

David M. Sommer*, Sebastian Meiser, and Esfandiar Mohammadi

Privacy Loss Classes: The Central Limit Theorem in Differential Privacy

Abstract: Quantifying the privacy loss of a privacy-preserving mechanism on potentially sensitive data is a complex and well-researched topic; the de-facto standard for privacy measures are ϵ -differential privacy (DP) and its versatile relaxation (ϵ, δ) -approximate differential privacy (ADP). Recently, novel variants of (A)DP focused on giving tighter privacy bounds under continual observation. In this paper we unify many previous works via the *privacy loss distribution* (PLD) of a mechanism. We show that for non-adaptive mechanisms, the privacy loss under sequential composition undergoes a convolution and will converge to a Gauss distribution (the central limit theorem for DP). We derive several relevant insights: we can now characterize mechanisms by their *privacy loss class*, i.e., by the Gauss distribution to which their PLD converges, which allows us to give novel ADP bounds for mechanisms based on their privacy loss class; we derive *exact* analytical guarantees for the approximate randomized response mechanism and an *exact* analytical and closed formula for the Gauss mechanism, that, given ϵ , calculates δ , s.t., the mechanism is (ϵ, δ) -ADP (not an over-approximating bound).

Keywords: differential privacy, continuous observation, privacy loss, Gauss mechanism, composition

DOI 10.2478/popets-2019-0029

Received 2018-08-31; revised 2018-12-15; accepted 2018-12-16.

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 [24]. Modern systems, however, are increasingly reliant on personal data to provide the expected utility. Hence, privacy and utility are often diametrical, render-

ing perfect solutions impossible but leaving space for solutions that provide privacy under limited usage.

To quantify the privacy of a mechanism, Dwork et al. [10] proposed a strong privacy notion, called (ϵ, δ) -approximate differential privacy (ADP). By now, there is a rich literature on ADP guarantees (e.g., [13, 27]). These privacy guarantees naturally deteriorate under repeated adversarial observation, i.e., continued usage increases (ϵ, δ) to a point where using the mechanism is considered insecure. A tight assessment of this deterioration is essential since loose bounds can lead to underestimating how often a mechanism can be used or to using too much noise. Finding tight bounds is a challenging task and has inspired a rich body of work [2, 8, 12, 16, 20, 21, 26].

The literature contains adaptive sequential composition bounds [16, 26] that are *mechanism-oblivious* in the following sense: given a sequence of ADP parameters $(\epsilon_i, \delta_i)_i$, the adversary may in round j adaptively choose any mechanism that satisfies (ϵ_j, δ_j) -ADP. It has been shown [16, 22] that analyzing the approximate randomized response (ARR) mechanism, i.e., analyzing two *worst-case (output) distributions* parametric solely in a (ϵ_j, δ_j) pair, exactly yields optimal mechanism-oblivious bounds. These results have been used [27] to analyze a mechanism by deriving (ϵ, δ) before composition and then computing an adaptive composition bound.

Often we are interested in quantifying the privacy of a particular mechanism under composition instead of the privacy of adversarially chosen mechanisms. Recent results show that better fitting worst-case distributions can lead to significantly tighter bounds under composition (Concentrated DP [8, 12], moments accountant & Rényi DP [2, 21], and Privacy Buckets [20]).

These methods started to more intensely use the *privacy loss* of a mechanism that has been proposed by a seminal work by Dinur and Nissim [9]. Most of these approaches, however, introduced novel privacy notions derived from characterizing the moments of the privacy loss (CDP, MA, RDP) and only derived loose bounds for well-established privacy notions, such as ADP. A notable exception is the iterative and numerical PB approach that can fall prey to numerical errors, memory limitations, and discretization problems.

*Corresponding Author: David M. Sommer: ETH Zurich, E-mail: david.sommer@inf.ethz.ch

Sebastian Meiser: UCL, E-mail: s.meiser@ucl.ac.uk

Esfandiar Mohammadi: ETH Zurich, E-mail: mohammadi@inf.ethz.ch

1.1 Contribution

This work directly leverages the privacy loss [9] by constructing a probability distribution out of it, the *privacy loss distribution* (PLD). The PLD is similar to the privacy loss random variable used by Dwork and Rothblum [12], but we also consider corner cases where the loss becomes infinite. Our analysis of the PLD, particularly under sequential composition of the mechanism it resulted from, deepens our understanding of privacy-deterioration under composition and yields a list of foundational and practical contributions.

(a) We show that the PLD can be used for deriving the following differential privacy metrics: pure differential privacy (DP), approximate differential privacy (ADP), concentrated differential privacy (CDP), Rényi differential privacy (RDP), and probabilistic differential privacy (PDP). We show that the PLD is a unique canonical representation for RDP and ADP.

(b) We prove that the PLD of any mechanism evolves under independent non-adaptive sequential composition as a convolution and thus, as an application of the central limit theorem, converges to a Gauss distribution, which we call the *privacy loss class* of the mechanism. This Gauss distribution has a variance and mean that directly follows from the variance and mean of the PLD and both values linearly grow under composition. We can extend this insight from non-adaptive composition to some adaptive mechanisms (such as adaptive query-response mechanisms), by finding worst-case distributions and analyzing the PLD of these worst-case distributions. The leniency regarding adaptive choices naturally reduces the tightness of the resulting bounds and is not the main focus of this work.

(c) As a practical application, we use these privacy loss classes to derive analytical formulas for ADP and PDP for the Gauss mechanism and the randomized response mechanism. For both mechanism, we provide tight ADP- and PDP-parameters under sequential composition that can be efficiently calculated.

(d) Our analysis shows that the Gauss mechanism clearly outperforms the one-dimensional Laplace mechanism under composition in terms of a variance to privacy trade-off: A Gauss 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 delta, which in our example ($\sigma = 40$) can be considered negligible even by cryptographic standards: less than 2^{-80} after 128 compositions.

(e) We further use the privacy loss class of a given non-adaptive mechanism to prove upper and lower ADP

bounds for n -fold sequential composition. We apply the Berry-Esseen and Nagaev normal approximation theorems to the privacy loss class and approximate the PLD after n convolutions (for n -fold sequential composition). We, thus, pave the way for future research on tight normal approximation bounds for PLDs, which would result in tight bounds for n -fold sequential composition.

Next, we characterize the PLD under sequential composition (Contribution (b)).

Informal Theorem 4 (The CLT for ADP): *Let M be a mechanism and x_0, x_1 be two inputs yielding the privacy loss distribution ω with finite variance σ^2 and finite mean μ . Then, the privacy loss distribution ω_n of M on x_0 and x_1 after n non-adaptive sequential compositions has variance $n \cdot \sigma^2$ and mean $n \cdot \mu$.*

Moreover, if $\sigma^2 > 0$ and the third absolute moment of ω is finite, then ω_n converges against a Gauss distribution with variance $n \cdot \sigma^2$ and mean $n \cdot \mu$.

2 Overview

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

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

This privacy loss spans a real-valued random variable obtained by sampling $o \sim M(x_0)$ and outputting $\mathcal{L}_{M(x_0)/M(x_1)}(o)$, which in turn defines the privacy loss distribution (PLD).

Worst-case distributions: The privacy loss is computed for two distributions, but is not restricted to special cases. For many mechanisms M there are so-called worst-case distributions A and B with a privacy loss maximally as great as that of $M(x_0)$ and $M(x_1)$ for all pairs of neighboring inputs x_0 and x_1 . We give some intuition on how worst-case distributions work and why they typically exist, but refer to Meiser and Mohammadi's recent work [20] for a more detailed discussion.

For non-adaptive mechanisms, i.e., for mechanisms that do not change structurally from one execution to the next, there is always such a pair of worst-case distributions [2, 11]. In most cases, the worst-case distribution is defined by the worst possible (in terms of privacy) pair of inputs that is still considered neighboring. If mechanisms behave structurally differently on different inputs, the worst-case distributions have to be artificially cre-




Fig. 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.

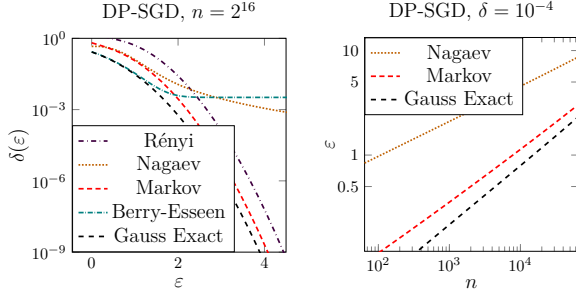


Fig. 2. Comparing bounds for differentially private stochastic gradient descent mechanism with noise parameter $q = 0.01$ and $\sigma = 4$. Left: after $n = 2^{16}$ compositions, right: minimal ε values over the number of compositions n for $\delta \leq 10^{-4}$. In the right graph, the Rényi-DP and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

ated by combining the PLDs for all neighboring pairs of inputs, which might be computationally challenging.

Adaptive queries can often be captured as well: queries can be considered part of the input, and they are neighboring if, e.g., without adding noise, the results at most differ by the application-specific sensitivity. As an illustrative example, consider a database-query-response system that adds noise to its real-valued answers to queries $q : X \rightarrow \mathbb{R}$ 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 q has a sensitivity of 1, i.e., for all allowed pairs of inputs x_0, x_1 we have $|q(x_0) - q(x_1)| \leq 1$, as is the case for sum-queries, then the distributions obtained by $M(0)$ and $M(1)$ are worst-case distributions. If a subsequent query q' differs from q (potentially depending on the mechanism's output for q), the worst-case distributions remain unchanged, as long as $|q'(x_0) - q'(x_1)| \leq 1$. Hence, the results we give for the Gauss mechanism, specifically our analytical formula for differential privacy (Theorem 5), holds in the light of adaptive queries.

For analyzing a mechanism, a pair of worst-case distributions has to exist for M and for all inputs that fall into the neighboring relation. In doing so, we abstract

away from the concepts of utility and sensitivity and require the privacy analyst interested in applying our results to provide worst-case distributions. In the remainder of this work we concentrate on a pair of distributions $M(x_0)$ and $M(x_1)$ for a mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$ and two concrete inputs $x_0, x_1 \in \mathcal{X}$. All our results also apply for a pair of worst-case distributions.

The privacy loss distribution: Given a pair of distributions, we can consider the corresponding privacy loss distribution. 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. By the central limit theorem, a privacy loss distribution converges to a predictable Gauss distribution under sufficiently many compositions (Theorem 4).

The privacy loss distribution of the Gauss mechanism also is a Gauss distribution and under convolution again remains a Gauss distribution. We give an analytical and efficiently computable formula for ADP and PDP composition bounds for any number of compositions. Note that these are not approximate bounds, but indeed precise characterizations (Theorem 5, Figure 5).

We also provide bounds for arbitrary mechanisms (for which a worst-case reduction exists); after many ($n > 2^{22}$) compositions, our bounds outperform previous work. Our representation with PLDs directly shows that the moments accountant and the RDP bound actually are applications of the Markov inequality to compute a PDP bound. 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 [25], see Figure 6. 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 $\delta(\varepsilon)$ such that the mechanism is still $(\varepsilon, \delta(\varepsilon))$ -ADP 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.

3 Related Work

Vadhan et al. [26] examined the same kind of n -fold adaptive composition as this work. Roughly speaking, they have shown that privacy will deteriorate as $\sqrt{n\varepsilon} + n\varepsilon^2$, rather than the (trivial) worst-case $n\varepsilon$ known in the literature. Meiser and Mohammadi [20] 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. Computing their bounds has 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 leads to significantly improved results for very high n values, which is highly relevant for a system like CoverUp.

Kairouz et al. [16] 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 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 (Example 1) for the ARR under sequential composition that directly computes such minimal bounds.

Recent work on concentrated differential privacy (CDP) [8, 12] directly focuses on the privacy loss for deriving tighter ADP and PDP bounds. This line of work

provides interesting insights into differential privacy and into improved bounds for the Gauss mechanism; for other mechanisms, however, these results either provide very loose bounds (e.g., the truncated Laplace mechanism) or no bounds at all (e.g., [2]). Our work, in contrast, identifies the variance, the mean, and the mass of the distinguishing events of the privacy loss distribution before composition (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.

Rényi differential privacy (RDP) [21] is a privacy notion based on the log normalized-moments of the privacy loss distribution (the Rényi divergence). RDP is a generalization of the moments account bound (MA) [2]. We evaluate MA in Section 6 and show an equivalence between RDP's moments, the PLD, and ADP (Theorem 2), which exceeds the RDP to ADP bound in [21].

In a concurrent work, Balle et al. [5] revisited the Gauss mechanism for optimal denoising in differential privacy. Interestingly, their concurrent work results in the same exact ADP-bound of the Gauss mechanism, without any composition results, however. Additionally, Balle et al. [4] concurrently proved that ADP bounds imply RDP bounds but not the converse direction.

