infty - \varepsilon_4) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \ldots & g(\varepsilon_\infty - \varepsilon_{n-1}) \end{bmatrix} \cdot \begin{bmatrix} \omega(\varepsilon_1) \\ \omega(\varepsilon_2) \\ \vdots \\ \omega(n-1) \\ \omega(\infty) \end{bmatrix} = \begin{bmatrix} C \\ \delta_A(\varepsilon_1) \\ \delta_A(\varepsilon_2) \\ \delta_A(\varepsilon_3) \\ \delta_A(\varepsilon_4) \\ \vdots \\ \delta_A(\varepsilon_{n-1}) \end{bmatrix}$$

The first matrix is a upper $n \cdot n$-triangle matrix with strictly positive entries on its diagonal ($\forall j \in \{1, \ldots, n\} : m_{jj} > 0$) as $\forall i \in \{2, \ldots, |\mathcal{Y}| - 1, \infty\}$: $\varepsilon_i - \varepsilon_{i-1} > 0$ and $\forall y > 0 : g(y) > 0$. It is a well known result in linear algebra that a upper triangular matrix with non-zero entries on the diagonal is invertible and thereby is a bijection[24]. This leads to the conclusion that any discrete privacy loss distribution ω with finite number of elements in its support has a unique representation $\{\delta_A(\varepsilon_y)\}_{y \in \mathcal{Y}}$ and vice versa. □

[3]In practice this is anyway typically the case due to discretization and finite representations of numbers.

4.4 Probabilistic Differential Privacy

Probabilistic differential privacy [17, 13] is a very intuitive variant of approximate differential privacy (see Figure 3). The main idea is to require that with probability 1-δ pure ε-differential privacy holds. While this definition has a clear semantics and is easy to understand, it is not closed under post-processing [18], which is a crucial property for practical applications; hence, this work concentrates on ADP. Nevertheless, we show that the privacy loss distribution is sufficient for precisely computing PDP bounds.

Definition 4.6 (PDP)**.** *Given a neighboring relation, a mechanism M is (ε, δ)-probabilistically differentially private (PDP), where $\varepsilon \geq 0$ and $\delta \geq 0$, if for all neighboring $x_0, x_1 \in \mathcal{X}$ there are sets $S_0^\delta, S_1^\delta \subseteq [M]$ with $\Pr\left[M(x_0) \in S_0^\delta\right] \leq \delta$ and $\Pr\left[M(x_1) \in S_1^\delta\right] \leq \delta$, s.t., for all sets $S \subseteq [M]$, where $[M]$ is the range of M, the following in-equations hold:*

$$\begin{aligned} \Pr\left[M(x_0) \in S \setminus S_0^\delta\right] &\leq e^\varepsilon \cdot \Pr\left[M(x_1) \in S \setminus S_0^\delta\right] \\ \wedge \Pr\left[M(x_1) \in S \setminus S_1^\delta\right] &\leq e^\varepsilon \cdot \Pr\left[M(x_0) \in S \setminus S_1^\delta\right]. \end{aligned} \tag{14}$$

A mechanism M is tightly (ε, δ)-PDP if δ is minimal for ε, i.e., if for all δ' such that M is (ε, δ')-PDP, $\delta' \geq \delta$.

The conditions of PDP can be directly translated to the privacy loss space as it requires each of tails with $y \geq \varepsilon$ of a PLD ω and its dual PLD $\text{ᵐ}\omega$ to be smaller than δ:

Lemma 4 (Connection to PDP)**.** *Let ω be a privacy loss distribution and $\text{ᵐ}\omega$ its dual PLD. Then*

$$\omega \text{ is } (\varepsilon, \delta)\text{-PDP} \iff \begin{matrix} \sum_{y>\varepsilon, y\in\mathcal{Y}} \omega(y) \leq \delta \\ \sum_{y>\varepsilon, y\in\mathcal{X}} \text{ᵐ}\omega(y) \leq \delta \end{matrix} \tag{15}$$

Proof. Let ω be created by $M(x_0)$ and $M(x_1)$. First, notice that Equation (14) in the PDP definition is equal to the privacy loss function for $i \in \{0, 1\}$:

$$\begin{aligned} &\mathcal{L}_{M(x_i)/M(x_{1-i})}\left(S \setminus S_i^\delta\right) \\ =&\log \frac{\Pr\left[M(x_i) \in S \setminus S_i^\delta\right]}{\Pr\left[M(x_{1-i}) \in S \setminus S_i^\delta\right]} \leq \varepsilon \end{aligned}$$

Let us create two sets

$$\begin{aligned} S_1' &:= \{o \mid \mathcal{L}_{M(x_0)/M(x_1)}(o) > \varepsilon, \forall o \in S_0^\delta\} \\ S_0' &:= \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) > \varepsilon, \forall o \in S_1^\delta\} \end{aligned}$$

As $S_i' \subseteq S_i^\delta \Rightarrow \Pr[M(x_i) \in S_i'] \leq \delta$. Moreover, $\forall o \in S_i^\delta \setminus S_i' : \mathcal{L}_{M(x_i)/M(x_{1-i})}(o) \leq \varepsilon$ by construction. Therefore,

$$\mathcal{L}_{M(x_i)/M(x_{1-i})}(o) \leq \varepsilon \quad \forall o \in S \setminus S_i' = (S \setminus S_i^\delta) \cup (S_i^\delta \setminus S_i') \tag{16}$$

which means that all $o \in S$ with $\mathcal{L}_i(o) > \varepsilon$ are in S_i'. Let us denote ω_0 as the privacy loss distribution created by $M(x_0)$ and $M(x_1)$ with support $\mathcal{Y}_0$ and ω_1 by $M(x_1)$ and $M(x_0)$ with support $\mathcal{Y}_1$ respectively. Then,

$$\begin{aligned} \delta &\geq \Pr[M(x_i) \in S_i'] \\ &\overset{I}{=} \sum_{o \in S_i'} \Pr[o \leftarrow M(x_i)] \\ &\overset{II}{=} \sum_{\{o \mid \mathcal{L}_{M(x_i)/M(x_{1-i})}(o) > \varepsilon,\, o \in S\}} \Pr[o \leftarrow M(x_0)] \\ &\overset{III}{=} \sum_{y>\varepsilon} \omega_i(y) \qquad y \in \mathcal{Y}_i \end{aligned}$$

where we have used independence of elementary events (I), Equation (16) (II) and the privacy distribution definition (III). Next, notice that $\omega_1(y) = \text{ᵐ}\omega(y)$ is the dual distribution to $\omega(y)$. This proves $\omega \Rightarrow$PDP. For the other direction, note that we have only used equalities, that $S_i' \subseteq S_i^\delta$, and that $\mathcal{L}_{M(x_i)/M(x_{1-i})}(S \setminus S_i') \leq \varepsilon \Rightarrow \mathcal{L}_{M(x_i)/M(x_{1-i})}\left(S \setminus S_i^\delta\right) \leq \varepsilon$. Therefore, the statement is proven.

□

4.5 Rényi Differential Privacy & Concentrated Differential Privacy

Recent work introduced novel ADP bounds that are based on the Rényi divergence (the higher moments of the exponentiated privacy loss random variable $e^{\mathcal{L}}$): concentrated DP (CDP) [10, 6], Rényi DP (RDP) [20], and the moments accountant [2]. These bounds were motivated as more comprehensively capturing the privacy guarantees of mechanisms. In fact, the work on concentrated differential privacy can be seen as a direct predecessor of the present work.

The Rényi divergence of two distributions can be directly derived from their PLD. By our Theorem 2, we can hence show that for distributions with finite support RDP and CDP can be determined from the set of all $(\varepsilon, \delta(\varepsilon))$ ADP bounds. We begin with defining the Rényi divergence, RDP and CDP.

Definition 4.7 (Rényi Divergence)**.** *The Rényi divergence $\mathcal{D}_\alpha(M(x_0)|M(x_1))$ with $\alpha > 1$ for a mechanism M and two inputs x_0 and x_1 is defined as*

$$\mathcal{D}_\alpha(M(x_0)|M(x_1)) = \frac{1}{\alpha-1} \log \mathop{\mathbb{E}}_{x\sim M(x_1)} \left(e^{\mathcal{L}_{M(x_0)/M(x_1)}}\right)^{\alpha\text{-}1}$$
$$\mathcal{D}_1(M(x_0)|M(x_1)) = \mathop{\mathbb{E}}_{x\sim M(x_0)} \left(\mathcal{L}_{M(x_0)/M(x_1)}\right)$$

Kullback-Leibler Divergence. Computing the Kullback-Leibler (KL) divergence from the PLD is straight-forward. We recall the definition of the KL divergence from $M(x_1)$ to $M(x_0)$ and directly see that it is a natural property of the PLD, if and only if no output $o \in \mathcal{U}$ has an infinite privacy loss:

$$\begin{aligned} &D_{\mathrm{KL}}(M(x_1)||M(x_0)) \\ &= \sum_{o\in\mathcal{U}} \Pr\left[o \leftarrow M\left(x_0\right)\right] \ln\left(\frac{\Pr\left[o \leftarrow M\left(x_0\right)\right]}{\Pr\left[o \leftarrow M\left(x_1\right)\right]}\right) \\ &= \sum_{o\in\mathcal{U}} \Pr\left[o \leftarrow M\left(x_0\right)\right] \mathcal{L}_{M(x_0)/M(x_1)}\left(o\right) = \sum_{y\in\mathcal{Y}} \omega(y)\cdot y, \end{aligned}$$

where ω is the PLD of $M(x_0)$ over $M(x_1)$ with support $\mathcal{Y}$. Analogously, we get the KL divergence from $M(x_1)$ to $M(x_0)$ by the dual PLD $\overline{\omega}$.

Rényi differential privacy directly characterizes the privacy as the sequence of Rényi divergences: $(\alpha, D_\alpha)_\alpha$.

Definition 4.8 (RDP)**.** *A randomized mechanism $M : X \to \mathcal{U}$ has ε Rényi differential privacy of order $\alpha > 1$, or (α, ε)-RDP for short, if for all neighboring $x_0, x_1 \in \mathcal{X}$ we have $\varepsilon = \mathcal{D}_\alpha(M(x_0)|M(x_1))$.*

Rényi differential privacy can be translated to (ε, δ)-PDP by using a logarithmic version of the Markov bound as follows: whenever $(\alpha, D_\alpha)_\alpha$, then also $(\varepsilon, \alpha D_\alpha - \alpha\varepsilon)$-ADP holds. The moments accountant uses the same characterization and proposes $(\varepsilon, \min_\alpha(\alpha D_\alpha - \alpha\varepsilon))$ as ADP bounds (as (ε, δ)-PDP implies (ε, δ)-ADP).

A pair of distributions A, B satisfies (ξ, ρ)-concentrated DP if the Rényi divergence is bounded by an affine linear function: $D_\alpha \leq \xi + \rho\alpha$ (for all $\alpha \geq 0$).

Definition 4.9 (CDP)**.** *A mechanism M satisfies* (ξ, ρ)-concentrated differential privacy *if for all $\alpha > 0$, and all x_0, x_1*

$$\mathcal{D}_\alpha(M(x_0)|M(x_1)) \leq \xi + \rho\alpha. \tag{17}$$

4.5.1 Connection to Rényi Differential Privacy

Rényi differential privacy is closely connected to the moments of the privacy loss distribution. In fact, the α-Rényi-divergence D_α are the (α-1)-log-moments of ω. If the moments ρ_λ of ω are are not growing too fast, $|\rho_\lambda| < cd^\lambda\lambda!$ for a $\lambda > 0$, then we have equivalence. Consequently, for privacy loss distributions on bounded support, we get equivalence always.

Lemma 5 (Equivalence to Rényi-DP)**.** *Let the privacy loss distribution ω be created by $M(x_0)$ over $M(x_1)$. Let $\omega(\infty) = 0$. Then, its λ-log-moment is*

$$m_\lambda = \log\left(\mathop{\mathbb{E}}_{y\sim\omega} y^\lambda\right)^{\frac{1}{\lambda}} = \mathcal{D}_{\lambda+1}(M(x_0)|M(x_1)) \tag{18}$$

with $\lambda > 0$. Moreover, if $\exp(\lambda\cdot|m_\lambda|) < cd^\lambda\lambda!$ for two positive constants c, d, then there exists a bijection between the Rényi-sequence $(\alpha, D_\alpha)_\alpha$ and ω.

Proof. First, let us show the equality between the moments m_λ and the Rényi-Divergence D_α. For simplicity, let us denote $A(o) := \Pr[o \leftarrow M(x_0)]$ and $B(o) := \Pr[o \leftarrow M(x_1)]$. As $\omega(\infty) = 0$, there is no $o \in \mathcal{U}$ where $B(o) = 0$ and $A(x) \neq 0$. Therefore, we can do the following:

$$\begin{aligned}
\log \left(\mathop{\mathbb{E}}_{y \sim \Omega} y^\lambda \right)^{\frac{1}{\lambda}} &= \log \left(\mathop{\mathbb{E}}_{o \sim A} \left(\frac{A(o)}{B(o)} \right)^\lambda \right)^{\frac{1}{\lambda}} \\
&= \frac{1}{\lambda} \log \sum_{o \in \mathcal{X}} A(o) \left(\frac{A(o)}{B(o)} \right)^\lambda \\
&= \frac{1}{\lambda} \log \sum_{o \in \mathcal{X}} B(o) \left(\frac{A(o)}{B(o)} \right)^{\lambda+1} \\
&= \frac{1}{\lambda} \mathop{\mathbb{E}}_{o \sim B} \left(\frac{A(o)}{B(o)} \right)^{\lambda+1} \\
&= \mathcal{D}_{\lambda+1}(A|B)
\end{aligned}$$

For the second statement, we need to prove that ω leads to a unique sequence $(\alpha, D_\alpha)_\alpha$ and vice versa. The first direction $\omega \Rightarrow (\alpha, D_\alpha)_\alpha$ follows immediately by applying the previously proven transformation.