4 Privacy Loss Space

We review the privacy loss, a representation of the privacy leakage introduced by Dinur and Nissim [9]. 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 (ADP) [10], probabilistic differential privacy (PDP) [14, 18], and Rényi differential privacy (RDP) [21], and concentrated differential privacy (CDP) [8, 12]. For a generalization of ADP (Definition 4.5) and for RDP, the PLD is even a canonical, unique, and succinct representation of the leakage (Theorem 2 and Corollary 1). We further prove that a sequential composition translates to convolution of the respective privacy loss distributions (Theorem 1). For any proofs, we refer to Appendix B.

Notation. See Table 1 for a summary of our notation. Formally, a *probabilistic mechanism* M from X to Y describes function $M : X \rightarrow (\Omega \rightarrow Y)$, with $\Omega := \bigcup_{x \in X} \Omega_x$ and Ω_x being the set of measurable sets on which the random variable $M(x)$ (for $x \in X$) is defined.

$M(x_0),$ $M(x_1)$	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 over 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 $\mathcal{L}_{A/B}(\mathcal{U})$
$\varpi, \varpi(y), \mathcal{C}$	dual PLD of ω (Definition 4.3)

Table 1. Notation table

4.1 Privacy Loss Variables / 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 $M : \mathcal{X} \rightarrow \mathcal{U}$, 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} .

We require the universe \mathcal{U} to be countable. For continuous distributions we generalize our results to Lebesgue measurable sets (c.f. Section 5.2).

Definition 4.2 (Privacy Loss Distribution (PLD)).

Let A and B be two probability distributions over the countable universe \mathcal{U} . The privacy loss distribution ω of A over B is defined as follows:

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

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

The support \mathcal{Y} of ω additionally includes the symbol¹ $-\infty$: $\text{supp}(\omega) := \{y \mid \omega(y) \neq 0\} \cup \{-\infty\}$. 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. $\omega(\infty) = \sum_{\{x \mid \Pr[o \leftarrow B] = 0\}} \Pr[o \leftarrow A]$
5. $\omega(-\infty) = 0$

These properties directly follow from the definition; we refer to Appendix B for our proofs. 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.

¹ We are aware that the support of a probability mass function $\omega(y)$ is usually defined as the set of y with $\omega(y) \neq 0$. The inclusion of $-\infty$ simplifies notation.

Theorem 1 (Composition). *Let $M : \mathcal{X} \rightarrow \mathcal{U}$ and $M' : \mathcal{X}' \rightarrow \mathcal{U}'$ 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' by $M'(x'_0)$ over $M'(x'_1)$ with support \mathcal{Y}' 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:*

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

So, $\forall y_c \in \mathcal{Y}_c \setminus \{-\infty, \infty\}$ we have

$$\begin{aligned} \omega_c(y_c) &= (\omega * \omega')(y_c) = \sum_{\{y, y' \mid y+y'=y_c\}} \omega(y) \cdot \omega(y'), \\ \omega_c(\infty) &= 1 - [1 - \omega(\infty)] \cdot [1 - \omega'(\infty)], \quad \omega_c(-\infty) = 0 \end{aligned}$$

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

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 probabilistic mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$, 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) ϖ with support \mathcal{L} is defined as*

$$\begin{aligned} \mathcal{L} &= \{-y \mid y \in \mathcal{Y}\} \\ \varpi(y) &= \omega(-y) e^y \quad \forall y \in \mathcal{L} \\ \varpi(\infty) &= 1 - \sum_{y \in \mathcal{L} \setminus \{-\infty, \infty\}} \varpi(y) \\ \varpi(-\infty) &= 0 \end{aligned}$$

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

4.3 Inner Privacy Loss Distribution

Most of the privacy bounds in this work and in the literature do not consider distinguishing events, i.e., $\omega(\infty) = 0$. Hence, these events with $\mathcal{L}(o) = \infty$ have to be treated differently. We examine the distribution conditioned on excluding such events.

Definition 4.4 (Inner Distribution). *The inner distribution $\bar{\omega}$ of a privacy loss distribution ω is the normalized distribution without $\omega(-\infty)$ and $\omega(\infty)$. $\forall y \in \mathcal{Y} \setminus \{-\infty, \infty\}$*

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

We can define a mechanism M' that leads to the inner distribution directly.

Lemma 3 (Mechanism for Inner Distribution). *Let $M : \mathcal{X} \rightarrow \mathcal{U}$ be a probabilistic mechanism and x_0, x_1 be inputs with common support $O = \{o \mid \forall i \in \{0, 1\}, \Pr[o \leftarrow M(x_i)] \neq 0, o \in \mathcal{U}\}$, leading to a privacy loss distribution ω . Let $M'_{M,O} : \mathcal{X} \rightarrow \mathcal{U}$ be a probabilistic mechanism with $\Pr[o \leftarrow M'_{M,O}(x)] = \Pr[o \leftarrow M(x) \mid o \in O]$. Then, the privacy loss distribution created by $M'_{M,O}(x_0)$ over $M'_{M,O}(x_1)$ is equal to the inner distribution $\bar{\omega}$ of ω .*

The following lemma shows that many of the bounds that do not consider distinguishing events can be generalized if the bound is considered to constrain only the inner distribution.

Lemma 4 (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 $g : \mathcal{Y} \rightarrow \mathbb{R}$ and for $\gamma \in \mathcal{Y} \setminus \{-\infty, \infty\}$, $\sum_{y \geq \gamma} g(y) \bar{\omega}(y) \leq \mathcal{B}(\gamma)$, then the bound can be expressed for the full distribution: $\sum_{y \geq \gamma} g(y) \omega(y) \leq \omega(\infty) + [1 - \omega(\infty)] \mathcal{B}(\gamma)$ with $g(\infty) = 1$.*

4.4 Approximate Differential Privacy

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

Definition 4.5 (ADP). *Let $M : \mathcal{X} \rightarrow \mathcal{U}$ be a probabilistic mechanism and $x_0, x_1 \in \mathcal{X}$. We say M is (ε, δ) -differentially private (or (ε, δ) -ADP) for x_0, x_1 if we have for all sets $S \subseteq \mathcal{U}$*

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

We say that δ is tight for ε and x_0, x_1 if there is no $\delta' < \delta$ such that the mechanism is (ε, δ') -ADP for x_0, x_1 . We write $\delta(\varepsilon)$ for this tight δ of an ε . The ADP-graph is defined as $(\varepsilon, \delta(\varepsilon))_{\varepsilon \in \mathbb{R}}$. Given a neighboring relation, we call the mechanism M (ε, δ) -ADP if M is (ε, δ) -ADP for all neighboring $x_0, x_1 \in \mathcal{X}$.

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 an illustration of ADP on two probability distributions, see Figure 4, following a depiction in [20].

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 [20].

Definition 4.6. For a probabilistic mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$ with 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

$$\delta_{M(x_0), M(x_1)}^*(\varepsilon) = \omega(\infty) + \sum_{y > \varepsilon, y \in \mathcal{Y} \setminus \{-\infty, \infty\}} (1 - e^{\varepsilon - y}) \omega(y)$$

We now show that Definitions 4.5 and 4.6 are equivalent.

Lemma 5. For every probabilistic mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$, $x_0, x_1 \in \mathcal{X}$, and for any values $\varepsilon, \delta \geq 0$, M is $(\varepsilon, \delta(\varepsilon))$ -tightly ADP for x_0, x_1 as in Definition 4.5 iff we have $\delta(\varepsilon) = \left(\delta_{M(x_0), M(x_1)}^*(\varepsilon) \right)$ (c.f., Definition 4.6).

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 [16] for (ε, δ) -ADP.

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

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

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 5, for n compositions, we get the exact result

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

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




Fig. 3. Uniform noise and its privacy loss distribution before composition ($n = 1$) and after a few compositions ($n = 4$).

Example 2 (Non-DP uniform noise). Consider a mechanism that adds uniform noise to its input. We refer to Figure 3 for a graphical depiction. Let $M : X \rightarrow U$, with $U = \{-5, \dots, 5\}$, $x_1, x_2 \in X$, and $\Pr[o \leftarrow M(x_0)] = p_0(o)$, $\Pr[o \leftarrow M(x_1)] = p_1(o)$ be

$$p_0(o) = \begin{cases} \frac{1}{9} & o \in \{-5, \dots, 4\} \\ 0 & o = 5 \end{cases} \quad p_1(o) = \begin{cases} 0 & o = -4 \\ \frac{1}{9} & o \in \{-4, \dots, 5\} \end{cases}$$

leading to

$$\omega(y) = \begin{cases} \frac{8}{9} & y = 0 \\ \frac{1}{9} & y = \infty \end{cases} \quad \omega_n(y) = \begin{cases} \left(\frac{8}{9}\right)^n & y = 0 \\ 1 - \left(\frac{8}{9}\right)^n & y = \infty \end{cases}$$

The privacy loss distribution ω is not $(\varepsilon, 0)$ -ADP as $\forall \varepsilon > 0, \varepsilon < \mathcal{L}_{M(x_0)/M(x_1)}(o = -5) = \infty$, i.e. $\omega(\infty) \neq 0$, but it is $(\varepsilon, \frac{1}{9})$ -ADP $\forall \varepsilon \geq 0$. Moreover, ω_n is the PLD after n compositions, according to Theorem 1.

Equivalence of PLD and ADP-Graph. We now show that an ADP-graph is as expressive as the privacy loss distribution. For distributions with finite support, it is possible to reconstruct the PLD from the $(\varepsilon, \delta(\varepsilon))_\varepsilon$ sequence. In practice, finite support is typically the case due to discretization and finite representations of numbers. From Lemma 5, the opposite direction then follows. This is a significant result, as the privacy loss distribution is sufficiently strong for other important privacy notions.

Lemma 6 (Equivalence of ADP and PLD). There exists a bijection R such that the following holds. For any probabilistic mechanism M on inputs x_0 and x_1 with a privacy loss distribution (\mathcal{Y}, ω) with finite cardinality $|\mathcal{Y}| = k$ (for $k \in \mathbb{N}$), we get the tight ADP-graph $(\varepsilon, \delta(\varepsilon))_\varepsilon$ for x_0, x_1 (as in Definition 4.5) with $R(\mathcal{Y}, \omega) = (\varepsilon, \delta(\varepsilon))_\varepsilon$ and backwards $R^{-1}((\varepsilon, \delta(\varepsilon))_\varepsilon) = (\mathcal{Y}, \omega)$.




Fig. 4. A graphical depiction of the (truncated) Gauss mechanism on two inputs, $A = \mathcal{N}(0, \sigma^2)$, $B = \mathcal{N}(1, \sigma^2)$, and of how to compute ADP $\delta_{ADP}(\epsilon)$ and PDP $\delta_{PDP}(\epsilon)$ for a given value ϵ . Note that $e^\epsilon \cdot A$ is not a probability distribution.

4.5 Probabilistic Differential Privacy

Probabilistic differential privacy [14, 18] is a very intuitive variant of approximate differential privacy (see Figure 4). The main idea is to require that with probability $1-\delta$ pure ϵ -differential privacy holds. While this definition has a clear semantics and is easy to understand, it is not closed under post-processing [19], 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 4.7 (PDP). A probabilistic mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$ is (ϵ, δ) -probabilistically differentially private (PDP) for $x_0, x_1 \in \mathcal{X}$, where $\epsilon \geq 0$ and $\delta \geq 0$, if there are sets $S_0^\delta, S_1^\delta \subseteq \mathcal{U}$ with $\Pr[M(x_0) \in S_0^\delta] \leq \delta$ and $\Pr[M(x_1) \in S_1^\delta] \leq \delta$, s.t., for all sets $S \subseteq \mathcal{U}$, the following in-equations hold:

$$\begin{aligned} \Pr[M(x_0) \in S \setminus S_0^\delta] &\leq e^\epsilon \cdot \Pr[M(x_1) \in S \setminus S_0^\delta] \\ \wedge \Pr[M(x_1) \in S \setminus S_1^\delta] &\leq e^\epsilon \cdot \Pr[M(x_0) \in S \setminus S_1^\delta]. \end{aligned} \quad (1)$$

M is tightly (ϵ, δ) -PDP for x_0, x_1 if δ is minimal for ϵ , i.e., if for all δ' such that M is (ϵ, δ') -PDP for x_0, x_1 , $\delta' \geq \delta$. Given a neighboring relation, if M is (ϵ, δ) -PDP for any neighboring $x_0, x_1 \in \mathcal{X}$ then M is (ϵ, δ) -PDP.