The other direction $\omega \Leftarrow (\alpha, D_\alpha)_\alpha$ is more tricky as there are cases where more than one distribution have the same moments (Hausdorff moments problem). First, let us define $\rho_\lambda := \exp(\lambda \cdot |m_\lambda|)$ and notice that the condition $|\rho_\lambda| < CD^\lambda \lambda!$ is sufficient such that the power series $\sum_{\lambda > 0} \rho_\lambda \frac{r^\lambda}{\lambda!}$ has a positive convergence radius. Now we apply Theorem 30.1 from [4] which states that for a series of moments a unique probability measure exists if the previous power series has a positive convergence radius. As both, the Rényi-sequence $(\alpha, D_\alpha)_\alpha$ and the privacy loss distribution ω are generated by the same mechanisms $M(x_0)$ and $M(x_1)$, the unique probability measure derived from the Rényi-sequence is a valid privacy loss distribution. This concludes the proof. □

Example 2. *For example, the Gauss distribution $\mathcal{N}(0, \sigma^2)$, which is on infinite support, can be derived by their moments ρ_λ[21] uniquely:*

$$|\rho_\lambda| = \left| \mathop{\mathbb{E}}_{y \sim \mathcal{N}(0,\sigma^2)} y^\lambda \right| = \begin{cases} 0 & \lambda \text{ odd} \\ (\lambda-1)!! \cdot \sigma^\lambda & \lambda \text{ even} \end{cases} \leq c d^\lambda \lambda! \tag{19}$$

for two constants c and d, which means by Lemma 5 that there exist a unique probability density function derived from the moments, i.e. $\mathcal{N}(0, \sigma^2)$.

4.6 Markov-ADP Bound

Next, we refine an ADP bound introduced by Abadi et al. [2], called the moments accountant. Our viewpoint with privacy loss distributions enables us to elegantly improve the moments accountant bound, which we named *Markov-ADP* bound.

Theorem 3 (Markov-ADP)**.** *Given a neighboring relation, a mechanism $M : \mathcal{X} \to \mathcal{U}$ with two neighboring inputs $x_0, x_1 \in \mathcal{X}$, and a privacy loss distribution ω with support $\mathcal{Y}$ created by $M(x_0)$ over $M(x_1)$. Let $\mathcal{P}$ be any finite partition of $\mathcal{Y}_n$, $\mathcal{P} = \{y_0, \ldots, y_k\} \subseteq \mathbb{R}^{k+1}$ with $y_i < y_{i+1} \forall i$. Then, after n compositions and $\varepsilon \in \mathcal{P}$, $\varepsilon < y_0$*

$$\delta_{M(x_0)}(\varepsilon) \leq \mathcal{T}(y_k) + \sum_{y_i \geq \varepsilon, y_i \in \mathcal{P}} \left(1 - e^{\varepsilon - y_i}\right) \cdot [\mathcal{T}(y_i) - \mathcal{T}(y_{i-1})] \tag{20}$$

and

$$\mathcal{T}(y) = \min_\lambda \mathop{\mathbb{E}}_{o \sim M(x_0)} \left[e^{\lambda \cdot \log\left(\frac{\Pr[o \leftarrow M(x_0)]}{\Pr[o \leftarrow M(x_1)]} \right)} \right]^n \cdot e^{-\lambda \cdot y} \tag{21}$$

is a upper bound for tight-ADP and smaller or equal to the Rényi-DP bound.

Proof. Let ω_n generated by $M^n(x_0)$ and $M^n(x_1)$ be the distribution ω after n independent self-compositions. The beginning of this proof is inspired by THEOREM 2 of [2] which has already proven the composability of the log moments

$$\alpha_{A^n, B^n}(\lambda) \leq \sum_{i=1}^{n} \alpha_{A,B}(\lambda) = n \cdot \alpha_{A,B}(\lambda) \tag{22}$$

with

$$\alpha_{A,B}(\lambda) = \log \mathop{\mathbb{E}}_{x \sim A} e^{\lambda \log \frac{\Pr[x \leftarrow A]}{\Pr[x \leftarrow B]}} = \log \mathop{\mathbb{E}}_{y \sim \mathcal{Y}_n} e^{\lambda \log \omega(y)} \tag{23}$$

for all $\lambda > 0$. Moreover, by applying Markov's inequality, they have proven for all $\gamma > 0, \lambda > 0$,

$$\Pr_{y\sim\omega}[y \geq \gamma] = \sum_{y\geq\gamma, y\in\mathcal{Y}_n} \omega(y) \;\leq\; \exp(\alpha_{A,B}(\lambda) - \lambda\gamma) \tag{24}$$

From which follows for $\omega_n = (\mathcal{Y}_n, \{\omega(y)_n\})$

$$\sum_{y\geq\gamma, y\in\mathcal{Y}_n} \omega_n(y) \leq \min_{\lambda>0} \exp\left(n \cdot \log \mathop{\mathbb{E}}_{y\sim\mathcal{Y}_n}\left[e^{\lambda \log \omega(y)}\right] - \lambda\gamma\right) \tag{25}$$

$$= \min_{\lambda>0} \frac{\mathbb{E}_{x\sim A}\left[e^{\lambda\cdot\log\left(\frac{A(x)}{B(x)}\right)}\right]^n}{e^{\lambda\cdot\gamma}} \tag{26}$$

$$= \mathcal{T}(\gamma) \tag{27}$$

as this is valid for all λ, the term can be minimized.

W.l.o.g, we can assume $\mathcal{T}(\gamma)$ to be monotone decreasing ($\forall\, \eta > 0,\, \mathcal{T}(\gamma) \geq \mathcal{T}(\gamma + \eta)$), else we just set $\mathcal{T}(\gamma) = \mathcal{T}(\gamma + \eta)$ as a probability mass cannot increase while we reduce the evaluated events.

For every $\varepsilon \in \mathcal{P}$ we have

$$\sum_{y>\varepsilon, y\in\mathcal{Y}_n} (1 - e^{\varepsilon - y})\omega_n(y) \tag{28}$$

$$\leq \mathcal{T}(y_k) + \sum_{y_j\geq\varepsilon, y_j\in\mathcal{P}} (1 - e^{\varepsilon - y_j})\left[\mathcal{T}(y_{j-1}) - \mathcal{T}(y_j)\right]$$

Due to the $(1 - e^{\varepsilon-y})$ terms, which are smaller than 1 (as in the RDP formula), this bound is less or equal to Rényi-DP.

To see why Equation (28) is true, we first investigate properties of ω_n. Note that in general, $a < b \Rightarrow (1 - \frac{1}{e^a}) \leq (1 - \frac{1}{e^b})$. Thus, for every $f \geq 0$ and for all numbers $a_0 \leq a_1$,

$$\sum_{a_0\leq a<a_1} (1 - e^{\varepsilon - a}) f(a) \leq (1 - e^{\varepsilon - a_1}) \sum_{a_0\leq a<a_1} f(a).$$

We split $\mathcal{Y}_n$ into several chunks $\mathcal{P} = \{y_0, \ldots, y_k\} \subseteq \mathbb{R}^{k+1}$ with $y_i < y_{i+1} \forall i$. We define, for $i \in \{0, \ldots, k\}$,

$$\begin{aligned} T''(i) &:= \sum_{y\in\mathcal{Y}_n, y\geq y_i} \omega_n(y) \\ T'(k) &:= T''(k) \\ \text{for i ¡ k:}\quad T'(j) &:= T''(i) - T''(i+1) \end{aligned}$$

We retain for every $y_i \in \mathcal{P}$,

$$\sum_{y\geq y_i, y\in\mathcal{Y}_n} \omega_n(y) = \sum_{j\geq i, j\in\{0,\ldots,k\}} T'(j)$$

We retain for every $y_i \in \mathcal{P}$,

$$\sum_{y_i\leq y, y\in\mathcal{Y}_n} (1 - e^{\varepsilon - y})\omega_n(y)$$

$$= \sum_{j\geq i, j\in\{0,\ldots,k-1\}} \left(\sum_{y_j\leq y<y_{i+1}, y\in\mathcal{Y}_n} (1 - e^{\varepsilon - y})\omega_n(y)\right) + \sum_{y_k\leq y, y\in\mathcal{Y}_n} (1 - e^{\varepsilon - y})\omega_n(y)$$

$$\leq \sum_{j\geq i, j\in\{0,\ldots,k-1\}} \left((1 - e^{\varepsilon - y_{j+1}}) \sum_{y_j\leq y<y_{j+1}, y\in\mathcal{Y}_n} \omega_n(y)\right) + \sum_{y_k\leq y, y\in\mathcal{Y}_n} \omega_n(y)$$

$$= \sum_{j\geq i, j\in\{0,\ldots,k-1\}} \left((1 - e^{\varepsilon - y_{j+1}})T'(j)\right) + T'(k)$$

Claim: For functions f_1, f_2 s.t. for all $i \in \{0, \dots, k\}$: $\sum_{i \le j, j \in \{0,\dots,k\}} f_1(j) \ge 0, \sum_{i \le j, j \in \{0,\dots,k\}} f_2(j) \ge 0$ and $\sum_{i \le j, j \in \{0,\dots,k\}} f_1(j) \le \sum_{i \le j, j \in \{0,\dots,k\}} f_2(j)$ and for all monotonously increasing functions $g \ge 0$,

$$\sum_{i \le j, j \in \{0,\dots,k\}} f_1(j) g(j) \le \sum_{i \le j, j \in \{0,\dots,k\}} f_2(j) g(j)$$

To see why the claim is true, let f_1, f_2, g be functions as above. We know that

$$\begin{aligned} & \sum_{i \le j, j \in \{0,\dots,k\}} f_1(j) \le \sum_{i \le j, j \in \{0,\dots,k\}} f_2(j) \\ \Leftrightarrow & \sum_{i \le j, j \in \{0,\dots,k\}} (f_1(j) - f_2(j)) \le 0 \end{aligned}$$

We start with this statement, but use it for other values of i subsequently.

$$\begin{aligned} & \sum_{i \le j, j \in \{0,\dots,k\}} (f_1(j) - f_2(j)) \le 0 \\ \Rightarrow & g(i) \cdot \sum_{i \le j, j \in \{0,\dots,k\}} (f_1(j) - f_2(j)) \le 0 \\ \Leftrightarrow & g(i)(f_1(i) - f_2(i)) + g(i) \cdot \sum_{i+1 \le j, j \in \{0,\dots,k\}} (f_1(j) - f_2(j)) \le 0 \end{aligned}$$

Since $\sum_{i+1 \le j, j \in \{0,\dots,k\}} (f_1(j) - f_2(j)) \le 0$ and $g(i) \le g(i+1)$ we thus know that

$$g(i)(f_1(i) - f_2(i)) + g(i+1) \cdot \sum_{i+1 \le j, j \in \{0,\dots,k\}} (f_1(j) - f_2(j)) \le 0.$$

We apply this argument repeatedly to yield

$$\begin{aligned} & \sum_{i \le j, j \in \{0,\dots,k\}} g(i) \cdot (f_1(j) - f_2(j)) \le 0 \\ & \Leftrightarrow \sum_{i \le j, j \in \{0,\dots,k\}} g(i) f_1(j) \le \sum_{i \le j, j \in \{0,\dots,k\}} g(i) f_2(j). \end{aligned}$$

This shows the claim. We split the Markov tails $\mathcal{T}$ into $\mathcal{T}'$ analogously to how we have split T into T':

$$\begin{aligned} \mathcal{T}'(k) &:= \mathcal{T}(y_k) \\ \text{for i ¡ k:} \quad \mathcal{T}'(i) &:= \mathcal{T}(y_i) - \mathcal{T}(y_{i+1}) \end{aligned}$$

We again retain for every $y_i \in \mathcal{P}$,

$$\mathcal{T}(y_i) = \sum_{y_j \ge y_i, y_j \in \mathcal{P}} \mathcal{T}'(i)$$

Note that for all $i \in \{0, \dots, k\}$,

$$\sum_{i \le j, j \in \{0,\dots,k\}} T'(j) \le \sum_{i \le j, j \in \{0,\dots,k\}} \mathcal{T}'(j)$$

and that furthermore we are now finally able to apply our property. Given $y_i \in \mathcal{P}$, we get

$$\begin{aligned} & \sum_{y_i \le y, y \in \mathcal{Y}_n} (1 - e^{\varepsilon - y}) \omega_n(y) \\ \le & \sum_{j \ge i, j \in \{0,\dots,k-1\}} \left((1 - e^{\varepsilon - y_{j+1}}) T'(j)\right) + T'(k) \\ \le & \sum_{j \ge i, j \in \{0,\dots,k-1\}} \left((1 - e^{\varepsilon - y_{j+1}}) \mathcal{T}'(j)\right) + \mathcal{T}'(k) \\ \le & \mathcal{T}(y_k) + \sum_{y_j \ge y_i, y_j \in \mathcal{P}} (1 - e^{\varepsilon - y_j}) \left[\mathcal{T}(y_{j-1}) - \mathcal{T}(y_j)\right] \end{aligned}$$

□

This concludes our discussion and review of individual bounds from our perspective on the privacy loss distribution. We see that considering privacy loss distributions is insightful and allows for a comprehensive investigation of a wide variety of privacy aspects. Next we turn to a central insight from our analyses of the privacy loss space: the privacy loss under sequential composition inevitably acquires the shape of a Gauss distribution. This insight then enables us to present a novel, elegant, and expressive characterizations for various mechanisms, which we call privacy loss classes.

5 Privacy Loss Classes

We characterize the approximate behavior of mechanisms the privacy loss distribution of $M(x_0)$ and $M(x_1)$ under sequential composition with a notion of a *privacy loss class*. The privacy loss class has the following property: with an increasing number of compositions n, two different privacy loss distributions (PLD) in the same privacy loss class converge to the same Gauss PLD. As composition excluding distinguishing events translates to convolution of PLDs (Theorem 1), the convergence is implied by the central limit theorem.

The central limit theorem further implies that it suffices to know the variance and mean of the two PLDs before convolution (i.e., composition). The variance and mean of the combined privacy loss distribution (i.e., after convolution) is the sum of the respective values. Thus, by classifying each mechanism by the variance and mean of the respective privacy loss distribution, we can (in the limit) describe the privacy loss of the mechanism and approximately calculate its privacy loss. Section 6 highlights that this description is actually very accurate in practical cases.

This convergence to a Gauss distribution is contrasted by events that have an infinite privacy loss ∞. These events have to be treated differently, so we first strip them away from the distribution. We renormalize the distribution afterwards, but remember the magnitude of the removed events.

Definition 5.1 (Inner Distribution)**.** *The inner distribution $\bar{\omega}$ of a privacy loss distribution ω is the normalized distribution without $\omega(-\infty)$ and $\omega(\infty)$. $\forall y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}$*

$$\bar{\omega}(y) = \Pr_{y\sim\omega}[y \mid y \neq \infty] = \frac{\omega(y)}{1-\omega(\infty)} \tag{29}$$

Therefore, privacy loss classes have three defining elements: the mean and the variance of the *inner distribution*, and the distinguishing events $\omega(\infty)$. In short: privacy loss class $(\mu, \sigma^2, \omega(\infty))$.

For the remaining inner distribution we now define *privacy loss classes*, i.e., functions that describe the privacy loss against which a given privacy loss distribution converges.

Definition 5.2 (Privacy Loss Classes)**.** *A privacy loss distribution ω with support $\mathcal{Y}$ belongs to the $(\mu, \sigma^2, \omega(\infty))$-privacy loss class*

$$\mu = \sum_{y\in\mathcal{Y}\setminus\{\text{-}\infty,\infty\}} y \cdot \bar{\omega}(y) \tag{30}$$

$$\sigma^2 = \sum_{y\in\mathcal{Y}\setminus\{\text{-}\infty,\infty\}} (y-\mu)^2 \cdot \bar{\omega}(y) \tag{31}$$

if $\omega(\infty) \neq 1$, or to the privacy loss class $(0, 0, 1)$ else.

Note that the privacy loss class of every privacy loss distribution coincides (by definition) with the mean and variance of the inner distribution.

All the privacy bounds discussed in the previous section, and most privacy bounds in literature as well, do not consider distinguishing events, i.e, $\omega(\infty) = 0$. The following lemma shows that all of them can be generalized if the bound is considered to constrain only the inner distribution.

Lemma 6 (Bound Conversion)**.** *Let ω be a privacy loss distribution with support $\mathcal{Y}$. If there exists a bound $\mathcal{B}(\gamma)$ on the inner distribution $\bar{\omega}$ for a positive function $g : \mathcal{Y} \to \mathbb{R}$ and for $\gamma \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}$*

$$\sum_{y\geq\gamma} g(y)\,\bar{\omega}(y) \;\leq\; \mathcal{B}(\gamma) \tag{32}$$

then the bound can be expressed for the full distribution:

$$\sum_{y \geq \gamma} g(y)\,\omega(y) \quad \leq \quad \omega(\infty) + [1 - \omega(\infty)]\,\mathcal{B}(\gamma) \tag{33}$$

with $g(\infty) = 1$.

Proof. Let the variables be defined as in the lemma. The statement follows immediately from the definition:

$$\begin{aligned} \sum_{y \geq \gamma} g(y)\,\bar{\omega}(y) &= \frac{1}{1 - \omega(\infty)} \sum_{y \geq \gamma,\, y \neq \infty} g(y)\,\omega(y) \\ &\leq \mathcal{B}(\gamma) \\ \iff \quad \sum_{y \geq \gamma,\, y \neq \infty} g(y)\,\omega(y) &\leq [1 - \omega(\infty)]\,\mathcal{B}(\gamma) \\ \iff \quad \sum_{y \geq \gamma} g(y)\,\omega(y) &\leq \omega(\infty) + [1 - \omega(\infty)]\,\mathcal{B}(\gamma) \end{aligned}$$

with setting $g(\infty) = 1$. □

5.1 The Central Limit Theorem of ADP

We now show our main theoretical result: all privacy loss distributions converge to Gauss privacy loss distributions.

Theorem 4 (The Central Limit Theorem for ADP)**.** *Let* $M(x_0)$ *and* $M(x_1)$ *be two probabilistic mechanisms. Let* ω_1 *be the corresponding privacy loss distribution with support* $\mathcal{Y}_1$ *and privacy loss class* $(\mu, \sigma^2, \omega(\infty))$ *where* μ *and* σ^2 *are finite. Let* ω_n *be the privacy loss distribution with support* $\mathcal{Y}_n$ *after* n *repeated independent compositions of* $M(x_0)$ *and* $M(x_1)$*. Then*

$$\mathcal{Y}_n = \left\{ y \;\middle|\; y = \sum_{i=1}^{n} \tilde{y}_i,\, \forall \tilde{\mathbf{y}} \in \mathcal{Y}^n \right\} \tag{34}$$

$$\omega_n(y) = (\circledast_{i=1}^{n}\, \omega_1)\,[y] \qquad \forall y \in \mathcal{Y}_n \setminus \{\text{-}\infty,\, \infty\} \tag{35}$$

$$\omega_n(\infty) = 1 - [\,1 - \omega_1(\infty)\,]^n \tag{36}$$

$$\omega_n(\text{-}\infty) = 0 \tag{37}$$

with privacy loss class $(n\mu, n\sigma^2, \omega_n(\infty))$ *where* $\circledast$ *denotes convolution. Moreover, if* $\sigma^2 > 0$ *and the third absolute moment of the inner distribution* $\gamma = \mathbb{E}\,|\bar{\omega}_1(y)|^3 < \infty$*, then the inner distribution* $\bar{\omega}_n(y)$ *converges in distribution against a normalized Gauss with*

$$\left| \Pr_{y \sim \omega_n} [y \leq z \,|\, y \neq \infty] - \Phi\left(\frac{z - n\mu}{\sqrt{n}\sigma}\right) \right| < c_u \cdot \frac{\gamma}{\sqrt{n}\sigma^3} \tag{38}$$

$\forall z \in \mathbb{R}$*, or equivalently*

$$\sum_{y \leq z,\, y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}} \omega_n(y) \quad \xrightarrow{d} \quad [\,1 - \omega_n(\infty)\,] \cdot \Phi\left(\frac{z - n\mu}{\sqrt{n}\sigma}\right) \tag{39}$$

where $\Phi(z)$ *denotes the cumulative distribution function of* $\mathcal{N}(0,1)$ *and* $c_u = 0.4748$*.*

Proof. For any $i \in \mathbb{N}$, let ω_i denote the privacy loss distribution with support $\mathcal{Y}_i$ after i compositions and let $(\mu_i, \sigma_i^2, \omega_i(\infty))$ be the corresponding privacy loss class. Note that ω_1 is the original distribution. The proof for this theorem is split into three parts: first, we prove the properties of $\omega_n(y)$ under composition, second we approach the privacy loss class, and as a third, we apply the central limit theorem implied by Berry-Esseen to $\omega_n(y)\,\forall y \in \mathcal{Y}_n \setminus \{\text{-}\infty, \infty\}$ for the Gauss shape. To ease readability we write $\mathbf{y}$ for a vector of elements $y_1, \ldots, y_k$ and we omit the exact declaration $\mathbf{y} = y_1, \ldots, y_k$ if that is clear from the context.

The first part will be proven by induction based on Theorem 1. If we use the same privacy loss distribution $\omega 1$ twice for Theorem 1, we get directly

$$\begin{aligned}
\mathcal{Y}_2 &= \{y \mid y = \tilde{y}_1 + \tilde{y}_2,\, \forall \tilde{\mathbf{y}} \in \mathcal{Y} \times \mathcal{Y}\} \\
\omega_2(y) &= \left(\circledast_{i=1}^{2}\, \omega_1\right)[y] \qquad \forall y \in \mathcal{Y}_2 \setminus \{\text{-}\infty,\, \infty\} \\
\omega_2(\infty) &= 1 - [\,1 - \omega_1(\infty)\,]^2 \\
\omega_2(\text{-}\infty) &= 0
\end{aligned}$$

and as Theorem 1 allows different privacy distributions as input, we use there ω_1 and n independent compositions of ω_1 (creating ω_n). Then by the theorem

$$\begin{aligned}
\mathcal{Y}_{n+1} &= \left\{\hat{y} \mid \hat{y} = y + \sum_{i=1}^{n} \tilde{y}_i + y,\, \forall y \in \mathcal{Y},\, \forall \tilde{\mathbf{y}} \in \mathcal{Y}^n\right\} \\
&= \left\{\hat{y} \mid \hat{y} = \sum_{i=1}^{n+1} \tilde{y}_i,\, \forall \tilde{\mathbf{y}} \in \mathcal{Y}^{n+1}\right\} \\
\omega_{n+1}(y) &= (\omega * \omega_n)[y] \qquad \forall y \in \mathcal{Y}_{n+1} \setminus \{\text{-}\infty,\, \infty\} \\
&= \left(\circledast_{i=1}^{n+1}\, \omega_1\right)[y] \\
\omega_{n+1}(\infty) &= 1 - [\,1 - \omega_n(\infty)\,] \cdot [\,1 - \omega_1(\infty)\,] \\
&= 1 - [\,1 - \omega_1(\infty)]^{n+1} \\
\omega_{n+1}(\text{-}\infty) &= 0
\end{aligned}$$

which is exactly privacy loss distribution after $n+1$ compositions.

For the rest of this proof, we omit $\omega_i(\text{-}\infty)$ as they are always zero and do not cause any problems. For the second part, we use the well known fact that for the inner distribution $\forall y \in \mathcal{Y}_i \setminus \{\text{-}\infty,\, \infty\}$

$$\bar{\omega}_i(y) = \Pr_{y \sim \omega_i}[y \,|\, y \neq \infty] = \frac{\omega_i(y)}{1 - \omega_i(\infty)}$$

which sums up to 1 and with finite mean and variance, we can add mean and variance. For any $i, j \in \mathbb{N}^+$:

$$\begin{aligned}
\mu_{i+j} &= \mathop{\mathbb{E}}_{y \sim \bar{\omega}_{i+j}} y \\
&= \sum_{y \in \mathcal{Y}_{i+j}} \bar{\omega}_{i+j}(y)\, y \\
&= \sum_{y \in \mathcal{Y}_{i+y}} \sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) \cdot \bar{\omega}_j(y - y_i)\, y \\
&\overset{I}{=} \sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) \cdot \sum_{y \in \mathcal{Y}_j} \bar{\omega}_j(y_j)\, (y_i + y_j) \\
&\overset{II}{=} \sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) \cdot (y_i + \mu_j) \\
&\overset{III}{=} \mu_i + \mu_j
\end{aligned}$$

where we have used a variable shift $y \to y_i + y_j$ and the absolute convergence property to re-order the summands (I), and the property $\sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) = 1$ and the definition of μ_i (II, III). Exactly the same way one proves $\sigma^2_{m+l} = \sigma^2_l + \sigma^2_m$, which we omit here. As these μ and σ^2 and $\omega_n(\infty)$ coincide with the definition of the privacy loss class, the theorem statement about the obtained privacy loss class follows directly by induction.

For the third part, we apply Berry-Esseen as stated in definition 11 directly on the normalized distribution $\Pr_{\Omega_n}[y \,|\, y \neq \infty] = \bar{\omega}_n(y)$. All its requirements, namely finite $\gamma, \sigma^2 < \infty$ and IID composition of ω_1, are met by the theorem assumptions. Therefore, $\forall z \in \mathbb{R}$

$$\left|\Pr_{\omega_n}[y \leq z \,|\, y \neq \infty] - \Phi\left(\frac{z - n\mu}{\sqrt{n}\sigma}\right)\right| \leq c_u \frac{\gamma}{\sqrt{n}\sigma^3}$$

The last theorem statement follows by Lemma 6 and by the fact that

$$\Pr_{\omega_n}[y \leq z \,|\, y \neq \infty] = \sum_{y \leq z,\, y \in \mathcal{Y} \setminus \{\text{-}\infty, \infty\}} \omega_n(y)$$

□

It should be mentioned that privacy loss distributions of two different independent mechanisms can converge to a Gauss as well, if they satisfy the so called Lindenberg condition [16]. Informally, the Lindenberg condition requires that no variance of the composing independent distributions dominates the other variances too much. This allows us to combine arbitrary privacy loss distributions while predicting their privacy loss class and therefore their privacy loss as long as they fulfill the Lindenberg condition.

5.2 Generalization to Lebesgue-Integrals

So far we have only considered discrete random variables. Now we extend our analysis to the continuous case, which formally requires us to consider Lebesgue integrals. This will eventually lead us to the analysis of the Gauss mechanism and its exact ADP-bound.