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

Lemma 7 (Connection to PDP). Let $M : \mathcal{X} \rightarrow \mathcal{U}$ be a probabilistic mechanism and $x_0, x_1 \in \mathcal{X}$ two inputs with the PLD ω and let ϖ be its dual PLD, then

$$M \text{ is } (\epsilon, \delta)\text{-PDP for } x_0, x_1 \iff \begin{aligned} \sum_{y > \epsilon, y \in \mathcal{Y}} \omega(y) &\leq \delta \\ \sum_{y > \epsilon, y \in \mathcal{X}} \varpi(y) &\leq \delta \end{aligned}$$

4.6 Rényi Differential Privacy

Recent work introduced novel ADP bounds that are based on the Rényi divergence (the logarithm of the higher moments of the exponentiated privacy loss random variable $e^{\mathcal{L}}$): concentrated DP (CDP) [8, 12], Rényi DP (RDP) [21], and the moments accountant [2]. This Rényi divergence can be defined using the PLD. In particular, as CDP and RDP are based on the Rényi divergence, Lemma 6 implies that RDP and CDP can be determined from the tight ADP-graph $(\epsilon, \delta(\epsilon))_\epsilon$ (for distributions with finite support).

Definition 4.8 (Rényi Divergence & RDP). The Rényi divergence $\mathcal{D}_\alpha(M(x_0)|M(x_1))$ with $\alpha > 1$ for a probabilistic mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$ and two inputs $x_0, x_1 \in \mathcal{X}$ is defined as

$$\begin{aligned} \mathcal{D}_\alpha(M(x_0)|M(x_1)) &= \frac{1}{\alpha-1} \ln \mathbb{E}_{o \sim M(x_1)} \left(e^{\mathcal{L}_{M(x_0)/M(x_1)}(o)} \right)^{\alpha-1} \\ \mathcal{D}_1(M(x_0)|M(x_1)) &= \mathbb{E}_{o \sim M(x_0)} \left(\mathcal{L}_{M(x_0)/M(x_1)}(o) \right) \end{aligned}$$

Rényi differential privacy characterizes privacy as the sequence of Rényi divergences: $(\alpha, D_\alpha)_\alpha$. Given a neighboring relation, M has ϵ -Rényi differential privacy of order $\alpha > 1$ ((α, ϵ) -RDP) if $\mathcal{D}_\alpha(M(x_0)|M(x_1)) \leq \epsilon$ for all neighboring $x_0, x_1 \in \mathcal{X}$.

Note that $\mathcal{D}_1(M(x_0)|M(x_1))$ coincides with the Kullback-Leibner (KL) divergence from $M(x_0)$ to $M(x_1)$. This is a natural property of the PLD, if and only if no output $o \in \mathcal{U}$ has an infinite privacy loss. Analogously, we get the KL divergence from $M(x_1)$ over $M(x_0)$ by the dual PLD ϖ .

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 $(\epsilon, \alpha D_\alpha - \alpha \epsilon)$ -ADP holds [21]. The moments accountant uses the same characterization and proposes $(\epsilon, \min_\alpha (\alpha D_\alpha - \alpha \epsilon))$ as ADP bounds (as (ϵ, δ) -PDP implies (ϵ, δ) -ADP).

Equivalence of PLD and RDP. RDP is closely connected to the moments of the privacy loss distribution [28]. In fact, the α -Rényi-divergence D_α is the α -1-root of the logarithm of the $(\alpha-1)$ -moments of the exponentiated distribution of ω . If the moments ρ_λ of the exponentiated ω are not growing too fast, $|\rho_\lambda| < c d^\lambda \lambda!$ for a $\lambda > 0$, then we have equivalence, i.e., we can compute the moments from the privacy loss distribution of a mechanism and vice versa. For privacy loss distributions on a bounded support we always have equivalence.

Lemma 8 (Equivalence to RDP). *There exists a bijection R_{RDP} such that the following holds. For any probabilistic mechanism M with a countable support on inputs x_0 and x_1 with a privacy loss distribution (\mathcal{Y}, ω) , s.t., $\omega(\infty) = 0$, the Rényi Divergence of order λ with $\lambda > 0$ of $M(x_0)$ and $M(x_1)$ is*

$$m_\lambda = \frac{1}{\lambda} \ln \left(\mathbb{E}_{y \sim \omega} e^{\lambda y} \right) = \mathcal{D}_{\lambda+1}(M(x_0) \| M(x_1)) \quad (2)$$

Moreover, if $\exp(\lambda \cdot m_\lambda) < cd^\lambda \lambda!$ for two positive constants c, d , then we get the Rényi-Divergence-sequence $R_{RDP}((\mathcal{Y}, \omega)) = (\alpha, D_\alpha)_\alpha$ and the PLD $R_{RDP}^{-1}((\alpha, D_\alpha)_\alpha) = (\mathcal{Y}, \omega)$.

4.7 Equivalence of Rényi-DP and ADP


From Lemma 6, we know that for any mechanism M and any two inputs x_0, x_1 the ADP-graph can be bijectively mapped to the PLD and vice versa, if the common support is discrete and finite. From Lemma 8, we know that for any mechanism M and any two inputs x_0, x_1 the Rényi Divergences can be bijectively mapped to the PLD and vice versa, as long as the sequence of Rényi Divergences satisfy a technical sanity condition. As a result, we can conclude that the series of moments $(D_\alpha)_\alpha$, the ADP-graph and the PLD contain the same information and can be transformed into one another.

Theorem 2 (Equivalence of ADP, RDP and PLD).

There exists bijections R_{ADP} and R_{RDP} such that the following holds. Given a mechanism M . Let (x_0, x_1) be an arbitrary but fixed pair of inputs such that

- $\forall o \in \mathcal{U}, i \in \{0, 1\}, \Pr[o \leftarrow M(x_i)] \neq 0$
- *the support \mathcal{Y} of the PLD of $M(x_0)$ and $M(x_1)$ has finite cardinality $|\mathcal{Y}| = k$ (for $k \in \mathbb{N}$),*
- *the support of $M(x_0)$ and of $M(x_1)$ is countable, and*
- $\forall \lambda > 0, \mathbb{E}_{o \sim M(x_0)} \left[\frac{\Pr[o \leftarrow M(x_0)]}{\Pr[o \leftarrow M(x_1)]} \right]^\lambda < cd^\lambda \lambda!$ *for two positive constants c, d .*

Let $(\varepsilon, \delta(\varepsilon))_{\varepsilon \in \mathbb{R}}$ be the ADP-Graph for x_0, x_1 (Definition 4.5), let $(\alpha, D_\alpha(M(x_0) \| M(x_1)))$ be the Rényi-Divergence-sequence ([21]), let ω be the PLD (Definition 4.2). Then, the following diagram commutes:



The previous result can be extended to mechanisms with non-equal image space, i.e. with $\omega(\infty) \neq 0$.

Corollary 1 (Equivalence with distinguishing events). *Given M on two inputs x_0, x_1 with distinguishing events and $M'_{M,O}(x)$ as in Lemma 3, we can apply Theorem 2, resulting in a bijection for the inner distribution $\bar{\omega}$ of $M(x_0)$ over $M(x_1)$ (see Definition 4.4), which allows us to consider the distinguishing events separately.*

4.8 Markov-ADP Bound

Next, we refine an ADP bound introduced by the moments accountant [2], which we coin *Markov-ADP* bound. We use Markov's inequality to limit the privacy loss. In contrast to [2], we discretize the tail and incorporate its contribution in a more fine-grained manner.

Theorem 3 (Markov-ADP bound). *A mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$ with two 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 non-empty sub-set of \mathcal{Y}_n , i.e., $\mathcal{P} = \{y_0, \dots, y_k\} \subseteq \mathbb{R}^{k+1}$ with $y_i < y_{i+1} \forall i$. Then, after n compositions and for $\varepsilon \in \mathcal{P}$, $\varepsilon < y_0$,*

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

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

is a upper bound for tight-ADP for x_0, x_1 and smaller or equal to the PDP bound given by RDP [21].

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 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 does not apply to events that have an infinite privacy loss ∞ . Consequently, we focus on the inner distribution (c.f., Section 4.3). Recall that the inner distribution is the renormalized privacy loss distribution without distinguishing events. Under composition, the probability that a distinguishing event occurs changes in a simple and straight-forward way. Consequently, privacy loss classes have three defining elements: $(\mu, \sigma^2, \omega(\infty))$, the mean and the variance of the *inner distribution*, and the probability of distinguishing events $\omega(\infty)$.

Definition 5.1 (Privacy Loss Classes). *A privacy loss distribution ω with support \mathcal{Y} belongs to the $(\mu, \sigma^2, \omega(\infty))$ -privacy loss class*

$$\begin{aligned}\mu &= \sum_{y \in \mathcal{Y} \setminus \{-\infty, \infty\}} y \cdot \bar{\omega}(y), \\ \sigma^2 &= \sum_{y \in \mathcal{Y} \setminus \{-\infty, \infty\}} (y - \mu)^2 \cdot \bar{\omega}(y)\end{aligned}$$

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.

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 $M : \mathcal{X} \rightarrow \mathcal{U}$ be a mechanism with two inputs $x_0, x_1 \in \mathcal{X}$. 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

$$\begin{aligned}\mathcal{Y}_n &= \left\{ y \mid y = \sum_{i=1}^n \tilde{y}_i, \forall \tilde{y} \in \mathcal{Y}^n \right\} \\ \omega_n(y) &= (\otimes_{i=1}^n \omega_1)[y] \quad \forall y \in \mathcal{Y}_n \setminus \{-\infty, \infty\} \\ \omega_n(\infty) &= 1 - [1 - \omega_1(\infty)]^n \\ \omega_n(-\infty) &= 0\end{aligned}$$

with privacy loss class $(n\mu, n\sigma^2, \omega_n(\infty))$ where \otimes 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

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

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

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

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

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 [17]. 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 9 (Lebesgue-Generalization). *Let $M : \mathcal{X} \rightarrow \tilde{\mathcal{U}}$ be a mechanism with two inputs $x_0, x_1 \in \mathcal{X}$ with continuous universe $\tilde{\mathcal{U}}$. Define a Lebesgue–Rokhlin probability space $(\tilde{\mathcal{U}}, B(\mathbb{R}), \lambda_i)$ where $\tilde{\mathcal{U}} \in \mathbb{R}$, $B(\tilde{\mathcal{U}})$ denotes the Borel set and $\lambda_i(\mathcal{O} \subseteq \tilde{\mathcal{U}}) = \Pr[M(x_i) \in \mathcal{O}]$ is a Lebesgue measure (for $i \in \{0, 1\}$).*

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

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

Let $\tilde{\mathcal{Y}} = \mathcal{L}(\tilde{\mathcal{U}}) \setminus \{-\infty, \infty\}$. 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

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

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

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

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

$$\bar{\omega}(y) = \frac{\omega(y)}{1 - \omega(\infty)} \quad \forall y \in B(\tilde{\mathcal{Y}}) \quad (3)$$

$$\mu = \int_{\tilde{\mathcal{Y}}} y d\bar{\omega}(y)$$

$$\sigma^2 = \int_{\tilde{\mathcal{Y}}} (y - \mu)^2 d\bar{\omega}(y)$$

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

One advantage of the continuous perspective is the ability to derive sometimes an 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 10 (Density Transformation). *Let $M : \mathcal{X} \rightarrow \tilde{\mathcal{U}}$ be a mechanism with two inputs $x_0, x_1 \in \mathcal{X}$. Let the probability density $A(u) := \Pr[u \leftarrow M(x_0)]$ be continuous. Let privacy loss distribution ω be created by $M(x_0)$ over $M(x_1)$ with support $\tilde{\mathcal{Y}}$. For a subset $y \subseteq \tilde{\mathcal{Y}}$, let $O = \mathcal{L}_{M(x_0)/M(x_1)}^{-1}(y)$. 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*

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

5.3 ADP for the Gauss Mechanism

The Gauss mechanism applies Gauss distributed noise to a real-valued deterministic function, e.g., a query-response mechanism that gets as input a database D and a query q and outputs $q(D) + \mathcal{N}(0, \sigma^2)$. We abstract away from the use case and analyze the privacy loss of $M(x) = x + \mathcal{N}(0, \sigma^2)$. Consequently, we can focus on a simple neighboring relation: x_0 and x_1 are neighboring iff $|x_0 - x_1| \leq s$, where $s \in \mathbb{R}$ is the (limited) sensitivity. In the query-response example, $x = q(D)$; note that we can easily replace q : queries of subsequent runs can be chosen adaptively (as long as the sensitivity is not exceeded) and we will still analyze the same mechanism on inputs x_0, x_1 with $|x_0 - x_1| \leq s$. Our analysis also applies to other use cases as long as the query on neighboring inputs has a sensitivity bounded by s . We here present a tight analytic formula for $\delta(\varepsilon)$ for the Gauss mechanism. This result is a significant contribution, as it allows to compute (not just approximate) the exact privacy loss. We show that PLD of the Gauss mechanisms also is a Gauss distribution and under composition remains a Gauss distribution. Thus, we yield a tight analytic formula for the Gauss mechanism after an arbitrary number of compositions. If the actual sensitivity of the underlying function differs from s our bounds naturally lose their tightness.