Lemma 7 (Lebesgue-Generalization)**.** *Let the continuous probabilistic mechanism $M : \mathcal{X} \to \widetilde{\mathcal{U}}$ define a Lebesgue–Rokhlin probability space $(\widetilde{\mathcal{U}}, B(\mathbb{R}), \lambda_i)$ where $\widetilde{\mathcal{U}} \in \mathbb{R}$, $B(\widetilde{\mathcal{U}})$ denotes the Borel set and $\lambda_i(\mathcal{O} \subseteq \widetilde{\mathcal{U}}) = \Pr[M(x_i) \in \mathcal{O}]$ is a Lebesgue measure.*

Let the privacy loss function $\mathcal{L} : \widetilde{\mathcal{U}} \to \mathcal{B}(\mathbb{R}) \cup \{\text{-}\infty, \infty\}$ be generalized to sets as follows: $\forall \mathcal{O} \in B(\widetilde{\mathcal{U}})$:

$$\mathcal{L}_{M(x_0)/M(x_1)}(\mathcal{O}) = \{y \,|\, y = \mathcal{L}_{M(x_0)/M(x_1)}(o), \forall o \in \mathcal{O}\} \tag{40}$$

Let $\widetilde{\mathcal{Y}} = \mathcal{L}\left(\widetilde{\mathcal{U}}\right) \setminus \{\text{-}\infty, \infty\}$. Let $\mathcal{L}$ to be integrable in respect to λ. Then we define the pushforward measure $\omega(y)$ for a Lebesgue integrable function g as

$$\forall A \in B(\widetilde{\mathcal{Y}}) : \quad \int_A g \, d\omega(y) := \int_{\mathcal{L}^{-1}(A) \subseteq \widetilde{\mathcal{U}}} g \circ \mathcal{L}(u) \, d\lambda(u) \tag{41}$$

if $g \circ \mathcal{L}$ is integrable with respect to λ. Moreover,

$$\forall y \in B(\widetilde{\mathcal{Y}}) : \quad \omega(y) = \int_{\mathcal{L}^{-1}(y)} d\lambda(u) \tag{42}$$

Additionally, let $\omega(\infty) = \int_{\mathcal{L}^{-1}(\infty)} d\lambda(u)$ and $\omega(\text{-}\infty) = 0$. This together gives us a measure space $(\widetilde{\mathcal{Y}}, B(\widetilde{\mathcal{Y}}), \omega)$ with the finite measure ω, on which we are able to rewrite the previous quantities:

$$\bar{\omega}(y) = \frac{\omega(y)}{1 - \omega(\infty)} \qquad \forall y \in B(\widetilde{\mathcal{Y}}) \tag{43}$$

$$\mu = \int_{\widetilde{\mathcal{Y}}} y \, d\bar{\omega}(y) \tag{44}$$

$$\sigma^2 = \int_{\widetilde{\mathcal{Y}}} (y - \mu)^2 \, d\bar{\omega}(y) \tag{45}$$

$$\delta_{M(x_0)}(\varepsilon) = \omega(\infty) + \int_{[\varepsilon, \infty) \cap \widetilde{\mathcal{Y}}} (1 - e^{\varepsilon - y}) \, d\omega(y) \tag{46}$$

Proof. First, note that $\lambda(u)$ is a σ-finite measure. The push-forward measure we can define as $\widetilde{\mathcal{Y}}$ and $\widetilde{\mathcal{U}}$ are both a subset of $\mathbb{R}$ [5].

Second, as $\omega(y)$ and $\lambda(u)$ are σ-finite measures and $\lambda(u) = 0 \Rightarrow \omega(y) = 0$, the loss random variable is a valid Radon–Nikodym derivative by the Radon-Nikodym theorem[5], and we can write $\omega(y)$.

To the generalized statements: The inner distribution (Equation (43)) is just a multiplication with a positive constant (the normalization) to the measure $\lambda(u)$ which is valid as $\lambda(u) \in \mathbb{R}$ everywhere. The mean and variance are defined as $\forall y \in \mathbb{R} : \bar{\omega}(y) \in \mathbb{R}$, and $\widetilde{\mathcal{Y}} \subseteq \mathbb{R}$ without -∞ and ∞. The derivation of $\delta_{M(x_0)}(\varepsilon)$ identical to Lemma 3 except that the set $\widetilde{\mathcal{Y}}$ does not include the distinguishing events. □

One advantage of the continuous perspective is the ability to derive sometimes a analytic form of the privacy loss distribution directly from the mechanism distribution itself. If the privacy loss variable $\mathcal{L}$ is bijective and derivable, then we can apply integration by substitution.

Lemma 8 (Density Transformation)**.** *Let $M(x_0)$ and $M(x_1)$ be probabilistic mechanisms. Let the probability density $A(u) := \Pr[M(x_0) \leftarrow u]$ be continuous. Let privacy loss distribution ω be created by $M(x_0)$ over $M(x_1)$ with support $\widetilde{\mathcal{Y}}$. For a subset $y \subseteq \widetilde{\mathcal{Y}}$, let $O = \mathcal{L}^{-1}_{M(x_0)/M(x_1)}(y)$. Let $\mathcal{L}_{M(x_0)/M(x_1)}$ be bijective on O, and let the derivative of the inverse $\frac{\partial \mathcal{L}^{-1}}{\partial y}$ be integrable on y. Then*

$$d\omega(y) = A(\,\mathcal{L}^{-1}(y)\,)\,\left(\frac{\partial \mathcal{L}^{-1}}{\partial y}\right)(y)\,dy \tag{47}$$

Proof. As we can evaluate a continuous function $f(y)$ in the bucket space as follows

$$\int_{\widetilde{\mathcal{U}}} f(\,\mathcal{L}(o)\,)\,A(o)\,do$$

we can apply integration by substitution with $\mathcal{L}^{-1}$:

$$\begin{aligned}\int_O f(\mathcal{L}(o))A(o)\,do &= \int_{\mathcal{L}(O)} f(y)\,\underbrace{A(\,\mathcal{L}^{-1}(y)\,)\left(\frac{\partial \mathcal{L}^{-1}}{\partial y}\right)(y)\,dy}_{d\omega(y)} \\ &= \int_{\mathcal{L}(O)} f(y)\,d\omega(y)\end{aligned}$$

□

5.3 ADP for the Gauss Mechanism

We here present a tight analytic formula for $\delta(\varepsilon)$ for the Gauss mechanism. This result is a significant contribution in its own right, as it allows to compute (not just approximate a sound bound) the exact privacy loss and thus ADP for Gauss noise under an arbitrary number of compositions. It is based on the observation that the PLD of a two Gauss mechanisms is a Gauss distribution again and the fact that we can compute exact tight-ADP for Gauss PLD after an arbitrary number of compositions.

Lemma 9 (PLD of Gauss Mechanism)**.** *Let $M(x_0) : \mathcal{X} \to \widetilde{\mathcal{U}}$ and $M(x_1) : \mathcal{X} \to \widetilde{\mathcal{U}}$ be a Gauss mechanism with*

$$M(x) \quad \sim \quad \mathcal{N}\big(x, \sigma^2\big)$$

for $\sigma^2 > 0$, then the privacy loss distribution ω generated by $M(x_0)$ and $M(x_1)$ is a Gauss distribution

$$\omega \quad \sim \quad \mathcal{N}\left(\frac{(x_0-x_1)^2}{2\sigma^2}, \frac{(x_0-x_1)^2}{\sigma^2}\right) \tag{48}$$

and $\omega(\infty) = 0$ for $x_0, x_1 \in \mathcal{X}$ with privacy loss class $\left(\frac{(x_0-x_1)^2}{2\sigma^2}, \frac{(x_0-x_1)^2}{\sigma^2}, 0\right)$.

Proof. Let the variables be defined as in the lemma statement. Let $u \in \widetilde{\mathcal{U}}$. This is a application of lemma 8. The privacy loss function $\mathcal{L} : \widetilde{\mathcal{U}} \to \mathbb{R}$ is

$$\begin{aligned}\mathcal{L}_{\mathcal{N}(x_0,\sigma^2)/\mathcal{N}(x_1,\sigma^2)}(u) &= \log \frac{\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(u-x_0)^2}{2\sigma^2}}}{\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(u-x_1)^2}{2\sigma^2}}} \\ &= \frac{2u(x_0-x_1) - (x_0^2 - x_1^2)}{2\sigma^2}\end{aligned}$$

Note $\forall u \in \widetilde{\mathcal{U}} : \mathcal{L}_{\mathcal{N}(x_0,\sigma^2)/\mathcal{N}(x_1,\sigma^2)}(u) \neq \infty \Rightarrow \omega(\infty) = 0$. Let us denote $y := \mathcal{L}_{\mathcal{N}(x_0,\sigma^2)/\mathcal{N}(x_1,\sigma^2)}(u)$. This function is invertible

$$\mathcal{L}^{-1}(y) = \frac{y\sigma^2 + (x_0^2 - x_1^2)}{2(x_0 - x_1)}$$

and it is derivable. As all involved functions are continuous, we can use Riehmann-integrals. Let $A(u) := \Pr[u \leftarrow M(x_0)]$. Using Lemma 8,

$$\begin{aligned} d\omega(y) &= A(\,\mathcal{L}^{-1}(y)\,)\,\left(\frac{\partial \mathcal{L}^{-1}}{\partial y}\right)(y)\,dy \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\mathcal{L}^{-1}(y)-x_0)^2}{2\sigma^2}}\,\left(\frac{\sigma^2}{2(x_0-x_1)}\right)\,dy \\ &= \frac{1}{\sqrt{2\pi\left[\frac{(x_0-x_1)^2}{\sigma^2}\right]}} \exp\left(-\frac{\left(y-\left[\frac{(x_0-x_1)^2}{2\sigma^2}\right]\right)^2}{2\left[\frac{(x_0-x_1)^2}{\sigma^2}\right]}\right) dy \\ &\sim \mathcal{N}\left(\frac{(x_0-x_1)^2}{2\sigma^2}, \frac{(x_0-x_1)^2}{\sigma^2}\right) \end{aligned}$$

This proves the first statement. In regard of the privacy loss class, note that $\mu = \mathbb{E}_{y\sim\omega}\, y = \frac{(x_0-x_1)^2}{2\sigma^2}$ and $\sigma^2 = \mathbb{E}_{y\sim\omega}\, y^2 = \frac{(x_0-x_1)^2}{\sigma^2}$ can be read out by inspection immediately. With priorly proven $\omega(\infty) = 0$, the privacy loss class of this distribution is $(\frac{(x_0-x_1)^2}{2\sigma^2}, \frac{(x_0-x_1)^2}{\sigma^2}, 0)$. □

It is worth mentioning that Bun and Steinke have already calculated the absolute moments of the Gauss mechanism [6] Lemma 2.4, which implies the result of Lemma 9 as well.

Lemma 10 (Tight ADP for Gauss PLD)**.** *Let ω be a continuous privacy loss distribution in the shape of a Gauss distribution*

$$\omega(y) = \frac{1-\omega(\infty)}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-\mu)^2}{2\sigma^2}}$$

and with privacy loss class $(\mu,\, \sigma^2,\, \omega(\infty))$ for any $0 \le \omega(\infty) \le 1$. Then

$$\delta_{M(x_0)}(\varepsilon) = \omega(\infty) + \frac{1-\omega(\infty)}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon-\mu}{\sqrt{2}\sigma}\right) - e^{\varepsilon-\mu+\frac{\sigma^2}{2}} \operatorname{erfc}\left(\frac{\varepsilon-\mu+\sigma^2}{\sqrt{2}\sigma}\right)\right] \tag{49}$$

where $\operatorname{erfc}(z) = \frac{2}{\pi}\int_z^\infty \exp(-t^2)dt$ is the well studied complementary error function[3].

Proof. First, use the definition

$$\begin{aligned} \delta(\varepsilon) &= \omega(\infty) + \int_\varepsilon^\infty (1-e^{\varepsilon-y})\, d\omega(y) \\ &= \omega(\infty) + [1-\omega(\infty)] \int_\varepsilon^\infty (1-e^{\varepsilon-y})\, \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{(y-\mu)^2}{2\sigma^2}}\, dy \end{aligned}$$

Let us split the integral in two parts and solve them separately.

$$\begin{aligned} \int_\varepsilon^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx &= \int_{\frac{\varepsilon-\mu}{\sqrt{2}\sigma}}^\infty \frac{1}{\sqrt{\pi}} e^{-u^2} du = && \frac{1}{2}\operatorname{erfc}\left(\frac{\varepsilon-\mu}{\sqrt{2}\sigma}\right) \\ \int_\varepsilon^\infty e^{\varepsilon-x} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx & && \\ &= e^\varepsilon \int_\varepsilon^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2-2x\sigma^2}{2\sigma^2}} dx && \\ &= e^\varepsilon \int_\varepsilon^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{-x^2+2x(\mu-\sigma^2)-\mu^2-\sigma^4-2\mu\sigma^2+\sigma^4+2\mu\sigma^2}{2\sigma^2}} dx && \\ &= e^\varepsilon \int_\varepsilon^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{-(x-(\mu-\sigma^2))^2}{2\sigma^2}} e^{\frac{-\sigma^4+2\mu\sigma^2}{2\sigma^2}}\, dx && \\ &= \frac{1}{2}\, e^{\varepsilon+\frac{\sigma^2}{2}-\mu} \operatorname{erfc}\left(\frac{\varepsilon-\mu+\sigma^2}{\sqrt{2}\sigma}\right) && \end{aligned}$$

The lemma statement follows directly by combining everything. □

Note: for numerical stability, the second term in Equation (49) should be evaluated in log-space as

$$e^{\varepsilon+\frac{\sigma^2}{2}-\mu}\cdot\operatorname{erfc}\left(\frac{\varepsilon-\mu+\sigma^2}{\sqrt{2}\sigma}\right)=\exp\left(\operatorname{log\_erfc}\left(\frac{\varepsilon-\mu+\sigma^2}{\sqrt{2}\sigma}\right)+\varepsilon+\frac{\sigma^2}{2}-\mu\right) \tag{50}$$

The GNU Scientific library offers such a function named `gsl_sf_log_erfc`.