Lemma 11 (PLD of Gauss Mechanism). *Let M be a probabilistic Gauss mechanism with $M : \mathcal{X} \rightarrow \tilde{\mathcal{U}}$ and $M(x) \sim \mathcal{N}(x, \sigma^2)$ for $\sigma^2 > 0$ and let $x_0, x_1 \in \mathcal{X}$. Then the privacy loss distribution ω generated by $M(x_0)$ and $M(x_1)$ is a Gauss distribution $\omega \sim \mathcal{N}\left(\frac{(x_0 - x_1)^2}{2\sigma^2}, \frac{(x_0 - x_1)^2}{\sigma^2}\right)$ and $\omega(\infty) = 0$ with privacy loss class $\left(\frac{(x_0 - x_1)^2}{2\sigma^2}, \frac{(x_0 - x_1)^2}{\sigma^2}, 0\right)$.*

Bun and Steinke have already derived the absolute moments of the Gauss mechanism [8, Lemma 2.4] which implies the result of Lemma 11 as well.

Lemma 12 (Tight ADP for Gauss PLD). *Let ω be a continuous privacy loss distribution in the shape of a Gauss distribution $d\omega(y) = \frac{1 - \omega(\infty)}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y - \mu)^2}{2\sigma^2}} dy$ and with privacy loss class $(\mu, \sigma^2, \omega(\infty))$ for any $0 \leq \omega(\infty) \leq 1$. Then $\delta_{M(x_0)}(\varepsilon) =$*

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

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

Recall from Lemma 11 that for the Gauss mechanism with noise parameter σ and sensitivity $|x_0 - x_1|$, the mean μ_{pld} and variance σ_{pld}^2 of their respective privacy loss distribution are related: $\mu_{\text{pld}} = \sigma_{\text{pld}}^2/2 = \frac{|x_0 - x_1|^2}{\sigma^2}/2$.

Theorem 5 (Tight ADP for the Gauss Mechanism).

A Gauss mechanism $M : \mathcal{X} \rightarrow \tilde{\mathcal{U}}$ with sensitivity s and $M(x) \sim \mathcal{N}(x, \sigma^2)$ for $\sigma^2 > 0$ has exactly

$$\delta(\varepsilon) = \frac{1}{2} \left[\operatorname{erfc} \left(\frac{\varepsilon - n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}} \right) - e^\varepsilon \cdot \operatorname{erfc} \left(\frac{\varepsilon + n\mu_{\text{pld}}}{\sqrt{2n}\sigma_{\text{pld}}} \right) \right] \quad (4)$$

after n compositions, with $\sigma_{\text{pld}} = \frac{s}{\sigma}$ and $\mu_{\text{pld}} = \sigma_{\text{pld}}^2/2$ and is tightly $(\varepsilon, \delta(\varepsilon))$ -ADP as in Definition 4.5.

Corollary 2 (Tight PDP for the Gauss Mechanism).

A Gauss mechanism $M : \mathcal{X} \rightarrow \tilde{\mathcal{U}}$ with $M(x) \sim \mathcal{N}(x, \sigma^2)$ for $\sigma^2 > 0$ has for $x_0, x_1 \in \mathcal{X}$ after n compositions exactly

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

with $\sigma_{\text{pld}} = \frac{|x_0 - x_1|}{\sigma}$ and $\mu_{\text{pld}} = \sigma_{\text{pld}}^2/2$ and is tightly $(\varepsilon, \delta_{\text{PDP}}(\varepsilon))$ -PDP as in Definition 4.7.

It is of practical interest to find the smallest degree of noise σ for a s.t. (ε, δ) -PDP still holds after n -fold sequential composition. For the Gauss mechanism, such a formula directly follows from Corollary 2.

Corollary 3 (Optimal σ for Gauss-Mechanism PDP).

A Gauss mechanism $M : \mathcal{X} \rightarrow \tilde{\mathcal{U}}$ with $M(x) \sim \mathcal{N}(x, \sigma^2)$ with $\mu = x_0 - x_1$, $x_0, x_1 \in \mathcal{X}$ requires for a privacy loss $\delta := \delta(\varepsilon) \leq \frac{1}{4}$ after n compositions

$$\sigma(\varepsilon, \delta, n) = \frac{\mu\sqrt{n}}{\sqrt{2\varepsilon}} \left(\operatorname{erfc}^{-1}(2\delta) + \sqrt{(\operatorname{erfc}^{-1}(2\delta))^2 + \varepsilon} \right)$$

For ADP, however, there might not be a analytic formula for σ because the inversion of two erfc functions simultaneously is difficult; as $\delta(\cdot)$ is strong monotonically decreasing, a unique solution can be found numerically.

5.4 ADP for Arbitrary 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.

Theorem 6 (ADP under composition). *Let $\varepsilon \geq 0$ and n be arbitrary but fixed. Let $M : \mathcal{X} \rightarrow \mathcal{U}$ be a mechanism with two inputs $x_0, x_1 \in \mathcal{X}$. 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 $\gamma = \mathbb{E}|\bar{\omega}_1(y)|^3 < \infty$. 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*

$$r_u := c_u \frac{\gamma}{\sigma^3}, \quad r_t(z) := \begin{cases} c_t \frac{\gamma}{\sigma^3(1+z^3)} & \text{if } z \geq 0 \\ \infty & \text{else} \end{cases}$$

$$\begin{aligned} \omega_n(\infty) &:= 1 - [1 - \omega_1(\infty)]^n \\ \Delta_\omega &:= \omega_n(\infty) + \frac{1 - \omega_n(\infty)}{2} \left[\operatorname{erfc} \left(\frac{\varepsilon - n\mu}{\sqrt{2n}\sigma} \right) - e^{\varepsilon - n\mu + n\frac{\sigma^2}{2}} \operatorname{erfc} \left(\frac{\varepsilon - n\mu + n\sigma^2}{\sqrt{2n}\sigma} \right) \right] \\ \beta_\omega &:= \frac{[1 - \omega_n(\infty)]}{\sqrt{n}} \min \left[r_u, r_t \left(z = \frac{\varepsilon - n\mu}{\sqrt{n}\sigma^2} \right) \right] \end{aligned}$$

with $c_u = 0.4748$ and $c_t = 25.80$. Then, M is $(\varepsilon, \max(\Delta_\omega + \beta_\omega, \Delta_\sigma + \beta_\sigma))$ -ADP for x_0, x_1 .

Remark: In this paper we have used privacy definitions for M on a concrete pair of inputs x_0, x_1 . If this pair of inputs is worst-case (i.e., $M(x_0)$ and $M(x_1)$ are worst-case distributions), the results immediately generalize to the whole mechanism. Otherwise, if a pair of worst-case distributions A and B can be found for that particular mechanism and sensitivity notion, we can replace all occurrences of $M(x_0)$ by A and all occurrences of $M(x_1)$ by B and our results and proofs still apply word by word.

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 former has key advantages.

6.1 Evaluating Our Bounds

We apply our various theoretical results to several mechanisms 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. For certain mechanisms, concentrated differential privacy (CDP)² provides compelling bounds. Our figures use the approximate zCDP ADP-bound [8] only for the Gauss mechanism, as zCDP requires to prove that the log-normalized-moments of the privacy loss distribution can be bounded by an affine linear function. While zCDP provides compelling ADP bounds for higher epsilons, it provides grossly inaccurate values for $\varepsilon = 0$ (i.e., total variation) and very small ε values.³

We use the numerical lower bound provided by the privacy buckets [20] as a benchmark in the right graph, but omit it in the left graphs to ease readability. In Figure 8 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.

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 5 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 data⁴, which results in a pair of Gauss mixture distributions. CoverUp [25] is 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

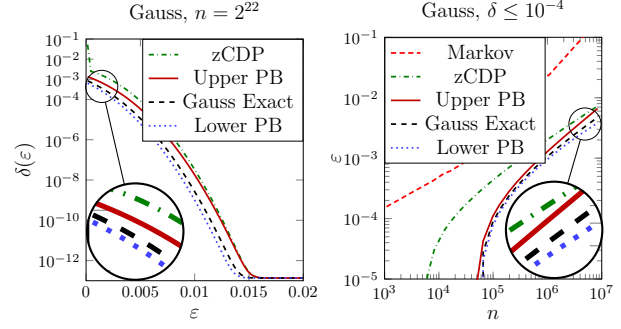


Fig. 5. Comparison of Gauss mechanism to known bounds with noise parameter $\sigma^2 = 900000^2$, left: 2^{22} compositions, right: minimal ε values over the number of compositions n for $\delta \leq 10^{-4}$. Comparing the exact Gauss-ADP formula with various bounds. In the right graph, Berry-Esseen bound did not fall into the plotting range and were omitted.

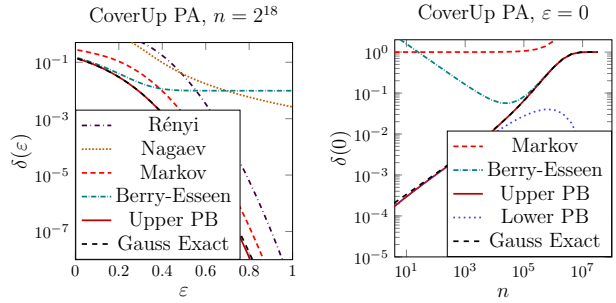


Fig. 6. Comparing bounds for the CoverUp measurement data with noise parameter `width_of_noise` = 100, left: after $n = 2^{18}$ compositions, right: $\delta(0)$ (i.e., $\varepsilon = 0$) over the number of compositions n . In the right graph, the Rényi-DP bound and the Nagaev bound did not reach meaningful values of $\delta(0) \leq 1$.

is not installed. Figure 6 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.

For Abadi et al.’s differentially private stochastic gradient descent (DP-SGD) mechanism [2], analyzing the following worst-case distributions suffices: a Gauss distribution $\mathcal{N}(0, \sigma^2)$ and a Gauss mixture distribution $q\mathcal{N}(0, \sigma^2) + (1-q)\mathcal{N}(1, \sigma^2)$ (with $q \in [0, 1]$). Figure 7 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 [16] 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 8.

2 Concentrated-DP [12]: A mechanism $M : \mathcal{X} \rightarrow \mathcal{U}$ satisfies (ξ, ρ) -CDP if for all $\alpha > 0$, and all neighboring $x_0, x_1 \in \mathcal{X}$ (for a neighboring relation), $\mathcal{D}_\alpha(M(x_0) \| M(x_1)) \leq \xi + \rho\alpha$.

3 The case $\varepsilon = 0$ is important: the total variation, $\delta(0)$, is used in the statistical indistinguishability notion when deniability ($\varepsilon > 0$) is irrelevant and only pure indistinguishability ($\varepsilon = 0$) matters, as, e.g., in the timing analysis of CoverUp [25].

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

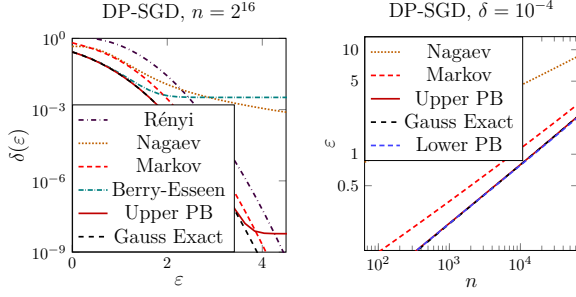


Fig. 7. Comparing bounds for differentially private stochastic gradient descent mechanism (DP-SGD) with noise parameters $q = 0.01$ and $\sigma = 4$, left: after $n = 2^{16}$ compositions, right: minimal ϵ values over the number of compositions n for $\delta \leq 10^{-4}$. In the right graph, the Berry-Esseen bound did not fall into the plotting range and were omitted.

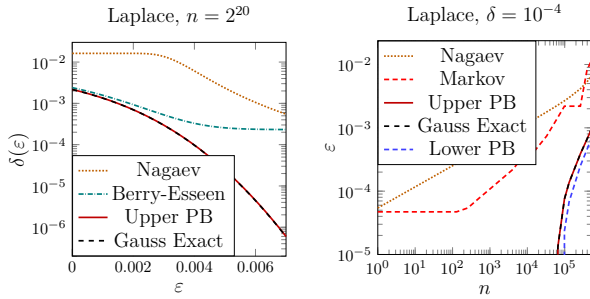


Fig. 8. Comparison of Laplace mechanism to known bounds with noise parameter $\lambda = 1.26 \cdot 10^6$, left: 2^{20} compositions, right: minimal ϵ values over the number of compositions n for $\delta \leq 10^{-4}$. In the right graph, Berry-Esseen bound did not fall into the plotting range and were omitted.

Markov-ADP. In Section 4.8 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 5 and 8 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 [20].

For CoverUp and DP-SGD in Figures 6 and 7, Markov-ADP clearly outperforms the other bounds, except the numerical privacy buckets. In particular, the Markov-ADP bound outperforms the Rényi DP bound.

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 Nagaev bound.

For CoverUp (Figure 6), the left graph shows that the Berry-Esseen bound is pretty tight until $\delta(\epsilon) < 10^{-2}$, similarly for DP-SGD (Figure 7) and Laplace (Figure 8) 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 show that the approach of using tail-bounds (such as the Nagaev Theorem) for normal approximations is a promising direction.

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 5 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 6 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.

6.2 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 .

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

$$\begin{aligned} \mu_{L,\text{pld}} &> \mu_{G,\text{pld}}, & \sigma_{L,\text{pld}} &\stackrel{(a)}{>} \sigma_{G,\text{pld}}, \\ (\mu_{L,\text{pld}}, \sigma_{L,\text{pld}}) &\xrightarrow{\frac{|x_0-x_1|}{\lambda} \rightarrow 0} (\mu_{G,\text{pld}}, \sigma_{G,\text{pld}}) \end{aligned}$$

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 $\omega(\infty)$.

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 2), while the Laplace 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 9 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 5, the upper and lower bounds from privacy buckets' numerical approximation [20] are as expected very close to the exact bound, yet they start to lose tightness for very high amount of compositions. This effect can be credited to numerical errors, memory constraints, and discretization errors. Our exact analytical




Fig. 9. 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.

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 [15] on the multiplication for numerical robustness. We can achieve high numerical stability with our implementation by rescaling the privacy loss distribution. Recall that the PLD of the Gauss mechanism is a Gauss distribution with mean μ_{pld} and variance σ_{pld}^2 . By computing $\mu_0 := \mu_{\text{pld}}/\mu = 1$ and $\sigma_0 := \sigma_{\text{pld}}/\mu$ and evaluating $\delta(\varepsilon)$ as $\delta(\varepsilon/\mu)$, we can avoid an overflow in computing the exponential function.

7 Conclusion and Future Work

We have analyzed the privacy loss of mechanisms and in doing so unified several 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, the central limit theorem applies and every privacy loss distribution converges to a Gauss distribution under composition. We categorize each mechanism into a privacy loss class by the parameters of this Gauss distribution.

For future work, we encourage finding a tight embedding of novel mechanisms into their respective privacy loss classes, in addition to the mechanisms for which we already give exact formulas: Laplace, Gauss

and randomized response; and searching for better convergence bounds, which obviously excludes the Gauss mechanism for which we provided an exact formula.

In practice, the privacy loss distribution typically converges to a Gauss distribution faster than the mechanism oblivious Berry-Esseen (BE) bound indicates. However, for some worst-case examples, the BE bound is tight. Using a more mechanism aware approximation bound is an interesting direction for future research.

Finally, we encourage examining which other distributions are closed under convolution and to find out whether the Gauss mechanism is provably optimal in the sense that the variance and mean of its privacy loss class are the smallest w.r.t. its initial variance.

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)*. ACM, 2016, pp. 308–318.
- [3] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, 1st ed. New York: Dover, 1972.
- [4] B. Balle, G. Barthe, and M. Gaboardi, “Privacy amplification by subsampling: Tight analyses via couplings and divergences,” in *Neural Information Processing Systems (NIPS)*, 2018.
- [5] B. Balle and Y. Wang, “Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal de-noising,” in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018, pp. 403–412.
- [6] P. Billingsley, *Probability and measure*. John Wiley & Sons, 2008.
- [7] V. I. Bogachev, *Measure theory*. Springer Science & Business Media, 2007, vol. 1.
- [8] M. Bun and T. Steinke, “Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds,” in *Theory of Cryptography (TCC)*. Springer, 2016, pp. 635–658.
- [9] 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)*. ACM, 2003, pp. 202–210.
- [10] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, “Our Data, Ourselves: Privacy Via Distributed Noise Generation,” in *Advances in Cryptology - EURO-CRYPT 2006*. Springer, 2006, pp. 486–503.
- [11] C. Dwork and A. Roth, “The Algorithmic Foundations of Differential Privacy,” *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [12] C. Dwork and G. N. Rothblum, “Concentrated Differential Privacy,” *CoRR*, vol. abs/1603.01887, 2016.
- [13] 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)*. ACM, 2014.
- [14] M. Götz, A. Machanavajjhala, G. Wang, X. Xiao, and J. Gehrke, “Privacy in search logs,” *CoRR*, vol. abs/0904.0682, 2009.
- [15] B. Gough, *GNU Scientific Library Reference Manual - Third Edition*, 3rd ed. Network Theory Ltd., 2009.
- [16] P. Kairouz, S. Oh, and P. Viswanath, “The composition theorem for differential privacy,” *IEEE Transactions on Information Theory*, vol. 63, no. 6, pp. 4037–4049, 2017.
- [17] J. W. Lindeberg, “Eine neue herleitung des exponentialgesetzes in der wahrscheinlichkeitsrechnung,” *Mathematische Zeitschrift*, vol. 15, no. 1, pp. 211–225, 1922.
- [18] 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*, April 2008, pp. 277–286.
- [19] S. Meiser, “Approximate and Probabilistic Differential Privacy Definitions,” <https://eprint.iacr.org/2018/277>, 2018.
- [20] 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.
- [21] I. Mironov, “Rényi Differential Privacy,” in *Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF)*. IEEE, 2017, pp. 263–275.
- [22] J. Murtagh and S. Vadhan, “The complexity of computing the optimal composition of differential privacy,” in *Proceedings, Part I, of the 13th International Conference on Theory of Cryptography (TCC)*. Springer, 2016, pp. 157–175.
- [23] I. Pinelis, “Chapter 4 - on the nonuniform berry-esseen bound,” in *Inequalities and Extremal Problems in Probability and Statistics*. Academic Press, 2017, pp. 103 – 138.
- [24] EU Regulation, “European data protection regulation (GDPR),” *Off J Eur Union*, vol. L119, pp. 1–88, 4th May 2016.
- [25] 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.
- [26] S. Vadhan, G. N. Rothblum, and C. Dwork, “Boosting and differential privacy,” in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS)*, 2010, pp. 51–60.
- [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)*. ACM, 2015, pp. 137–152.
- [28] Y.-X. Wang, B. Balle, and S. Kasiviswanathan, “Subsampled Rényi differential privacy and analytical moments accountant,” *arXiv preprint arXiv:1808.00087*, 2018.

A Examples

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

A.1 Approximate Randomized Response

event o	$M(x_0)$	$M(x_1)$	$\mathcal{L}_{M(x_0)/M(x_1)}(o)$
$o = 1$	δ	0	∞
$o = 2$	$\frac{(1-\delta)e^\varepsilon}{e^\varepsilon+1}$	$\frac{(1-\delta)}{e^\varepsilon+1}$	ε
$o = 3$	$\frac{(1-\delta)e^\varepsilon}{e^\varepsilon+1}$	$\frac{(1-\delta)e^\varepsilon}{e^\varepsilon+1}$	$-\varepsilon$
$o = 4$	0	δ	$-\infty$

$y \in \mathcal{Y}$	$\omega(y)$	$\bar{\omega}(y)$
$y = -\infty$	0	$-$
$y = -\varepsilon$	$\frac{(1-\delta)}{e^\varepsilon+1}$	$\frac{1}{e^\varepsilon+1}$
$y = \varepsilon$	$\frac{(1-\delta)e^\varepsilon}{e^\varepsilon+1}$	$\frac{e^\varepsilon}{e^\varepsilon+1}$
$y = \infty$	δ	$-$

Note that $\bar{\omega}(y) = \binom{2}{k} p^{k_{2,y}} (1-p)^{n-k_{2,y}} = B_{k_{2,y}}(n=2, p)$ with $k_{2,y} = \frac{y+2\varepsilon}{2\varepsilon}$, $p = \frac{1}{e^\varepsilon+1}$, where $B_k(n, p)$ denotes the total success probability of k successes with n trials and individual trial success probability p according to a Binomial distribution $B(n, p)$. Moreover, the convolution of two Binomial distributions is again a Binomial:

$$B(n, p) \otimes B(m, p) = B(n+m, p)$$

which gives us after n compositions by Theorem 4

$$\begin{aligned} \bar{\omega}_n(y) &= \binom{n}{k_{n,y}} p^{k_{n,y}} (1-p)^{n-k_{n,y}} \\ \omega_n(y) &= \begin{cases} 0 & y = -\infty \\ 1 - (1-\delta)^n & y = \infty \\ (1-\delta)^n \cdot \bar{\omega}_n(y) & \text{else} \end{cases} \\ \text{with } k_{n,y} &= \frac{y+n\varepsilon}{2\varepsilon}, \quad p = \frac{1}{e^\varepsilon+1} \end{aligned}$$

From there follows immediately by definition

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

with $y(k) = \varepsilon(2k-n)$ and $\lceil \cdot \rceil$ 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

$\tilde{\mathcal{U}} = \mathbb{R}, \Pr[o \leftarrow M(x)] = \frac{e^{-\frac{(o-x)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma}$. Given x_0 and x_1 we assume for simplicity $x_0 < x_1$

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

$$\omega(\infty) = 0, \quad \omega(-\infty) = 0, \quad \forall y \in \mathbb{R}, \omega(y) = \bar{\omega}(y)$$

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

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

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

$\tilde{\mathcal{U}} = \mathbb{R}, \Pr[o \leftarrow M(x)] = \frac{1}{2b} e^{\frac{|o-x|}{b}}$. Given x_0 and x_1 we assume for simplicity $x_0 < x_1$

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

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

$$\mu = \int_{-\infty}^{\infty} \mathcal{L}(o) A(o) do = e^{\frac{x_0-x_1}{b}} - \frac{b + x_0 - x_1}{b}$$

$$\begin{aligned} \sigma^2 &= \int_{-\infty}^{\infty} (\mathcal{L}(o) - \mu)^2 A(o) do \\ &= 3 - \frac{2e^{\frac{x_0-x_1}{b}}(b - 2(x_0 - x_1))}{b} - e^{\frac{2(x_0-x_1)}{b}} \end{aligned}$$

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

$$\begin{aligned} &\begin{cases} \int_{x_1}^{\infty} A(x) = \frac{1}{2} e^{\frac{x_1-x_0}{b}} & y = \frac{x_0-x_1}{b} \\ A(x(y)) \frac{\partial x}{\partial y} dy = \frac{1}{4} e^{\frac{by-x_0+x_1}{2b}} & \frac{x_0-x_1}{b} < y \leq \frac{x_1-x_0}{b} \\ \int_{-\infty}^0 A(x) = \frac{1}{2} & y = \frac{x_1-x_0}{b} \\ 0 & \text{else} \end{cases} \end{aligned}$$

$$\delta_A(\varepsilon) = \int_{\varepsilon}^{\infty} (1 - e^{\varepsilon-y}) d\omega(y) \quad \text{with } \varepsilon \geq 0$$

$$\begin{aligned} &\begin{cases} \frac{1}{2} e^{-\frac{x_0}{b}} \left(e^{\frac{b\varepsilon+x_0}{2b}} - e^{\frac{x_1}{2b}} \right)^2 \\ + \frac{1}{2} \left(1 - e^{\varepsilon - \frac{x_1-x_0}{b}} \right) & \varepsilon \leq \frac{x_1-x_0}{b} \\ 0 & \text{else} \end{cases} \end{aligned}$$

$\delta_{M(x_1)}(\varepsilon) = \delta_{M(x_0)}(\varepsilon) = \delta_A(\varepsilon)$ 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 11). Using our exact formulas of the privacy loss class of the Laplace mechanism (Appendix A.3), we get $\mu_{L,pld} = e^z - 1 - z$. As e^z can be represented as a Taylor expansion, $\sum_{k=i}^{\infty} \frac{z^k}{k!} =: T(i, z)$.

$$\mu_{L,pld} = T(1, z) - 1 - z = \underbrace{\frac{z^2}{2}}_{\mu_{G,pld}} + \underbrace{T(3, z)}_{>0} \geq \mu_{G,pld}$$

Similarly for the variance:

$$\begin{aligned} \sigma^2 &= 3 - \frac{2e^{\frac{x_0 - x_1}{\lambda}} (\lambda - 2(x_0 - x_1))}{\lambda} - e^{2\frac{(x_0 - x_1)}{\lambda}} \\ &= 3 - e^z (2 - 4z) - e^{2z} = 3 - e^z (2 - 4z + e^z) \\ &= 3z^2 - \left(\sum_{i=2}^{\infty} \frac{z^i}{i!} \right) (3 - 3z + e^z) \\ &= z^2 + z^3 - \frac{z^2}{2} \sum_{i=2}^{\infty} \frac{z^i}{i!} - \left(\sum_{i=3}^{\infty} \frac{z^i}{i!} \right) \left(4 - 2z + \sum_{i=2}^{\infty} \frac{z^i}{i!} \right) \\ &\stackrel{(a)}{\geq} z^2 + z^3 - \frac{z^2}{2} \sum_{i=2}^{\infty} \frac{z^i}{i!} - \left(\sum_{i=3}^{\infty} \frac{z^i}{i!} \right) \cdot (4 - z) \\ &= z^2 + \frac{2}{6} z^3 - \frac{z^4}{4} - \frac{z^5}{12} - \frac{z^2}{2} T(z, 4) - 4T(z, 4) + \frac{z^4}{6} + \frac{3}{2} zT(z, 4) \\ &\geq z^2 + \frac{2}{6} z^3 - \frac{z^4}{4} - \frac{z^5}{12} - \frac{z^2}{2} T(z, 4) - 4T(z, 4) + \frac{z^4}{4} + \frac{3}{2} zT(z, 4) \\ &\geq z^2 + \frac{2}{6} z^3 - \frac{1}{3} z^4 - 4T(z, 5) \geq z^2 + \frac{1}{12} z^3 \end{aligned}$$

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 $z^2 = \sigma_{G,pld}$.