Recall from Lemma 9 that for the Gauss mechanism with noise parameter σ and sensitivity $|x_0-x_1|$,[4] the mean μ_{pld} and variance σ^2_{pld} of their respective privacy loss distribution are related: $\mu_{pld}=\sigma^2_{pld}/2=\frac{|x_0-x_1|}{\sigma}^2/2$. Hence, Lemma 10 directly implies the following theorem.

Theorem 5 (Tight ADP for the Gauss Mechanism)**.** *A Gauss mechanism $M:\mathcal{X}\to\widetilde{\mathcal{U}}$ with*

$$M(x)\sim\mathcal{N}\big(x,\sigma^2\big)$$

for $\sigma^2>0$ has for $x_0,x_1\in\mathcal{X}$ after n compositions exactly

$$\delta(\varepsilon)=\frac{1}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon-n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}}\right)-e^{\varepsilon}\cdot\operatorname{erfc}\left(\frac{\varepsilon+n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}}\right)\right] \tag{51}$$

with $\sigma_{\text{pld}}=\frac{|x_0-x_1|}{\sigma}$ and $\mu_{\text{pld}}=\sigma^2_{\text{pld}}/2$ and is tightly $(\varepsilon,\delta(\varepsilon))$-ADP as in Definition 4.4.

Proof. Let the variables be as in the theorem statement. By Lemma 9 we know that the probabilistic mechanisms $M(x_0)$ and $M(x_1)$ are again depicted as a Gauss in the privacy loss space with the privacy loss class $(\frac{(x_0-x_1)^2}{2\sigma^2},\frac{(x_0-x_1)^2}{\sigma^2},0)$. It is well known that a convolution of two Gauss is a Gauss again

$$\mathcal{N}\big(y_1,\sigma^2\big)+\mathcal{N}\big(y_2,\sigma^2\big)=\mathcal{N}\big(y_1+y_2,2\sigma^2\big)\qquad y_1,y_2\in\mathbb{R}$$

which can be generalized to

$$\bigoplus_{i=0}^{n}\mathcal{N}\big(y,\sigma^2\big)=\mathcal{N}\big(ny,n\sigma^2\big)$$

Applying this, the privacy loss class, and Theorem 4 (CLT for differential privacy) gives us after n composition a Gauss shaped probability distribution ω_n created by $M^n(x_0)$ and $M^n(x_1)$ with

$$\omega_n\sim\mathcal{N}\left(n\frac{(x_0-x_1)^2}{2\sigma^2},n\frac{(x_0-x_1)^2}{\sigma^2}\right)$$

and privacy loss class $(n\frac{(x_0-x_1)^2}{2\sigma^2},n\frac{(x_0-x_1)^2}{\sigma^2},0)$. As ω_n is Gauss shaped, we can apply Lemma 10 and get $\delta_{M^n(x_0)}(\varepsilon)$

$$=\frac{1}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon-n\frac{(x_0-x_1)^2}{2\sigma^2}}{\sqrt{2n}\frac{|x_0-x_1|}{\sigma}}\right)-e^{\varepsilon}\cdot\operatorname{erfc}\left(\frac{\varepsilon+n\frac{(x_0-x_1)^2}{2\sigma^2}}{\sqrt{2n}\frac{|x_0-x_1|}{\sigma}}\right)\right]$$

where we assumed the root of the variance in the privacy loss space to be positive. As the discussed problem is symmetric in $M(x_0)$ and $M(x_1)$, we get

$$\delta_{M^n(x_0)}(\varepsilon)=\delta_{M^n(x_1)}(\varepsilon)$$

which results according to Lemma 3 in tight $(\varepsilon,\delta(\varepsilon))$-ADP. □

Corollary 1 (Tight PDP for the Gauss Mechanism)**.** *A Gauss mechanism $M:\mathcal{X}\to\widetilde{\mathcal{U}}$ with*

$$M(x)\sim\mathcal{N}\big(x,\sigma^2\big)$$

for $\sigma^2>0$ has for $x_0,x_1\in\mathcal{X}$ after n compositions exactly

$$\delta_{PDP}(\varepsilon)=\frac{1}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon-n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}}\right)\right] \tag{52}$$

with $\sigma_{\text{pld}}=\frac{|x_0-x_1|}{\sigma}$ and $\mu_{\text{pld}}=\sigma^2_{\text{pld}}/2$ and is tightly $(\varepsilon,\delta_{PDP}(\varepsilon))$-PDP as in Definition 4.6.

Proof. The corollary follows analogously to Theorem 5 by considering only the tail bound; we simply do not subtract the terms within the tail that are captured by e^{ε} bound. □

[4] As discussed in Section 2, for a large class of real-valued queries, the sensitivity can be represented as $|x_0-x_1|$.

5.4 ADP for Arbitrary Distributions

We extract practical utility from the theoretical observation of privacy loss classes and from our analytical formula for Gauss privacy loss distributions. We provide a generic way to compute a novel ADP bound for arbitrary distributions. First, we recall bounds on the distance between probability distributions under convolution and the Gauss distribution. Second, we combine these bounds with our analytical formula to derive ADP upper and lower bounds.

Lemma 11 (Berry-Esseen and Nagaev Bound,[22]). *Let $X_1, \dots, X_n$ be independent and identically distributed zero mean random variables with*

$$S := X_1 + \cdots + X_n,\ \gamma = \mathbb{E}\left|X_i\right|^3 < \infty,\ \textit{and}\ \sigma := \sqrt{\mathbb{E}\left|X_i\right|^2}$$

then

$$\left|\Pr\left[S > n\sigma z\right] - \Pr\left[Z > z\right]\right| \leq c_u \frac{\gamma}{\sqrt{n}\sigma^3} \qquad \textit{(Berry-Esseen)} \tag{53}$$

$$\left|\Pr\left[S > n\sigma z\right] - \Pr\left[Z > z\right]\right| \leq c_t \frac{\gamma}{\sqrt{n}\sigma^3(1+z^3)} \qquad \textit{(Nagaev)} \tag{54}$$

where $Z \sim \mathcal{N}(0,1)$, $z \geq 0$, $c_u = 0.4748$, $c_t = 25.80$, and $\omega_n(\infty) = 1 - [1-\omega_1(\infty)]^n$.

There exist similar forms of the Berry-Esseen theorem for non-iid random variables with slightly worse $c_u \leq 0.5600$ and $c_t < 31.935$ [22].

Theorem 6 (ADP for Arbitrary PLD). *Let $\varepsilon \leq 0$ and n be arbitrary but fixed. Let ω_1 be a privacy loss distribution created by $M(x_0)$ over $M(x_1)$ with privacy loss class $(\mu, \sigma^2, \omega_1(\infty))$ where $0 < \sigma^2 < \infty$ and finite third absolute moment of the inner distribution $\gamma = \mathbb{E}\left|\bar{\omega}_1(y)\right|^3 < \infty$. Let ω_n be the privacy loss distribution after n independent compositions of ω_1. Let the same be valid for the dual distribution ω_1. Let*

$$\omega_n(\infty) = 1 - [1-\omega_1(\infty)]^n, \quad r_u := c_u \frac{\gamma}{\sigma^3}, \quad r_t(z) := \begin{cases} c_t \frac{\gamma}{\sigma^3(1+z^3)} & \textit{if } z \geq 0 \\ \infty & \textit{else} \end{cases}$$

$$\Delta_\omega := \omega_n(\infty) + \frac{1-\omega_n(\infty)}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon - n\mu}{\sqrt{2n}\sigma}\right) - e^{\varepsilon - n\mu + n\frac{\sigma^2}{2}}\ \operatorname{erfc}\left(\frac{\varepsilon - n\mu + n\sigma^2}{\sqrt{2n}\sigma}\right)\right]$$

$$\beta_\omega := \frac{[1-\omega_n(\infty)]}{\sqrt{n}} \min\left[r_u,\ r_t\left(z = \frac{\varepsilon - n\mu}{\sqrt{n\sigma^2}}\right)\right]$$

with $z \geq 0$, $c_u = 0.4748$ and $c_t = 25.80$. Then

$$\left|\delta_{M^n(x_0)}(\varepsilon) - \Delta_\omega\right| \leq \beta_\omega \tag{55}$$

Moreover, it is $(\varepsilon, \max\left(\Delta_\omega + \beta_\omega,\ \Delta_{\omega} + \beta_{\omega}\right))$-ADP.

Proof. Let the variables be defined as in the theorem statement. Let $\Phi_n(z)$ be the cumulative distribution function of $\mathcal{N}(n\mu, n\sigma)$. We operate by definition of a Lebesgue integrable privacy loss density on a measurable space $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \omega)$. By definition we have $\mu = \mathbb{E}\,\bar{\omega}_1(y)$ and finite $\sigma^2 = \mathbb{E}\left|\bar{\omega}_1(y)\right|^2$. First, we prove that $\forall \varepsilon > 0$ we have

$$\left|\Pr_{y \sim \omega_n}[y \geq \varepsilon \,|\, y \neq \infty] - \Pr\left[Z_n \geq \varepsilon\right]\right| \leq \frac{1}{\sqrt{n}} r_{u/t}\left(z = \frac{\varepsilon - n\mu}{\sqrt{n\sigma^2}}\right)$$

where $Z_n \sim \mathcal{N}(n\mu, n\sigma)$ and where $r_{u/t}(z)$ denotes either r_u or $r_t(z)$. For r_u the statement follows directly from the definition of the Berry-Esseen bound (Lemma 11).

For $r_t(z)$, the insight that $\forall z,\ r_t(z) \leq \infty$ let us use Lemma 11 as well for Nagaev. As $r_u < \infty$, we obtain always a valid bound if we take the minimum for r_u and r_t.

Second, we examine the function

$$g(\varepsilon - y) := (1 - e^{\varepsilon - y}) \qquad \forall y, \varepsilon \in \mathbb{R}.$$

Note that $\forall \varepsilon \in \mathbb{R},\ \forall y \geq \varepsilon,\ 0 \leq g(\varepsilon - y) < 1$.

For simplification, let us denote $\bar{\delta}_\omega(\varepsilon) := \frac{\delta_\omega(\varepsilon)-\omega(\infty)}{1-\omega(\infty)}$. Then

$$\begin{aligned}\left|\bar{\delta}_{\omega_n}(\varepsilon)-\bar{\delta}_{\Phi_n}(\varepsilon)\right| &= \left|\int_\varepsilon^\infty g(\varepsilon-y)d\bar{\omega}_n(y) - \int_\varepsilon^\infty g(\varepsilon-y)d\Phi_n(y)\right| \\ &\overset{I}{\leq} \left|\int_\varepsilon^\infty d\bar{\omega}_n(y) - \int_\varepsilon^\infty d\Phi_n(y)\right| \\ &= \left|\Pr_{y\sim\omega_n}\left[y \geq \varepsilon \,|\, y \neq \infty\right] - \Pr\left[Z_n \geq \varepsilon\right]\right| \\ &\overset{II}{\leq} \frac{1}{\sqrt{n}} r_{u/t}\left(frac\varepsilon - n\mu\sqrt{n\sigma^2}\right)\end{aligned}$$

where we have used the fact that $\forall\varepsilon \in \mathbb{R}, \forall y \geq \varepsilon,\ 0 \leq g(\varepsilon - y) < 1$ (I), and (II) we have proven beforehand.

Now let us include $\omega_n(\infty)$. Theorem 4 gives us immediately $\omega_n(\infty) = 1 - [1-\omega_1(\infty)]^n$. If we multiply by $[1-\omega_n(\infty)]$ and add zero, we get

$$\begin{aligned}&[1-\omega_n(\infty)]\cdot\left|\bar{\delta}_{\omega_n}(\varepsilon)-\bar{\delta}_{\Phi_n}(\varepsilon)\right| \leq \frac{[1-\omega_n(\infty)]}{\sqrt{n}}\cdot r_{u/t}\left(\frac{\varepsilon-n\mu}{\sqrt{n\sigma^2}}\right) \\ &\Leftrightarrow \left|\omega_n(\infty)+[1-\omega_n(\infty)]\cdot\bar{\delta}_{\omega_n}(\varepsilon) - \omega_n(\infty)\right. \\ &\quad \left.+[1-\omega_n(\infty)]\cdot\bar{\delta}_{\Phi_n}(\varepsilon)]\right| \leq \frac{[1-\omega_n(\infty)]}{\sqrt{n}}\cdot r_{u/t}\left(\frac{\varepsilon-n\mu}{\sqrt{n\sigma^2}}\right) \\ &\Leftrightarrow |\delta_{\omega_n}(\varepsilon) - \delta_{\Phi_n}(\varepsilon)| \leq \frac{[1-\omega_n(\infty)]}{\sqrt{n}}\cdot r_{u/t}\left(\frac{\varepsilon-n\mu}{\sqrt{n\sigma^2}}\right)\end{aligned}$$

Together with the definition of $\delta_{\omega_n}(\varepsilon) = \delta_{M^n(x_0)}(\varepsilon)$ and Lemma 10, we can directly obtain

$$\left|\delta_{M^n(x_0)}(\varepsilon) - \Delta^{n,\varepsilon}_{\mu,\sigma^2,\omega_n(\infty)}\right| \leq \beta^{n,\varepsilon}_{\mu,\sigma^2,\omega_n(\infty)}$$

Obviously, this defines an upper bound: $\delta_{M^n(x_0)}(\varepsilon) \leq \Delta^{n,\varepsilon}_{\mu,\sigma^2,\omega_n(\infty)} + \beta^{n,\varepsilon}_{\mu,\sigma^2,\omega_n(\infty)}$. For the ADP property, apply all of the proof before to the dual distribution $\mathcal{m}_n$ with privacy loss class $(\mathcal{H}, \mathcal{o}^2, \mathcal{m}(\infty))$ an then invoke Lemma 3 on the upper bounds. □

6 Evaluation

We apply our derived ADP-bounds to different differentially private mechanisms from the literature. In particular, we compare the Gauss mechanism with the Laplace mechanism and see that the Gauss mechanism has several key advantages.