B Proofs

Proof of Lemma 1. 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} , $o \in \mathcal{U}$:

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

Proof of Theorem 1. Let $M : \mathcal{X} \rightarrow \mathcal{U}$ and $M' : \mathcal{X}' \rightarrow \mathcal{U}'$ be two probabilistic mechanisms, let $x_0, x_1 \in \mathcal{X}$ and $x'_0, x'_1 \in \mathcal{X}'$. For ease of readability we write $\mathcal{U}^2 = \mathcal{U} \times \mathcal{U}'$. 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 $o \in \mathcal{U}, o' \in \mathcal{U}'$. We split $\mathcal{U}^2 = \mathcal{U} \times \mathcal{U}'$ into three sets as follows

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

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}_∞^2 (for $\omega(\infty)$ and partly $\omega(-\infty)$), and last \mathcal{U}_0^2 (leftover $\omega(-\infty)$).

First, we examine the set \mathcal{U}_+^2 . This will lead to the convolution property for $y \neq -\infty, \infty$. 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 $\omega_c(-\infty)$ in this part. For all events $(o, o') \in \mathcal{U}_+^2$, the privacy loss is additive under composition: $\forall (o, o') \in \mathcal{U}_+^2$

$$\begin{aligned} \mathcal{L}_{(M(x_0), M'(x'_0)) / (M(x_1), M'(x'_1))} (o, o') \\ &= \ln \left(\frac{\Pr[(o, o') \leftarrow (M(x_0), M'(x'_0))]}{\Pr[(o, o') \leftarrow (M(x_1), M'(x'_1))]} \right) \\ &= \ln \left(\frac{\Pr[o \leftarrow M(x_0)] \Pr[o' \leftarrow M'(x'_0)]}{\Pr[o \leftarrow M(x_1)] \Pr[o' \leftarrow M'(x'_1)]} \right) \\ &= \ln \left(\frac{\Pr[o \leftarrow M(x_0)]}{\Pr[o \leftarrow M(x_1)]} \right) + \ln \left(\frac{\Pr[o' \leftarrow M'(x'_0)]}{\Pr[o' \leftarrow M'(x'_1)]} \right) \\ &= \mathcal{L}_{M(x_0)/M(x_1)}(o) + \mathcal{L}_{M'(x'_0)/M'(x'_1)}(o') \end{aligned}$$

since M and M' are independent. Let us define

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

$$\omega_c(y_c) = \sum_{\{(o, o') \mid \mathcal{L}_c(o, o') = y_c\}} \Pr[(o, o') \leftarrow (M(x_0), M'(x'_0))]$$

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

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 $\omega_c(-\infty)$ here.

In the second part, we prove the composition of $\omega_c(\infty)$ and show that all events in \mathcal{U}_∞^2 which add to $\omega_c(-\infty)$ are zero. Next, we define four sets $\mathcal{U}_\infty, \mathcal{U}'_\infty, \mathcal{U}_+, \mathcal{U}'_+$ from the set \mathcal{U}_∞^2 and then subtract combinations of these four sets from \mathcal{U}_∞^2 to define \mathcal{U}_\perp^2 :

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

First, let us argue about $\omega(-\infty)$: 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 left-over 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 $\omega_c(\infty)$.

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

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

where we have separated the infinite sums as before (independence and Cauchy products) and we have used $\sum_{o \in \mathcal{U}_+} \Pr[o \leftarrow M(x_0)] = 1 - \omega(\infty)$ (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 $-\infty$, but any occurrence-probabilities are zero leads to the conclusion that its contribution to any event in ω_c is 0.

We show $\mathcal{Y}_c = \{y_c \mid y_c = y + y' \ \forall y \in \mathcal{Y}, \forall y' \in \mathcal{Y}'\}$. 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 $y = -\infty$. Together with the addition rules in Definition 4.2, it is valid to define $\mathcal{Y}_c = \mathcal{Y}_+ \cup \{-\infty, \infty\}$. 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 $\{-\infty, \infty\}$ are countable. this concludes the proof. \square

Proof of Lemma 2. Let us split \mathcal{U} in three sets

$$\begin{aligned}
 \mathcal{U}_+ &= \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) \in \mathcal{A} \setminus \{-\infty, \infty\}, o \in \mathcal{U}\} \\
 \mathcal{U}_\infty &= \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) = \infty, o \in \mathcal{U}\} \\
 \mathcal{U}_0 &= \{o \mid \mathcal{L}_{M(x_1)/M(x_0)}(o) = -\infty, o \in \mathcal{U}\}
 \end{aligned}$$

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 $\mathcal{A}^+ = \{-y \mid \forall y \in \mathcal{Y} \setminus \{-\infty, \infty\}\}$. So,

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

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

$$\varpi(\infty) = \sum_{o \in \mathcal{U}_\infty} \Pr[o \leftarrow M(x_1)] = \sum_{o \in \mathcal{U} \setminus \mathcal{U}_0, \mathcal{U}_+} \Pr[o \leftarrow M(x_1)]$$

$$= 1 - \underbrace{\varpi(-\infty)}_{=0} - \sum_{y \in \mathcal{L}^+} \varpi(y)$$

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

Proof of Lemma 3. Let \mathcal{Y} be the support of ω , $y \in \mathcal{Y} \setminus \{-\infty, \infty\}$. The PLD of $M'_{M,O}(x_0)$ over $M_{M,O'}(x_1)$ is

$$\begin{aligned} \omega'(y) &= \sum_{o \in \mathcal{L}^{-1}_{M'_{M,O}(x_0)/M'_{M,O}(x_1)}(y)} \Pr[o \leftarrow M'_{M,O}(x_0)] = \sum_{o \in \mathcal{L}^{-1}_{M(x_0)/M(x_1)}(y)} \frac{\Pr[o \leftarrow M(x_0)]}{\Pr[M(x_0) \in O]} \\ &= \frac{1}{1 - \Pr[M(x_0) \notin O]} \sum_{o \in \mathcal{L}^{-1}_{M(x_0)/M(x_1)}(y)} \Pr[o \leftarrow M(x_0)] \\ &= \frac{\omega(y)}{1 - \omega(\infty)} = \bar{\omega}(y). \end{aligned} \quad \square$$

Proof of Lemma 4. Let the variables be defined as in the lemma. The statement follows immediately from the definition, setting $g(\infty) = 1$:

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

Proof of Lemma 5. Let M be a probabilistic mechanism and $x_0, x_1 \in \mathcal{X}$ be two inputs. For simplicity, let us denote $A(o) := \Pr[o \leftarrow M(x_0)]$ and $B(o) := \Pr[o \leftarrow M(x_1)]$, and let $\mathcal{L}^{-1}_{A/B}(y) = \{o \mid y = \mathcal{L}_{A/B}(o), o \in \mathcal{U}\}$ be the pre-image of y . First, we show that

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

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 $y \in \mathcal{Y} \setminus \{-\infty, \infty\}$ and $\forall o \in \mathcal{L}^{-1}_{A/B}(y)$

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

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

where we have used the fact that $\forall o \in \mathcal{U}, A(o) \geq 0$. 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 $o \in \mathcal{U}$.

$$\sum_{o \in \mathcal{U}} \max(0, A(o) - e^\varepsilon B(o))$$

$$\begin{aligned} &= \sum_{o \in \mathcal{L}^{-1}(\infty)} \max(0, A(o) - \underbrace{e^\varepsilon}_{=0} \underbrace{B(o)}_{=0}) \\ &+ \sum_{o \in \mathcal{L}^{-1}(-\infty)} \max(0, \underbrace{A(o)}_{=0} - \underbrace{e^\varepsilon B(o)}_{\leq 0}) \\ &+ \sum_{y \in \mathcal{Y} \setminus \{-\infty, \infty\}} \sum_{o \in \mathcal{L}^{-1}(y)} \max(0, [1 - e^{\varepsilon-y}] \cdot A(o)) \\ &= \sum_{o \in \mathcal{L}^{-1}(\infty)} A(o) + \sum_{y > \varepsilon, y \neq \infty} \sum_{o \in \mathcal{L}^{-1}(y)} [1 - e^{\varepsilon-y}] \cdot A(o) \\ &= \omega(\infty) + \sum_{y > \varepsilon, y \in \mathcal{Y} \setminus \{-\infty, \infty\}} [1 - e^{\varepsilon-y}] \omega(y) = \delta^*_{M(x_0), M(x_1)} =: \delta^* \end{aligned}$$

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. What is left is the connection to tight-ADP; we use Lemma 1 in [20]:

Claim ([20, Lemma 1], Connection to tight-ADP).

For every ε , two distributions A and B over a finite universe \mathcal{U} are tightly (ε, δ) -ADP with

$$\delta = \max \left(\sum_{o \in \mathcal{U}} \max(\Pr[o \leftarrow A] - e^\varepsilon \Pr[o \leftarrow B], 0), \sum_{o \in \mathcal{U}} \max(\Pr[o \leftarrow B] - e^\varepsilon \Pr[o \leftarrow A], 0) \right),$$

which in application directly concludes the proof. \square

Proof of Lemma 6. To show bijectivity of R , we need to prove injectivity and surjectivity. First, some general considerations. According to Definition 4.5 and Lemma 5, for all mechanisms $M : \mathcal{X} \rightarrow \mathcal{U}$ and all inputs $x_0, x_1 \in \mathcal{X}$, $\delta(\varepsilon)$ of the tight ADP-graph is of the following form, where (\mathcal{Y}, ω) is the PLD of $M(x_0)$ and $M(x_1)$: $\delta(\varepsilon) = \sum_{y > \varepsilon, y \in \mathcal{Y}} (1 - e^{\varepsilon-y}) \omega(y)$. Moreover, for any PLD and for all $y \in \mathcal{Y}, \omega(y) > 0$ and thus each single summand $(1 - e^{\varepsilon-y}) \omega(y)$ that is included in the sum (i.e., $y > \varepsilon$) is always positive.

By definition, for all PLDs (\mathcal{Y}, ω) generated by all M on all inputs x_0, x_1 , the image of the map

$$R : (\mathcal{Y}, \omega) \rightarrow \left(\varepsilon, \sum_{y > \varepsilon, y \in \mathcal{Y}} (1 - e^{\varepsilon-y}) \omega(y) \right)_\varepsilon \quad (5)$$

contains all valid tight ADP-graphs. Therefore, R is surjective.

We now prove injectivity by contradiction. Assume there are two non-equal PLDs $(\mathcal{Y}, \omega), (\mathcal{Y}', \omega')$ for which R outputs the same tight ADP-graph $R(\mathcal{Y}, \omega) = (\varepsilon, \delta(\varepsilon))_\varepsilon$ and $R(\mathcal{Y}', \omega') = (\varepsilon, \delta'(\varepsilon))_\varepsilon$ with $R(\mathcal{Y}, \omega) = R(\mathcal{Y}', \omega')$.

If the PLDs consist of a single and identical y , i.e. $\mathcal{Y} = \mathcal{Y}' = \{y\}$, from $\delta(\varepsilon) = \delta'(\varepsilon)$ follows immediately $\omega(y) = \omega'(y)$ as the term $(1 - e^{\varepsilon-y})$ is identical in both. This is a contradiction. Otherwise, if we

have more than a single and identical y , we can find the minimal distance between any two y in $\mathcal{Y} \cup \mathcal{Y}'$: As k is finite and $\mathcal{Y}, \mathcal{Y}'$ are discrete, the minimal distance is $\eta = \min_{y_0, y_1 \in \mathcal{Y} \cup \mathcal{Y}', y_0 \neq y_1} |y_0 - y_1| > 0$. This means that $\forall y_0, y_1 \in \mathcal{Y} \cup \mathcal{Y}'$ with $y_0 \neq y_1$, $y_1 \notin (y_0 - \eta, y_0)$.

Let $y_n = \max_{y \in \mathcal{Y} \cup \mathcal{Y}'} y$. We prove the statement by induction. Base step: As $\delta'(y_n) = \delta(y_n)$, we get $\omega'(\infty) = \omega(\infty)$ immediately by Definition 4.6. Induction step: Let $y_n \in \mathcal{Y}$, otherwise, switch (\mathcal{Y}, ω) and (\mathcal{Y}', ω') . If $y_n \notin \mathcal{Y}'$, then $\forall \varepsilon \in (y_n - \eta, y_n)$, $\delta(\varepsilon) - \delta'(\varepsilon) = (1 - e^{\varepsilon - y_n})\omega'(y_n) > 0$. Therefore, $\delta(\varepsilon) \neq \delta'(\varepsilon)$. This is a contradiction. If $y_n \in \mathcal{Y}'$, then let $\varepsilon_1 = y_n - \frac{\eta}{2}$ and

$$\begin{aligned} \delta'(\varepsilon_1) &= (1 - e^{\varepsilon_1 - y_n})\omega'(y_n) + \sum_{y' > y_n, y' \in \mathcal{Y}'} (1 - e^{\varepsilon_1 - y'})\omega'(y') + \omega'(\infty) \\ \delta(\varepsilon_1) &= (1 - e^{\varepsilon_1 - y_n})\omega(y_n) + \sum_{y > y_n, y \in \mathcal{Y}} (1 - e^{\varepsilon_1 - y})\omega(y) + \omega(\infty) \end{aligned}$$

But by previous induction steps:

$$\sum_{y' > y_n, y' \in \mathcal{Y}'} (1 - e^{\varepsilon_1 - y'})\omega'(y') + \omega'(\infty) = \sum_{y > y_n, y \in \mathcal{Y}} (1 - e^{\varepsilon_1 - y})\omega(y) + \omega(\infty)$$

Therefore, by $\delta'(\varepsilon_1) = \delta(\varepsilon_1) \Rightarrow \omega'(y_n) = \omega(y_n)$ as the term $(1 - e^{\varepsilon_1 - y_n})$ is identical. Now we define $y_{n-1} = \max_{y \in \mathcal{Y} \cup \mathcal{Y}' \setminus \{y_m | m \geq n\}} y$ and repeat until $\mathcal{Y} \cup \mathcal{Y}' \setminus \{y_m | m \geq n\} = \{\}$. This will give us $\mathcal{Y} = \mathcal{Y}'$ and $\omega' = \omega$, which is in contradiction to our assumptions. This proves injectivity. As we have proven injectivity and surjectivity for R , we can conclude bijectivity. \square

Proof of Lemma 7. Let ω with support \mathcal{Y} be created by $M(x_0)$ and $M(x_1)$. First, notice that Equation (1) in the PDP definition is equal to the privacy loss function: for $i \in \{0, 1\}$:

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

Let us create two sets

$$S'_i := \{o \mid \mathcal{L}_{M(x_i)/M(x_{1-i})}(o) > \varepsilon, o \in S_i^\delta\}$$

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,

$$\mathcal{L}_{M(x_i)/M(x_{1-i})}(o) \leq \varepsilon \quad \forall o \in S \setminus S'_i = (S \setminus S_i^\delta) \cup (S_i^\delta \setminus S'_i) \quad (6)$$

which means that all $o \in S$ with $\mathcal{L}_i(o) > \varepsilon$ are in S'_i .

$$\begin{aligned} \delta &\geq \Pr[M(x_0) \in S'_0] \stackrel{I}{=} \sum_{o \in S'_0} \Pr[o \leftarrow M(x_0)] \\ &\stackrel{II}{=} \sum_{\{o \mid \mathcal{L}_{M(x_0)/M(x_1)}(o) > \varepsilon, o \in S\}} \Pr[o \leftarrow M(x_0)] \stackrel{III}{=} \sum_{y > \varepsilon} \omega(y) \quad y \in \mathcal{Y} \end{aligned}$$

where we have used independence of elementary events (I), Equation (6) (II), and the privacy distribution definition (III). The argument for ϖ follows analogously. This proves one direction of the lemma. For the other direction, note that we have only used equalities, that $S'_i \subseteq S_i^\delta$, and that $\forall o \in S \setminus S'_i, \mathcal{L}_{M(x_i)/M(x_{1-i})}(o) \leq \varepsilon \Rightarrow \forall o \in S \setminus S_i^\delta, \mathcal{L}_{M(x_i)/M(x_{1-i})}(o) \leq \varepsilon$. \square

Proof of Lemma 8. This proof is separated in three parts. First, we show Equation (2). Second, we show that there exists a bijection R' between Rényi-DP and the exponentiated PLD (exp-PLD), defined as $(\{exp(y) \mid y \in \mathcal{Y}\}, \omega \circ \ln)$, where (\mathcal{Y}, ω) is the PLD of A and B. Third, we show the existence of a bijection R between the exp-PLD and the PLD itself.

First, let us show the equality between 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 $A(o) \neq 0$. Therefore, we can do the following:

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

For the second part, the bijection between RDP and exp-PLD, we derive for any λ the corresponding moment as in the calculation from above. This is an algebraic identity, i.e. any PLD ω generated by two distributions A and B results in one specific Rényi sequence $(\alpha, D_\alpha(A|B))_\alpha$.

The other direction $\omega \circ \ln \Leftarrow (\alpha, D_\alpha)_\alpha$ is more tricky as there are cases where more than one distribution have the same moments (Hausdorff moments problem). First, let us define $\rho_\lambda := \exp(\lambda \cdot m_\lambda) > 0$ 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. More formally, if $\exists c, d, d' > 0$, s.t. $0 < \rho_\lambda < cd^\lambda \lambda!$ and $d' > d$, then $\forall r$ with $0 < r < \frac{1}{d'}$ the power series $0 < \sum_{\lambda} \rho_\lambda \frac{r^\lambda}{\lambda!} < \sum_{\lambda} cd^\lambda \lambda! \frac{r^\lambda}{\lambda!} < c \sum_{\lambda} \left(\frac{d}{d'} \right)^\lambda < \infty$ as $0 < \frac{d}{d'} < 1$ leads to a geometric series. Now we apply the following claim about probability measures and moments:

Claim ([6, Theorem 30.1]). *Let μ be probability measure on the line having finite moments $a = \int_{-\infty}^{\infty} x^k \mu(dx)$ of all orders. If the power series $\sum_k \alpha_k r^k / k!$ has a positive radius of convergence, then μ is the only probability measure with the moments $(\alpha_i)_{i \in \mathbb{N}}$.*

By this claim, we know that, if the previous power series has a positive convergence radius, then there exists a unique probability measure μ for a given series of moments $(D_\alpha(A|B))_\alpha$. As shown above, the moments of the exp-PLD are exactly the Rényi divergences; hence, this uniqueness of the measure implies that μ equals the exp-PLD $(\mathcal{Y}', \omega \circ \ln)$.

For the third part, we then observe that we can transform any PLD into the exp-PLD $(\mathcal{Y}', \omega \circ \ln)$ and vice versa. This is a bijective step, since the exponentiation $\exp : (-\infty, \infty) \rightarrow (0, \infty)$ is bijective and the logarithm $\ln : (0, \infty) \rightarrow (-\infty, \infty)$ is bijective on the domain of strictly positive real values. More precisely, for a probability space (\mathcal{Y}, E, ω) the following function R is bijective, where \mathcal{Y} is the set of atomic events, $E := 2^\mathcal{Y}$ is the set of all events on \mathcal{Y} , and ω is the probability measure:

$$R(\mathcal{Y}, \omega) := (\underbrace{\{\exp(y) \mid y \in \mathcal{Y}\}}_{=: \mathcal{Y}'}, \underbrace{\omega \circ \ln}_{=: \omega'})$$

Next, we will show that, with $E' = 2^{\mathcal{Y}'} = \{\{\exp(y) \mid y \in v\} \mid v \in E\}$, $(\mathcal{Y}', E', \omega')$ is a probability space. In particular, we show that in this new probability space, $\omega \circ \ln$ is countably additive: for any countable collection $\{E'_i\}_{i \in I}$ with pairwise disjoint events $E'_i \in E'$, we know that there are events $E_i \in E$ s.t., $E'_i = \{\exp(v) \mid v \in E_i\}$ for all i . We show that ω' on $\{E'_i\}_{i \in I}$ can be expressed using ω and $\{E_i\}_{i \in I}$ as follows.

$$\begin{aligned} \omega' \left(\bigcup_{i \in I} E'_i \right) &= \omega \left(\ln \left(\bigcup_{i \in I} E'_i \right) \right) = \omega \left(\left\{ \ln(v') \mid v' \in \bigcup_{i \in I} E'_i \right\} \right) \\ &= \omega \left(\left\{ \underbrace{\ln(\exp(v))}_{=: v} \mid v \in \bigcup_{i \in I} E_i \right\} \right) = \omega \left(\bigcup_{i \in I} \{v \mid v \in E_i\} \right) \end{aligned}$$

As \mathcal{Y} is countable and ω is countably additive, we can write this as a sum (*). Since the sets E_i are pairwise disjoint, equality (**) holds. Plugging in the definition of ω' and E_i , we get the following.

$$\stackrel{(*)}{=} \sum_{v \in \bigcup_{i \in I} E_i} \omega(v) \stackrel{(**)}{=} \sum_{i \in I} \omega(E_i) = \sum_{i \in I} \omega \circ \ln(E'_i) = \sum_{i \in I} \omega'(E'_i)$$

Hence, we obtain that (\mathcal{Y}', ω') (together with $E' = 2^{\mathcal{Y}'}$) is a probability space. Now, we can create a bijection R_{RDP} . As $R(\mathcal{Y}, \omega) = (\mathcal{Y}', \omega \circ \ln)$ and $R'(\mathcal{Y}', \omega \circ \ln) = (\alpha, D_\alpha)_\alpha$ are both bijections, $R_{\text{RDP}} := R' \circ R$ is also a bijection and $R_{\text{RDP}}(\mathcal{Y}, \omega) = R'(R(\mathcal{Y}, \omega)) = R'(\mathcal{Y}', \omega \circ \ln) = (\alpha, D_\alpha)_\alpha$. \square

Proof of Theorem 2. Lemma 6 states that there is a bijection R_{ADP} such that for all mechanisms and all pairs of inputs x_0, x_1 such that the support of the PLD \mathcal{Y} has finite cardinality $|\mathcal{Y}| = k$ (for $k \in \mathbb{N}$) we have

$R_{\text{ADP}}((\varepsilon, \delta(\varepsilon))_{\varepsilon \in \mathbb{R}}) = \omega$. Lemma 8 states that there is a bijection R_{RDP} such that for all mechanisms with the support of $M(x_0)$ and of $M(x_1)$ is countable and $\exp(\lambda \cdot |m_\lambda|) < cd^\lambda \lambda!$ for two positive constants c, d we have $R_{\text{RDP}}((\alpha, D_\alpha(M(x_0) || M(x_1)))_\alpha) = \omega$. Note that

$$m_\lambda = \frac{1}{\lambda} \ln \left(\mathbb{E}_{y \sim \omega} e^{\lambda y} \right) = \frac{1}{\lambda} \ln \left(\mathbb{E}_{o \sim M(x_0)} e^{\lambda \ln \frac{\Pr[\alpha \leftarrow M(x_0)]}{\Pr[\alpha \leftarrow M(x_1)]}} \right)$$

leads to $\mathbb{E}_{o \sim M(x_0)} \left[\frac{\Pr[\alpha \leftarrow M(x_0)]}{\Pr[\alpha \leftarrow M(x_1)]} \right]^\lambda < cd^\lambda \lambda!$. As both mappings are bijections, the diagram commutes. \square