6.1 Evaluating Our Bounds

We apply our various theoretical results to several mechanism from the literature. For each mechanism, we display a pair of graphs: an ADP-graph after n compositions (left) and the growth of the minimal ε such that $\delta(\varepsilon) \leq 10^{-4}$ over the number of compositions leading up to the number in the left graph; as an exception, for the CoverUp mechanism we display the growth of $\delta(0)$ over the number of compositions. In all figures, the labels are ordered by the values of the respective bounds. We only show bounds that yield reasonable results for the respective graph, e.g., we omit the Berry-Esseen bound in the right graphs where $\delta(\varepsilon) \leq 10^{-4}$ is required. Our figures use approximate zCDP only for the Gauss mechanism, as zCDP requires an to prove that the log-normalized-moments of the privacy loss distribution can be bounded by an affine linear function.

We use the numerical lower bound provided by the privacy buckets [19] as a benchmark in the right graph, but omit it in the left graphs to ease readability. In Figure 7 we additionally omit Rényi DP and Markov-ADP, as computing them lead to numerical problems in the underlying optimization problem.

We discuss each of our bounds separately and refer to different aspects of each of the graphs. We also portray ADP values directly derived from the privacy loss class of the mechanism (i.e., our Gauss formula applied to $(\mu, \sigma^2, \omega(\infty))$ to compare them with the bounds.

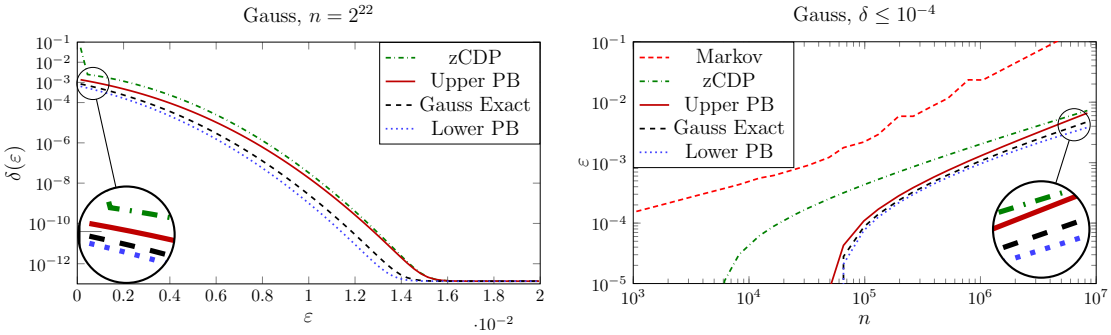

Figure 4: Comparison of Gauss mechanism to known bounds with noise parameter $\sigma^2 = 900000^2$, left: 2^{22} compositions, right: minimal ε values over the number of compositions n for $\delta \leq 10^{-4}$. Comparing the exact Gauss-ADP formula with various bounds. In the right graph, Berry-Esseen bound did not fall into the plotting range and were omitted.

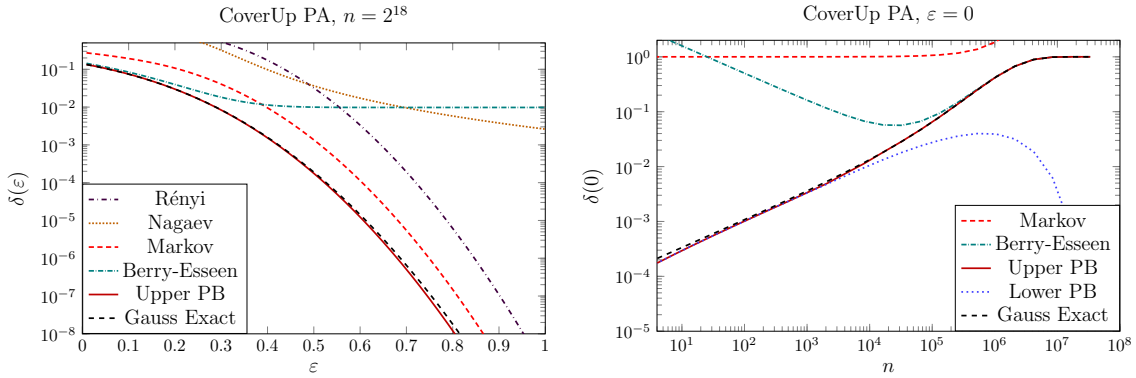

Figure 5: Comparing bounds for the CoverUp measurement data with noise parameter width_of_noise = 100, left: after $n = 2^{18}$ compositions, right: $\delta(0)$ (i.e., $\varepsilon = 0$) over the number of compositions n. In the right graph, the Rényi-DP bound and the Nagaev bound did not reach meaningful values of $\delta(0) \leq 1$.

6.1.1 The Mechanisms in Our Evaluation

We evaluate our bounds with the following mechanisms:

The truncated Gauss mechanism that adds truncated Gauss distributed noise to the result of a computation (see Section 2 for more). Figure 4 compares previous bounds with our exact characterization of the ADP-graph at and up to $n = 2^{22}$ compositions.

Gauss distributed noise applied to two histograms based on CoverUp data[5], which results in a pair of Gauss mixture distributions. CoverUp [26] is a recent work on anonymous communication which measured timing-leakage-histograms of network-level delays for a scenario where a particular browser extension is installed versus a scenario where that browser extension is not installed. Figure 5 displays the ADP-graph after $n = 2^{18}$ compositions and illustrates the growths of $\delta(0)$ (i.e., total variation) over the number of compositions n. The authors argue that deniability ($\varepsilon > 0$) is not reasonable for their scenario; hence, total variation is considered. The graph shows that our theoretical insights lead to promising approaches for deriving valuable bounds.

Abadi et al.'s differentially private stochastic gradient descent (DP-SGD) mechanism [2]; they showed that analyzing the following worst-case distributions suffices: a Gauss distribution $\mathcal{N}(0, \sigma)$ and a Gauss mixture distribution $q\mathcal{N}(0, \sigma) + (1-q)\mathcal{N}(1, \sigma)$ (for some $q \in [0, 1]$ and variation σ^2). Figure 6 displays the ADP graph after and up to $n = 2^{16}$ compositions (i.e., around 600 ANN training epochs).

The truncated Laplace mechanism. We omit the KOV bound [15] as the privacy buckets bounds offer similarly tight bounds and can be computed for a higher number of compositions, which is required for our choice of $n = 2^{20}$ in Figure 7 .

[5] We use the data set Linux periodic loading active from the CoverUp measurements found at [1].

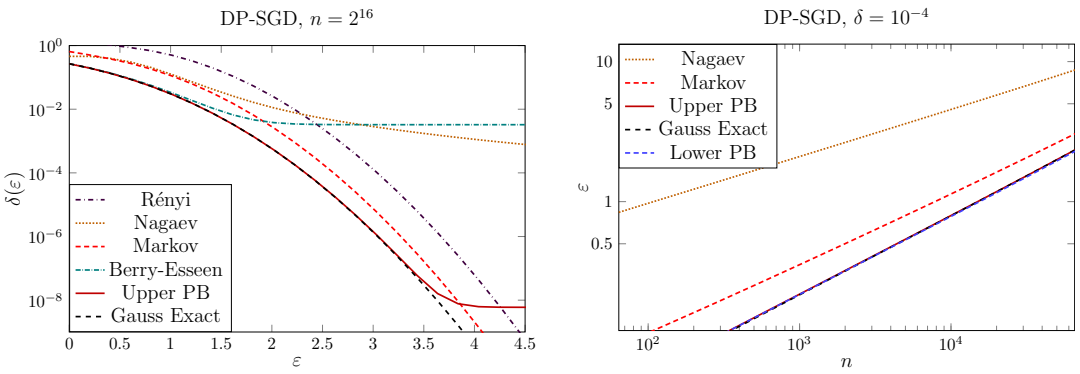

Figure 6: Comparing bounds for differentially private stochastic gradient descent mechanism (DP-SGD) with noise parameters $q = 0.01$ and $\sigma = 4$, left: after $n = 2^{16}$ compositions, right: minimal ε values over the number of compositions n for $\delta \leq 10^{-4}$. In the right graph, the Berry-Esseen bound did not fall into the plotting range and were omitted.

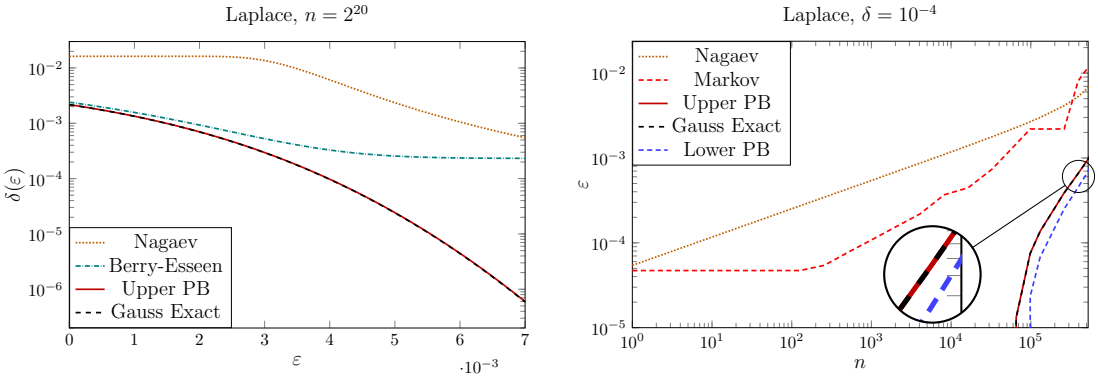

Figure 7: Comparison of Laplace mechanism to known bounds with noise parameter $\lambda = 1.26 \cdot 10^6$, left: 2^{20} compositions, right: minimal ε values over the number of compositions n for $\delta \leq 10^{-4}$. In the right graph, Berry-Esseen bound did not fall into the plotting range and were omitted.

6.1.2 Markov-ADP

In Section 4.6 we improved the Rényi DP bound (or moments accountant) that was previously tailored towards PDP for ADP and called it the Markov-ADP bound. In Figures 4 and 7 both the Markov-ADP bound and the Rényi DP bound are far behind the other bounds; hence, we do not display them. For both mechanisms, this effect is expected: zCDP is tailored to the Gauss mechanism and we have an exact characterization for the Gauss mechanism; for the Laplace mechanism this observation is consistent with previous results about Rényi DP [19].

For CoverUp and DP-SGD in Figures 5 and 6, Markov-ADP clearly outperforms the other bounds, except the numerical privacy buckets. In particular, the Markov-ADP bound outperforms the Rényi DP bound.

6.1.3 Normal Approximation Bounds

We have shown in Theorem 4 and illustrated in Figure 1 that every PLD converges to a Gauss distribution after sufficiently many observations; Theorem 6 provides two separate upper and lower bounds for ADP under n-fold sequential composition, based on the Berry-Esseen and the Nagaev bound respectively.

For CoverUp (Figure 5), the left graph shows that the Berry-Esseen bound is pretty tight until $\delta(\varepsilon) < 10^{-2}$, similarly for DP-SGD (Figure 6) and Laplace (Figure 7) where it is tight almost until 10^{-3}. The reason for this decline becomes apparent if we look at the Berry-Esseen bound: it decreases with a factor of $1/\sqrt{n}$ with the number of convolutions. For a higher number of convolutions, the Berry-Esseen bound provides an even tighter bound. For the Nagaev-based ADP bound, the DP-SGD and the Laplace figures[6] shows that the approach of using tail-bounds (such as the Nagaev Theorem) for normal

[6] We omitted the Nagaev-based bound in CoverUp (Figure 5), since the Nagaev is not tight for small ε values.

approximations is a promising direction.

6.1.4 Convergence to ADP of the Privacy Loss Class

We evaluate the accuracy of ADP derived directly from the privacy loss class of a mechanism $(\mu, \sigma^2, \omega(\infty))$. While this characterization is exact for the Gauss mechanism, it is only approximate for other mechanisms. Figure 4 shows that even the privacy buckets, which we use as a benchmark, diverge from our exact formula for a very large number of compositions. Figure 5 shows that Gauss-ADP is astonishingly accurate in predicting the ADP bounds, already after little more than 10 compositions. This gives evidence that the privacy loss class, already after a few compositions, is a good characterization of the privacy loss of a mechanism. It appears that the imprecision of our normal approximation bounds thus mainly stems from the looseness of these approximation bounds more than from an imprecision of the ADP values calculated from the privacy loss class. We leave it for future work to prove tighter ADP-bounds from this privacy loss class.

6.1.5 zCDP

We only use the approximate zCDP bounds for the Gauss mechanism, as the authors provide explicit bounds for Gauss mechanism. While zCDP provides compelling bounds for higher epsilons, it provides grossly inaccurate values for $\varepsilon = 0$ (i.e., total variation) and very small ε values. This observation is important, as $\varepsilon = 0$ is an important special case: the total variation, $\delta(0)$, is used in the statistical indistinguishability notion. This notion is useful when deniability ($\varepsilon > 0$) is irrelevant and only pure indistinguishability ($\varepsilon = 0$) matters, as, e.g., in the timing analysis of the CoverUp paper [26].