Proof of Corollary 1. This follows directly from Lemma 3 and Theorem 2. \square

Proof of Theorem 3. Let ω_n generated by $A = M^n(x_0)$ and $B = M^n(x_1)$ 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

$$\alpha_{A^n, B^n}(\lambda) \leq \sum_{i=0}^n \alpha_{A, B}(\lambda) = n \cdot \alpha_{A, B}(\lambda)$$

with

$$\alpha_{A, B}(\lambda) = \ln \mathbb{E}_{o \sim A} e^{\lambda \ln \frac{\Pr[\alpha \leftarrow A]}{\Pr[\alpha \leftarrow B]}}$$

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

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

From which follows for ω_n and the corresponding \mathcal{Y}_n

$$\begin{aligned} \sum_{y \geq \gamma, y \in \mathcal{Y}_n} \omega_n(y) &\leq \min_{\lambda > 0} \exp(\alpha_{A^n, B^n}(\lambda) - \lambda \gamma) \\ &\leq \min_{\lambda > 0} \exp(n \cdot \alpha_{A, B}(\lambda) - \lambda \gamma) \\ &= \min_{\lambda > 0} \exp \left(n \cdot \ln \mathbb{E}_{o \sim A} \left[e^{\lambda \ln \frac{\Pr[\alpha \leftarrow A]}{\Pr[\alpha \leftarrow B]}} \right] - \lambda \gamma \right) \\ &= \min_{\lambda > 0} \frac{\mathbb{E}_{o \sim A} \left[e^{\lambda \cdot \ln \frac{\Pr[\alpha \leftarrow A]}{\Pr[\alpha \leftarrow B]}} \right]^n}{e^{\lambda \cdot \gamma}} = \mathcal{T}(\gamma) \end{aligned}$$

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 want to show

$$\begin{aligned} &\sum_{y > \varepsilon, y \in \mathcal{Y}_n} (1 - e^{\varepsilon - y}) \omega_n(y) \\ &\leq \mathcal{T}(y_k) + \sum_{y_j \geq \varepsilon, y_j \in \mathcal{P}} (1 - e^{\varepsilon - y_j}) [\mathcal{T}(y_{j-1}) - \mathcal{T}(y_j)] \end{aligned} \quad (7)$$

Due to the $(1 - e^{\varepsilon - y})$ terms, which are smaller than 1 (as in the RDP bound formula [21]), this bound is less or equal to the Rényi-DP bound. To see why Equation (7) 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 chunks with boundaries $\mathcal{P} = \{y_0, \dots, y_k\} \subseteq \mathbb{R}^{k+1}$ with $y_i < y_{i+1} \forall i$. We define, for $i \in \{0, \dots, k\}$, $T''(i) := \sum_{y \in \mathcal{Y}_n, y \geq y_i} \omega_n(y)$, $T'(k) := T''(k)$, for $i < k$: $T'(j) := T''(i) - T''(i+1)$. We retain for every $y_i \in \mathcal{P}$,

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

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

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

Note that for functions f_1, f_2 s.t. for all $i \in \{0, \dots, k\}$: $\sum_{i \leq j, j \in \{0, \dots, k\}} f_1(j) \geq 0$, $\sum_{i \leq j, j \in \{0, \dots, k\}} f_2(j) \geq 0$ and $\sum_{i \leq j, j \in \{0, \dots, k\}} f_1(j) \leq \sum_{i \leq j, j \in \{0, \dots, k\}} f_2(j)$ and for all monotonously increasing functions $g \geq 0$,

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

We split the Markov tails \mathcal{T} into \mathcal{T}' analogously to how we have split T into T' : $\mathcal{T}'(k) := \mathcal{T}(y_k)$ and for $i < k$: $\mathcal{T}'(i) := \mathcal{T}(y_i) - \mathcal{T}(y_{i+1})$. We again retain for every $y_i \in \mathcal{P}$, $\mathcal{T}(y_i) = \sum_{y_j \geq y_i, y_j \in \mathcal{P}} \mathcal{T}'(j)$. Note that for all $i \in \{0, \dots, k\}$, $\sum_{i \leq j, j \in \{0, \dots, k\}} \mathcal{T}'(j) \leq \sum_{i \leq j, j \in \{0, \dots, k\}} \mathcal{T}'(j)$, and that furthermore we are now finally able to apply our property. Given $y_i \in \mathcal{P}$, we get

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

Proof of Theorem 4. 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 $\omega_n(y) \forall y \in \mathcal{Y}_n \setminus \{-\infty, \infty\}$ for the Gauss shape. To ease readability we write \mathbf{y} for a vector of elements y_1, \dots, y_k and we omit the exact declaration $\mathbf{y} = y_1, \dots, y_k$ 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 ω_1 twice for Theorem 1, we get directly

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

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

$$\begin{aligned} \mathcal{Y}_{n+1} &= \left\{ \hat{y} \mid \hat{y} = y + \sum_{i=1}^n \tilde{y}_i + y, \forall y \in \mathcal{Y}, \forall \tilde{\mathbf{y}} \in \mathcal{Y}^n \right\} \\ &= \left\{ \hat{y} \mid \hat{y} = \sum_{i=1}^{n+1} \tilde{y}_i, \forall \tilde{\mathbf{y}} \in \mathcal{Y}^{n+1} \right\} \end{aligned}$$

$$\begin{aligned} \omega_{n+1}(y) &= (\omega * \omega_n)[y] \quad \forall y \in \mathcal{Y}_{n+1} \setminus \{-\infty, \infty\} \\ &= (\otimes_{i=1}^{n+1} \omega_1)[y] \\ \omega_{n+1}(\infty) &= 1 - [1 - \omega_n(\infty)] \cdot [1 - \omega_1(\infty)] \\ &= 1 - [1 - \omega_1(\infty)]^{n+1} \\ \omega_{n+1}(-\infty) &= 0 \end{aligned}$$

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

For the rest of this proof, we omit $\omega_i(-\infty)$ 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 $\forall y \in \mathcal{Y}_i \setminus \{-\infty, \infty\}$

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

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

$$\begin{aligned} \mu_{i+j} &= \mathbb{E}_{y \sim \bar{\omega}_{i+j}} y = \sum_{y \in \mathcal{Y}_{i+j}} \bar{\omega}_{i+j}(y) y \\ &= \sum_{y \in \mathcal{Y}_{i+j}} \sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) \cdot \bar{\omega}_j(y - y_i) y \end{aligned}$$

$$\begin{aligned}
 & \stackrel{I}{=} \sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) \cdot \sum_{y_j \in \mathcal{Y}_j} \bar{\omega}_j(y_j) (y_i + y_j) \\
 & \stackrel{II}{=} \sum_{y_i \in \mathcal{Y}_i} \bar{\omega}_i(y_i) \cdot (y_i + \mu_j) \stackrel{III}{=} \mu_i + \mu_j
 \end{aligned}$$

where we have used a variable shift $y \rightarrow y_i + y_j$ 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 $\sigma_{m+l}^2 = \sigma_l^2 + \sigma_m^2$, 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 13 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 4 and by the fact that $\Pr_{\omega_n}[y \leq z | y \neq \infty] = \sum_{y \leq z, y \in \mathcal{Y} \setminus \{-\infty, \infty\}} \omega_n(y)$. \square

Proof of Lemma 9. First, note that $\lambda(u)$ is a σ -finite measure. The push-forward measure we can define as $\tilde{\mathcal{Y}}$ and $\tilde{\mathcal{U}}$ are both a subset of \mathbb{R} [7].

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[7], and we can write $\omega(y)$.

To the generalized statements: The inner distribution (Equation (3)) 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 $\tilde{\mathcal{Y}} \subseteq \mathbb{R}$ without $-\infty$ and ∞ . The derivation of $\delta_{M(x_0)}(\varepsilon)$ identical to Lemma 5 except that the set $\tilde{\mathcal{Y}}$ does not include the distinguishing events. \square

Proof of Lemma 10. As we can evaluate a continuous function $f(y)$ in the privacy loss space as $\int_{\tilde{\mathcal{U}}} f(\mathcal{L}(o)) A(o) do$, we can apply integration by substitution with \mathcal{L}^{-1} :

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

\square

Proof of Lemma 11. Let the variables be defined as in the lemma statement. Let $u \in \tilde{\mathcal{U}}$. This is an application of lemma 10. The privacy loss function $\mathcal{L} : \tilde{\mathcal{U}} \rightarrow \mathbb{R}$ is

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

Note $\forall u \in \tilde{\mathcal{U}} : \mathcal{L}_{\mathcal{N}(x_0, \sigma^2)/\mathcal{N}(x_1, \sigma^2)}(u) \neq \infty \Rightarrow \omega(\infty) = 0$. 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 Riemann-integrals. Let $A(u) := \Pr[u \leftarrow M(x_0)]$. Using Lemma 10,

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

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 $(\frac{(x_0 - x_1)^2}{2\sigma^2}, \frac{(x_0 - x_1)^2}{\sigma^2}, 0)$. \square

Proof of Lemma 12. 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.

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

The lemma statement follows directly by combining everything. \square

Proof of Theorem 5. Let the variables be as in the theorem statement and let $x_0, x_1 \in \mathbb{R}$ with $|x_0 - x_1| = s$. By Lemma 11 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 $\left(\frac{(x_0 - x_1)^2}{2\sigma^2}, \frac{(x_0 - x_1)^2}{\sigma^2}, 0\right)$. It is well known that a convolution of two Gauss is a Gauss again

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

which can be generalized to $\bigoplus_{i=0}^n \mathcal{N}(y, \sigma^2) = \mathcal{N}(ny, n\sigma^2)$. 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 12 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 5 in tight $(\varepsilon, \delta(\varepsilon))$ -ADP. \square

Proof of Corollary 2. 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. \square

Lemma 13 (Berry-Esseen and Nagaev Bound, [23]).

Let X_1, \dots, X_n be independent and identically distributed zero mean random variables with $S := X_1 + \dots + X_n$, $\gamma = \mathbb{E}|X_i|^3 < \infty$, and $\sigma := \sqrt{\mathbb{E}|X_i|^2}$, then

$$|\Pr[S > n\sigma z] - \Pr[Z > z]| \leq c_u \frac{\gamma}{\sqrt{n}\sigma^3} \quad (\text{Berry-Esseen})$$

$$|\Pr[S > n\sigma z] - \Pr[Z > z]| \leq c_t \frac{\gamma}{\sqrt{n}\sigma^3(1 + z^3)} \quad (\text{Nagaev})$$

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 $c_t < 31.935$ [23].

Proof of Theorem 6. Let the variables be defined as in the theorem statement. Let $\Phi_n(z)$ be the cumulative distribution function of $\mathcal{N}(n\mu, n\sigma^2)$. We use 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 $\sigma^2 = \mathbb{E} |\bar{\omega}_1(y)|^2$. First, we prove that $\forall \varepsilon > 0$ we have

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

where $Z_n \sim \mathcal{N}(n\mu, n\sigma^2)$ and where $r_{u/t}(z)$ denotes either r_u or $r_t(z)$ from Lemma 13. As $\forall z \geq 0$, $r_t(z) \leq \infty$ and $r_u \leq \infty$, we obtain always a valid bound if we take the minimum for r_u and r_t . Second, let $\forall y, \varepsilon \in \mathbb{R}$: $g(\varepsilon - y) := (1 - e^{\varepsilon - y})$. Note that $\forall \varepsilon \in \mathbb{R}$, $\forall y \geq \varepsilon$, $0 \leq g(\varepsilon - y) < 1$. For simplicity, let $\bar{\delta}_\omega(\varepsilon) := \frac{\delta_\omega(\varepsilon) - \omega(\infty)}{1 - \omega(\infty)}$.

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

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.

Next, we include $\omega_n(\infty)$. Theorem 4 implies $\omega_n(\infty) = 1 - [1 - \omega_1(\infty)]^n$. Multiplying $[1 - \omega_n(\infty)]$ and adding zero results in

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

Together with the definition of $\delta_{\omega_n}(\varepsilon) = \delta_{M^n(x_0), M^n(x_1)}^*(\varepsilon)$ and Lemma 12, we obtain

$$\left| \delta_{M^n(x_0), M^n(x_1)}^*(\varepsilon) - \Delta_{\omega_n(\infty)} \right| \leq \beta_{\omega_n(\infty)}$$

This defines an upper bound: $\delta_{M^n(x_0), M^n(x_1)}^*(\varepsilon) \leq \Delta_{\omega_n(\infty)} + \beta_{\omega_n(\infty)}$. By applying the same proof before to the dual distribution ω_n , we can bound $\delta_{M^n(x_0), M^n(x_1)}^*(\varepsilon)$ and $\delta_{M^n(x_0), M^n(x_1)}^*(\varepsilon)$. By taking the maximum, we get that M is $(\varepsilon, \max(\Delta_\omega + \beta_\omega, \Delta_\gamma + \beta_\gamma))$ -ADP. \square