6.2 Gauss vs. Laplace Mechanism

We now compare our results for the Gauss mechanism and the Laplace mechanism. First, we draw a comparison between the privacy loss classes of both mechanisms, showing that they indeed are related. Second, show that the Gauss mechanism has a better variance to privacy trade-off, even if pure DP is preferred, as long as we can tolerate a cryptographically negligible δ.

6.2.1 Comparing the Privacy loss classes

We compare the privacy loss class of a Laplace mechanism with parameter λ (and thus with variance $\sigma^2_{\mathrm{L,ev}} = 2\lambda^2$) with that of a Gauss mechanism with parameter $\sigma_{\mathrm{G,ev}} = \lambda$ (thus half the variance $\sigma^2_{\mathrm{G,ev}} = \lambda^2$). Using our exact formulas for the mean $\mu_{\mathrm{L,pld}}$ and variance $\sigma^2_{\mathrm{G,pld}}$ of the privacy loss class of the Laplace mechanism (Appendix A.3), we can show (Appendix A.4)

$$\mu_{\mathrm{L,pld}} > \mu_{\mathrm{G,pld}} \tag{56}$$

$$\sigma_{\mathrm{L,pld}} \overset{(a)}{>} \sigma_{\mathrm{G,pld}} \tag{57}$$

$$(\mu_{\mathrm{L,pld}}, \sigma_{\mathrm{L,pld}}) \xrightarrow{\frac{|x_0 - x_1|}{\lambda} \to 0} (\mu_{\mathrm{G,pld}}, \sigma_{\mathrm{G,pld}}) \tag{58}$$

where (a) requires $\frac{|x_0 - x_1|}{\lambda} \leq \frac{1}{2}$, which is the case whenever a meaningful degree of privacy is provided. Note that higher values for μ and σ^2 describe a greater privacy loss and result in higher values for $\delta(\varepsilon)$.

As a result, for relevant sensitivity to noise ratios $|x_0 - x_1|/\lambda$, a Gauss mechanism with parameter $\sigma_{ev} = \lambda$ has a strictly, although slightly, better privacy loss class than a Laplace mechanism (resulting in twice the variance, λ^2 vs. $2\lambda^2$). When the sensitivity to noise parameter approaches zero, the privacy loss classes converge. We consider this observation surprising, as the Gauss distribution has much steeper falling tail than the Laplace distribution, which comes with a potential advantage: a truncated Gauss distribution has far less mass in the tail than a Laplace distribution and hence comes with a smaller inherent distinguishing event $\omega(\infty)$.

6.2.2 Sacrificing Pure DP for Gauss?

The Laplace mechanism is a very popular mechanism for achieving differential privacy. The most important argument of the Laplace mechanism over the Gauss mechanism is that the latter cannot achieve pure differential privacy, i.e., $\delta_G(\varepsilon) > 0$ for all ε (cf. Theorem 5 and Corollary 1), while the Laplace

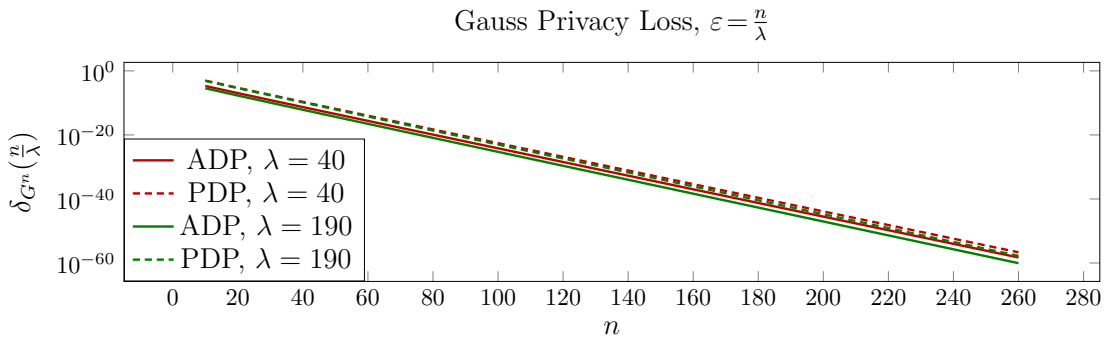

Figure 8: Pure DP vs. both ADP and PDP of a Gauss mechanism: Given a Laplace mechanism with λ, which for n compositions has $\delta_{L^n}(n/\lambda) = 0$ (ADP and PDP), compared to tight ADP-δ_{G^n} and tight PDP-δ_{G^n} of a Gauss mechanism with $\sigma = \lambda$. $\delta_{G^n}(n/\lambda)$ becomes negligible quickly, renders it comparable to Laplace, with half the variance and therefore potentially higher utility.

mechanism can, e.g., with scale factor λ we get $\delta_L(1/\lambda) = 0$. Under n-fold composition, however, the Laplace mechanism can only achieve $\delta_{L^n}(n/\lambda) = 0$.

We compare different Laplace mechanisms with noise parameter λ and with variance $2\lambda^2$ to Gauss mechanisms with half the variance $\sigma^2 = \lambda^2$ and thus a potentially higher utility. Figure 8 illustrates that for $\varepsilon = n/\lambda$ (where $\delta_{L^n}(n/\lambda) = 0$) the $\delta_{G^n}(n/\lambda)$ values fall extremely fast (for ADP and PDP) and for $n = 256$ compositions even negligibly small in the (concrete) cryptographic sense ($< 10^{-50} < 2^{-150}$). These PDP-results can be interpreted as achieving pure differential privacy with $\varepsilon = 256/\lambda$ with probability $1 - 2^{-150}$ with the Gauss mechanism ($\lambda = 40$) after 256 compositions.

6.3 Implementation Considerations

In Figure 4, the upper and lower bounds from privacy buckets' numerical approximation [19] are as expected very close to the exact bound, yet they start to loose tightness for very high amount of compositions. This effect can be credited to numerical errors, memory constraints, and discretization errors. Our exact analytical bound, in contrast, can be directly evaluated for number of compositions and any noise parameters sigma without the need to discretize the Gauss distribution:

$$\delta(\varepsilon) = \frac{1}{2}\left[\operatorname{erfc}\left(\frac{\varepsilon - n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}}\right) - e^{\varepsilon} \cdot \operatorname{erfc}\left(\frac{\varepsilon + n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}}\right)\right]$$

where $\sigma_{\text{pld}} = \frac{|x_0 - x_1|}{\sigma}$ and $\mu_{\text{pld}} = \sigma_{\text{pld}}^2/2$.

We use the `gsl_sf_log_erfc` function from the GNU Scientific Library [14] on the multiplication for numerical robustness. We can achieve very high numerical stability with our implementation by rescaling the privacy loss distribution. Recall that the privacy loss distribution of the Gauss mechanism is again a Gauss distribution with mean μ_{pld} and variance σ_{pld}^2. By computing $\mu_0 := \mu_{\text{pld}}/\mu = 1$ and $\sigma_0 := \sigma_{\text{pld}}/\mu$, we can avoid an overflow in computing the exponential function. Of course, in this case we need to divide the ε of interest by μ as well to achieve the same results.

7 Conclusion and Future Work

In this paper, we have analyzed the privacy loss of privacy-preserving mechanisms and in doing so unified several seemingly different perspectives in the (differential) privacy literature, including Rényi-DP, the moments accountant, (z)CDP, ADP and PDP. We have shown that the non-adaptive composition of mechanisms corresponds to the convolution of their respective privacy loss distribution. Consequently, we were able to apply the central limit theorem, showing that every privacy loss distribution converges to a Gauss distribution under composition.

Our theoretical analysis shifts the perspective of creating new privacy analysis from finding seemingly unrelated upper bounds of the privacy loss a mechanism has to a much clearer procedure: we now understand that all mechanisms fall into straight-forward privacy loss classes and that their privacy loss

converges towards a Gauss distribution either fast or more slowly. Consequently, we see the following directions for future work: 1) finding a tight or even exact embedding of novel mechanisms into their respective privacy loss classes, except for the three widely used mechanisms for which we already give exact formulas: Laplace, Gauss and randomized response, and 2) searching for better convergence bounds, which obviously excludes the Gauss mechanism for which we have an exact formula and hence don't require a bound. One interesting candidate is the staircase mechanism [12], which has been proven to be optimal for differential privacy (without considering composition for ADP) for a class of utility functions. We expect this mechanism to have a privacy loss distribution that is very similar to the Laplace mechanism; hence, the Gauss mechanism would also offer an improved variance over the staircase mechanism. We deem it furthermore particularly interesting to examine which other distributions are closed under convolution and, ultimately, to find out whether the Gauss mechanism is merely a very good mechanism that is better than Laplace, or whether it is provably optimal in the sense that the variance and (particularly) the mean of its privacy loss class is the smallest w.r.t. its initial variance σ? In other words, is there a mechanism of the form $M(x) = q(x) + N$ where N has a smaller variance, but achieves the same or even a better privacy loss class than the Gauss mechanism?

Acknowledgement

This work has been partially supported by the European Commission through H2020-DS-2014-653497 PANORAMIX, the EPSRC Grant EP/M013-286/1, and the Zurich Information Security Center (ZISC).

References

[1] CoverUp Measurement Data. `http://e.mohammadi.eu/paper/coverup_measurements.zip`, 2018. [Online].

[2] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 308–318. ACM, 2016.

[3] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover, 1 edition, 1972.

[4] P. Billingsley. *Probability and measure*. John Wiley & Sons, 2008.

[5] V. I. Bogachev. *Measure theory*, volume 1. Springer Science & Business Media, 2007.

[6] M. Bun and T. Steinke. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Theory of Cryptography (TCC)*, pages 635–658. Springer, 2016.

[7] I. Dinur and K. Nissim. Revealing Information While Preserving Privacy. In *Proceedings of the Twenty-second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 202–210. ACM, 2003.

[8] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *Advances in Cryptology - EUROCRYPT 2006*, pages 486–503. Springer, 2006.

[9] C. Dwork and A. Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[10] C. Dwork and G. N. Rothblum. Concentrated Differential Privacy. *CoRR*, abs/1603.01887, 2016.

[11] U. Erlingsson, V. Pihur, and A. Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, pages 1054–1067, New York, NY, USA, 2014. ACM.

[12] Q. Geng and P. Viswanath. The optimal mechanism in differential privacy. In *2014 IEEE International Symposium on Information Theory (ISIT)*, pages 2371–2375. IEEE, 2014.

[13] M. Götz, A. Machanavajjhala, G. Wang, X. Xiao, and J. Gehrke. Privacy in search logs. *CoRR*, abs/0904.0682, 2009.

[14] B. Gough. *GNU Scientific Library Reference Manual - Third Edition*. Network Theory Ltd., 3rd edition, 2009.

[15] P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017.

[16] J. W. Lindeberg. Eine neue herleitung des exponentialgesetzes in der wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 15(1):211–225, 1922.

[17] A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber. Privacy: Theory meets practice on the map. In *2008 IEEE 24th International Conference on Data Engineering*, pages 277–286, April 2008.

[18] S. Meiser. Approximate and Probabilistic Differential Privacy Definitions. `https://eprint.iacr.org/2018/277`, 2018.

[19] S. Meiser and E. Mohammadi. Tight on Budget? Tight Bounds for r-Fold Approximate Differential Privacy. In *Proceedings of the 25th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2018.

[20] I. Mironov. Renyi Differential Privacy. In *Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

[21] A. Papoulis. Probability, random variables, and stochastic processes. 1965.

[22] I. Pinelis. Chapter 4 - on the nonuniform berry–esseen bound. In I. Pinelis, editor, *Inequalities and Extremal Problems in Probability and Statistics*, pages 103 – 138. Academic Press, 2017.

[23] E. Regulation. European data protection regulation (GDPR). *Off J Eur Union*, L119:1–88, 4 May 2016.

[24] S. Roman. Advanced linear algebra, volume 135 of graduate texts in mathematics, 2008.

[25] I. Roy, S. T. V. Setty, A. Kilzer, V. Shmatikov, and E. Witchel. Airavat: Security and privacy for mapreduce. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, NSDI'10, pages 20–20, 2010.

[26] D. Sommer, A. Dhar, L. Malitsa, E. Mohammadi, D. Ronzani, and S. Capkun. Anonymous Communication for Messengers via "Forced" Participation. Technical report, available under `https://eprint.iacr.org/2017/191`, 2017.

[27] J. van den Hooff, D. Lazar, M. Zaharia, and N. Zeldovich. Vuvuzela: Scalable private messaging resistant to traffic analysis. In *Proceedings of the 25th Symposium on Operating Systems Principles (SOSP)*, pages 137–152. ACM, 2015.

A Examples

This section lists common examples. The use of the symbols are according to their definition earlier.

A.1 Approximate Randomized Response

$$\mathcal{U} = \{1, 2, 3, 4\} \tag{59}$$

$$\Pr\left[o \leftarrow M(x_0)\right] = \begin{cases} \delta & o = 1 \\ \frac{(1-\delta)e^{\varepsilon}}{e^{\varepsilon}+1} & o = 2 \\ \frac{(1-\delta)}{e^{\varepsilon}+1} & o = 3 \\ 0 & o = 4 \end{cases} \tag{60}$$

$$\Pr\left[o \leftarrow M(x_1)\right] = \begin{cases} 0 & o = 1 \\ \frac{(1-\delta)}{e^{\varepsilon}+1} & o = 2 \\ \frac{(1-\delta)e^{\varepsilon}}{e^{\varepsilon}+1} & o = 3 \\ \delta & o = 4 \end{cases} \tag{61}$$

$$\mathcal{L}_{M(x_0)/M(x_1)}(o) = \begin{cases} \infty & o = 1 \\ \varepsilon & o = 2 \\ -\varepsilon & o = 3 \\ -\infty & o = 4 \end{cases} \tag{62}$$

$$\omega(y) = \begin{cases} \delta & y = \infty \\ \frac{(1-\delta)e^{\varepsilon}}{e^{\varepsilon}+1} & y = \varepsilon \\ \frac{(1-\delta)}{e^{\varepsilon}+1} & y = -\varepsilon \\ 0 & y = -\infty \end{cases} \tag{63}$$

$$\bar{\omega}(y) = \begin{cases} \frac{e^{\varepsilon}}{e^{\varepsilon}+1} & y = \varepsilon \\ \frac{1}{e^{\varepsilon}+1} & y = -\varepsilon \end{cases} \tag{64}$$

$$= \binom{2}{k} p^{k_{2,y}} (1-p)^{n-k_{2,y}} \tag{65}$$

$$= B_k(n = 2, p) \tag{66}$$

$$\text{with} \qquad k_{2,y} = \frac{y + 2\varepsilon}{2\varepsilon}, \quad p = \frac{1}{e^{\varepsilon}+1} \tag{67}$$

where $B_k(n, p)$ denotes the k^{th} summand of a Binomial distribution $B(n, p)$ with n trials and success probability p. Moreover, the convolution of two Binomial distributions is again a Binomial:

$$B(n, p) + B(m, p) = B(n + m, p) \tag{68}$$

which gives us after n compositions by Theorem 4

$$\bar{\omega}_n(y) = \binom{n}{k} p^{k_{n,y}} (1-p)^{n-k_{n,y}} \tag{69}$$

$$\omega_n(y) = \begin{cases} 0 & y = -\infty \\ 1-(1-\delta)^n & y = \infty \\ (1-\delta)^n \cdot \bar{\omega}_n(y) & \text{else} \end{cases} \tag{70}$$

$$\text{with} \qquad k_{n,y} = \frac{y+n\varepsilon}{2\varepsilon}, \quad p = \frac{1}{e^\varepsilon + 1}$$

From there follows immediately by definition

$$\begin{aligned} \delta_A(\xi) = [1-\omega_n(\infty)] \cdot \sum_{k=\lceil k_{n,\xi} \rfloor}^{n} \left[1-e^{\xi-y(k)}\right] \binom{n}{k} p^k (1-p)^{n-k} \\ + \omega_n(\infty) \end{aligned} \tag{71}$$

$$\begin{aligned} = \frac{(1-\delta)^n}{(1+e^\varepsilon)^n} \cdot \sum_{k=\lceil k_{n,\xi} \rceil}^{n} \binom{n}{k} \left[1-e^{\xi-\varepsilon(2k-n)}\right] e^{\varepsilon(n-k)} \\ + [1-(1-\delta)^n] \end{aligned} \tag{72}$$

with $y(k) = \varepsilon(2k-n)$ and $\lceil\rceil$ rounds up to nearest integer. Obviously, $k_{n,\xi}$ has to stay between 0 and n. Due to symmetry reasons, $\delta(\varepsilon)$ of the dual PLD is identical.

A.2 Gauss Mechanism

$$\widetilde{\mathcal{U}} = \mathbb{R} \tag{73}$$

$$\Pr\left[o \leftarrow M(x_0)\right] = \frac{e^{-\frac{(x-x_0)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} \tag{74}$$

$$\Pr\left[o \leftarrow M(x_1)\right] = \frac{e^{-\frac{(x-x_1)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} \tag{75}$$

for simplicity $x_0 < x_1$

$$\mathcal{L}_{M(x_0)/M(x_1)}(o) = \frac{(x_1-x_0)(x_0+x_1-2x)}{2\sigma^2} \tag{76}$$

$$\omega(\infty) = 0 \tag{77}$$

$$\omega(\text{-}\infty) = 0 \tag{78}$$

$$\omega(y) = \bar{\omega}(y) \tag{79}$$

$$= \frac{1}{\sqrt{2\pi\left[\frac{(x_0-x_1)^2}{\sigma^2}\right]}} \exp\left(-\frac{\left(y-\left[\frac{(x_0-x_1)^2}{2\sigma^2}\right]\right)^2}{2\left[\frac{(x_0-x_1)^2}{\sigma^2}\right]}\right) \tag{80}$$

$$\mu = \frac{(x_0-x_1)^2}{2\sigma^2} \tag{81}$$

$$\sigma^2 = \frac{(x_0-x_1)^2}{\sigma^2} \tag{82}$$

For $\delta_{M(x_0)}(\varepsilon)$ we refer to Theorem 5. Due to symmetry: $\delta_{M(x_0)}(\varepsilon) = \delta_{M(x_1)}(\varepsilon)$

A.3 Laplace Mechanism

$$\widetilde{\mathcal{U}} = \mathbb{R} \tag{83}$$

$$\Pr\left[o \leftarrow M(x_0)\right] = \frac{1}{2b} e^{\frac{|o-x_0|}{b}} \tag{84}$$

$$\Pr\left[o \leftarrow M(x_1)\right] = \frac{1}{2b} e^{\frac{|o-x_1|}{b}} \quad \text{with } x_0 < x_1 \tag{85}$$

$$\mathcal{L}_{M(x_0)/M(x_1)}(o) = \begin{cases} \frac{x_0-x_1}{b} & o \le x_0 \\ \frac{x_0+x_1-2o}{b} & x_0 \le o \le x_1 \\ \frac{x_1-x_0}{b} & o \le x_1 \end{cases} \tag{86}$$

$$\tag{87}$$

Let us denote $A(o) := \Pr\left[o \leftarrow M(x_0)\right]$

$$\mu = \int_{-\infty}^{\infty} \mathcal{L}(o)\, A(o)\, do \tag{88}$$

$$= e^{\frac{x_0-x_1}{b}} - \frac{b + x_0 - x_1}{b} \tag{89}$$

$$\sigma^2 = \int_{-\infty}^{\infty} \left(\mathcal{L}(o) - \mu\right)^2 A(o)\, do \tag{90}$$

$$= 3 - \frac{2e^{\frac{x_0-x_1}{b}}(b - 2(x_0 - x_1))}{b} - e^{\frac{2(x_0-x_1)}{b}} \tag{91}$$

$$\omega(y) = \int_{\mathcal{L}^{-1}(y)} A(o)\, do \tag{92}$$

$$= \begin{cases} \int_{x_1}^{\infty} A(x) & y = \frac{x_0-x_1}{b} \\ A(x(y))\, \frac{\partial x}{\partial y} & \frac{x_0-x_1}{b} < y \le \frac{x_1-x_0}{b} \\ \int_{-\infty}^{0} A(x) & y = \frac{x_1-x_0}{b} \\ 0 & \text{else} \end{cases} \tag{93}$$

$$= \begin{cases} \frac{1}{2} e^{\frac{x_1-x_0}{b}} & y = \frac{x_0-x_1}{b} \\ \frac{1}{4} e^{\frac{by-x_0+x_1}{2b}}\, dy & \frac{x_0-x_1}{b} < y \le \frac{x_1-x_0}{b} \\ \frac{1}{2} & y = \frac{x_1-x_0}{b} \\ 0 & \text{else} \end{cases} \tag{94}$$

$$\delta_A(\varepsilon) = \int_{\varepsilon}^{\infty} \left(1 - e^{\varepsilon - y}\right)\, d\omega(y) \qquad \text{with } \varepsilon \ge 0 \tag{95}$$

$$= \begin{cases} \frac{1}{2} e^{-\frac{x_0}{b}} \left(e^{\frac{b\varepsilon + x_0}{2b}} - e^{\frac{x_1}{2b}}\right)^2 & \\ \quad + \frac{1}{2}\left(1 - e^{\varepsilon - \frac{x_1-x_0}{b}}\right) & \varepsilon \le \frac{x_1-x_0}{b} \\ 0 & \text{else} \end{cases} \tag{96}$$

Finally, $\delta_{M(x_1)}(\varepsilon) = \delta_{M(x_0)}(\varepsilon) = \delta_A(\varepsilon)$ due to symmetry.

A.4 Gauss vs. Laplace σ^2 derivation

Let $z = \frac{x_0-x_1}{\lambda}$. The two mechanisms (Gauss and Laplace) are symmetric, therefore, w.l.o.g., $z > 0$. The Gauss mechanism has the privacy loss class $(\frac{z^2}{2}, z^2, 0)$ (see Lemma 9). Using our exact formulas of the privacy loss class of the Laplace mechanism (Appendix A.3), we get $\mu_{\text{L,pld}} = e^z - 1 - z$. As e^x can be represented as a Taylor expansion,

$$\sum_{k=i}^{\infty} \frac{x^k}{k!} =: T(i, x)$$

$$\mu_{\text{L,pld}} = T(1, z) - 1 - z = \overbrace{\frac{z^2}{2}}^{\mu_{\text{G,pld}}} + \overbrace{T(3, z)}^{>0} \ge \mu_{\text{G,pld}}$$

Similarly for the variance:

$$
\begin{aligned}
\sigma^2 &= 3-\frac{2e^{\frac{x_0-x_1}{\lambda}}(\lambda-2(x_0-x_1))}{\lambda}-e^{2\frac{(x_0-x_1)}{\lambda}} \\
&= 3-e^z(2-4z)-e^{2z} = 3-e^z(2-4z+e^z) \\
&= 3-(1+z+\sum_{i=2}^{\infty}\frac{z^i}{z!})(2-4z+e^z) \\
&= 3-(2-4z+e^z)-z(2-4z+e^z) \\
&-(\sum_{i=2}^{\infty}\frac{z^i}{z!})(2-4z+e^z) \\
&= 1+4z-e^z-2z+4z^2-ze^z-(\sum_{i=2}^{\infty}\frac{z^i}{z!})(2-4z+e^z) \\
&= 1+2z-(1+z+\sum_{i=2}^{\infty}\frac{z^i}{z!})+4z^2-z(1+z+\sum_{i=2}^{\infty}\frac{z^i}{z!}) \\
&-(\sum_{i=2}^{\infty}\frac{z^i}{z!})(2-4z+e^z)
\end{aligned}
$$

$$
\begin{aligned}
&= -(\sum_{i=2}^{\infty}\frac{z^i}{z!})+3z^2-z(\sum_{i=2}^{\infty}\frac{z^i}{z!}) \\
&-1\cdot(\sum_{i=2}^{\infty}\frac{z^i}{z!})(2-4z+e^z) \\
&= 3z^2-(\sum_{i=2}^{\infty}\frac{z^i}{z!})(3-3z+e^z) \\
&= z^2+2z^2-(\sum_{i=2}^{\infty}\frac{z^i}{z!})\left(3-3z+1+z+\sum_{i=2}^{\infty}\frac{z^i}{z!}\right) \\
&= z^2+2z^2-(\sum_{i=2}^{\infty}\frac{z^i}{z!})\left(4-2z+\sum_{i=2}^{\infty}\frac{z^i}{z!}\right) \\
&= z^2+2z^2-\left(2z^2-z^3+\frac{z^2}{2}\sum_{i=2}^{\infty}\frac{z^i}{z!}\right) \\
&-(\sum_{i=3}^{\infty}\frac{z^i}{z!})\left(4-2z+\sum_{i=2}^{\infty}\frac{z^i}{z!}\right) \\
&= z^2+z^3-\frac{z^2}{2}\sum_{i=2}^{\infty}\frac{z^i}{z!}-(\sum_{i=3}^{\infty}\frac{z^i}{z!})\left(4-2z+\sum_{i=2}^{\infty}\frac{z^i}{z!}\right)
\end{aligned}
$$

$$
\begin{aligned}
&\overset{(a)}{\geq} z^2+z^3-\frac{z^2}{2}\sum_{i=2}^{\infty}\frac{z^i}{z!}-(\sum_{i=3}^{\infty}\frac{z^i}{z!})\cdot(4-z)\\
&= z^2+z^3-\frac{z^4}{4}-\frac{z^5}{12}-\frac{z^2}{2}T(z,4)-4T(z,3)+zT(z,3)\\
&= z^2+\frac{2}{6}z^3-\frac{z^4}{4}-\frac{z^5}{12}-\frac{z^2}{2}T(z,4)-4T(z,4)+\frac{z^4}{6}+\frac{3}{2}zT(z,4)\\
&\geq z^2+\frac{2}{6}z^3-\frac{z^4}{4}-\frac{z^5}{12}-\frac{z^2}{2}T(z,4)-4T(z,4)+\frac{z^4}{4}+\frac{3}{2}zT(z,4)\\
&= z^2+\frac{2}{6}z^3-\frac{z^5}{12}-\frac{z^2}{2}T(z,4)-4T(z,4)+\frac{3}{2}zT(z,4)\\
&= z^2+\frac{2}{6}z^3-\frac{z^6}{24}-\frac{z^2}{2}T(z,5)-4T(z,4)+\frac{z^5}{24}+\frac{3}{2}zT(z,5)\\
&\geq z^2+\frac{2}{6}z^3-\frac{z^2}{2}T(z,5)-4T(z,4)+\frac{3}{2}zT(z,5)\\
&\geq z^2+\frac{2}{6}z^3-\frac{1}{3}z^4-4T(z,5)\\
&\geq z^2+\frac{1}{12}z^3
\end{aligned}
$$

Inequality (a) holds since for $z \leq \frac{1}{2}$, $\frac{1}{2}z \geq T(z,2)$ (and the term we removed is overall positive). Note $z^2 = \sigma_{\mathrm{G,pld}}$.